

INVESTIGATING THE QUANTUM PROPERTIES OF JETS AND
THE SEARCH FOR A SUPERSYMMETRIC TOP QUARK
PARTNER WITH THE ATLAS DETECTOR

A DISSERTATION
SUBMITTED TO THE DEPARTMENT OF PHYSICS
AND THE COMMITTEE ON GRADUATE STUDIES
OF STANFORD UNIVERSITY
IN PARTIAL FULFILLMENT OF THE REQUIREMENTS
FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

Benjamin Philip Nachman

July 2016



© 2016 by Benjamin P. Nachman. All Rights Reserved.

Re-distributed by Stanford University under license with the author.



This work is licensed under a Creative Commons Attribution-Noncommercial 3.0 United States License.

<http://creativecommons.org/licenses/by-nc/3.0/us/>

This dissertation is online at: <http://purl.stanford.edu/qp879cw5022>

I certify that I have read this dissertation and that, in my opinion, it is fully adequate in scope and quality as a dissertation for the degree of Doctor of Philosophy.

Dong Su, Primary Adviser

I certify that I have read this dissertation and that, in my opinion, it is fully adequate in scope and quality as a dissertation for the degree of Doctor of Philosophy.

Ariel Schwartzman, Co-Adviser

I certify that I have read this dissertation and that, in my opinion, it is fully adequate in scope and quality as a dissertation for the degree of Doctor of Philosophy.

Michael Peskin

I certify that I have read this dissertation and that, in my opinion, it is fully adequate in scope and quality as a dissertation for the degree of Doctor of Philosophy.

Lauren Tompkins

Approved for the Stanford University Committee on Graduate Studies.

Patricia J. Gumpert, Vice Provost for Graduate Education

This signature page was generated electronically upon submission of this dissertation in electronic format. An original signed hard copy of the signature page is on file in University Archives.

Abstract

Quarks and gluons are the fundamental building blocks of matter responsible for most of the visible energy density in the universe. However, they cannot be directly observed due to the confining nature of the strong force. The Large Hadron Collider (LHC) uses proton-proton collisions to probe the highest energy reactions involving quarks and gluons happening at the smallest distance scales ever studied in a terrestrial laboratory. The observable consequence of quark and gluon production in these reactions is the emergent phenomenon known as the *jet*: a collimated stream of particles traveling at nearly the speed of light. The quantum properties of the initiating quarks and gluons are encoded in the distribution of energy inside and around jets. These *quantum properties of jets* can be used to study the high energy nature of the strong force and provide a way to tag the hadronic decays of heavy boosted particles. The ATLAS detector at the LHC is well-suited to perform measurements of the internal structure of high energy jets. A variety of novel techniques utilizing the unique capabilities of the ATLAS calorimeter and tracking detectors are introduced in order to probe the experimental and theoretical limits of the quantum properties of jets.

Studying quarks and gluons may also be the key to understanding the fundamental problems with the Standard Model (SM) of particle physics. In particular, the top quark has a unique relationship with the newly discovered Higgs boson and as such could be a portal to discovering new particles and new forces. In many extensions of the SM, the top quark has a partner with similar relationships to other SM particles. For example, a scalar top partner (stop) in Supersymmetry (SUSY) could solve the Higgs boson mass hierarchy problem. Miraculously, a SUSY neutralino could also

account for the dark matter observed in the universe and may be copiously produced in stop decays. High-energy top quarks from stop decays result in jets with a rich structure that can be identified using the techniques developed in the study of the quantum properties of jets. While there is no significant evidence for stop production at the LHC, the stringent limits established by this search have important implications for SUSY and other models.

In memory of Isabella Threlkeld

Acknowledgements

Like any good story, a thesis describes a grand adventure with many twists and turns. I leave the quality of this story for you to judge, but I cannot proceed without thanking the many characters behind the scenes who have supported my adventure and contributed to my ongoing development as a physicist. My first high energy physics experiment was part of the Cosmic Ray Observatory Project (CROP) when I was in High School. Thanks to the guidance of John Rogers and the support of Dan Claes and Greg Snow, I had fun building and operating muon detectors on the roof of my school and my house. In college, I had many inspiring teachers including Csaba Csaki, Yuval Grossman, Flip Tanedo, Camil Muscalu, and Ravi Ramakrishna. I had the pleasure of working with and learning from Itai Cohen about fluid dynamics and from Keith Dennis about group theory. I am indebted to Jim Alexander for showing me how to think like an experimental particle physicist. I learned a lot from and with Adam Dishaw and Nathan Mirman about statistics and top quark physics. Across the ocean, I learned from Christopher Lester that science really does happen on the back of an envelope. Studying with Ben Allanach, I built a foundation in Supersymmetry theory.

Since formally starting at SLAC and with ATLAS, I have been blessed to work with many kind, passionate, and knowledgeable physicists. There are simply too many wonderful people to thank all of them here - I apologize! First of all, I would like to thank all of the (sub)group conveners, editorial board (chairs), and members at large for their insightful comments on my work. It was hard to make small list, but I'm grateful to these people in particular for advice, opportunities, and extensive feedback: Alison Lister, Mark Owen, Tancredi Carli, Andreas Hocker, Jan Kretzschmar,

Bogdan Malaescu, Monica D’Onofrio, Jamie Boyd, Tomasso Lari, Iacopo Vivarelli, Michael Begel, David Miller, Cigdem Issever, Judith Katzy, Ian Hinchliffe, Luciano Mandelli, Tony Doyle, and Mathieu Benoit. Till Eifert has been an amazing mentor, colleague, and friend. I have thoroughly enjoyed a camaraderie over jets with Max Swiatlowski, David Lopez Mateos, and Nurfikri Norjoharuddeen. Thank you to Matt Schwartz, Jesse Thaler, and Andrew Larkoski for teaching me jet phenomenology. The SLAC/Stanford ATLAS group has been extremely supportive, especially Charlie Young, Philippe Grenier, Rainer Bartoldus, Lauren Tompkins, Pascal Nef, Michael Kagan, Francesco Rubbo, Qi Zeng, Zihao Jiang, and Aviv Cukierman. I am very grateful to Michael Peskin for many fun conversations about particle physics - I have never met someone so knowledgeable, clear, and enthusiastic. All of my high energy physics adventures would not have been possible without generous funding from the National Science Foundation, the Department of Energy, and the Stanford Graduate Fellowship. This funding would have been useless without the strong support, guidance, and mentorship from my fantastic advisers Su Dong and Ariel Schwartzman. Thank you Lester Mackey for many fun discussions about connecting machine learning and particle physics and for chairing my thesis defense. Thank you Su Dong, Ariel, Michael P., Lauren, and Till for providing feedback on this thesis (and for your patience with its length).

My parents, Beth (to the muons, moon and back!) and Gary, and brothers, Marty and Lev, have always supported me in all my endeavors. No finite amount of text could describe how much I am indebted to them. This thesis is a discussion about the smallest distance scales ever studied (on Earth), so my enormous gratitude for them must wait until I see them in person. I am additionally grateful for my extended Bay Area family: Bart, Carrie, Nathan, Brynna, Cole, Ali, Oli, Felix, and Maceo.

Miracles do happen. Three days after defending this thesis, I married my soulmate Hannah Joo. Not only is she the most beautiful, wise, caring, and careful person I have ever met, she has been my anchor through the ups and downs of graduate student life. There are no gloomy days when my sunshine is forever nearby.

Contents

Abstract

Acknowledgements

0.1	Preface	1
0.1.1	Units	1
0.1.2	Coordinates	2
0.1.3	Statistics	2
I The Theory of Experimental Particle Physics		
1	The Standard Model of Particle Physics	
1.1	Particles and Forces	7
1.2	Successes and Limitations	18
2	Experimental Apparatus	
2.1	The Large Hadron Collider	22
2.2	Interactions of Particles with Matter	26
2.3	The ATLAS Detector	32
2.3.1	Inner Detector	33
2.3.2	Electromagnetic and Hadronic Calorimeters	35
2.3.3	Muon Spectrometer	36
2.3.4	Trigger System	37
3	Monte Carlo Simulation	
3.1	Matrix Elements	39

3.2	Fragmentation and the Underlying Event	40
3.3	Material Interactions and Detector Simulation	42
4	Event Reconstruction	
4.1	Calorimeter-cell Clusters	44
4.2	Charged Particle Tracks	44
4.3	Jets	46
4.4	Electrons, Photons, Muons, and Taus	50
4.5	Missing Transverse Momentum	51
II	The Quantum Properties of Jets	
1	Jet Charge	
1.1	Background	58
1.1.1	Constructing the jet charge	58
1.1.2	Jet Charge Properties	61
1.1.2.1	Sensitivity to Soft Radiation	61
1.1.2.2	Lorentz Invariance	64
1.1.3	From Parton Charge to Jet Charge	67
1.1.4	Charge Tagging	73
1.2	Analysis Design	76
1.2.1	Data and simulated samples	78
1.2.1.1	Jet Triggers	78
1.2.2	Object reconstruction and event selection	81
1.2.2.1	Object reconstruction at detector level	82
1.2.2.2	Object definitions at particle level	83
1.3	Performance Studies	86
1.3.1	Comparisons Between Data and Simulation	86
1.3.2	Modeling and Tagging Performance with QCD Jets	93
1.3.2.1	Single Jet Charge	93
1.3.2.2	Re-examining the Definition of Jet Charge	101
1.3.3	Simulation and In-situ Studies with W Bosons	103

1.3.3.1	Dataset and Simulation Samples	103
1.3.3.2	Charge Reconstruction Performance	107
1.3.3.3	Charge Tagging in a Boosted Topology	111
1.4	Unfolding	115
1.4.1	Unfolding for the Jet Charge	120
1.4.1.1	Binning	122
1.4.1.2	Correction Factors	125
1.4.1.3	Response Matrix	127
1.4.1.4	The Number of Iterations	129
1.4.1.5	Unfolded Data	130
1.5	Systematic uncertainties	133
1.5.1	Correction Factors	137
1.5.2	Calorimeter jet uncertainties	139
1.5.3	Tracking uncertainties	140
1.5.3.1	Truth Matching	141
1.5.3.2	Isolated Track Reconstruction Efficiency	143
1.5.3.3	Track Reconstruction Inside Jets	145
1.5.3.4	The ζ Method	159
1.5.3.5	Track Momentum Resolution	165
1.5.3.6	Track Charge Identification	167
1.5.3.7	Fake Tracks	170
1.5.3.8	Charged Particle Multiplicity	173
1.5.4	Unfolding Non-closure	177
1.6	Results	181
1.6.1	Unfolded Jet Charge Spectrum	181
1.6.2	Sensitivity of PDF Modeling	183
1.6.3	Sensitivity of QCD Models and Tunes	184
1.6.4	Model comparison overview	188
1.6.5	The average up-quark and down-quark jet charges	191
1.6.6	Up- and down-quark jet charge dependence on $\mathbf{p}_T$	193
1.7	Summary	196

2	Color flow	
2.1	Introduction	198
2.1.1	Color flow in QCD	200
2.1.2	Jet Pull	203
2.2	Analysis Design	205
2.2.1	Dataset and Simulation	205
2.2.1.1	Data-driven backgrounds	206
2.2.1.2	Simulating Exotic Colorflow	210
2.2.2	Object Reconstruction and Event Selection	216
2.2.3	Particle-level Event Selection	216
2.2.3.1	W Boson Candidate Selection	219
2.2.3.2	Comparisons Between Data and Simulation	223
2.3	Jet Pull Reconstruction Performance	226
2.3.1	Detector effects	226
2.3.2	Jet Pull Angle Response	228
2.3.3	Jet Kinematics and the Jet Pull Angle Response	233
2.3.4	Relationship Between the Jet Pull Angle Response and Jet Constituents	236
2.3.5	Jet Pull Angle and Event Properties	238
2.3.5.1	Jet Labeling in $t\bar{t}$ Events	238
2.3.5.2	Pileup	239
2.3.6	Comparisons between Data and Simulation	241
2.3.6.1	Cluster and Jet Origin Corrections	244
2.4	Unfolding	252
2.4.1	Unfolding Parameters	252
2.4.2	Correction Factors	259
2.4.3	Response Matrix	262
2.5	Systematic uncertainties	264
2.5.1	Tracking	264
2.5.2	Calorimeter Cell Clusters	266
2.5.2.1	Cluster Reconstruction Efficiency	267

2.5.2.2	Cluster Energy Scale	268
2.5.2.3	Cluster Angular Resolution	271
2.5.3	Jet Angular Resolution	274
2.5.3.1	In-situ method using track jets	275
2.5.4	Other Experimental Uncertainties	279
2.5.5	Background Processes	279
2.5.6	Top Quark Pair Production Modeling	280
2.5.6.1	ME Generator and Fragmentation Model	280
2.5.6.2	Color Reconnection	280
2.5.6.3	Initial and Final State Radiation	282
2.5.6.4	Top p_T	284
2.5.6.5	Color flow Model	284
2.5.6.6	Other	285
2.5.7	Correction Factors	285
2.5.8	Non-closure	286
2.5.9	Summary	288
2.6	Results	293
2.7	Summary	296
3	Constituent Multiplicity	
3.1	Motivation	299
3.1.1	QCD Predictions for Multiplicity	299
3.1.2	Quark and Gluon Tagging	303
3.2	Analysis Design	306
3.3	Unfolding	311
3.4	Systematic uncertainties	318
3.4.1	Response Matrix	318
3.4.1.1	Charged-energy loss in the dense core of jets	318
3.4.2	Correction Factors	320
3.4.3	Unfolding Procedure	322
3.4.4	Summary	324

3.5	Results	326
3.5.1	Quark and Gluon Multiplicity	329
3.6	Summary	334
4	Boson and Top Quark Jets	
4.1	Jet Mass	339
4.1.1	The Mass of Quark and Gluon Jets	339
4.1.2	Reconstructing the Calorimeter Jet Mass	346
4.1.2.1	Track-jet Method	353
4.1.2.1.1	Interpreting r_{track} Uncertainties	358
4.1.2.2	Bottom-up Method	360
4.1.2.3	Resonance Method	362
4.1.2.3.1	Dataset and Event Reconstruction	363
4.1.2.3.2	Extracting the Jet Mass Scale and Resolution	364
4.1.2.3.3	Systematic Uncertainties	369
4.1.2.3.4	Results	370
4.1.2.3.5	Intermediate Conclusions	375
4.1.2.3.6	Improvements and Prospects with 13 TeV data	376
4.1.3	Re-clustered Jet Mass	383
4.1.3.1	Re-clustering Jets	385
4.1.3.2	Particle-level Jet Mass Performance	387
4.1.3.3	Re-clustered Jet Substructure	391
4.1.3.4	Detector-level Jet Tagging	393
4.1.3.4.1	Low p_T : 200 GeV - 350 GeV	393
4.1.3.4.2	High p_T : 500 GeV - 1000 GeV	397
4.1.3.5	Re-clustered Jet Resolution	400
4.1.3.6	Close-by Effects	404
4.1.3.6.1	Triple Ratio Technique	405
4.1.3.6.2	In-situ Close-by Results for the Jet Mass	407
4.1.3.7	Summary and Outlook	409
4.1.4	Track-assisted Jet Mass	410

4.1.5	Conclusions and Future Outlook	418
4.2	Boson Type Tagger	419
4.2.1	Introduction	419
4.2.2	Datasets	421
4.2.3	Distinguishing a Z boson from a W boson	423
4.2.4	Definitions of reconstructed objects	426
4.2.5	Tagger performance	434
4.2.6	Systematic uncertainties	441
4.2.7	Validation of tagging variables using data	447
4.2.8	Outlook	458
4.2.9	Conclusions	459
4.3	Machine Learning for Jet Tagging	461
4.3.1	Fuzzy Jets	462
4.3.1.1	Mixture Model Jets	462
4.3.1.2	Clustering Fuzzy Jets: the EM Algorithm	466
4.3.1.3	Comparisons with Sequential Recombination and Jet Tagging	468
4.3.1.3.1	Details of the Simulation	468
4.3.1.3.2	Kinematic Properties of Fuzzy Jets	470
4.3.1.3.3	New Information from Fuzzy Jets	472
4.3.1.3.4	Fuzzy Jets for Tagging	474
4.3.1.4	Underlying Event and Pileup	475
4.3.1.5	Conclusions	478
4.3.2	Jet Images	479
4.3.2.1	Pre-processing and the Symmetries of Spacetime	479
4.3.2.2	Network Architecture	485
4.3.2.3	Performance and Visualization	487
4.3.2.4	Outlook and Conclusions	493
4.3.3	Conclusions and Future Outlook	494

III The Search for a Light Stop Quark

1 Introduction and Motivation

1.1	Superspace as an Extension of Spacetime	500
1.2	Broken Supersymmetry	508
1.3	The Hierarchy Problem and Weak-Scale SUSY	509
1.4	The Minimal Supersymmetric Standard Model	517
1.5	Related Models	530

2 Analysis Strategy

2.1	Phenomenology	533
2.2	The Control Region Method	549
2.2.1	Single Bin Signal Regions	551
2.2.2	Multibin (Shape Fit) Signal Region	552
2.2.3	Drawbacks of the CR Method and Alternatives	555
2.3	Dataset and Monte Carlo Samples	560

3 Object and Variable Definitions

3.1	Object Selection	564
3.2	Discriminating Variables	572
3.2.1	Transverse Mass Variables	573
3.2.1.1	Transverse Mass	575
3.2.1.2	Multiple Missing Particles: m_{T2}	585
3.2.1.2.1	Perpendicular m_{T2}	588
3.2.1.2.2	Numerical Methods	592
3.2.1.2.3	Tailoring m_{T2} for the stop search	598
3.2.1.2.4	Comparisons between transverse mass variables	603
3.2.1.2.5	Additional Considerations	607
3.2.2	Significance Variables	611
3.2.2.1	Constructing Significance Variables	614
3.2.2.2	Examples of Optimal Significance Variables	615

	3.2.2.3	Choosing the Separation Scale M	617
	3.2.2.4	Empirical Examples	621
	3.2.2.5	Examples with Full Detector Simulation	624
	3.2.2.5.1	Approximate m_T Significance	624
	3.2.2.5.2	$H_{T,\text{sig}}^{\text{miss}}$ Significance	627
	3.2.3	Tau veto	632
	3.2.4	Hadronic Top Mass Reconstruction	638
4		Signal Regions	
	4.1	Preselection	643
	4.2	Optimization Procedure	645
	4.3	Single Bin Regions	648
	4.4	Compressed Signal Region	664
5		Background Estimation	
	5.1	Top Quark Pair Production	673
	5.1.1	Modeling Dilepton Events	676
	5.1.1.1	Tau Validation Region	677
	5.1.1.2	Dilepton Validation Region	680
	5.2	W +jets	683
	5.2.1	Extrapolating in b -jet Multiplicity	685
	5.2.2	Modeling the m_T Tail	685
	5.3	Single Top	689
	5.3.1	A Data-driven Approach	691
	5.4	Top Quark Pair Production with a Z Boson	696
	5.4.1	Estimation from Simulation	697
	5.4.2	A Data-driven Method with Photons	698
	5.4.2.1	Motivation and Overview	698
	5.4.2.2	Relating $t\bar{t} + \gamma$ to $t\bar{t} + Z$ at Leading Order	699
	5.4.2.3	Simulation and Matrix Element Photons	705
	5.4.2.4	Higher Order QCD Corrections	710
	5.4.2.5	Event Selection	714

5.5	Dibosons	718
5.6	Overview	719
6	Systematic Uncertainties	
6.1	Experimental Systematic Uncertainties	727
6.1.1	Jet Energy Scale	730
6.1.2	Jet Energy Resolution	734
6.1.3	b-tagging	736
6.1.4	Missing Transverse Momentum	738
6.1.5	Other	739
6.2	Theoretical Modeling Uncertainties	740
6.2.1	Top Quark Pair Production	742
6.2.2	Single Top	744
6.2.3	Top Quark Pair Production with a Z Boson	749
6.2.3.1	Simulation-based	749
6.2.3.2	Data-driven	751
6.2.4	W+jets	755
6.2.5	Dibosons	756
6.3	Summary	758
7	Search Results	
7.1	Statistical Methods	761
7.2	CR-only Fit	767
7.3	Validation Regions	770
7.4	Exclusion Limits	771
7.4.1	Early $\sqrt{s} = 8$ TeV Results	772
7.4.2	Full $\sqrt{s} = 8$ TeV Results	774
7.4.2.1	Recasting Stop Limits	779
7.4.2.1.1	Equivalent Models	782
7.4.2.1.2	Limits on Compressed Gluino Mediated Stop Production	783
7.4.2.1.3	Derived Limits	788

7.4.2.1.4	Transitioning to $\sqrt{s} = 13$ TeV	790
7.4.3	Early $\sqrt{s} = 13$ TeV Results	791
7.5	The LHC Run I SUSY Epilogue	792
7.5.1	Constructing the Dataset	792
7.5.2	Statistical Analysis	795
7.5.3	Results and Discussion	797
8	Conclusions and Future Outlook	
A	Radiation Damage	
B	Boson Polarizations	
B.1	Angular Distributions in W and Z Boson Decays	816
B.2	Polarization of W Bosons	818
C	Additional Statistical Considerations	
C.1	Uncertainty Ellipses	822
C.2	Wrapped Gaussian	824
C.3	The EM algorithm	825
C.4	A Leading Order Description of Fuzzy Jet σ	828
C.5	Controlling Jet Multiplicity with p_T	829
C.6	Computation of Significance Variables	831
C.7	The Non-closure of Numerical Inversion	832
C.8	Gaussian Invariance Lemma	835
	Bibliography	

List of Tables

1	Natural units. There are multiple ways to define the electric charge. In these natural units, $e = \sqrt{4\pi\alpha}$ so that one unit represents ~ 0.3 of an elementary charge.	1
1.1	The particle content of the SM prior to electroweak symmetry breaking. The values under $U(1)$ are the Abelian charge (the actual representation is one-dimensional) whereas the entries under $SU(2)$ and $SU(3)$ are the representation of the field in the first column. For example a bold eight denotes the octet (adjoint) representation of $SU(3)$	10
1.2	The particle content of the SM after electroweak symmetry breaking. All couplings are given to one significant figure at the scale m_Z . Parameter values are from Ref. [27].	12
1	The various quantum properties of jets studied in Part II along with the corresponding chapter. Even though some of the quantum properties are due to forces other than the strong force, their realization in jets is due to the quantum evolution from partons to hadrons via QCD. .	55
1.1	For each algorithm, the probability that the b - and $\bar{b}$ -jet assignment is correct (reliability R) and the fraction of events to which the algorithm can be applied (efficiency E). Note that the value of R in the second row is less than 50% of the time, due to cases in which neither jet has an associated lepton ($Q_L = 0$).	75

1.2	The single-jet trigger menu used to collect dijet events with the 2012 dataset. The first column is the level-three (Event Filter) jet p_T threshold and the second column is the offline leading-jet p_T range corresponding to the given trigger. The luminosity collected with each trigger is in the last column. The total 2012 dataset was 20.3 fb^{-1} ; the highest- p_T trigger is not prescaled.	80
1.3	The data and MC signal and background yields after all selections for the 5.8 fb^{-1} sample, shown separately for μ^+ and μ^- final states. The MC uncertainties are purely statistical and included solely for the purposes of illustrating the sample composition.	105
1.4	A summary of all the systematic uncertainties and their impact on the mean jet charge for $\kappa = 0.5$ and the more forward jet. The correction factors are the fake and inefficiency corrections applied before/after the response matrix. The Other Tracking category includes uncertainty on the track reconstruction efficiency, track momentum resolution, charge misidentification, and fake track rate. All numbers are given in percent. As a result, the uncertainty appears artificially large in the first p_T bin where the average jet charge is nearly zero (small compared to the resolution).	135
1.5	A summary of all the systematic uncertainties and their impact on the jet charge distribution's standard deviation for $\kappa = 0.5$ and the more forward jet. The correction factors are the fake and inefficiency corrections applied before/after the response matrix. The Other Tracking category includes uncertainty on the track reconstruction efficiency, track momentum resolution, charge misidentification, and fake track rate. All numbers are given in percent.	136
1.6	An overview of the method and relative size of the various tracking systematic uncertainties.	140
1.7	A summary of the track reconstruction efficiency uncertainties.	143
1.8	A summary of the momentum (scale and) resolution uncertainties, taken from Ref. [131].	166

2.1	Monte Carlo samples used in this analysis. The abbreviations ME, PS, PDF, MPI, LO and NLO respectively stand for matrix element, parton shower, parton distribution function, multiple parton interactions, leading order and next-to-leading order in QCD. Tune refers to the used set of tunable MC parameters. Those samples marked with a † are used as alternative $t\bar{t}$ samples to evaluate uncertainties due to the modeling of $t\bar{t}$ events. The 4-flavor scheme (4f) is used for CT10 only for the t-channel single top production.	207
2.2	Estimated composition of the selected event sample. The uncertainties are the sum in quadrature of the statistical uncertainties and either the uncertainties of the normalisation method (for the data driven W+jet and multi-jet estimates) or the uncertainties of the cross-section estimates.	217
2.3	A summary of the tracking systematic uncertainty and their impact on the charged-particle pull angle measurement. Values are given in percent. For comparison, the data statistical uncertainty is the last line.	265
2.4	A summary of the cluster systematic uncertainty and their impact on the all-particle pull angle measurement. See Sec. 2.5.2.2 for an explanation of the two options for the cluster energy scale uncertainty. Values are given in percent. For comparison, the data statistical uncertainty is the last line.	266
2.5	A summary of the α and β coefficients in Eq. 2.10 used to bound the differences between data and simulation.	268
2.6	The impact on the unfolded jet pull angle distribution from re-weighting the jet p_T spectrum to match the data.	285
2.7	Uncertainties in each bin of the all-particle pull angle. The “Other” category includes uncertainties due to the non- $t\bar{t}$ backgrounds.	288
3.1	The coefficients of various parts of the N ³ LO prediction (Ref. [330, 331]) for the quark and gluon constituent multiplicity distributions from Eq. 3.13 and 3.14.	302

3.2	Monte Carlo samples used for measuring and studying the charged particle multiplicity inside jets.	307
3.3	A summary of all the systematic uncertainties and their impact on the n_{track} mean for $p_T^{\text{track}} > 0.5$ GeV and the more forward jet. Uncertainties are given in percent. The last row is the measured average charged particle multiplicity. A value of 0.0 is quoted if the uncertainty is below 0.05%.	325
3.4	A summary of the systematic uncertainties on the average charged multiplicity extraction for gluons. (*) NNPDF versus CT10, used only as a cross-check. (**) Using a cone size of $\Delta R < 0.2$ instead of the nominal 0.4 in the q/g identification. Used only as a cross-check. (***) Matching the jets with the outgoing partons in the ME to do the q/g ID. Used only as a cross-check. The uncertainties are in units of n_{charged}	331
3.5	A summary of the systematic uncertainties on the average charged multiplicity extraction for quarks. (*) NNPDF versus CT10, used only as a cross-check. (**) Using a cone size of $\Delta R < 0.2$ instead of the nominal 0.4 in the q/g identification. Used only as a cross-check. (***) Matching the jets with the outgoing partons in the ME to do the q/g ID. Used only as a cross-check. The uncertainties are in units of n_{charged}	332
4.1	A summary of the measured relative jet mass scale and jet mass resolution using both the subtraction and the forward folding methods. Uncertainties are given as a fraction of the nominal. The jet energy scale, ISR/FSR, and MC Normalization background uncertainties are treated as asymmetric but the maximum of the two variations are reported in this table.	374
4.2	Summary of the systematic uncertainties for the relative jet mass or energy scales ($s_{\text{data}}^{\text{MC}}$) and resolutions ($r_{\text{data}}^{\text{MC}}$). The first column states which observable is used to extract the relative jet mass (first four rows) or jet energy (rows 5-10) scale and resolutions.	382

1.1	A summary of the MSSM field content in terms of electroweak eigenstates before symmetry breaking. Even though the SUSY partners of the left- and right-handed fermions are scalars, they still carry the L or R subscript to emphasize their relationship to the SM particles. The superpartners of the fermions are called <i>sfermions</i> (squarks and sleptons) and the superpartners of the bosons are called <i>bosinos</i> (bino, wino, gluino, and higgsino). The right-handed fields are specified in terms of the charge conjugate of left-handed fields because chiral superfields only contain left-handed fermions.	520
2.1	Bias in the measured top quark mass and $t\bar{t}$ cross-section due to the presence of a light stop ($m_t = 170$ GeV) that decays via the three-body process. All masses are in GeV and all cross-sections are in pb. The measured top quark mass is biased low from the true mass which results in the true cross-section at the measured top mass, true $\sigma_{t\bar{t}}(m_t^{\text{measured}})$ to be higher than the true cross-section at the true mass, true $\sigma_{t\bar{t}}(m_t^{\text{true}})$. The former quantity is what would be predicted under the SM-only hypothesis in the presence of the 170 GeV stop. The measured $\sigma_{t\bar{t}}$ is the sum of true $\sigma_{t\bar{t}}(m_t^{\text{true}})$ and true $\sigma_{t\bar{t}}$, corrected for the lower acceptance for the three-body decay.	545
2.2	Overview of the nominal simulated samples. The last row indicates the order and reference for the inclusive cross-section to which the (lower order) MC simulations are normalized. The blue indicates a setup at $\sqrt{s} = 8$ that differs from the one used at $\sqrt{s} = 13$ TeV. More information about the $t\bar{t} + \gamma$ generation can be found in Sec. 5.4.2. .	562
3.1	An overview of the objects definitions used for the stop search. The color blue indicates the criteria at $\sqrt{s} = 13$ TeV whereas black is for $\sqrt{s} = 8$ TeV. See the text for details.	569

3.2	A summary of the procedure to resolve ambiguous object labels at $\sqrt{s} = 8$ TeV. The first two rows list the overlapping objects: electrons (e), muons (μ), electron or muon (l), jets (j), photons (γ), and hadronically decaying τ leptons (τ). The third and fourth rows give the overlap condition and the last row lists which label is given to the ambiguous object. The procedure is applied from left to right. (*) Jets are removed only in the $t\bar{t} + \gamma$ validation region (see Sec. 5.4.2). . . .	571
3.3	A summary of the procedure to resolve ambiguous object labels at $\sqrt{s} = 13$ TeV. The first two rows list the overlapping objects: electrons (e), muons (μ), electron or muon (l), jets (j), photons (γ), and hadronically decaying τ leptons (τ). The third and fourth rows give the overlap condition and the last row lists which label is given to the ambiguous object. The procedure is applied from left to right.	571
3.4	The separation power and correlation with E_T^{miss} and m_T for a variety of variables described in the text. The background is $t\bar{t}$ and the signal is a stop model with $(m_{\text{stop}}, m_{\text{LSP}}) = (500, 1)$. The variant P means that p_T is used to pick the b-jets while B means that the b-tagging weight is used. The symbol $\perp$ denotes the perpendicular variant of the variable in the first column. A boost correction is indicated by the letters BC. The symbols ΔR and min represent the scheme for addressing the matching ambiguity between the lepton and the b-jet by using the closest pair or considering the minimum of both possible pairings.	606
3.5	A summary of the separation power and correlation with E_T^{miss} for $H_{T,\text{sig}}^{\text{miss}}$ and $E_T^{\text{miss}}/\sqrt{H_T}$ based on Fig. 3.46 and Fig. 3.47.	631
4.1	Benchmark signal models used for optimizing the single bin regions. The regions SR1-3 were optimized using the entire grid and so the chosen models are representative of the regions of parameter space that the three regions target. SR13 was optimized using a GMS model, but the equivalent stop model is the one in the table above.	649

4.2	A summary of the six single bin signal region event selections. Dashed lines indicate that there is no requirement on the given variable. Even though the same symbol might be used for all six regions, some of the variables have a different meaning across columns. For example, the jets used in the first five columns have the LCW calibration while the jets in the last column are at the EM scale but with the GS calibration applied. Furthermore, the choice of the τ candidate differs between the first five columns and the last column. The m_{top,X^2} variable is an explicit tri-jet mass reconstruction in the first five columns (using jet resolution information in columns 4 and 5) and the large radius jet mass in the last column (see Sec. 3.2.4). (*) There is also an upper threshold of 250 GeV for this loose selection.	650
4.3	Yields for the SR13 defined in Table 4.2 but with various b-tagging working points. The last row is an approximate significance using the RooStat [595] NumberCountingUtils routine BinomialExpZ. A p-value is computed with a likelihood given by the product of a Poisson term for the statistical uncertainty and a Poisson term for the systematic uncertainty treated as a statistical uncertainty from an auxiliary measurement, i.e. a Poisson with mean $\tau = 1/(30\%)^2$ so that the fractional uncertainty of the auxiliary measurement is $1/\sqrt{\tau} = 30\%$. The conversion from p-value to σ is given by $\sigma = \Phi^{-1}(1 - p)$, for Φ the Gaussian cumulative distribution function.	661
4.4	The definition of tN_diag. In addition to the variables shown above, the signal region is defined by a four-jet requiremenet with $p_T > 80, 60, 40, 25$ GeV, $m_{\text{top}}^{\text{had}} \in [130, 205]$ GeV, $E_T^{\text{miss}}/\sqrt{H_T} > 5$ GeV ^{1/2} , and $\Delta\phi(\text{jet}_i, \vec{p}_T^{\text{miss}}) > 0.8$ for $i = 1$ and 2.	667

4.5	CL _s values (see Sec. 7.1) computed with the shape-fit (multibin) and single bin setups described in the text. These values are approximately p-values for a hypothesis test (smaller values are better). The E _T ^{miss} and m _T values are the thresholds for the tightest region of the shape fit and define the single bin region. The * denotes the CL _s value for the shape fit without any systematic uncertainties. As expected, systematic uncertainties have a big impact on the sensitivity. In two cases (marked ‘Fit Failed’) the multibin fit did not converge.	668
4.6	A summary of the multibin shape fit region, tN_diag. This region was released with a preliminary selection in Ref. [413] with a small change for the final result in Ref. [414]. Dashed lines indicate that there is no requirement on the given variable.	669
5.1	Estimates for the yield of Z+jets events relative to the W → lv yield. The number of additional jets only contributes to the last column when it differs from four, which is the necessary number of extra jets already needed by W+jets events to pass the event selection. If a tau decays leptonically, it is considered to be reconstructed as a jet, reducing the number of extra jets required. The rate of fake or non-prompt leptons ε _f ≲ 1% [277]. The probability for events with no real E _T ^{miss} to be misreconstructed as events with large E _T ^{miss} , ρ _f , is negligible because the resolution scales as √Σ E _T GeV ^{1/2} ∼ 20 GeV. Events with tau decays naturally have E _T ^{miss} , but this needs to be severally mis-measured (with rate ρ _m) to pass at least the m _T requirements.	672

5.2	The definition of the $t\bar{t}$ control regions for each signal region presented in Chapter 4. Only the requirements that differ from the corresponding signal region are indicated in the table, with a ‘–’ if there is no change between the signal and control region. The lower two rows show the total background yield and the fraction of $t\bar{t}$ events in both the signal and control region using the CR-only fit, described in Sec. 7.1. The upper $\mathbf{am}_{T2}$ requirement in TCR13 is to ensure orthogonality from STCR13, described in Sec. 5.3.1.	675
5.3	An event selection requiring a hadronic τ candidate to study the modeling of $1L1\tau$ events.	677
5.4	An event selection requiring two reconstructed signal leptons. The variable q_l denotes the charge of lepton l	680
5.5	The definition of the W control regions for each signal region presented in Chapter 4. Only the requirements that differ from the corresponding signal region are indicated in the table, with a ‘–’ if there is no change between the signal and control region. All changes highlighted in red are different from the definition of the corresponding $t\bar{t}$ control region. The lower two rows show the total background yield and the fraction of W +jets events in both the signal and control region using the CR-only fit, described in Sec. 7.1.	684
5.6	The definition of the single top control region for SR13 compared with the definitions of SR13 and the corresponding $t\bar{t}$ control region. Only the requirements that differ from the corresponding signal region are indicated in the table, with a ‘–’ if there is no requirement on the given variable. The lower panel indicates pre-fit predictions from the simulation compared with the data.	694
5.7	The composition of photon origins in the $t\bar{t} + \gamma$ and $t\bar{t}$ samples at $\sqrt{s} = 8$ TeV. The order of photon precedence is described in the first column. Except for the first row, all other rows are labeled $t\bar{t}$ and not $t\bar{t} + \gamma$. The last column quantifies the fraction of the various categories in the $t\bar{t} + \gamma$ validation region described in Sec. 5.4.2.5	709

5.8	NLO QCD corrections to the $t\bar{t}+\gamma$ and $t\bar{t}+Z$ cross-sections. Ref. [544] is for 14 TeV and has a second k-factor given for a second selection that has a harder jet requirement. Since the calculation is fixed-order, this artificially increases the k-factor to the 1.9 value that was used for the 8 TeV analysis. The k-factor in Ref. [543] is for 8 TeV. Both Ref. [619] and Ref. [620] show the differential (in Z p_T) k-factor up to 200 GeV, which appears to be relatively flat in that range. The value [621](1) is for the zero-width approximation while [621](2) is for a narrow-width approximation and uses the MSTW08 PDF set for both LO and NLO. The third value [621](3) is for a narrow-width approximation and mixes CTEQ6L1 at LO with CT10 at NLO. Electroweak corrections have also been reported in Ref. [622].	711
5.9	LO and NLO cross-sections for $t\bar{t}\gamma$ and $t\bar{t}Z$ as a function of the boson p_T threshold computed with SHERPA+OPENLOOPS by Stefan Hoche. All cross-sections are in pb. The numbers in parentheses are the statistical uncertainties. The PDF is CT14. A scale of $H_T = \sum p_T$ over all final state objects is used (the difference between the scalar sum of p_T and m_T was found to be negligible in this range).	712
5.10	LO and NLO cross-sections for $t\bar{t}\gamma$ and $t\bar{t}Z$ as a function of the boson p_T threshold computed with MG5_AMC. A custom fortran filter is used to isolate $t\bar{t}+Z$ events with a fixed Z boson threshold. All cross-sections are in pb. The uncertainties on the NLO cross section are from variations of the factorization and renormalization scale. The numbers in parentheses are the statistical uncertainties. The PDF is NNPDF2.3 LO for the LO calculations and NNPDF2.3NLO for the NLO calculations. The for both LO and NLO scale is half the scalar sum of the transverse mass of all out-going partons (default for NLO and scale option 3 for LO [623]). The impact of adding a PS was found to be small ($\lesssim 10\%$).	712

5.11	The requirements for the $t\bar{t} + \gamma$ VR ($\sqrt{s} = 8$ TeV) and CR ($\sqrt{s} = 13$ TeV). In both regions, exactly one signal lepton is required with no other baseline leptons. Furthermore the event selections require at least one b -tagged jet. The tilde variables include the photon in the $\vec{p}_T^{\text{miss}}$ as described in the text.	714
5.12	Expected and observed event yields in the $t\bar{t} + \gamma$ validation/control regions. All MC numbers are normalized to 20.3 fb^{-1} for the VR at $\sqrt{s} = 8$ TeV and 3.32 fb^{-1} for the CRs at $\sqrt{s} = 13$ TeV. The $t\bar{t}$ sample at $\sqrt{s} = 8$ TeV is reweighted according to the standard procedure described in Section 5.1. The displayed uncertainties are due to limited statistics.	715

6.1	Performance metrics for the various reconstructed objects and their systematic uncertainties at $p_T = 25$ GeV (upper number) and $p_T = 100$ GeV (lower number). The symbol σ denotes the energy resolution (width / mean), μ the energy scale, and ϵ the reconstruction efficiency. Many of the object reconstruction algorithms changed between Run 1 and Run 2; the selection in this table mixes the defaults between the two runs to give an idea of the overall performance. The jet reconstruction efficiency is based on the <i>loose</i> quality criteria [546] (and refs. therein) which is over 99.9% efficient and a jet vertex tagger (JVT) threshold of 0.59, which corresponds to a pileup jet efficiency of just over 1% [124]. The b-tagging efficiency is determined using the 70% working point of the MV1 algorithm [75]. Combined (CB) + segment tagged (ST) muons [131] are considered without any further quality criteria and the scale and resolution (uncertainties) are on $m_{\mu\mu}$ from various resonance decays. Electrons are reconstructed with the <i>loose</i> criteria. Unconverted <i>tight</i> photons [553] are used for photon reconstruction properties. All numbers refer to central objects only (except for E_T^{miss} , which uses objects across η). Medium one-prong tau reconstruction at $\sqrt{s} = 13$ TeV is used for illustration [550]. The last line indicates the properties of the component of the track-based soft term parallel to p_T^{hard} [138] (and refs. therein).	728
6.2	A decomposition of theory modeling uncertainties into several categories.	741
6.3	A summary of the theoretical modeling uncertainties for $t\bar{t}$ at $\sqrt{s} = 8$ TeV and $\sqrt{s} = 13$ TeV. The * indicates that this uncertainty was studied but not applied because of worse agreement between the alternative generators and data in an inclusive event selection. P stands for POWHEG-BOX and M stands for MG5_aMC.	743
6.4	A summary of the theoretical modeling uncertainties for single top at $\sqrt{s} = 8$ TeV and $\sqrt{s} = 13$ TeV. The * indicates that this uncertainty was studied but not applied (see the text). P stands for POWHEG-BOX.	744

6.5	A summary of the theoretical modeling uncertainties for $t\bar{t} + Z$ at $\sqrt{s} = 8$ TeV and $\sqrt{s} = 13$ TeV.	749
6.6	A summary of the theoretical modeling uncertainties for W +jets at $\sqrt{s} = 8$ TeV and $\sqrt{s} = 13$ TeV.	756
6.7	Theoretical modeling uncertainties for dibosons.	756
6.8	A summary of the uncertainty in the total signal region yield after the background-only fit from the control region method (see Sec. 7.1). If an uncertainty is less than 1%, it is marked with ‘-’ while uncertainties that are not applicable are labeled N/A. PDF uncertainties are included in the HS uncertainty. tN_{xy} is the $(x + 1)^{\text{th}}$ E_T^{miss} bin and $(y + 2)^{\text{th}}$ m_T bin of the shape fit. When there is more than one JES nuisance parameter, there are at least three, but only the two biggest ones are shown here. Due to correlations in the uncertainties after the fit, the total systematic uncertainty is not the sum in quadrature of the individual uncertainties.	759
7.1	Example models that result in $t\bar{t} + E_T^{\text{miss}}$ categorized by their spin, color charge, electric charge (has to be 2/3 if the missing particles are neutral), and the relative cross-section. The acronym GMS stands for gluino mediated stop. The particle T exists in a variety of models and could decay via $T \rightarrow tA_0$ for a new weakly interacting particle A_0 (see Ref. [500] and the references therein), or it could be a vector-like quark decaying via $T' \rightarrow tZ(\rightarrow \nu\bar{\nu})$, which also has the $t\bar{t} + E_T^{\text{miss}}$ topology. A ‘-’ indicates that there is no standard model with these properties.	779
7.2	A set of direct stop models that are equivalent to GMS models. The third and fourth columns are from Ref. [503].	784

7.3	The expected increase in yields for the direct stop search and the re-interpreted gluino search from $\sqrt{s} = 8$ to 13 TeV. The first column is the stop mass in GeV, the second column is the stop cross section at $\sqrt{s} = 8$ TeV in pb from Ref. [503]. The third column is in pb and also uses Ref. [503] to solve $\sigma(m_{\tilde{t}}) = \sigma(m_{\tilde{g}})$. The fourth and fifth columns give the cross sections for stop and gluino production at $\sqrt{s} = 13$ TeV from Ref. [502]. The last two columns give the ratio of the increase in yields for direct stop and GMS production, respectively.	790
7.4	An overview of the signal regions used in the meta-analysis from 8 TeV CMS searches.	793
7.5	An overview of the signal regions used in the meta-analysis from 8 TeV ATLAS searches. Note that OS = opposite sign. The stop search results are part of the ‘One Lepton’ category.	794
7.6	Results for statistical hypothesis tests on combined ATLAS and CMS excesses, under the assumptions of Gaussian and lognormal error distributions.	799
7.7	Results for statistical hypothesis tests on combined ATLAS and CMS deficits, under the assumptions of Gaussian and lognormal error distributions.	799
A.1	Physical constants describing the mobility of charge carriers in silicon. The first three rows are reformatted from Ref. [685] and the Hall scale factor is from Ref. [686].	804

List of Figures

1	A comparison of η and y for a massive particle with $\phi = 0$	3
1.1	A comparison of all dimensionless SM parameters. The nine Yukawa couplings are on the left in red and are organized by generation. The three gauge couplings g_i are $g_1 = e/\cos(\theta_W)$, $g_2 = e/\sin(\theta_W)$, $g_3 = \sqrt{4\pi\alpha_s}$, for the weak mixing angle θ_W , e is the electron charge $e = \sqrt{4\pi\alpha_{EM}}$, and $\alpha_s \sim 0.1$ is the strong coupling constant. The magnitude of all elements of the CKM matrix are in blue on the right. Parameter values are from Ref. [27].	12
1.2	An illustration of the strength of the three forces through the lifetimes of various elementary and composite particles. The color of the labels is determined by the dominant decay mode. Values of the mass and lifetimes are from Ref. [27]. The lifetime is defined as the inverse of the decay width (for the strong force resonances, this is what is measured).	13
1.3	PDF sets from the NNPDF collaboration [32] at $Q^2 = (100 \text{ GeV})^2$ extracted using the HepData [33] interface. The left and right plots are the same aside from the scale on the horizontal axis.	15
1.4	Cross section measurements at $\sqrt{s} = 7, 8, 13 \text{ TeV}$ from the ATLAS experiment compared to theoretical measurements. From Ref. [39].	18
2.1	Diagram illustrating the length scales probes by various ‘microscopes.’	21
2.2	An overview of the CERN accelerator complex. The LEP experimental areas are indicated in red while the four LHC experiments are labeled in blue.	23

2.3	Diagrams illustrating the trajectory of charged particles in a magnetic field. Both views are cross sections of the detector with the beam axis and B-field perpendicular to the page. Circles on the left plot indicate the locations of various tracking detector elements for the ATLAS detector, discussed in Sec. 2.3. A 2 T solenoid magnet is used to determine R	27
2.4	A schematic diagram used to derive the momentum resolution of a charged particle track measurement. The length L (<i>lever arm</i>) is the distance over which measurement of the trajectory can be made. The three dots indicate discrete samplings (measurements) of the trajectory.	28
2.5	The average transverse distance traveled by various particles labeled on the horizontal axis as a function of their p_T	31
2.6	An overview of the measurement pattern for various particle types in the ATLAS detector. A dashed line means that the given particle leaves no trace in the sub-detector. A solid horizontal line indicates insignificant energy loss (ionization and small levels of multiple scattering for electrons).	31
2.7	The ATLAS detector and its subsystems (to scale). Image from Ref. [78].	32
2.8	Left: A schematic view of the subsystems of the ATLAS inner detector [80]. The IBL was added between Runs 1 and 2. The red line indicates the trajectory of a hypothetical particle with $p_T = 10$ GeV at $\eta = 0.3$. Right: an enlarged view of the pixel detector prior to the insertion of the IBL [81].	34
2.9	Left: A picture of the IBL just after being inserted into the detector. The yellow cables are wrapped in order to pass through to the far side. These cables host data transmission, high/low voltage, and many other services. Right: The difference in the measured resistance in the data pins before and after a practice wrapping.	34
2.10	A schematic diagram of the ATLAS calorimeters. Image from Ref. [82].	35

2.11	A candidate $H \rightarrow WW^* \rightarrow 4\mu$ event. The long purple lines indicate reconstructed muon tracks and each muon tracking chamber contributing to those tracks is highlighted in green (MDT) or purple (CSC). The invariant mass of the four muons is 145.8 GeV (inconsistent with the now known $m_H \sim 125$ GeV).	36
4.1	An event display in the early Run 2 data showing charged particle tracks reconstructed from the ID. The fitted track trajectory is shown as a continuous colored line and the hits in the three sub-detectors are shown as points. From Ref. [113].	45
4.2	A simulated $Z' \rightarrow t\bar{t}$ event clustered with anti- k_t (left), C/A (middle) and k_t (right). Particles are colored according to the jet they are clustered in (highest p_T jet is black, then red, then green). The particle size is proportional to $\log(10 \times p_T/\text{GeV})$. The bottom panel shows the clustering history. Each proto jet merger is recorded with the size of each pre-merger proto jet proportional to $\log(10 \times p_T/\text{GeV})$, where p_T is the momentum of the merged proto jet. Mergers are colored from dark (earlier) to light (later). Note that ϕ is 2π periodic so the particles at $-\pi$ are geometrically close to the particles at $+\pi$	48
4.3	Jet clustering with two high p_T particles and a grid of ghost particles. Particles clustered in the higher p_T jet are colored red and those in the lower p_T jet are colored blue. All other ghosts are gray. The ghosts marker size is arbitrary and the radius of the two high p_T particle markers $\propto p_T$. When the two particles are far away (left), the jet areas are πR^2 , while if they are within $ \Delta\vec{\rho} < 2R$, then one is circular and the other is a crescent. If $ \Delta\vec{\rho} < R$ (right), then the jets merge.	49
1.1	Left: the process $pp \rightarrow u\bar{u}$ is simulated with PYTHIA 8 once; Right: the same hard-scatter process with fragmentation re-simulated 10,000 times. Each pixel intensity represents the charge-weighted sum of the energies of all particles produced within the ϕ and η covered by the pixel area.	57

1.2	Top left (right): The particle- (detector-)level jet charge distribution for various jet flavors in a sample of jets with $p_T > 500$ GeV for $\kappa = 0.3$. Bottom: the distribution of the jet-by-jet difference between the particle-level and detector-level jet charge distributions. The shaded region is used to fit a Gaussian function to extract the bulk response resolution, which is $\sigma \sim 0.5 e$, where e is the positron charge. See Sec. 1.2.1 for details about the simulation.	60
1.3	The distribution of the jet charge built from the leading n tracks ($Q_{J,n}$) for (left) $\kappa = 0.3$ and (right) $\kappa = 0.7$ or $\kappa = 1.0$ for (top) $500 \text{ GeV} < p_T < 600 \text{ GeV}$ and (bottom) $70 \text{ GeV} < p_T < 100 \text{ GeV}$. In the top (bottom) plots, the mean number of tracks is about 15 (7). See Sec. 1.2.1 for details about the simulation.	62
1.4	The joint distribution of the jet charge defined with different values of the weighting factor $\kappa \in \{0.3, 0.5, 0.7\}$. The left plots show the joint distributions of $\kappa = 0.3$ and $\kappa = 0.5$ while the right plots show the joint distributions of $\kappa = 0.3$ and $\kappa = 0.7$. Jets in the top plots have $50 \text{ GeV} < p_T < 100 \text{ GeV}$ while those in the bottom plots have $500 \text{ GeV} < p_T < 600 \text{ GeV}$. See Sec. 1.2.1 for details about the simulation.	63
1.5	A schematic diagram to illustrate the impact of Lorentz boosts on the ‘jet’ charge. A Z boson decays at rest at the origin in the lab frame and decays into electrons that are along the x direction (left). For a boost along the $-x$ direction, the e^+ has a higher p_T than the e^- and vice versa for a boost along $+x$	64
1.6	The decay of a simulated W boson event with PYTHIA 8. The length of the arrow is proportional to the energy of the decay product; red arrow denote negatively charged hadrons, blue arrows mark positively charged hadrons and neutral hadrons and photons are in gray. In the right plot, the W has received a large boost in the $+x$ direction.	66

1.7	Left: for fixed $\beta_W = 0.65$, the dependence of the ‘jet’ (= all W boson decay products) charge on κ for the event shown in Fig. 1.6. Right: for $\kappa = 1$, the dependence of the jet charge on the boost direction and magnitude for the same event as in the left plot.	66
1.8	A representative set of leading order QCD Feynman diagrams with an up quark in the initial state. Every diagram that has a down quark could be replaced with any other parton. There are also the additional diagrams related to the t -channel ones by crossing symmetry.	70
1.9	The predicted p_T dependence of the jet charge with input from the CT10 PDF set and assuming $\langle Q_g \rangle = 0, \langle Q_q \rangle = -\langle Q_{\bar{q}} \rangle, \langle Q_u \rangle = -2\langle Q_d \rangle$. For the red, green, and blue lines, the impact of a p_T -dependent fragmentation function is added on top of the PDF-dependence. This is the more forward of the two jets in dijet events (see Sec. 1.2) and therefore more likely to be initiated from a quark. The fragmentation functions (FF) do not depend on κ , but the energy-dependence of their κ -moments do depend on κ	72
1.10	A schematic of the $t\bar{t}$ decay topology in the semi-leptonic channel. The clouds depict the fragmentation process that ultimately lead to observable tracks (dotted lines) and hadronic jets (cones). The final state is characterized by two b -tagged jets, two hadronic jets from the W decay, an isolated lepton, and missing momentum from the undetectable neutrino. In current kinematic fits, only the momentum of the jets are used and not the charge properties of the tracks. . . .	74
1.11	For a given jet flavor, (a) shows the fraction f of jets with that flavor in events passing the particle-level event selection and (b) shows the p_T evolution of the flavor fractions weighted by charge-sign: $f_{\text{up}} + f_{\text{anti-down}} - f_{\text{anti-up}} - f_{\text{down}}$. The CT10 PDF set is combined with matrix elements from PYTHIA 8. The forward-central differences between the flavor fractions are largest at low p_T , but the highest quark-jet purity occurs at high jet p_T . The markers for the more forward and central jets are distinguished by their blue and red colors, respectively. . . .	77

1.12	Trigger Efficiencies for the various single jet triggers used in this analysis. The numbers in the legend show the point at which the trigger is fully efficient.	79
1.13	The average number of interactions per bunch crossing in data and in Pythia (after re-weighting) after the event selection described in Sec. 1.2.2.	81
1.14	The distribution of the number of tracks associated with a jet in two example jet p_T ranges.	83
1.15	The distribution of the distance ΔR between the selected detector-level jet and the selected particle-level jet using a momentum scheme (left) and a rapidity scheme (right).	85
1.16	The impact of adding a particle-level jet p_T threshold of 500 MeV. .	85
1.17	Reconstructed jet p_T spectrum for the more forward jet (left) and the more central jet (right). Note that the pre-scales are applied to the data to arrive at a smooth and steeply falling distribution of the jet p_T . .	88
1.18	The ratio of the more forward to the more central jet p_T distributions. .	88
1.19	The differences in η (left) and ϕ (right) between the more forward and more central jet.	89
1.20	The η of the more forward (left) and more central (right) jet.	89
1.21	The distance in ΔR to the nearest jet with $p_T > 25$ GeV and the p_T spectrum of this jet (right).	90
1.22	The dependance of the average (left) and standard deviation (right) of the more forward jet charge distribution on the p_T of the nearest jet above 25 GeV. Events are only plotted if the ΔR to the nearest such jet is < 0.8 . Uncertainties are statistical only.	90
1.23	The top (bottom) row shows the sum of the charges of the two leading jets in dijet events for two bins of the leading jet p_T . Two values of the p_T weighting factor are shown: $\kappa = 1.0$ on the left and $\kappa = 0.3$ on the right. The lower panels show the ratios between data and MC. The gray band in the ratio includes jet p_T and track reconstruction efficiency uncertainties, described in Sec. 1.5.	91

1.24	The dependence of the dijet charge on the dijet mass for two different values of κ in data and MC for a dijet sample. The uncertainty band includes preliminary jet p_T and track isolation uncertainties.	92
1.25	Truth-level jet charge distribution for simulated dijet events for different parton flavours. Each distribution is normalised to unit area, except the gluon probability distribution function which is normalised to 0.3 for easier comparison. Distributions are shown in two different p_T bins for the leading jet and for two values of the weighting factor κ	95
1.26	Discrimination power of the jet charge to differentiate quark jets of opposite charge (left) and quark from gluon jets (right). Results are obtained from MC truth information in simulated dijet samples.	96
1.27	The parton flavour tagging performance of the jet charge, defined as in Fig. 1.26. The left plot shows the positive parton charge jet efficiency versus the negative parton charge rejection (inverse efficiency) for various κ values in a fixed p_T bin. The right plot fixes the positive charge efficiency at 0.5 and then shows the distribution of the negative charge rejection with κ	96
1.28	A summary of the information shown in Fig. 1.27 for many p_T bins. The horizontal axis is the jet p_T and the vertical axis is the maximum negative charge rejection for a fixed positive charge efficiency (ϵ_+) of 0.25, 0.5 or 0.75. For a fixed positive charge efficiency, the optimal κ value and the maximum negative charge rejection vary little. The p_T bins are chosen based on trigger thresholds.	97
1.29	Jet charge distributions for heavy-flavour quarks and inclusive quark types and their ratios (lower panels). The left plot shows the comparison for bottom type quarks and the right plots for charm type quarks. Results are obtained from MC truth information in simulated dijet samples.	97

1.30	Average (left) and RMS (right) jet charge response ($Q^{\text{reco}} - Q^{\text{truth}}$) as a function of the leading jet p_T for a sample of simulated dijet events and for two different bins in track multiplicity. Uncertainties are from the limited size of simulated samples.	98
1.31	Average (left) and RMS (right) jet charge response ($Q^{\text{reco}} - Q^{\text{truth}}$) as a function of the number of tracks used to compute the charge in different p_T bins of the leading jet in dijet simulated events. Uncertainties are from the limited size of simulated samples.	99
1.32	The RMS of the jet charge response ($Q^{\text{reco}} - Q^{\text{truth}}$) as a function of the average number of interactions per crossing for three values of κ and in two bins of p_T , as obtained in simulated dijet events. Uncertainties are from the limited size of simulated samples.	99
1.33	The dependance of the response resolution (measured using RMS) as a function of the jet p_T for three values of the track p_T threshold and two p_T weighting values $\kappa = 1$ (left) and $\kappa = 0.3$ (right). One can see that for both κ values, there is a big increase in the RMS at low jet p_T when moving to a track threshold of 5 GeV. For the lower value of $\kappa = 0.3$, the trend persists for high p_T since in this case, more weight is given to lower p_T tracks. The lower plots show the ratio of the 5 GeV track threshold distribution with the 500 MeV track threshold.	100
1.34	Variations on the definition of the jet charge, see the text for details. The left plot shows $\kappa = 0.3$ and the right shows $\kappa = 0.7$. The variants on the colors are for jet originating from a quark with a positive charge versus a negative charge. The small differences in modes between positive and negative quark initiated jets is likely due to statistical fluctuations from the limited MC sample size.	102
1.35	ROC curves for the distributions shown in Fig. 1.34. The left plot shows $\kappa = 0.3$ and the right shows $\kappa = 0.7$	102

1.36	Schematic representation of the object selection. At least four jets are required: two b -tagged jets B_1, B_2 and at least two non b -tagged jets, labelled J_1, J_2 . The jets J_1 and J_2 are those non b -tagged jets with invariant mass closest to the W boson mass. The charged lepton is used to trigger and a cut on the missing energy from the neutrino is used to purify the sample in $t\bar{t}$ events.	105
1.37	The distribution of the sum of the jet charges from the two daughter candidates in hadronic W boson decays in semileptonic $t\bar{t}$ events. The plot on the left is obtained with p_T weighting factor $\kappa = 1.0$ and the right plot with $\kappa = 0.3$. Events with a $\mu^\pm$ correspond to a hadronically-decaying $W^\mp$. The bottom panels show the data/MC ratios with the bands giving the systematic uncertainties described in the text.	107
1.38	The power to reject W^- as a function of the efficiency to tag W^+ as measured in semileptonic $t\bar{t}$ events. Each point on the curve corresponds to a cut value on the charge distribution shown in Fig. 1. The left plot shows results for inclusive jet multiplicity, and the right plot for several multiplicities and $\kappa = 1.0$	108
1.39	For $t\bar{t}$ simulated events, the mean (top) and RMS (bottom) distributions of the W daughter dijet charge response as a function of the total number of tracks used to compute the charge for two values of κ . The insets show the (arbitrarily normalised) distribution of the number of tracks. The error bars indicate the statistical uncertainties on the number of MC events.	109
1.40	For $t\bar{t}$ simulated events, the mean (top) and RMS (bottom) distributions of the W daughter dijet charge response as a function of the transverse momentum of the dijet system for two values of κ . The insets show the (arbitrarily normalised) distribution of the W candidate p_T . The error bars indicate the statistical uncertainties on the number of MC events.	110

1.41	The charge of a boosted hadronically-decaying W boson in simulated semileptonic $t\bar{t}$ events for $\kappa = 1.0$ (left) and $\kappa = 0.3$ (right). The hashed distributions are for the extension of the jet charge definition to large- R jets. The solid line histograms show the distribution of the large- R jet charge after trimming and the dashed lines show the sum of the charge of the two leading $R = 0.3$ k_t subjets. Note that the domain of the two jet charge distributions are not the same, but are plotted with the same x -axis range, in contrast to Fig. 1.37.	113
1.42	The inverse efficiency (rejection) of boosted W^- bosons as a function of the efficiency for boosted hadronic W^+ bosons for simulated semileptonic $t\bar{t}$ events and $\kappa = 1.0$ (left) and $\kappa = 0.3$ (right). The three curves indicate the performance of three different ways of measuring the jet charge for boosted W bosons.	114
1.43	A Gaussian fit to the particle-level jet charge distribution in three bins of jet p_T for the more forward of the two selected jets and $\kappa = 0.3$. The distribution is not exactly Gaussian, which is why the fitted curve is slightly shifted with respect to the histograms.	121
1.44	Various binning choices for the quantization of jet charge and jet p_T . The overflow fraction is the fraction of events with a jet charge outside the range set by the x -axis. Vertical lines indicate the values used for $\kappa = 0.3$ and $\kappa = 0.5$	123
1.45	Various binning choices for the quantization of jet charge and jet p_T for the jet charge distribution average (left) and standard deviation (right). For the RMS, there is a significant dependence up to about 10 iterations, which is why only a few of the small-iteration cases are shown.	123
1.46	The discretization bias for various binning choices for the quantization of jet charge and jet p_T for the jet charge distribution average (left) and standard deviation (right).	124

1.47	The distribution of the one-dimensional transformation of the jet charge and the jet p_T (see Sec. 1.4.1) for $\kappa = 0.5$ for the more forward jet. All distributions are scaled to have the same normalization. . . .	124
1.48	The jet p_T and jet charge distribution is transformed into a one-dimensional variable (see Sec. 1.4.1) and the inefficiency (left) and fake (right) factors are shown as a function of the bin number. The jet p_T increases from left to right and the jet p_T bin edges are marked by vertical dashed lines. A horizontal line at one indicates that no correction is applied. The above plots are for $\kappa = 0.5$ and the more forward jet, but the distributions for the other cases are qualitatively similar.	125
1.49	For events which pass the truth-based selection but fail the reconstructed event selection, this plot shows $ 1 - p_T^1/p_T^2 $ for the reconstructed jet p_T . The horizontal axis is the p_T bin number, from 1-10. In only the first p_T bin, there are events which pass the p_T asymmetry cut but fail the p_T cut. In every other bin, the event fails the truth selection due to the asymmetry cut. The histogram is normalized per p_T bin.	126
1.50	The conditional distribution of the detector-level one-dimensional transformation of the jet charge and jet p_T (see Sec. 1.4.1) in bins of the particle-level analogue (response matrix) for the $\kappa = 0.5$ for the more forward jet. The response matrix for the other variants is qualitatively similar.	127
1.51	The response matrix for $\kappa = 0.5$ for the more forward jet in various jet p_T bins.	128
1.52	Herwig unfolded with Pythia (left) and the statistical uncertainty (right) for various iteration multiplicity settings in the iterative Bayesian unfolding algorithm for the mean (top) and RMS (bottom).	129

1.53	The detector-level (data and simulation) and particle-level jet charge distribution's (a) average and (b) standard deviation as a function of the jet p_T for the more forward jet. The ratios in the bottom panel are constructed from the simulation, and show the prediction of detector-level PYTHIA over the data (top ratio), and detector-level PYTHIA over particle-level PYTHIA (bottom ratio). Bars on the data markers represent only the statistical uncertainties. For both (a) and (b), $\kappa = 0.5$	130
1.54	The distribution of the 1D transformation of the 2D binned jet charge and jet p_T distribution for the more forward jet for $\kappa = 0.5$. All distributions are scaled to have the same normalization.	132
1.55	A visualization of the systematic uncertainties for the jet charge mean (left) and standard deviation (right). As a result of displaying the uncertainty as a fraction of the mean or RMS, the uncertainty appears artificially large in the first p_T bin in the left plot where the average jet charge is nearly zero (small compared to the resolution).	134
1.56	The uncertainty on the fake factors by comparing the predictions from PYTHIA 8 and HERWIG for the more forward jet and $\kappa = 0.5$. The left plot shows the differences in all bins of the combined jet p_T and jet charge distributions. The middle and right plots show the uncertainty on the jet charge and jet charge distribution standard deviation, respectively.	138
1.57	The uncertainty on the inefficiency factors by comparing the predictions from PYTHIA 8 and HERWIG for the more forward jet and $\kappa = 0.5$. The left plot shows the differences in all bins of the combined jet p_T and jet charge distributions. The middle and right plots show the uncertainty on the jet charge and jet charge distribution standard deviation, respectively.	138

1.58	Left: The efficiency of the track $\chi^2/\text{NDF} \geq 3$ requirement given all other track requirements as a function of the jet p_T . The track p_T dependence for three p_T bins is shown in the right plot. Note that no statistical uncertainties are included in the right plot (the size of the bin-to-bin fluctuations indicates the size of these uncertainties). The low jet p_T efficiency is poor at high track p_T due to the large contribution from fakes.	144
1.59	The CES as a function of the particle-level jet p_T	146
1.60	Left: The fractional calo loss in data and simulation. Right: the fractional loss in simulation.	148
1.61	The detector- and particle-level CTER as a function of the jet p_T for Pythia, Herwig, and the data.	149
1.62	The uncertainty in the fractional loss determined from the difference in the data/MC of the fractional calo loss and the JES uncertainties.	150
1.63	The fractional calo loss for the nominal simulation, the data, and various parameterizations of the tracking inefficiency uncertainty.	152
1.64	Left: The ratio of the number of charged particles matched to reconstructed tracks divided by the total number of charged particles inside particle-level jets as a function of jet p_T and charged particle p_T . By construction, this efficiency is between 0 and 1. Right: A similar ratio, but the numerator is replaced with the numerator replaced with all reconstructed tracks. Due to resolution effects, the ‘efficiency’ in the right plot can exceed unity.	154
1.65	The distribution of tracks (left) and tracking efficiency (right) as a function of ΔR between the track and the jet axis for various p_T bins and inclusive in track $p_T > 500$ MeV.	155
1.66	The (un)corrected distribution of $\sum p_{T,i}/p_T^{\text{jet}}$ in two jet p_T bins (before computing CTER as the average) for low p_T jets (left) and high p_T jets (right).	155
1.67	The CTER for various particle-level and detector-level definitions (see text for details).	156

1.68	The uncertainty on the average jet charge for the more forward (left) and more central (right) jet with $\kappa = 0.3$ (top), $\kappa = 0.5$ (middle) and $\kappa = 0.7$ (bottom). The uncertainty from tracking in dense environments is assumed to be negligible below 400 GeV, where nuclear interactions are the dominant source of inefficiency.	157
1.69	The uncertainty on the jet charge distribution standard deviation for the more forward (left) and more central (right) jet with $\kappa = 0.3$ (top), $\kappa = 0.5$ (middle) and $\kappa = 0.7$ (bottom).	157
1.70	The fraction of the sum track $\mathbf{p}_T$ that is due to secondary tracks. . .	158
1.71	A schematic diagram of the jet core, $R \lesssim 0.02$. The pixel is drawn approximately to scale assuming it is from the third pixel layer in the central region.	159
1.72	A schematic diagram of the distribution of ζ , accounting for merging.	161
1.73	The joint distribution of ζ and the ΔR between tracks for jet $\mathbf{p}_T \sim 400$ GeV (left) and $\mathbf{p}_T \sim 1.2$ TeV (right).	161
1.74	The conditional distribution of ζ given jet $\mathbf{p}_T$ for data (right) and MC (left).	163
1.75	Left: the $\mathbf{p}_T$ dependence of $\partial_\zeta f(\zeta)$, for $f(\zeta)$ the probability distribution of ζ . Right: the impact on the ζ distribution from reducing the loss in simulation.	163
1.76	The distribution of ζ in bins of the jet $\mathbf{p}_T$ for the data as well as simulation. The top left bin corresponds to $140 \text{ GeV} < \mathbf{p}_T < 200 \text{ GeV}$ and each subsequent plot going from left to right, top to bottom, increases this range by 120 GeV.	164
1.77	For an ensemble of tracks of the same $\mathbf{p}_T$ and η , the standard deviation of the distribution of the track $\mathbf{p}_T$ divided by the track $\mathbf{p}_T$ by applying the smearing procedure described in the text for the track momentum resolution uncertainty.	166
1.78	The charge mis-id rate as a function of the jet $\mathbf{p}_T$ (left) and $ \eta $ (right).	168

1.79	The charge mis-id rate as a function of the track p_T . The charge mis-id rate increases rapidly for the lowest p_T bin due to fake tracks that happen to have a truth matching probability of $> 50\%$ but have a mis-id rate of $\sim 50\%$ because the reconstructed charge is random. .	168
1.80	The MC matching probability for tracks with a flip and without a flip (left) as well as the p_T difference for flipped and non-flipped tracks between the ‘true’ and reconstructed track p_T (right). The fact that the MC probability is in general lower for tracks with a flipped charge and the p_T is usually very different from the truth p_T indicates that many of these tracks are actually fake tracks.	169
1.81	Left: The joint distribution of the fake track p_T and the jet p_T . The dashed line indicates the $p_T^{\text{track}} = p_T^{\text{jet}}$ line. Right: the normalized distribution of the track p_T in five bins of jet p_T . The dashed horizontal lines indicate the jet p_T thresholds. The ratio only shows points for which the track p_T exceeds the jet p_T	170
1.82	The fraction of tracks inside a jet that are classified as fake in bins of jet p_T . The circle markers are for central $ \eta < 0.75$ jets while the squares are for all tracks within the tracker acceptance. The markers for the squares are offset by 30 GeV, but the actual fake rate is computed with the same p_T binning as for the circles. The markers indicate the median of the fake fraction distribution in a given jet p_T bin and the error bar is the inter-quartile range.	171
1.83	The p_T spectrum of fake tracks inclusively in η (left) and for $ \eta < 0.75$ (right).	172
1.84	The track multiplicity for jets in data and in Pythia for various p_T bins.	173
1.85	The conditional distribution of the jet charge distribution given n_{track} . This is for the more forward jet and $\kappa = 0.5$	174
1.86	The impact of n_{track} re-weighting on the jet charge distribution average (left) and standard deviation (right) as a function of jet p_T for the more forward jet and $\kappa = 0.5$	175

1.87	The uncertainty on the jet charge distribution average (left) and standard deviation (right) due to the n_{track} modeling. These plots are for the more forward jet and $\kappa = 0.5$	175
1.88	The track p_T spectrum in a particular jet p_T bin (left) and the conditional distribution of the jet charge given the track p_T (right); see the text for details. This is for the more forward jet and $\kappa = 0.5$. . .	176
1.89	The uncertainty due to the track p_T on the jet charge distribution average (left) and standard deviation (right).	176
1.90	The distribution of the weights w_i used to re-weight the MC distribution for the non-closure test (left) and the Data/MC ratio with the re-weighted truth distribution (labeled after) $\tilde{t}_i$ (right) for the more forward jet with $\kappa = 0.5$	179
1.91	The impact of the non-closure uncertainty on the jet charge mean (left) and the jet charge distribution standard deviation (right) for the more forward jet and $\kappa = 0.5$. Bins without a black and red point indicate that one of the two is smaller than 10^{-3}	179
1.92	The fractional difference on the jet charge mean (left) and standard deviation (right) when unfolding the PYTHIA simulation with a HERWIG++ response matrix and vice versa.	180
1.93	The measured average (standard deviation) of the jet charge distribution on top (bottom) in units of the positron charge as a function of the jet p_T for $\kappa = 0.3, 0.5$, and 0.7 for the more forward jet (left) and the more central jet (right). The crossed lines in the bars on the data indicate the systematic uncertainty and the full extent of the bars is the sum in quadrature of the statistical and systematic uncertainties. The solid continuous line is a smooth approximation to the PYTHIA prediction.	182

1.94	The average jet charge ($\kappa = 0.5$) in units of the positron charge for (a) the more forward jet and (b) the more central jet compared with theory predictions due to various PDF sets. The crossed lines in the bars on the data indicate the statistical uncertainty and the full extent of the bars is the sum in quadrature of the statistical and systematic uncertainties.	183
1.95	The standard deviation of the jet charge ($\kappa = 0.5$) distribution in units of the positron charge for (a) the more forward jet and (b) the more central jet compared with theory predictions due to various PDF sets. The crossed lines in the bars on the data indicate the statistical uncertainty and the full extent of the bars is the sum in quadrature of the statistical and systematic uncertainties.	184
1.96	The average of the jet charge distribution in units of the positron charge for (a) $\kappa = 0.3$, (b) 0.5, and (c) 0.7 comparing various QCD MC models and tunes for the more forward jet. The crossed lines in the bars on the data indicate the statistical uncertainty and the full extent of the bars is the sum in quadrature of the statistical and systematic uncertainties.	186
1.97	The standard deviation of the jet charge distribution in units of the positron charge for (a) $\kappa = 0.3$, (b) 0.5, and (c) 0.7 comparing various QCD MC models and tunes for the more forward jet. The crossed lines in the bars on the data indicate the statistical uncertainty and the full extent of the bars is the sum in quadrature of the statistical and systematic uncertainties.	187

1.98	The measured average of the jet charge distribution (top), and the standard deviation (bottom), in units of the positron charge as a function of the jet p_T for $\kappa = 0.3, 0.5$, and 0.7 for the more forward jet (left) and the more central jet (right) using CT10 as the PDF set. The markers in the lower panel are artificially displaced horizontally to make distinguishing the three κ values easier. The POWHEG+PYTHIA curves are nearly on top of the PYTHIA curves. The crossed lines in the bars on the data indicate the systematic uncertainty and the full extent of the bars is the sum in quadrature of the statistical and systematic uncertainties.	189
1.99	The measured average of the jet charge distribution (top), and the standard deviation (bottom), in units of the positron charge as a function of the jet p_T for $\kappa = 0.3, 0.5$, and 0.7 for the more forward jet (left) and the more central jet (right) using CTEQ6L1 as the PDF set. The markers in the lower panel are artificially displaced horizontally to make distinguishing the three κ values easier. The crossed lines in the bars on the data indicate the systematic uncertainty and the full extent of the bars is the sum in quadrature of the statistical and systematic uncertainties.	190
1.100	The extracted value of up- and down-quark jet charges in units of the positron charge in bins of jet p_T for $\kappa = 0.3, 0.5$, and 0.7 . The error bars include statistical, experimental systematic, and CT10 PDF uncertainties added in quadrature. The thick part of the error bar indicates the PDF contribution to the total uncertainty and the horizontal line on each error bar indicates the contribution from the statistical uncertainty. The first two p_T bins in the left plot are excluded due to their very large uncertainties.	192

1.101	The extracted values of the scale violation parameter c_κ from the data compared to theoretical calculations [172,240]. The error bars include statistical, experimental systematic, and PDF uncertainties added in quadrature. The thick part of the error bar indicates the PDF contribution to the total uncertainty and the horizontal line on each error bar indicates the contribution from the statistical uncertainty (each shown without adding in quadrature any other source of uncertainty).	195
2.1	An illustration of coherence. Photons emitted at large angles $\theta_{e\gamma}$ are not able to resolve the e^+e^- pair, leading to a suppression of radiation outside of θ_{ee} .	201
2.2	The same high energy parton undergoes fragmentation many times (<i>re-showered</i>). Each time the parton fragments, the stable hadrons in the event are clustered into a $R = 1.0$ jet and trimmed using k_t subjects with $R_{\text{sub}} = 0.3$ which are removed using the parameter $f_{\text{cut}} = 0.05$. The high energy parton is a color singlet (left) or color octet (right) Higgs boson h with $m_h = 125$ GeV and $p_T = 500$ GeV. To ensure no other significant radiation in the event, momentum is conserved by balancing the h against a $Z(\rightarrow \nu\bar{\nu})$. The histograms are the average jet image (see Sec. 4.3.2) over all re-showers with p_T intensity and the L^2 norm.	202
2.3	The average difference between the left and right distributions in Fig. 2.2.	202
2.4	The construction of the jet pull angle for jet J_1 with respect to J_2 .	204
2.5	The total statistical uncertainty and the charge-asymmetry statistical uncertainty of a pure sample of W +jets events with 30% more W^+ than W^- .	209
2.6	Diagram illustrating the color connections for the nominal sample with a color-singlet W (left) and the flipped sample with a color-octet W (right).	210
2.7	One LHE event for a nominal (top) and inverted (bottom) color flow.	212

2.8	The invariant mass of the leading non b -jets with (right) and without (left) a reweighting to the invariant mass spectrum.	213
2.9	The ΔR between the non b -tagged jets with (right) and without (left) a reweighting to the invariant mass and ΔR spectra.	213
2.10	The p_T of the leading non b -tagged jet with (right) and without (left) a reweighting to the invariant mass and ΔR spectra.	214
2.11	The pull angle with (right) and without (left) a reweighting to the invariant mass and ΔR spectra.	214
2.12	The pull vector magnitude with (right) and without (left) a reweighting to the invariant mass and ΔR spectra.	215
2.13	The charged particle pull angle (left) and the charged pull vector magnitude (right). No reweighting is applied to the invariant mass spectrum.	215
2.14	Left: The fraction of the W boson energy carried by the jets selected with three methods: Best, Baseline ('Base.'), and Alternative ('Alt.'). The baseline method uses the leading non- b tagged jets, the alternative method uses an invariant mass constraint, and the best method uses truth information; see the text for details. Right: The difference in fractions between the various methods.	220
2.15	The difference in the W boson energy fractions between the leading (in p_T) and subleading jets, $f_{\text{jet},i}^W(\text{lead}) - f_{\text{jet},i}^W(\text{sublead})$	221
2.16	Top: The distribution of f_W^{jet} (Jet energy fraction of W) and f_{jet}^W (W energy fraction of jet) for the leading (left) and sub-leading (right) jets. Bottom: the distribution between $f_W^{\text{jet},1}$ and $f_W^{\text{jet},2}$ (left) and between $f_{\text{jet},1}^W$ and $f_{\text{jet},2}^W$ (right).	222
2.17	Left (right): The leading (sub-leading) non b -tagged jet p_T . The uncertainty band includes the detector-related experimental uncertainties described in Sec. 2.5. The final bin includes overflow.	223

2.18	The distribution of the dijet p_T constructed from the leading two non b -tagged jets (hadronically decaying W boson candidate). The uncertainty band includes the detector-related experimental uncertainties described in Sec. 2.5. The final bin includes overflow.	224
2.19	Left (right): The leading (sub-leading) non b -tagged jet η . The uncertainty band includes the detector-related experimental uncertainties described in Sec. 2.5. The final bin includes overflow.	224
2.20	Left (right): The ΔR (invariant mass) distribution of the leading to non b -tagged jets. The uncertainty band includes the detector-related experimental uncertainties described in Sec. 2.5. The final bin includes overflow.	225
2.21	The jet pull angle $\theta_P(X, Y)$ distribution for various choices of X and Y for particle-level jets and also for detector-level jets matched to the particle-level jets.	227
2.22	The distribution of the jet pull angle response, $R(\theta_P)$, for both $\theta_P(J_1, J_2)$ and $\theta_P(J_1, B_1)$ as well as for the calorimeter pull angle and the track pull angle. Statistical uncertainties on the mean and RMS are an order of magnitude less than the values shown.	229
2.23	The $\theta_P(J_1, B_1)$ distribution for the calorimeter (left) and track (right) pull angle after replacing detector-level jet axes by particle-level jet axes.	230
2.24	The jet pull angle constructed in a pseudo MC with various resolution settings for the jet constituents.	232
2.25	A χ^2 -like scan between the simulated charged-particles pull angle distribution and templates using toy MC that are with varying asymmetry α and resolution r . The minimum χ^2 is for $\alpha = 1.5$ and $r = 1.0$. .	232
2.26	Relationship between the θ_P (response) and ΔR between the jets. . .	233
2.27	The jet pull angle distribution in three bins of p_T of J_1	234
2.28	The RMS of the jet pull angle response as a function of the jet p_T . The ΔR between jets is restricted between 1.4 and 1.6 in the right plot. The track-based pull angle response is nearly flat by $p_T \sim 100$ GeV.	235

2.29	The pull vector magnitude (left) and the relationship (right) between the jet pull angle response and $ \mathbf{v}_P(J_1) $ in a particular bin of jet $\mathbf{p}_T$. For the track (charged particle) pull magnitude in the left plot, at least one track (charged particle) is required.	237
2.30	The RMS of the jet pull angle response as a function of the fraction of events that pass a cut on the pull vector magnitude (left) and the $\mathbf{p}_T$ dependence of the average pull vector magnitude (right).	237
2.31	The jet pull angle response as a function of the number of jet constituents for J_1	238
2.32	Pairwise all-particles pull angle correlations using particle-level jets.	239
2.33	Un-matched event (treated as a background) rejection versus the matched (treated as a signal) efficiency. The optimal combination is constructed from the full joint likelihood.	240
2.34	The RMS of the jet pull angle response as a function of μ for $50 \text{ GeV} < \mathbf{p}_T^{J_1} < 60 \text{ GeV}$	240
2.35	The pull vector magnitude for both calorimeter pull and track pull. For the track pull, at least two tracks are required. Uncertainty bands include uncertainties on the jet energy scale and uncertainty as well as on the $t\bar{t}$ component of the MC. An uncertainty on the tracking efficiency is added for the track pull. No uncertainty is included for individual calorimeter clusters or for jet angular resolutions.	242
2.36	The distribution of the jet pull angle $\theta_P(J_1, J_2)$ for both calorimeter cluster constituents and track constituents in both data and MC. The left plot has a 25 GeV requirement for the jets while the right plot has a tight threshold placed on the $\mathbf{p}_T$ of the dijet system.	243
2.37	Same as Fig. 2.36 only with $\theta_P(J_1, B_1)$ instead of $\theta_P(J_1, J_2)$	243
2.38	A schematic diagram of the origin correction. The quantity θ^{detector} is the angle in the $\mathbf{zx}$ plane measured with respect to the geometric center of the detector, whereas the ‘true’ θ^{physics} is offset and based from the primary vertex (PV).	244

2.39	The η (left) and ϕ (right) jet angular response before and after the jet origin correction. The response is defined as the RMS of the $\Delta\eta$ or $\Delta\phi$ between the reconstructed jet and the $\Delta R < 0.4$ matched particle-level jets.	245
2.40	The pull angle response before and after applying the cluster origin correction. The jet origin correction is applied in both cases.	247
2.41	The distribution of the all-particles pull vector magnitude with both the jet axis and the cluster axes origin corrected (left) and only the jet axis corrected (right).	247
2.42	Left: All-particles pull angle response. Right: the all-particles pull angle distribution after various origin corrections.	248
2.43	The ϕ (left) and η (right) angular response for various choices of the jet axis. The ‘baseline’ configuration uses the calorimeter jet axis without the origin correction. The ‘calo axis’ uses the origin correction and the ‘track axis’ and ‘truth track’ use the four-vector sum of tracks. In all cases except the last one, the particle-level reference object is the full particle-level jet axis while in the last case, the reference is the four-vector sum of the charged particles only. The number in parenthesis is the RMS.	249
2.44	Left: Charged-particles pull angle response. Right: the charged-particles pull angle distribution after various origin corrections.	250
2.45	The charged-particles pull vector magnitude using the track four-vector sum for the jet axis.	250
2.46	The all-particles (left), charged-particles (middle), and pull angle response (right) in simulation using the nominal color singlet $t\bar{t}$ model and additionally with the color octet model (left and middle only).	251

2.47	The detector-level all-particles (left) and charged-particles pull (right) angles in data and in simulation. The uncertainty band includes only the experimental uncertainties on the inputs to the event selection and the jet pull calculation (see Sec. 2.5). A large part of the uncertainty displayed here affects the overall normalization and is correlated between the individual bins. This component of the uncertainty is cancelled in the unfolded measurement of the unit-normalized pull angle distribution (see Sec. 2.4).	251
2.48	The number of iterations required for the unfolded results to saturate using a Toy MC. The unfolding is said to saturate if the unfolded bin contents do not change by more than 0.01% between iterations. In this model, $\theta_p \sim \text{Uniform}(0, \text{Range})$ and $\theta_p^{\text{reco}} - \theta_p^{\text{true}} \sim \mathcal{N}(0, \sigma^2)$, where $\text{Range} = \pi$ and the smearing is done modulo π	253
2.49	The sum in quadrature of the data statistical uncertainty and the color flow model uncertainty (See Sec. 2.5) as a function of the number of iterations in the IB unfolding method for the three bins of the all-particles pull angle (left) and for the four bins of the charged-particles pull angle (right). For the first bin, the open markers show how the total uncertainty is broken down into the two components.	255
2.50	A schematic diagram of the jet pull angle response which illustrates that the pull angle resolution is strongly correlated with the pull angle itself. For example, if you know that the difference between the particle-level and detector-level (normalized by π) pull angles is greater than 0.5 (blue shaded box on the right), then the truth pull angle must have been greater than 0.5.	256
2.51	The truth distribution for the toy experiment described in Sec. 2.4.1. The distributions are indexed by the base of a triangle which varies between $X = 1$ and 5. When the base length is 5, the distribution is close to uniform and when it is 1, the distribution is strongly peaked at zero.	257

2.52	The bin averaged fractional uncertainty from comparing the truth and unfolded toy data ($X = 5$) using different response matrices (with $X = X_{MC}$). The value of σ is used in both the toy data and the response matrix. The left plot uses three iterations while the right plot uses 15 iterations.	258
2.53	The background composition as a function of the all-particles pull angle (left) and the all-particles pull angle (right).	259
2.54	The pull-angle dependence of the fake factors for the all-particles pull angle (left) and for the charged-particles pull angle (right).	260
2.55	The pull-angle dependence of the inefficiency factors for the all-particles pull angle (left) and for the charged-particles pull angle (right). . . .	261
2.56	The ratio of the fake (inefficiency) factor for muon events to electron events on the left (right).	262
2.57	The response matrix for the all-particles pull angle on the left and for the charged-particles pull angle on the right. The z-axis is the probability for an event to be reconstructed in a detector-level bin given that it started in a fixed particle-level bin on the horizontal axis (i.e. the columns are normalized to unity).	263
2.58	The difference between the response matrices for the muon and electron channels divided by the simulation statistical uncertainty. . . .	263
2.59	The χ^2/NDF distribution of tracks before the $\chi^2/NDF < 3$ requirement. 265	
2.60	The probability for not finding a calorimeter cell cluster matched to a track as a function of the track momentum. Reproduced from Ref. [232]. 267	
2.61	The average LCW E/p for $0 < \eta < 0.6$ (left) and $0.6 < \eta < 1.1$ (right), using the same analysis framework as in Ref. [287] (but LCW is not in the public note - thank you B. Axen for the inputs). The blue band in the ratio shows the estimated uncertainty used for the cluster energy scale uncertainty.	269
2.62	The distribution of cluster energy in three bins of the cluster isolation, with less isolated clusters on the left and more isolated cluster on the right. See the text for the definition of the cluster isolation.	270

2.63	Left (Right): The ΔR between isolated low (high) momentum tracks and clusters in $Z \rightarrow \mu\mu$ events in the barrel of the detector. Inputs from C. Young.	272
2.64	Left (Right): The ΔR between isolated low (high) momentum tracks and clusters in $Z \rightarrow \mu\mu$ events in the endcap of the detector. Inputs from C. Young.	272
2.65	Left (Right): The RMS of the $\Delta\eta$ ($\Delta\phi$) between isolated single particle tracks and clusters for tracks extrapolated to the second layer of the calorimeter in the barrel of the detector.	273
2.66	Left (Right): The RMS of the $\Delta\eta$ ($\Delta\phi$) between isolated single particle tracks and clusters for tracks extrapolated to the second layer of the calorimeter in the endcap of the detector.	273
2.67	A comparison between the track-jet and cluster-induced JAR uncertainties (baseline) for the jet η (left) and the jet ϕ (right).	274
2.68	The p_T dependence of the jet angular resolution for η (left) and ϕ (right) for the leading non b -tagged jets in $t\bar{t}$ events The error bars reflect the statistical uncertainty and the dashed line is a fit to $a + b/p_T$	276
2.69	The η dependence of the jet angular resolution for η (left) and ϕ (right) for the leading non b -tagged jets in $t\bar{t}$ events.	277
2.70	The jet p_T dependence of the resolution of the $\Delta\eta$ (left) and $\Delta\phi$ (right) between track jets and calorimeter jets in simulated dijet events. The error band is described in the text. A dashed line is a fit to $a + b/p_T$. Inputs from F. Guescini.	277
2.71	The same as Fig. 2.70, only with η instead of ϕ . The dashed line is a fit to $a + b\eta^2$. Inputs from F. Guescini.	278
2.72	The relative jet angular resolution for ϕ and η (left) and the amount by which the jet angle should be smeared to conservatively cover the uncertainty in the angular resolution (right).	278

2.73	The fractional difference in the unfolded result when the Wt is replaced with the flipped $t\bar{t}$ scaled to the Wt yield compared with the data statistical uncertainty for the all particles pull angle (left) and the charged particles pull angle (right).	280
2.74	Particle-level distributions comparing PYTHIA 6 and HERWIG for the all particles pull angle (left) and the charged particles pull angle (right).	281
2.75	Particle-level distributions comparing PYTHIA 6 and HERWIG for the all particles pull vector magnitude (left) and the charged particles pull vector magnitude (right).	282
2.76	Truth level distributions comparing the nominal and low CR tunes of Pythia for the all particles pull angle (left) and the charged particles pull angle (right).	283
2.77	The ISR/FSR variations at truth level for the all particle pull (left) and the charged particles pull (right). Electron and muon channels combined.	283
2.78	The fractional difference between the nominal particle-level distribution and the nominal detector-level distribution unfolded with a response matrix constructed from the flipped sample (<i>unfolding difference</i>). The dashed line shows the fractional difference at particle-level between the two models. All distributions are normalized to unity before computing fractional differences.	286
2.79	The fake factors for the all-particles pull angle (left) and the charged-particles pull angle (right). The (1) and (8) in the legend refers to the nominal and flipped sample, respectively.	287
2.80	The inefficiency factors for the all-particles pull angle (left) and the charged-particles pull angle (right). The (1) and (8) in the legend refers to the nominal and flipped sample, respectively.	287

2.81	The systematic uncertainty covariance matrix for all experimental and background normalization uncertainties. If the matrix were proportional to the identity matrix, then the systematic uncertainty in bin i of the all-particles pull angle would be $\sqrt{\Sigma_{ii}}$ and in bin i of the charged-particles pull angle would be $\sqrt{\Sigma_{i+3,i+3}}$ for Σ a matrix representing the plot above. The matrix Σ is the sum of the matrices from all individual sources of uncertainty.	290
2.82	The same as Fig. 2.81, but broken into four categories: cluster and tracking (top left), background normalization (top right), jet energy resolution (bottom left), and jet energy scale (bottom right).	291
2.83	Representations of correlations between and within bins of the all-particles and charged-particles pull angles induced from the unfolding (top left), normalization (bottom left), and jet-by-jet statistical fluctuations (top and bottom right). The unfolding correlation includes normalization. These matrices are generated by running the unfolding procedure on the data 10,000 times, each time generating <i>pseudo-data</i> $X_{ij} \sim \text{Poisson}(\lambda_{ij})$, where λ_{ij} is the number of events measured in bin (i, j) . For each pseudo-dataset, the backgrounds are subtracted before (normalization, unfolding, and) computing the correlation.	292
2.84	Left (Right): the all-particles (charged-particles) pull angle distribution for the unfolded data and three particle-level simulations. The orange inner band on the data represents the statistical uncertainty while the yellow band is the sum in quadrature of the statistical and systematic uncertainty. Final version of this plot is from T. Neep. . .	294
2.85	The distribution of the log-likelihood ratio test statistic under the SM and flipped color flow models. To generate the distributions, the full covariance matrix is resampled ten million times. In order to impose unitarity of the pseudo-data, only the first n bins are sampled and the $(n - 1)^{\text{th}}$ is fixed by normalization. Additionally removing this last bin from the log-likelihood ratio has little impact on the approximate significances quoted in the text.	295

3.1	The response after applying an inclusive jet energy calibration (left) and after the residual correction (right). The lower panel is the distribution of the number of tracks inside jets for three jet p_T ranges. Reproduced from Ref. [125].	304
3.2	The two-dimensional likelihood ratio for the track multiplicity (vertical axis) and track width (horizontal axis) quark/gluon tagger in data (left) and simulation (right). Reproduced from Ref. [167].	305
3.3	The dijet invariant mass spectrum for the full jet selection without the n_{track} requirement (left) and with all requirements (right). Reproduced from Ref. [298].	305
3.4	The number of reconstructed tracks associated with a jet in three example jet p_T ranges for data and for PYTHIA 8 and HERWIG++ predictions. The data points have statistical uncertainties which in all bins are smaller than the marker size.	306
3.5	The simulated fraction of jets originating from gluons as a function of jet p_T for the more forward jet (down triangle), the more central jet (up triangle), and the difference between these two fractions (circle). The fractions are derived from PYTHIA 8 with the CT10 PDF set and the error bars represent the PDF and matrix element uncertainties, further discussed in Sect. 4.1.2.3.4. The uncertainties on the fraction difference are computed from propagating the uncertainties on the more forward and more central fractions, treating as fully correlated.	308
3.6	The p_T dependence of $\langle n_{\text{charged}}^f \rangle$, $\langle n_{\text{charged}}^c \rangle$, $\langle n_{\text{charged}}^g \rangle$, and $\langle n_{\text{charged}}^q \rangle$ in simulation for PYTHIA 8 AU2 as well as the extracted values of $\langle n_{\text{charged}}^g \rangle$, and $\langle n_{\text{charged}}^q \rangle$ using the procedure described in the text.	310
3.7	The impact of binning the n_{charged} distribution at particle-level. The notation 1-1 means that there is one bin per charged particle multiplicity.	312
3.8	The fraction of events with more than 60 charged particles.	312
3.9	For each bin of the combined jet p_T and n_{track} distribution, the inefficiency factors (left) and the fake factors (right) for the more forward jet.	313

3.10	The fractional bias induced when unfolding PYTHIA 8 simulation with a HERWIG++ response matrix for various numbers of iterations in the IB technique.	314
3.11	The total response matrix (left) and the response matrix for the 61 n_{track} bins averaged over the 11 jet p_T bins. The z-axis is truncated at 1%.	315
3.12	The response matrix for the 61 n_{track} bins in the ninth jet p_T bin ($1.0 \text{ TeV} < p_T < 1.2 \text{ TeV}$) on the left and the last jet p_T bin ($p_T > 1.5 \text{ TeV}$) on the right. The lower panel shows the average detector-level n_{track} divided by the particle-level n_{track} in each particle-level n_{track} bin. The z-axis is truncated at 1%.	315
3.13	The jet p_T dependence of (a) the average reconstructed track multiplicity for uncorrected data and detector-level simulation, (b) the average reconstructed track multiplicity for the detector-level simulation and the average charged-particle multiplicity for the particle-level simulation, (c) the average charged-particle multiplicity for the unfolded data and the particle-level simulation, and (d) the average charged-particle multiplicity divided by the average reconstructed track multiplicity in simulation. For the data, only statistical uncertainties are included in the error bars (which are smaller than the markers for most bins).	317
3.14	The impact of the charged energy loss uncertainty on the average number of charged particles in each p_T bin. The vertical axis is the relative uncertainty on the average unfolded n_{charged} . For each n , tracks are dropped randomly with a probability given by $\text{Pr}(\text{drop track } i) = \alpha p_{Ti}^n$, where α is fixed such that the MC is the same as the data in Fig. 1.61. Below 400 GeV, nuclear interactions dominate the track reconstruction efficiency uncertainty.	319
3.15	The bin-by-bin inefficiency factor fractional uncertainty before unfolding (left) and the inefficiency factor fractional uncertainty on $\langle n_{\text{track}} \rangle$ as a function of the jet p_T bin (right) for track $p_T > 500 \text{ MeV}$	321

3.16	The bin-by-bin fake factor fractional uncertainty before unfolding (left) and the fake factor fractional uncertainty on $\langle n_{\text{track}} \rangle$ as a function of the jet p_T bin (right) for track $p_T > 500$ MeV.	321
3.17	The distribution of the weights used to re-weight the MC distribution for the non-closure test (left) and the re-weighted truth distribution (labeled after) (middle) for the more forward jet. The $\langle n_{\text{track}} \rangle$ in data and simulation before and after the re-weighting is shown in the right plot.	322
3.18	The non-closure uncertainty for $\langle n_{\text{charged}} \rangle$ compared with the raw data/MC difference in the same quantity (left) and the relative difference in $\langle n_{\text{charged}} \rangle$ when unfolding PYTHIA 8 with itself (= particle-level PYTHIA by construction) and with a HERWIG++ response matrix (right).	323
3.19	A visualization of the systematic uncertainties for the $\langle n_{\text{charge}} \rangle$ distribution.	324
3.20	The measured average charged-particle multiplicity as a function of the jet p_T , combining the more forward and the more central jets for $p_T^{\text{track}} > 0.5$ GeV. The band around the data is the sum in quadrature of the statistical and systematic uncertainties. Error bars on the data points represent the statistical uncertainty (which are smaller than the markers for most bins).	327
3.21	The measured average charged-particle multiplicity as a function of the jet p_T , combining the more forward and the more central jets for $p_T^{\text{track}} > 2$ GeV (left) and $p_T^{\text{track}} > 5$ GeV (right). The band around the data is the sum in quadrature of the statistical and systematic uncertainties. Error bars on the data points represent the statistical uncertainty (which are smaller than the markers for most bins).	328

3.22	The jet p_T dependence of the difference in the average charged-particle multiplicity ($p_T^{\text{track}} > 0.5$ GeV) between the more forward and the more central jet. The band for the data is the sum in quadrature of the systematic and statistical uncertainties and the error bars on the data points represent the statistical uncertainty. Bands on the simulation include MC statistical uncertainty.	330
3.23	The jet p_T dependence of the average charged-particle multiplicity ($p_T^{\text{track}} > 0.5$ GeV) for quark- and gluon-initiated jets, extracted with the gluon fractions from PYTHIA 8.175 with the CT10 PDF. In addition to the experimental uncertainties, the error bands include uncertainties in the gluon fractions from both the PDF and ME uncertainties. The MC statistical uncertainties on the open markers are smaller than the markers. The uncertainty band for the N ³ LO pQCD prediction is determined by varying the scale μ by a factor of two up and down. The markers are truncated at the penultimate p_T bin in the right because within statistical uncertainty, the more forward and more central jet constituent charged-particle multiplicities are consistent with each other in the last bin.	333
4.1	A schematic illustration of the setup described in the text to compute the p_T dependence of ΔR . The black dot represents the W boson and the arrows from the dot represent the momentum of the quark decay products.	336
4.2	Left: the joint distribution of ΔR and p_T and Right: the fraction of events in which the decay products ($W \rightarrow qq'$ or $t \rightarrow bqq'$) are within $\Delta R < 1$ of the parent particle.	337
4.3	A schematic diagram for the emission of one soft and collinear gluon off of a quark. In this section, $z \ll 1$ and $\theta \ll 1$	340

- 4.4 A schematic diagram of the $\log(1/z)$ versus $\log(R/\theta)$ plane in which the probability for the emission of a gluon is approximately uniform. The dashed lines show contours of constant m^2 , which increase from upper right to lower left. The blue triangle corresponds to mass m_0^2 . 344
- 4.5 Left: The leading logarithm distribution of the jet mass for $(E,R) = (200 \text{ GeV}, 1)$ and $(E,R) = (400 \text{ GeV}, 0.4)$ separately for quark and gluon jets. The quark distribution is suppressed at zero, but it increases for finite mass much faster than the gluon distribution. Right: The average jet mass as a function of jet p_T for the same four settings as the left plot. Horizontal lines indicate the W boson, Higgs boson and top quark masses. 345
- 4.6 The p_T dependence of the average leading anti- k_t $R = 1.0$ trimmed with $f_{\text{cut}} = 0.05$ and $R_{\text{sub}} = 0.2$ jet mass. See Sec. 4.1.2 for details. . 345
- 4.7 The impact of a jet constituent on that jet's mass as a function of ΔR between the constituent and the jet axis and the p_T fraction carried by the constituent. W bosons are generated with $p_T = 200 \text{ GeV}$ in a random direction and decay via a scalar two-body phase space into quarks which subsequently decay into 10 massless particles each also with the scalar phase space. The 'jet' is the four-vector sum of all these particles, which has mass 80 GeV and $p_T = 200 \text{ GeV}$ 346
- 4.8 An event display of a dijet event in the early Run 2 data. The gray and colored filled circles correspond to calorimeter cell clusters where the circle radius is proportional to the log of the cluster energy. Colored circles represent the constituents of the trimmed jets. The blue-gray circles are the ungroomed anti- k_t $R=1.0$ jets. The dark gray circles are the constituent clusters removed by trimming. The remaining gray area underneath the colored circles is the ghost area of the trimmed jet. 348
- 4.9 The distribution of the jet mass response (R) for boosted hadronically decaying W and Z bosons in four boson p_T ranges from 200 GeV up to 2 TeV. See the text for a description of the various fitting methods. 350

4.10	The distribution of the jet p_T response (R) for boosted hadronically decaying W and Z bosons in four boson p_T ranges from 200 GeV up to 2 TeV. See the text for a description of the various fitting methods.	351
4.11	A summary of the p_T -dependence of the jet mass (left) and jet p_T (right) resolution. See the text for a description of the various fitting methods.	352
4.12	Left: The dependence of $\langle f_Q^{-1} \rangle$ on the standard deviation of the charged-to-neutral fluctuations. The charged fraction f_Q is modeled as a log normal distribution with mean $2/3$ and variable standard deviation σ . The horizontal dashed line is at $2/3$. The vertical dashed line is at $\sigma = 0.47 \approx \sqrt{\frac{2}{3}(1 - \frac{2}{3})}$, which is what one expects if each particle inside a jet carries an equal fraction of the jet energy and has a probability $2/3$ of being charged. Right: the p_T dependence of $\langle f_Q \rangle$ for both jet mass and jet p_T . The band is the standard deviation of the f_Q distribution using p_T .	354
4.13	The (a) ratio of the calorimeter jet mass to track jet mass and (b) average value of this ratio as a function of the calorimeter jet mass for leading anti- k_t $R = 1.0$ trimmed with $f_{\text{cut}} = 0.05$ and $R_{\text{sub}} = 0.2$ jets. Calorimeter jets which contain only a single cluster have a mass of zero. MC is normalized to the number of events observed in data. The last bin includes overflow events.	356
4.14	(a) The distribution of r_{track} in the data for semi-leptonic $t\bar{t}$ events with the selected jet in the range $50 \text{ GeV} < m^{\text{jet}} < 120 \text{ GeV}$. (b) The median of the r_{track} distribution as a function of the jet p_T . (c) The inter-quantile range as a measure of the width. See Sec. 4.2 for a description of the event selection.	357
4.15	The joint distribution of the response (R) and the product of the inverse charged ratio f_Q^{-1} and the inverse track mass response R_T^{-1} using simulated W boson jets with $p_T^{\text{truth}} > 200 \text{ GeV}$ from POWHEG-BOX+PYTHIA 6 $t\bar{t}$ production.	358
4.16	Left: The gap between $\langle R \rangle \langle f_Q^{-1} R_T^{-1} \rangle$ and $\langle r_{\text{track}} \rangle = \langle R f_Q^{-1} R_T^{-1} \rangle$.	359

4.17	The average jet p_T (left) and mass (right) response as a function of p_T using simulated detector-level jets and propagating single particle responses via Eq. 4.23. The points marked data use E/p as measured with the ATLAS detector. Error bars on the points indicate the MC statistical uncertainty on the mean.	361
4.18	Left: the mass distribution of the W boson (parton-level) and boosted and isolated W boson jets at particle-level and the same jet at detector-level. The jets are required to have $p_T > 200$ GeV. See Sec. 2.2.1 for details of the simulation. Right: the particle-level jet mass spectrum for candidate boosted W boson jets from $t\bar{t}$ events using the default Perugia 2012 PYTHIA 6 tune or the radHi/radLo variations. The one-loop running of the strong coupling-constant is $\alpha_s \propto 1/\ln(Q^2/\Lambda^2)$ with $\Lambda = 0.26$ in the default Perugia 2012 tune. This value is double or halved in the radHi/radLo variations.	363
4.19	The jet mass distribution for events passing the selection described in Sec. 4.1.2.3.1 for electron events (left) and muon events (right). . . .	365
4.20	The jet p_T distribution for events passing the selection described in Sec. 4.1.2.3.1 for electron events (left) and muon events (right). . . .	365
4.21	Left: dependence of the average response $\langle m_{\text{reco}}/m_{\text{true}} m_{\text{true}}, p_T^{\text{reco}} \rangle$ as a function of p_T^{reco} and m_{true} . Right: the fitted distribution when using the detector-level simulation in place of the data. By construction, the fitted distribution is identical to the input detector-level distribution ($r = s = 1$).	369
4.22	The value of the χ^2 per degree of freedom (NDF) at a given relative jet mass scale (left) or jet mass resolution (right) minimized over the variable not shown.	371
4.23	A comparison between the post- and pre-fit simulation and the data. The vertical dotted lines indicate the fit range.	372
4.24	The one (solid) and two (dashed) σ uncertainty ellipses for the relative jet mass scale and the relative jet mass resolution.	372

4.25	The probability distribution (left) and the integral of the probability distribution above a given threshold (strictly speaking, this is one minus the cumulative distribution) on the top quark p_T for $\sqrt{s} = 8$ and 13 TeV. The increase in the inclusive $t\bar{t}$ cross-section between these two center-of-mass energies is about 3.5.	376
4.26	The jet mass computed directly from calorimeter-cell clusters (top) and with the track-assisting procedure (bottom) for a lower p_T (left) and a higher p_T (right). The track-assisted jet mass is defined in Sec. 4.1.4. Note that the ATLAS default trimming parameters for Run 2 are slightly different than Run 1: R_{sub} is now 0.2 instead of 0.3 in order to improve the resolution at high p_T . The event selection is identical to the one at $\sqrt{s} = 8$ TeV, but only the muon channel is used in this section (negligible multijet contribution). The bands include detector-level jet and particle-level modeling systematic uncertainties. For illustration, the top left plot shows also the particle-level distribution. Note that all distributions are normalized to the data integral.	378
4.27	The value of the χ^2 per degree of freedom (NDF) at a given relative jet mass scale (left) or jet mass resolution (right) minimized over the variable not shown.	379
4.28	Left: The statistical (solid) and total (dashed) 1σ uncertainty ellipses for the relative jet mass scale and the relative jet mass resolution at $\sqrt{s} = 8$ TeV (red) and $\sqrt{s} = 13$ TeV (black). Right: the relative scale and resolutions for the two mass reconstruction algorithms in Fig. 4.26.	379
4.29	The ratio of the large-radius jet p_T to the lepton p_T (left), the two-vector sum of the lepton and close-by jet transverse momentum (middle), and the full leptonically decaying top quark using the lepton, the close-by jet, and the $\vec{p}_T^{\text{miss}}$ (right).	380

4.30	The value of the χ^2 per degree of freedom at a given relative jet $\mathbf{p}_T$ scale minimized over the jet $\mathbf{p}_T$ resolution for three different reference objects: the lepton $\mathbf{p}_T$ (left), the two-vector sum of the lepton and close-by jet transverse momentum (middle), and the full leptonically decaying top quark (right).	381
4.31	The value of the χ^2 per degree of freedom at a given relative jet $\mathbf{p}_T$ resolution minimized over the jet $\mathbf{p}_T$ scale for three different reference objects: the lepton $\mathbf{p}_T$ (left), the two-vector sum of the lepton and close-by jet transverse momentum (middle), and the full leptonically decaying top quark (right).	381
4.32	An example event which has been clustered using the anti- k_t $R = 1.0$ (left) and with anti- k_t $R = 1.0$ re-clustered $r = 0.3$ anti- k_t jets (right). The shaded regions show the jet area determined by clustering ghost particles. Only large radius jets with $\mathbf{p}_T > 50$ GeV are shown and small radius jets are required to have $\mathbf{p}_T > 15$ GeV. As with groomed large radius jets, re-clustered jets can have non-circular shapes. . . .	386
4.33	Various re-clustered grooming parameters for anti- k_t $r = 0.3$ jets for $\text{NPV} = 20$ on the left and $\text{NPV} = 80$ on the right.	389
4.34	Mean, mass resolution, and mass window efficiency of the mass distribution as a function of the number of additional vertices for various re-clustered jet grooming schemes.	389
4.35	Various small radii for a fixed algorithm of anti- k_t for $\text{NPV} = 20$ on the left and $\text{NPV} = 80$ on the right.	390
4.36	Mean, mass resolution, and mass window efficiency of the mass distribution as a function of the number of additional vertices for various small jet radii.	390

4.37	The performance of a bottom-up approach to jet substructure where the radius r jets are inputs to substructure variables. Solid lines show the performance curves for large radius trimmed jets ($R_{\text{sub}} = 0.3, f_{\text{cut}} = 0.05$) and the dashed lines show the analogous re-clustered variable. Random tagger denotes a classifier which picks signal and background with equal probability. The variable N_{sub} is the number of re-clustered subjets. The curves are determined by placing threshold requirements on the variable likelihoods.	392
4.38	Left: The efficiency for signal and background jets to be reconstructed in the smallest mass window containing 68% of the signal. Right: the size of the window from the left plot. Both plots compare large radius trimmed jets with $R = 1.0, f_{\text{cut}} = 0.05$ and $R_{\text{sub}} = 0.05$ with re-clustered jets that have analogous parameters.	394
4.39	The small-radius jet constituent multiplicity for several choices of R and a fixed $r = 0.2$	395
4.40	Left: The efficiency for signal and background jets to be reconstructed in the smallest mass window containing 68% of the signal. Right: the size of the window from the left plot. In both plots, $r = 0.2$ and $f_{\text{cut}} = 0.05$	396
4.41	The jet mass distribution for the signal (background) for various R values. In both plots, $r = 0.2$ and $f_{\text{cut}} = 0.05$	396
4.42	Left: The efficiency for signal and background jets to be reconstructed in the smallest mass window containing 68% of the signal. Right: the size of the window from the left plot. In both plots, $r = 0.2$ and $R = 1.0$	397
4.43	Left: Window efficiency for various r values. Right: Size of the smallest 68% window. Both plots compare large radius trimmed jets with $R = 1.0, f_{\text{cut}} = 0.05$ and $R_{\text{sub}} = 0.05$ with re-clustered jets that have analogous parameters.	398
4.44	The small-radius jet constituent multiplicity for several choices of f_{cut} and a fixed $r = 0.2$ (left) and $r = 0.4$ (right).	399

4.45	The size of the smallest 68% window as a function of f_{cut} for $r = 0.2$ (left) and $r = 0.4$ (right).	399
4.46	The p_T (left) and mass (right) response of large radius trimmed jets and re-clustered jets for $200 \text{ GeV} < p_T^V < 400 \text{ GeV}$. In this p_T bin, over 80% of the re-clustered jets have at least two constituents. . .	402
4.47	The p_T response, with the p_T^V in the denominator instead of the truth jet p_T , of large radius trimmed jets and re-clustered jets for various bins of boson p_T	402
4.48	Left: the ratio of the re-clustered jet p_T to the large radius trimmed jet p_T as a function of the boson p_T . Middle (right): the joint distribution of the large-radius jet response and the re-clustered jet mass response in a low (high) p_T^V bin.	403
4.49	The p_T response as a function of the ΔR between the re-clustered jet constituents in events where the re-clustered jet has exactly two constituents and $200 \text{ GeV} < p_T^{\text{re-clustered jet}} < 400 \text{ GeV}$	403
4.50	The joint distributions of f_{closeby} computed with nearby ($1 < \Delta R < 2$) small-radius jets or all non-jet clusters within $\Delta R < 2$ (left) and two cluster-based definition that either include or exclude clusters dropped during trimming (right).	406
4.51	Left: the jet m/p_T and trimmed jet area regions used in this analysis. Right: the distribution of the cluster f_{closeby} in all the nine bins of jet m/p_T and trimmed jet area, evenly spaced between 3×3 bins in the range 0.2 and 0.5 in m/p_T and in the range 0.2π and 0.5π in the trimmed jet area.	406
4.52	The ratio of r_{track} between isolated (lose f_{closeby}) and non-isolated jets (high f_{closeby}) as a function of the jet area and m/p_T in data (left) and simulation (right). The smaller number in the data plot is the statistical uncertainty.	407
4.53	The ratio of the left and right plots from Fig. 4.52. The smaller number in each bin indicates the data statistical uncertainty.	408

4.54	The same as Fig. 4.53, but with the JMS varied down (left) or up (right).	408
4.55	The distribution of the ratio of the leading large-radius jet mass to small-radius jet mass as a function of W boson p_T	409
4.56	Left: the average multiplicity of tracks and clusters in $R = 0.4$ boosted W and Z boson jets as a function of the particle-level jet p_T . Right: an event display of the clusters inside one particular high p_T boosted boson jet. The TA mass is the track-assisted jet mass (see text for details).	411
4.57	The jet mass distribution of reconstructed boosted hadronically decaying W and Z boson jets with $p_T > 1.5$ TeV. A jet collection is calibrated if the average value of $m_{\text{reco jet}}/m_{\text{truth jet}}$ is unity.	412
4.58	The track-assisted and calorimeter jet mass distribution for boosted W and Z boson jets clustered using the $R = 0.4$ anti- k_t algorithm for $600 < p_T < 800$ GeV (left) and $p_T > 2$ TeV (right).	413
4.59	The p_T dependence of the inter-quantile range divided by the median of the track-assisted and calorimeter jet mass distributions. The 10%, 20%, and 30% inter-quantile ranges are centered around the median.	414
4.60	The calorimeter jet momentum resolution uncertainty (circle), the jet p_T resolution-induced track-assisted jet mass resolution uncertainty (up triangle), and the difference in the track-assisted jet mass resolution when comparing PYTHIA and HERWIG++. The resolution uncertainty is defined as $\sqrt{v^2 - n^2}$, where n is the nominal response resolution and v is the response resolution from the alternative sample. The alternative sample for the first two cases is created by smearing the jet energy within the resolution uncertainty and the alternative sample for the last case is the HERWIG++ simulation.	415

4.61	Left: the joint distribution of the track-assisted jet mass response and the calorimeter jet mass response. Right: the p_T dependence of the normalized jet mass resolution for the track-assisted jet mass, the calorimeter jet mass, and the optimal combination of the two mass definitions.	416
4.62	Distributions of the track-assisted jet mass, the calorimeter jet mass, and the optimal combination are shown in a few jet p_T bins.	417
4.63	The p_T spectrum of the simulated W and Z bosons from $W' \rightarrow WZ$ decays before applying any p_T re-weighting.	422
4.64	(a) Breit–Wigner resonances for the W (red) and Z (blue) bosons, (b) angular distribution of the decay products of transversely polarized W/Z bosons with respect to the spin direction in the boson rest frame, (c) hadronic branching fractions of the W^+ boson, and (d) of the Z boson. In (c) and (d), <i>light</i> stands for decay modes not involving c and b quarks.	425
4.65	The boson mass at generator level (top left), the ‘truth’ jet mass at particle-level after parton fragmentation (top right), and the detector-level jet mass distributions (bottom) for W and Z boson jets separately (left) and for an even admixture of the two jet types (right). The parton-level plot has a different vertical scale than the other plots and also has no p_T requirement.	427
4.66	The p_T dependence of the jet mass distribution. The points are the mean value in a given boson p_T bin and the shaded region is the standard deviation for the Z boson distribution.	428
4.67	The jet charge distribution for jets originating from $W^\pm$ and Z bosons in simulated W' decays. Each distribution is normalized to unity. The parameter κ controls the p_T -weighting of the tracks in the jet charge sum.	429
4.68	The p_T dependence of the jet charge distribution. The points are the mean value in a given boson p_T bin and the shaded region is the standard deviation for the Z boson distribution.	430

4.69	The distribution of the MV1 discriminant for the leading (left) and subleading (right) small-radius jets matched to the leading large-radius jet. The spike at -1 in the right plot corresponds to cases in which there is not a second small-radius jet. The other features in the distribution correspond to transitions in the dominant input algorithms to MV1 [75].	431
4.70	The efficiency-binned MV1 distribution for small-radius jets associated with large-radius jets resulting from W and Z boson decays. The left (right) plot shows the leading (sub-leading) small-radius jet MV1 distribution. The bins correspond to exclusive regions of b -jet efficiency. As such, the bin content of the black line (b -tagging for b -jets) should be proportional to the size of the efficiency window: about 50% for the rightmost bin, 10% for the three middle bins and 20% for the second bin.	432
4.71	Left: The fraction of events with $\Delta R(\mathbf{q}, \mathbf{q}) < 0.4$ in a toy MC simulation as a function of the simulated boson mass. Right: The p_T spectrum of the softer of the two decay products for a simulated (scalar) boson with mass 1 TeV when one of the decay products is below 25 GeV or the two decay products are within $\Delta R < 0.4$	433
4.72	The p_T dependence of the number of small radius b -tagged jets (70% working point) associated to the large radius jet. The points are the mean value in a given boson p_T bin and the shaded region is the standard deviation for the Z boson distribution.	433
4.73	The distribution of the jet mass for Z boson jets with 0, 1, or 2 matched muons before (left) and after (right) adding the muon to the jet four-vector before computing the mass.	435

4.74	Left: The relationship between the leading and sub-leading MV1 distributions. The distributions are normalized per bin of the leading small-radius jet (horizontal axis). The two MV1 values are independent if the distribution in each bin along the vertical axis does not change as a function of the horizontal axis (which is true within the MC statistical uncertainties). Right: A validation of the templates for the combined binned MV1 B (template is dashed).	436
4.75	(a) The jet mass $p(M \mathcal{F}, V)$ and (b) jet charge $p(Q \mathcal{F}, V)$ templates conditioned on the flavor $\mathcal{F}$ of the boson V decay for jets with $200 \text{ GeV} < p_T < 400 \text{ GeV}$. The solid lines are for Z boson decays and the dashed lines are for W boson decays.	437
4.76	The joint distribution of the jet mass and jet charge for W boson jets (left) and Z boson jets (right). The peak of the distribution along the jet mass axis is shifted toward higher values for the Z due to its higher mass. The linear correlation is less than 1% in both cases and the two distributions are nearly independent.	437
4.77	The full likelihood ratio for the tagger formed from jet mass, jet charge, and a small-radius jet b -tagging discriminant. The black histogram shows the likelihood ratio for Z bosons and the red histogram is the likelihood ratio for W^+ bosons. The shaded areas show the region of the likelihood ratio corresponding to 90%, 50%, and 10% working points of the Z -boson tagging efficiency.	439
4.78	The tradeoff between Z efficiency and (a) $1 - (W^+ \text{ efficiency})$ (b) or $1/(W^+ \text{ efficiency})$ on (a) a linear scale and (b) a logarithmic scale. Each curve is constructed by placing thresholds on the likelihood constructed from the inputs indicated in the legend. Since the b -tagging discriminant is binned in efficiency, there are only discrete operating points for the tagger built only from B	440

4.79	The tradeoff between W^+ efficiency and (a) $1 - (\mathbf{Z}$ efficiency) or (b) $1/(\mathbf{Z}$ efficiency) on (a) a linear scale and (b) a logarithmic scale. Each curve is constructed by placing thresholds on the likelihood constructed from the inputs indicated in the legend. Since the $\mathbf{b}$ -tagging discriminant is binned in efficiency, there are only discrete operating points for the tagger built only from $\mathbf{B}$	440
4.80	The $\mathbf{p}_T$ spectrum of the small-radius $\mathbf{b}$ -jets and the $\Delta\mathbf{R}$ between them when there are two.	443
4.81	The impact of selected systematic uncertainties on benchmark working points of the boson-type tagger. (a) a jet-mass-only tagger, for 50% (left) and 90% $\mathbf{Z}$ efficiency benchmarks. (b) a jet-charge-only tagger, for 50% (left) and 90% $\mathbf{Z}$ efficiency benchmarks. The point marked HERWIG uses the alternative shower and hadronization model for the simulation, with the likelihood template from PYTHIA . See the text for an explanation of the notation in the legend.	444
4.82	The impact of selected systematic uncertainties on benchmark working points of a $\mathbf{b}$ -tagging-only tagger at a 10% $\mathbf{Z}$ efficiency benchmark. The $\mathbf{b}$ -tagging discriminant is binned, so there are only discrete operating points. The point marked HERWIG uses the alternative shower and hadronization model for the simulation, with the likelihood template from PYTHIA . The $\mathbf{b}$ -tagging scale factor uncertainties are determined separately for $\mathbf{b}$ -, $\mathbf{c}$ -, and light-quark jets. Variations are added in quadrature for each ‘truth’ jet flavor. There is no contribution from the $\mathbf{b}$ -jet scale factor uncertainties on the W rejection because there are no ‘truth’ $\mathbf{b}$ -jets. Conversely, the $\mathbf{c}$ - and light-jet scale factor uncertainties do not impact the $\mathbf{Z}$ bosons because at this low efficiency, all the selected $\mathbf{Z}$ bosons decay into $\mathbf{b}\bar{\mathbf{b}}$	445

4.83	The impact of uncertainties on the jet-mass scale and resolution for 50% (a) and 90% (b) Z efficiency working points of the full boson-type tagger. The point marked HERWIG uses the alternative shower and hadronization model for the simulation, with the likelihood template from PYTHIA	446
4.84	The distribution of p_T^{jet}/p_T^W (left) and $\Delta R(\text{jet}, W)$ (right) for W' and $t\bar{t}$ events. The large tails in $t\bar{t}$ events are due to cases in which more than the W boson hadronic decay products are merged inside the large-radius jet.	448
4.85	The joint distribution of p_T^{jet}/p_T^W and $\Delta R(\text{jet}, W)$ for $t\bar{t}$ events (left) and W' events (right). Dashed lines indicate the selection for the Boosted W category.	449
4.86	The p_T distribution of the selected large-radius jets. The uncertainty band includes all the experimental uncertainties on the jet p_T and jet mass described in Sec. 4.2.6.	450
4.87	(a) The jet-mass distribution of the selected jets in semi-leptonic $t\bar{t}$ events. (b) The median of the mass distribution as a function of the jet p_T for events with the selected jet in the range $50 \text{ GeV} < m^{\text{jet}} < 120 \text{ GeV}$. This includes the contributions from events which are not classified as Boosted W . (c) For the same events as in (b), the inter-quantile range as a measure of spread. The quantiles are centred at the median. The uncertainty band includes all the experimental uncertainties on the jet p_T and jet mass described in Sec. 4.2.6. The inter-quantile range of size $0\% < X < 50\%$ is defined as the difference between the $50\%+X\%$ quantile and the $50\%-X\%$ quantile. Statistical uncertainty bars are included on the data points but are smaller than the markers in many bins.	453

- 4.88 (a) The distribution of r_{track} in the data for semi-leptonic $t\bar{t}$ events with the selected jet in the range $50 \text{ GeV} < m^{\text{jet}} < 120 \text{ GeV}$. (b) The median of the r_{track} distribution as a function of the jet p_T . This includes the contributions from events that are not classified as Boosted W . (c) The inter-quantile range as a measure of the width. The quantiles are centred at the median. The uncertainty band includes all the experimental uncertainties on the jet p_T and jet mass described in Sec. 4.2.6. The inter-quantile range of size $0\% < X < 50\%$ is defined as the difference between the $50\% + X\%$ quantile and the $50\% - X\%$ quantile. Statistical uncertainty bars are included on the data points but are smaller than the markers in many bins. 454
- 4.89 (a) The distribution of the number of tracks associated with the selected large-radius jet in the semi-leptonic $t\bar{t}$ data for events with the selected jet in the range $50 \text{ GeV} < m^{\text{jet}} < 120 \text{ GeV}$. (b) The median of the distribution of the number of tracks as a function of the jet p_T . This includes the contributions from events that are not classified as Boosted W . (c) The inter-quantile range as a measure of the width. The quantiles are centred at the median. The uncertainty band includes all the experimental uncertainties on the jet p_T and jet mass described in Sec. 4.2.6. The inter-quantile range of size $0\% < X < 50\%$ is defined as the difference between the $50\% + X\%$ quantile and the $50\% - X\%$ quantile. Statistical uncertainty bars are included on the data points but are smaller than the markers in many bins. 455

4.90	(a) The distribution of the jet charge in the data for semi-leptonic $t\bar{t}$ events with the selected jet in the range $50 \text{ GeV} < m^{\text{jet}} < 120 \text{ GeV}$. The ratio uses the positive lepton charge. (b) The median of the jet charge distribution as a function of the jet p_T . This includes the contributions from events that are not classified as Boosted W (except for the blue triangles, for which only the Boosted W is included). (c) The inter-quartile range as a measure of the width. The quantiles are centred at the median. The uncertainty band includes all the experimental uncertainties on the jet p_T and jet mass described in Sec. 4.2.6. The inter-quartile range is defined as the difference between the 75% quantile and the 25% quantile. Statistical uncertainty bars are included on the data points but are smaller than the markers in many bins.	456
4.91	The efficiency-binned MV1 distribution for the (a) leading and (b) sub-leading matched small-radius in semi-leptonic $t\bar{t}$ events. If there is no second small-radius jet with $p_T > 25 \text{ GeV}$ and $\Delta R < 1$ to the selected large-radius jet axis, the event is put in the ‘No jet’ category in (b). The uncertainty band includes all the experimental uncertainties on the jet p_T and jet mass and those related to the b -tagging described in Sec. 4.2.6. Statistical uncertainty bars are included on the data points but are smaller than the markers in many bins.	457
4.92	The joint distribution of the jet mass and 2-subjettiness for W boson jets (left) and Z boson jets (right). The peak of the distribution along the jet mass axis is shifted toward higher values for the Z due to its higher mass. The linear correlation is less than $\pm 10\%$ in both cases and the two distributions are nearly independent.	460
4.93	The joint distribution of the jet charge and 2-subjettiness for W boson jets (left) and Z boson jets (right). The linear correlation is less than $\pm 2\%$ in both cases and the two distributions are nearly independent.	460

4.94	An example of the learned per-particle probability density specified in Eq. 4.37 with $k = 3$ and Gaussian $f = \Phi$ in $m = 2$ dimensions. One cluster is associated with each component density $\Phi_i = \Phi(\cdot \mid \mu_i, \Sigma_i)$, where the dot $\cdot$ is a placeholder for the function argument.	464
4.95	An illustration of of the EM algorithm for $k = 2$. The circles represent data points, the triangles represent the estimated cluster locations μ_j , and the ellipsoids are equidensity contours describing the shapes Σ_j of the learned cluster distributions. In the E-step, bluer colors correspond to higher value of $p_{i, \text{blue jet}}$	467
4.96	A $Z' \rightarrow t\bar{t}$ event clustered with fuzzy jets (left) or anti- k_t $R = 1$ (right) without pileup ($n_{PU} = 0$). The circles indicate the 1σ contour (fuzzy jets) or the ungroomed jet area (anti- k_t). The small filled colored circles are the particles, with the color and size indicating their energy. The top quark locations from the generator-record are indicated by red stars. In the left plot, anti- k_t jet locations are shown with gray crosses where the long tail points towards the mGMM jet for which it was a seed. In the right plot, darker discs correspond to higher p_T jets. The highest p_T fuzzy jet has a black 1σ contour while all others are shown in gray.	469
4.97	The jet p_T for the leading anti- k_t jet (left) and leading fuzzy jet under the HML particle assignment scheme (right). All the processes are re-weighted so that the anti- k_t p_T spectra are the same.	471
4.98	The jet mass for the leading anti- k_t (left) and leading fuzzy jet under the HML particle assignment scheme (right), in an anti- k_t leading jet p_T window of 350 to 450 GeV. All the processes are re-weighted so that the anti- k_t p_T distributions are the same. The dashed white lines mark $m_W \approx 80$ GeV and $m_{top} \approx 175$ GeV.	471
4.99	The learned value of σ for all fuzzy jets (left) and for the highest p_T jet under the HML scheme (right).	473

4.100	The joint distribution of the leading fuzzy jet σ and the leading anti- k_t jet m/p_T for Z' events (left) and QCD multijet events (right). The Pearson correlation coefficient is shown in the bottom right of both plots.	473
4.101	The tradeoff between signal efficiency versus QCD multijet rejection ($=1/\text{efficiency}$) when the signal process is $Z' \rightarrow t\bar{t}$ (left) or $W' \rightarrow WZ$ (right). The <i>random</i> tagger line the curve signal efficiency = background efficiency.	474
4.102	The same Z' event with $n_{PU} = 0$ (left) and $n_{PU} = 40$ (right). The grid lines show the 0.1×0.1 tower size and the vertical dashed lines show the range over which the charged pileup energy is subtracted from each tower. The top quark locations from the generator-record are indicated by red stars and anti- k_t jet locations are shown with gray crosses where the long tail points towards the mGMM jet for which it was a seed. The fuzzy jets themselves are represented by their 1σ contour.	475
4.103	The same schematic $k = 3$ per-particle probability density from Fig. 4.94, but with a constant likelihood added to represent the event jet. . . .	476
4.104	The same events as in Fig. 4.102, but with event-jet and tower-based pileup subtraction applied.	477
4.105	The mean and standard deviation of the σ distribution in Z' events as a function of n_{PU} with and without the event jet and tower-based pileup corrections.	477
4.106	The distribution of the leading jet σ using the same event selection as in Sec. 4.2.7. Only the muon channel is included (negligible QCD multijets contribution).	478
4.107	The distributions of the jet mass (left), 2-subjettiness (τ_{21}) (middle) and the ΔR between subjets (right) for signal (blue) and background (red) jets.	480

4.108	The average jet image for signal W jets (top) and background QCD jets (bottom) before (left) and after (right) applying the rotation, re-pixelation, and inversion steps of the pre-processing. The average is taken over images of jets with $240 \text{ GeV} < p_T < 260 \text{ GeV}$ and $65 \text{ GeV} < \text{mass} < 95 \text{ GeV}$	482
4.109	The distribution of the image mass after various states of pre-processing for signal jets (left) and background jets (right). The naive translation and the I^2 normalization image masses are both multiplied by constants so that the centers of the distribution are roughly in the same location as for the other distributions.	484
4.110	The tradeoff between W boson (signal) jet efficiency and inverse QCD (background) efficiency for various pre-processing algorithms applied to the jet (images).	485
4.111	Left: The distribution of the fraction of pixels (occupancy) that have a nonzero entry (blue) or at least 1% of the scalar sum of the pixel intensities from all pixels (red). Right: A typical single W jet image.	486
4.112	Left: ROC curves for individual physically motivated features as well as three deep neural network discriminants. Right: the DNNs are compared with pairwise combinations of the physically motivated benchmarks.	487
4.113	ROC curves that combined the DNN outputs with physics motivated features MaxOut architecture.	488
4.114	The distribution of τ_{21} (left), ΔR (middle) and the jet mass conditioned on the CNN output for background jets. Distributions for the MaxOut network look qualitatively the same.	489
4.115	$W' \rightarrow WZ$ (top) and QCD (bottom) average jet-images in three small windows of τ_{21} : $[0.19, 0.21]$ (left), $[0.39, 0.41]$ (middle), and $[0.59, 0.61]$ (right). In all cases, jet mass is restricted to be between 79 GeV and 81 GeV and the jet p_T is required to be in the interval $[250, 260] \text{ GeV}$	490

4.116	The average difference between $W' \rightarrow WZ$ jet-images in same small windows of τ_{21} as Fig. 4.115. Red (blue) colors are more signal- (background-)like.	491
4.117	Pearson correlation coefficient for pixel intensity and the CNN output for $W' \rightarrow WZ$ and QCD (combined) in the same small windows of τ_{21} as Fig. 4.115.	491
4.118	ROC curves including the jet pull angle θ_P in a restricted phase space (left) and using an inclusive event selection, but with events weighted so that the joint jet mass and τ_{21} distributions are identical between signal and background.	492
1.1	A one-loop diagram contributing to the Higgs boson self energy at next-to-leading order.	510
1.2	SUSY NLO corrections to the Higgs boson self energy. While the right diagram is topologically the same as the leading NLO fermion diagram in Fig. 1.1, it is suppressed with respect to the left diagram due to the two scalar propagators ($\sim \int d^4k/k^4$) compared with one ($\sim \int d^4k/k^2$). There are also two powers of the coupling constant for the right diagram, but the Yukawa coupling is nearly one for the top/stop.	511
1.3	The one-loop renormalization group evolution of the inverse couplings as a function of energy. The $U(1)$ coupling is scaled by $\sqrt{5/3}$, which is the quantity that unifies with the others in grand unified theories. .	516
1.4	The Feynman diagram that is the source of the chiral anomaly in the SM. Any electroweakly interacting fermion contributes to the loop. .	518
1.5	The Feynman diagram illustrating proton decay with RPV couplings.	519
1.6	An example CMSSM model specified with sfermion mass $m_0 = 400$ GeV, gaugino mass $m_{1/2} = 500$ GeV, zero trilinear couplings, $\tan(\beta) > 0$ and a positive μ . The running of the masses is calculated at NLO using SOFTSUSY 3.7.1 [455].	523

1.7	Left: the number of electroweak scale spectra consistent with the CMSSM parameters as a function of m_0 and $M_{1/2}$ for fixed $A_0 = 0, \tan \beta < 0$ and $\mu < 0$. Right: the difference in select phenomenological parameters along the strip of two solutions from the left plot just below where A^0 is tachyonic. Sparticle masses are nearly identical between the two spectra, but stop and sbottom branching ratios vary by more than 10% and the predicted dark matter relic density differs by more than 100% for $M_{1/2} \sim 840$ GeV. See Ref. [460] for more detail.	524
1.8	One model from a recent ATLAS summary [463] of SUSY searches interpreted in the context of a random scan in the pMSSM [462, 464–466]. This particular model is not excluded by direct stop searches but is ruled out by other searches with a broader scope. One reason the direct searches do not rule out this model, and one powerful use of the pMSSM, is the model complexity: there are multiple light neutralinos/charginos with cascade decays. This plot was created with PySLHA [467]. All sparticles not shown are heavier than 700 GeV.	525
1.9	The maximum Higgs mass $\sqrt{m_Z^2 c_{2\beta}^2 + \Delta(m_h^2)}$ from Eq. 1.39 for no stop mixing $\theta_{\tilde{t}} = 0$ (left) and mixing that maximizes the correction (right).	528
1.10	Schematic diagrams illustrating models with similar signatures as stop pair production (top left). Top right: the gluino mediated stop process with nearly mass-degenerate stop and neutralino such that the stop decay products are too soft to be reconstructed. Bottom left: pair production of vector-like quarks T' decaying into Z bosons that decay into neutrinos that are undetected just like neutralinos. Bottom right: leptoquarks carry both lepton and baryon numbers and can decay to a top quark and a neutrino. All of these models are constrained by the search presented in Part III.	531

2.1	Left: The stop pair production cross section at $\sqrt{s} = 8$ TeV and $\sqrt{s} = 13$ TeV as a function of the stop mass compared with important SM background processes $t\bar{t}$ and $t\bar{t} + Z(\rightarrow \nu\bar{\nu})$. The SUSY cross section is calculated at NLO+NLL [502, 503] while the $t\bar{t}$ cross section has NNLO+NNLL accuracy using TOP++2.0 [218] with the PDF4LHC prescription [504]. The $t\bar{t} + Z$ cross section is computed at NLO from Ref. [505]. Right: the $t\bar{t}$ cross section as a function of the particle-level E_T^{miss}	534
2.2	The average invariant mass of the b-quark and SM fermions from the decay $\tilde{t} \rightarrow b f f' \tilde{\chi}_0^1$ as a function of the stop mass and neutralino mass.	535
2.3	A leading order Feynman diagram for the simplified stop model considered in Part III. The pair production of stops is followed by the subsequent decay to top quarks and neutralinos. The top quark decays nearly 100% of the time into a W boson and a b-quark. One of the W bosons decays leptonically and the other decays hadronically (into quarks).	536
2.4	For the kinematically tightest event selection from the Run 1 stop search in the one-lepton channel (see Sec. 4), the acceptance is plotted as a function of the top quark momentum given by Eq. 2.2. The acceptance is defined as the fraction of simulated signal events that pass a particle-level version of the analysis (no detector simulation). The ratio of the particle-level acceptance to the acceptance using the full detector simulation are all within a few percent of one in the relevant region and so can be safely ignored. In the region beyond 600 GeV, the acceptance is well described by a straight line $\epsilon = 0.027 \frac{\%}{\text{GeV}} \times p - 4.3\%$. Every point with the same color corresponds to a model with the same stop mass. The spread in the top/neutralino (LSP) momentum is due to the spread in the neutralino masses. . . .	537

2.5	The top (and neutralino) momentum given by Eq. 2.2 as a function of the stop mass and neutralino mass. The momentum is normalized to one at $m_{\text{LSP}} = 0$ for a direct comparison of the acceptance different stop masses.	538
2.6	A combination of the acceptance from Fig. 2.5 and the cross section from Fig. 2.1. The z-axis is normalized to unity at $(m_{\text{stop}}, m_{\text{LSP}}) = (400, 0)$	538
2.7	Diagrams indicating the various spin configurations of the stop decay products. Single black arrow indicate the direction of the momentum in the rest frame of the particle above the origin of the arrows and double-lined colored arrows indicate the production of the spin along the momentum direction. There are two other possible collinear diagrams per panel where the top begins moving to the right, but the conclusions are the same as for these two.	539
2.8	The dependence of $\cos(\theta_{\text{eff}})$ on the stop mixing angle for a pure bino (black dashed) and a pure higgsino (red) LSP.	540
2.9	The charged lepton p_T spectrum in stop decays for three values of θ_{eff} for $(m_{\text{stop}}, m_{\text{LSP}}) = (500, 0)$	541
2.10	The ratio of the efficiency for a $p_T^{\text{charged lepton}} > 25$ GeV requirement between the mostly left-handed top quark ($\theta_{\text{eff}} = 0$) and the mostly right-handed top quark ($\theta_{\text{eff}} = \pi/2$).	542
2.11	For a fixed stop mass $m_{\text{stop,R}}$ (LSP is massless), plotted is $m_{\text{stop,R}} - m_{\text{stop,L}}$ for $\sigma(m_{\text{stop,L}}) \times \epsilon_L = \sigma(m_{\text{stop,R}}) \times \epsilon_R$. The value $m_{\text{stop,R}}$ is the stop mass for $\theta_{\text{eff}} = \pi/2$ and $m_{\text{stop,L}}$ is the stop mass for $\theta_{\text{eff}} = 0$, so the plot shows how much the stop mass needs to be reduced to have the same predicted number of events with a most left-handed top quark configuration compared to a given mostly right-handed top quark configuration.	542

- 2.12 Left: Unit normalized distributions of the m_{jjj} variable for $t\bar{t}$ with $m_t = 172.5$ GeV, and for $\tilde{t}$ pair production with a two-body $\tilde{t} \rightarrow tN$ decay with $m_{\tilde{t}} = 175$ GeV (and $m_t = 172.5$ GeV), and a three-body decay $\tilde{t} \rightarrow bWN$ for $m_{\tilde{t}} = 170$ GeV. The neutralino is assumed massless. Jets are assigned to the hadronically decaying top quark by minimizing $\chi^2 = (m_{j_1 j_2 b_1} - m_{b_2 l\nu})^2 / (20 \text{ GeV})^2 + (m_{j_1 j_2} - m_W)^2 / (10 \text{ GeV})^2$ for $\{j_i\}$ the set of jets not identified as originating from a b -quark. Even though the neutrino p_z is unmeasured, it can be inferred by solving $m_{l\nu} = m_W$. The simulations are performed using HERWIG++ 2.7 [195, 335] and analyzed using the RIVET 1.8.2 framework [297] with FASTJET 3.0.6 [203] for clustering anti- k_t jets with $R = 0.4$ [117]. Right: Calibration curve that relates the measured value $\langle m_{jjj} \rangle$ to the (MC) top quark mass, m_{top} in $t\bar{t}$ events. See Ref. [525] for more detail. 546
- 2.13 Summary of the effects leading to the sneaky stop scenario: the shifts in the measured $t\bar{t}$ cross-section and measured top quark mass. The solid line corresponds to an unbiased measurement of the $t\bar{t}$ cross-section as a function of the top quark mass. The dot-dashed line is what would be measured in the presence of a $\tilde{t} \rightarrow bWN$ with $m_{\tilde{t}} = 170$ GeV for an unbiased top quark mass measurement. However, under the SM+SUSY hypothesis the top quark mass measurement would be *biased* which translates into what would actually be observed shown in the dashed line. For all three lines, the band reflects the $\sim 5 - 6\%$ theory uncertainty on the cross-section. For comparison, the measured top quark mass and $t\bar{t}$ cross-section are shown from recent CMS [526] and ATLAS [520] results. 547
- 2.14 The measured top quark mass as a function of the true top quark mass. The bias in the measurement arises from the presence of a light $\tilde{t}$ with $m_{\tilde{t}} = 170$ GeV and decaying via the three-body process (left) or with $m_{\tilde{t}} = 175$ GeV and decaying via the two-body process. . . . 548

2.15	A schematic diagram to illustrate the control region method. The red distribution represents $N_B p_B(V)$ and the blue is $N_S p_S(V)$ (see the text for details). The left and right distributions are identical, only with a logarithmic scale for the vertical axis in the right plot. Control, validation, and signal regions are denoted by CR, VR, and SR, respectively. The symbol S is shorthand for $N_S p_S(V)$ and B represents $N_B p_V(B)$	550
2.16	The single bin control and validation region setup for the $\sqrt{s} = 13$ TeV analysis. All regions share a common selection and then are further distinguished by requirements on m_T , am_{T2} and the number of b-jets. The prefixes stand for ST = single top, T = $t\bar{t}$, W = W +jets, and TZ = $t\bar{t} + Z$. The exact definitions of the regions are given in Sec. 5.	551
2.17	The tradeoff between Type II (do not reject SM when SUSY is true) and Type I errors (reject SM when SM is true) for three scenarios. In all cases, the number of background events is 200 in one bin and 100 in a second bin. The total number of signal events is 30. In the left plot, both bins have 10% signal; in the middle plot the fraction is twice as high in the first bin; in the right plot, the fraction is four times as high in the first bin compared with the second bin. The errors are computed by scanning threshold requirements on the (log) likelihood ratio distribution. For the two-bin case, the likelihood ratio is a product of the individual bin likelihood ratios. Only statistical uncertainties are included in the likelihood. The vertical dashed line is at 5%.	552

2.18	The various bins of the shape fit signal region tN_diag from the $\sqrt{s} = 8$ TeV analysis. All the bins share a common selection and are distinguished by requirements on m_T , E_T^{miss} and the number of b-jets. The red numbers are the SM prediction, the blue numbers are the observed events, and the black numbers are the predicted signal yield. The top right bin is the most signal sensitive bins, the lowest row is the most sensitive to W +jets and the row just above that is the most sensitive to the $t\bar{t}$ normalization. These last two regions behave similarly to the control-regions from the one bin regions and as such the fitted background yield is nearly identical to the observed data. The selection is described in Sec. 4.4 and the implications for light stops are discussed in Sec. 7.4.2.	554
2.19	The probability for a 2σ excess when there is only background as a function of the number of expected events in the control region and in the signal region using the standard control region method. A 2σ excess is defined as a case when the probability for the observed number of events in the signal region to exceed the number of predicted events in the signal region to be less than 5%. Only statistical uncertainties are included in this calculation.	556
2.20	The probability for a 3σ excess (with ∞ events in the control region) to be measured as a $< 3\sigma$ excess.	556
2.21	The probability for a 2σ excess under the null (only known particles) hypothesis as a function of the number of expected events in the control region and in the signal region using the Bayesian control region method described in the text. A 2σ excess is defined as a case when the probability for the observed number of events in the signal region to exceed the number of predicted events in the signal region to be less than 5%. A systematic uncertainty of 30% is assumed when choosing β	558

2.22	A graphical representation of the hierarchical Bayes model described in the text. For model M_j , the predicted background yield (N_B^{MC}) is denoted λ_i^j	558
2.23	Sample posteriors using the hierarchical Bayes model described in the text. The nominal simulation predicts $(N_{CR}^{raw}, N_{SR}^{raw}) = (10, 5)$ and the two alternative models have $(N_{CR}^{raw}, N_{SR}^{raw}) = (10.5, 5.5)$ and $(N_{CR}^{raw}, N_{SR}^{raw}) = (9.5, 4.5)$. The number of observed events in the control region is 11 and the systematic uncertainty, used to construct β is 30%.	559
2.24	Left: the trigger efficiency for various algorithms as a function of the stop mass. Right: The additional efficiency provided by various triggers beyond the efficiency already provided by the triggers to the left of it in the legend.	560
3.1	The joint distribution of the ΔR between the lepton and b -quark from the same top quark decay for $\tilde{t} \rightarrow t\tilde{\chi}^0$. At $m_{stop} = 500$ GeV, about 2% (10%) of events have $\Delta R < 0.4$ (0.8) and at $m_{stop} = 800$ GeV, this fraction increases to $\sim 5\%$ (25%).	567
3.2	Left: The transverse mass distribution for the discovery of the W boson by the UA1 collaboration [59]. Right: the transverse mass distribution used by the CDF collaboration for one of the most precise determinations of the W boson mass [557].	573
3.3	The distribution of m_{T2} in dilepton $t\bar{t}$ events using the leptons (b -tagged jets) as the visible particles in the left (right) plot. These distributions are used to measure the top quark, W boson, and neutrino masses. Published in Ref. [384].	574

3.4	The distribution of the transverse mass for the isotropic decay of a W boson for $p_T^W = 0$ GeV on the left and $p_T^W = 200$ GeV on the right. The red histogram shows the m_T distribution for the decay $W \rightarrow e\nu$ and the blue histogram shows the distribution for $W \rightarrow \tau\nu$ where the τ decays to an electron and neutrinos. The black histograms correspond to the pair production of $W \rightarrow e\nu$ where the second lepton is either included (lost) or not included (mis-id) in the E_T^{miss} . The dotted line is the analytic formula derived in the text.	578
3.5	The transverse mass distribution for a scalar W decay (spin 0) and a vector W decay (spin 1). Each distribution is generated by MG5_aMC 2.1.1. A Higgs boson with $m_H = m_W$ is used for the spin 0 line, which is identical to the lines from Fig. 3.4 using a phase-space only generator and the analytic formula.	579
3.6	The distribution of the transverse mass for signal $\tilde{t} \rightarrow t\tilde{\chi}^0$ events and dilepton $t\bar{t}$ events where one lepton is lost or mis-identified. The signal histograms are identical in the left and right plots. In the left plot, the p_T of the $t\bar{t}$ system is boosted in an arbitrary direction with a fixed p_T . In the right plot, the $t\bar{t}$ pair is produced at rest, but with a large $m_{t\bar{t}}$ so that the top quarks have significant boost.	580
3.7	The linear correlation between the E_T^{miss} and m_T as a function of the p_T or mass scale for various processes. The $p_T^{t\bar{t}}$ and $m_{t\bar{t}}$ are setup in the same way as for Fig. 3.20. The correlation is set to zero for unphysical parameter values such as $m_{t\bar{t}} < 2m_t$	581
3.8	A comparison of data and simulation using a loose selection requiring exactly one signal lepton, four jets with $p_T > 25$ GeV and at least one b -tagged jet. The left and right plots differ only in the scaling of the vertical axis. The uncertainty band includes jet energy scale and resolution uncertainties (see Sec. 6).	582

3.9	A comparison of data and simulation using a loose selection requiring exactly one signal lepton, four jets with $p_T > 25$ GeV and at exactly no b -tagged jets. The left and right plots differ only in the scaling of the vertical axis. The uncertainty band includes jet energy scale and resolution uncertainties (see Sec. 6).	583
3.10	Identical m_T distributions as in Fig. 3.4, except the neutrino mass parameter in the transverse mass calculation is set to 20 GeV (left) and 50 GeV (right)	584
3.11	A schematic diagram showing the generic m_{T2} setup.	586
3.12	A comparison of m_T and m_{T2} for events with an identified hadronically decaying τ	587
3.13	An illustration of the setup for the projected m_{T2} variable. The upstream momentum is the sum of all particles not associated with the visible particles V_i and child particles C_i . In this $t\bar{t}$ topology, the upstream momentum is due to the b -jets directly from the top quark decay and any FSR produced before the $t\bar{t}$ prouction.	590
3.14	Diagrams illustrating the two possible orientations of projected transverse momenta (top) and the corresponding $m_{T,\perp}$ graphs for both branches of the decay (bottom). The circles represent a transverse cross-section of the detector - the beam is into and out of the page. In the left configuration, the visible transverse momenta are on opposite sides of the upstream momentum and the resulting $m_{T2,\perp}$ value is equal to zero. In the right diagrams, the visible momenta are on the same side of the upstream momentum which allows for a finite $m_{T2,\perp}$ value.	591
3.15	Graphs illustrating the two possible configurations leading to the value of m_{T2} at the intersection of the two m_T conic sections (left) or at the minimum of one of the sections, if it is above the other curve (right).	593

3.16	The fraction of events for which the m_{T2} value is $m_{V_1} + m_{C_1}$ (unbalanced case) for $t\bar{t}$ events as a function of m_{C_1} and m_{C_2} . The visible object V_1 is the four-vector sum of the b -quark and lepton from the same top quark decay and V_2 is the b -quark from the other top quark decay. A white star indicates the choice for the am_{T2} variable described in Sec. 3.2.1.2.3.	593
3.17	Three views of the m_T^2 surfaces from the same event, rotated so illustrate the structure of the intersection. The value of m_{T2}^2 is indicated in the middle graph by a white arrow.	595
3.18	Level surfaces of the graphs from Fig. 3.17. As noted above, these curves are ellipses and the point at which they are tangent (indicated by a black star) corresponds to the value of m_{T2}	596
3.19	A demonstration of the numerical accuracy of the m_{T2} calculator based on the overlap of conic sections instead of the intersection of their boundaries. Plotted is the difference d between the numerical calculation and the analytic formula using the projected m_T from Sec. 3.2.1.2.1 in $t\bar{t}$ events for the variable am_{T2} described in Sec. 3.2.1.2.3. The scale of $m_T \sim \mathcal{O}(100 \text{ GeV})$. See Ref. [570] for more detail.	597
3.20	A schematic diagram of dileptonic $t\bar{t}$ decay where one of the charged leptons is lost. Lost particles are circled with a dashed line. For the am_{T2} variable, the visible particle on the top (bottom) branch is the b -jet (sum of the b -jet and lepton). The missing particle in the top (bottom) branch is a W boson (neutrino).	599
3.21	The distribution of am_{T2} for dileptonic $t\bar{t}$ and stop events with various configurations as indicated in the legend. The p_T spectrum of the top quarks in $t\bar{t}$ events is chosen to be identical to the distribution for stop events.	600

3.22	A comparison of data and simulation using a loose selection requiring exactly one signal lepton, four jets with $p_T > 25$ GeV, at least one b -tagged jet, and $m_T > 100$ GeV to enrich the dilepton contribution. The b -tag weight is used to select the two b -jet candidates and the minimum value over both b -lepton pairings is used to resolve the matching ambiguity. The uncertainty band includes jet energy scale and resolution uncertainties (see Sec. 6).	601
3.23	A dileptonic $t\bar{t}$ event with one hadronically decaying τ . Lost particles are circled with a dashed line. For the m_{T2}^τ variable, the visible particle on the top (bottom) branch is the hadronic τ (reconstructed lepton). The missing particle in the top (bottom) branch is the sum of τ neutrinos (e or μ neutrino).	601
3.24	The distribution of m_{T2}^τ in dilepton $t\bar{t}$ events with top quark p_T spectra that match that of a 800 GeV stop and $m_{LSP} = 0$. For the filled histogram, the τ is decayed using a scalar three-body phase space to mimic $\tau \rightarrow \nu_\tau \pi^0 \pi^\pm$. The ‘pions’ are added to form the visible τ and the ν_τ is added to the $\vec{p}_T^{\text{miss}}$	602
3.25	A diagram of dileptonic the $t\bar{t}$ decay chain and kinematic variables targeting various components of the chain. The dotted lines highlight which aspects of the top quark decay are involved with the construction of the variables with the same color. Particle inputs to each variable are circled with the same colors.	604
3.26	The distributions of m_T (top left), E_T^{miss} (top right), $\mathbf{a}m_{T2}$ (bottom left), and $\mathbf{a}m_{T2,\perp}$ (bottom right) for $t\bar{t}$ and stop events with $(m_{\text{stop}}, m_{LSP}) = (500, 1)$. The separation (<i>sep</i>) is defined in Eq. 3.20. The b -tagging weight (MV1) is used to select the b -jets for $\mathbf{a}m_{T2}$ and a ΔR scheme matches the b -jet with the lepton.	605
3.27	The distribution of $\mathbf{a}m_{T2}$ for dileptonic $t\bar{t}$ where one lepton is lost and a three-body stop decay (left) with $(m_{\text{stop}}, m_{LSP}) = (200, 100)$ GeV and for a flavor-changing $\mathbf{b} + \tilde{\chi}^\pm$ decay and flavor-neutral two-body decay $\mathbf{t} + \tilde{\chi}^0$ (right) with $(m_{\text{stop}}, m_{\text{chargino}}, m_{LSP}) = (800, 300, 150)$ GeV.	608

3.28	The distribution of m_{T2} in diboson ($ZZ/WW \rightarrow l^+l^-\nu\bar{\nu}$) and stop events using the selected lepton as one visible particle and the next highest p_T muon for the second visible particle. For the stop model, $(m_{\text{stop}}, m_{\text{LSP}}) = (800, 1)$ GeV.	609
3.29	The distribution of $m_{T2}^{\tilde{t}}$ (left) and αm_{T2} (right) for $t\bar{t} + Z(\rightarrow \nu\bar{\nu})$ and $\tilde{t}\tilde{t}$ events with $(m_{\text{stop}}, m_{\text{LSP}}) = (800, 1)$ GeV. The left plot also includes the distributions using particle-level top quarks for the visible particles.	610
3.30	The joint distribution of $m_{T2}^{\tilde{t}}$ and αm_{T2} for stop events (left) and $t\bar{t} + Z(\rightarrow \nu\bar{\nu})$ events (right).	610
3.31	The green and red annuli represent the ATLAS calorimeters and the yellow arrows represent the measured jet directions. The direction of the right jet is mis-measured leading to an apparent E_T^{miss} represented by the dashed arrow. Such events could be identified if the jet kinematic covariance matrices show that the jets are statistically consistent with being back-to-back.	611
3.32	The Gumbel distribution and an equivalent Gaussian distribution.	617
3.33	Contours of constant $V_{\text{opt}}^{(\text{Gumbel})}$ in the (m, β) plane for $M_b = 80$ and $M_s = 85$. The contours are drawn at powers of 2 from $-2^3, \dots, -2, 0, 2, \dots, 2^8$	618
3.34	The dependance of X_M on M for two events with $\Delta \equiv (m_2\sigma_1 - m_1\sigma_2)/(\sigma_1 - \sigma_2) > 0$ in the left plot and $\Delta \in [-\infty, 0) \cup \{\infty\}$ in the right plot.	619
3.35	These plots illustrate the distributions of m , X_M and $\hat{s}$ for a simple model in which m is always ‘on shell’ at 80 for the background and 90 for the signal. The resolutions can take one of two values with probability 1/2, independent of the physics process. The optimal (V_{opt}) variable is based on the likelihood ratio detailed in Sec. 3.2.2.2 and Eq. 3.26.	620

3.36	In each row, the left plot compares the transverse mass distribution for a Standard Model W and a W' with mass 100 GeV. The middle plot is the corresponding distributions of X_M with $M = 80$ GeV. The right plot shows the rejection $s\sqrt{b}$ as a function of the signal efficiency, in arbitrary units. The bands show the statistical uncertainty due to limited Monte Carlo statistics. The top row has a boson mass width of 0, the middle has a width of 20% of the natural width, and the bottom row has the full width of about 2 GeV [27]. We can see that for this fixed value of M , the performance of X_M is better than m_T for a narrow width and then worse at higher width. By construction, X_M cannot be worse than m_T and thus the optimal M in the last row must be different than 80. The inset plot shows X_M for $M = 100$, for which the performance of X_M and m_T is the same.	622
3.37	The left plot is the m_T distribution for dileptonic $Z \rightarrow \tau\tau$ and $H \rightarrow \tau\tau$ for a 125 GeV Higgs. The middle plot is the corresponding X_M curve with $M=70$ and the right plot is the rejection versus efficiency relationship.	623
3.38	The left plot is the m_{T2} distribution for dileptonic $t\bar{t}$ and $\tilde{t} \rightarrow t + \text{LSP}$ for a 350 GeV stop and 170 GeV LSP. The middle plot is the corresponding X_M curve with $M=80$ and the right plot is the rejection versus efficiency relationship.	623
3.39	The distributions of m_T (left) and X_{m_T} (with $M = 100$ GeV) (right) for $t\bar{t}$ and stop pair production.	625
3.40	A comparison of m_T (left), σ (middle), and X_{m_T} (right) in data and simulation at $\sqrt{s} = 8$ TeV.	625
3.41	A comparison of the statistical significance after a threshold requirement on m_T and the (approximate) m_T significance, X_{m_T} . The background is $t\bar{t}$ and the signal is a stop model with $(m_{\text{stop}}, m_{\text{LSP}}) = (350, 150)$ GeV. This statistical significance is the ‘discovery significance’ because it quantifies the number of standard deviations a signal would be above the background-only noise.	626

3.42	The joint distribution of m_T and X_{m_T} with $M = 100$. When $m_T = 100$, $X_{m_T} = 0$ by construction.	626
3.43	The distribution of E_T^{miss} conditioned on H_T^{miss} in bins of 20 GeV for each variable in $t\bar{t}$ events. All events have $E_T^{\text{miss}} > 100$ GeV.	628
3.44	The joint distribution of the approximate E_T^{miss} resolution, $\sqrt{H_T}$ and the H_T^{miss} resolution computed via Eq. 3.33.	628
3.45	A comparison of data and simulation using a loose selection requiring exactly one signal lepton, four jets with $p_T > 25$ GeV and at least one b -tagged jet. The left and right plots differ only in the scaling of the vertical axis. The uncertainty band includes jet energy scale and resolution uncertainties (see Sec. 6).	629
3.46	The distributions of $H_{T,\text{sig}}^{\text{miss}}$ with $M = 50$ GeV (bottom left), $M = 100$ GeV (top left), and $M = 0$ GeV (top right) along with the distribution of $E_T^{\text{miss}}/\sqrt{H_T}$ (bottom right) for $t\bar{t}$ and stop events.	630
3.47	The joint distribution of $H_{T,\text{sig}}^{\text{miss}}$ or $E_T^{\text{miss}}/\sqrt{H_T}$ with E_T^{miss} . The linear correlation is indicated in each plot.	631
3.48	The pseudorapidity (left) and p_T (right) distributions of hadronically decaying τ leptons from W decays in $t\bar{t}$ events decomposed by how the τ is reconstructed (if at all). A τ lepton in the simulation is <i>matched</i> to a jet if $\Delta R < 0.4$. If the hadronic decay products of the τ constitute $\leq 90\%$ of the jet p_T , the τ is labeled as <i>hiding in a jet</i> . Events are required to have exactly one signal lepton, at least four signal jets, at least one b -tagged jet, $E_T^{\text{miss}} > 200$ GeV, $m_T > 150$ GeV, and at least one large-radius jet with $p_T > 150$ GeV (see Sec. 3.2.4).	632
3.49	The frequency for a hadronically decaying τ lepton to be reconstructed as one of a variety of objects, described by the labels on the horizontal axis.	634
3.50	Given an event has a hadronically decaying τ , the frequency that a given τ identification technique selects the τ correctly (left) or incorrectly (right).	635

3.51	Given an event does not have a hadronically decaying τ , the frequency that a given τ identification algorithm produces a reconstructed τ candidate.	635
3.52	A summary figure with all the considered hadronic τ vetoes based on the above considerations. The (black) red line shows the efficiency for $t\bar{t}$ events with(out) a hadronically decaying τ lepton. The signal $(m_{\text{stop}}, m_{\text{LSP}}) = (800, 1)$ GeV is has mostly semi-leptonic $t\bar{t}$. The ΔR is between the hadronically decaying τ candidate and the leading large-radius jet. In addition to the loose, medium, and tight reco τ working points, a simple $n_{\text{track}} = n_{\text{trk}} < 5$ identification scheme is part of the comparison. The dashed line is at an efficiency of 90%. .	636
3.53	The distribution of m_{T2}^{τ} with an event selection enriched in dilepton $t\bar{t}$ events with a hadronically decaying τ . Events are required to have at least one reconstructed reco τ passing the loose identification. In addition, events must have at least four jets with $p_T > 80, 50, 40, 25$ GeV, $E_T^{\text{miss}} > 200$ GeV and at least one b-tagged jet. In order to suppress semi-leptonic $t\bar{t}$ events, $m_T > 100$ GeV. See Sec. 5.1.1.1 for details.	637
3.54	A comparison of the two resolved hadronic top quark mass tagging techniques described in the text. The stop model is $(m_{\text{stop}}, m_{\text{LSP}}) = (600, 250)$ GeV.	640
3.55	The large-radius jet mass in a one-lepton $t\bar{t}$ selection requiring $\Delta\phi(\text{jet}_i, \vec{p}_T^{\text{miss}}) > 0.4$, at least four jets with $p_T > 80, 50, 40, 40$ GeV, $E_T^{\text{miss}} > 200$ GeV, $m_T > 30$ GeV, at least one b-jet, at least one large-radius jet with $R = 1.2$ with $p_T > 200$ GeV and $m_{\text{jet}} > 50$ GeV. Leptons are not included in the re-clustering and small-radius jets are dropped (trimmed) if they have $p_T < 5\%$ of the jet p_T . The jet mass distribution is decomposed by constituent multiplicity.	640

4.1	An overview of the various signal regions described in Chapter 4. The horizontal axis is the stop mass and the vertical axis is the neutralino mass; each point in this plane corresponds to a simplified stop model. Model cross-sections are set by the stop mass and decrease from left to right. The dashed line corresponds to the kinematic boundary above which stop decay to an on-shell top quark is forbidden. The various lines and shaded region are the expected exclusion limits using the statistical procedures documented in Sec. 7.1. Seven signal regions are optimized to ensure a broad sensitivity across the plane. Names of the signal regions are placed in the locations of parameter space where they add the most unique sensitivity.	642
4.2	The distribution of the leading (top left), second leading (top right), third leading (bottom left) and fourth leading (bottom right) jet $\mathbf{p}_T$ spectra after the preselection at $\sqrt{s} = 13$ TeV. All plots are normalized to unity. Note that the horizontal axes have been chosen per distribution so that the distributions fill out the entire plot.	644
4.3	An example scan in the number of signal events. For each fixed S , the background yield is minimized and the CL_s value is computed (see Sec. 7.1 for details).	647
4.4	The distribution of the E_T^{miss} for the event preselection (top) and after all SR13 requirements except E_T^{miss} (bottom). Figures with all signal region requirement but the one displayed are called N-1 plots. Both the total background and signal yields are normalized to unity in the right plots. An arrow indicates the signal region requirement. The first number in parenthesis is the expected yield without applying any normalization factors and the second number is the raw event count in simulation (an indication of the statistical uncertainty). The Wt component includes all single top processes, but is dominated by the single production of a top quark in association with a W boson. . . .	651
4.5	The same as Fig. 4.4, but for m_T instead of E_T^{miss}	652
4.6	The same as Fig. 4.4, but for $a m_{T2}$ instead of E_T^{miss}	653

4.7	The distribution of the m_T for the event preselection (top) and after all SR13 requirements except m_T (bottom). Both the total background and signal yields are normalized to unity in the right plots. An arrow indicates the signal region requirement.	654
4.8	The same as Fig. 4.4, but for $H_{T,\text{sig}}^{\text{miss}}$ instead of E_T^{miss}	655
4.9	The distribution of the large-radius $R = 1.2$ jet mass for the event preselection (top) and after all SR13 requirements except the jet mass (bottom). Both the total background and signal yields are normalized to unity in the right plots. An arrow indicates the signal region requirement.	657
4.10	The distribution of the large-radius $R = 1.2$ jet p_T after all SR13 event selections except the large-radius jet p_T . For large-radius jets with $p_T > 500$ GeV, the trimming actively removes low p_T signal jets constituents.	657
4.11	Top: the $\Delta\phi$ between the leading large radius jet and the p_T^{miss} . Bottom: the same as the top, but for the subleading large radius jet if one exists. Large radius jets are required to have $p_T > 150$ GeV, $m > 50$ GeV, and $ \eta < 2.5$	658
4.12	Same as Fig. 4.10, but with the ΔR between the highest p_T b-jet and the selected lepton instead of the jet mass.	659
4.13	The normalized distribution of the $m_{T_2}^{\tau}$ for the event preselection (left) and after all SR13 requirements except m_T (right). An arrow indicates the signal region requirement in the right plot.	660
4.14	The b-tagged jet multiplicity for W +jets events in tNhigh broken down by jet flavor. The defining efficiency of the b-tagging working point is evaluated in inclusive $t\bar{t}$ events.	662
4.15	The distribution of the number of b-tagged jets at preselection (left) and in SR13 (right).	662

4.16	The median significance = $\Phi^{-1}(1 - \text{CL}_s)$ (see Sec. 7.1) as a function of the number of background and signal events, assuming zero systematic uncertainty. The numbers show the gain in significance when weighting each event by 2 if there are ≥ 2 b-jets. With this simple setup, 2 is nearly optimal, though it depends on s and b (see Sec. 2.2.2). The relative frequencies of b-tagged jets is taken from Fig. 4.15.	663
4.17	The distribution of E_T^{miss} , m_T , $a m_{T2}$, and $m_{\text{had}}^{\text{top}}$ with a one-lepton, four jets at $p_T > 80, 60, 40, 25$ GeV and $E_T^{\text{miss}} > 100$ GeV preselection.	665
4.18	The distribution of E_T^{miss} , $a m_{T2}$, and $m_{\text{had}}^{\text{top}}$ after the preselection from Fig. 4.17 with the additional requirement $m_T < 50$ GeV.	666
4.19	The yields for the preliminary version of tN_diag. An analogous table for the final version appears in Fig. 2.18. All background yields are after the post-fit - see Sec. 7.1 for details.	669
5.1	Left: The yields in each signal region broken down by background type after the CR-only fit setup described in Sec. 7.1. Right: yields are normalized to unity to show the fractional background composition. The numbers in the right plot show the fraction of non- $t\bar{t}$ background. tNdiagxy is the $(x + 1)^{\text{th}}$ E_T^{miss} bin and $(y + 2)^{\text{th}}$ m_T bin of the shape fit region.	671
5.2	Optimization of the $t\bar{t}$ CR associated with tNhigh. The left plot shows the control region composition for various fractional E_T^{miss} thresholds after changing the m_T window and removing the m_T requirements. The signal region requirement on E_T^{miss} is 320 GeV. The $H_{T,\text{sig}}^{\text{miss}}$ threshold is varied coherently with the E_T^{miss} requirement.	674
5.3	The number of signal events in TCR13 as a function of the stop mass and neutralino mass. The benchmark model for SR2 has a stop mass of 800 GeV and a neutralino mass of 0 GeV. The number of fitted SM events in TCR13 is 102.	675

5.4	Feynman diagrams illustrating dilepton $t\bar{t}$ events passing the four-jet selection. In the left diagram, one of the electron or muon is not identified or reconstructed as a jet. In the right diagram, the tau is reconstructed as a jet.	676
5.5	Left: the m_T distribution in the $1L1\tau$ validation with all requirements except the m_T threshold, which is indicated by an arrow. Right: the am_{T2} distribution in the $1L1\tau$ validation region. Jet energy scale and resolution uncertainties are included in the error band. The last bin contains overflow.	678
5.6	The number of signal jets in the $1L1\tau$ VR. Jet energy scale and resolution uncertainties are included in the error band. The last bin contains overflow.	679
5.7	The leading (left) and sub-leading (right) non b -tagged jets in the $1L1\tau$ validation region. When there are more than two b -tagged jets in the event, the plotted jets are those after the leading two b -tagged jets, ordered by p_T . Jet energy scale and resolution uncertainties are included in the error band. The last bin contains overflow.	679
5.8	Left: the m_T distribution in the 2L validation region where the softer lepton is treated as measured but not reconstructed. Right: the am_{T2} distribution with the same lepton treatment as the left plot. Jet energy scale and resolution uncertainties are included in the error band. The last bin contains overflow.	681
5.9	The jet multiplicity in the 2L VR. Jet energy scale and resolution uncertainties are included in the error band. The last bin contains overflow.	682
5.10	The leading (left) and sub-leading (right) non b -tagged jets in the 2L validation region. When there are more than two b -tagged jets in the event, the plotted jets are those after the leading two b -tagged jets, ordered by p_T . Jet energy scale and resolution uncertainties are included in the error band. The last bin contains overflow.	682
5.11	The distribution of the number of b -jets with the inclusive preselection.	684

5.12	An illustration of the extrapolation from the W +jets control region to the signal region decomposed into two steps. The top left plot is the m_T distribution in a preselection with a b -jet veto. The bottom left plot shows the transfer factor (TF) that relates the yield in the region $30 \text{ GeV} < m_T < 90 \text{ GeV}$ region to various regions with m_T greater than the value on the x -axis. The right column of plots are analogous to the first column, but with a b -jet requirement instead of a veto. The middle column plots are ratios of the left and right columns. The error bars represent statistical uncertainty.	686
5.13	The m_T (top left), E_T^{miss} (top right), $a m_{T2}$ (bottom left), and $H_{T,\text{sig}}^{\text{miss}}$ (bottom right) distributions in the WVR-tail validation region. . . .	687
5.14	The leading (top left), subleading (top right), third leading (bottom left) and fourth leading (bottom right) p_T distributions in the WVR-tail validation region.	688
5.15	Feynman diagrams for the s -, t -, and Wt -channels on the left, middle, and right, respectively. Even though these are leading order in α_s , these are not necessarily the dominant diagrams because the b -quark PDF is highly suppressed due to the b -quark mass.	689
5.16	The distribution of the invariant mass of the W boson and b -quark from the non-top quark side. Diagrams with doubly resonant top quarks are explicitly removed. About 50% of the distribution has $m(Wb) > m_{\text{top}} \sim 175 \text{ GeV}$	690
5.17	The average number of b -jets, the probability for at least and for exactly one b -jet at particle-level from POWHEG-BOX+PYTHIA 6. .	691
5.18	Next-to-leading-order (up to $\alpha_s^4 \alpha_w^2$) Feynman diagrams from MG5_AMC that have the same outgoing particles as $t\bar{t}$ at tree level. The top left diagram overlaps with $t\bar{t}$ when the two intermediate top quarks go on-shell. The top right diagram has no top quarks at all and the bottom diagrams contain top quarks but do not interfere with top quark pair production even when the intermediate top quark(s) go on-shell.	692

5.19	The value of $\mathbf{am}_{T2}$ from a simple scalar parton-level simulation. In the correct assignment, the two $\mathbf{b}$ -quarks from the top decay are used to calculate $\mathbf{am}_{T2}$. When the charm quark from the W decay is used in place of one of the $\mathbf{b}$ -quarks, then the second $\mathbf{b}$ -quark used in the calculation can either be from the same top quark as the lepton or the opposite top quark. In all cases, the plotted value is the minimum over both pairings of the lepton and ‘ $\mathbf{b}$ -jets’.	693
5.20	The distribution of $\mathbf{am}_{T2}$ in the single top control region with all selections applied except the $\mathbf{am}_{T2}$ requirement. The left and right plots are identical except for the difference in scale.	695
5.21	The distribution of $\mathbf{n}_{\mathbf{b}\text{-jets}}$ and $\Delta\mathbf{R}(\mathbf{b}_1, \mathbf{b}_2)$ in the single top control region with all selections applied except the ones shown in the plots.	695
5.22	Feynman diagrams for stop pair production and decay (left) and top quark pair production in association with a Z boson that decays to neutrinos (right). The dashed circles show the dominant contribution to the E_T^{miss} . There are many other leading order Feynman diagrams for the $t\bar{t} + Z$ process, which are described in Sec. 5.4.2.	696
5.23	Distributions of (clockwise) E_T^{miss} , $\mathbf{m}_T$, $\mathbf{m}_{\text{jet}}$, and $\mathbf{am}_{T2}$. The stop model has $(\mathbf{m}_{\text{stop}}, \mathbf{m}_{\text{LSP}}) = (800, 0)$. Events are required to have four jets with $\mathbf{p}_T > 50, 50, 25, 25$ GeV, $E_T^{\text{miss}} > 300$ GeV, $\mathbf{m}_T > 120$ GeV and at least one $R = 1.2$ large radius re-clustered jet with $\mathbf{p}_T > 300$ GeV. The last bin contains overflow.	697
5.24	Leading order Feynman diagrams for $t\bar{t} + Z$ (left) and $t\bar{t} + \gamma$ (right) with the boson produced in the final state. In addition to these diagrams, there are two additional sets of FSR diagrams with gluons in the initial state: one with the same setup as above and one with a t -channel exchange of top quarks fusing into the boson (see Fig. 5.22).	698
5.25	Leading order Feynman diagrams for $t\bar{t} + Z$ (left) and $t\bar{t} + \gamma$ (right).	699

5.26	A lowest order (α_w^5) diagram that exists for $t\bar{t} + Z$ that has no $t\bar{t} + \gamma$ version (left) and the strong production ($\alpha_w\alpha_s^4$) analogue (right). The left diagram is suppressed with respect to the right one by $(\alpha_w/\alpha_s)^4 \sim 10^{-4}$	700
5.27	The distribution of the Z p_T given the particle level E_T^{miss} in $t\bar{t} + Z(\rightarrow \nu\bar{\nu})$ events. All neutrinos contribute to the particle level E_T^{miss}	701
5.28	The cross section ratio for $t\bar{t}+Z$ and $t\bar{t}+\gamma$ for various sub-processes described in the text. Unlike the simulation setups described in Sec. 2.3, the $t\bar{t} + Z$ and $t\bar{t} + \gamma$ use exactly the same setup: MG5_AMC 2.1.1 with PDF set CTEQ6L1. No extra partons are generated in the ME and photon radiation from top decay products is not included. The ISR only processes are generated with the syntax <code>generate p p > t t~ t~ a / t t~</code>	703
5.29	Given a $t\bar{t} + V$ ISR event, this is the fraction of events originating from a $u\bar{u}$ initial state. The dotted lines show the $p_T = 0$ fractions. .	704
5.30	The full cross-section ratio between $t\bar{t} + Z$ and $t\bar{t} + \gamma$ at 8 TeV (left) and 13 TeV (right).	704
5.31	Representative diagrams from MADGRAPH where the photon originates from one of the charged decay products of the top quark. . . .	705
5.32	The fraction of photons radiated from the charged top quark decay products (Top Decay) versus the fraction of photons directly from the hard-scatter (Not from Top Decay) as a function of the photon p_T . Events are generated with MG5_AMC at leading order with the full process <code>generate p p > t t~ > l+ vl b b~ l- vl~ a</code> . Photons are categorized as originating from the charged decay products of the top quark by parsing the ancestry stored in the LHE output of MG5_AMC. The band is the MC statistical uncertainty.	706
5.33	An example (abridged) event record from PYTHIA showering a $t\bar{t}$ event in which a high p_T photon is added as ISR (red).	707

5.34	The composition of the leading photon origin in POWHEG-BOX+PYTHIA 6+PHOTOS $t\bar{t}$ events with at least one photon at particle-level with $p_T > 80$ GeV. The ME- $t\bar{t}$ system category includes photons radiated from the decay products of the top quarks.	708
5.35	Example particle ancestries for the categories shown in Fig. 5.34. Photons are identified with their PDG ID (22) and then the MC event ancestry is parsed to find the origin of the photon. Each line contains four items: PDG ID, particle name, HepMC status code [616], and location in the event record. Not all numbers are used for the location in the event record, but higher numbers do correspond to later in the simulation.	709
5.36	The distribution of E_T^{miss} (left) and $\tilde{E}_T^{\text{miss}}$ (right) constructed from $\vec{p}_T^{\text{miss}}$ and $\vec{p}_T^\gamma$. Both distributions use data and simulation in the VR at $\sqrt{s} = 8$ TeV. Only statistical uncertainties are included in the error bars and bands.	716
5.37	The distribution of m_T^{miss} (left) and $\tilde{m}_T^{\text{miss}}$ (right) constructed from $\vec{p}_T^{\text{miss}}$ and $\vec{p}_T^\gamma$. Both distributions use data and simulation in the VR at $\sqrt{s} = 8$ TeV. Only statistical uncertainties are included in the error bars and bands.	716
5.38	The distribution of photon p_T and photon η in the $t\bar{t} + \gamma$ CR at $\sqrt{s} + 13$ TeV. Jet energy scale and resolution uncertainties in addition to statistical uncertainties are included in the error bars and bands. A normalization factor of 1.42 is applied. The last bin includes overflow.	717
5.39	The distribution of $\tilde{E}_T^{\text{miss}}$ and $\tilde{m}_T$ built from $\tilde{\vec{p}}_T^{\text{miss}} = \vec{p}_T^{\text{miss}} + \vec{p}_T^\gamma$ in the $t\bar{t} + \gamma$ CR at $\sqrt{s} + 13$ TeV. Uncertainties are the same as Fig. 5.38. A normalization factor of 1.42 is applied. The last bin includes overflow.	717
5.40	The m_T distribution of various diboson subprocesses with SHERPA 2.1.	718

5.41	The distribution of the leading jet p_T in the $t\bar{t}$ enriched region (left), the W +jets enriched region (middle), and the single top enriched region (right). See the text for the event selections. The first number in parenthesis after the process in the legend is the estimated event yield and the second number is the number of raw MC events used to make that prediction.	720
5.42	The distribution of $H_{T,\text{sig}}^{\text{miss}}$ in the $t\bar{t}$ enriched region (left), the W +jets enriched region (middle), and the single top enriched region (right). See the text for the event selections. The first number in parenthesis after the process in the legend is the estimated event yield and the second number is the number of raw MC events used to make that prediction.	720
5.43	The distribution of E_T^{miss} in the $t\bar{t}$ enriched region (left), the W +jets enriched region (middle), and the single top enriched region (right). See the text for the event selections. The first number in parenthesis after the process in the legend is the estimated event yield and the second number is the number of raw MC events used to make that prediction	721
5.44	The distribution of the leading large-radius ($R = 1.2$) jet mass in the $t\bar{t}$ enriched region (left), the W +jets enriched region (middle), and the single top enriched region (right). See the text for the event selections. The first number in parenthesis after the process in the legend is the estimated event yield and the second number is the number of raw MC events used to make that prediction	721
5.45	The dependence of the normalization factors on E_T^{miss} (left) and $H_{T,\text{sig}}^{\text{miss}}$ (right). The error band represents the statistical uncertainty derived from bootstrapping the data in the enriched regions and resolving Eq. 5.3.	722

5.46	The dependence of the normalization factors on the leading jet p_T (left) and the leading large radius ($R = 1.2$) jet mass (right). The error band represents the statistical uncertainty derived from bootstrapping the data in the enriched regions and resolving Eq. 5.3.	723
5.47	Statistical correlations between the various normalization factors determined from bootstrapping. The data yields are set to the prediction so that the scale factors are centered at $\mu = 1$. In the middle row, the $t\bar{t}$ contribution to the single top enriched region is set to zero and in the bottom row the single top contribution to the $t\bar{t}$ enriched region is set to zero.	724
6.1	The normalized significance as a function of the number of background events and the fractional systematic uncertainty for a one-bin fit with log-normal uncertainties and no profiling. The ‘significance’ is computed as $\Phi^{-1}(1 - CL_s)$ for Φ the cumulative distribution function of the normal distribution. The number of signal events is set to the number of background events (10 events) in the left (right) plot and each column is normalized to be one in the first row. The dashed line represents the data statistical uncertainty: $1/\sqrt{\text{Background Events}}$. .	726

6.2	An overview of the impact of the experimental uncertainties on the predicted yield in SR13. The Raw MC prediction does not include any data constraints from the control regions. The vertical dashed line indicates the nominal SM prediction after normalizing the simulation to the data in each of the four control regions described in Chapter 5. The black error bands show the size of the uncertainties before the control region constraint and the red bands show the error after using the control region method. Sections 6.1.1, 6.1.2, 6.1.3, and 6.1.4 describe how the various uncertainties are applied. To make them visible on the plot, the $\mathbf{b}$ -tagging and E_T^{miss} soft term uncertainties are multiplied by 10. The data and CR stats error bands represent the symmetric 68% inter-quantile range centered about the median determined from bootstrapping the data.	729
6.3	The signal jet multiplicity in $t\bar{t}$ events after the preselection before (left) and after (middle) normalizing the distributions to unity. The ratio of the JES $\pm 1\sigma$ to the nominal in the left plot is shown in the right plot.	732
6.4	The leading jet p_T in $t\bar{t}$ events after the preselection before (left) and after (middle) normalizing the distributions to unity. The ratio of the JES $\pm 1\sigma$ to the nominal in the left plot is shown in the right plot.	732
6.5	The large-radius ($R = 1.2$) jet mass in $t\bar{t}$ events after the preselection before (left) and after (middle) normalizing the distributions to unity. The ratio of the JES $\pm 1\sigma$ to the nominal in the left plot is shown in the right plot.	733
6.6	The size of the total JES uncertainty on each bin of the tNshape signal region before (left) and after (right) normalizing the total yields in each E_T^{miss} bin.	733
6.7	Left: The large-radius ($R = 1.2$) jet mass distribution in $t\bar{t}$ events after the preselection with the nominal JER and the JER inflated within its 1σ uncertainty. Right: the ratio of the two histograms in the left plot.	734

6.8	Left: The m_T distribution in $t\bar{t}$ events after the preselection with the nominal JER and the JER inflated within its 1σ uncertainty. Right: the ratio of the two histograms in the left plot.	735
6.9	The size of the total JER uncertainty on each bin of the tNshape signal region before (left) and after (right) normalizing the total yields in each E_T^{miss} bin.	735
6.10	The b -jet multiplicity in $t\bar{t}$ events after the preselection before (left) and after (middle) normalizing the distributions to unity. The ratio of the $\pm 1\sigma$ variations to the nominal in the left plot is shown in the right plot. The dashed lines in the right plot correspond to the JES uncertainty while the solid lines are the b -jet efficiency scale factor uncertainties.	737
6.11	The relative b -tagging efficiency scale factor uncertainty for various components of the uncertainty as a function of the b -tagged jet multiplicity. The ‘extrapolation from charm’ uncertainty is for τ jets that are b -tagged and the ‘extrapolation’ uncertainty is for extending the in-situ constraints to high p_T where the data statistics are limited. .	737
6.12	E_T^{miss} in $t\bar{t}$ events after the preselection before (left) and after (middle) normalizing the distributions to unity. The ratio of the $\pm 1\sigma$ variations to the nominal in the left plot is shown in the right plot. The dashed lines in the right plot correspond to the JES uncertainty while the solid lines are the E_T^{miss} soft term scale uncertainties.	738
6.13	The particle-level E_T^{miss} (left) and the number of particle-level jets (right) for POWHEG+PYTHIA 8 and MG5_aMC+HERWIG++. . . .	743
6.14	The inclusive $WWb\bar{b}$ process includes Feynman diagrams with doubly resonant (left), singly resonant (middle) and non-resonant top quark contributions. Feynman diagrams from MG5_aMC.	746

6.15	The particle-level E_T^{miss} for events passing a one lepton (a) and a two lepton (b) selection. Both selections require at least four jets with $p_T > 25$ GeV, of which two must have originated from b -quarks. All distributions are normalized to have the same integral in the above range. The gray band in the ratio is the statistical uncertainty from the simulation using the DR scheme and the uncertainty on the markers is from the simulation used in the numerator of the ratio. Most of these uncertainties are smaller than the markers.	747
6.16	Same as Fig. 6.15 but for the leading jet p_T instead of the E_T^{miss} . . .	747
6.17	Same as Fig. 6.15, but comparing the NLO calculation with interference removal with an inclusive $WWb\bar{b}$ sample generated with MG5_aMC.	748
6.18	Same as Fig. 6.17 but for the leading jet p_T instead of the E_T^{miss} . . .	748
6.19	The αm_{T2} (left) and E_T^{miss} (right) distributions when using one or two extra partons in the ME for $t\bar{t} + Z$ with MADGRAPH 5 + PYTHIA 6.	750
6.20	The slope parameter from a linear fit to the ratio of the variations listed on the horizontal axis. For ‘finite partons’, the comparison is between one and two extra partons in the ME; for ‘FSR’, PYTHIA 6 parameters are varied (see the text), for ‘Scales’, the μ_f and μ_r are simultaneously doubled and halved; for α_s , ISR parameters in MADGRAPH 5 and PYTHIA 6 are varied (see the text); for ‘Matching’, the MLM matching scale is varied. The error bars are due to MC statistical uncertainty. The units of the slope parameter are 1/GeV.	751
6.21	Variations of the leading order renormalization and factorization MadGraph scales coherently by factors of two and one-half. The default scale is the transverse mass (see Table 13 in Ref. [623]). The $t\bar{t} + \gamma$ sample has a ME cut of 100 GeV. The $\tilde{E}_T^{\text{miss}}$ is the regular E_T^{miss} in the $t\bar{t} + V$ case and the magnitude of the vector sum of the missing momentum and the photon momentum. The relative normalization of the $t\bar{t} + V$ and $t\bar{t} + \gamma$ samples is arbitrary. The upper ratio compares the scale up with the scale down for the $t\bar{t} + V$ and $t\bar{t} + \gamma$ separately. The lower panel is the ratio of the ratios: $r = \mu_{\text{up}}/\mu_{\text{down}}$	753

6.22	Left: The k -factor for $t\bar{t} + \gamma$ (left) and $t\bar{t} + Z$ (right) as a function of the boson p_T for SHERPA+OPENLOOPS and MG5_aMC. The PDF and scale choice are given in the legend.	754
6.23	The k -factor ratio between $t\bar{t} + \gamma$ and $t\bar{t} + Z$ as a function of the boson p_T for SHERPA+ OPENLOOPS and MG5_aMC. The PDF and scale choice are given in the legend.	754
6.24	Various comparisons of jet related quantities between SHERPA and POWHEG-BOX+PYTHIA 8 for the process $ll\nu\nu$. There is a $\sim 50\%$ difference between the predicted yields at preselection and an additional $\sim 20\%$ from extrapolating between the preselection and the SR. . . .	757
7.1	A comparison of the p -value under H_1 (CL_b) and under H_0 (CL_{s+b}) computed using numeric methods (10,000 toys) and the asymptotic formula given in Eq. 7.4. The vertical dashed line indicates the observed value of t and the dashed line is at the p -value corresponding to the observed t -value using the asymptotic formula.	764
7.2	Left: The power for three exclusion techniques (see the text for details) as a function of the number of signal events, for a fixed number of background events. Right: The minimum number of signal events that are needed to exclude a model given the number of estimated background events. A log-normal constraint is used to model the background uncertainty and 100,000 toys are used for each p -value calculation.	766
7.3	The μ values from the CR-only fit for all single bin signal regions. There is one normalization factor per E_T^{miss} bin for the tN1shape fit. The outer error bars indicate the total post-fit uncertainty while the inner error bars show the statistical uncertainty only (see the text for details).	768

7.4	The post-fit JES and JER nuisance parameters for the background-only all-bins fit for tNmed, tNhigh, and tN1shape. Red (blue) lines indicate the JER (JES) parameters. Single bin (shape fit) regions are on the left (right).	769
7.5	A comparison of the data and simulation in the validation and signal regions using the CR-only background fit. The error bar in the ratio panel shows the <i>pull</i> , defined as the difference between the data and the prediction, divided by the uncertainty. In this case, the uncertainty is the sum in quadrature of the data statistical uncertainty and the total background uncertainty.	770
7.6	A schematic diagram demonstrating how the exclusion limits are presented. See the text for details.	771
7.7	The exclusion contour of simplified stop models using the early $\sqrt{s} = 8$ TeV data. The observed limit is computed using the SR with the best expected sensitivity (lowest expected CL_s) for the given model as shown in Fig. 7.8. For comparison, the exclusion limits with the full $\sqrt{s} = 7$ TeV Run are overlaid with a gray dashed line [501]. The numerical values at each signal mass point show the smallest cross-section that would be excluded for a model with exactly the same acceptance. This is computed by scanning the signal cross-section and re-running the fit.	773
7.8	The signal region used for each mass point in the $(m_{\text{stop}}, m_{\text{LSP}})$ mass plane to compute the expected and observed limits in Fig. 7.7. . . .	773
7.9	The exclusion contour and SR mapping of simplified stop models using the full $\sqrt{s} = 8$ TeV dataset based on the lowest expected CL_s value. By construction, tNmed is the most sensitive for intermediate m_{stop} and large m_{LSP} and tNhigh is the most sensitive at m_{stop}	775

- 7.10 The observed limits using the best-expected map from Fig. 7.9 for two different branching ratio assumptions ($\mathcal{BR}(\tilde{t} \rightarrow t\tilde{\chi}_1^0) + \mathcal{BR}(\tilde{t} \rightarrow b\tilde{\chi}_1^\pm) = 1$). The green lines are computed by comparing the observed cross-section limits from Fig. 7.9 to 25%, 50%, or 75% of the predicted cross-section. To make a smooth contour, the significance is set to $2\sigma_{\text{excluded}}/\sigma_{\text{predicted}}$. The wavy line on the left-hand side is due to the lack of signal models above the limit. 776
- 7.11 Top: the acceptance for the tNmed (left) and tNhigh (right) event selections as a function of m_{stop} and m_{LSP} using a *particle-level* event selection analogous to the ones described in Sec. 4. The particle-level objects are similar to the ones used throughout Part II and are detailed in the appendix of Ref. [646]. Bottom: the ratio of the acceptance using detector-level objects to the acceptance using particle-level objects. This ratio is a correction for detector-effects in the event selection and is mostly uniform in the sensitive regions of parameter space. The acceptance in the upper plots is defined using particle-level objects in order to facilitate comparisons with other models for which a detector-simulation is not available. 777
- 7.12 Excluded regions in the $(m_{\text{stop}}, m_{\text{stop}} - m_{\text{top}} - m_{\text{LSP}})$ plane. The plots are identical aside from a vertical inversion. The region inside the bounded contour is excluded using the $\text{CL}_s < 0.05$ criteria (N.B. this is **not** the 95% confidence level). In the region below (above) the contour in the left (right) plot, the LSP is a Tachyon ($m < 0$) and in the region above (below) the contour, models are allowed by the data. 778
- 7.13 Left: the excluded region of the $(m_{\text{stop}}, m_{\text{stop}} - m_{\text{top}} - m_{\text{LSP}})$ mass plane (linear-scale version of Fig. 7.12). Right: a one-dimensional projection of the significance $\sigma = \Phi^{-1}(1 - \text{CL}_s)$ as a function of the mass gap $m_{\text{stop}} - m_{\text{top}} - m_{\text{LSP}}$ for $m_{\text{stop}} = 250$ GeV. Crosses indicate grid points and the solid/dashed lines interpolate between points. The horizontal dashed line is the exclusion threshold of $\Phi^{-1}(0.95)$ 778

7.14	The cross-section for various models highlighted in Table 7.1. The parameter k for the vector lepto-quarks corresponds to the κ model parameter in Ref. [482].	780
7.15	All possible mass hierarchies for $\tilde{g} \rightarrow t\bar{t}$, $\tilde{t} \rightarrow b\chi_1^\pm$. The arrows indicate which, if any, high energy particles are produced in the cascade decay. Decay (d) is phenomenologically the same as the compressed GMS and (g) is equivalent to the gluino mediated sbottom. The other possibilities are better covered by dedicated searches than direct stop/sbottom production. The gluino-mediated sbottom is similar, but with $\tilde{t} \leftrightarrow \tilde{b}$ and $t \leftrightarrow b$	781
7.16	Kinematic distributions for two distinct, but equivalent SUSY models. The model labeled <i>stop</i> is direct stop pair production with a 700 GeV stop mass and a massless neutralino. The model labeled <i>gluino</i> is a gluino mediated compressed stop model with a 1.1 TeV gluino, a 652 GeV stop and a 650 GeV neutralino. The generation is performed with MadGraph5_aMC@NLO version 5.2.1.1 [505] for the matrix element and Pythia 6.428 [96] for the parton shower and hadronization. The 2 GeV difference between the stop mass and neutralino mass in the gluino mediated stop model is due to a 1.5 GeV charm mass in Pythia (for $\tilde{t} \rightarrow c\tilde{\chi}^0$). A detector simulation is modeled with Delphes v3.1.2 [657]. Jets are clustered with the anti- k_t algorithm [117] with $R = 0.4$ using the fastjet program [203]. The top plot shows the distribution of the leading four jet p_T , the bottom left plot shows the magnitude of the missing transverse momentum and the bottom right plot shows the number of jets with $p_T > 25$ GeV.	785

7.17	The efficiency of a 315 GeV threshold on the magnitude of the missing transverse momentum (square markers) and the efficiency of a four jet requirement with transverse momentum thresholds (75, 65, 40, 25) GeV (circles). The markers indicate the efficiency for the compressed gluino model with $m_{\tilde{g}} = 1.1$ TeV and $m_{\tilde{t}} = 650$ GeV. The yellow band is the efficiency for the equivalent stop model with $m_{\tilde{t}} = 700$ GeV and a massless neutralino. The band and the error bars represent statistical uncertainties from finite simulated datasets. The generation is performed with MadGraph5_aMC@NLO version 5.2.1.1 [505] for the matrix element and Pythia 6.428 [96] for the parton shower and hadronization. A detector simulation is modeled with Delphes v3.1.2 [657]. Jets are clustered with the anti- k_t algorithm [117] with $R = 0.4$ using the FASTJET program [203].	786
7.18	A comparison of existing limits and the re-interpretation of the direct stop search limits at $\sqrt{s} = 8$ TeV (see Sec. 7.4.2). The expected limits (based on the CL_s procedure [642]) are used to control for statistical fluctuations in the observations. The blue shaded region is the re-interpretation of the direct stop search. Above the marked diagonal line, every gluino model has an equivalent stop model. Below this line, there is no equivalent stop model and the exclusion limits are estimated by extrapolating the signal region acceptance as described in the text. The hatched region is from the $\sqrt{s} = 8$ TeV ATLAS search for compressed direct stop production via an ISR monojet [513]. The red line is from the ATLAS inclusive one lepton search [650] (Fig. 18a) and the black line is from the ATLAS same-sign lepton search [648].	789
7.19	The exclusion limits using SR13 at $\sqrt{s} = 13$ TeV for direct stop pair production (left) and GMS with $m_{\tilde{t}} - m_{\tilde{\chi}^0} = 5$ GeV (right). The blue filled area in the right plot is from the early Run 2 ATLAS mono-jet search [659].	791

7.20	The distribution of p -values for excesses (left) and deficits (right) for both Gaussian and log-normal uncertainty distributions. For a continuous probability distribution, one expects the distribution of p -values to be uniform on $[0, 1]$	798
A.1	A schematic diagram illustrating the components of the digitizer model described in this section. While included in the model, thermal diffusion is a generally small effect and is not discussed in the following section.	803
A.2	The average (over x and y) electric field as a function of the depth (z) inside an unirradiated planar $200\text{ }\mu\text{m}$ deep planar sensor with a bias voltage of 80 V (1000 V) for an unirradiated ($5 \times 10^{15}\text{ n}_{\text{eq}}/\text{cm}^2$) sensor on the left (right).	805
A.3	The time for an electron or hole to drift to the top (collecting electrode) or bottom of the sensor as a function of the depth (z) using the averaged E fields shown in Fig. A.2 for a planar sensor with a bias voltage of 80 V (1000 V) for an unirradiated ($5 \times 10^{15}\text{ n}_{\text{eq}}/\text{cm}^2$) sensor on the left (right).	806
A.4	The Ramo potential at $x = y = 0$ (centered on the collecting electrode ϵ) as a function of the distance z away from ϵ for three sizes of ϵ : $W = L/10$, $W = L/2$, and $W = \infty$, where L is the sensor thickness.	807
A.5	The Ramo potential at $y = 0$ (centered on the collecting electrode ϵ) as a function of the distance z away from ϵ for three sizes of ϵ : $W = L/10$, $W = L/2$, and $W = \infty$, where L is the sensor thickness.	807

A.6	A cross-section of the surface charge density from a point charge a distance z away from a series of grounded strips. The ‘primary electrode’ is the strip centered at 0 (extending to $\pm\infty$ in the direction into and out of the page) and the ‘neighboring electrode’ is the one shown between 2 and 3. The middle plot shows the total charge on the neighboring electrode as a function of the distance the charge is from the origin and the plot on the right shows the fraction of the charge that is induced on the neighbor using the Ramo potential. . .	808
A.7	The final location (based only on drift) for electrons starting at a depth z and traveling a time given by the vertical axis. The coordinate z is measured with respect to the back-side (away from the collecting electrode) of the sensor using the averaged E fields shown in Fig. A.2 for a planar sensor with a bias voltage of 80 V (1000 V) for an unirradiated ($5 \times 10^{15} \text{ n}_{\text{eq}}/\text{cm}^2$) sensor on the left (right). . .	809
A.8	The final location (based only on drift) for holes starting at a depth z and traveling a time given by the vertical axis. The coordinate z is measured with respect to the back-side (away from the collecting electrode) of the sensor using the averaged E fields shown in Fig. A.2 for a planar sensor with a bias voltage of 80 V (1000 V) for an unirradiated ($5 \times 10^{15} \text{ n}_{\text{eq}}/\text{cm}^2$) sensor on the left (right).	810
A.9	The charge induced on the primary (left) and neighboring electrodes (short direction in the middle and long direction on the right) computed using the Ramo potential.	810

A.10	An illustration of the impact of representing multiple fundamental charges by a single multi-charged chunk. Displayed is the fraction of collected charge as a function of the exponential trapping time constant. N chunks are propagated to the electrode and some fraction f of them reach the electrode before a random exponential time. If a chunk reaches the electrode, a charge of Q is recorded. The points show $\langle fQN \rangle$ and the error bars are the standard deviation. Red corresponds to $Q = 1$ and blue corresponds to $Q = 100$ (left), $Q = 10$ (middle), and $Q = 1$ (right).	811
A.11	Pictures of the testbeam setup. The left picture shows the six telescope planes, three on each side of the DUT. The unirradiated module does not need to be cooled and sits on a stage that can be automatically raised, lowered, and tilted. The beam goes from left to right, piercing the center of the black squares shown in the middle of each telescope plane. The blue, red, and black cables from the DUT are for the data, low voltage, and high voltage, respectively. The right picture shows a close-up of the irradiated module inside the box used to keep it cool (with dry ice). The orange cables are temperature sensors.	812
A.12	A illustration of the tilted sensor configuration. A charged particle travels from the bottom left to the top right and traverses about $50\text{ }\mu\text{m}$ of silicon in each sensor. The tilt angle is chosen so that the particle will traverse about 15 pixels, probing different depths along its path.	813

A.13	Representative event displays from the unirradiated (left) and the irradiated (right) sensors. In both cases, there are three clusters. The sensors are tilted (Fig. A.12) so that one expects about 15 pixel clusters per particle. The unirradiated sensor uses a 2000 electron threshold and 8 TOT is tuned to 11000 collected electrons, while the unirradiated sensor uses a 1500 electron threshold with 8000 electrons corresponding to 5 TOT. The unirradiated sensor is biased with 60 V while the irradiated sensor is biased by 1.1 kV and cooled to about -35°C	814
A.14	The TOT distribution as a function of depth inside clusters of length exactly 15 (left) or 10 (right) for the unirradiated and irradiated sensors from Fig. A.13. The tilt angle is the same for both modules, but the pixel size distribution was peaked at lower values (10 versus 15) for the irradiated sensor, indicating that it is not fully depleted (even at 1.1 kV).	814
B.1	A diagram illustrating the setup for the calculation described in the text. The boson spin is along the z-axis.	817
C.1	A schematic diagram of an uncertainty ellipse for a bivariate normal distribution. See the text for details.	823
C.2	A three-particle event display illustrating the results of fuzzy jet clustering using a Gaussian density for ϕ (left) and a wrapped Gaussian density approximation for ϕ (right). Figure from C. Stansbury. . . .	824
C.3	Changing the choice of the p_{T} cut used to select seeds can make a vast difference in the values of the constructed variables, like σ . In this event, clustered on the left with a cut of 5 GeV resulting in five jets, and on the right with a cut of 50 GeV resulting four jets. Fewer degrees of freedom in the four jet case means a much larger learned value for the σ variable. Figures from C. Stansbury.	830

C.4	The calculated non-closure for numerical inversion assuming $Y X = x \sim \mathcal{N}(f(x), \sigma(x))$, where $\sigma(x)/x = 1/\sqrt{x}$ estimated from Ref. [289] and $f(x)$ has several possibilities. The blue line is $f(x) = ax+b$ (which closes identically), red is $f(x) = a + b \log(x)$ (derived by estimating the curve from the lowest $ \eta $ bin in Ref. [289] - $a = 0.5$ and $b = 0.09$), purple is $f(x) = \sqrt{x}$, and green is $f(x) = ax^2 + bx + c$ with $b = .73, b = 0.002, c = 0.0$	834
-----	--	-----

0.1 Preface

The following sections summarize useful nomenclature and background information.

0.1.1 Units

All physical results can be presented in any unit system, but not all systems are equivalently useful. The familiar meters-kilograms-seconds (SI) unit system will be used to express dimensions of the detector. Most other discussions will use *natural units* in which the rulers are not distance, mass, and time, but instead speed, angular momentum, and energy. The rulers of time have length c , the rulers of angular momentum have length $\hbar$ and the rulers of energy have length giga-electron-volts (GeV). In SI, an object has ‘length 1’ if it is one meter long. Equivalently, in natural units an object has ‘speed 1’ if it is going at the speed of light. This nomenclature is used throughout - the c and $\hbar$ will be henceforth implied and not stated explicitly for all dimensionful quantities. For example, masses, momenta, and energies are all given in units of GeV (the $1/c^2$ and $1/c$ for mass and momentum are implied) and lengths and time are given in units of $1/\text{GeV}$ (the $\hbar c$ and $\hbar$ are implied). Table 1 gives a representative set of useful units for high energy physics and their abbreviations.

Quantity	Abbreviated Units	Full Units	SI (approximate)	Comment
Speed	1	c	$3 \times 10^8 \text{ m/s}$	
Angular Momentum	1	$\hbar$	$10^{-34} \text{ m}^2 \text{ kg/s}$	
Energy	GeV	GeV	$1.6 \times 10^{-10} \text{ J}$	
Momentum	GeV	GeV/c	$10^{-19} \text{ kg} \cdot \text{m/s}$	
Mass	GeV	GeV/c^2	$1.8 \times 10^{-27} \text{ kg}$	
Time	$1/\text{GeV}$	$\hbar/\text{GeV}$	$6.6 \times 10^{-25} \text{ s}$	
Length	$1/\text{GeV}$	$\hbar c/\text{GeV}$	$2 \times 10^{-16} \text{ m}$	
Charge	1	$e/\sqrt{4\pi\alpha}$	$5.3 \times 10^{-19} \text{ C}$	$e = 1.6 \times 10^{-19} \text{ C}$
Magnetic Field	$(\text{GeV})^2$	$\text{GeV}^2/(\hbar c^2)$	$5 \times 10^{16} \text{ T}$	$T = (\text{kg})/(\text{C} \cdot \text{s})$

Table 1: Natural units. There are multiple ways to define the electric charge. In these natural units, $e = \sqrt{4\pi\alpha}$ so that one unit represents ~ 0.3 of an elementary charge.

0.1.2 Coordinates

Two sets of coordinates will be used interchangeably: $(\mathbf{p}_x, \mathbf{p}_y, \mathbf{p}_z)$ and $(\eta, \phi, \mathbf{p}_T)$, where the z -axis is along the beam (longitudinal) direction, ϕ is the azimuthal angle, $\mathbf{p}_T^2 = \mathbf{p}_x^2 + \mathbf{p}_y^2$ is the *transverse momentum*, and η is the *pseudo-rapidity*:

$$\eta = -\ln \left(\tan \left(\frac{\theta}{2} \right) \right) = \frac{1}{2} \ln \left(\frac{|\vec{p}| + p_z}{|\vec{p}| - p_z} \right) = \tanh^{-1} \left(\frac{p_z}{|\vec{p}|} \right), \quad (1)$$

where θ is the angle between the z -axis and the transverse plane. Particles with $\eta = 0$ point in the transverse plane and $\eta = \pm\infty$ are moving along the z -axis. These coordinates are particularly useful because after a Lorentz boost along the z -axis with magnitude β , a massless particles with $(\eta, \phi, \mathbf{p}_T)$ is described by $(\eta + \tanh^{-1}(\beta), \phi, \mathbf{p}_T)$. In particular, the difference $\Delta\eta$ between two massless particles is invariant under a boost along z . This motivates the distance metric $\Delta R^2 = \Delta\eta^2 + \Delta\phi^2$, which is invariant under longitudinal boosts for massless particles. At a hadron collider, the partonic $\mathbf{p}_z$ of a collision is in general not known, so the invariance of η is crucial. For massive particles, the generalization of η is the *rapidity* ($\mathbf{y}$), which is defined using the second or third equality Eq. 1, but replacing $|\vec{p}|$ with E . Rapidity is not determined solely by geometry (no equivalent to the first equality in Eq. 1), but does transform additively under a boost along z . Figure 1 compares η and $\mathbf{y}$. For a particle with $(\mathbf{p}_T, m) \approx (200, 100)$ GeV at $\eta \sim 1$, the difference is about 10%.

0.1.3 Statistics

The distribution $\mathcal{D}(\theta)$ of a random variable X will be denoted $X \sim \mathcal{D}(\theta)$ for some parameters θ . For example, $X \sim \mathcal{N}(\mu, \sigma^2)$ means that X follows a normal distribution with mean μ and variance σ^2 . The probability distribution $f_X(\mathbf{x})$ is related to the cumulative distribution $F_X(\mathbf{x}) = \Pr(X \leq \mathbf{x})$ by $f_X(\mathbf{x}) = \partial_{\mathbf{x}} F_X(\mathbf{x})$. If $\mathbf{y} = \mathbf{g}(\mathbf{x})$, then $f_Y(\mathbf{y}) = f_X(\mathbf{g}^{-1}(\mathbf{y})) |\partial_{\mathbf{y}} \mathbf{g}^{-1}(\mathbf{y})|$. The quantity $\int_{-\infty}^{\infty} d\mathbf{x} \mathbf{x} f_X(\mathbf{x})$ will be interchangeably called the expected value, mean, or average and is denoted $E[X]$ or $\langle X \rangle$. The square root of the variance $\langle X^2 \rangle - \langle X \rangle^2$ uses the symbol $\sigma(X)$ and is referred to as the




Figure 1: A comparison of η and y for a massive particle with $\phi = 0$.


standard deviation. A distribution's mean and standard deviation are sensitive to outliers and in general do not carry any probability content. Therefore, frequently used alternatives are the median $\mathfrak{m}$, defined by $\int_{-\infty}^{\mathfrak{m}} dx f_X(x) = \int_{\mathfrak{m}}^{\infty} dx f_X(x)$, and the inter-quantile range, which is a symmetric interval around $\mathfrak{m}$ that contains a specified fraction of the distribution $f_X(x)$.

A technique that is used extensively to numerically estimate the uncertainty in the measured statistics of X is the *bootstrap* [1]. Let $x_1, \dots, x_n$ be independent and identically distributed measurements from a random variable X . A bootstrap dataset $x'_1, \dots, x'_n$ is generated by picking $j_1, \dots, j_n$ with $j_i \sim \text{Uniform}(1, \dots, n)$ and settings $x'_i = x_{j_i}$. Note that the same measurement x_i may appear multiple times in the bootstrap dataset. Many such datasets are generated and then the uncertainty on a statistic is estimated by computing moments or percentiles of the distribution of the statistic over the ensemble of bootstrap datasets. For proofs about the bootstrap, see Ref. [2] and references therein.

Additional statistical tools and techniques are introduced when needed in later sections and in Appendix C.

Part I

The Theory of Experimental Particle Physics



A schematic diagram illustrating the production and measurement of stop quark pair production, from 10^{-19} m up to 10^1 m. Modeled after Fig. 1 in Ref. [3].

The ultimate goal of particle physics is to uncover the equations of motion of the elementary degrees of freedom: the *fundamental laws of the universe*. Due to quantum mechanics, there is inherent randomness in these fundamental laws. Therefore, the elementary degrees of freedom are not particles in the classical sense. Instead, they are described by *quantum fields* which are operator functions on spacetime from which one can compute probability distributions. Fields, and not individual particle wave functions, are elementary in order to rectify quantum mechanics and special relativity: fields preserve causality and allow for particle creation and destruction. *Quantum field theory* (QFT) is the language of modern particle physics.

The definition of *elementary* has evolved over time. Perhaps the first ‘particle physicists’ were the early chemists at the turn of the 19th century who discovered various distinct elements that were postulated to be built from atoms (literally Greek for ‘uncuttable’). Subatomic particle physics began with the discovery of the electron by J. J. Thomson in 1897. Arguably the beginning of modern particle physics was the first *fixed target* experiment by Geiger, Marsden, and Rutherford using gold foil in ~ 1910 to show that the positive charge inside atoms is concentrated in a point-like center called the nucleus. The same idea was used to discover the internal structure of the nucleus 50 years later at SLAC with a much more sophisticated apparatus and significantly higher energy probe particles. The last 100 years of nuclear and elementary particle physics have been filled with many stories of discovery, confusion, prediction, and success. Some of these will be introduced throughout Part I; for a more detailed account, there are many excellent references¹.

Part I serves as a brief introduction to experimental particle physics, with the goal of explaining all aspects of the schematic diagram from p.4. Chapter 1 introduces the Standard Model of particle physics, which describes all known elementary particles and interactions. The detection of particles, specifically with the ATLAS detector, is described in Chapter 2. Simulation and modeling of particle production and detection is documented in Chapter 3 and Chapter 4 describes how final states are reconstructed with the ATLAS detector.

¹A thorough account can be found in Ref. [4] up to the mid 1980s. Reference [5] is a detailed account of QED. Various dedicated articles on the LEP program at CERN, the discovery of the top quark, and the discovery of the Higgs boson can be found in e.g. [6–8].

Chapter 1

The Standard Model of Particle Physics

Due to technical advances in the 1950s, there was an explosion of new unstable subatomic particles discovered with lifetimes ranging from 10^{-23} to 10^{-6} seconds. Just as the periodic table of elements reduced the number of degrees of freedom from patterns in atomic spectra, the quark model based on $\text{SU}(3)$ (gauge) symmetry was postulated to explain the structure of the newly discovered *hadrons* [9–11]. The $\text{SU}(3)$ gauge theory built on the success of the slightly earlier foundational work on the quantum theory of electromagnetism (QED) based on a $\text{U}(1)$ gauge group¹. Around the same time as the strong force, a coherent quantum theory of the weak force and electromagnetic forces was assembled involving a $\text{U}(1) \times \text{SU}(2)$ gauge theory [13–15] that contained a mechanism for generating masses for SM particles [16–19].

The Standard Model (SM)² is a composite theory combining the strong and electroweak forces. This chapter summarizes the particle content and interactions of the SM (Sec. 1.1). Section 1.2 provides an overview of the success of the SM as well as its limitations, which have lead to an extensive literature on extensions of the SM. One important class of models is Supersymmetry (SUSY), which is the focus of Part III.

¹See Ref. [12] and references therein by Tomonaga, Schwinger, and Feynman.

²There are many excellent books on QFT that describe the technicalities of the SM. See for example, Ref. [20] and Ref. [21] (based on S. Coleman’s lectures). There are also many dedicated books on the SM or various parts of the SM such as Ref. [22] (QCD).

1.1 Particles and Forces

The SM describes three elementary forces: electromagnetic, weak, and strong. The electromagnetic interaction at the subatomic level is the same long range force that is familiar at everyday distance scales; it is responsible for most aspects of daily life (protein structure, batteries, friction, etc.). Note that the SM does not describe the other force relevant to daily life: gravity. This is irrelevant for terrestrial particle physics probed thus far because the strength of gravity is incredibly weak (but is revisited in Sec. 1.2). The weak and strong forces are mostly unfamiliar to daily life because they are not long range forces. At very short distance scales, the weak force is actually *stronger* than the electromagnetic force. However, the weak force analogue to electromagnetism's photon is about 100 times as massive as the proton (1 GeV). For distances comparable to or larger than this mass $\hbar c/(100 \text{ GeV}) \sim 10^{-18} \text{ m}$, the weak force is highly suppressed. Despite this small distance scale, the weak force is responsible for radioactive decay. The large mass of the mediator allows many unstable particles to survive for macroscopic times before decaying. In contrast, the strong force is mediated by a massless particle called the gluon. The reason the strong force is not long range is because it is *too strong*. Unlike electromagnetism, the strong force strengthens with distance; this means that the energy stored in the field of two objects participating in the strong force will be sufficient by $E = mc^2$ to make more particles. These additional particles *screen* the original force. Beyond about 10^{-15} m , the primary strong force is highly suppressed, but there is a residual *strong nuclear force* that is responsible for binding protons and neutrons inside the nucleus. After a few femtometers, the strong force is negligible compared with electromagnetism.

All matter particles participate in the weak interaction³, while only *quarks* feel the strong force, and both quarks and *charged leptons* interact via the electromagnetic force. Neutrinos are matter particles that only interact via the weak force. Each force is mediated by the exchange of force-carrying particles. Both the electromagnetic and strong forces are mediated by massless spin 1 bosons (photon and gluon) while the weak force is carried by three massive particles called the $W^\pm$ and Z bosons. The spin

³Right-handed neutrinos are ignored. If they exist, they do not interact within the SM.

and mass of a particle are associated with their representation of the Poincaré group, which is the symmetry group of spacetime (rotations, translations, and boosts). In the context of QFT, the electromagnetic, weak, and strong forces are described by *internal* (in contrast to spacetime) symmetries of the SM Lagrangian. Electromagnetism is the result of a dimension one unitary group $U(1)$ symmetry while the weak force and the strong force are described by the special unitary groups $SU(2)$ and $SU(3)$, respectively. The SM Lagrangian is given by:

$$\mathcal{L} = -\frac{1}{4}(F_{\mu\nu}^a)^2 + \bar{\psi}(i\mathcal{D})\psi + y_{ij}\bar{\psi}_i\psi_j\phi + |D_\mu\phi|^2 + \mu^2\phi^\dagger\phi - \lambda(\phi^\dagger\phi)^2, \quad (1.1)$$

where $F_{\mu\nu}^a = \partial_\mu A_\nu^a - \partial_\nu A_\mu^a + gf^{abc}A_\mu^b A_\nu^c$ and $D_\mu = \partial_\mu - igA_\mu^a t_r^a$ for fermion fields ψ , scalar field ϕ , gauge fields A , representation matrices t_r^a , and structure constants f^{abc} . Using the Feynman slash notation, $\mathcal{D} = \gamma^\mu D_\mu$, where γ^μ form a matrix representation of the Clifford algebra. The *Yukawa* couplings y_{ij} and the parameters λ and μ describe the interactions of the Higgs boson with itself and with the fermions. Equation 1.1 has an implicit sum over gauge groups, fermion types, group indices, and implicit Hermitian conjugates. Expanding the Lagrangian in Eq. 1.1 would fill multiple pages; however, there are already some interesting observations one can make: (1) $f^{abc} \neq 0$ for non-Abelian groups and therefore the bosons of the weak and strong forces *interact with themselves*, i.e. participate in the force they mediate. In contrast, the structure constant for electromagnetism is zero so photons do not directly interact with other photons via electromagnetism⁴. (2) There are no explicit mass terms in Eq. 1.1. The fields participating in the weak force are *Weyl* fermions with definite (left-handed) chirality. Therefore, mass terms which link left- and right-handed fermions such as $m\bar{\psi}_L\psi_R + \text{h.c.}$ are not allowed by symmetry. This is a serious problem because fermions certainly have mass. The Lagrangian above is written *before electroweak symmetry breaking*, which is the process by which fermions and the bosons acquire

⁴This is only true at tree-level in perturbation theory; there are virtual corrections that lead to non-trivial photon-photon interactions. Due to the smallness of the electromagnetic coupling, these interactions are highly suppressed, which is why at everyday energy scales this is negligible.

a mass⁵ (described below). Note that there is no problem writing down a mass term for the *Higgs field* ϕ since $\phi^\dagger\phi$ is invariant under an $\text{SU}(2)$ rotation. (3) The parameters y_{ij} are dimensionless. Integrating the Lagrange density from Eq. 1.1 must be dimensionless so $[y] + [\phi] + 2[\psi] - 4 = 0$, where $[*]$ is the mass dimension of $*$ and $0 = [\int d^4x \mathcal{L}] = [\mathcal{L}] - 4$. The mass term for a scalar is $m^2\phi^2$ and for Dirac fermion is $m\bar{\psi}\psi$ so $2[\phi] - 2 = 0$ and $2[\phi] - 3 = 0$. Solving these three equations gives $[y] = 0$. Naively, one might expect these Yukawa couplings to be $\mathcal{O}(1)$; it is therefore a surprise that they span 6 orders of magnitude.

Table 1.1 lists all of the fields and their representations as part of Eq. 1.1 (prior to electroweak symmetry breaking). There are three families each of right-handed up-type and down-type quarks. The left-handed up- and down-type quarks are grouped into an $\text{SU}(2)$ doublet. Similarly, there are three generations of right-handed charged leptons and three generations of left-handed charged leptons grouped with neutrinos as $\text{SU}(2)$ doublets. One additional $\text{SU}(2)$ doublet ϕ is occupied by two complex scalar fields. This field will play a critical role in electroweak symmetry breaking. The other SM fields are the gauge boson for $\text{U}(1)$ called the B , the $W^i, i = 1, 2, 3$ bosons for $\text{SU}(2)$, and the gluons for $\text{SU}(3)$. There are eight gluon fields, one for each generator of the $\text{SU}(3)$ Lie algebra (the Gell-Mann matrices). Likewise, there are three quarks for each entry in Table 1.1, one for each dimension of the fundamental representation of $\text{SU}(3)$. To build an analogy to QED where particles have electrical charge, the three possible $\text{SU}(3)$ options for each quark are called *color charge* and labeled red, green, and blue. These have nothing to do with actual color, but are useful because like visible light, a triple of quarks covering all three colors (red, green, blue) acts as if it were colorless. The eight gluons can be considered as having one color charge and one anti-color charge (the ‘anti-’ refers to the electric charge, since gluons are electrically neutral). Color charge in $\text{SU}(3)$ is discussed in more detail in Part II.

The idea of electroweak symmetry breaking is that the potential for the Higgs field $V(\phi) = -\mu^2\phi^\dagger\phi + \lambda(\phi^\dagger\phi)^2$ can have a classical nonzero minimum if $\mu^2, \lambda > 0$.

⁵In fact, the $\text{U}(1)$ in Eq. 1.1 is not exactly electromagnetism. After electroweak symmetry breaking, the B field combines with the neutral W field to form *both* the photon and Z boson.

Field	Content	Spin	U(1)	SU(2)	SU(3)	Comment
Q_i	$(u_L \ d_L)$	$\frac{1}{2}$	$\frac{1}{6}$	2	3	3 generations
$u_{R,i}$	u_R	$\frac{1}{2}$	$\frac{2}{3}$	1	$\bar{\mathbf{3}}$	3 generations
$d_{R,i}$	d_R	$\frac{1}{2}$	$-\frac{1}{3}$	1	$\bar{\mathbf{3}}$	3 generations
L_i	$(e_L \ \nu_L)$	$\frac{1}{2}$	$\frac{1}{2}$	2	1	3 generations
$e_{R,i}$	e_R	$\frac{1}{2}$	-1	1	1	3 generations
ϕ	$(\phi^+ \ \phi^0)$	0	$\frac{1}{2}$	2	1	
B	B	1	0	1	1	
W	$(W_1 \ W_2 \ W_3)$	1	0	3	1	
g	g	1	0	1	8	

Table 1.1: The particle content of the SM prior to electroweak symmetry breaking. The values under U(1) are the Abelian charge (the actual representation is one-dimensional) whereas the entries under SU(2) and SU(3) are the representation of the field in the first column. For example a bold eight denotes the octet (adjoint) representation of SU(3).

The field ϕ has four real degrees of freedom, but one can write the minimum as $\frac{1}{\sqrt{2}}(0 \ \nu)$, with $\nu = \sqrt{\frac{\mu^2}{\lambda}}$ and then all other points in the minimum are related to this one by SU(2) transformations. One can re-write the field ϕ as a fluctuation about this minimum: $\phi = \frac{1}{\sqrt{2}}(0 \ \nu + h(x))$, where h is a real-valued scalar field⁶. With this formulation of ϕ , expanding Eq. 1.1 gives rise to terms of the form $\mathcal{L}_{\psi\bar{\psi}\nu} = -(y\nu/\sqrt{2})\psi\bar{\psi}$, resulting in masses for the fermions $m = y\nu/\sqrt{2}$, and

$$\mathcal{L}_{\text{gauge-boson-}\nu} = \frac{1}{2} \frac{\nu^2}{4} \left[g_2^2 (W_\mu^1)^2 + g_2^2 (W_\mu^2)^2 + \left(-g_2 W_\mu^3 + \frac{1}{2} g_1 B_\mu \right)^2 \right], \quad (1.2)$$

which are the mass terms for the electroweak Gauge bosons. There are three vector boson mass eigenstates from Eq. 1.2: $W_\mu^\pm = \frac{1}{\sqrt{2}}(W_\mu^1 \pm iW_\mu^2)$ with mass $m_W = g\nu/2$, $Z_\mu = \cos(\theta_W)W_\mu^3 - \sin(\theta_W)B_\mu$ with mass $m_Z = m_W/\cos(\theta_W)$ and the massless photon field $A_\mu = \sin(\theta_W)W_\mu^3 + \cos(\theta_W)B_\mu$. The *weak mixing angle* is $\theta_W =$

⁶This choice is the *unitary gauge*.

$\cos^{-1}(g_2/\sqrt{g_1^2 + g_2^2})$. Three of the four real degrees of freedom from the field ϕ have been absorbed by the massive gauge bosons, adding a longitudinal polarization and allowing them to be massive. The fourth real degree of freedom is $h(x)$, known as the Higgs boson. This field couples to all massive particles and was the last particle of the SM to be discovered [23, 24]. Table 1.2 summarizes the mass eigenstates of the SM after electroweak symmetry breaking. There are six quark fields and six lepton fields, organized into three families of increasing mass. The Yukawa couplings are $y_f = \sqrt{s}m_f/v$ where m_f is the fermion mass and v is the Higgs field vacuum expectation value ~ 250 GeV. In total, there are 19 free parameters of the SM⁷, including nine Yukawa couplings (fermion masses), one Higgs mass parameter, three gauge couplings, one Higgs self coupling, and one 3×3 matrix V (the Cabibbo-Kobayashi-Maskawa (CKM) matrix [25, 26]) to describe transitions between quark types from weak decay. This last item is the result of defining the quark fields in Table 1.2 as the mass eigenstates: this induces off-diagonal components in the electroweak basis. The matrix V has four independent real number degrees of freedom. Except for the Higgs mass parameter, all other SM parameters are dimensionless. Some of the numerical values are given in Table 1.2 and Figure 1.1 graphically compares all of the dimensionless values. The Yukawa couplings (and therefore the SM masses) span six orders of magnitude. In contrast, the range for the gauge couplings is less than one order of magnitude. The CKM matrix is nearly diagonal.

Another, ‘practical’ way of visualizing the strength of the three forces is through the decay times of various particles. Figure 1.2 shows the lifetimes and masses of various elementary and composite particles. The decay rate Γ scales with g^2 , for gauge coupling g . For charged current weak decays, there is also a factor of $|V_{ij}|^2$ (CKM matrix), which is near unity when the transition is near the diagonal of the CKM matrix. The decays mediated by the strong force are the fastest, with typical lifetimes $\sim 1/(1 \text{ GeV}) \sim 10^{-24}$ s. Admixture of electromagnetic decays and phase space factors can increase these lifetimes. In contrast, the electromagnetic decays are

⁷This depends on how one counts. For example, the number of families could be a free parameter. The electric charge could be viewed as a free parameter, but it is basically fixed by the coupling structure of the SM (including the anomaly cancellation - see Sec. 1.4). There are also several terms which are allowed by symmetry but are so close to zero that they are neglected (see Sec. 1.2).

Field	Q	SU(3)	Yukawa Coupling	Other Couplings
u, c, t	$\frac{2}{3}$	3	$10^{-5}, 7 \times 10^{-3}, 1$	—
d, s, b	$-\frac{1}{3}$	3	$3 \times 10^{-5}, 5 \times 10^{-4}, 0.03$	—
e, μ, τ	-1	1	$3 \times 10^{-6}, 6 \times 10^{-4}, 0.01$	—
ν_e, ν_μ, ν_τ	0	1	—	—
h	0	1	—	$\mu = 90 \text{ GeV}, \lambda = 0.1$
γ	0	1	—	$\alpha = 1/127$
Z	0	1	—	$\sin(\theta_W) = 0.5$
$W^\pm$	± 1	1	—	V
g	0	8	—	$\alpha_s = 0.1$

Table 1.2: The particle content of the SM after electroweak symmetry breaking. All couplings are given to one significant figure at the scale m_Z . Parameter values are from Ref. [27].




Figure 1.1: A comparison of all dimensionless SM parameters. The nine Yukawa couplings are on the left in red and are organized by generation. The three gauge couplings g_i are $g_1 = e/\cos(\theta_W)$, $g_2 = e/\sin(\theta_W)$, $g_3 = \sqrt{4\pi\alpha_s}$, for the weak mixing angle θ_W , e is the electron charge $e = \sqrt{4\pi\alpha_{EM}}$, and $\alpha_s \sim 0.1$ is the strong coupling constant. The magnitude of all elements of the CKM matrix are in blue on the right. Parameter values are from Ref. [27].

much slower, with $\alpha_{\text{EM}} < \alpha_s$ reducing Γ . Slower still are the weak decays, which can persist for macroscopic timescales. Even though $g_2 > g_1$ (from Fig. 1.1), the weak decays are highly suppressed because the matrix element squared scales as g^4/m_W^4 (often called the Fermi constant, G_F up to an $\mathcal{O}(1)$ constant) when $m \ll m_W$. The lifetime of the heavier b - and c -mesons (B and D) as well as the τ are less suppressed than for the lighter mesons and baryons. In particular, the mass splitting between the proton and the neutron is so small ($\mathcal{O}(0.001)$ GeV) that a free neutron survives for about 15 minutes on average. However, Fig. 1.2 also reflects the fact that g_2 is not inherently small; the top quark decays via the weak interaction and has a lifetime comparable to the strong force resonances. This is because $m_{\text{top}} > m_W$ so there is little phase space suppression.




Figure 1.2: An illustration of the strength of the three forces through the lifetimes of various elementary and composite particles. The color of the labels is determined by the dominant decay mode. Values of the mass and lifetimes are from Ref. [27]. The lifetime is defined as the inverse of the decay width (for the strong force resonances, this is what is measured).

Aside from the top quark, the other five quarks are not present in Fig. 1.2. This is because the strong force exhibits *asymptotic freedom* [28, 29] at high energy and *confinement* at low energy. The effective coupling for the three forces are energy-scale dependent, as governed by the Callan-Symanzik equations [30, 31]:

$$\mu \frac{\partial g}{\partial \mu} = \beta(g), \quad (1.3)$$

where the β -function on the right-hand side is computed in perturbation theory. For QCD⁸, $\beta(g_3) = -7g^3/(4\pi)^2 < 0$. This means that the coupling strength of the strong force is stronger at lower energies and weaker at higher energies. The top quark lifetime is sufficiently short that it decays before the strong force confines. In contrast, the other quarks live long enough so that after $\sim 10^{-24}$ s, the strong force is so strong that quarks and gluons are created from the potential energy to surround the bare quarks in color neutral configurations (hadrons). Aside from the μ , τ , and top quark, all of the particles in Fig. 1.2 are hadrons. Hadrons built from a three-quark configuration are called *baryons* while those constructed from two-quark configurations are called *mesons*. One well-known baryon is the proton, which is composed of two up-quarks and one down-quark. The mass of the proton is about 1 GeV even though its constituent quarks have masses in the MeV range. This is due to the binding energy from the strong force⁹ that is realized by continuous exchange of soft gluons between the quarks. These gluons are also ‘in’ the proton and due to quantum fluctuations from gluon splitting, an entire *sea* of quarks and gluons are also ‘in’ the proton. The precise statement that a proton is composed of two up quarks and a down quark (called valence quarks) is

$$\int_0^1 (f_u(x) - f_{\bar{u}}(x)) dx = 2 \quad \int_0^1 (f_d(x) - f_{\bar{d}}(x)) dx = 1 \quad \int_0^1 (f_q(x) - f_{\bar{q}}(x)) dx = 0, \quad (1.4)$$

⁸The general equation for the three gauge couplings is given in Sec. 1.3. The general form for SU(N) given in e.g. Sec. 16.7 of Ref. [20] only applies to QCD because the other gauge couplings (before electroweak symmetry breaking) involve a scalar field. The factor $7 = 11 - \frac{2}{3}n_f$ assumes there are only 6 quarks.

⁹Ironically, even though the Higgs boson has been called the ‘God particle’ that gives all particles their mass, most of the mass around you is due to the binding energy from the strong force and not the Higgs mechanism. More appropriately, the Higgs mechanism gives rise to radioactivity through bestowing an electroweak scale mass to the W and Z bosons.

where $q \in \{s, c, b, t\}$ and $f_q(x)$ are *parton distribution functions* (PDF) that describe the probability density for a parton of type q to carry a momentum fraction x of the proton. These functions also depend on the energy scale $|Q|$ at which the proton is probed: $f_q(x) = f_q(x, Q^2)$. The timescale for the dynamics of parton creation and destruction inside the proton is bounded by $\sim 1/(1 \text{ GeV}) \sim 10^{-24}$ seconds. However, when two protons collide at the LHC with energies of $\sim 10 \text{ TeV}$, the protons pass through each other on the scale of $\sim 1/(10 \text{ TeV}) \sim 10^{-29}$ seconds. With this separation of scales, the proton-proton collision is really parton-parton scattering where all of the relevant non-perturbative information about the proton dynamics are neatly bundled into the PDFs. Figure 1.3 shows representative PDF sets at a hard-scatter scale of $Q^2 = (100 \text{ GeV})^2$. The valence u and d PDFs (defined as $f_q - f_{\bar{q}}$) dominate at high x and then all the sea quark PDFs approach each other (and diverge) at low x . The gluon PDF dominates below $x \sim 0.2$.




Figure 1.3: PDF sets from the NNPDF collaboration [32] at $Q^2 = (100 \text{ GeV})^2$ extracted using the HepData [33] interface. The left and right plots are the same aside from the scale on the horizontal axis.

As they are inherently non-perturbative, PDFs cannot be calculated from an α_s expansion in QCD, but the *energy-dependence* of PDFs can be calculated and is an essential ingredient to cross-section predictions at a hadron collider. The energy (Q^2) dependence is governed by the Dokshitzer-Gribov-Lipatov-Altarelli-Parisi (DGLAP) [34–36] equations:

$$\begin{aligned}\mu \frac{d}{d\mu} f_q(x, Q^2) &= \frac{\alpha_s(Q^2)}{\pi} \int_x^1 \frac{dz}{z} \left[P_{q \leftarrow q}(z) f_q\left(\frac{x}{z}, Q^2\right) + P_{q \leftarrow g}(z) f_g\left(\frac{x}{z}, Q^2\right) \right] \\ \mu \frac{d}{d\mu} f_{\bar{q}}(x, Q^2) &= \frac{\alpha_s(Q^2)}{\pi} \int_x^1 \frac{dz}{z} \left[P_{q \leftarrow q}(z) f_{\bar{q}}\left(\frac{x}{z}, Q^2\right) + P_{g \leftarrow q}(z) f_g\left(\frac{x}{z}, Q^2\right) \right] \\ \mu \frac{d}{d\mu} f_g(x, Q^2) &= \frac{\alpha_s(Q^2)}{\pi} \int_x^1 \frac{dz}{z} \left[P_{g \leftarrow g}(z) f_g\left(\frac{x}{z}, Q^2\right) \right. \\ &\quad \left. + P_{g \leftarrow q}(z) \sum_{q'} \left(f_{q'}\left(\frac{x}{z}, Q^2\right) + f_{\bar{q}'}\left(\frac{x}{z}, Q^2\right) \right) \right],\end{aligned}\quad (1.5)$$

where μ is the running scale and the functions $P_{p_2 \leftarrow p_1}$ are the Altarelli-Parisi splitting functions that encode the probability for a parton p_1 to split or radiate parton p_2 . At leading order in α_s , the splitting functions are given by

$$P_{q \leftarrow q}(x) = C_F \left[\frac{1+x^2}{1-x} \right]_+ \quad (1.6)$$

$$P_{g \leftarrow q}(x) = C_F \frac{1+(1-x)^2}{x} \quad (1.7)$$

$$P_{q \leftarrow g}(x) = T_R [x^2 + (1-x)^2] \quad (1.8)$$

$$P_{g \leftarrow g}(x) = 2C_A \left[\frac{x}{(1-x)_+} + \frac{1-x}{x} + x(1-x) \right] + \delta(1-x) \frac{11C_A - 4n_f T_R}{6}, \quad (1.9)$$

where the plus-notation $[f(x)]_+$ is defined in the context of an integral equation $\int_0^1 dx g(x) [f(x)]_+ = \int_0^1 dx (g(x) - g(1)) f(x)$. The ‘color factors’ $C_A = 3, C_F = 4/3$ and $T_F = 1/2$ are properties of the $SU(3)$ QCD algebra.

The intuition for e.g. the first line of Eq. 1.5 is as follows: a quark of type q with momentum fraction x could be due to either a quark of the same type with

proton momentum fraction $x' \geq x$ that has radiated a gluon or due to a gluon with momentum fraction $x' \geq x$ that has split into a $q\bar{q}$ pair. The energy fraction of the initial quark or gluon carried by the final quark or gluon is $z = x/x'$ which means that the proton momentum fraction of the initial quark or gluon is $x' = x/z$. The probability for the final quark to be from an initial quark is (heuristically) $\text{Pr}(q \text{ from } q) = \text{Pr}(\text{final } q | \text{initial } q) \text{Pr}(\text{initial } q) = P_{q \leftarrow q}(z) f_q(x/z)$. Similarly for the gluon term, $\text{Pr}(q \text{ from } g) = \text{Pr}(\text{final } q | \text{initial } g) \text{Pr}(\text{initial } g) = P_{q \leftarrow g}(z) f_g(x/z)$. The integral in the first line in Eq. 1.5 is over all emissions from the initial quark or gluon and the dz/z is the phase space for these emissions (see Sec. 4.1.1). The DGLAP equations and the splitting functions will be revisited in more detail in Part II.

Even though proton-proton collisions are well-described by parton-parton scattering, no out-going parton has ever been directly observed. As the out-going partons travel away from the interaction point, the same processes that fill protons with a sea of quarks and gluons generate a shower of partons described in the soft and collinear limits by the Altarelli-Parelli splitting functions. Once this *parton shower* has cooled to an energy $\lesssim 1$ GeV, the partons hadronize due to confinement. The resulting collimated spray of hadrons is known as a *jet*. Jets are ubiquitous at the LHC because of the prevalence of quark and gluon radiation from both the initial partons in the proton (initial-state radiation) as well from out-going quarks and gluons participating in the hard-scatter process (final-state radiation). Information about the initiating quark or gluon is embedded in the complex radiation pattern within jets. Part II is dedicated to study of this radiation.

1.2 Successes and Limitations

The Standard Model is incredibly successful. Increasingly precise calculations in the context of perturbation theory have accurately predicted and matched cross section measurements over 10 orders of magnitude at the LHC alone - see Fig. 1.4. These calculations and measurements span a wide range of processes probing all three fundamental forces. In addition to the LHC measurements shown in Fig. 1.4, there are numerous collider- and non-collider-based experiments that probe various aspects of the SM. One of the most impressive single measurement is the anomalous magnetic moment of the muon, g_μ . This quantity can be calculated and measured to nine significant figures. The value of $g - 2$ is dominated by QED, but the accuracy is such that there are non-negligible contributions from electroweak and hadronic processes in loops¹⁰. A related success story is the use of complementary measurements across colliders and energies to indirectly determine parameters that are not measured directly (global fits) [38]. Without m_h , the m_{top} uncertainty is ~ 4 GeV (assuming there is only the SM) when the direct measurement has ~ 1 GeV uncertainty.




Figure 1.4: Cross section measurements at $\sqrt{s} = 7, 8, 13$ TeV from the ATLAS experiment compared to theoretical measurements. From Ref. [39].

¹⁰For a nice overview, see Ref. [37]. There is some tension between $g - 2$ and the SM prediction, but the agreement spans many orders of magnitude.

Despite its great success, there are significant issues as well. Some of these problems are due to an inability of the SM to describe known phenomena while others are mostly aesthetic problems, but suggestive of the SM's incompleteness. There are no known logical inconsistencies with the theory. One important physical sector not described by the SM is gravity. Currently, there is no coherent quantum theory of gravity. For most physical phenomena that are known or will be studied in the near future, the SM augmented with general relativity is an accurate description of nature. However, there are extreme phenomena where quantum gravity effects are non-negligible (see Sec. 1.3 about the Planck scale). New approaches to QFT such as string theory are a promising direction, but there are currently no unique testable predictions from such models (see e.g. Ref. [40]).

Another aspect of nature not described by the SM is dark matter. There is overwhelming evidence¹¹ that most of the gravitationally interacting matter in the universe is not composed of SM particles. There is a small component of the non-luminous matter due to neutrinos, but they explain less than one percent of the total dark matter relic density (See e.g. Ref. [44]). Massive weakly interacting particles (WIMP) are an excellent dark matter candidate and these particles are a natural aspect of supersymmetric (SUSY) extensions of the SM. This will be revisited in more detail with Sec. 1.3.

Another phenomenon¹² not explained by the SM is the neutrino mass. In the SM introduced in Sec. 1.1, neutrinos are massless, but it is now known that neutrinos have a nonzero mass [45]. One could readily accommodate neutrino masses by adding a Yukawa coupling y_ν to Eq. 1.1. However, the neutrino masses are known to be less than about 0.3 eV [46], so there is an enormous hierarchy $y_{\nu_\tau}/y_\tau \sim 10^{-10}$ that is unexplained. Also, the off-diagonal elements of the corresponding mixing matrix (the CKM matrix analogue) are experimentally constrained to be much smaller than the off-diagonal elements of the CKM matrix. There is a plethora of theories to extend the SM to naturally explain the smallness of the neutrino mass - see Ref. [47] for a

¹¹See for instance the evidence from the velocity profile of galaxies [41] and from the analysis of colliding galaxy clusters [42, 43].

¹²There are others, including dark energy (related to gravity) and the imbalance between matter and anti-matter.

recent review.

The problem of neutrino masses is mostly aesthetic because there is a mechanism in the SM for generating neutrino masses, but the associated parameters (y_ν) are small. Another issue with the SM of this type is called the *strong CP problem*. In principle, there could be a term in Eq. 1.1 of the form $\theta F_{\mu\nu}^a \epsilon_{\mu\nu\rho\sigma} F^{\rho\sigma a}$, where θ is a dimensionless parameter and F is the gluon field strength tensor (see e.g. Ref. [48]). The parameter θ is experimentally constrained by neutron electric dipole moment measurements [49–51] to be $\theta < 10^{-10}$. A set of popular theories to naturally explain why θ is so small is to augment the SM with particles called axions [52–55].

Arguably the issue with the SM that has received the most theoretical and experimental attention is the size of the Higgs boson mass. Unlike the fermions and gauge bosons, there is no symmetry principle which protects the mass of the Higgs boson from quantum corrections. These corrections make the mass sensitive to particles and forces at the highest energy scales. As a result, there is an enormous unnatural hierarchy between the measured Higgs boson mass and the *Planck scale* (10^{19} GeV) where quantum gravity must be important. This *hierarchy problem* is one of the main motivations for SUSY and therefore an entire section is devoted to describe it properly (Sec. 1.3).

The SM has other (minor) aesthetic and practical problems that may suggest it is incomplete. For example, there is a large unexplained mass hierarchy for the known SM particles, $y_e/y_t \sim 10^{-6}$. In fact, there is no reason within the SM for any of the 18 dimensionless parameters. As a result of $g_1 \ll g_3$, perturbation theory results in extremely precise predictions for electrodynamic processes but low energy hadron spectra are incalculable in perturbation theory (and difficult to calculate with lattice techniques - see Ref. [56]). In addition, the non-perturbative nature of low energy QCD requires the introduction of phenomenological models for e.g. hadronization that have many new (non-fundamental) parameters.

The sensitivity to precise SM measurements to non-perturbative modeling will be revisited in Part II and all of Part III is dedicated to a search for new particles *beyond the SM* (BSM).

Chapter 2

Experimental Apparatus

Measuring the properties of nature at the smallest distance scales ever recorded with a terrestrial apparatus requires the highest energy particles accelerator ever built. The size of a structure is related to the probe energy via the de Broglie relation $\lambda = 1/p$. The same principle governs optical microscopes, limiting their resolution to hundreds of nanometers. The Large Hadron Collider (LHC) produces protons with energies up to $\sqrt{s} = 13$ TeV. At this energy, **A Toroidal LHC ApparatuS** (ATLAS) is able to capture the byproducts of the proton-proton collisions to probe distance scales as small as 10^{-20} m. For comparison, Fig. 2.1 shows various technologies that have been used to measure increasingly smaller distance scales. This chapter explores how the LHC works and how particles are measured by the ATLAS detector.




Figure 2.1: Diagram illustrating the length scales probes by various ‘microscopes.’

2.1 The Large Hadron Collider

The life of a proton at the LHC begins as hydrogen. After being stripped of its electron, the remaining protons proceed through a series of accelerators to successively increase their speed. To begin, a linear accelerator increases the proton energy to 50 MeV ($\beta \approx 5\%$). Then, a small (25 m in radius) circular accelerator called the Proton Synchrotron (PS) Booster increases the energy to 1.4 GeV ($\beta \approx 80\%$) after which a larger synchrotron (100 m in radius), the PS, increases the energy to 25 GeV ($\beta \approx 99.9\%$). While never itself used as a particle collider, the PS has a rich history [57] providing a variety of beams to other experiments such as a neutrino beam to the Gargamelle bubble chamber where weak neutral currents were discovered in 1974 [58]. Following the PS, protons are accelerated to 450 GeV in the 7 km (in circumference) Super Proton Synchrotron (SPS). Like the PS, the SPS provides beams for a variety of experiments. The SPS has also played an important role as a collider in its own right, such as facilitating the discovery of the W and Z bosons by UA1 [59,60] and UA2 [61,62]. The SPS directly injects into the 27 km LHC where the energy is ramped up to $\sqrt{s} = 8$ TeV (Run 1) or $\sqrt{s} = 13$ TeV (Run 2). There are a series of crossing points where the beams collide. An overview of the CERN accelerator complex is shown in Fig. 2.2. Prior to its use as a proton-proton collider, the LHC tunnel was filled with an e^+e^- accelerator called the Large Electron Positron (LEP) collider, which had four experiments. The LHC ring also has four collision points, with two multipurpose experiments ATLAS and the Compact Muon Solenoid (CMS) as well as two special detectors ALICE and LHCb.

Particle acceleration occurs via radio frequency (RF, 400 MHz) cavities driven by high-power klystrons. Each cavity in the LHC provides a gradient of 5 MV/m. There are eight such cavities, each supplying 2 MV for a total of 16 MeV added per beam per revolution. The RF must be an integer multiple of the revolution frequency (asynchronous protons slow down/speed up accordingly) which induces a *bucket structure* whereby there are about 36,000 possible locations along the beam for packets of protons. For protons traveling at nearly the speed of light, this means that the buckets are separated by about 2.5 ns. Only 10% of the possible buckets




Figure 2.2: An overview of the CERN accelerator complex. The LEP experimental areas are indicated in red while the four LHC experiments are labeled in blue.

are filled resulting in a nominal collision rate of 40 MHz. The buckets that are filled with *bunches* contain about 100 billion protons. Protons are steered around the LHC ring using dipole magnets (and focused using quadrupole magnets). The relationship between the (dipole) magnet strength B , the radius of the accelerator R , and the energy of the protons E is given by¹

$$E = \sqrt{4\pi\alpha} \times \left(\frac{B}{T}\right) \left(\frac{R}{m}\right) \approx 800 \times \left(\frac{B}{T}\right). \quad (2.1)$$

Therefore, at $\sqrt{s} = 8$ TeV, dipole magnets at ~ 5 T are required and at $\sqrt{s} = 13$, the magnets need to be powered to about 8 T. To achieve such high field magnets, superconducting Niobium-Titanium coils must be cooled down to 1.9 K using superfluid helium-4. From Eq. 2.1, for a fixed radius accelerator, the collision energy is set by the strength of the dipole magnetic field. The availability of robust high (enough)

¹This is $qvB = mv^2/R$ from equating the magnetic and centripetal forces, but accounting for the relativistic factor γ so that $p = \gamma mv = qBR$. For protons, $E \sim p$, and the elementary charge is $e = \sqrt{4\pi\alpha}$. See Table 1 for the unit conversions.

temperature superconductors was a limiting factor to the design energy of the LHC. There are promising alternatives to Ni-Ti such as Ni₃-Sn, but this is still an area of active research. One of the main limitations in the number of protons per bunch is from the cooling of the magnets. Accelerating charged particles radiate, resulting in an emitted power per proton (synchrotron radiation) given by²

$$P \approx (4 \times 10^{-15} \text{ W}) \left(\frac{B}{T} \right)^2 \left(\frac{E}{\text{TeV}} \right)^2 \left(\frac{m}{\text{GeV}} \right)^{-4}, \quad (2.2)$$

where $m \sim 1 \text{ GeV}$ for protons. Due to high power of m in the denominator of Eq. 2.2, this is a severe limitation for electron beams. For proton beams, the synchrotron radiation is highly suppressed but if there are $N_b \sim 10^{11}$ protons per bunch, then the total power at $\sqrt{s} = 13 \text{ TeV}$ per bunch is $N_b P \sim 1 \text{ W}$ ³. There are about 3,000 bunches (just below 10% of the available buckets) so the total power per beam per meter is about 0.1 W/m. This is one of the main challenges to the cryogenics [66] and leads to an important justification for keeping $N_b \lesssim 10^{11}$. Another factor is the collision rate, which is set by the size of the bunches and the number of protons per bunch. The transverse size of a bunch σ is given by $\sqrt{\beta\epsilon}$ where ϵ is the area of the beam in phase space (*beam emittance*) and β (*betatron function*) captures the changes in the beam due to focusing magnets. By Louisville's theorem, ϵ does not depend on the position along the LHC. In a region without a magnetic field, the betatron function has the form [67]

$$\beta(z - z_0) = \beta(z_0) + \frac{(z - z_0)^2}{\beta(z_0)}. \quad (2.3)$$

When z_0 is the collision point, $\beta(z_0)$ is called β^* and because ϵ is constant along the beam, by Eq. 2.3, β^* measures the distance from the collision point at which the

²For a derivation, see e.g. Chapter 8 in Ref. [63] or Chapter 14 in Ref. [64]. The full formula is $P = e^4 \gamma^2 B^2 / (6\pi\epsilon_0 m^2 c)$. See Table 1 for unit conversions.

³The synchrotron radiation may become dominant in the not-to-distant future if $\sqrt{s}$ is increased by a factor of 10 at a future collider [65].

transverse size σ doubles. At the LHC, $\beta^* \sim 0.5$ m and the *normalized emittance* $\epsilon_N = \beta\gamma\epsilon \sim 3\mu\text{m}$ [68]. The emittance itself is actually not conserved along the beam for nonzero acceleration, but the normalized emittance does obey Louisville's theorem. At $\sqrt{s} = 13$ TeV, $\beta\gamma \sim 13000$ so $\epsilon \sim 2 \times 10^{-10}$ m resulting in the physical beam size in the lab frame $\sigma \sim 10 \mu\text{m}$. Note that this is significantly larger than the 'size' of the proton, which is about one femtometer $\sigma_p \sim 10^{-15}$ m. The probability for one proton-proton collision could be estimated by $p \sim \sigma_p^2/\sigma^2 \ll 1$. For $N_b \sim 10^{11}$ protons per bunch, the average number of collisions is $p \sim N_b^2 \sigma_p^2/\sigma^2 \sim 100$, which is another reason to keep $N_b \sim 10^{11}$ ⁴. This quick calculation demonstrates that the average number of events from a particular process in a given bunch crossing can be calculated as the product of a process-dependent *cross section* and a quantity related to the rate of collisions. The later quantity will be called the *instantaneous luminosity* and is given in full by

$$\mathcal{L} = \frac{N_b^2 f n_{\text{bunches}} F}{4\pi\epsilon\beta^*}, \quad (2.4)$$

where f is the revolution frequency (40 MHz) and F is a $\mathcal{O}(1)$ geometric factor to correct for an off-axis crossing angle. Van Der Meer [69] scans are combined with a variety of techniques for measuring $\mathcal{L}$ in-situ [70]. For the data collected so far at the LHC, $\mathcal{L} \sim 10^{34}\text{-}10^{35} \text{ cm}^{-2}\text{s}^{-1}$.

The *integrated luminosity*, $\int \mathcal{L} dt$ is used to quantify the amount of collected data. Units of the integrated luminosity are *inverse barns* ($\text{b}^{-1} \approx 10^{28}/\text{m}^2$); the full $\sqrt{s} = 8$ TeV dataset was about 20 fb^{-1} . For any process $pp \rightarrow X$, the average number of predicted events for that process is given by $\int dt \mathcal{L} \sigma_{pp \rightarrow X}$, for $\sigma_{pp \rightarrow X}$ calculated⁵ in barns.

⁴Not all of these collisions result in interesting *inelastic* scattering in which the protons dissociate (about 50%). The actual number of collisions per bunch crossing will be called μ and is discussed in the context of *pileup* in Sec. 2.3.

⁵From QFT and corrected for various detector effects, discussed in Parts II and III.

2.2 Interactions of Particles with Matter

Typical proton-proton collisions delivered by the LHC result in hundreds of particles scattering away from the interaction point. There are two ways to measure the properties of these particles⁶. One possibility is to passively observe secondary particle production without disturbing the trajectory of the primary particle. These techniques are available for charged particles, which can interact electromagnetically with a detector without losing a significant fraction of their energy. The distribution of secondary particles contains information about the momentum and type of the original particle. If a series of such measurements are made along the trajectory of the particle, a fit can reconstruct the particle trajectory with high precision. The transverse momentum of a particle is related to the curvature of its trajectory in a magnetic field perpendicular to its motion by⁷

$$\frac{p_T}{\text{GeV}} \approx 0.3 \left(\frac{\text{B}}{\text{T}} \right) \left(\frac{\text{R}}{\text{m}} \right), \quad (2.5)$$

where one unit of the electric charge is about 0.3 in natural units and $\text{Tm} \sim 10^{-1}$ (see Table 1). In particular, for particles with $p_T \lesssim 0.3$ GeV in a 2 T magnetic field, they will never travel further than one meter. Figure 2.3 shows the trajectory of charged particles in a solenoidal magnetic field with the same setup as the ATLAS detector, discussed in Sec. 2.3. Note that charged particles of the opposite charge would bend *down* instead of up in Fig. 2.3. The momentum resolution from track fitting is determined by how well the *sagitta* can be measured. Figure 2.4 shows an example charged particle trajectory in the same coordinates as Fig. 2.3 where three measurements (dots) of the *track* have been measured. The sagitta s is related to the radius R by $s = R(1 - \cos \alpha)$. When $\alpha \ll 1$, $s \approx \frac{1}{R}\alpha^2$. Also in this approximation, $\alpha \approx \frac{1}{2}L/R$. Using Eq. 2.5, this gives the result $s \approx \frac{1}{8} \frac{L^2 e B}{p_T}$. Linear propagation of errors

⁶This section will briefly introduce some of the main concepts of particle detection techniques. There are many books on this subject; see e.g. Ref. [71, 72] and the PDG review [27].

⁷This is the same formula used to derive Eq. 2.1, only now the magnetic field is parallel to the beam and perpendicular to the particle trajectory.

shows that $\sigma_{p_T}/p_T \approx \sigma_s/s$. Therefore, $\sigma_{p_T}/p_T \propto \sigma_s p_T/L^2 B$; the resolution is worse at high p_T and can be improved with a longer lever arm L and a higher magnetic field B . The resolution σ_s is independent of L and B and scales as $\sigma_s \propto 1/\sqrt{N}$ for enough hit measurements N ⁸. In addition to measurement uncertainty, there is a contribution to the resolution from *multiple scattering* of the primary particle in the detector material. This term is approximately independent of momentum and scales as [71, 72] $\sigma_{p_T}/p_T \propto \frac{1}{BL\beta} \sqrt{L/X_0}$, where X_0 is the *radiation length* of the detector material⁹.




Figure 2.3: Diagrams illustrating the trajectory of charged particles in a magnetic field. Both views are cross sections of the detector with the beam axis and B-field perpendicular to the page. Circles on the left plot indicate the locations of various tracking detector elements for the ATLAS detector, discussed in Sec. 2.3. A 2 T solenoid magnet is used to determine R .

A second possibility for particle detection is to stop the primary particle and measure the heat deposited in the material (*calorimeter*)¹⁰. Both charged and neutral

⁸The exact formula is derived in Ref. [73], which also has an interesting discussion about the optimal spacing of measurements.

⁹The radiation of a material is the characteristic length for energy loss via Bremsstrahlung; quantitatively, $dE/dx = E/X_0$. After X_0 , the particle has only $1/e$ of its original energy.

¹⁰A detailed description of calorimetry is in the dedicated textbook Ref. [74].




Figure 2.4: A schematic diagram used to derive the momentum resolution of a charged particle track measurement. The length L (*lever arm*) is the distance over which measurement of the trajectory can be made. The three dots indicate discrete samplings (measurements) of the trajectory.

particles can be measured this way and the location of the energy deposition provides information about the particle's momentum direction. For electrons, energy loss at high energy is dominated by the same Bremsstrahlung that is a nuisance for tracking detectors. Photons undergo pair production, which is also set by the radiation length X_0 ¹¹. The radiation length for muons is much larger than for electrons. Muons deposit energy via ionization (governed by Bethe-Bloch), but are not stopped by reasonably sized calorimeters. Hadrons lose energy by a combination of ionization (charged particles) and nuclear (i.e. via the strong force) interactions. Electromagnetically decaying hadrons subsequently lose energy via the processes listed above for leptons and photons. These interactions are characterized by the *hadronic interaction length* λ . For a given material λ is often much larger than X_0 ; for example in liquid argon (used by the ATLAS calorimeters), $\lambda \sim 6X_0$.

Energy in a calorimeter is lost via a cascade of collisions. Radiative and hadronic

¹¹Though the exact dependence for photons is different than for electrons: the probability for a photon to survive a distance x before pair production is $e^{-\frac{7}{9} \frac{x}{X_0}}$.

processes result in significant energy loss until ionization or other low energy phenomena dominate and the remaining particles slowly lose energy and are absorbed. The crossover energy is called the *critical energy* E_c . If each collision occurs after time δt and results in a reduction of the primary particle energy by $\frac{1}{2}$, then the timespan of a particle shower in the calorimeter is proportional to $\log_2(E/E_c)$. In general, the depth of a shower scales *logarithmically* with the energy. For this reason, calorimeters of a fixed depth can detect particles over many decades in energy. Showers initiated by electromagnetic particles are shallower than those started by nuclear processes. For this reason, calorimeters specifically optimized for detecting electromagnetic showers are closer to the interaction point than *thicker* calorimeters aimed at stopping nuclear showers from hadrons. The transverse size of a shower is also significantly larger for hadronic showers compared with electromagnetic showers, which scale with λ and X_0 (called the Molière radius), respectively.

In contrast to tracking detectors, the resolution of a calorimeter *decreases* with energy. The energy in a calorimeter is related to the number of particles produced in the shower; as such, the energy follows a Poisson distribution: $\sigma_E/E \propto 1/\sqrt{E}$. Estimating the proportionality constant is complicated because the hadronic/electromagnetic composition plays a significant role in determining the resolution. As with tracking detectors, calorimeters also have a constant term due to a variety of sources, such as differences in behavior for electromagnetic and hadronic showers [71]. Additional sources of (e.g. electronic) noise result in a constant energy resolution that is independent of the primary particles. This results in a term $\sigma_E/E \propto 1/E$.

In addition to measuring the momentum or energy of a particle, detectors can be used to infer the particle type. Figure 2.5 shows the average distance that various particles travel in the lab frame before decaying, $\beta\gamma\tau$, as a function of p_T . Charged pions and muons travel well past any detector element before decaying. For pions, this is largely irrelevant because of nuclear interactions that stops them in the calorimeters. Muons lose only a small amount of energy in the tracking detectors and calorimeters. Therefore, one can identify muons by placing an additional set of tracking detectors *beyond* the calorimeters. Except for occasional *punch-through* hadrons and low energy sources of radiation around the detector, particles measured in these outer tracking

chambers can be identified as muons. On the other end of the spectrum, neutral pions decay nearly immediately after production into two photons. The angular distribution between the two photons scales as $2m_\pi/p_T$ (see Chapter 4). One of the reasons that the electromagnetic calorimeter needs to be very finely segmented is to separate high p_T isolated photons from two photons produced collinearly from a pion decay. There are a class of particles in Fig. 2.5 that can be produced at the primary collision, but decay after macroscopic distances in the detector. The existence and properties of the secondary decay vertices for these particles are powerful observables for separating b -quark jets, c -quark jets, and hadronically decaying τ -lepton jets from light(er) quark and gluon jets. Since top quark jets decay nearly 100% of the time to b -quarks, b -quark jet tagging will be a critical aspect of the measurements and search presented in Parts II and III. Reference [75] describes b -quark jet tagging with the ATLAS detector, which is also discussed when used in subsequent chapters.

Various other particle type identification techniques exist that require specialty detector elements or work only in a limited kinematic region. For example, as part of the ATLAS tracking detector, there is a special subdetector for transition radiation emitted by electrons when they traverse different materials. This is useful for separating charged pion tracks from electron tracks. The tracking detector can also measure the amount of energy deposited per unit distance, dE/dx , which depends on the mass of the primary particle. This information can be used to separate proton, pion, kaon, and possibly new massive particle tracks from each other as long as $\beta\gamma \lesssim 1$ [76].

Figure 2.6 presents a schematic overview for the detector signature of various particle classes. The majority of hadrons in jets are pions because they are the lightest hadron. The mass of the light hadrons is insignificant compared with their momenta when $p_T \gg 1$ GeV and so there is no distinction between various charged hadron types. Section 2.3 presents an overview of all the ATLAS detector elements, which will follow the pattern in Fig. 2.6.




Figure 2.5: The average transverse distance traveled by various particles labeled on the horizontal axis as a function of their p_T .




Figure 2.6: An overview of the measurement pattern for various particle types in the ATLAS detector. A dashed line means that the given particle leaves no trace in the sub-detector. A solid horizontal line indicates insignificant energy loss (ionization and small levels of multiple scattering for electrons).

2.3 The ATLAS Detector

ATLAS is a general-purpose detector designed to measure the properties of particles produced in high-energy pp collisions with nearly a full 4π coverage in solid angle¹². In order to provide shielding from cosmic rays (and reduce costs), the LHC and the cavern containing the ATLAS detector are about 100 m below ground. The innermost subsystem of the detector is a series of tracking devices used to measure charged-particle trajectories bent in a 2 T axial field provided by a solenoid whose axis is parallel with the beam direction. This inner detector (ID) consists of a silicon pixel detector surrounded by a semiconductor microstrip detector (SCT) and a straw-tube tracker that can detect electron transition radiation (TRT) (Sec. 2.3.1). Surrounding the ID are electromagnetic and hadronic calorimeters that use liquid argon and scintillating tile as active media (Sec. 2.3.2). Beyond the calorimeters is a 4 T toroidal magnetic field and a multi-component tracking system for muon detection (Sec. 2.3.3). Section 2.3.4 discusses the data acquisition including the *trigger*. A diagram of the subsystems of ATLAS is shown in Fig. 2.7. For scale, people are shown walking on the cavern floor as well as between the muon chamber wheels.




Figure 2.7: The ATLAS detector and its subsystems (to scale). Image from Ref. [78].

¹²This section is intended to be a brief overview - for many more details, see Ref. [77].

2.3.1 Inner Detector

The innermost layer of ATLAS is a series of tracking detectors with three different technologies, illustrated by Fig. 2.8. Closest to the beam pipe is a pixel detector, which is composed of 3 (4) layers in Run 1 (2). A new pixel layer (insertable b-layer, or IBL) was inserted closer to the collision point between Runs 1 and 2. This was a significant technical challenge because the beampipe had to be removed and replaced with a smaller radius pipe upon which the IBL was mounted and inserted with all of its services into the small space inside the original pixel detector. The left image of Fig 2.9 shows the just-inserted IBL before the service cables were unwound. This winding was needed in order to connect tables to the side of the IBL opposite the insertion. These service cables (assembled at SLAC) were thoroughly tested at every stage of processing, including before and after mechanical stress tests such as a practice winding. The right plot of Fig. 2.9 shows the difference in resistances on all the pins of one data cable before and after a practice winding. As desired, the resistance is unchanged.

The original three pixel layers are composed of 250 μm thick planar sensors most with a $50 \times 400 \mu\text{m}^2$ surface area. In order to cope with a higher radiation dose, the IBL sensors are smaller $50 \times 250 \mu\text{m}^2$ (also thinner) planar sensors in the central region and 3D sensors [79] at high $|\eta|$ with charges drifting perpendicular to the sensor depth instead of parallel. Beyond the pixel detector are four layers of silicon microstrips (SCT). In order to provide a (crude) measurement along the z direction, each SCT module has two sensors that are rotated by ± 20 mrad with respect to each other. Each module provides about $20 \mu\text{m}$ resolution in the azimuthal direction and about $600 \mu\text{m}$ resolution along z . Surrounding both the silicon-based detectors is an annulus between about 50 cm to 1 m filled with 2 mm radius drift tubes. The region around the tubes is filled with a material that enhances the electron transition radiation. Charged particles leave ionization energy in an average of 36 tubes of this transition radiation tracker (TRT). Dedicated low and high thresholds are used to measure minimum ionizing particles and energy from X-ray photons due electron transition radiation, respectively.




Figure 2.8: Left: A schematic view of the subsystems of the ATLAS inner detector [80]. The IBL was added between Runs 1 and 2. The red line indicates the trajectory of a hypothetical particle with $p_T = 10$ GeV at $\eta = 0.3$. Right: an enlarged view of the pixel detector prior to the insertion of the IBL [81].




Figure 2.9: Left: A picture of the IBL just after being inserted into the detector. The yellow cables are wrapped in order to pass through to the far side. These cables host data transmission, high/low voltage, and many other services. Right: The difference in the measured resistance in the data pins before and after a practice wrapping.

2.3.2 Electromagnetic and Hadronic Calorimeters

Surrounding the ID and solenoid are electromagnetic and hadronic calorimeters to measure showers from charged and neutral particles. A high-granularity lead/liquid-argon (LAr) sampling electromagnetic calorimeter is located just beyond the solenoid and spans the range $|\eta| < 3.2$. Beyond the electromagnetic calorimeter is a two-component hadronic calorimeter that uses steel absorbers and scintillator-tile sampling technology in the range $|\eta| < 1.7$ and copper/LAr sampling technology for $1.5 < |\eta| < 3.2$. Additional calorimetry is provided up to $|\eta| = 4.9$ using copper (tungsten)/LAr in the electromagnetic (hadronic) sections. Figure 2.10 is a schematic diagram of the various calorimeter components.

The bulk of energy deposited in the electromagnetic calorimeter is deposited in the second layer which contains about $17X_0$ (out of about $23X_0$) with a granularity of 0.025×0.025 in $\Delta\eta \times \Delta\phi$. In contrast, most of the hadronic energy is deposited in the first two layers of the hadronic calorimeter with about 5.5 (out of about 7.5) hadronic interaction lengths λ with a granularity of 0.1×0.1 in $\Delta\eta \times \Delta\phi$. The total detector thickness is about 10 hadronic interaction lengths at $\eta = 0$. The inner detector material accounts for about $0.5X_0$ and 0.2λ at $\eta = 0$ and about $2X_0$ and 0.7λ just beyond the edge of the ID ($\eta \sim 1.5$). Most of this material is in the form of support structures, coolant, electronics, and cables.




Figure 2.10: A schematic diagram of the ATLAS calorimeters. Image from Ref. [82].

2.3.3 Muon Spectrometer

Surrounding the calorimeters is a muon spectrometer with trigger and precision chambers. Four different detector technologies are used for these purposes. Monitored Drift Tubes (MDT) provide precision tracking in the central region (except for a small gap at $|\eta| \approx 0$ for services) and Cathode Strip Chambers (CSC) cover the forward region $2 < |\eta| < 2.7$. The MDTs have a $35 \mu\text{m}$ resolution along z while the CSCs measure both the z and radial position with $40 \mu\text{m}$ and 5 mm resolutions, respectively. The main reason for using CSCs in the forward region is the higher particle flux so the second coordinate measurement is important for resolving track ambiguities. A long drift time in the MDT ($\mathcal{O}(100)$ ns) makes them unusable for triggering (25 ns crossings). Therefore, two additional detectors are dedicated to triggering: Resistive Plate Chambers (RPC) in the central region ($|\eta| < 1.05$) and Thin Gap Chambers (TGC) up to $|\eta| = 2.4$. RPCs are parallel plate capacitors filled with gas and separated radially for a crude but fast momentum measurement. TGCs are multi-wire proportional chambers with a finer granularity than RPCs in order to cope with the higher multiplicity and reduced track bending (for a fixed p_T) in the forward region.




Figure 2.11: A candidate $H \rightarrow WW^* \rightarrow 4\mu$ event. The long purple lines indicate reconstructed muon tracks and each muon tracking chamber contributing to those tracks is highlighted in green (MDT) or purple (CSC). The invariant mass of the four muons is 145.8 GeV (inconsistent with the now known $m_H \sim 125 \text{ GeV}$).

2.3.4 Trigger System

Due to the large event rate, not every collision can be recorded for processing offline. Events are selected using a three- (two-)level trigger system that is hardware-based at the first level and software-based for the (two) following level(s) in Run 1 (2) [83, 84]. The nominal interaction rate is 40 MHz ($=1/25$ ns). It is physically not feasible and undesirable to read out all detector elements at this rate. First of all, the cross-section for interesting hard-scatter events is significantly below the total cross-section. For example, the W +jets cross-section (highest rate non-QCD process) is about 100 nb [85] while the total inelastic (any collision other than $pp \rightarrow pp$) cross-section is about 70 mb and the total cross-section is about 95 mb [86]. With $\mathcal{O}(10)$ simultaneous pp collisions (pileup) per bunch crossing, this means that only one bunch crossing in 10^5 produces anything potentially interesting; the rate for $t\bar{t}$, Higgs, etc. is even lower. Another problem is storage space: an entire event is $\mathcal{O}(1)$ MB. A readout of 40 MHz would result in about 100 TB/s. These challenges are circumvented by quickly deciding if an event should be saved or discarded. An event must satisfy all three (two) trigger levels to be recorded for further processing. The hardware-based trigger system has about $2.5 \mu\text{s}$ to make this decision and reduces the rate from 40 MHz to 70 (Run 1) or 100 (Run 2) kHz. The total software based trigger operates on an $\mathcal{O}(1)$ s timescale with an output of 400 (Run 1) or 1000 (Run 2) Hz readout with a high efficiency for physics processes of interest. The highest level (software) trigger uses offline-like algorithms while the hardware-based trigger uses crude approximations to object reconstruction in order to increase the speed. There is redundancy built into the trigger system in order to use one trigger to study another and ensure a high efficiency for processes of interest. An event can fire multiple triggers, though only one is required to record the event. As many interesting processes (and rarely any uninteresting processes) contain leptons in the final state, the largest trigger bandwidth is dedicated to single lepton triggers ($\gtrsim 20\%$). Some triggers are *pre-scaled* in order to artificially reduce the rate. A pre-scale is implemented by randomly keeping only a fraction of events that pass a given trigger. The pre-scale rates are known, so the luminosity of the data can be corrected to account for the rate reduction (see Sec. 1.2.1.1 for more detail).

Chapter 3

Monte Carlo Simulation

In order to interpret the ATLAS data in the context of the SM or any other theory, precise predictions for the detector output are required. This output depends on physical processes occurring on length scales spanning 10^{-20} m up to the macroscopic size of the detector at $\mathcal{O}(10)$ m. High precision simulation of this entire process is possible because the physical laws *factorize*: in order to understand the behavior at one length scale, it is only necessary to know what happened previously at one length scale smaller. Therefore, each length scale is simulated in series. At the core of each step is a Monte Carlo (MC) integration. The basic idea of a MC simulation is that the expected value of a function can be approximated by computing the average value of the function when sampling many times from the underlying probability distribution. Factorization is realized by using *Markov Chain* MC in which the stochastic evolution of a simulated event at one stage only depends on the previous stage. Some aspects of a simulated event are unphysical, these parts are often called the *MC truth* because they are unknowable in reality. At the stage when the simulated events represent the same information that is present in a real data event, the simulated event is treated exactly as if it were a real event when reconstructing the final state in terms of high level objects (see Chapter 4). The only difference is that one simulated event often represents far fewer than one real event; in order for the averaging to be useful (have small uncertainty), the number of simulated events needs to (greatly) exceed the number of real events.

The following sections briefly introduce the various stages of simulation. At the smallest distance scales, perturbative calculations are combined with MC techniques to generate the *hard-scatter* process, which describes the collision of partons to produce the process of interest (Sec. 3.1). Next, perturbative scale evolution takes the outgoing colored particles through radiation down to $\mathcal{O}(1)$ GeV where QCD is no longer well-described by perturbation theory (electromagnetic effects are also included). Phenomenological models are then used to convert the quarks and gluons into hadrons as well as describe the (relatively) soft processes related with additional radiation in the event (*underlying event* and *multiple parton interactions*) (Sec. 3.2). Any particle with $\tau \lesssim 30$ ps is decayed before modeling the interaction of the remaining particles with the various detector elements, including inactive components. The last step in the simulation chain is to model the detector response by converting the energy deposited into digital signals, including the effect of noise (Sec. 3.3).

3.1 Matrix Elements

Matrix element (ME) calculations describe the hard-scatter process of interest and are computed at fixed order in α_s . In order for such a calculation to be useful for later stages of the simulation, these calculations must be interfaced with another generator that simulates the parton shower (PS) described at the end of Sec. 1.1. The fundamental challenge of the combined ME+PS simulation is how to treat the overlapping soft and collinear regions of phase space. At lowest order in α_s , this problem is manifest when additional quarks and gluons are included in the ME calculation. This is solved by *merging* ME calculations with a PS simulation. There are several approaches to merging (see Ref. [87] for a comparative review), but the idea used in all of them is to veto emissions in the PS that overlap the ME phase space and then apply event weights based on the probability of the ordering of ME emissions (*Sudakov form factors*) involving the splitting functions from Sec. 1.1. The ME emissions are preferred to the PS ones because they better describe hard and wide angle radiation. At next-to-leading-order (NLO) in α_s , there is a phase space overlap between the first real emission with the radiation from the PS. There are two common

schemes for subtracting the overlap from the ME calculation (MC@NLO [88]) or from the PS (POWHEG [89]) while still maintaining NLO accuracy in the *matched* simulation. See Ref. [90] for a review of these methods. Algorithms have also recently been developed and (partially) automated to simultaneously match and merge with extra out-going partons in the ME at NLO. Three actively developed approaches are UNLOPS [91], FxFx [92], and MEPS@NLO [93, 94].

3.2 Fragmentation and the Underlying Event

The role of a PS simulation is to evolve outgoing colored partons from a starting scale $t_{\text{hard}} \sim \mathcal{O}(10)\text{--}\mathcal{O}(100)$ GeV down to a cutoff t_{cutoff} at which QCD perturbation theory is no longer valid. The most common MC implementations of the PS are based on a Markov Chain of $1 \rightarrow 2$ splittings from scale t_i to t_{i+1} with no-emission probabilities between these two scales given by exponentiating the leading order splitting functions. This is a *leading logarithm* (LL) approximation which numerically accounts for the resummation of logs of the opening angle ρ of the radiation $(\alpha_s \log^2 \rho)^n$ to all orders in perturbation theory (see Sec. 4.1.1 for more detail). The shower is produced in the limit that there is an infinite number of color charges ($N_c = \infty$) to avoid complicated non-local effects. Corrections to this picture are suppressed by $1/N_c^2 \sim 1/10$. Many modern PS generators include effects beyond LL and also beyond leading color. The two most widely used PS generators¹ are PYTHIA [96] and HERWIG [97], which are distinguished by their choice of t . PYTHIA uses a p_T -ordered shower [98] while HERWIG++ uses angular ordering [99] in order to explicitly account for coherence effects (see Sec. 2.1.1). Electromagnetic radiation is also included in the modeling of fragmentation (sometimes with dedicated generators like PHOTOS [100]), but it is suppressed by $\alpha/\alpha_s \sim 1/10$.

After the PS, the remaining partons are combined into color-neutral hadrons.

¹SHERPA [95] is also a widely used generator, but its PS and hadronization models are conceptually similar to HERWIG.

There is no first-principles model of hadronization², so PYTHIA and HERWIG implement physically-inspired phenomenological models with various tunable parameters that can be adjusted to match data. A model based on color strings with a tension to represent the non-perturbative strong force (Lund string model [101]) is used by PYTHIA while a cluster model is used by HERWIG [102]. The two generators also differ in other aspects of non-perturbative modeling such as for the underlying event [103, 104], which is the production of radiation from the same pp collision as the hard-scatter process, but not directly involving the two scattering partons. Phenomenological models designed to describe non-perturbative physical effects have many parameters which are tuned to data. Specific sets of parameter tunes are described in later sections. During the hadronization process, unstable particles are decayed, including B-hadrons and τ -leptons, often with decay tables from dedicated programs like EvtGen [105] and Tauola [106], respectively.

Multiple simultaneous interactions (pileup) are modeled by overlaying independent *minimum bias* (inelastic) events on top of the hard-scatter event. The number of such collisions is stochastic and modeled to match the pileup level in data. This only accounts for the *in-time pileup*: radiation resulting from collisions that occurred in the same bunch crossing as the primary hard-scatter event. *Out-of-time* pileup from bunch crossings before or after the primary one are modeled in the same way, but are offset in time in the simulation to allow for an accurate model of the signal processing that can take > 25 ns [107].

Each simulation setup in the subsequent chapters will be specified by the ME and PS generators as well as the various perturbative and non-perturbative tunable parameters.

²The words *fragmentation* and *hadronization* are often used interchangeably, but can also mean different processes depending on the context. In this document, hadronization will refer to the transition between the end of the parton shower and the formation of hadrons, whereas fragmentation includes both the parton shower and hadronization.

3.3 Material Interactions and Detector Simulation

Up to this stage, all the steps of the event generation only depend on the beam type (pp) and beam energy ($\sqrt{s} = 8$ or 13 TeV). After fragmentation, hadrons begin to interact with the detector material and so all subsequent stages are tailored to the ATLAS detector composition and geometry. A detailed model of each detector element, including inactive material, is constructed and imported into the GEANT4 generator [108]. Particles produced from fragmentation are propagated through each subdetector and the nuclear interactions are modeled using a variety of physically-inspired models³. Custom algorithms for each subdetector then transform the energy deposited into an analogue and/or digital signal and model the readout [110]. For example, when a high energy pion traverses the doped silicon inside a planar pixel detector, GEANT4 stochastically calculates the energy deposited along the path length. This energy is assigned to low energy electrons or holes⁴ that are propagated (including thermal diffusion) to the collecting electrodes. The collected charge is converted into a time over threshold (TOT), which is digitized into 4 (IBL) or 8 bits. The various voltage and tuning parameters of the sensor and readout are part of the simulation. One important condition that is not currently part of the pixel simulation is the radiation level, which can degrade charge collection. The Run 1 and early Run 2 dosages are likely not sufficient for a significant degradation in performance, but this will be an important phenomenon to model in the future. See Appendix A for further details about modeling radiation damage.

³Analyses that are not particularly sensitive to local fluctuations in the energy deposited in the calorimeter use a parameterized description that significantly speeds up the simulation time [109].

⁴Holes are gaps in the electron Fermi distribution that propagate as if they were a positive charged particle.

Chapter 4

Event Reconstruction

Once the data are collected (or simulated events are generated), pattern recognition algorithms are employed to reconstruct basic physical objects. The first step in this process is to build low-level objects representing individual particles. In the inner detector, tracks are constructed from space point hits (Sec. 4.2) and in the calorimeter, calorimeter-cell clusters are formed (Sec. 4.1). In order to reject tracks that do not originate from particles produced in the primary collision, various *quality criteria* are imposed for tracks used in subsequent analysis. The energy of calorimeter-cell clusters is corrected (*calibrated*) based on shower properties so that it is an unbiased measurement of the initiating particle energy. From tracks and calorimeter-cell clusters, electrons, photons, muons, taus (Sec. 4.4), and jets (Sec. 4.3) are constructed. Object properties are used to construct *particle identification* schemes intended to reject objects of one type mis-identified as another type. Various corrections are applied to ensure that the energy or momentum of the reconstructed objects are calibrated. All of the aforementioned objects are then used to construct the missing transverse momentum, which is a measure of the momentum carried away by undetected particles such as neutrinos (Sec. 4.5).

4.1 Calorimeter-cell Clusters

Energy deposits in the calorimeter that are likely to have originated from a single hadron shower are grouped into calorimeter-cell clusters called *topo-clusters* [82]. First, cells with energy exceeding 4σ above the noise are labeled as seeds¹. Any neighboring cells (or neighbors of the neighbors) with energy exceeding the noise by 2σ are added to the seeds. This second step is repeated, but with a lower threshold of zero energy. The remaining *topologically connected* clusters with local maxima are split into multiple pieces, resulting in the final topo-clusters. Calorimeter noise is highly η dependent as a result of changes in detector technology and at $\mu = 30$ ranges from about 70 MeV per layer of the Tile calorimeter for $|\eta| < 1.5$ to 1-10 GeV in the forward calorimeter at $|\eta| > 4.5$.

A local cluster weighting (LCW) scheme is used to correct for biases in the energy assigned to each topo-cluster [82]. Corrections are applied to each cluster to correct for energy in the calorimeter but outside the cluster, for energy lost in inactive material, and for the different response to the EM and hadronic components of the shower.

4.2 Charged Particle Tracks

Charged-particle tracks are reconstructed from all three inner detector components, providing measurements of the transverse momentum of tracks with a resolution $\sigma_{p_T}/p_T \approx 0.05\% \times p_T / \text{GeV} \oplus 1\%$, where $\oplus$ indicates a sum in quadrature. The track reconstruction algorithm fits five track parameters: d_0 , z_0 , ϕ , θ , and q/p , where d_0 and z_0 are the transverse and longitudinal impact parameters, respectively, q is the track charge and p is the track momentum. Reference [111] provides a detailed explanation of the various algorithms used to build tracks and Fig. 4.1 shows an example event display from the early Run 2 data where tracks are reconstructed from all three ID subdetectors including the IBL. Excellent spatial precision is required to

¹There is a subtle point that the absolute value of the energy is used. Due to the shaping function in the LAr calorimeter, out of time pileup from previous bunch-crossings can result in negative energy. Retaining these clusters can be useful for canceling positive energy fluctuations from pileup.

maintain a well-performing track reconstruction out to and exceeding charged-particle p_T of 1 TeV, where track sagittas are $\lesssim 0.2$ mm. The large particle density in the core of high p_T jets is a challenge for track reconstruction. At low p_T , *fake tracks* that are due to combinations of hits from several particles can be suppressed by reducing the number of tracks with shared hits in the pixel detector. However, at high p_T , real tracks can also have shared hits on the pixel detector. At $\eta \approx 0$, the innermost (non-IBL) pixel layer has a size of approximately 0.001×0.008 in $\Delta\phi \times \Delta\eta$. If there are $n \sim \mathcal{O}(10)$ particles in the $\Delta R < 0.02$ core of a jet, then multiple particles can deposit hits in the same pixel and one of the resulting tracks can be lost. Despite this challenge², tracking inside jets will play an important role in Part II. Methods to measure the reconstruction efficiency in jet cores are presented in Sec. 1.5.




Figure 4.1: An event display in the early Run 2 data showing charged particle tracks reconstructed from the ID. The fitted track trajectory is shown as a continuous colored line and the hits in the three sub-detectors are shown as points. From Ref. [113].

²Between Runs 1 and 2, a new method for resolving ambiguities in the assignment of hits in the pixel detector to tracks significantly improved the track reconstruction in jets [112]. This may help to improve the resolution of the tracks-in-jets based algorithms presented in Part II.

4.3 Jets

There is no unique way to construct a jet: jets are *defined by a jet clustering algorithm*³. For a jet algorithm to be useful experimentally and theoretically it must be *IRC safe*. Let $\vec{\rho} = (\mathbf{y}, \phi)$, where $\mathbf{y}$ is the rapidity⁴. For an algorithm to be IRC safe:

1. Infrared safe (IR): if a particle i is added with $|\mathbf{p}_T| \rightarrow 0$, the jets are unaffected.
2. Collinear safe (C): if a particle i with momentum $\mathbf{p}_i$ is replaced with two particles j and k with momenta $\mathbf{p}_j + \mathbf{p}_k = \mathbf{p}_i$ such that $|\vec{\rho}_i - \vec{\rho}_j| = 0$, then the jets are unaffected.

The most widely used algorithms are categorized by *sequential recombination* [114]. These IRC safe schemes require metrics $\mathbf{d}$ on momenta $\mathbf{d}_{ij} = \mathbf{d}(\mathbf{p}_i, \mathbf{p}_j) : (\mathbf{p}_i, \mathbf{p}_j) \rightarrow \mathbb{R}^+$, $\mathbf{d}_{iB} = \mathbf{d}(\mathbf{p}_i) : \mathbf{p}_i \rightarrow \mathbb{R}^+$ and proceed as follows:

1. Assign each particle as a proto-jet.
2. Repeat until there are no proto-jets left: Let $(\mathbf{k}, \ell) = \text{argmin}_{i,j} \mathbf{d}(\mathbf{p}_i, \mathbf{p}_j)$. If $\mathbf{d}_{mB} < \mathbf{d}_{k\ell}$ for $\mathbf{m} = \text{argmin}_i \mathbf{d}(\mathbf{p}_i)$, then declare proto-jet $\mathbf{m}$ a jet and remove it from the list. Otherwise, combine proto-jets $\mathbf{k}$ and ℓ into a new proto-jet with momentum $\mathbf{p}_{\text{new}} = \mathbf{p}_\ell + \mathbf{p}_k$.

One common widely used set of algorithms use the k_t family of metrics, $\mathbf{d}_{ij}(\mathbf{k}) = \min(\mathbf{p}_{Ti}^{2k}, \mathbf{p}_{Tj}^{2k}) |\vec{\rho}_i - \vec{\rho}_j|^2 / R^2$ and $\mathbf{d}_{iB}(\mathbf{k}) = \mathbf{p}_{Ti}^{2k}$. The parameter R is roughly the size of the jet in $(\mathbf{y}, \phi)$. When $\mathbf{k} = 0$, the clustering procedure is called the Cambridge-Aachen (C/A) algorithm [115, 116] and the distance metric is independent of $\mathbf{p}_T$. By far, the most ubiquitous jet algorithm used at the LHC is the anti- k_t algorithm [117] with $\mathbf{k} = -1$ (the k_t algorithm has $\mathbf{k} = +1$). Figure 4.2 shows an example $Z' \rightarrow t\bar{t}$ simulated event clustered with the k_t , C/A and anti- k_t algorithms. The core of the

³The connection between jet clustering and unsupervised machine learning is discussed in Sec. 4.3.

⁴Then $|\vec{\rho}|^2$ is the same as ΔR^2 for massless particles, but a different notation is used to distinguish them in the massive case. Jet masses are often quite small (see Sec. 4.1), but there can be small differences in the jet $\mathbf{p}_T$ when using $\mathbf{y}$ or η for clustering. As a geometric coordinate η is useful for relating to rigid detector boundaries, but the full rapidity has the desired Lorentz covariance and is thus used for jet clustering (see Sec. 0.1.2).

highest p_T jets is the same for all three algorithms⁵. However, the soft radiation on the outside of the jets varies between the three approaches. One way to visualize the origin of these differences is in the bottom plots of Fig. 4.2 which shows the clustering history. Because of the negative power of p_T , the anti- k_t algorithm clusters higher p_T particles first. The k_t algorithm clusters the softest particles first and the C/A algorithm clusters the closest particles first, independent of p_T . As a result, anti- k_t jets have the most regular catchment area which makes them easier to calibrate. This is quantified with the notion of the *jet area* [120], which is defined by⁶

$$A_J = \lim_{a_g \rightarrow 0} \lim_{p_{T,g} \rightarrow 0} \sum_{g \in G} a_g \mathbb{I}(g \in J), \quad (4.1)$$

where G is a set of *ghost particles* uniformly spread over $|\eta| < \eta_{\max}$ and ϕ , and $\mathbb{I}$ is the indicator function that is 1 when its argument is true and 0 otherwise. The area of a ghost g is $a_g = 4\eta_{\max}\pi/|G|$ (so the limit $a_g \rightarrow 0$ is the same as $|G| \rightarrow \infty$). Each ghost particle has a small but finite $p_{T,g}$; for an IRC safe algorithm, as $p_{T,g} \rightarrow 0$, the ghost particles do not influence the clustering. A ghost particle $g \in J$ when after running jet clustering with the ghost particles, g is clustered in the jet J (which coincides with the jets clustered without ghosts by IRC safety). In practice, the ghost particle transverse momenta are set to a small number and the number of ghosts is fixed, but large enough so that edge effects are negligible. The catchment area of an anti- k_t jet is a circle with area πR^2 , except when two anti- k_t jets are within $|\Delta\vec{p}| < 2R$ in which case the higher p_T one is a circle and the other is a crescent. This is demonstrated by Fig. 4.3. The regular shape of anti- k_t jets makes the their calibration (defined below) more universal, i.e. less dependent on the event topology.

With the ATLAS detector, jets are built from calorimeter-cell clusters. If locally calibrated calorimeter-cell clusters are used for the jet clustering, then the resulting jets are at the LCW-scale and otherwise are at the EM-scale. Jets are calibrated so

⁵This statement can be quantified: in perturbation theory, there is no difference between algorithms in the k_t family at leading order [118, 119].

⁶This is called the *active area*; there are other less used possibilities such as the Voronoi area [120].




Figure 4.2: A simulated $Z' \rightarrow t\bar{t}$ event clustered with anti- k_t (left), C/A (middle) and k_t (right). Particles are colored according to the jet they are clustered in (highest p_T jet is black, then red, then green). The particle size is proportional to $\log(10 \times p_T/\text{GeV})$. The bottom panel shows the clustering history. Each proto jet merger is recorded with the size of each pre-merger proto jet proportional to $\log(10 \times p_T/\text{GeV})$, where p_T is the momentum of the merged proto jet. Mergers are colored from dark (earlier) to light (later). Note that ϕ is 2π periodic so the particles at $-\pi$ are geometrically close to the particles at $+\pi$.

that on average they have the same energy as a *particle-level* jet clustered from all detector-stable particles prior to reaching the detector, excluding muons and weakly interacting particles such as neutrinos [121, 122]. The default clustering scheme is the anti- k_t $R = 0.4$ algorithm. In the first step of the calibration⁷, the average amount of pileup energy is subtracted from each jet using the jet areas method [123, 124]. As a diffuse source of noise, the amount of pileup energy contributing to a jet is

⁷At the end of Run 1, an *origin correction* was added as a first step to improve the angular resolution. See Sec. 2.3.6.1.




Figure 4.3: Jet clustering with two high p_T particles and a grid of ghost particles. Particles clustered in the higher p_T jet are colored red and those in the lower p_T jet are colored blue. All other ghosts are gray. The ghosts marker size is arbitrary and the radius of the two high p_T particle markers $\propto p_T$. When the two particles are far away (left), the jet areas are πR^2 , while if they are within $|\Delta\vec{p}| < 2R$, then one is circular and the other is a crescent. If $|\Delta\vec{p}| < R$ (right), then the jets merge.

proportional to the jet area. The jet-by-jet correction is given by $p_{T,J} \mapsto p_{T,J} - \rho A_J$, where ρ is the median pileup density⁸ $\rho = \text{median}_J(p_{T,J}/A_J)$. Additional corrections based on μ and the number of reconstructed vertices (NPV) remove the residual pileup dependence⁹. The next step is the core calibration, which corrects the energy and η of the jets using numerical inversion (See Appendix C.7) based on simulation. At this stage, the reconstructed jet energy in simulation is inclusively unbiased. At the end of Run 1, an additional MC-based calibration was introduced (*global sequential calibration* [125]) to reduce residual biases depending on the jet flavor and energy leaked beyond the hadronic calorimeter (see Sec. 3.1.2). A variety of object balancing techniques are used to validate the calibration procedure in data and result in a residual correction applied to the data [121, 122, 126, 127].

Basic quality criteria are also used to remove jets from pileup and other sources of noise [121, 124]. A quantitative description of jets reconstruction, including systematic uncertainties, is found in Sec. 6.1. Jets and their internal structure are the main focus of Part II and will therefore be discussed in much more detail in subsequent chapters.

⁸ $R = 0.4$ k_t jets are used to calculate ρ . Hard-scatter jets have little impact on this median.

⁹NPV is only sensitive to in-time pileup whereas μ additionally reflects the out-of-time pileup.

4.4 Electrons, Photons, Muons, and Taus

Lepton and photon reconstruction all rely on inner detector tracks. Electrons are built from single tracks matched to a cluster of electromagnetic calorimeter-cells [128–130]. Photons are either matched to two tracks if a *conversion* happens in the ID¹⁰ or zero tracks if no conversion happens before the calorimeter [128]. Tracking for particles from secondary vertices (such as conversion electrons) extends as far as 80 cm into the ID. In addition to stricter matching requirements between tracks and clusters (including energy/momentum and ΔR), further background rejection is achieved by requiring that the electromagnetic shower and the amount of transition radiation from the TRT be consistent with expectations for electrons and photons.

Muons are constructed from tracks in the ID matched to tracks in the MS [131–133]. Additional muons beyond the ID acceptance are built entirely out of MS tracks. Furthermore, the efficiency for muons is recovered for $|\eta| < 0.1$, where the MS is only partially instrumented due to calorimeter and ID services, by using ID tracks matched to either a calorimeter energy deposit consistent with a minimum ionizing particle or a track segment in the MS.

Hadronically decaying tau leptons are constructed from jets [134,135]. Tau leptons decay hadronically about 2/3 of the time and of those, about 80% have one charged pion (*one-prong*) while about 20% have three charged pions (*three-prong*). A series of calorimeter and tracking observables such as the fraction of EM energy, the width in ΔR , and the jet mass are used to discriminate taus from electrons and jets.

The reconstruction efficiency and energy scale of leptons and photons are calibrated using simulation and corrected based on in-situ studies. Low mass resonances and Z bosons are used for both the energy calibration and the efficiency measurement. The latter uses a *tag-and-probe* method where one object $\mathbf{o}_1$ passes a strict selection (*tag*) and another object $\mathbf{o}_2$ with $m_{\mathbf{o}_1\mathbf{o}_2}$ near the resonance mass is probed to see if it passes the particle identification. A quantitative comparison of the reconstruction efficiencies and resolutions, including systematic uncertainties, can be found in Sec. 6.1. Specific particle identification requirements are specified when used.

¹⁰There is about $0.5X_0$ for $|\eta| < 0.8$ and about $1.5X_0$ for $0.8 < |\eta| < 1.5$.

4.5 Missing Transverse Momentum

Particles that only decay weakly such as neutrinos are not measured directly. However, the sum of the transverse momenta from all such particles can be inferred using conservation of momentum in the transverse plane. In the absence of a detector, $\vec{p}_T^{\text{miss}} = -\sum \vec{p}_{T,i}^{\text{visible}}$ would be equal to $\vec{p}_T^{\text{non-interacting}}$. The reconstructed $\vec{p}_T^{\text{miss}}$ in ATLAS is built from all objects described in this chapter. Each object's unique calibration is used to improve the overall missing momentum resolution. The measured energy that is not assigned to jets¹¹, electrons, photons, etc. is called the *soft-term*. At $\sqrt{s} = 8$ TeV, the baseline soft-term was constructed from calibrated calorimeter-cell clusters not assigned to jets or other objects [136, 137]. For the early Run 2 data, this default has shifted to a track-based soft-term [138, 139]. Information about neutral particles is lost when only using tracks, but the neutral contribution cancels on average because charge-to-neutral fluctuations are symmetric in azimuth. The main motivation for the track-based term is the robustness to pileup. Tracks from collisions other than the primary hard-scatter vertex can be readily identified and removed from the soft-term. By construction, the contribution from the soft-term is subdominant to the contribution from other high p_T objects in events with real sources of missing particles.


The magnitude of the missing transverse momentum¹² is called E_T^{miss} and is a powerful discriminating variable for identifying events with neutrinos, such as the pair production of top quarks, $pp \rightarrow t\bar{t} \rightarrow b\bar{b}W^+W^- \rightarrow b\bar{b}l\nu qq'$ or the production of new particles that do not interact with the detector. The particles targeted by the SUSY search presented in Part III decay via undetectable particles that can have $p_T \sim \mathcal{O}(100)$ GeV and so the E_T^{miss} will be one of the most important observables.

¹¹All topo-clusters belong to a jet, but jets below a threshold of ~ 20 GeV are not calibratable (no correlation between detector-level and particle-level energy). Measured energy is ‘not in a jet’ if it is in one of these low p_T jets.

¹²This quantity uses an ‘E’ because it has historically been called the *missing transverse energy*. This is a misnomer because energy is a scalar, but has been used because it is mostly due to calorimeter energy measurements (as opposed to momentum measurements from tracks).

Part II

The Quantum Properties of Jets



A schematic (oversimplified) diagram illustrating the transmission of quark charge to the quantum properties of jets. The u and $\bar{d}$ quarks are not directly observable, but their electric and color charge have observable consequences for the pattern of hadrons.

The *quantum properties of jets* are the observable consequences of the quantum properties of the initiating quarks and gluons. Quarks are the only elementary particles that are charged under all forces. Due to confinement, these quantum properties are not directly observable and are instead embedded in the radiation pattern within and around a jet. The strong coupling constant is sufficiently small that many aspects of this transmission of quark and gluon charge to inter- and intrajet radiation can be understood in the context of perturbation theory. However, there are important non-perturbative aspects of jet formation. There are additional theoretical challenges due to the rich structure of QCD; for example the gluon carries color charge unlike the photon in QED. Measuring the quantum properties of jets also presents a significant experimental challenge. Differences in radiation patterns between different quark and gluon charges are often subtle and thus require precise measurements of jet constituent energies and locations. Furthermore, there are several sources of diffuse noise such as pileup that complicate both the measurements and their interpretations.

A parton that initiates a jet is uniquely determined by its charges under all symmetry groups of the Standard Model. The most basic property is the parton three-momentum, a charge of the Poincaré group. For most applications involving jets, this is the only relevant quantum property, as the jet is viewed as a noisy proxy of the parton; the internal structure is a nuisance. The jet three-momentum also has the least quantum noise of all quantum properties of jets; the average jet $\mathbf{p}_T$ is the same as the average parton $\mathbf{p}_T$ within 5-10% [119]. A related quantum property is the parton mass. Jets produced by light quarks and gluons can acquire non-negligible mass resulting from relatively hard wide-angle radiation. This mass encodes information about color charge of the initiating quark or gluon and is unrelated to the on-shell quark and gluon mass. On the other hand, jets initiated by the hadronic decays of genuinely heavy particles such as W/Z or Higgs bosons have significant mass that is correlated with the parton mass. The last charge of the Poincaré group is the particle spin. This information is lost for light quark and gluon jets as a result of the hadronization process. However, the angular distribution of subjets within top quark and heavy boson jets does contain some information about the polarization of the initiating parton. The quantum properties of jets related to the Poincaré group

symmetries are studied in Chapter 4.

The other charges of the Standard Model are associated with the internal $U(1) \times SU(2) \times SU(3)$ symmetry. Partons are most likely to fragment into hadrons of the same electric charge. Therefore the electric charge of the hadrons inside a jet encodes information about the parton electric charge. This is complicated by the finite acceptance of the detector for low p_T particles and also the fact that additional charge must flow into quark jets in order to make the net charge an integer. Higher energy hadrons carry more information about the parton charge and low energy hadrons are subject to threshold effects, so a *jet charge* can be constructed by using the energy-weighted charge of hadrons inside a jet as a proxy of the parton charge. Chapter 1 is an extensive study of jet charge, both as a probe of jet formation and for charge tagging. For example, the energy-dependence of the jet charge is studied for evidence of *scale violation*. Additionally, jet charge in boosted boson jets is studied in Chapter 4, where no additional charge needs to flow into the jet due to the color singlet nature of the initial state.

Analogous to the electric charge for the electroweak force is the color charge for the strong force. Partons can either be in the singlet (no charge), triplet, or octet representations of $SU(3)$. Partons in the triplet representation carry one color while partons in the octet representation carry one color and one anti-color. The radiation pattern from hadronic jets resulting from singlet partons tends to be enhanced within the core of the jet relative to jets from octets that are *color-connected* to other partons in the event. This is particularly important to study because of the applications to jet tagging as a boosted $H \rightarrow b\bar{b}$ jet gives rise to a singlet radiation pattern while the background $g \rightarrow b\bar{b}$ process should resemble the octet radiation pattern. Chapter 2 is a study of colorflow in and around boson jets. In addition to differentiating singlet-induced jets from octet-induced jets, it is essential to study the differences between octet-induced jets (gluons) and triplet-induced jets (quarks). Quark and gluon tagging is ubiquitous (if only implicit) at the LHC and despite being well-studied, is still an area of active research theoretically and experimentally. Chapter 3 presents a measurement of constituent multiplicity, an observable that is directly proportional to strength of the quark and gluon color charges C_F and C_A .

One final quantum property is flavor. Quark and gluon flavor are uniquely specified by color charge, but there is a further distinction into the various quark types. Jet charge is sensitive to the up versus down type of the initiating parton, but there are an entire class of observables sensitive to heavy flavor quarks. The use of **b**-tagging to probe quark flavor is studied in the context of boson jet tagging in Chapter 4.

Part II will explore the *substructure* and *superstructure* [140] of high energy jets in order to understand how quark and gluon quantum charges are realized within the observable pattern of hadrons. These quantum properties of jets probe the detailed nature of the strong force as well as provide tools for discovering new particles and forces beyond those described by the Standard Model. Table 1 summarizes quantum properties of jets discussed above as well as which chapters cover them.

Quantum Property	Charge	Observable	Chapter
Electric Charge	$\pm 2/3, \pm 1/3, 0$	Jet Charge	1
Color Charge	1, 8	Jet Pull	2
Color Charge	1, 3	Constituent Multiplicity	3
Mass	$\alpha_s R p_T, m_W, m_Z$	Jet Mass	4
Electric Charge	$\pm 1, 0$	Large-radius Jet Charge	4
Flavor	b/c/light	b -tagging	4

Table 1: The various quantum properties of jets studied in Part II along with the corresponding chapter. Even though some of the quantum properties are due to forces other than the strong force, their realization in jets is due to the quantum evolution from partons to hadrons via QCD.

Chapter 1

Jet Charge

Quarks and gluons produced in high-energy particle collisions hadronize before their electric charge can be directly measured. However, information about the electric charge is embedded in the resulting collimated sprays of hadrons. One jet observable sensitive to the electric charge of quarks and gluons is the momentum-weighted charge sum constructed from charged-particle tracks in a jet [141]. Called the *jet charge*, this observable was first used experimentally in deep inelastic scattering studies [142–148] to establish a relationship between the quark model and hadrons. Since then, jet charge observables have been used in a variety of applications, including tagging the charge of b-quark jets [149–159] and hadronically decaying W bosons [160–165] as well as distinguishing hadronically decaying W bosons from jets produced in generic quantum chromodynamic (QCD) processes [166] and quark jets from gluon jets [164, 167–169].

As will be a reoccurring theme for the quantum properties of jets, the charge information embedded in the radiation pattern of jets is subtle. The left plot of Fig. 1.1 shows the pixelated energy distribution for a simulated event $pp \rightarrow u\bar{u}$. The pixel intensity is the charge-weighted sum of the energy of all hadrons resulting from fragmentation, prior to interactions with the detector. There are clearly two nodes of localized energy deposits (jets), but it is not possible to deduce which jet was initiated by the up quark and which was initiated by the anti-up quark. In contrast, after re-simulating the parton shower and hadronization 10,000 times, the right plot

of Fig. 1.1 clearly shows the upper jet is from the up quark (positive charge) and the bottom jet is from the anti-up quark. In the data, a given hard-scatter event fragments only once. The tools developed in this chapter are intended to apply to individual events, but they are most useful when considering an ensemble of events.




Figure 1.1: Left: the process $pp \rightarrow u\bar{u}$ is simulated with PYTHIA 8 once; Right: the same hard-scatter process with fragmentation re-simulated 10,000 times. Each pixel intensity represents the charge-weighted sum of the energies of all particles produced within the ϕ and η covered by the pixel area.

This chapter presents¹ performance studies related to the detector reconstruction and charge tagging performance of the jet charge as well as a precision measurement of the jet charge moments as a function of jet p_T with the ATLAS detector. The chapter begins in Sec. 1.1 with some background information.

¹The performance studies presented in this chapter are published in Ref. [164] (with technical help from M. Swiatlowski and manuscript help from A. Arce) and the precision measurement is published in Ref. [170] (with help from M. Schwartz on the theory calculation).

1.1 Background

The jet charge is defined in Sec. 1.1.1 and its important properties are reviewed in Sec. 1.1.2. Section 1.1.3 describes the theoretical predictions for the jet charge distribution. The section ends in Sec. 1.1.4 with some comments about charge tagging.

1.1.1 Constructing the jet charge

There is no unique way to define the jet charge. The most naïve construction is to add the charge of all tracks associated with a jet. However, this scheme is very sensitive to lost radiation and diffuse soft radiation that contaminates the jet. Therefore, a weighting scheme is introduced to suppress fluctuations. The matching of tracks with the calorimeter-based jets is performed via the ghost-association technique [120]: the jet clustering process is repeated with the addition of *ghost* versions of measured tracks that have the same direction but infinitesimally small $\mathbf{p}_T$, so that they do not change the properties of the calorimeter jets. A track is associated with a jet if its ghost version is contained in the jet after reclustering. Using such tracks, the jet charge Q_J of a jet J is calculated using a transverse-momentum-weighting scheme [141]:

$$Q_J = \frac{1}{(\mathbf{p}_{T,J})^\kappa} \sum_{i \in \mathbf{Tracks}} q_i \times (\mathbf{p}_{T,i})^\kappa, \quad (1.1)$$

where $\mathbf{Tracks}$ is the set of tracks associated with jet J , q_i is the electric charge of track i in units of the positron charge, $\mathbf{p}_{T,i}$ is transverse momentum of track i , κ is a free regularization parameter, and $\mathbf{p}_{T,J}$ is the transverse momentum of the calorimeter jet. The distributions of Q_J for various jet flavors are shown in Fig. 1.2 for $\kappa = 0.3$. In the simulation, there is a clear relationship between the jet charge and the initiating parton's charge, as up-quark jets tend to have a higher jet charge than gluon jets. Furthermore, gluon jets tend to have a higher jet charge than down-quark jets. However, the jet charge distribution is already broad at particle level and the jet charge response ($Q_{\text{particle-level}} - Q_{\text{detector-level}}$) resolution is comparable to the differences in the means of the distributions for different flavors, so one can expect

only small changes in the inclusive jet charge distribution for changes in the jet flavor composition. The three narrow distributions on top of the bulk response distribution in Fig. 1.2(b) are due to cases in which only one or two charged particles dominate the jet charge calculation at particle level. The two off-center peaks are due to cases in which one of the two high- $\mathbf{p}_T$ -fraction tracks is not reconstructed and the widths of the two off-center and central peaks are due to the (single) track and jet $\mathbf{p}_T$ resolutions. The bulk response is fit to a Gaussian function with standard deviation $\sigma \sim 0.5\ e$ (units of the positron charge).




Figure 1.2: Top left (right): The particle- (detector-)level jet charge distribution for various jet flavors in a sample of jets with $p_T > 500$ GeV for $\kappa = 0.3$. Bottom: the distribution of the jet-by-jet difference between the particle-level and detector-level jet charge distributions. The shaded region is used to fit a Gaussian function to extract the bulk response resolution, which is $\sigma \sim 0.5$ e, where e is the positron charge. See Sec. 1.2.1 for details about the simulation.

1.1.2 Jet Charge Properties

By using the calorimeter jet p_T in the denominator of Eq. 1.1, there is some sensitivity to the charge-to-neutral fraction in the jet, which contains useful information about the parton charge. Alternative definitions using $(\sum_{i \in \text{Tracks}} (p_{T,i}))^\kappa$ which lead to a bounded jet charge are studied in Sec. 1.3.2.2. Sections 1.1.2.1 describes how κ regulates the sensitivity to soft radiation within a jet and Sec. 1.1.2.2 shows how the jet charge transforms under Lorentz boosts.

1.1.2.1 Sensitivity to Soft Radiation

The parameter κ in Eq. 1.1 regulates the sensitivity of the jet charge to soft radiation. Low values of κ enhance the contribution to the jet charge from low- p_T particles while in the $\kappa \rightarrow \infty$ limit, only the highest- p_T track contributes to the sum in Eq. 1.1. The dependence on the highest- p_T tracks is demonstrated using the plots in Fig. 1.3 with the variable $Q_{J,n}$, which is the jet charge in Eq. 1.1, but built from the leading n tracks. The variable $Q_{J,1}$ is simply the weighted fragmentation function of the leading-track p_T to the jet p_T with weight κ . The usual Q_J is recovered in the limit $n \rightarrow \infty$. Figure 1.3 shows the sequence $Q_{J,n}$ for $\kappa = 0.3$ and $\kappa = 0.7$. For lower values of κ , many tracks are required for the sequence of distributions to converge to the full jet charge. However, for $\kappa \gtrsim 0.7$, the distribution converges quickly, indicating that only the highest- p_T tracks are contributing. The peaks in the distributions in Fig. 1.3 are due to the discrete nature of hadron electric charge: if there is only one track, then the peaks are at $\pm \langle p_{\text{track}}^\kappa / p_{\text{jet}}^\kappa \rangle$ while if there are two tracks, then a peak at zero develops from the cases in which the two tracks have opposite charge. If the charge of the tracks are chosen at random, it is twice as likely that the tracks have opposite charge compared with the case that both have positive charge and therefore the peak at zero is taller than the peaks at larger values of $|Q_{J,n}|$.

All reconstructed tracks are henceforth used when computing the jet charge, but the plots in Fig. 1.3 give an indication of the contribution of (relatively) high- and low- p_T tracks. Figure 1.4 shows the joint distribution of jet charges with different κ values. While the distributions are peaked along the diagonal, there is a significant

off-diagonal spread that is bigger when the two κ values are further apart. The stripes in the low p_T bin are due to cases in which there is only one track; in those cases the jet charge for one value of κ is uniquely specified by the jet charge at any other $\kappa \neq 0$ value. The studies presented in this chapter use a range of κ values in order to maintain a broad sensitivity to both hard and soft radiation inside jets. The impact of low p_T tracks on the jet charge reconstruction is revisited in Sec. 1.3.2.




Figure 1.3: The distribution of the jet charge built from the leading n tracks ($Q_{J,n}$) for (left) $\kappa = 0.3$ and (right) $\kappa = 0.7$ or $\kappa = 1.0$ for (top) $500 \text{ GeV} < p_T < 600 \text{ GeV}$ and (bottom) $70 \text{ GeV} < p_T < 100 \text{ GeV}$. In the top (bottom) plots, the mean number of tracks is about 15 (7). See Sec. 1.2.1 for details about the simulation.




Figure 1.4: The joint distribution of the jet charge defined with different values of the weighting factor $\kappa \in \{0.3, 0.5, 0.7\}$. The left plots show the joint distributions of $\kappa = 0.3$ and $\kappa = 0.5$ while the right plots show the joint distributions of $\kappa = 0.3$ and $\kappa = 0.7$. Jets in the top plots have $50 \text{ GeV} < p_T < 100 \text{ GeV}$ while those in the bottom plots have $500 \text{ GeV} < p_T < 600 \text{ GeV}$. See Sec. 1.2.1 for details about the simulation.

1.1.2.2 Lorentz Invariance

The electric charge of a particle is a Lorentz invariant quantity, but the jet charge is not Lorentz invariant. It is even possible (sometimes) to flip the sign of the jet charge by performing a suitable Lorentz transformation. To illustrate the problem, consider a simplified case where $Z \rightarrow e^+e^-$, as illustrated in the left plot of Fig. 1.5. Define the ‘jet charge’ as $Q = \frac{1}{m_Z^\kappa}(\mathbf{p}^\kappa - \mathbf{p}^\kappa) = 0$, where $\mathbf{p} = \mathbf{m}_Z/2$. Now, suppose that the Z has some transverse boost with speed β along the $\pm x$ direction. Then,

$$Q = \frac{1}{(\gamma m_Z)^\kappa} \left((\gamma p(1 \pm \beta))^\kappa - (\gamma p(1 \mp \beta))^\kappa \right). \quad (1.2)$$




Figure 1.5: A schematic diagram to illustrate the impact of Lorentz boosts on the ‘jet’ charge. A Z boson decays at rest at the origin in the lab frame and decays into electrons that are along the x direction (left). For a boost along the $-x$ direction, the e^+ has a higher p_T than the e^- and vice versa for a boost along $+x$.

If $\kappa = 1$, then $Q = \pm\beta$. If $\kappa \ll 1$, then $Q = \pm\kappa\beta$. In either case, one can make Q arbitrarily positive or negative depending on the direction of the boost. Now, slightly more generally consider the decay of a color singlet, like a W boson, which decays into n particles and has speed β in the $\hat{\mathbf{r}}$ direction in a particular frame. Then,

$$Q = \frac{1}{(\gamma m_{\text{boson}})^\kappa} \sum_{i=1}^n q_i \gamma^\kappa (E_i - \beta \vec{P}_i \cdot \hat{\mathbf{r}})^\kappa. \quad (1.3)$$

When $\kappa = 1$,

$$\begin{aligned}
Q &= Q(\text{rest frame}) - \frac{\beta}{(m_{\text{boson}})} \sum_{i=1}^n q_i \vec{p}_i \cdot \hat{r} \\
&= Q_{\text{boson}} - \frac{\beta}{m} \sum_{i=1}^n q_i \hat{p}_i \cdot \hat{r}
\end{aligned} \tag{1.4}$$

For a given event, the second term in Eq. 1.4 will not be zero, unless all the tracks are perpendicular to the boost. For an ensemble of events, the non-closure term (second term in Eq. 1.4) will have average zero since the particles are randomly oriented and so the average jet charge is the boson charge. However, the standard deviation of the non-closure term is not zero and so there is an induced smearing of the boson charge due to the distribution of boosts. This is illustrated with a simulated W boson in Fig. 1.6. Because the W boson is a color singlet, it is possible to uniquely associate final state hadrons with the decay of the W boson (See Sec. 2.2.3.1). The jet charge using all of the W^+ decay products in the lab frame is positive, but the sign changes after a large boost in the $-x$ direction. This change of sign can be understood by the dominance of one high p_T negatively charged hadron, shown in left plot of Fig. 1.6. Figure 1.7 shows how the jet charge sign depends on the value of κ as well as on the direction and magnitude of the boost. In the left plot of Fig. 1.7, $\beta_W = 0.65$. For $\kappa = 0$, $Q = 1$ by construction as all W boson decay products are part of the ‘jet’. Even though the jet charge is positive for a large range of κ values, as $\kappa \rightarrow \infty$ the jet charge is driven negative because the leading charged particle has a negative charge. The right plot of Fig. 1.7 has $\kappa = 1$. The horizontal axis begins at $p_T = 10$ GeV because the jet charge is ∞ as $\beta_W \rightarrow 0$. The blue region in the center, where the jet charge is negative, corresponds to a boost in the $-x$ direction, resembling a configuration as in the right plot of Fig. 1.6.

The jet charge is invariant under longitudinal boosts, which is critical at a hadron collider where there is a large range of p_z values that contain much less information about the scale of jet formation compared with p_T . Even though generic quark and gluon jets do not have a well-defined decay frame, the above example illustrates how the jet charge changes between frames.




Figure 1.6: The decay of a simulated W boson event with PYTHIA 8. The length of the arrow is proportional to the energy of the decay product; red arrow denote negatively charged hadrons, blue arrows mark positively charged hadrons and neutral hadrons and photons are in gray. In the right plot, the W has received a large boost in the $+x$ direction.




Figure 1.7: Left: for fixed $\beta_W = 0.65$, the dependence of the ‘jet’ (= all W boson decay products) charge on κ for the event shown in Fig. 1.6. Right: for $\kappa = 1$, the dependence of the jet charge on the boost direction and magnitude for the same event as in the left plot.

1.1.3 From Parton Charge to Jet Charge

In general, the jet charge is not an infrared and collinear safe observable. Infrared safety is guaranteed for $\kappa > 0$, since the contribution of an arbitrarily soft particle is suppressed by $\mathbf{p}_{T,\text{soft}}^\kappa$. However, the charge-weighting in Eq. 1.1 spoils collinear safety. To see this, suppose that a positively charged particle $\mathbf{P}$ carries momentum fraction $z_{\mathbf{P}} = \mathbf{p}_{T,\mathbf{P}}/\mathbf{p}_{T,\text{jet}}$ and splits into a charged particle $\mathbf{P}_+$ and a neutral particle $\mathbf{P}_0$ whose momenta are collinear. The contribution to the jet charge before the splitting is $z_{\mathbf{P}}^\kappa$, while after the splitting it is $z_{\mathbf{P}_+}^\kappa$. For $z_{\mathbf{P}_+} < z_{\mathbf{P}}$, these two contributions are not identical. As a result of collinear sensitivity, hadronization must be included in any reliable description of the jet charge. This information cannot be described perturbatively within QCD, but the non-perturbative components can be quantified and isolated [171, 172]. For a parton of type $\mathbf{p}$ with energy E , the probability for a hadron of type $\mathbf{h}$ to carry a fraction $[z, z + dz]$ of the parton's momentum is given by the *fragmentation function* $D_{\mathbf{p}}^{\mathbf{h}}(z, E)dz$. The normalization of $D_{\mathbf{p}}^{\mathbf{h}}(z, E)$ is the average number of hadrons $\mathbf{h}$ produced by a jet initiated by a parton of energy E , $\langle n_{\mathbf{p}}^{\mathbf{h}}(E) \rangle$. This can be shown by dividing the interval $[0, 1]$ into N pieces so that the probability for multiple hadrons of the same type (e.g. π^+) to have $z \in [i, i+1]/N$ is small:

$$\begin{aligned}
\langle n_{\mathbf{p}}^{\mathbf{h}}(E) \rangle &= \sum_{i=0}^{N-1} \sum_{k=0}^{\infty} \text{Pr}(k \text{ hadrons of type } \mathbf{h} \text{ with } z \in [i, i+1]/N) \\
&= \sum_{i=0}^{N-1} \text{Pr}(\text{one hadron of type } \mathbf{h} \text{ with } z \in [i, i+1]/N) + \mathcal{O}(1/N^2) \\
&= \sum_{i=0}^{N-1} \frac{1}{N} D_{\mathbf{p}}^{\mathbf{h}}(i/N, E) + \mathcal{O}(1/N^2) \\
&\stackrel{N \rightarrow \infty}{=} \int_0^1 dz D_{\mathbf{p}}^{\mathbf{h}}(z, E).
\end{aligned} \tag{1.5}$$

The average multiplicity will be revisited in Chapter 3. Ignoring non-strong force processes, conservation of energy requires that the first moment of the fragmentation function summed over all hadron species is equal to one: $\sum_{\mathbf{h}} \int_0^1 dz z D_{\mathbf{p}}^{\mathbf{h}}(z, E) = 1$. The

average jet charge follows a related form:

$$\langle Q_p(E, \kappa) \rangle = \sum_h Q_h \int_0^1 dz z^\kappa D_p^h(z, E) \equiv \sum_h Q_h \tilde{D}_p^h(\kappa, E), \quad (1.6)$$

where Q_h is the charge of hadron h and $\tilde{D}(\kappa, E)$ is the *Mellin transform* of D at $\kappa + 1$. One can include perturbative contributions to Eq. 1.6 within the context of Soft Collinear Effective Theory (SCET) [173–176] with the Fragmenting Jet Function $\mathcal{G}_p^h$ [177–179]. The average jet charge is given by [171, 172]:

$$\langle Q_i(E, R, \kappa, \mu) \rangle = \sum_h Q_h \int_0^1 dz z^\kappa \frac{\mathcal{G}_p^h(E, R, z, \mu)}{2(2\pi)^3 J_p(E, R, \mu)}, \quad (1.7)$$

where

$$\mathcal{G}_p^h(E, R, z, \mu) = \sum_{p'} \int_z^1 \frac{dz'}{z'} \mathcal{J}_{pp'}(E, R, z', \mu) D_{p'}^h\left(\frac{z}{z'}, \mu\right). \quad (1.8)$$

The factors $\mathcal{J}_{pp'}$ are defined as $\mathcal{J}_{pp'} = 2(2\pi)^3 \delta(1 - z) \delta_{pp'} + \mathcal{O}(\alpha_s)$. Therefore, $\mathcal{G}_p^h = 2(2\pi)^3 D_p^h + \mathcal{O}(\alpha)$. Similarly, the *jet function* $J_p(E, R, \mu) = 1 + \mathcal{O}(\alpha_s)$ [180] and so Eq. 1.6 and Eq. 1.7 are the same up to $\mathcal{O}(\alpha_s)$ corrections. These corrections are less than 10%, but are also not known precisely due to large uncertainties in the fragmentation functions [172]. The fragmenting jet function is the extension of the inclusive fragmentation function in the context of a jet with finite size. The intuition for Eq. 1.8 is that the parton p radiates the parton p' which then in turn fragments into hadron h . The parton p' has energy fraction $z' > z$ and the hadron has energy fraction z/z' of this energy which is a fraction $z = z' \times z/z'$ of the initial parton's energy. The factor dz'/z' is the phase space for parton p to emit p' and the perturbatively calculable functions $\mathcal{J}_{pp'}$ are related to the QCD splitting functions². Higher moments of the jet charge distribution can be computed in a similar fashion,

²See Sec. 4.1.1 for a discussion of the phase space and the QCD splitting functions.

but in general depend on additional non-perturbative information encoded in the multi-hadron fragmentation functions [171, 172].

The jet charge distribution depends on the jet energy due to two related effects. First, since the jet charge depends on the initiating parton type p , the jet charge distribution varies as the parton distribution functions, $f_p(z, \mu)$, change with energy. Figure 1.8 shows a representative set of leading order QCD Feynman diagrams for $2 \rightarrow 2$ scattering with an up-quark in the initial state. In all cases except for the annihilation diagram, the up quark is also an out-going parton. As discussed earlier, the average gluon jet charge is zero. The average up quark jet charge is positive, since the probability for an up quark to fragment into a positively charged hadron is larger than the probability for an up quark to fragment into a negatively charged hadron. If there were only up quarks and gluons, then the average inclusive jet charge would be proportional to the fraction of up quark jets. As discussed in Sec. 1.1, the fraction of up quarks increases with momentum fraction. The momentum fractions of the two initial partons x_1, x_2 , and the proton and parton center-of-mass energies $\sqrt{s}$ and $\sqrt{\hat{s}}$ are related by $\sqrt{\hat{s}} = \sqrt{x_1 x_2 s}$. In central dijet events, the jet $p_T \sim \sqrt{\hat{s}}/2$. Therefore, the fraction of up quark jets increases with jet p_T . The fraction of down quark jets also increases, but the fraction of up quark jets is expected to be larger. These considerations predict that the average jet charge should increase with jet p_T ³.

A second effect that makes the jet charge distribution p_T -dependent is the energy-dependence of the fragmentation functions. The magnitude of the jet charge for a given parton flavor decreases with p_T because of an increasing contribution from electrically neutral gluons. The fragmentation functions evolve with energy scale in an analogous manner to the evolution for parton density functions. In particular, the DGLAP equation also applies⁴:

³Taking into account the fact that there are twice as many valence up quarks versus down quarks, there is residual enhancement of the up quark PDF relative to the down quark one at high x because the mass of the spectator valence quarks is larger.

⁴Note that in some textbooks and papers, there is a factor of two in this equation which depends on if μ is the energy scale or the virtual mass squared (the factor of two is the Jacobian).




Figure 1.8: A representative set of leading order QCD Feynman diagrams with an up quark in the initial state. Every diagram that has a down quark could be replaced with any other parton. There are also the additional diagrams related to the t -channel ones by crossing symmetry.

$$\mu \frac{\partial}{\partial \mu} D_p^h(z, \mu) = \sum_{p'} \int_z^1 \frac{dz'}{z'} \frac{\alpha_s P_{p' \leftarrow p}(z')}{\pi} D_{p'}^h\left(\frac{z}{z'}, \mu\right) \quad (1.9)$$

The integral of Eq. 1.9 gives the evolution equation for the Mellin moment $\tilde{D}_p^h$:

$$\begin{aligned} \mu \frac{\partial}{\partial \mu} \tilde{D}_p^h(\kappa, \mu) &= \frac{\alpha_s}{\pi} \sum_{p'} \int_0^1 dz z^\kappa \int_z^1 \frac{dz'}{z'} P_{p' \leftarrow p}(z') D_{p'}^h\left(\frac{z}{z'}, \mu\right) \\ &= \frac{\alpha_s}{\pi} \sum_{p'} \int_0^1 \frac{dz'}{z'} P_{p' \leftarrow p}(z') \int_0^{z'} dz z^\kappa D_{p'}^h\left(\frac{z}{z'}, \mu\right) \\ &\stackrel{x=z/z'}{=} \frac{\alpha_s}{\pi} \sum_{p'} \int_0^1 dz' (z')^\kappa P_{p' \leftarrow p}(z') \int_0^1 dx x^\kappa D_{p'}^h(x, \mu) \\ &= \frac{\alpha_s}{\pi} \sum_{p'} \tilde{P}_{p' \leftarrow p}(\kappa) \tilde{D}_i^h(\kappa, \mu). \end{aligned} \quad (1.10)$$

For the average jet charge, the transition $g \leftarrow q$ is irrelevant for quark jets because the net charge from a fragmenting gluon is zero. The average gluon jet charge is zero

by symmetry. This means that the only relevant term in Eq. 1.10 is for $\mathbf{p} = \mathbf{q}$ and $\mathbf{p}' = \mathbf{q}$. At *leading power* (the approximation of narrow jets), the ratio of the jet mass to the jet energy m/E is small and so $p_T \propto E$ at fixed η . Therefore Eq. 1.10 and 1.6 can be used to compute the p_T dependence of the average jet charge for a particular quark type jet:

$$\begin{aligned} \frac{p_T}{\langle Q_q(\kappa) \rangle} \frac{d\langle Q_q(\kappa) \rangle}{dp_T} &= \frac{1}{\sum_h \tilde{D}_i^h(\kappa, p_T)} \sum_h p_T \frac{d}{dp_T} \tilde{D}_i^h(\kappa, p_T) \\ &= \frac{\alpha_s}{\pi} \tilde{P}_{q \leftarrow q}(\kappa) \end{aligned} \quad (1.11)$$

The righthand side of Eq. 1.11 can be computed numerically:

$$\begin{aligned} \frac{\alpha_s}{\pi} \tilde{P}_{q \leftarrow q}(\kappa) &= \frac{\alpha_s C_F}{\pi} \int_0^1 dz z^\kappa \left[\frac{1+z^2}{1-z} \right]_+ \\ &= \frac{\alpha_s C_F}{\pi} \int_0^1 dz (z^\kappa - 1) \frac{1+z^2}{1-z} \approx \begin{cases} -0.024 \pm 0.004 & \kappa = 0.3 \\ -0.038 \pm 0.006 & \kappa = 0.5, \\ -0.049 \pm 0.008 & \kappa = 0.7 \end{cases} \end{aligned} \quad (1.12)$$

where the last form are numerical approximations varying the scale of α_s between 50 and 500 GeV, with the average giving the central value ($\alpha_s(50 \text{ GeV}) = 0.130$ and $\alpha_s(500 \text{ GeV}) = 0.094$). The solution is of the form $\langle Q_q(\kappa) \rangle \propto p_T^{c(\kappa)}$, where $c(\kappa)$ is the factor computed in Eq. 1.12.

Figure 1.9 shows the relative size of the two sources of p_T dependence. Assuming $\langle Q_g \rangle = 0, \langle Q_q \rangle = -\langle Q_{\bar{q}} \rangle, \langle Q_u \rangle = -2\langle Q_d \rangle$, the only free parameter is $\langle Q_u \rangle$, which is removed by normalizing the jet charge at a fixed $p_T = 75 \text{ GeV}$. By construction, the relative jet charge is 1 in the first p_T bin. The relative change in the average jet charge is a factor of 10 due to PDFs with just over a factor of 10 increase in the jet p_T . In contrast, the additional impact of p_T -dependent fragmentation functions results in a $\sim 10\%$ change in the average jet charge.




Figure 1.9: The predicted p_T dependence of the jet charge with input from the CT10 PDF set and assuming $\langle Q_g \rangle = 0$, $\langle Q_q \rangle = -\langle Q_{\bar{q}} \rangle$, $\langle Q_u \rangle = -2\langle Q_d \rangle$. For the red, green, and blue lines, the impact of a p_T -dependent fragmentation function is added on top of the PDF-dependence. This is the more forward of the two jets in dijet events (see Sec. 1.2) and therefore more likely to be initiated from a quark. The fragmentation functions (FF) do not depend on κ , but the energy-dependence of their κ -moments do depend on κ .

1.1.4 Charge Tagging

Section 1.1.3 showed that there are several interesting theoretical aspects of the jet charge that make non-trivial predictions for the p_T -dependence. This section discusses a practical aspect of studying the jet charge: charge tagging. In the one-jet-one-parton paradigm, it is often necessary to resolve ambiguities in the matching between partons and jets that could be solved with an additional handle based on the electric charge information. As an example, consider ambiguity solving in $t\bar{t}$ events in the $t\bar{t} \rightarrow bW(\rightarrow l\nu)bW(\rightarrow qq')$ channel. Such events can be isolated with high purity due to the leptonically decaying W boson. However, there are many applications where one needs to directly identify the selected jets with the top decay products. One example is the measurement presented in Chapter 2. Figure 1.10 schematically illustrates the setup: jets need to be assigned to partons in the top decay topology. The W boson and top quark masses provide powerful constraints on the jet momenta, but the jet charge could provide additional information. In particular, the jet charge could help resolve the matching of the b -tagged jets with the b or $\bar{b}$ quark⁵.

In addition to jet charge, b -quark jets offer additional handles to identify the parton charge from the semi-leptonic decays of B and D hadrons. These leptons are often too soft to measure⁶ and even though the lepton charge is highly correlated with the quark charge, there is some contamination from light hadron decays. Denote the lepton charge as Q_L . In the first paper on jet charge, Field and Feynman [141] describe two criteria for evaluating a charge tagging algorithm:

- Efficiency (E): The percentage of jets to which the algorithm can be applied.
- Reliability (R): Given that the algorithm is applicable, the probability that the assignment is correct.

For example $R[Q_L]$ is expected to be close to one, but since not all jets have an identified soft lepton ($Q_L \neq 0$), $E[Q_L] < 1$. Table 1.1 shows the values of reliability

⁵This idea has now been implemented as a dedicated tagger - see Ref. [181] for details. Vertex charge tagging is also available, when secondary and tertiary charged hadron decay vertices are reconstructed.

⁶The threshold used here is $p_T > 4$ GeV. This requirement comes from the soft-lepton b -tagging algorithm used in ATLAS [75], where there is about a 50% chance of identifying a soft-lepton.




Figure 1.10: A schematic of the $t\bar{t}$ decay topology in the semi-leptonic channel. The clouds depict the fragmentation process that ultimately lead to observable tracks (dotted lines) and hadronic jets (cones). The final state is characterized by two **b**-tagged jets, two hadronic jets from the W decay, an isolated lepton, and missing momentum from the undetectable neutrino. In current kinematic fits, only the momentum of the jets are used and not the charge properties of the tracks.

and efficiency for several variations of the jet charge using particle-level simulation with PYTHIA 8. The algorithm with the best reliability is the lepton charge, but the efficiency to have two reconstructed semileptonic B or D decays is low. On the other hand, the jet charge performs well (reliability 66%) and applies to every jet (efficiency 100%). There is not a strong dependence on κ for $\kappa \sim 0.5$, but this will be revisited in Sec. 1.3.2.1 with the full ATLAS detector simulation.

The jet charge reliability is comparable to simple kinematic techniques. For example, when the initial top quarks are produced with some initial momentum, the resulting **b** quark and W bosons will tend to be closer in ΔR than to the anti-top decay products. Particle-level simulation predicts that a ΔR -based scheme has a similar reliability to the jet charge and is rather uncorrelated; combining the two results in a $\sim 10\%$ increase in reliability. In addition to **b**-jet charge identification, the jet charge can aid in the assignment of jets to the hadronic W boson decay. For example, by

Algorithm	R	E	$E \times R$
$Q_L^{\bar{b}} > Q_L^b$ and $Q_L^{\bar{b}}, Q_L^b \neq 0$	77%	3.7%	2.8%
$Q_L^{\bar{b}} > Q_L^b$	42%	66%	28%
$Q_J^{\bar{b}}(\kappa = 1) > Q_J^b(\kappa = 1)$	62%	100%	62%
$Q_J^{\bar{b}}(\kappa = 0.3) > Q_J^b(\kappa = 0.3)$	62%	100%	62%

Table 1.1: For each algorithm, the probability that the b - and $\bar{b}$ -jet assignment is correct (reliability R) and the fraction of events to which the algorithm can be applied (efficiency E). Note that the value of R in the second row is less than 50% of the time, due to cases in which neither jet has an associated lepton ($Q_L = 0$).

requiring the dijet charge to be opposite the charge of the lepton, the jet selection based only on the invariant mass of the two jets can be improved by $\sim 15\%$ ⁷. Charge tagging for hadronic W boson decays will be revisited in Sec. 1.3.3.

Topological assignments in top quark pair production is only one example where charge tagging could improve the performance of existing methods. Other examples include high b -quark multiplicity final states (e.g. $t\bar{t}H, \tilde{g} \rightarrow t\bar{t}\tilde{\chi}_1^0$, and $T' \rightarrow Ht$) and quark versus gluon tagging. It is therefore important to study the jet charge performance in order to validate and improve the inputs to jet charge-based tagging techniques. Charge tagging will be revisited in Sec. 1.3.2.1 for small-radius jets and in Sec. 1.3.3.3 for charge tagging large-radius jets.

⁷There is no unique way to declare an assignment *correct*, but the results stated here are nearly the same when using a ΔR matching between the W boson and the two jets and an energy fraction method, described in Sec. 2.2.3.1. The efficiency of this requirement is 60%.

1.2 Analysis Design

The main purpose of this chapter is to present a precision measurement of the jet p_T -dependence of the jet charge distribution. As part of this analysis, the jet charge reconstruction is studied in order to improve the measurement as well as the understanding of charge tagging. The jet charge distribution is measured in inclusive dijet events from pp collisions at $\sqrt{s} = 8$ TeV. Inclusive dijet events provide a useful environment for measuring the jet charge as they are an abundant source of gluon-initiated and quark-initiated jets. There are fewer theoretical ambiguities associated with assigning the jet flavor in events with two jets than in events with higher jet multiplicities. Furthermore, the transverse momentum (p_T) range accessible in dijet events is broad, $\mathcal{O}(10)$ GeV up to $\mathcal{O}(1000)$ GeV. As discussed in Sec. 1.1.3, the jet charge distribution is expected to change significantly over this kinematic range due to changes in the PDF. The PDFs are fairly well constrained [182–186] in the momentum fraction range relevant for this study, 0.005–0.5. However, because the jet charge is directly sensitive to the parton flavor, its p_T dependence can provide a consistency check using new information beyond the jet p_T , which is currently used in PDF fits. The PDFs are not the only nonperturbative input needed to model the jet charge distribution and its evolution with $\sqrt{\hat{s}}$. As a momentum-weighted sum over jet constituents, the jet charge is sensitive to the modeling of fragmentation. Previous studies have shown that there are qualitative differences between the charged-particle track multiplicities of jets in data and as predicted by the leading models of hadron production [167]. Thus, a measurement of the jet charge distribution with a range of quark/gluon compositions can provide a constraint on models of jet formation. Furthermore, the high energy dataset can be used to probe the sub-leading p_T dependence of the jet charge due to the $\sqrt{\hat{s}}$ -dependence of the fragmentation functions. This requires new techniques for extracting the jet charge for individual jet flavors. The average jet charge is extracted for both the leading and subleading jet and they are distinguished based on their relative orientation in rapidity. The more forward of the two jets has a larger energy and is associated with the incoming parton that had a higher momentum fraction of the proton. As this parton is more likely than

the lower momentum fraction parton to be a(n up) quark, the difference in the average jet charge between the more forward and more central jets provides a way to extract the jet charge per jet flavor. Figure 1.11(a) shows the flavor fraction for the more forward and more central particle-level jets that are well-balanced in p_T (see Sec. 1.2.2). The fraction of gluon jets decreases with p_T for both the more forward and the more central jet, but the quark jet purity is higher for the more forward jet. The p_T evolution of the sum of the flavor fractions weighted by the sign of the parton charge is shown in Fig. 1.11(b).




Figure 1.11: For a given jet flavor, (a) shows the fraction f of jets with that flavor in events passing the particle-level event selection and (b) shows the p_T evolution of the flavor fractions weighted by charge-sign: $f_{\text{up}} + f_{\text{anti-down}} - f_{\text{anti-up}} - f_{\text{down}}$. The CT10 PDF set is combined with matrix elements from PYTHIA 8. The forward-central differences between the flavor fractions are largest at low p_T , but the highest quark-jet purity occurs at high jet p_T . The markers for the more forward and central jets are distinguished by their blue and red colors, respectively.

Further details about the analysis setup are described in subsequent sections. The dataset and simulated samples are detailed in Sec. 1.2.1 and the object reconstruction and event selection are in Sec. 1.2.2.

1.2.1 Data and simulated samples

This measurement uses the full dataset of pp collisions recorded by the ATLAS detector in 2012, corresponding to an integrated luminosity of 20.3 fb^{-1} at a center-of-mass energy of $\sqrt{s} = 8 \text{ TeV}$. Events are only considered if they are collected during stable beam conditions and satisfy all data-quality requirements [187]. To reject noncollision events, there must be a primary vertex reconstructed from at least two tracks each with $p_T > 400 \text{ MeV}$ [188]. Due to the high instantaneous luminosity and the large total inelastic proton-proton cross section, on average there are about 21 simultaneous (*pileup*) collisions in each bunch crossing.

1.2.1.1 Jet Triggers

A set of single-jet triggers is used to collect dijet events with high efficiency. Due to the large rate for jet production at the LHC and the limited bandwidth, these triggers are *pre-scaled*. For a given trigger T , the prescale $\frac{1}{p(T)} = \text{Pr}(\text{save event}|\text{pass } T)$. A trigger is not prescaled if $p = 1$. The values p are chosen a priori; the collected data are *un-prescaled* by weighting an event by $p(T)$ if the highest p_T trigger that the event passes is T . A standard method for measuring the trigger efficiency is to use a reference trigger that is fully efficient well below the region of interest and then compute the fraction of events passing the reference trigger that also pass the probe trigger. The challenge with this method is that by construction the reference trigger will have a larger pre-scale than the probe trigger, and thus a smaller sample size in data. Another possibility is to *emulate* the trigger offline on all collected events. An event is said to pass the emulated trigger T if the corresponding trigger jet objects all pass the corresponding L1, L2, and Event Filter thresholds. Figure 1.12 shows the trigger efficiency as a function of the offline jet p_T threshold using the emulation method. There are some clear differences between data and simulation in the turn-on region of the trigger, but all offline jet thresholds are chosen to avoid this region. Table 1.2 shows the collected luminosity for each trigger as well as the offline jet p_T ranges used, chosen such that the trigger is fully efficient. The highest- p_T trigger is not prescaled. The prescale factor is the ratio of total luminosity to the collected




Figure 1.12: Trigger Efficiencies for the various single jet triggers used in this analysis. The numbers in the legend show the point at which the trigger is fully efficient.

luminosity for a given trigger.

Monte Carlo (MC) simulated events are generated in p_T slices in order to ensure a large number of events over a broad range of reconstructed jet p_T , given constraints on the available computing resources. The p_T slices span the interval 0 to 5 TeV in ranges that approximately double with each increasing slice, starting with a range of size 8 GeV and ending with a range of size 2240 GeV. The baseline sample used for the measurement is generated with PYTHIA 8.175 [189] with the AU2 [190] set of tuned parameters (tune) and the next-to-leading-order (NLO) PDF set⁸ CT10 [193, 194]. Another large sample of events is generated with HERWIG++ 2.63 [195, 196] with tune EE3 [197] and leading-order (LO) PDF set CTEQ6L1 [198] (particle-level samples with CT10 and EE4 are also used for comparisons). Both PYTHIA and HERWIG++ are LO in perturbative QCD for the $(2 \rightarrow 2)$ matrix element and resum the leading

⁸A discussion on the use of NLO PDF sets with LO matrix elements is given in Refs. [191, 192].

Trigger threshold [GeV]	Offline Selection [GeV]	Luminosity [fb^{-1}]
25	[50,100]	7.84×10^{-5}
55	[100, 136]	4.42×10^{-4}
80	[136, 190]	2.32×10^{-3}
110	[190, 200]	9.81×10^{-3}
145	[200, 225]	3.63×10^{-2}
180	[225, 250]	7.88×10^{-2}
220	[250, 300]	2.61×10^{-1}
280	[300, 400]	1.16
360	≥ 400	20.3

Table 1.2: The single-jet trigger menu used to collect dijet events with the 2012 dataset. The first column is the level-three (Event Filter) jet $\mathbf{p}_T$ threshold and the second column is the offline leading-jet $\mathbf{p}_T$ range corresponding to the given trigger. The luminosity collected with each trigger is in the last column. The total 2012 dataset was 20.3 fb^{-1} ; the highest- $\mathbf{p}_T$ trigger is not prescaled.

logarithms (LL) in the parton shower. However, the ordering of emissions in the MC resummation in the shower differs between these two generators: PYTHIA implements $\mathbf{p}_T$ -ordered showers [98] whereas HERWIG++ uses angular ordering [99]. The phenomenological modeling of the non-perturbative physics also differs between PYTHIA and HERWIG++. In addition to different underlying-event models (Ref. [103] for PYTHIA and an eikonal model [104] for HERWIG++) the hadronization models differ between PYTHIA (Lund string model [101]) and HERWIG++ (cluster model [102]). These two schemes are known [167] to predict different numbers of charged particles within jets and different distributions of the charged-particle energies within jets, both of which are important for the jet charge. All tunes of the underlying event that are used with PYTHIA and HERWIG++ in this analysis use LHC data as input. As discussed in Sec. 1.2, the corrected data are compared to models with various PDF sets; for consistency, each set has a dedicated underlying-event tune constructed in the same way from a fixed set of data inputs (AU2) described in detail in Ref. [190]. The PDF sets include LO sets CTEQ6L1 [198] and MSTW08LO [182] as well as

NLO sets CT10 [193, 194], NNPDF21 NLO [199], and MSTW2008NLO [182]. A sample generated with a NLO matrix element from POWHEG-BOX R2262 [89, 200–202] (henceforth referred to as POWHEG) with PDF set CT10 interfaced with PYTHIA 8.175 and the AU2 tune is also used for comparisons.

Pileup is simulated by overlaying minimum bias events generated with PYTHIA 8 on top of the hard scatter. The distribution is re-weighted to match the data as shown in Fig. 1.13. All MC samples are processed using the full ATLAS detector simulation [110] based on GEANT4 [108].




Figure 1.13: The average number of interactions per bunch crossing in data and in Pythia (after re-weighting) after the event selection described in Sec. 1.2.2.

1.2.2 Object reconstruction and event selection

The reconstructed objects used for the jet charge as well as for the event selection are described in Sec. 1.2.2.1. The fiducial definition of the measurement, unfolded to particle level, is given in Sec. 1.2.2.2.

1.2.2.1 Object reconstruction at detector level

Jets are clustered using the anti- k_t jet algorithm [117] with radius parameter $R = 0.4$ implemented in FastJet [203] from topological calorimeter-cell clusters [204], calibrated using the local cluster weighting (LCW) algorithm [205, 206]. An overall jet energy calibration accounts for residual detector effects as well as contributions from pileup [207] in order to make the reconstructed jet energy an unbiased measurement of the particle-level jet energy. Jets are required to be central ($|\eta| < 2.1$) so that their charged particles are within the $|\eta| < 2.5$ coverage of the ID.

When more than one primary vertex is reconstructed, the one with the highest $\sum p_T^2$ of tracks is selected as the hard-scatter vertex. Events are further required to have at least two jets with $p_T > 50$ GeV and only the leading two jets are considered for the jet charge measurement. To select dijet topologies, the two leading jets must have $p_T^{\text{lead}}/p_T^{\text{sublead}} < 1.5$, where p_T^{lead} and p_T^{sublead} are the transverse momenta of the jets with the highest and second-highest p_T , respectively. The jet with the smaller (larger) absolute pseudorapidity $|\eta|$ is classified as the more central (more forward) jet. A measurement of the more forward and more central jet charge distributions can exploit the rapidity-dependence of the jet flavor to extract information about the jet charge for a particular flavor. This is discussed in more detail in Sec. 1.2.2.2.

Tracks used to calculate the jet charge are required to have $p_T \geq 500$ MeV, $|\eta| < 2.5$, and a χ^2 per degree of freedom (resulting from the track fit) less than 3.0. Additional quality criteria are applied to select tracks originating from the collision vertex and reject fake tracks reconstructed from random hits in the detector. In particular, tracks must be well-matched to the hard-scatter vertex with $|z_0 \sin(\theta)| < 1.5$ mm and $|d_0| < 1$ mm, where z_0 and d_0 are calculated with respect to the primary vertex. Tracks must furthermore have at least one hit in the pixel detector and at least six hits in the SCT. The distribution of the number of tracks in jets in two representative jet p_T ranges is shown in Fig. 1.14. The number of tracks increases with jet p_T and the data fall between the predicted distributions of PYTHIA and HERWIG++.




Figure 1.14: The distribution of the number of tracks associated with a jet in two example jet p_T ranges.

1.2.2.2 Object definitions at particle level

The measurement is carried out within a fiducial volume matching the experimental selection to avoid extrapolation into unmeasured kinematic regions that have additional model-dependence and related uncertainties. Particle-level (prior to a real or simulated detector) definitions corresponding to the reconstructed objects are chosen to be as close as possible to those described in Sec. 1.2.2.1. Particle-level jets are clustered from generated stable particles with a mean lifetime $\tau > 30$ ps, excluding muons and neutrinos⁹. As with the detector-level jets, particle-level jets are clustered with the anti- k_t $R = 0.4$ algorithm. In analogy to the ghost-association of tracks to jets performed at detector level, any charged particle clustered in a particle-level jet is considered for the jet charge calculation¹⁰. There must be at least two jets

⁹Only particles prior to the detector simulation are used in the unfolding. For example, tracks from photon conversions ($\gamma \rightarrow e^+e^-$) in the inner detector or $K_s \rightarrow \pi^+\pi^-$ decays may be reconstructed as detector-level tracks, but excluded as particle-level tracks.

¹⁰There is no $p_T > 500$ MeV threshold applied to charged particles. The impact of applying such a threshold is negligible for all p_T bins except the first two where effects of up to 1% are observed in the mean and standard deviation of the jet charge. See Fig. 1.16.

with $|\eta| < 2.1$ and $p_T > 50$ GeV. The two highest- p_T jets must satisfy the same p_T -balance requirement between the leading and subleading jet as at detector level ($p_T^{\text{lead}}/p_T^{\text{sublead}} < 1.5$). Due to the high-energy and well-separated nature of the selected jets, the hard-scatter quarks and gluons can be cleanly matched to the outgoing jets. While it is possible to classify jets as quark- or gluon-initiated beyond leading order in $m_{\text{jet}}/E_{\text{jet}}$ [208], the classification is algorithm-dependent and unnecessary for the present considerations (in part because of the large experimental uncertainty). In this analysis, the flavor of a jet is defined as that of the highest energy parton in simulation within a $\Delta R < 0.4$ cone around the particle-jet axis. The jet flavor depends on rapidity and so the two selected jets are classified as either more forward or more central; the more forward jet tends to be correlated to the higher- x parton and is less likely to be a gluon jet. Another benefit of this pseudorapidity-based scheme is that the particle-level and detector-level jets are more often the same objects. This is because by conservation of momentum, the transverse momentum of the two jets in dijet events is similar. Figure 1.15 quantifies this effect; the fraction of events where the selected jets are swapped is half as large under the rapidity scheme compared with the momentum scheme.




Figure 1.15: The distribution of the distance ΔR between the selected detector-level jet and the selected particle-level jet using a momentum scheme (left) and a rapidity scheme (right).




Figure 1.16: The impact of adding a particle-level jet p_T threshold of 500 MeV.

1.3 Performance Studies

For both charge tagging and precision measurements with jet charge, a detailed understanding of the impact of the ATLAS detector on the jet charge reconstruction is critical for improving performance. Define the *jet charge response* as the difference between detector-level jet charge and the particle-level jet charge from jets prior to detector-simulation. The figures of merit used in this section are the mean and standard deviation of the jet charge response as well as the tradeoff between positive parton type jet efficiency and negative parton type jet efficiency for a given jet charge threshold (charge tagging performance). The jet charge response is defined as a difference and not a ratio because the jet charge can be close to zero compared with its resolution and thus the ratio with the particle-level jet charge can naturally be large compared to one.

Two complementary samples are used to study the response and the charge tagging performance in a variety of settings. One selection identifies $t\bar{t}$ events to obtain a high-purity sample of hadronically decaying W bosons. W boson decays are particularly clean because the color singlet W boson is hadronically isolated from the rest of the event. Additionally, $t\bar{t}$ events in the one lepton final state offer a unique opportunity to use a tag-and-probe technique to study the charge tagging capabilities of the jet charge in-situ. A second selection targets generic quark and gluon jets in order to probe high p_T jets and allow for a simulation study of the tagging capabilities of individual quark and gluon jets.

1.3.1 Comparisons Between Data and Simulation

This section contains various comparisons between the reconstructed MC and the data, using the event selection described in Sec. 1.2.2.1 that targets generic quark and gluon jets. Figure 1.17 shows the jet p_T spectrum for the more forward and the more central of the two leading jets in dijet events. Over nearly two orders of magnitude in jet p_T , the distribution of events drops by nearly ten orders of magnitude. The overall shape is well-described by the leading order MC, though there is a small trend at low p_T in the ratio between data and simulation. Qualitatively, the left and right

plots of Fig. 1.17 are similar - this is quantified by the ratio between the more forward and more central jets in Fig. 1.18. As expected, the distribution is peaked at one and is nearly symmetric about the peak (cutoff at 0.5 and 1.5 due to the p_T symmetry requirement). The difference in η and ϕ between the more forward and more central jet are shown in Fig. 1.19. The two jets are nearly back-to-back in the transverse plane and are on average close in η . Figure 1.20 shows the η distribution separately for the more forward and the more central jets. Even though the average $\Delta\eta$ is zero, most of the more central jets are within $|\eta| < 1$ and most of the more forward jets have $|\eta| > 1$. There is no explicit isolation requirement, but the requirement for nearly p_T balanced jets indirectly leads to the two leading jets to be relatively isolated. The left plot of Fig. 1.21 shows the distance in ΔR to the nearest jet with $p_T > 25$ GeV, excluding the other selected jet. A significant fraction of events have only the two selected jets with $p_T > 25$ GeV, which accounts for the spike in the overflow bin. The p_T of the closest jet is shown in the right plot of Fig. 1.21, excluding the other selected jet. For close-by jets, the p_T spectrum is steeply falling away from 25 GeV. There is no significant evidence for an impact of these close-by jets on the jet charge distribution. This is demonstrated by Fig. 1.22, which shows the average jet charge and the standard deviation of the jet charge as a function of the p_T of the close-by jet. Within the statistical uncertainties of the data and simulation, the jet charge distribution is independent of the p_T of the close-by jet.

The actual jet charge distribution is shown in Fig. 1.23 for low and high jet p_T and for $\kappa = 0.3$ and $\kappa = 0.7$. Across jet p_T , the jet charge distribution is roughly symmetric around zero, with a small shift to positive values in the higher p_T bin. The distribution is wider for $\kappa = 0.3$ than for $\kappa = 0.7$. To see this, note that $\partial_\kappa x^\kappa = x^\kappa \log(x) < 0$ for $0 < x < 1$. Each term in the defining sum for the jet charge has the form $x = p_{T,\text{track}}/p_{T,\text{jet}}$ and therefore by decreasing κ , the absolute value of each contribution to the sum increases. This effects the mean in addition to the width of the jet charge distribution, as seen by the jet p_T -dependence of the dijet charge in Fig. 1.24. As expected from Sec. 1.1.3, the average jet charge increases with the energy scale. There are also qualitative systematic differences between data and simulation. These observations are revisited in more detail in Sec. 1.4.




Figure 1.17: Reconstructed jet p_T spectrum for the more forward jet (left) and the more central jet (right). Note that the pre-scales are applied to the data to arrive at a smooth and steeply falling distribution of the jet p_T .




Figure 1.18: The ratio of the more forward to the more central jet p_T distributions.




Figure 1.19: The differences in η (left) and ϕ (right) between the more forward and more central jet.




Figure 1.20: The η of the more forward (left) and more central (right) jet.




Figure 1.21: The distance in ΔR to the nearest jet with $p_T > 25$ GeV and the p_T spectrum of this jet (right).




Figure 1.22: The dependance of the average (left) and standard deviation (right) of the more forward jet charge distribution on the p_T of the nearest jet above 25 GeV. Events are only plotted if the ΔR to the nearest such jet is < 0.8 . Uncertainties are statistical only.




Figure 1.23: The top (bottom) row shows the sum of the charges of the two leading jets in dijet events for two bins of the leading jet p_T . Two values of the p_T weighting factor are shown: $\kappa = 1.0$ on the left and $\kappa = 0.3$ on the right. The lower panels show the ratios between data and MC. The gray band in the ratio includes jet p_T and track reconstruction efficiency uncertainties, described in Sec. 1.5.




Figure 1.24: The dependence of the dijet charge on the dijet mass for two different values of κ in data and MC for a dijet sample. The uncertainty band includes preliminary jet p_T and track isolation uncertainties.

1.3.2 Modeling and Tagging Performance with QCD Jets

1.3.2.1 Single Jet Charge

Using the parton-matching scheme described earlier, the truth charge distributions are separated according to the jet flavour in Fig. 1.25. Figure 1.25 is similar to Fig. 1.2, but for multiple p_T bins and multiple κ values. As observed with the earlier figure, in the simulation there is significant discrimination between the charge sign of the quark at the generator level, but not between flavour types with the same charge. The rejection of jets with a flavour corresponding to a negative charged parton as a function of the efficiency for jets with a flavour corresponding to a positively charged parton is shown in the left plot of Fig. 1.26. For an positive quark jet efficiency of about 50%, there is a rejection(= 1/negative quark jet efficiency) of about 6, independent of κ . The discrimination between quark and gluon jets is quantified in the right plot of Fig. 1.26. A rejection of about 3 against gluon jets is expected for an efficiency of 50% for quark jets. While not competitive with dedicated quark/gluon taggers [209] on its own, the jet charge could be used as an additional discriminating variable within a multivariate approach.

A more extensive scan in κ for quark charge tagging performance is shown in Fig. 1.27. A value $\kappa \sim 0.5$ is optimal for the chosen p_T bin; Fig. 1.28 shows that the optimality of this value is nearly independent of p_T . For $p_T \lesssim 500$ GeV, the charge tagging performance is also relatively independent of p_T . For $p_T \gtrsim 500$ GeV, the performance begins to degrade as the jet charge resolution significantly worsens as discussed below.

The discrimination is slightly degraded for heavy-flavor jets. The jet charge distributions for positive and negative charm and bottom quarks is shown in Fig. 1.29 where the inclusive samples used for comparison are the down type quark jets for the b-quark jets and the up type quark jets for c-quark jet charge. Both plots of Fig. 1.29 show that the ratio of heavy-quark jet charge to the inclusive jet charge of the same charge type is low for positive flavor in the positive region (thus high in the negative tail) and vice versa. This means that the heavy-flavor distributions are shifted towards the center and thus the separation between positive and negative charge is

reduced. This shift is quantified by noting that the difference between the means of the two bottom-flavor distributions is 0.35 ± 0.02 (statistical uncertainty only) while the difference for the inclusive sample is 0.42 ± 0.01 . Likewise for charm-flavor jets, the difference in means is 0.40 ± 0.02 while for the inclusive sample the difference is 0.58 ± 0.01 . This effect cannot be due to the selection on the track vertices, as the d_0 requirement is much larger than the decay length of heavy-flavor mesons. It might be due to differences in the fragmentation, although further investigations are needed to draw firmer conclusions. However since the effect is relatively small, the degradation in separation is expected to be small; this may be important for $W^\pm$ discrimination as one of the decay products is a charm quark about 50% of the time. The flavor dependence of the jet charge is re-investigated in Sec. 4.2.3 in the context of boosted W and Z boson jets.

The MC can be used to investigate the performance of jet charge reconstruction. As above, jet charge response is defined as the difference between the (MC) reconstructed jet charge and the truth jet charge. The mean of the jet charge response is shown as a function of jet p_T in the left plot of Fig. 1.30 and as a function of the number of charged tracks (n_{track}) within the jet in the left plot of Fig. 1.31. The response is nearly independent of both the p_T and the number of tracks. There is a slight decreasing trend in the response with p_T . A contributing factor to the trend is the loss of tracks in the core of high p_T jets so that $|Q^{\text{reco}}| < |Q^{\text{true}}|$. Since the fraction of positive quark jets increases with p_T , this also means that there will be a trend towards $Q^{\text{reco}} < Q^{\text{true}}$, which is a negative response. The spread of the response (as measured by the RMS) does depend both on the p_T and number of tracks (right plots in Figs. 1.30 and 1.31). For low p_T , the RMS of the charge response distribution decreases with p_T and for jet p_T above about 100 GeV, the response RMS increases with p_T . This high p_T trend is consistent with the degradation of the relative momentum resolution as tracks become less curved and also begin to merge in the dense jet core. As expected, the RMS tends to decrease with the number of tracks as fluctuations about the mean are suppressed. However, this trend is less evident at lower κ where the individual contribution to the jet charge from any one track is decreased. There is also a strong correlation between p_T and number of tracks, which can further weaken




Figure 1.25: Truth-level jet charge distribution for simulated dijet events for different parton flavours. Each distribution is normalised to unit area, except the gluon probability distribution function which is normalised to 0.3 for easier comparison. Distributions are shown in two different p_T bins for the leading jet and for two values of the weighting factor κ .




Figure 1.26: Discrimination power of the jet charge to differentiate quark jets of opposite charge (left) and quark from gluon jets (right). Results are obtained from MC truth information in simulated dijet samples.




Figure 1.27: The parton flavour tagging performance of the jet charge, defined as in Fig. 1.26. The left plot shows the positive parton charge jet efficiency versus the negative parton charge rejection (inverse efficiency) for various κ values in a fixed p_T bin. The right plot fixes the positive charge efficiency at 0.5 and then shows the distribution of the negative charge rejection with κ .




Figure 1.28: A summary of the information shown in Fig. 1.27 for many p_T bins. The horizontal axis is the jet p_T and the vertical axis is the maximum negative charge rejection for a fixed positive charge efficiency (ϵ_+) of 0.25, 0.5 or 0.75. For a fixed positive charge efficiency, the optimal κ value and the maximum negative charge rejection vary little. The p_T bins are chosen based on trigger thresholds.




Figure 1.29: Jet charge distributions for heavy-flavour quarks and inclusive quark types and their ratios (lower panels). The left plot shows the comparison for bottom type quarks and the right plots for charm type quarks. Results are obtained from MC truth information in simulated dijet samples.




Figure 1.30: Average (left) and RMS (right) jet charge response ($Q^{\text{reco}} - Q^{\text{truth}}$) as a function of the leading jet p_T for a sample of simulated dijet events and for two different bins in track multiplicity. Uncertainties are from the limited size of simulated samples.

the decreasing trend at high track multiplicity.

In the 2012 LHC data, pileup has a non-negligible effect on reconstruction. However, since the jet charge is built mostly of tracks associated to the primary collision vertex, the performance of this variable is expected to be independent of the number of pileup vertices. This expectation is confirmed in Fig. 1.32 where for two bins of jet p_T and three values of the p_T -weighting factor κ , it is shown that the RMS of the jet charge response is independent of the average number of interactions per crossing ($\langle\mu\rangle$). Related to the dependence of pileup is the choice of track quality criteria used in constructing the charge. The track p_T threshold (500 MeV) and quality cuts are not expected to have an impact on the jet charge response, as shown in Fig. 1.33 for the p_T threshold.




Figure 1.31: Average (left) and RMS (right) jet charge response ($Q^{\text{reco}} - Q^{\text{truth}}$) as a function of the number of tracks used to compute the charge in different p_T bins of the leading jet in dijet simulated events. Uncertainties are from the limited size of simulated samples.




Figure 1.32: The RMS of the jet charge response ($Q^{\text{reco}} - Q^{\text{truth}}$) as a function of the average number of interactions per crossing for three values of κ and in two bins of p_T , as obtained in simulated dijet events. Uncertainties are from the limited size of simulated samples.




Figure 1.33: The dependance of the response resolution (measured using RMS) as a function of the jet p_T for three values of the track p_T threshold and two p_T weighting values $\kappa = 1$ (left) and $\kappa = 0.3$ (right). One can see that for both κ values, there is a big increase in the RMS at low jet p_T when moving to a track threshold of 5 GeV. For the lower value of $\kappa = 0.3$, the trend persists for high p_T since in this case, more weight is given to lower p_T tracks. The lower plots show the ratio of the 5 GeV track threshold distribution with the 500 MeV track threshold.

1.3.2.2 Re-examining the Definition of Jet Charge

This section considers some variations on the definition of jet charge. As noted in Sec. 1.1.2.2, the jet charge is not Lorentz invariant. One possible Lorentz invariant definition uses the jet ‘rest frame’ (jets are massive, so this is sensible - see Sec. 4.1). This variation and others on the jet charge definition are studied in Fig. 1.34 and the performance is quantified in Fig. 1.35. One important variation that has been used in some of the analyses mentioned in the chapter introduction is the one labeled ‘tracks’. For this definition, the denominator of Eq. 1.1 is replaced by the scalar sum of track p_T , raised to the κ , i.e. $Q = \sum Q_i p_{T,i}^\kappa / (\sum p_{T,i})^\kappa$. When $\kappa = 1$, this track-only definition is bounded by 1 and there are spikes in the left plot of Fig. 1.34 at ± 1 corresponding to cases where there is only one track in the jet.

The performance of the track-only jet charge in Fig. 1.35 is nearly the same as for the ‘nominal’ definition, in which the calorimeter jet p_T is used instead, except at low efficiency where the calorimeter-based definition is superior. The Lorentz invariant definition is clearly worse than the other definitions, in part because the mass of generic QCD jets is highly sensitive to diffuse soft radiation. These soft tracks that may be independent from the initiating parton can have significant momentum in the jet rest frame. Only the definition from Eq. 1.1 is considered in the rest of this chapter.




Figure 1.34: Variations on the definition of the jet charge, see the text for details. The left plot shows $\kappa = 0.3$ and the right shows $\kappa = 0.7$. The variants on the colors are for jet originating from a quark with a positive charge versus a negative charge. The small differences in modes between positive and negative quark initiated jets is likely due to statistical fluctuations from the limited MC sample size.




Figure 1.35: ROC curves for the distributions shown in Fig. 1.34. The left plot shows $\kappa = 0.3$ and the right shows $\kappa = 0.7$.

1.3.3 Simulation and In-situ Studies with W Bosons

1.3.3.1 Dataset and Simulation Samples

The studies presented in this section use a subset of the $\sqrt{s} = 8$ TeV data from Run 1 corresponding to 5.8 fb^{-1} . Single lepton triggers are used to select the data. The MC setup is similar to Sec. 1.2.1, except that instead of inclusive dijets as the main process, the target is $t\bar{t}$ production.

Top quark pair production is simulated with two next-to-leading-order (NLO) generators. When studying $W^\pm$ discrimination in $t\bar{t}$ events, MC@NLO [88] is used with the NLO parton density function (PDF) set CT10 [193, 194], and parton showering and underlying event modelled with HERWIG [210] and JIMMY [211], respectively. For jet charge studies in W +jets, $t\bar{t}$ is simulated with POWHEG [89, 200, 201] using the PDF set CT10 and PYTHIA 6.4 [96] for fragmentation and hadronization with the Perugia2011C [212] tune that employs the LO CTEQ6L1 PDF set [198]. In all $t\bar{t}$ MC events, events are filtered by requiring at least one lepton consistent with the lepton trigger selection used for the measurements in which $t\bar{t}$ is relevant. Before filtering, the $t\bar{t}$ cross section is $\sigma_{t\bar{t}} = 238^{+22}_{-24} \text{ pb}$ for a top quark mass of 172.5 GeV. It has been calculated at next-to-next-to leading-order (NNLO) in QCD including resummation of next-to-next-to-leading logarithmic (NNLL) soft gluon terms with $\text{top}++2.0$ [213–218]. The PDF and α_s uncertainties are calculated using the PDF4LHC prescription [219] with the MSTW2008 68% CL NNLO [220, 221], CT10 NNLO and NNPDF2.3 5f FFN [32] PDF sets, and added in quadrature to the scale uncertainty. W +jets production is based on ALPGEN [222], with the parton shower modelled with PYTHIA 6.4 and the Perugia2011C tune; for these samples the production of heavy quarks is modelled separately, and overlapping phase space produced in the inclusive samples is removed.

The single top (s - and Wt -channel) backgrounds are modelled with the same MC@NLO setup as $t\bar{t}$ while the t -channel is modelled with ACERMC [223] and the CTEQ6L1 PDF set interfaced with PYTHIA using the Perugia2011C tune. Like W +jets, the Z +jets backgrounds are modelled with ALPGEN, PYTHIA 6.4 showering, and the Perugia2011C tune. Dibosons are generated with HERWIG using the

CTEQ6L1 PDF set.

The $t\bar{t}$ events are selected with exactly one leptonic $W \rightarrow \mu + \nu$ decay to obtain a high-purity source of hadronically-decaying bosons with known charge in $t\bar{t} \rightarrow (W \rightarrow \ell\nu)(W \rightarrow qq')b\bar{b}$ final-states. Candidate events are chosen by requiring a $p_T > 25$ GeV muon with $|\eta| < 2.5$ and missing transverse momentum $E_T^{\text{miss}} > 20$ GeV; in addition, the sum of the missing transverse momentum and the transverse mass¹¹ of the W boson reconstructed from the lepton and missing momentum is required to be greater than 60 GeV, as expected for leptonic W decays. Muons from heavy-flavour decays are suppressed by requiring the muon to be isolated in both the tracker and calorimeter from unclustered objects as well as from jets. Events must also have at least four jets with $|\eta| < 2.5$ and $p_T > 25$ GeV. Exactly two of these jets must be identified as b -quark jets using the multivariate discriminant ‘MV1’ [224] which includes impact parameter and secondary vertex information as inputs. The chosen MV1 working point corresponds to an average b -tagging efficiency of 70% for b -jets in simulated $t\bar{t}$ events. Among the jets not selected by the b -tagger, there must exist a pair each with $|\eta| < 2.1$ and a dijet invariant mass within 30 GeV of the W boson mass. The two jets with invariant mass closest to the W boson mass are chosen as the W daughter candidates. This procedure selects a sample that is expected to contain more than 90% $t\bar{t}$ production, as shown in Table 1.3 for the positive muon and negative muon channels separately. Figure 1.36 illustrates the object selection.

The charge of the hadronically decaying W boson can be inferred from the measured charge of the selected muon. Therefore, the discrimination power of the jet charge can be directly determined from data. Figure 1.37 shows the distributions of dijet charge from W^+ and W^- decays for two different values of the p_T weighting factor κ from Eq. 1.1. The dijet charge is computed as the sum of the jet charges from the W boson daughter candidates. In μ^+ events, the dijet charge tends to be negative, while the opposite is true for μ^- events, indicating that the dijet charge shows correlation with the charge of the hadronically-decaying W .

¹¹The transverse mass is defined as $m_T^2 = 2p_T^{\text{lep}}E_T^{\text{miss}}(1 - \cos(\Delta\phi))$, where $\Delta\phi$ is the azimuthal angle between the lepton and the missing transverse momentum direction.




Figure 1.36: Schematic representation of the object selection. At least four jets are required: two b -tagged jets B_1, B_2 and at least two non b -tagged jets, labelled J_1, J_2 . The jets J_1 and J_2 are those non b -tagged jets with invariant mass closest to the W boson mass. The charged lepton is used to trigger and a cut on the missing energy from the neutrino is used to purify the sample in $t\bar{t}$ events.

Process	N_{events} with μ^+	N_{events} with μ^-
$t\bar{t}$	3575 ± 29	3522 ± 20
Single Top	126 ± 3	97 ± 3
W +jets	170 ± 29	91 ± 15
Z +jets	23 ± 5	18 ± 3
Dibosons	3 ± 0.4	3 ± 0.3
Total MC	3895 ± 36	3729 ± 25
2012 Data	4095	3893

Table 1.3: The data and MC signal and background yields after all selections for the 5.8 fb^{-1} sample, shown separately for μ^+ and μ^- final states. The MC uncertainties are purely statistical and included solely for the purposes of illustrating the sample composition.

The ratio between data and MC in the lower panels of Fig. 1.37 includes the relevant systematic uncertainties: the jet energy scale uncertainty (JES), the jet energy resolution uncertainty (JER), tracking efficiency, b -tagging related uncertainties, and the uncertainty on the background normalisation. The JES and its uncertainty are determined from a combination of test-beam data, LHC collision data, and MC simulation [225]. The $\pm 1\sigma$ variations are computed as a function of the p_T and η of each reconstructed jet and are then propagated through to the jet charge distributions. To model the impact of the uncertainty on the energy resolution, reconstructed jet energies are smeared by a Gaussian function such that the new width incorporates a $+1\sigma$ variation of the JER uncertainty. The effect on the charge distribution is symmetrized by taking the difference between the nominal and the shifted as 1σ . The JES and JER uncertainties are comparable and amount to about 20% of the nominal in the $|Q| < 1$ region for $\kappa = 1.0$. Track reconstruction efficiency [226] and b -tagging uncertainties are much smaller than the JES and JER contributions (less than percent level). The b -tagging only affects acceptance and not the charge itself, unlike JES and JER which contribute to both acceptance and the actual charge via the jet calorimeter energy in the denominator of the jet charge definition. The uncertainty on the background normalization is taken to be the same as the cross section uncertainty for $t\bar{t}$ stated in Section 2, namely about 6%. This is justified because the combinatorial background from $t\bar{t}$ represents more than 90% of the total background. All uncertainties are added in quadrature for each bin.

To quantify the discriminating power of the dijet charge, the rejection of negatively-charged W bosons is computed against the efficiency for selecting positively-charged W bosons. This relationship is shown in the left plot of Fig. 1.38. The points along the solid (dashed) lines correspond to cuts on the charge distribution in data (MC) for two values of κ . Since $\mu^\pm$ events correspond to hadronic $W^\mp$, the horizontal axis is computed as the fraction of μ^- events beyond a given cut value with respect to all μ^- events and the vertical axis values are the reciprocals of the fraction from μ^+ events past the same cut value, with respect to all μ^+ events. A negative W rejection of about 6 is expected for a positive W boson efficiency of 50%, almost independent of κ . Some degradation of the separation power between positive and negative W bosons is




Figure 1.37: The distribution of the sum of the jet charges from the two daughter candidates in hadronic W boson decays in semileptonic $t\bar{t}$ events. The plot on the left is obtained with p_T weighting factor $\kappa = 1.0$ and the right plot with $\kappa = 0.3$. Events with a $\mu^\pm$ correspond to a hadronically-decaying $W^\mp$. The bottom panels show the data/MC ratios with the bands giving the systematic uncertainties described in the text.

expected to come from the combinatorial background, i.e. the two W daughter candidates may not have originated from or contain all of the partons associated with the W decay. Such an effect appears in both data and MC. Its impact can be estimated by selecting a purer sample, which reduces the combinatorial background. The right plot in Fig. 1.38 shows the positive W efficiency for different numbers of jets in the event. Jets are required to be above 25 GeV in p_T and have $|\eta| < 2.5$. For example, for a fixed positive W efficiency of 50%, the rejection of negative W increases by 20% when the jet multiplicity decreases from six to four.

1.3.3.2 Charge Reconstruction Performance

The detector response for jet charge in W events is qualitatively similar to the response studied earlier in generic quark and gluon jets. The top row of Fig. 1.39 shows the mean dijet charge response versus the dijet track multiplicity for two values of κ in $t\bar{t}$ MC events. The response is close to zero and constant with respect to the




Figure 1.38: The power to reject W^- as a function of the efficiency to tag W^+ as measured in semileptonic $t\bar{t}$ events. Each point on the curve corresponds to a cut value on the charge distribution shown in Fig. 1. The left plot shows results for inclusive jet multiplicity, and the right plot for several multiplicities and $\kappa = 1.0$.

number of tracks. However, the resolution, parameterized by the distribution RMS, does depend on the number of tracks, as can be seen for the same values of κ in the bottom row of Fig. 1.39. The response RMS decreases with the track multiplicity, as in the case of generic quark and gluon jets studied in Sec. 1.3.2.1. The top row of Fig. 1.40 shows the mean response as a function of the hadronic W p_T , defined as the transverse momentum of the dijet system formed from the W daughter candidates. As with track multiplicity, the response is constant around zero (indicating very good agreement between the reconstructed and the true values), while the RMS (bottom row of Fig. 1.40) decreases with the W boson p_T .




Figure 1.39: For $t\bar{t}$ simulated events, the mean (top) and RMS (bottom) distributions of the W daughter dijet charge response as a function of the total number of tracks used to compute the charge for two values of κ . The insets show the (arbitrarily normalised) distribution of the number of tracks. The error bars indicate the statistical uncertainties on the number of MC events.




Figure 1.40: For $t\bar{t}$ simulated events, the mean (top) and RMS (bottom) distributions of the W daughter dijet charge response as a function of the transverse momentum of the dijet system for two values of κ . The insets show the (arbitrarily normalised) distribution of the W candidate p_T . The error bars indicate the statistical uncertainties on the number of MC events.

1.3.3.3 Charge Tagging in a Boosted Topology

In $t\bar{t}$ events, when the hadronic W has a large Lorentz boost, its decay products become merged in the lab frame, obscuring the resolution of the $R = 0.4$ jets that are usually associated with the W decay. In a classical two-body decay of a boosted object, the separation ΔR scales as $2m/p_T$, where m (p_T) is the mass (transverse momentum) of the boosted object; see Chapter 4 for more detail. For a W boson ($m_W = 80.4$ GeV [27]) with a p_T of 200 GeV, a $R = 1.0$ (*large-R*) jet often captures most of the hadronic decay products. The jet charge is considered here also in this boosted scenario. The same MC@NLO semileptonic $t\bar{t}$ MC events as described before are used as a source of boosted W bosons. In each event, the hadronically-decaying W is identified at truth level and its p_T is required to be above 200 GeV. The anti- k_t algorithm is used to cluster the hadronic decay products of the W using an $R = 1.0$ radius parameter. A first definition of jet charge is the simple extension of the procedure described previously: tracks are assigned to the $R = 1.0$ jets in the event using ghost association and then Eq. 1.1 is used with the large- R (calorimeter) jet p_T in the denominator. The distribution of this *large-R jet charge* is shown in Fig. 1.41 for $\kappa = 1.0$ and $\kappa = 0.3$. Large- R jets are chosen as the closest $R = 1.0$ jet in ΔR to the truth $p_T > 200$ GeV hadronic W and $\Delta R(\text{jet}, W) \leq 1.0$ is required. Since jets are only matched geometrically to the truth W boson, a momentum and mass threshold are imposed: only reconstructed large- R jets with $p_T > 100$ GeV and mass above 30 GeV are considered. The jet four-vector ($\mathbf{p}$) is corrected for pileup using the area correction [120, 207] $\mathbf{p} \mapsto \mathbf{p} - \rho \times \mathbf{A}$, where $\mathbf{A}$ is the four-vector jet area determined from ghost four-momenta and ρ is the median p_T density per unit area in $\eta - \phi$ space.

A modification of the large- R jet charge definition can be obtained from *trimming* [227]. To form a trimmed large- R jet, first the jet constituent topo-clusters are grouped using the k_t algorithm with a distance parameter of $R = 0.3$. Then, the clusters (and ghosts) of all the subjets that carry less than 5% of the total jet momentum are removed. The remaining clusters determine the trimmed jet. The tracks associated to the trimmed jet are determined by the ghost tracks that remain after subjets removal. The *trimmed large-R jet charge* is defined, as above, by summing over the tracks according to Eq. 1.1, with the (calorimeter) trimmed jet p_T in

the denominator. The trimmed large- R jet charge is shown in Fig. 1.41 for $\kappa = 1.0$ and 0.3 for the same selection as for the untrimmed distribution. The trimmed and untrimmed jets have similar distributions, with the untrimmed distributions being slightly wider. The reason why there is not much difference in the charge distributions is that trimming only removes 20% of tracks, all of which have a low $\mathbf{p}_T^{\text{track}}/\mathbf{p}_T^{\text{jet}}$ weight ($\lesssim \mathcal{O}(1\%)$) and thus do not contribute significantly to the charge. Trimming removes more than 20% of calorimeter clusters, but the tracks are required to match to the primary vertex and are thus significantly protected against pileup.

In the process of trimming, $R = 0.3$ subjets associated with each large- R jet are clustered together. This gives rise to a third natural definition of the hadronic W charge: the sum of the charge of the two leading k_t subjets. Tracks are already matched with subjets from the ghost association to the $R = 1.0$ jets. This *subjet charge* is shown in Fig. 1.41 for $\kappa = 1.0$ and 0.3. The subjet charge is more spread out than the (trimmed) large- R jet charge. Part of the stretching is from the definitions. To see this, consider an example in which the large- R jet transverse momentum P is parallel to the two subjet momenta $\mathbf{p}$ and $\mathbf{q}$ and assume that $\mathbf{p} + \mathbf{q} = \mathbf{P}$. Then, $1/p + 1/q > 1/(\mathbf{p} + \mathbf{q}) = 1/P$, so the subjet charge will tend to have a larger spread.

The performance of hadronic W charge-tagging in the boosted regime is shown in Fig. 1.42 in terms of the inverse efficiency (rejection) to identify a W^- as a function of the efficiency to identify a W^+ . Since the large- R and trimmed large- R charge distributions are similar, the performance is similar. For a 50% W^+ efficiency, a factor of four rejection is obtained. The subjet charge performs significantly worse than the (trimmed) large- R jet charge. There are several factors that contribute to the difference in performance. For example, there are many selected large- R jets with three or more subjets. In these cases, the jet charge calculation does not include information beyond what is contained in the two leading subjets. In addition, the decay products of the W may not be fully merged into the $R = 1.0$ cone. The large- R jet charge may take this partial contribution into account, but the subjet charge may miss important information from tracks on the edge of the jet.

Jet charge for boosted object tagging is revisited in Sec. 4.2.3 in the context of the W versus Z tagger.




Figure 1.41: The charge of a boosted hadronically-decaying W boson in simulated semileptonic $t\bar{t}$ events for $\kappa = 1.0$ (left) and $\kappa = 0.3$ (right). The hashed distributions are for the extension of the jet charge definition to large- R jets. The solid line histograms show the distribution of the large- R jet charge after trimming and the dashed lines show the sum of the charge of the two leading $R = 0.3$ k_t subjets. Note that the domain of the two jet charge distributions are not the same, but are plotted with the same x-axis range, in contrast to Fig. 1.37.




Figure 1.42: The inverse efficiency (rejection) of boosted W^- bosons as a function of the efficiency for boosted hadronic W^+ bosons for simulated semileptonic $t\bar{t}$ events and $\kappa = 1.0$ (left) and $\kappa = 0.3$ (right). The three curves indicate the performance of three different ways of measuring the jet charge for boosted W bosons.

1.4 Unfolding

In order to facilitate direct comparisons between the data and particle-level simulations and predictions, it is necessary to remove distortions from detector effects. Let $\mathbf{h}_D$ be the detector-level histogram corresponding to measured values $\mathbf{d}_1, \dots, \mathbf{d}_n$ of a particular observable where each event $\mathbf{i} = 1, \dots, n$ passed a selection based on detector-level objects. Furthermore, let $\mathbf{h}_T$ be the histogram of particle-level values of another observable $\mathbf{t}_1, \dots, \mathbf{t}_m$ for events that pass a particle-level event selection for a particular process. Unfolding is the name given to an algorithm for estimating $\mathbf{h}_T$ given $\mathbf{h}_D$. The name *unfolding* is used because *folding* is a procedure for distorting particle-level measurements to simulate the effects of a detector with finite acceptance and resolution. Ideally, the detector-level and particle-level definitions of the observables and event selections are close - a notion that will be quantified below. Intuitively, the stronger the correlation between the detector-level and particle-level object and event selections, the more information the detector-level observations contains about the particle-level quantities and thus the more precise the measurement. It is useful to focus on measurements of *general, well-defined* observables, i.e. that do not depend on the details of the particle-level model or of the particular detector. For example, parton-level momenta are ill-defined¹² because colored objects are not directly observable and their properties in simulation depend on particular (unphysical) parameter values. When unfolding to events chosen with a non-trivial particle-level selection, the measurement is called a *fiducial measurement* and the particle-level selection is called the *fiducial volume* of the measurement. Such measurements are useful from the point of view of unfolding because they require less extrapolation to uninstrumented regions of the detector or to unmeasurable kinematic values. However, fiducial measurements can be a challenge for making theoretical predictions which often require additional assumptions/precision to reduce the calculation to a specified region of phase space.

¹²This has not stopped many analyses from measuring parton-level quantities, treating fragmentation as a ‘detector’. However, the meaning of the measurement depends on the fragmentation model used in the unfolding and thus is not general.

In general, unfolding has to correct for many interrelated effects:

Acceptance and Efficiency Not every particle produced is measured because of the finite coverage of the detector and even those particles, jets, etc. that are detected are not recorded as such because of quality and identification criteria. This effect decreases n relative to m (in particular, n and m need not be equal).

Detector Noise Some of the objects measured in the detector have no particle-level sources. For instance, tracks can be formed from random hits in the inner detector. This category also includes jets and tracks from pileup interactions, which do have a particle-level source but not from the hard-scatter particle-level event. This effect increases n relative to m .

Background Processes A measurement is usually made on a give process and not a particular final state. For example, one may be interested in a property of $t\bar{t}$ events, so the W +jets background needs to be subtracted.

Combinatorics Objects chosen based on some criteria at detector-level may not correspond to the objects chosen at particle-level based on the same criteria. For example, the highest p_T detector-level jet need not originate from the highest p_T particle-level jet.

Detector Scale Detector-level quantities are not always unbiased measurements of the corresponding particle-level quantities. For example, the average jet energy is not exactly the same as the average particle jet energy due to a non-closure in the jet energy scale calibration.

Detector Resolution The finite resolution of the detector smears out particle-level quantities when measured at detector-level.

The first two points account for both per-object acceptance and efficiencies as well as the overall (particle- and detector-level) event selection efficiency.

As a starting point for constructing an unfolding algorithm, consider the *folding equation*:

$$h_{D,i} = \sum_{j=1}^m \text{Pr}_{D|T}(i|j) h_{T,j} + h_{F,i}, \quad (1.13)$$

where $\text{Pr}_{D|T}(i|j)$ is the probability for an event in bin $h_{T,j}$ to be measured and recorded in bin i of h_D and h_F is a histogram containing events at detector-level that were not produced by events of the target process passing the particle-level selection. It is customary to further decompose $h_{F,i} = h_{D,i}f_i + h_{B,i}f_i + h_{B,i}$, where h_B is a histogram containing events at detector-level that pass the particle-level selection but originate from a process that is not the target one¹³ and f_i is the *fake-factor* that accounts for the fraction of the events at detector-level that do not correspond to particle-level events that pass the selection.

Letting $R_{ij} = \text{Pr}_{D|T}(i|j)$ and representing the histograms as vectors, the Eq. 1.13 can be written as a matrix equation $\tilde{h}_D = R h_T$, where $\tilde{h}_D = h_D - h_F$ or equivalently, $\tilde{h}_{D,i} = (1 - f_i)(h_{D,i} - h_{B,i})$. The matrix R is called the *response matrix* and is estimated from simulation. In the matrix form, one may be tempted to solve for $h_T = R^{-1} \tilde{h}_D$. However, even if R is a square matrix and is invertible, $R^{-1} \tilde{h}_D$ may not be the best estimator for h_T because matrix inversion can enhance statistical fluctuations in both h_D , due to a finite dataset, and R , due to a finite simulation, when there are significant off-diagonal transition probabilities in R . For example, consider a simple response matrix

$$R = \begin{pmatrix} 1 - \epsilon & \epsilon \\ \epsilon & 1 - \epsilon \end{pmatrix}, \quad (1.14)$$

where $0 \leq \epsilon < 0.5$ in order to make the matrix invertible by satisfying $\text{Det}(R) =$

¹³In this chapter, there are no relevant background processes. However, this will not be true in Chapter 2.

$1 - 2\epsilon > 0$. The problem is that the variance of $\mathbf{R}^{-1}\tilde{\mathbf{h}}_{\mathbf{D}}$ is proportional to $1/\text{Det}(\mathbf{R})$, which diverges as $\epsilon \rightarrow 0.5$. Ideally, ϵ is as small as possible, but there are many cases where it is not small compared to 0.5 due to a large detector resolution. As an alternative method¹⁴, consider the following Bayesian approach [229]. Using the law of total probability ($\mathbf{T}$ = truth, $\mathbf{D}$ = detector):

$$\mathbf{h}_{\mathbf{T},i} = \sum_{j=1}^n \text{Pr}_{\mathbf{T}|\mathbf{D}}(i|j) \tilde{\mathbf{h}}_{\mathbf{D},j} \quad (1.15)$$

Then, the probability $\text{Pr}_{\mathbf{D}|\mathbf{T}}(i|j)$ can be inverted using Bayes Theorem:

$$\mathbf{h}_{\mathbf{T},i} = \sum_{j=1}^n \frac{\text{Pr}_{\mathbf{D}|\mathbf{T}}(i|j) \text{Pr}_{\mathbf{T}}(i)}{\sum_{i'} \text{Pr}_{\mathbf{D}|\mathbf{T}}(i'|j) \text{Pr}_{\mathbf{T}}(i')} \tilde{\mathbf{h}}_{\mathbf{D},j} = \sum_{j=1}^n \frac{\mathbf{R}_{ij} \bar{\mathbf{h}}_{\mathbf{T},i}^1}{\mathbf{R} \bar{\mathbf{h}}_{\mathbf{T}}^1} \tilde{\mathbf{h}}_{\mathbf{D},j} \equiv \mathbf{B}_1 \tilde{\mathbf{h}}_{\mathbf{D}}, \quad (1.16)$$

where $\bar{\mathbf{h}}_{\mathbf{T}}^1$, with $\sum_{i=1}^n \bar{\mathbf{h}}_{\mathbf{T},i}^1 = 1$, is a *prior density* for $\mathbf{h}_{\mathbf{T}}$. By construction, when $\bar{\mathbf{h}}_{\mathbf{T}}^1 \propto \mathbf{h}_{\mathbf{T}}$, the solution $\mathbf{B}_1 \tilde{\mathbf{h}}_{\mathbf{D}}$ is unbiased (when it exists, $\mathbf{R}^{-1}\tilde{\mathbf{h}}_{\mathbf{D}}$ is also unbiased). In the example above, the matrix $\mathbf{B}_1$ is given by

$$\mathbf{B}_1 = \begin{pmatrix} \frac{(1-\epsilon)\bar{\mathbf{h}}_{\mathbf{T},1}^1}{(1-\epsilon)\bar{\mathbf{h}}_{\mathbf{T},1}^1 + \epsilon\bar{\mathbf{h}}_{\mathbf{T},2}^1} & \frac{\epsilon\bar{\mathbf{h}}_{\mathbf{T},1}^1}{\epsilon\bar{\mathbf{h}}_{\mathbf{T},1}^1 + (1-\epsilon)\bar{\mathbf{h}}_{\mathbf{T},2}^1} \\ \frac{\epsilon\bar{\mathbf{h}}_{\mathbf{T},2}^1}{\epsilon\bar{\mathbf{h}}_{\mathbf{T},2}^1 + (1-\epsilon)\bar{\mathbf{h}}_{\mathbf{T},1}^1} & \frac{(1-\epsilon)\bar{\mathbf{h}}_{\mathbf{T},2}^1}{(1-\epsilon)\bar{\mathbf{h}}_{\mathbf{T},2}^1 + \epsilon\bar{\mathbf{h}}_{\mathbf{T},1}^1} \end{pmatrix}. \quad (1.17)$$

An important property of $\mathbf{B}_1(\epsilon)$ is that its components remain finite as $\epsilon \rightarrow 0.5$ and thus the variance of the estimate also remains finite in this limit. The price paid when the off-diagonal elements of $\mathbf{R}$ are large is a dependence on the prior density $\bar{\mathbf{h}}_{\mathbf{T}}$. This can be mitigated by *iterating* the above procedure. Let

¹⁴There are other unfolding techniques that address the matrix inversion challenge with alternative techniques. One common alternative to the Bayesian method is to apply regularized singular value decomposition (SVD) to the response matrix [228]. This Bayesian method is used exclusively for the rest of this chapter.

$$B_k = \sum_{j=1}^n \frac{R_{ij} \bar{h}_{T,i}^{k-1}}{R \bar{h}_T^{k-1}} \quad \text{and} \quad \bar{h}_T^k = B_{k-1} \tilde{h}_D. \quad (1.18)$$

Then the *Iterative Bayesian Unfolding Method* (IB) estimates $h_T = B_k \tilde{h}_D$. The free parameter k is the number of iterations and is a regularization parameter that, like $\bar{h}_T^1$, must be chosen ahead of time. The prior $\bar{h}_T^1$ is chosen to be the particle-level distribution in simulation and the number of iterations is chosen to compromise between bias and statistical uncertainty. One further step is required because there are some events that pass the particle-level selection but are not measured by the detector-level selection. This is accounted for by an *efficiency factor* ϵ_i that represents the fraction of events in bin i that pass the particle-level selection but not the detector-level selection. As was the case with the fake factor f_i , the efficiency factor ϵ_i is estimated from simulation. The unfolding procedure used for the rest of this chapter is summarized as follows:

1. Using simulation, estimate the response matrix R , the fake factors f , the efficiency factors ϵ , and the particle-level spectrum $\bar{h}_T^1$ used as the initial prior.
2. Correct the data for background processes and noise via

$$h_{D,i} \rightarrow \tilde{h}_{D,i} = (1 - f_i)(h_{D,i} - h_{B,i}).$$

3. Estimate the particle-level histogram h_T by iteratively applying the Bayes method described above:

$$h_{T,i} = \frac{1}{\epsilon_i} \sum_{j=1}^m B_{k,ij} \tilde{h}_{D,j}.$$

4. Estimate the uncertainty based on all the inputs from simulation and the finite statistics of the data. These are described in more detail in Sec. 1.5.

1.4.1 Unfolding for the Jet Charge

The unfolding procedure described in the previous section can be readily extended to histograms of multiple dimensions. The primary interest for the jet charge measurement is to extract the particle-level dependence of the jet charge spectrum on the particle-level jet $\mathbf{p}_T$. Since the jet charge and the jet $\mathbf{p}_T$ are not independent, it is important to simultaneously unfold them. A simple way to extend the IB method for this case is to transform the two-dimensional jet charge and jet $\mathbf{p}_T$ distribution into a one-dimensional histogram. This is accomplished as follows:

1. Bin the jet charge and jet $\mathbf{p}_T$ distributions. The jet charge bin centers are $Q_i, i = 1, \dots, N$ and the jet $\mathbf{p}_T$ bin centers are $P_i, i = 1, \dots, M$. In general, the jet charge and jet $\mathbf{p}_T$ bins can vary in size and the jet charge bin size can vary as a function of jet $\mathbf{p}_T$. The number of jet charge bins is the same for all jet $\mathbf{p}_T$ bins. Binning is described in more detail in Sec. 1.4.1.1.
2. Define the integer map $(i, j) \mapsto z(i, j) = N(i - 1) + j$, for jet charge bin i and jet $\mathbf{p}_T$ bin j . Transform the two-dimensional histogram of jet charge and jet $\mathbf{p}_T$ h_{ij}^{2D} into a one-dimensional histogram via $h_{ij}^{2D} \mapsto h_{z(i,j)}^{1D}$.
3. Use the IB unfolding algorithm described in Sec. 1.4 to unfold h_z^{1D} .
4. Transform the unfolded one-dimensional histogram back into a two-dimensional histogram via the inverse integer map $i = z \bmod N$ and $j = (z - i)/N + 1$ for jet charge bin i and jet $\mathbf{p}_T$ bin j .

The jet charge distribution in a fixed jet $\mathbf{p}_T$ bin is nearly Gaussian and so most of the information in the particle-level distribution is contained in the mean and standard deviation. Figure 1.43 shows the particle-level jet charge distribution in three jet $\mathbf{p}_T$ bins along with a χ^2 minimization to Gaussian distributions. Therefore, instead of measuring the full two-dimensional distribution of the jet charge and the jet $\mathbf{p}_T$, the focus is on the jet $\mathbf{p}_T$ dependence of the jet charge distribution mean and standard deviation.

The jet charge distribution average and standard deviation are extracted from the above procedure using Eq. 1.19:

$$\langle Q_{\text{jet}} \rangle_i = \frac{\sum_{j=1}^N n_j Q_j}{\sum_{j=1}^N n_j} \quad (1.19)$$

$$\sigma_{Q_{\text{jet}}}^2 = \frac{\sum_{j=1}^N n_j Q_j^2}{\sum_{j=1}^N n_j} - \langle Q_{\text{jet}} \rangle_i^2, \quad (1.20)$$

where $i = 1, \dots, M$ is the jet p_T bin and n_i is the content of jet charge bin i in jet p_T bin j . One of the important considerations for choosing the binning is to reduce the bias that the discretization procedure introduces in estimating the above moments, i.e. the difference between $\langle Q_{\text{jet}} \rangle_i$ at particle level and the true mean in jet p_T bin i .




Figure 1.43: A Gaussian fit to the particle-level jet charge distribution in three bins of jet p_T for the more forward of the two selected jets and $\kappa = 0.3$. The distribution is not exactly Gaussian, which is why the fitted curve is slightly shifted with respect to the histograms.

1.4.1.1 Binning

Increasing the number of bins reduces the bias from discretization, at the cost of decreasing the unfolding stability and increasing the uncertainty. The jet charge and jet p_T binning are chosen independently. Due to the falling p_T spectrum the bins in jet p_T increase in size as a function of p_T : $[50, 100)$, $[100, 200)$, $[200, 300)$, $[300, 400)$, $[400, 500)$, $[500, 600)$, $[600, 800)$, $[800, 1000)$, $[1000, 1200)$, $[1200, \infty)$ GeV. For displaying the p_T dependence of the jet charge distribution moments, the jet p_T bin mid-point is used to represent the p_T ¹⁵. The jet charge is steeply falling away from the mean, but it is important to have fine binning to reduce the discretization bias. Jet charge bins are equally spaced in the range $|Q_{\text{jet}}| < 1.8$ for $\kappa = 0.3$, $|Q_{\text{jet}}| < 1.2$ for $\kappa = 0.5$, and $|Q_{\text{jet}}| < 0.9$ for $\kappa = 0.7$. These values are roughly chosen to reduce the overflow fraction to less than about 1% as shown in Fig. 1.44. Events with a jet charge larger than this upper value are placed in the last bin.

An optimization for the number of jet charge bins is demonstrated with Fig. 1.45 and 1.46. The average jet charge is largely insensitive to the number of bins, as long as there are multiple bins on either side of zero. In contrast, the standard deviation of the jet charge distribution is significantly sensitive to the number of bins¹⁶. For 15 bins, the discretization bias is $\lesssim 1\%$; this is the value that is used for the remainder of the analysis.

For the given binning choices, Fig. 1.47 shows the one-dimensional transformation of the joint binned distribution of the jet charge and the jet p_T . The bumps represent individual jet p_T bins and the general decreasing trend is due to the steeply falling jet p_T spectrum. There are 150 bins in total which are the input to the unfolding algorithm.

¹⁵Another approach is to simultaneously unfold the p_T distribution with a finer binning in order to place the measured value at the mean of the p_T distribution. This is mostly an aesthetic change and was implemented, but the overhead was sufficiently cumbersome that it was dropped for the simpler presentation.

¹⁶One can in principle correct for this bias, but the correction depends on the distribution within a bin. If the bin sizes are small, this renders the impact of any systematic uncertainties on the shape subleading; however if the bins are small the need for a correction is also negligible.




Figure 1.44: Various binning choices for the quantization of jet charge and jet p_T . The overflow fraction is the fraction of events with a jet charge outside the range set by the x-axis. Vertical lines indicate the values used for $\kappa = 0.3$ and $\kappa = 0.5$.




Figure 1.45: Various binning choices for the quantization of jet charge and jet p_T for the jet charge distribution average (left) and standard deviation (right). For the RMS, there is a significant dependence up to about 10 iterations, which is why only a few of the small-iteration cases are shown.




Figure 1.46: The discretization bias for various binning choices for the quantization of jet charge and jet p_T for the jet charge distribution average (left) and standard deviation (right).




Figure 1.47: The distribution of the one-dimensional transformation of the jet charge and the jet p_T (see Sec. 1.4.1) for $\kappa = 0.5$ for the more forward jet. All distributions are scaled to have the same normalization.

1.4.1.2 Correction Factors

The correction factors, described in the introduction to Fig. 1.4 are shown in Fig. 1.48 as a function of the one-dimensional transformation of the jet p_T and jet charge described in Sec. 1.4.1. Both the fake and inefficiency factors are nearly one for high jet p_T . In the first p_T bin, there is a significant correction due to threshold effects from the 50 GeV jet p_T requirement. Within a given jet p_T bin, the fake and inefficiency factors are nearly independent of the jet charge.




Figure 1.48: The jet p_T and jet charge distribution is transformed into a one-dimensional variable (see Sec. 1.4.1) and the inefficiency (left) and fake (right) factors are shown as a function of the bin number. The jet p_T increases from left to right and the jet p_T bin edges are marked by vertical dashed lines. A horizontal line at one indicates that no correction is applied. The above plots are for $\kappa = 0.5$ and the more forward jet, but the distributions for the other cases are qualitatively similar.

The dominant contribution to the corrections outside of the first bin in Fig. 1.48 is the p_T symmetry requirement. This is demonstrated by Fig. 1.49. There are no events in simulation in the second p_T bin that pass the jet p_T symmetry requirement but fail the detector-level event selection. In contrast, there are such events in the first jet p_T bin. These jets fail the event selection due to the detector-level jet $p_T > 50$ GeV threshold.




Figure 1.49: For events which pass the truth-based selection but fail the reconstructed event selection, this plot shows $|1 - p_T^1/p_T^2|$ for the reconstructed jet p_T . The horizontal axis is the p_T bin number, from 1-10. In only the first p_T bin, there are events which pass the p_T asymmetry cut but fail the p_T cut. In every other bin, the event fails the truth selection due to the asymmetry cut. The histogram is normalized per p_T bin.

1.4.1.3 Response Matrix

The full response matrix for the one-dimensional transformation of the jet p_T and the jet charge is shown in Fig. 1.50. There are discrete regions along the diagonal corresponding to the 10 jet p_T bins. The large strip just above and below the main diagonal corresponds to events that migrate between jet p_T bins. The probability to migrate into a lower jet p_T bin (below the diagonal) is higher than the probability to migrate to a higher jet p_T bin because the distribution of jet p_T within the particle-level bin is steeply falling. Within one of these regions, the response matrix is peaked along the diagonal, but due to the broad jet charge resolution, the probability for a particle level event to remain in the same bin is $\lesssim 25\%$. This probability decreases as a function of jet p_T as the response matrix is more spread out away from the diagonal. This trend is more evident in Fig 1.51 which shows the response matrix in a given jet p_T bin.




Figure 1.50: The conditional distribution of the detector-level one-dimensional transformation of the jet charge and jet p_T (see Sec. 1.4.1) in bins of the particle-level analogue (response matrix) for the $\kappa = 0.5$ for the more forward jet. The response matrix for the other variants is qualitatively similar.




Figure 1.51: The response matrix for $\kappa = 0.5$ for the more forward jet in various jet p_T bins.

1.4.1.4 The Number of Iterations

The number of iterations is an arbitrary tuning parameter that can be used to tradeoff bias with statistical uncertainty. Figure 1.52 shows the average and RMS of the jet charge distributions as a function of jet p_T . The HERWIG++ MC is treated as data and the response matrix is derived from PYTHIA 8. About four iterations minimizes the bias and henceforth used as the nominal parameter setting.




Figure 1.52: Herwig unfolded with Pythia (left) and the statistical uncertainty (right) for various iteration multiplicity settings in the iterative Bayesian unfolding algorithm for the mean (top) and RMS (bottom).

1.4.1.5 Unfolded Data

Figure 1.53 displays the p_T -dependence of the jet charge distribution's mean and standard deviation for detector-level data and simulation and for particle-level simulation. The differences between the simulated detector- and particle-level distributions give a indication of the corrections required to account for detector acceptance and resolution effects in the unfolding procedure. The growing difference between the particle- and detector-level average jet charge is due to the loss of charged-particle momentum inside jets as a result of track merging. At particle level, the standard deviation of the jet charge distribution decreases with increasing p_T , but at detector level it increases with p_T due to resolution effects.




Figure 1.53: The detector-level (data and simulation) and particle-level jet charge distribution's (a) average and (b) standard deviation as a function of the jet p_T for the more forward jet. The ratios in the bottom panel are constructed from the simulation, and show the prediction of detector-level PYTHIA over the data (top ratio), and detector-level PYTHIA over particle-level PYTHIA (bottom ratio). Bars on the data markers represent only the statistical uncertainties. For both (a) and (b), $\kappa = 0.5$.

The data are unfolded using the iterative Bayesian technique [229], implemented in the RooUnfold framework [230]. Figure 1.54 shows the unfolded distribution over all

bins of the one-dimensional transformation of the jet p_T and jet charge distributions. Even though there are $\lesssim 10\%$ differences in the mean and $\lesssim 1\%$ differences in the standard deviation between the data and simulation (Fig. 1.53), there are many bins in Fig. 1.54 with large deviations from unity in the ratio. This is due to two effects:

1. The plots all have the same normalization. Since most events are in the first p_T bin, there is a compensating offset in the other p_T bins. This is an artifact of the normalization.
2. Small changes in the mean and RMS can result in large changes in the ratio of the raw distribution away from zero. To see this, suppose that the true and unfolded distributions in given p_T bin are exactly Gaussians with means zero and standard deviations σ_t and σ_u , respectively. Then, the ratio r will depend on the distance x of the jet charge bin from zero in the following way:

$$r = \frac{\sigma_u}{\sigma_t} \exp \left(-\frac{x^2}{2} \left(\frac{1}{\sigma_u^2} - \frac{1}{\sigma_t^2} \right) \right). \quad (1.21)$$

In particular, the ratio will go to zero or blow up to infinity (depending on the ordering of σ_u and σ_t) as $|x|$ becomes large. For some numerical values, for $\sigma \sim 0.5$ and $|\sigma_u/\sigma_t - 1| \sim 2\%$, the ratio will change by $\sim 6\%$ when $x = 0.5$ and $\sim 25\%$ when $x = 1$, which is consistent with the behavior in Fig. 1.53 and 1.54.




Figure 1.54: The distribution of the 1D transformation of the 2D binned jet charge and jet p_T distribution for the more forward jet for $\kappa = 0.5$. All distributions are scaled to have the same normalization.

1.5 Systematic uncertainties

All stages of the jet charge measurement are sensitive to sources of potential bias. The three stages of the measurement are listed below, with an overview of the systematic uncertainties that impact the results at each stage:

Correction Factors: Fake and inefficiency factors are derived from simulation to account for the fraction of events that pass either the detector-level or particle-level fiducial selection, but not both. These factors are generally between 0.9 and 1.0 except in the first p_T bin, where threshold effects introduce corrections that can be as large as 20%. Experimental uncertainties correlated with the detector-level selection acceptance, such as the jet energy scale uncertainty, result in uncertainties in these correction factors. An additional source of uncertainty on the correction factors is due to the explicit dependence on the particle-level jet charge and jet p_T spectra. A comparison of particle-level models (PYTHIA and HERWIG++) is used to estimate the impact on the correction factors.

Response Matrix: For events in simulation that pass both the detector-level and particle-level fiducial selections, the response matrix describes migrations between bins when moving between the detector level and the particle level. The response matrix is taken from simulation and various experimental uncertainties on the jet charge and jet p_T spectra result in uncertainties in the matrix. Uncertainties can be divided into two classes: those impacting the calorimeter jet p_T and those impacting track reconstruction inside jets.

Unfolding Procedure: A data-driven technique is used to estimate the potential bias from a given choice of prior and number of iterations in the IB method [231]. The particle-level spectrum is reweighted using the response matrix so that the simulated detector-level spectrum has improved agreement with data. The modified detector-level distribution is unfolded with the nominal response matrix and the difference between this and the reweighted particle-level spectrum is an indication of the bias due to the unfolding method.

The following subsections describe the above uncertainties in more detail. Uncertainties on the calorimeter jet p_T are described in Sec. 1.5.2 and the uncertainties related to tracking are described in Sec. 1.5.3. Summaries of the systematic uncertainties for the more forward jet and $\kappa = 0.5$ are found in Table 1.4 and Table 1.5 for the average jet charge and the jet charge distribution's standard deviation, respectively¹⁷. The uncertainties for the more central jet are similar. Figure 1.55 presents a visualization of the uncertainties in Tables 1.4 and 1.5.




Figure 1.55: A visualization of the systematic uncertainties for the jet charge mean (left) and standard deviation (right). As a result of displaying the uncertainty as a fraction of the mean or RMS, the uncertainty appears artificially large in the first p_T bin in the left plot where the average jet charge is nearly zero (small compared to the resolution).

¹⁷The uncertainties on the first p_T bin of the average jet charge are much larger than on the other bins because the mean is small compared to the resolution.

Average Jet Charge Systematic Uncertainty [%]	Jet p_T Range Lower Edge [100 GeV]									
	0.5	1	2	3	4	5	6	8	10	12
JES	+8.4 -13.6	+3.8 -3.5	+0.9 -5.0	+0.8 -0.3	+1.1 -1.6	+1.1 -1.1	+0.7 -1.0	+0.7 -0.9	+0.4 -0.7	+0.9 -0.3
JER	+6.8 -6.8	+2.3 -2.3	+0.7 -0.7	+0.7 -0.7	+0.3 -0.3	+0.3 -0.3	+0.1 -0.1	+0.1 -0.1	+0.1 -0.1	+0.3 -0.3
Charged Energy Loss	+0.0 -0.0	+0.0 -0.0	+0.0 -0.0	+0.0 -0.0	+1.7 -0.0	+1.5 -0.0	+1.5 -0.0	+1.5 -0.0	+1.6 -0.0	+3.6 -0.0
Other Tracking	+3.3 -1.6	+0.0 -0.4	+0.9 -0.2	+0.7 -0.1	+0.5 -0.4	+1.4 -0.6	+0.7 -0.9	+1.2 -1.2	+1.1 -1.3	+0.9 -1.7
Track Multiplicity	+0.0 -1.5	+0.1 -0.0	+0.0 -0.6	+0.0 -1.1	+0.0 -0.8	+0.0 -0.6	+0.0 -1.2	+0.0 -1.4	+0.0 -2.1	+0.0 -2.9
Correction Factors	+23 -23	+0.9 -0.9	+0.8 -0.8	+1.0 -1.0	+0.3 -0.3	+0.6 -0.6	+0.1 -0.1	+0.3 -0.3	+0.2 -0.2	+0.1 -0.1
Unfolding Procedure	+28 -28	+2.4 -2.4	+0.3 -0.3	+0.2 -0.2	+0.2 -0.2	+0.3 -0.3	+1.1 -1.1	+1.0 -1.0	+1.6 -1.6	+0.6 -0.6
Total Systematic	+38 -39	+5.1 -4.9	+1.7 -5.2	+1.6 -1.7	+2.1 -1.9	+2.4 -1.6	+2.1 -2.1	+2.3 -2.3	+2.6 -3.0	+3.8 -3.4
Data Statistics	28	7.4	1.4	0.7	0.3	0.6	0.9	2.0	4.2	7.0
Total Uncertainty	+47 -48	+9.0 -8.9	+2.2 -5.4	+1.8 -1.9	+2.1 -1.9	+2.5 -1.7	+2.3 -2.3	+3.0 -3.0	+5.0 -5.2	+8.0 -7.8
Measured Value [0.1e]	0.014	0.24	0.49	0.65	0.76	0.82	0.92	1.00	1.08	1.15

Table 1.4: A summary of all the systematic uncertainties and their impact on the mean jet charge for $\kappa = 0.5$ and the more forward jet. The correction factors are the fake and inefficiency corrections applied before/after the response matrix. The Other Tracking category includes uncertainty on the track reconstruction efficiency, track momentum resolution, charge misidentification, and fake track rate. All numbers are given in percent. As a result, the uncertainty appears artificially large in the first p_T bin where the average jet charge is nearly zero (small compared to the resolution).

Standard Deviation Systematic Uncertainty [%]	Jet p_T Range [100 GeV]									
	0.5	1	2	3	4	5	6	8	10	12
Total Jet Energy Scale	+1.9 -1.7	+1.5 -1.3	+1.1 -1.1	+1.1 -1.0	+0.9 -0.8	+1.0 -0.7	+0.8 -0.8	+0.7 -0.8	+0.5 -0.5	+0.5 -0.5
Jet Energy Resolution	+1.3 -1.3	+0.3 -0.3	+0.1 -0.1	+0.2 -0.2	+0.3 -0.3	+0.4 -0.4	+0.2 -0.2	+0.2 -0.2	+0.2 -0.2	+0.2 -0.2
Charged Energy Loss	+0.0 -0.0	+0.0 -0.0	+0.0 -0.0	+0.0 -0.0	+0.2 -0.0	+0.3 -0.0	+0.3 -0.0	+0.3 -0.0	+0.4 -0.0	+1.1 -0.0
Other Tracking	+0.0 -0.3	+0.1 -0.3	+0.2 -0.4	+0.3 -0.4	+0.4 -0.5	+0.5 -0.4	+0.5 -0.5	+0.5 -0.5	+0.5 -0.4	+0.4 -0.4
Track Multiplicity	+0.0 -0.2	+0.0 -0.3	+0.0 -0.2	+0.0 -0.1	+0.0 -0.0	+0.1 -0.0	+0.2 -0.0	+0.2 -0.0	+0.3 -0.0	+0.2 -0.0
Correction Factors	+0.9 -0.9	+0.1 -0.1	+0.0 -0.0	+0.1 -0.1	+0.0 -0.0	+0.1 -0.1	+0.0 -0.0	+0.0 -0.0	+0.0 -0.0	+0.0 -0.0
Unfolding Procedure	+1.9 -1.9	+0.4 -0.4	+0.0 -0.0	+0.1 -0.1	+0.2 -0.2	+0.0 -0.0	+0.1 -0.1	+0.3 -0.3	+0.4 -0.4	+1.7 -1.7
Total Systematic	+3.1 -3.0	+1.6 -1.5	+1.1 -1.2	+1.2 -1.1	+1.1 -1.0	+1.2 -0.9	+1.0 -0.9	+1.0 -1.0	+1.0 -0.8	+2.1 -1.8
Data Statistics	0.9	0.3	0.1	0.1	0.0	0.1	0.1	0.3	0.6	1.0
Total Uncertainty	+3.2 -3.1	+1.6 -1.5	+1.1 -1.2	+1.2 -1.1	+1.1 -1.0	+1.2 -0.9	+1.0 -1.0	+1.1 -1.0	+1.2 -1.0	+2.4 -2.1
Measured Value [0.1e]	4.10	3.87	3.75	3.72	3.70	3.69	3.68	3.67	3.62	3.55

Table 1.5: A summary of all the systematic uncertainties and their impact on the jet charge distribution's standard deviation for $\kappa = 0.5$ and the more forward jet. The correction factors are the fake and inefficiency corrections applied before/after the response matrix. The Other Tracking category includes uncertainty on the track reconstruction efficiency, track momentum resolution, charge misidentification, and fake track rate. All numbers are given in percent.

1.5.1 Correction Factors

There are two components to the uncertainty in the fake and inefficiency factors described in Sec. 1.4. Experimental uncertainties are estimated by re-computing the factors coherently with the variations in the response matrix, as described in Sec. 1.5.2 and 1.5.3. The correction factors encode differences between particle-level and detector-level selections. The experimental systematic uncertainties take into account variations in the detector-level event selection efficiency. In order to estimate the uncertainty on the particle-level selection efficiency, two particle-level models are compared. Fixing the response matrix, the fake and inefficiency factors in PYTHIA 8 are re-weighted to match the corresponding factors in HERWIG++. The left plots of Fig. 1.56 and 1.57 show the bin-by-bin difference when unfolding the nominal detector-level PYTHIA 8 sample with the nominal PYTHIA 8 response matrix but fake and inefficiency factors from HERWIG++. These differences are mostly below 1% but can be as high as 10% in the first p_T bin. The corresponding differences in the extracted jet charge average and jet charge distribution standard deviation are shown in the middle and right plots of Fig. 1.56 and 1.57. In all p_T bins aside from the first one, the uncertainties are less than 1%. For the jet charge distribution standard deviation, these uncertainties are mostly less than 0.1%.




Figure 1.56: The uncertainty on the fake factors by comparing the predictions from PYTHIA 8 and HERWIG for the more forward jet and $\kappa = 0.5$. The left plot shows the differences in all bins of the combined jet p_T and jet charge distributions. The middle and right plots show the uncertainty on the jet charge and jet charge distribution standard deviation, respectively.




Figure 1.57: The uncertainty on the inefficiency factors by comparing the predictions from PYTHIA 8 and HERWIG for the more forward jet and $\kappa = 0.5$. The left plot shows the differences in all bins of the combined jet p_T and jet charge distributions. The middle and right plots show the uncertainty on the jet charge and jet charge distribution standard deviation, respectively.

1.5.2 Calorimeter jet uncertainties

Jets are calibrated so that the detector-level $\mathbf{p}_T$ is an unbiased measurement of the particle-level jet $\mathbf{p}_T$ and various data-driven techniques are used to derive *in situ* estimates of the difference in this calibration between the data and the simulation. Uncertainties in the energy scale and resolution of calibrated jets impact the jet charge in the normalization of Eq. 1.1 (but preserve the jet charge sign) as well as the binning for the 2D distribution. Complete details of this source of uncertainty can be found in Ref. [121]. There are many components of the jet energy scale uncertainty. The *in situ* correction is derived from data using the momentum balance in events with Z bosons (low $\mathbf{p}_T$) or photons (moderate $\mathbf{p}_T$) produced in association with jets as well as the balance of multijet (high $\mathbf{p}_T$) and dijet (high $|\eta|$) systems. Uncertainties on this method stem from the modeling of these processes in simulation. There is also a contribution from the response to single hadrons [232]. Additional sources of uncertainty are due to the modeling of the in-time and out-of-time pileup corrections to the jet energy scale as well as differences in the response due to the flavor of the jet. To assess the impact of each component of the jet energy scale uncertainty, the jet energies in simulation are shifted according to the $\mathbf{p}_T$ - and η -dependent $\pm 1\sigma$ variations. For a fixed variation, the response matrix, and fake and inefficiency factors are recomputed and the unfolding procedure is repeated. The resulting uncertainty on the jet charge distribution's mean and standard deviation is about 1% or less for jet $\mathbf{p}_T$ above 200 GeV. While subdominant for the average jet charge, the JES uncertainty is dominant for $\mathbf{p}_T \lesssim 1$ TeV for the jet charge distribution RMS. This is because scaling the jet $\mathbf{p}_T$ by multiplicative factor directly scales the jet charge RMS by the same factor. The jet charge mean is less effected because of cancellations between positive and negative charges.

The jet energy resolution uncertainty is derived using data-driven techniques in dijet events [233]. To assess the impact of a slightly larger jet energy resolution, jet energies are smeared according to $\mathbf{p}_T$ - and η -dependent factors and propagated through the entire unfolding procedure, as for the jet energy scale uncertainty. The jet energy resolution uncertainty is subdominant to the jet energy scale uncertainty.

1.5.3 Tracking uncertainties

Uncertainties on tracking are broken down into contributions related to the efficiency of reconstructing tracks and measurements of those tracks that are successfully reconstructed. In particular, Sec. 1.5.3.2 describes the isolated track reconstruction efficiency due to the uncertainty in the inner detector material and Sec. 1.5.3.3 documents a novel technique for constraining the uncertainty on the modeling of track reconstruction inside the dense hit environment in the core of high p_T jets. Then, Sec. 1.5.3.5 discusses an estimate of the track momentum resolution due to the modeling of the detector material, magnetic field, and the spatial resolution and alignment of the various detector components. Additional uncertainties related to the identification of the charge of high p_T tracks are in Sec. 1.5.3.6 and the impact of tracks resulting from random combinations of hits is in Sec. 1.5.3.7. Table 1.6 gives an overview of the method and relative size of the various tracking systematic uncertainties. A common tool for studying and varying the track efficiency and momentum resolution is truth-matching tracks to charged particles, as discussed in Sec. 1.5.3.1.

Source	Method	Approximate Size	Section
Isolated Efficiency	Material Variations	$\lesssim 1\%$ for $ \eta < 2.1$	1.5.3.2
Tracking in Jets	r_{track} (Alternate: ζ)	$\lesssim 4\%$ at high p_T	1.5.3.3
Momentum Resolution	Resonance Decays	$\sim 2\%$ at high p_T	1.5.3.5
Charge Identification	Resonance Decays	negligible	1.5.3.6
Fake Tracks	Simulation Variations	$\lesssim 0.5\%$ for $\sigma(\text{Jet Charge})$	1.5.3.7

Table 1.6: An overview of the method and relative size of the various tracking systematic uncertainties.

1.5.3.1 Truth Matching

In the simulation, GEANT4 models the interaction of charged particles with the material of the inner detector. The deposited energy in each detector element is later digitized and forms the input for the pattern recognition for track reconstruction. By matching the deposited energy from an individual charged particle with hits on a track, one can associate charged particles to tracks. This is useful for studying the tracking momentum resolution and charge identification as well as the track reconstruction efficiency. For each track, define the variable Pr_{trk} as

$$\text{Pr}_{\text{trk}} = \frac{\sum_{i \in \text{ID}} W_i^{\text{matched}}}{\sum_{i \in \text{ID}} W_i}, \quad (1.22)$$

where

$$W_i = \begin{cases} w_i & \text{a hit on layer } i \text{ is part of the track} \\ 0 & \text{else} \end{cases}, \quad (1.23)$$

with hit-weight w_i that depends on the detector (defined below) and

$$W_i^{\text{matched}} = W_i \times \begin{cases} 1 & \text{the matched truth particle deposited energy} \\ 0 & \text{else} \end{cases}. \quad (1.24)$$

The truth matched particle is the particle in simulation that deposits energy (from GEANT4) in the most detector elements in common with the track. In other words, if T is the set of truth particles in the simulation, then the matched particle index i is given by

$$i = \operatorname{argmax}_{j \in T} \sum_{k \in ID} \begin{cases} 1 & \text{Particle } j \text{ left energy in } k \text{ and } k \text{ is part of the track.} \\ 0 & \text{else} \end{cases} \quad (1.25)$$

If multiple particles deposited energy in the same pixel, only the one that left the highest energy is considered. The set ID contains the various layers of the inner detector (the pixel detector, SCT, and TRT) and $w_i = 10$ for the pixel detector, 5 for the SCT and 1 for the TRT. The weight for the pixel detector is twice the weight for the SCT because two hits are required in the SCT to give 3D information about the track location, but one hit in the pixel detector already gives this information. A higher weight is used for the pixel detector over the TRT because even though the TRT is useful for the momentum measurement, many of the important track parameters (such as d_0, z_0) are defined in the pixels.

A track in simulation is considered *real* if $\Pr_{\text{trk}}$ is at least 0.5 and *fake* otherwise. One can remove the resolution of real tracks by replacing their momentum with the matched charged particle $\mathbf{p}_T$.

1.5.3.2 Isolated Track Reconstruction Efficiency

The uncertainty on the track reconstruction efficiency is mostly due to the uncertainty in the material in the inner detector. The material is known to within $\sim 5\%$ [234]. This precise modeling of the material in the ID has led to sub-percent level uncertainties in the track reconstruction efficiency for $|\eta| < 2.1$ [235]. These uncertainties are estimated as a function of p_T and η by comparing the track reconstruction efficiency in simulated detector geometries with various levels of material in the ID. In the forward region of the tracking acceptance, the material is less constrained and so older and larger uncertainties are still used to set the uncertainty based on the radius dependence of the $K_s^0 \rightarrow \pi^+\pi^-$ reconstructed invariant mass and the length of tracks reaching into the SCT [226]. Table 1.7 summarizes the track reconstruction efficiency uncertainties.

p_T [GeV]	$ \eta < 1.5$	$1.5 < \eta < 2.1$	$2.1 < \eta < 2.3$	$2.3 < \eta < 2.5$
[0.5,1]	0.7%	1.2%	4%	7%
≥ 1	0.5%	1.1%	3.2%	5.6%

Table 1.7: A summary of the track reconstruction efficiency uncertainties.

In order to estimate the impact of these uncertainties, tracks are randomly removed with η and p_T dependent probabilities as stated in Table 1.7. The studies used to determine the inclusive track reconstruction efficiency did not have an explicit track χ^2 requirement. Since this analysis requires $\chi^2/\text{NDF} \geq 3$, we must also check the data/MC differences of this further selection. Figure 1.58 shows that the χ^2 cut is very efficient, with a $\geq 99\%$ track efficiency in all jet p_T bins. The efficiency is generally higher in the simulation than in the data, with a $\lesssim 10\%$ difference in all p_T bins and there is no strong evidence for a p_T dependence in the level of the mis-modeling. In order to assess the impact of the χ^2 requirement mis-modeling, tracks are randomly removed with a jet p_T -dependent probability. The $\chi^2/\text{NDF} \geq 3$ requirement efficiency in the simulation (from the left plot of Fig. 1.58) is well approximated by $1 - f(p_T)$ where $f(x) = a + bx + cx^2$ for $a = 0.005$, $b = 4 \times 10^{-6}/\text{GeV}$,

and $c = -6.5 \times 10^{-10}/\text{GeV}^2$. Therefore, tracks are removed randomly with probability given by $10\% \times (1 - f(p_T))$.




Figure 1.58: Left: The efficiency of the track $\chi^2/\text{NDF} \geq 3$ requirement given all other track requirements as a function of the jet p_T . The track p_T dependence for three p_T bins is shown in the right plot. Note that no statistical uncertainties are included in the right plot (the size of the bin-to-bin fluctuations indicates the size of these uncertainties). The low jet p_T efficiency is poor at high track p_T due to the large contribution from fakes.

1.5.3.3 Track Reconstruction Inside Jets

In addition to the loss of tracks due to the material in the inner detector, tracks can be lost due to the high hit density inside the core of jets. A useful variable for quantifying this loss is the charged energy ratio (CER), computed from particle-level jets that are geometrically matched to detector-level jets:

$$\text{CER} = \left\langle \frac{\sum p_{\text{T}}^{\text{charged,matched}}}{\sum p_{\text{T}}^{\text{charged}}} \right\rangle, \quad (1.26)$$

where the denominator runs over all charged particles in the particle-level jet and the numerator runs over all tracks reconstructed inside the detector-level jet. To remove track resolution effects, the tracks in the numerator are replaced with the matched charged particle (fakes and secondaries¹⁸ are thus excluded). The CER is plotted in Fig. 1.59 as a function of the jet p_{T} . The CER decreases at low jet p_{T} due to decreasing importance of losses due to hadronic interactions inside the detector and decreases at high jet $p_{\text{T}} \gtrsim 500$ GeV due to track merging inside high density jet cores. The *loss* is defined

$$\text{loss}(\text{jet } p_{\text{T}}) = \max_{\text{jet } p'_{\text{T}}} \text{CER}(\text{jet } p'_{\text{T}}) - \text{CER}(\text{jet } p_{\text{T}}). \quad (1.27)$$

For most analyses using tracks, the uncertainty on the modeling of the loss in Eq. 1.27 is negligible because the jets have $p_{\text{T}} \lesssim 500$ GeV. Early Run 1 studies also suggest that in this low p_{T} regime, hit sharing is well modeled by the simulation [236]. However, the jet charge measurement is probing a new kinematic regime involving tracks inside jets: the loss is not small and therefore a careful assessment of the systematic uncertainty is critical. An early Run 2 method uses double peaks in the dE/dx distribution [237]. Such an approach is quite general, but neglects the impact of track p_{T} , which is important for the jet charge due to the track p_{T} weighting in the definition. The rest of this section describes a new method for constraining the loss modeling with data

¹⁸Tracks from material interactions such as photon conversions, $\gamma \rightarrow e^+e^-$.

using a detector-level analogue to the CER.




Figure 1.59: The CES as a function of the particle-level jet p_T .

An outline of the new method is as follows:

1. Demonstrate that the calorimeter loss from the charged-to-total energy ratio (CTER) can be used as a proxy for the loss derived from the CER.
2. Measure the CTER in the data to constrain the loss.
3. Translate the modeling of the loss to an uncertainty on the tracking efficiency. Consider the impact on the jet charge measurement from all parameterizations of the tracking inefficiency that reproduce the measured loss.
4. Prove that for the (average) jet charge, the tracking inefficiency can be treated as independent per track.

5. The systematic uncertainty on the jet charge is then given by the data/MC difference in the measured loss and is conservatively implemented using an additional tracking inefficiency that has the biggest impact on the jet charge.

The CTER is defined as

$$\text{CTER} = \left\langle \frac{\sum \mathbf{p}_T^{\text{charged}}}{\mathbf{p}_T^{\text{jet}}} \right\rangle, \quad (1.28)$$

where the denominator is the full (particle- or detector-level jet) and the numerator is the sum over tracks (detector-level) or charged particles (particle-level). Analogously to the loss, the calorimeter (calo) loss is

$$\text{calo loss}(\text{jet } \mathbf{p}_T) = \max_{\text{jet } \mathbf{p}'_T} \text{CTER}(\text{jet } \mathbf{p}'_T) - \text{CTER}(\text{jet } \mathbf{p}_T), \quad (1.29)$$

and the fractional calo loss is the calo loss divided by the maximum CTER. Analogously, the fractional loss is the loss divided by the maximum CER. The first observation is that in the simulation, the fractional calo loss is similar to the fractional loss. This makes sense heuristically, since the energy depositions in the calorimeter do not depend on how close the particles are when they reach the calorimeter and so fractional changes in the CTER as a function of jet $\mathbf{p}_T$ should be due to tracking inefficiencies in the core of jets. Empirical evidence for this similarity is shown in Fig. 1.60, in which the fractional loss and the fractional calo loss (detector-level) are nearly identical as a function of jet $\mathbf{p}_T$. Differences between the data and MC in the left plot of Fig. 1.60 indicate that the simulation underestimates the fractional loss by a relative $\sim 10\%$.

A few more details about the $\mathbf{p}_T$ -dependence of the CTER are required before performing a careful measurement of the loss. First of all, it is important that the fake (see Sec. 1.5.3.7) and secondary track (see Fig. 1.70) rates are negligible with the track quality criteria requirements¹⁹. Next, it is crucial that the particle-level

¹⁹It would have been sufficient for these rates to be independent of $\mathbf{p}_T$ for the charged-energy




Figure 1.60: Left: The fractional calo loss in data and simulation. Right: the fractional loss in simulation.

CTER does not depend on p_T . If it did, then changes in the detector-level CTER may simply be due to changes in the particle-level CTER. Fortunately, the particle-level CTER is p_T independent and is nearly $2/3$ due to isospin: there are nearly twice as many charged pions inside jets as neutral pions, with small deviations due to presence of heavier hadrons and bremsstrahlung photons. Furthermore, the exact value of the particle-level CTER is largely generator-independent, in part because it is highly constrained by low(er) energy physics. Figure 1.61 shows the detector- and particle-level CTER as a function of the jet p_T ; the difference between PYTHIA 8 and HERWIG++ is at or below the 0.1% level.

The final requirement for the CTER to be a useful proxy for the CES is that the jet p_T should be an unbiased measurement of the particle-jet p_T . If there is a p_T -dependent bias, then changes in the CTER as a function of p_T could be due to the change in the biased measurement of the jet p_T . This is not exactly satisfied. The jet energy scale uncertainty is not zero and does depend on p_T . However, it is small, and one can quantify its influence on the measurement by conservatively adding the JES

method to work.




Figure 1.61: The detector- and particle-level CTER as a function of the jet p_T for Pythia, Herwig, and the data.

uncertainty in quadrature to the uncertainty that is determined from the differences between data and MC in describing the calo loss.

All the ingredients are now ready to quantitatively measure the calo loss and by proxy determine an uncertainty on the fractional loss. The data and MC calo loss distributions have already been shown in Fig. 1.60 and Fig. 1.61. The data/MC fractional difference is shown in Fig. 1.62, also added in quadrature with the data statistical uncertainty and all of the JES uncertainty components. The total uncertainty is about 1% until about 1.2 TeV, after which it increases to about 2%.

The next challenge is to translate the fractional uncertainty in the loss to an uncertainty in the tracking (in)efficiency. To begin, assume that one can treat the inefficiency as independent per track. This means that a good model for increasing the loss in the simulation to match the loss in the data is to randomly remove tracks




Figure 1.62: The uncertainty in the fractional loss determined from the difference in the data/MC of the fractional calo loss and the JES uncertainties.

independent of the other tracks in the jet. This assumption will be justified below. Consider all possible tracking inefficiency uncertainty parameterizations that reproduce the loss:

$$\left\langle \frac{\sum_i p_{T,i}^{\text{kept}}}{\sum_i p_{T,i}} \right\rangle = 1 - \epsilon, \quad (1.30)$$

where ϵ is the uncertainty determined in Fig. 1.62, the sum runs over all tracks associated to the detector-level jet, and *kept* indicates that the track was retained after randomly removing tracks. Consider a generic parameterization: $\text{Pr}(\text{drop track } i) = \alpha p_{T,i}^n$, where n is a non-negative integer. For a fixed parameterization (n), there is one measurement (ϵ) and one unknown (α). Due to the form of Eq. 1.30, the relationship between α and ϵ is linear. Therefore, the solution for α is unique:

$$1 - \epsilon = \left\langle \frac{\sum_i p_{T,i}^{\text{kept}}}{\sum_i p_{T,i}} \right\rangle = \frac{\sum_i (1 - \alpha p_{T,i}^{\mathfrak{n}}) p_{T,i}}{\sum_i p_{T,i}} \implies \alpha = \epsilon \frac{\sum_i p_{T,i}}{\sum_i p_{T,i}^{\mathfrak{n}+1}}. \quad (1.31)$$

Any choice of $\mathfrak{n}$ with the above value of α will exactly reproduce the fractional calo loss observed in the data. The value of ϵ is not exactly the the values shown in Fig. 1.62, which are the absolute difference in fractional loss $\mathfrak{x}\%$:

$$\text{fractional loss (MC)} - \text{fractional loss (data)} = \mathfrak{x}\%, \quad (1.32)$$

where $\mathfrak{x} \sim 1\%$ for $p_T < 1.2$ TeV and $\mathfrak{x} \sim 2\%$ for $p_T > 1.2$ TeV. Let $\max_i = \max_{\text{jet } p_T} \text{CTER}(\text{jet } p_T')$ for $i \in \{\text{MC}, \text{data}\}$. Then the relationship between ϵ and $\mathfrak{x}$ is given by:

$$\begin{aligned} \epsilon &= \mathfrak{x}\% \times \frac{\max_{\text{MC}}}{\text{CTER}_{\text{MC}}} + \frac{\max_{\text{MC}} - \text{CTER}_{\text{MC}}}{\text{CTER}_{\text{MC}}} - \frac{\max_{\text{MC}}}{\text{CTER}_{\text{MC}}} \left(\frac{\max_{\text{data}} - \text{data}_{\text{MC}}}{\max_{\text{data}}} \right) \\ &\sim \mathfrak{x}\% \times \frac{\max_{\text{MC}}}{\text{CTER}_{\text{MC}}}. \end{aligned}$$

Thus, the value of ϵ is about 1% for $p_T < 1.2$ GeV and about 3% for $p_T > 1.2$ TeV. The closure for inefficiency parameterizations for $\mathfrak{n} = 0, 1, 2, 3, 10 \sim \infty$ are shown in Fig. 1.63. For all values of $\mathfrak{n}$, the fractional calo loss is the same as the data (by construcdtion) and higher than the nominal simulation.

The next step is to justify the assumption that the tracking inefficiencies can be treated as independent. In general, the tracking inefficiency will depend on the other tracks present in the jet, measured and unmeasured. However, the CTER and the average jet charge are special variables which is not sensitive to these effects. First, for the CTER:




Figure 1.63: The fractional calo loss for the nominal simulation, the data, and various parameterizations of the tracking inefficiency uncertainty.

$$\text{CTER} = \left\langle \frac{\sum_{i \in \text{Tr}} p_{T,i}}{p_T^{\text{jet}}} \right\rangle = \frac{\langle \sum_{i \in \text{Tr}} p_{T,i} \rangle}{p_T^{\text{jet}}} = \frac{\sum_{j \in J} w_j \sum_{i \in \text{Tr}} p_{T,i}}{p_T^{\text{jet}} \sum_{j \in J} w_j} \quad (1.33)$$

$$= \frac{\sum_{p_T} p_T n_{p_T}}{p_T^{\text{jet}} \sum_{j \in J} w_j}, \quad (1.34)$$

where J is the set of all jets and w_j is the event weight associated with jet j . The first equality is true in a fixed jet p_T bin and the last equality is true by exchanging the finite sums. The quantity $n_{p_T} = \sum_{j \in J} w_j n_{p_T,j}$, where $n_{p_T,j}$ is the number of tracks in jet j with transverse momentum p_T . Equation 1.33 shows that the CTER only depends on the total number of tracks of a given p_T in a particular jet p_T bin and the total weighted number of jets in all of the jet p_T bins. This means that a sufficient statistic for the CTER is the joint distribution of track p_T and jet p_T , i.e. one does not need know the distribution of track p_T inside each individual jet. As

a closure of the method, one can compute the tracking efficiency as the ratio of the detector-level and particle-level jet and track p_T joint distributions. The closure test is then to correct each track and see if the detector-level CTER after the track-by-track correction is given by the particle level CTER²⁰. The efficiencies are shown in Fig. 1.64. The left plot of Fig. 1.64 is the track reconstruction efficiency with the impact of jet and track resolutions removed, while the right plot of Fig. 1.64 is what is applied in practice. Due to resolution effects, the ‘efficiency’ can exceed one in the right plot of Fig. 1.64 while the z-axis is between 0 and 1 by definition in the left plot of Fig. 1.64. The track-by-track correction is applied by replacing $\sum p_{T,i}/p_T^{\text{jet}}$ for a given jet by $\sum (p_{T,i}/e(p_{T,i}, p_T^{\text{jet}}))/p_T^{\text{jet}}$, where e is the tracking efficiency, including resolution effects. Figure 1.65 shows the efficiencies as a function of ΔR between the track and the jet axis for various p_T bins. It is clear that the inefficiency is larger at lower ΔR (in the jet core) for higher p_T jets.

The efficiency corrected distribution of $\sum p_{T,i}/p_T^{\text{jet}}$ in two jet p_T bins (before computing CTER as the average) is shown in Fig. 1.66. Note that the corrected distributions can be larger than one in order to get the correct average (while the particle-level distributions never exceed one by construction).

The actual closure is seen in Fig. 1.67. Circles show the particle-level distribution of the CTER, which as already discussed is flat and nearly 2/3. Triangles and diamonds show various detector-level versions of the CTER, with(out) fakes, secondaries, and the track resolution. The crosses are the corrected detector-level CTER values, which nicely fall on top of the circles. A slight non-closure in the lowest bins is due to the finite binning of the 2D track and jet p_T distributions. The various other curves in Fig. 1.67 show the CTER computing using a subset of tracks, indicated by the requirements in the legend.

Now that we have shown the method works for the CTER, we quickly prove that the same principle holds for the average jet charge:

²⁰The tracking efficiency depends also on η , but since the jet charge is not measured as a function of η , all distributions are marginalized over η .




Figure 1.64: Left: The ratio of the number of charged particles matched to reconstructed tracks divided by the total number of charged particles inside particle-level jets as a function of jet p_T and charged particle p_T . By construction, this efficiency is between 0 and 1. Right: A similar ratio, but the numerator is replaced with all reconstructed tracks. Due to resolution effects, the ‘efficiency’ in the right plot can exceed unity.

$$\left\langle \frac{\sum_{i \in \text{Tr}} q_i p_{T,i}^k}{(p_T^{\text{jet}})^k} \right\rangle = \frac{\langle \sum_{i \in \text{Tr}} q_i p_{T,i}^k \rangle}{(p_T^{\text{jet}})^k} = \frac{\sum_{j \in J} w_j \sum_{i \in \text{Tr}} q_i p_{T,i}^k}{(p_T^{\text{jet}})^k \sum_{j \in J} w_j} \quad (1.35)$$

$$= \frac{\sum_{p_T} p_T^k n_{p_T}^+}{(p_T^{\text{jet}})^k \sum_{j \in J} w_j} - \frac{\sum_{p_T} p_T^k n_{p_T}^-}{(p_T^{\text{jet}})^k \sum_{j \in J} w_j}, \quad (1.36)$$

where as with the CTER, J is the set of all jets, w_j is the event weight associated with jet j , the first inequality is true in a fixed jet p_T bin, and the last inequality is true by exchanging the finite sums. The quantity $n_{p_T}^\pm = \sum_{j \in J} w_j n_{p_T,j}^\pm$, where $n_{p_T,j}^\pm$ is the number of tracks in jet j with transverse momentum p_T and charge $\pm$. So as with the CTER, the average jet charge does not depend on the correlations between the tracking (in)efficiencies of all the constituent tracks²¹.

²¹It is straight-forward to show that the jet charge distribution standard deviation does not share




Figure 1.65: The distribution of tracks (left) and tracking efficiency (right) as a function of ΔR between the track and the jet axis for various p_T bins and inclusive in track $p_T > 500 \text{ MeV}$.




Figure 1.66: The (un)corrected distribution of $\Sigma p_{T,i}/p_T^{\text{jet}}$ in two jet p_T bins (before computing CTER as the average) for low p_T jets (left) and high p_T jets (right).

the property of the CTER and the average jet charge - it depends on the first conditional distribution, i.e. depends on every pair (not every track in isolation). However, the uncertainties on the jet charge




Figure 1.67: The CTER for various particle-level and detector-level definitions (see text for details).

The last step is to pick a value of $\mathbf{n}$. Representative plots showing the uncertainty for various choices of $\mathbf{n}$ are shown in Fig. 1.68 and Fig. 1.69. The most conservative procedure seems to be $\mathbf{n} = 10 \sim \infty$, which is used for the final prescription.

One final note: the uncertainty described in this section is not a generic tracking-in-dense-environments uncertainty. It can only be applied out-of-the-box to quantities which have the same properties as the CTER and the average jet charge. It may be possible to use this method in the future to constrain a more general tracking-in-dense-environments uncertainty.

standard deviation are much smaller than the average so this subtlety is not considered further and the same prescription for the average is applied for the standard deviation.




Figure 1.68: The uncertainty on the average jet charge for the more forward (left) and more central (right) jet with $\kappa = 0.3$ (top), $\kappa = 0.5$ (middle) and $\kappa = 0.7$ (bottom). The uncertainty from tracking in dense environments is assumed to be negligible below 400 GeV, where nuclear interactions are the dominant source of inefficiency.




Figure 1.69: The uncertainty on the jet charge distribution standard deviation for the more forward (left) and more central (right) jet with $\kappa = 0.3$ (top), $\kappa = 0.5$ (middle) and $\kappa = 0.7$ (bottom).




Figure 1.70: The fraction of the sum track p_T that is due to secondary tracks.

1.5.3.4 The ζ Method

The charged-energy fraction method in Sec. 1.5.3.3 to determine the systematic uncertainty of charged particle losses inside the core of jets is broadly applicable to observables of the form $\sum_i f_i p_{T,i}^k$, where f_i contains information about the track i that does not depend on $p_{T,i}$. This section presents an alternative method²² based on the asymmetry of the pixels in the inner detector that could be applied for any observable. The setup is outlined in Fig. 1.71. The planar sensors in the pixel detector are about $50 \mu\text{m}$ in the ϕ direction and $400 \mu\text{m}$ in the z direction. The outermost pixel layer is at 122.5 mm from the center of ATLAS and so each pixel in that layer covers about $2\pi/[(\pi \times 122.5 \text{ mm})/40 \mu\text{m}] \sim 6.5 \times 10^{-4}$ radians in the ϕ direction. Consider two pairs of particles where $\Delta\phi_1 = \Delta\eta_2 = 0$ and $\Delta\eta_1$ and $\Delta\phi_2$ have the same distribution. Due to the asymmetry in the pixel dimensions, hit merging will occur for higher values of ΔR in the first pair with respect to the second pair. More generally, define $\zeta = |\text{atan}(\Delta\phi/\Delta\eta)|$. Assuming that the distribution of radiation is the same in the ϕ and η directions, if the pixel dimensions were symmetric, $\zeta \sim \text{Uniform}(0, \pi)$. However, due to the asymmetry and hit merging due to the high density environment at high p_T , the ζ distribution is not uniform.




Figure 1.71: A schematic diagram of the jet core, $R \lesssim 0.02$. The pixel is drawn approximately to scale assuming it is from the third pixel layer in the central region.

²²This method was first introduced by M. Begel, I. Hinchliffe, H. Ma, F. Paige, and M. Shapiro.

Before studying the distribution of ζ in the data, it is useful to estimate the predicted dependence of ζ on the loss. Since the track density is highest in the jet core and the most important impact is on high p_T tracks, the rest of the section constructs ζ only using tracks with $p_T > 10$ GeV and with $\Delta R < 0.02$ to the calorimeter jet axis. Define ϕ_0 to be the characteristic length of a pixel in radians in the ϕ direction and let r be the pixel aspect ratio between the η and ϕ dimensions of the pixel ($r \approx 8$). Consider the probability that the track from a particle at exactly the jet center merges with another track with $p_T > 10$ GeV. Assume that $\eta, \phi \sim \text{Uniform}$ in the jet core. Then, if there are n tracks with $p_T > 10$ GeV in the jet core,

$$\Pr(\text{merger}|\Delta\eta = 0) \sim 1 - \prod_{i=1}^n \Pr(\Delta\phi_i > \phi_0) \quad (1.37)$$

$$= 1 - \left(1 - \frac{\phi_0}{\Omega}\right)^n \sim \frac{n\phi_0}{\Omega} \equiv L(= \text{Loss}), \quad (1.38)$$

where $\Delta\phi$ is between the target track (at the jet center) and another track i . The size of the jet core is $\Omega \sim 0.02$. A similar calculation shows that $\Pr(\text{merger}|\Delta\phi = 0) \sim rL$. Therefore, the ratio of the distribution of ζ at $\zeta = 0$ versus $\zeta = 1$ is given by $(1 - rL)/(1 - L) \sim 1 + L(1 - r)$. In other words, the asymmetry of the ζ distribution depends on the total loss and the aspect ratio. When the aspect ratio is 1, there is no sensitivity to the loss. Since L increases with p_T , the asymmetry in the ζ distribution should also increase with p_T .

Figure 1.73 shows the joint distribution of ζ and the ΔR between tracks. Each event contributes multiple tracks to the histograms. While the distribution of the ΔR between tracks is qualitatively similar between the two distributions²³, but there is a clear difference in the distribution of ζ at low jet p_T and high jet p_T .

The p_T dependence of the ζ distribution is quantified in Fig. 1.74. The ζ distribution is nearly uniform for jet $p_T \lesssim 400$ GeV and the asymmetry in the distribution grows with p_T . Qualitatively this trend appears for both data and simulation. Representative

²³If the track locations were uniform in the jet core, one would expect a triangle probability distribution for their ΔR , which resembles the distribution in Fig. 1.73 projected onto the ΔR axis.




Figure 1.72: A schematic diagram of the distribution of ζ , accounting for merging.




Figure 1.73: The joint distribution of ζ and the ΔR between tracks for jet $p_T \sim 400$ GeV (left) and $p_T \sim 1.2$ TeV (right).

slices from Fig. 1.74 are shown in Fig. 1.76, with the simulation from PYTHIA 8 and HERWIG++ overlaid for comparison. Since the probability distribution for ζ in a fixed p_T bin is approximately linear in ζ , $f(\zeta) \propto \zeta$, a useful statistic of the distribution is the slope, $\partial_\zeta f(\zeta)$. The heuristic argument above suggests that this slope is proportional to the loss L . The fitted slopes are shown as a function of p_T

in Fig. 1.75. As expected from the p_T dependence of the loss, the slope increases with p_T . There is a small difference between data and simulation in the slope; the simulation seems to under-predict the loss, in agreement with the charged-energy fraction in Sec. 1.5.3.3. One way to estimate the relationship between the loss and $\partial_\zeta f(\zeta)$ in simulation is to decrease²⁴ the loss in simulation by adding truth particles without a reconstructed track to the jet. For a direct comparison, all reconstructed tracks are also replaced with their matched truth particles in order to remove the effect of the detector response. The difference between the solid squares and open circles in Fig. 1.75 quantifies the impact on the slope when removing these detector distortions. The other markers in Fig. 1.75 represent various levels of reduced loss (100% loss means 100% of the loss in the nominal simulation, not 100% of tracks lost). Slopes of linear fits are given in the legend of Fig. 1.75 and suggest that the slope is quadratically dependent on the loss. A simple fit yields $\partial_\zeta f(\zeta) \sim 0.8 - 0.1L + 0.0035L^2$, where L is the loss in percent. Inverting this relationship and using the slopes in the left plot of Fig. 1.75 results in $L_{\text{data}}/L_{\text{MC}} \sim 10 - 15\%$, which is quantitatively similar to the values derived using the charged-energy loss method in Sec. 1.5.3.3. Note that the fit is required because it is non-trivial to simulate more loss, as is the case in data (Fig. 1.75 shows the pattern for a *reduced loss* only).

²⁴Increasing the loss would be more relevant for matching to the data, but is highly non-trivial because it needs to respect the pixel geometry.




Figure 1.74: The conditional distribution of ζ given jet p_T for data (right) and MC (left).




Figure 1.75: Left: the p_T dependence of $\partial_\zeta f(\zeta)$, for $f(\zeta)$ the probability distribution of ζ . Right: the impact on the ζ distribution from reducing the loss in simulation.




Figure 1.76: The distribution of ζ in bins of the jet p_T for the data as well as simulation. The top left bin corresponds to $140 \text{ GeV} < p_T < 200 \text{ GeV}$ and each subsequent plot going from left to right, top to bottom, increases this range by 120 GeV.

1.5.3.5 Track Momentum Resolution

The momentum resolution of isolated tracks has been well-measured in $J/\psi \rightarrow \mu\mu$ and $Z \rightarrow \mu\mu$ events [131]. In addition to applying this method to muon tracks combining MS and ID information, this technique has been applied to ID-only tracks and is therefore applicable for the jet charge. The scale and resolution of reconstructed muon candidates are shifted and smeared in the MC to account for differences between the data and the simulation for $m_{\mu\mu}$. As these corrections are not applied to generic tracks, the correction factors are taken here as the systematic uncertainty on the momentum resolution. The momentum resolution is parameterized as

$$\frac{\sigma(p_T)}{p_T} = \frac{r_0}{p_T} \oplus r_1 \oplus r_2 \cdot p_T, \quad (1.39)$$

where $\oplus$ means ‘add in quadrature.’ The first term accounts for fluctuations in the energy loss in the detector material, the second term captures effects due to multiple scattering, and the third term accounts for the intrinsic resolution caused by misalignment and the finite spatial resolution of hits. Unlike for muon spectrometer tracks, inner detector tracks do not traverse a significant amount of material and so r_0 and its uncertainty are neglected. The uncertainties on r_1 , r_2 and the momentum scale s are estimated by smearing every track according to

$$p_T^{\text{track}} \mapsto \frac{p_T^{\text{track}} + s \cdot p_T^{\text{track}}}{1 + \sigma(r_1) \cdot z_1 + \sigma(r_2) \cdot p_T^{\text{track}} \cdot z_2}, \quad (1.40)$$

where z_i are independent random variables that are normally distributed with mean zero and standard deviation 1. The values of r_i and s as a function of η are shown in Table 1.8. A graphical representation of the uncertainties is shown in Fig. 1.77. The impact of this uncertainty is negligible for $p_T^{\text{track}} < 100$ GeV, but is significant for $p_T^{\text{tracks}} \sim 1$ TeV.

	$\sigma(r_1)$	$\sigma(r_2)$ [1/TeV]	s
$ \eta < 1.05$	0.0068	0.146	-0.92×10^{-3}
$1.05 < \eta < 2.0$	0.0105	0.302	-0.86×10^{-3}
$ \eta > 2.0$	0.0069	0.088	-0.49×10^{-3}

Table 1.8: A summary of the momentum (scale and) resolution uncertainties, taken from Ref. [131].




Figure 1.77: For an ensemble of tracks of the same p_T and η , the standard deviation of the distribution of the track p_T divided by the track p_T by applying the smearing procedure described in the text for the track momentum resolution uncertainty.

1.5.3.6 Track Charge Identification

Aside from the track p_T , the other track parameter that is relevant for the jet charge is the track charge. Especially at high p_T when the tracks are nearly straight, the probability for mis-identifying the track charge increases. The left plot of Fig. 1.78 shows the simulation probability for the charge mis-identification as a function of the jet p_T . The truth charge is the electric charge of the matched truth particle. At low jet p_T , the charge mis-id rate is less than 0.01% and even in the highest jet p_T bin, the charge mis-id rate is less than 1%. There is a small increase in the mis-id rate as a function of η , shown in the right plot of Fig. 1.78. The strongest dependence of the mis-id rate is on the track p_T , as shown in Fig. 1.79. The mis-id rate does not depend strongly on the jet p_T given the track p_T . Dedicated charge flipping studies in the data using leptonically decaying resonances suggest that the mis-modelling of the mis-id rate is much less than 50% [238]. Therefore, the charge mis-id uncertainty is conservatively estimated by randomly flipping the charge of tracks at 50% of the mis-id rate. The rate extracted from the simulation (Fig. 1.79) is $< 0.1\%$ for track $p_T < 100$ GeV, 0.5% for $100 \text{ GeV} < p_T < 200$ GeV, 1% for $200 \text{ GeV} < p_T < 300$ GeV, 2% for $300 \text{ GeV} < p_T < 400$ GeV and 4% for $p_T > 400$ GeV²⁵. The rate shown in Fig. 1.78 is likely to be very conservative. This is because tracks with a truth matching probability of $> 50\%$ that are actually fake (see Sec. 1.5.3.7) will have the wrong charge $\sim 50\%$ of the time as there is no relation between the track charge and truth particle charge. As a result, increasing the charge flipping rate has a larger impact (though still negligible) on the jet charge than reducing the mis-id rate. This is because increasing the mis-id rate mostly impacts correctly classified real tracks while decreasing the mis-id rate effects mostly mis-classified fake tracks. In other words, a large fraction of the tracks that are classified with a charge mis-id are likely fake tracks, while the majority of tracks with a truth matching probability of $> 50\%$ are not fake. Figure 1.80 supports the claim that many of the tracks with charge mis-id are actually fake. By construction, the tracks have a truth matching probability of $> 50\%$, but the distribution of probabilities is not as strongly peaked at one as for

²⁵Note that even in the highest p_T jets that pass the event selection, there are very few with tracks that have $p_T > 400$ GeV.

tracks with the correct charge. Furthermore, many tracks with a misclassified charge have a significantly different p_T than the truth-matched particle.




Figure 1.78: The charge mis-id rate as a function of the jet p_T (left) and $|\eta|$ (right).




Figure 1.79: The charge mis-id rate as a function of the track p_T . The charge mis-id rate increases rapidly for the lowest p_T bin due to fake tracks that happen to have a truth matching probability of $> 50\%$ but have a mis-id rate of $\sim 50\%$ because the reconstructed charge is random.




Figure 1.80: The MC matching probability for tracks with a flip and without a flip (left) as well as the p_{T} difference for flipped and non-flipped tracks between the ‘true’ and reconstructed track p_{T} (right). The fact that the MC probability is in general lower for tracks with a flipped charge and the p_{T} is usually very different from the truth p_{T} indicates that many of these tracks are actually fake tracks.

1.5.3.7 Fake Tracks

Random combinations of hits in the detector can be combined together to form a reconstructed track. Tracks resulting in particular from multi-particle trajectories that have kinks can result in a large reconstructed track p_T . The joint distribution of the fake track p_T (truth matching less than 50%) and jet p_T is shown in the left plot of Fig. 1.81. Tracks with p_T larger than the jet p_T are most likely from fakes and can be used to study the fake rate in data. The right plot of Fig. 1.81 shows the distribution of track p_T in five jet p_T bins. The rate of high p_T tracks is generally higher in the data than in the simulation, but this is especially relevant beyond the dashed lines where the track p_T exceeds the jet p_T . One contribution to the excess of high p_T tracks is from an underestimation of fake tracks in the simulation. Figure 1.81 suggests that this excess for high p_T tracks is less than 50%.




Figure 1.81: Left: The joint distribution of the fake track p_T and the jet p_T . The dashed line indicates the $p_T^{\text{track}} = p_T^{\text{jet}}$ line. Right: the normalized distribution of the track p_T in five bins of jet p_T . The dashed horizontal lines indicate the jet p_T thresholds. The ratio only shows points for which the track p_T exceeds the jet p_T .

To conservatively estimate the impact of fake tracks on the jet charge, fake tracks are randomly removed at a rate that is $\pm 50\%$ of the rate in simulation. The fraction of

fake tracks inside jets, integrating over all track momenta, is shown in Fig. 1.82. The fake rate is largely independent of the jet p_T and is $\lesssim 0.1\%$.




Figure 1.82: The fraction of tracks inside a jet that are classified as fake in bins of jet p_T . The circle markers are for central $|\eta| < 0.75$ jets while the squares are for all tracks within the tracker acceptance. The markers for the squares are offset by 30 GeV, but the actual fake rate is computed with the same p_T binning as for the circles. The markers indicate the median of the fake fraction distribution in a given jet p_T bin and the error bar is the inter-quartile range.

The distribution of the fake track p_T conditioned on the jet p_T is shown in Fig. 1.83. As expected, the fake track p_T spectrum is largely independent of the jet p_T . There is a small dependence, especially in the lowest p_T bin, because the hit density and thus fake rate increase monotonically with jet p_T .




Figure 1.83: The p_T spectrum of fake tracks inclusively in η (left) and for $|\eta| < 0.75$ (right).

1.5.3.8 Charged Particle Multiplicity

The tracking uncertainties described so far take into account the resolution and efficiency of the reconstruction of charged-particle momenta. One last source of systematic uncertainty is the number of charged particles. The unfolding procedure uncertainty takes into account the uncertainty on the prior due to the charged-particle multiplicity, but the jet charge resolution also changes with the charged-particle multiplicity. To assess the impact on the response matrix of the mismodeled charged-particle multiplicity, the distribution of n_{track} is reweighted in the simulation to match data per jet p_T bin and the relative difference when unfolding the nominal PYTHIA distribution with the reweighted PYTHIA distribution is taken as a systematic uncertainty²⁶. Fig. 1.84 shows the track multiplicity in three bins of jet p_T before any reweighting. These distributions will be the main focus of Chapter 3 and so are not discussed in more detail here.




Figure 1.84: The track multiplicity for jets in data and in Pythia for various p_T bins.

An example conditional distribution of the jet charge given the track multiplicity is shown in Fig. 1.85. When $n_{\text{track}} = 0$, the jet charge is zero by definition, resulting in a spike in the first row of Fig. 1.85. The resolution of the jet charge improves with n_{track} , which is why the higher rows in Fig. 1.85 have a jet charge distribution that is more peaked around the mean (see Fig. 1.31). The re-weighted jet charge

²⁶Since the prior is also changed, this uncertainty at least partially includes the unfolding procedure uncertainty.




Figure 1.85: The conditional distribution of the jet charge distribution given n_{track} . This is for the more forward jet and $\kappa = 0.5$.

distributions are shown in Fig. 1.86 as a function of jet p_T . Interestingly, the jet charge distribution average and standard deviation of the re-weighted PYTHIA distribution agree well with the data. This indicates that a significant contribution to the mis-modeling is from the track multiplicity. However, the full explanation must be more complicated because the re-weighted HERWIG++ is worse than the un-weighted simulation, though the effect is not as significant as with PYTHIA.

The uncertainty associated with the n_{track} re-weighting is shown in Fig. 1.87. This uncertainty is much smaller than other uncertainties for the standard deviation across p_T and also for the jet charge mean at low to moderate jet p_T . For the mean jet charge, the largest uncertainty is with the smallest κ and for large p_T , where it is 3–4% percent in the highest p_T bin for $\kappa = 0.3$ and $\kappa = 0.5$.

In principle, the uncertainty on the track multiplicity is actually part of a larger uncertainty on the full fragmentation. The remainder of this section explores the impact of the track p_T spectrum on the jet charge. The method non-closure includes some aspects of the full fragmentation mis-modeling, but there may be additional sources of uncertainty from variations in the response matrix due to differences in the track p_T spectrum. The left plot of Fig. 1.88 shows the track p_T spectrum inside jets with $200 \text{ GeV} < p_T < 300 \text{ GeV}$. Similar to the track multiplicity, the track p_T




Figure 1.86: The impact of n_{track} re-weighting on the jet charge distribution average (left) and standard deviation (right) as a function of jet p_T for the more forward jet and $\kappa = 0.5$.




Figure 1.87: The uncertainty on the jet charge distribution average (left) and standard deviation (right) due to the n_{track} modeling. These plots are for the more forward jet and $\kappa = 0.5$.

distribution from PYTHIA and HERWIG bracket the data. A re-weighting procedure analogous to the n_{track} reweighting is used to assess the impact of the mis-modeling. The right plot of Fig. 1.88 shows the conditional distribution of the jet charge given the track p_T . Every track contributes to the right plot of Fig. 1.88 and since each jet has many tracks, each jet contributes many times. The fork in the right plot

of Fig. 1.88 is due to events with one or a few tracks that carry a significant energy fraction and therefore the jet charge sign is set by the track charge. Figure 1.89 shows the uncertainty due to the track p_T after re-weighting to the data. In all bins, the uncertainty is $< 1\%$ and in most bins $\ll 1\%$; therefore it is ignored for the remainder of the analysis.




Figure 1.88: The track p_T spectrum in a particular jet p_T bin (left) and the conditional distribution of the jet charge given the track p_T (right); see the text for details. This is for the more forward jet and $\kappa = 0.5$.




Figure 1.89: The uncertainty due to the track p_T on the jet charge distribution average (left) and standard deviation (right).

1.5.4 Unfolding Non-closure

A standard method [231] for evaluating the systematic uncertainty from the procedure is to re-weight the MC to the data and take the difference between the unfolded re-weighted reconstructed MC to the truth MC of the same generator. The re-weighted truth is a reasonable prior with which one can estimate the bias from the choice of prior in the unfolding method. Heuristically, let $f(\mathbf{d}, \mathbf{p}, \mathbf{R})$ be a function that takes as inputs three histograms (data $\mathbf{d}$, prior $\mathbf{p}$, and the response matrix $\mathbf{R}$) and outputs another histogram (the unfolding function). By construction, $\mathbf{p} = f(\mathbf{R}\mathbf{p}, \mathbf{p}, \mathbf{R})$. Pick $\mathbf{t}$ such that $\mathbf{R}\mathbf{t} \sim \mathbf{d}$. Then, the non-closure uncertainty is the difference between $f(\mathbf{R}\mathbf{t}, \mathbf{p}, \mathbf{R})$ and $\mathbf{t}$. The non-closure is a method uncertainty and not an uncertainty on the prior per se, as the same $\mathbf{p}$ is used for $\mathbf{t}$ and $f(\mathbf{R}\mathbf{t}, \mathbf{p}, \mathbf{R})$. The following is a more detailed and careful description of the non-closure uncertainty, taking note of the proper normalizations for the various histograms and matrices.

Define the following histograms; x_i will interchangeably mean the histogram x and also the content in the i^{th} bin of x :

$\mathbf{d}_i$: The measured spectrum. There are 150 total bins (10 p_T bins and 15 jet charge bins) so $i = 1, \dots, 150$.

$\mathbf{R}_{ij}$: The unnormalized response matrix; $\mathbf{R}_{ij}$ is the number of events in the simulation that fall in the reconstructed bin i and the truth bin j .

$\mathbf{t}_i$: $\mathbf{t}_i = \sum_j \mathbf{R}_{ji}$: the particle-level spectrum for events that pass both particle- and detector-level selections.

$\mathbf{r}_i$: $\mathbf{r}_i = \sum_j \mathbf{R}_{ij}$: the detector-level spectrum for events that pass both particle- and detector-level selections

$\tilde{\mathbf{R}}_{ij}$: The normalized version of $\mathbf{R}_{ij}$ (earlier, this was just called *the* response matrix): $\mathbf{r}_i = \sum_j \tilde{\mathbf{R}}_{ij} \mathbf{t}_j$. Explicitly, $\tilde{\mathbf{R}}_{ij} = \mathbf{R}_{ij} / \sum_{i'} \mathbf{R}_{i'j}$. The entries of $\tilde{\mathbf{R}}_{ij}$ are the conditional probability for a truth event in bin j to be reconstructed in bin i .

The re-weighting procedure can only be applied to simulation events which pass both

the particle-level and detector-level event selections and so the first step is to take the data and apply the fake factors bin-by-bin:

$$(1 - f_i) = \frac{\text{Pass both reconstructed and truth selections}}{\text{Pass the reconstructed selection}}, \quad (1.41)$$

where i is the bin number. Define $\tilde{\mathbf{d}}_i = (1 - f_i)\mathbf{d}_i$ to be the corrected data histogram. A reasonable prior $\tilde{\mathbf{t}}_i$ is one such that $\tilde{\mathbf{R}}_{ij}\tilde{\mathbf{t}}_j$ is very close to $\tilde{\mathbf{d}}_i$. Since $\tilde{\mathbf{R}}_{ij}$ is not too far from a diagonal matrix, one way of generating (an approximate) $\tilde{\mathbf{t}}_i$ is to use weights built from the reconstructed simulation: $\mathbf{w}_i = \tilde{\mathbf{d}}_i/r_i$. Define $\tilde{\mathbf{t}}_i = \mathbf{w}_i\mathbf{t}_i$. The left plot of Fig. 1.90 shows the distributions of $\mathbf{w}_i$. In order to reduce the sensitivity to statistical fluctuations in the data in generating the weights $\mathbf{w}_i$, the histogram of weights is smoothed before generating $\tilde{\mathbf{t}}_i$. A standard median smoothing procedure implemented in ROOT with 20 iterations [239] is used for this purpose. There is a clear low-frequency trend in the weight histogram that increases monotonically with the bin number and corresponds to the $\mathbf{p}_T$ spectrum, while the high-frequency trends reflect the fact that the width of the charge distribution in each $\mathbf{p}_T$ bin changes. The right plot of Fig. 1.90 shows that the weights $\mathbf{w}_i$ are effective at improving the data/MC agreement of $\tilde{\mathbf{r}}_i = \sum_j \tilde{\mathbf{R}}_{ij}\tilde{\mathbf{t}}_j$ with respect to $\mathbf{r}_i$. In general, the trend is that the data/MC is greatly improved in all but the highest $\mathbf{p}_T$ bins, where the data/MC was already very good to begin with. Figure 1.91 shows the actual non-closure uncertainty for the jet charge and the standard deviation of the jet charge distribution, compared to the raw data/MC differences in the reconstructed version of these quantities. Except in the first two bins where fractional uncertainties have little meaning due to the small value of the jet charge compared with the uncertainty, the non-closure uncertainty for the jet charge mean is significantly smaller than the raw difference between the data and simulation. This is also mostly true for the jet charge distribution standard deviation, but the raw differences are already much smaller. As a comparison, the impact of unfolding the PYTHIA simulation with a HERWIG++ response matrix is shown in Fig. 1.92. The size of the differences shown in Fig. 1.92 are approximately compatible with those in Fig. 1.91.




Figure 1.90: The distribution of the weights w_i used to re-weight the MC distribution for the non-closure test (left) and the Data/MC ratio with the re-weighted truth distribution (labeled after) $\tilde{t}_i$ (right) for the more forward jet with $\kappa = 0.5$.




Figure 1.91: The impact of the non-closure uncertainty on the jet charge mean (left) and the jet charge distribution standard deviation (right) for the more forward jet and $\kappa = 0.5$. Bins without a black and red point indicate that one of the two is smaller than 10^{-3} .




Figure 1.92: The fractional difference on the jet charge mean (left) and standard deviation (right) when unfolding the PYTHIA simulation with a HERWIG++ response matrix and vice versa.

1.6 Results

The data satisfying the event selection criteria described in Sec. 3.1 are unfolded according to the procedure in Sec. 1.4 and the average and standard deviation of the jet charge distribution are computed as a function of the jet p_T . These results, along with the systematic uncertainties detailed in Sec. 1.5, are discussed in Sec. 1.6.1. The PDF uncertainty and jet formation uncertainties in the theory predictions are compared to the unfolded data in Secs. 1.6.2 and 1.6.3, respectively. Using PDF information as input, the average charge per jet flavor is extracted in Sec. 1.6.5 and its p_T -dependence is studied in Sec. 1.6.6.

1.6.1 Unfolded Jet Charge Spectrum

The unfolded jet charge mean is shown as a function of the jet p_T in the top plots of Fig. 1.93 for $\kappa = 0.3, 0.5$ and 0.7 . The average charge increases with jet p_T due to the increase in up-flavor jets from PDF effects. The average charge increases from $0.01e$ at $p_T \sim 100$ GeV to $0.15e$ at $p_T \sim 1.5$ TeV. Systematic uncertainties are generally a few percent, except at low jet p_T where the fractional uncertainty is large because the average jet charge in the denominator is small, and at high p_T where the tracking uncertainties are not negligible. The first bin suffers from large statistical uncertainties (up to 170%), but for the higher p_T bins the systematic uncertainty is dominant, except at the highest p_T bin where statistical and systematic uncertainty are of similar size (about 7%). The jet charge distributions of the more forward and more central jet differ in shape, in particular at low p_T , due to the different shape of the up/down flavor fractions in those bins as shown in Fig. 1.11(b).

Analogous results for the standard deviation of the jet charge distribution are shown in the bottom plots of Fig. 1.93. Even though the standard deviation of the reconstructed jet charge distribution increases with jet p_T (Fig. 1.53), the particle-level value decreases and approaches an asymptote for $p_T \gtrsim 300$ GeV.




Figure 1.93: The measured average (standard deviation) of the jet charge distribution on top (bottom) in units of the positron charge as a function of the jet p_T for $\kappa = 0.3, 0.5$, and 0.7 for the more forward jet (left) and the more central jet (right). The crossed lines in the bars on the data indicate the systematic uncertainty and the full extent of the bars is the sum in quadrature of the statistical and systematic uncertainties. The solid continuous line is a smooth approximation to the PYTHIA prediction.

1.6.2 Sensitivity of PDF Modeling

Variations in the PDF set impact the relative flavor fractions and thus in turn change the jet charge distribution. Such changes do not vary much with κ , since the PDF impacts the jet charge distribution mostly through the flavor fractions. Figures 1.94 and 1.95 compare the unfolded distributions of the jet charge distribution's average and standard deviation with several PDF sets, with tuned predictions for PYTHIA for each PDF, and with the same AU2 family of tunes. The sampling of PDF sets results in a significant spread for the average jet charge, but has almost no effect on the standard deviation. CTEQ6L1 describes the data best, although the data/MC ratio has a stronger p_T dependence. In particular, the data/MC differences with CTEQ6L1 are up to 10% (15%) at moderate p_T for the more forward (central) jet. For high p_T , differences between data and simulation are less significant. NLO PDFs such as CT10 are consistently below the data by about 10%-15%.




Figure 1.94: The average jet charge ($\kappa = 0.5$) in units of the positron charge for (a) the more forward jet and (b) the more central jet compared with theory predictions due to various PDF sets. The crossed lines in the bars on the data indicate the statistical uncertainty and the full extent of the bars is the sum in quadrature of the statistical and systematic uncertainties.




Figure 1.95: The standard deviation of the jet charge ($\kappa = 0.5$) distribution in units of the positron charge for (a) the more forward jet and (b) the more central jet compared with theory predictions due to various PDF sets. The crossed lines in the bars on the data indicate the statistical uncertainty and the full extent of the bars is the sum in quadrature of the statistical and systematic uncertainties.

1.6.3 Sensitivity of QCD Models and Tunes

The measurements presented in Sec. 1.6.1 show that there are qualitative differences between the data and the MC simulations, and comparisons in Sec. 1.6.2 suggest that variations in the PDF set cannot fully explain the differences. Differences in Sec. 1.6.1 between PYTHIA and HERWIG++ suggest that some aspect of the modeling of fragmentation could lead to the observed differences between the simulation and the data. One possible source is the hadronization modeling, which differs between PYTHIA (Lund-string fragmentation) and HERWIG++ (cluster fragmentation). The modeling of final-state radiation (FSR) is expected to have an impact on the jet charge distribution because variations in the radiation lead to different energy flow around the initial parton and hence different fragmentation of the jet. The plots in Fig. 1.96 and Fig. 1.97 show the measured average jet charge and the jet charge distribution's standard deviation, respectively, for $\kappa = 0.3, 0.5$, and 0.7 , compared to various models

for a fixed PDF set (CTEQ6L1). In addition to PYTHIA 8 and HERWIG++ model predictions, Figs. 1.96 and 1.97 contain the predictions from PYTHIA 6 using the Perugia 2012 tune [212] and the radHi and radLo Perugia 2012 tune variations. These Perugia tune variations test the sensitivity to higher/lower amounts of initial- and final-state radiation (via the scaling of α_s), although only variations of the FSR are important for the jet charge distribution. For the mean jet charge, PYTHIA 6 with the P2012 radLo tune is very similar to PYTHIA 8 with the AU2 tune. The spread in the average jet charge due to the difference between the radHi and radLo tunes increases with κ , since suppression of soft radiation makes the jet charge distribution more sensitive to the modeling of the energy fraction of the leading emissions. For the jet charge distribution's standard deviation, the sensitivity to the α_s scaling is large at both high and low κ . However, the sensitivity is inverted: radHi gives a larger standard deviation for $\kappa = 0.3$, but a lower standard deviation for $\kappa = 0.7$. Other Perugia 2012 tunes have been studied, testing the sensitivity to color-reconnection and multiple parton interactions, but the differences in the jet charge distribution's mean and standard deviation are small. The Perugia 2012 tunes may not fully capture the spread in nonperturbative effects, which is also suggested by the increasing difference between PYTHIA 8 and HERWIG++ for decreasing κ .




Figure 1.96: The average of the jet charge distribution in units of the positron charge for (a) $\kappa = 0.3$, (b) 0.5, and (c) 0.7 comparing various QCD MC models and tunes for the more forward jet. The crossed lines in the bars on the data indicate the statistical uncertainty and the full extent of the bars is the sum in quadrature of the statistical and systematic uncertainties.




Figure 1.97: The standard deviation of the jet charge distribution in units of the positron charge for (a) $\kappa = 0.3$, (b) 0.5, and (c) 0.7 comparing various QCD MC models and tunes for the more forward jet. The crossed lines in the bars on the data indicate the statistical uncertainty and the full extent of the bars is the sum in quadrature of the statistical and systematic uncertainties.

1.6.4 Model comparison overview

Figures 1.98 and 1.99 show comparisons of the unfolded jet charge distribution's mean and standard deviation for different QCD simulations using LO and NLO PDF sets. The predictions using the CT10 NLO PDF set as shown in Fig. 1.98 are generally about 10% below the data. Consistent with the expectation that the PDF and (nearly collinear) fragmentation are responsible for the jet charge distribution's mean and standard deviation, there does not seem to be an effect from the POWHEG NLO matrix element. For the jet charge distribution's standard deviation and $\kappa = 0.3$, the data falls between PYTHIA (larger standard deviation) and HERWIG++ (smaller standard deviation), but this trend is less evident for larger κ values, suggesting a difference due to soft tracks. As seen in Sec. 1.6.2, comparisons with CTEQ6L1 show it be to a better model for the p_T -dependence of the mean jet charge than CT10. The analogous plots to Fig. 1.93 but using CTEQ6L1 instead of CT10 are shown in Fig. 1.99. Generally, there is agreement between the simulation and the data with only a $\lesssim 5\%$ difference in the lower p_T bins.




Figure 1.98: The measured average of the jet charge distribution (top), and the standard deviation (bottom), in units of the positron charge as a function of the jet p_T for $\kappa = 0.3, 0.5$, and 0.7 for the more forward jet (left) and the more central jet (right) using CT10 as the PDF set. The markers in the lower panel are artificially displaced horizontally to make distinguishing the three κ values easier. The POWHEG+PYTHIA curves are nearly on top of the PYTHIA curves. The crossed lines in the bars on the data indicate the systematic uncertainty and the full extent of the bars is the sum in quadrature of the statistical and systematic uncertainties.




Figure 1.99: The measured average of the jet charge distribution (top), and the standard deviation (bottom), in units of the positron charge as a function of the jet p_T for $\kappa = 0.3, 0.5$, and 0.7 for the more forward jet (left) and the more central jet (right) using CTEQ6L1 as the PDF set. The markers in the lower panel are artificially displaced horizontally to make distinguishing the three κ values easier. The crossed lines in the bars on the data indicate the systematic uncertainty and the full extent of the bars is the sum in quadrature of the statistical and systematic uncertainties.

1.6.5 The average up-quark and down-quark jet charges

In addition to understanding the trends in the jet charge distribution from PDFs, one can use PDFs to extract information about jets of a particular flavor. These *exclusive* interpretations rely on flavor-fraction information in PDFs and matrix element calculations to extract the jet charge distribution for particular jet (anti-)flavors in each p_T bin. The required nonperturbative information is summarized in Fig. 1.11(a). Jets with flavors other than up/down/anti-up/anti-down/gluon are not included in Fig. 1.11(a) and give a negligible contribution ($\lesssim 2\%$) in the highest p_T bins.

One way of extracting the up- and down-flavor average jet charges is to exploit the difference in flavor fractions shown in Fig. 1.11(a) between the more forward and the more central jets. Due to the p_T -balance requirement between the leading and subleading jet in the event selection, to a good approximation, the p_T spectrum is the same for the more forward and the more central jet. Assuming that the average jet charge of the sum of flavors that are not up/down/anti-up/anti-down is zero, in each bin i of p_T :

$$\begin{aligned}\langle Q_J^{\text{forward}} \rangle_i &= (f_{\text{up},i}^{\text{forward}} - f_{\text{anti-up},i}^{\text{forward}}) Q_i^{\text{up}} + (f_{\text{down},i}^{\text{forward}} - f_{\text{anti-down},i}^{\text{forward}}) Q_i^{\text{down}} \\ \langle Q_J^{\text{central}} \rangle_i &= (f_{\text{up},i}^{\text{central}} - f_{\text{anti-up},i}^{\text{central}}) Q_i^{\text{up}} + (f_{\text{down},i}^{\text{central}} - f_{\text{anti-down},i}^{\text{central}}) Q_i^{\text{down}},\end{aligned}\tag{1.42}$$

where Q_J is the jet charge from Eq. 4.33, $f_{y,i}^x$ is the fraction of flavor y in p_T bin i for the jet $x \in \{\text{more forward, more central}\}$ and Q_i^y is the average jet charge for such jets (average gluon jet charge is zero). The values $f_{y,i}^x$ are taken from simulation (PYTHIA with CT10 PDF and AU2 tune), which then allows an extraction of Q_i^y by solving the system of equations in Eq. 1.42. This extraction is performed separately in each p_T bin. The left plot of Fig. 1.100 shows the extracted up- and down-flavor jet charges in bins of jet p_T . At very high jet p_T , the absolute quark flavor fractions are large (Fig. 1.11), but the difference between the more forward and more central jets is small and the statistical uncertainty is large. At low jet p_T , the difference between the more forward and more central jets is large (Fig. 1.11), but the absolute quark flavor fraction is small and the statistical uncertainty is once again large because the

mean jet charge is close to zero. In the limit that the flavor fractions are identical for the more forward and more central jet, the equations become degenerate and it is not possible to simultaneously extract the average up- and down-flavor jet charges. The uncertainties on the flavor fractions and on the measured average jet charges are propagated through the solutions of Eq. 1.42. Generally, the uncertainty is larger for the down-flavor jets because the fraction of these jets is smaller than the fraction of up-flavor jets.

The right plot of Fig. 1.100 compares the extracted up quark and down quark jet charges. The central value of the up quark jet charge is slightly less than twice the down quark jet charge, though this is not significant beyond one standard deviation for $\kappa = 0.5$ and $\kappa = 0.7$ and just beyond one standard deviation for $\kappa = 0.3$.




Figure 1.100: The extracted value of up- and down-quark jet charges in units of the positron charge in bins of jet p_T for $\kappa = 0.3, 0.5$, and 0.7 . The error bars include statistical, experimental systematic, and CT10 PDF uncertainties added in quadrature. The thick part of the error bar indicates the PDF contribution to the total uncertainty and the horizontal line on each error bar indicates the contribution from the statistical uncertainty. The first two p_T bins in the left plot are excluded due to their very large uncertainties.

1.6.6 Up- and down-quark jet charge dependence on p_T

Using the methods of Sec. 1.6.5, one can examine the residual p_T -dependence of the average jet charge *after* accounting for PDF effects. The inclusive jet charge has been shown to increase with p_T due to a mixing of jet flavors and the following subsection investigates the p_T -dependence of a fixed jet flavor. Using the theoretical predictions from Sec. 1.1.3, this section describes how the p_T -dependence is extracted from the data. Since $c_\kappa \ll 1$ from Eq. 1.12, one can approximate a linear dependence on c_κ :

$$\langle Q_J \rangle(p_T) = \bar{Q}(1 + c_\kappa \ln(p_T/\bar{p}_T)) + \mathcal{O}(c_\kappa^2), \quad (1.43)$$

where $\bar{Q} = \langle Q_J \rangle(\bar{p}_T)$ for some fixed (but arbitrary) transverse momentum, $\bar{p}_T$. Therefore, for a fixed p_T bin i , the measured charge is given as a superposition of the average jet charge for various jet flavors:

$$\langle Q_i \rangle \approx \sum_f \beta_{f,i} \bar{Q}_f (1 + c_\kappa \ln(p_{T,i}/\bar{p}_T)), \quad (1.44)$$

where $\beta_{f,i}$ is the fraction of flavor f in bin i , $\bar{Q}_f$ is the average jet charge of flavor f and $\bar{p}_T$ is a fixed transverse momentum. Fitting the model in Eq. 1.44 directly to the data to extract $\bar{Q}_f$ is not practical because there are three parameters and only 10 p_T bins, some of which have very little sensitivity due to low fractions β or large uncertainties on $\langle Q_J \rangle$. One way around this is to extract $\bar{Q}_f$ in one fixed bin of transverse momentum (denoted $\bar{p}_T$) as described in Sec. 1.6.5. Then Eq. 1.44 is highly constrained, with only one parameter for which each other bin of p_T gives an estimate. The systematic uncertainties are propagated through the fit treated as fully correlated between bins and the statistical uncertainty is treated coherently by bootstrapping²⁷. A weighted average is performed across all p_T bins and for both the

²⁷Pseudo-datasets are generated by adding each event in the nominal dataset j times where j is a Poisson random variable with mean 1. Since events are coherently added, this respects the correlations in the statistical uncertainty for the more forward and central jet charges.

more forward and the more central jet. The procedure is summarized below:

1. In the bin $600 \text{ GeV} < p_T < 800 \text{ GeV}$, extract the values $\bar{Q}_{\text{up}}$ and $\bar{Q}_{\text{down}}$. These values can be seen in the fifth p_T bin of Fig. 1.100.
2. With $\bar{Q}_{\text{up}}$ and $\bar{Q}_{\text{down}}$ fixed, extract the scale violation parameter estimate $c_{\kappa,i}$ in each p_T bin i by solving

$$\langle Q_i \rangle_{\text{measured}} = \sum_f \beta_{f,i} \bar{Q}_f (1 + c_{\kappa,i} \ln(p_{T,i}/\bar{p}_T)) \quad (1.45)$$

where $\bar{p}_T = 700 \text{ GeV}$ is the bin center from the previous step.

3. Repeat the procedure for all systematic variations and for all bootstrap pseudo-datasets to arrive at estimates of the uncertainty $\sigma(c_{\kappa,i})$ for each p_T bin i . The bin in step 2 is fixed, but the value in the bin varies.
4. The central value for the extracted scale violation parameter is

$$c_{\kappa} = \left(\sum_i c_{\kappa,i} / \sigma(c_{\kappa,i}) \right) / \sum_i (1 / \sigma(c_{\kappa,i})).$$

5. The uncertainty $\sigma(c_{\kappa})$ is determined by repeating step (3) with the nominal values $c_{\kappa,i}$ replaced by their systematic varied versions or the bootstrap pseudo-data values for the statistical uncertainty estimate.

The results are presented in Fig. 1.101. The data support the prediction that $c_{\kappa} < 0$ and $\partial c_{\kappa} / \partial \kappa < 0$. Linear correlations between κ values can be determined using the bootstrapped datasets: about 0.9 between $c_{0,3}$ and $c_{0,5}$ as well as between $c_{0,5}$ and $c_{0,7}$, while the correlation is about 0.7 between $c_{0,3}$ and $c_{0,7}$. Thus, the three points are quite correlated, but there is additional information from considering more than one κ value.




Figure 1.101: The extracted values of the scale violation parameter c_κ from the data compared to theoretical calculations [172, 240]. The error bars include statistical, experimental systematic, and PDF uncertainties added in quadrature. The thick part of the error bar indicates the PDF contribution to the total uncertainty and the horizontal line on each error bar indicates the contribution from the statistical uncertainty (each shown without adding in quadrature any other source of uncertainty).

1.7 Summary

This chapter presents a measurement of the particle-level p_T -dependence of the jet charge distribution's mean and standard deviation in dijet events from 20.3 fb^{-1} of $\sqrt{s} = 8 \text{ TeV}$ pp collision data recorded by the ATLAS detector at the LHC. The measured jet charge distribution is unfolded to correct for the detector acceptance and resolution for direct comparison to particle-level models. Comparisons are made at particle level between the measured jet charge distribution and various PDF sets and models of jet formation. Simulations with Pythia 8 using the CTEQ6L1 PDF set describe the average jet charge of the more forward jet within about 5% and the more central jet within about 10%. The jet charge distribution's standard deviation is described within 2%. HERWIG++ shows a similarly good agreement for $p_T > 500 \text{ GeV}$ and $\kappa = 0.7$. However, the HERWIG++ predictions decrease systematically for both the average and the standard deviation for decreasing κ . Predictions with the CT10 NLO PDF are systematically below the data across jet p_T for the average jet charge and systematically above for the jet charge distribution's standard deviation. Taking the PDFs as inputs, the average up- and down-flavor jet charges are extracted as a function of p_T and are compared with predictions for scale violation. The data show that the average up- and down-quark jet charges decrease slightly with p_T and this decrease increases with κ , as predicted. The particle-level spectra are publicly available [33] for further interpretation and can serve as a benchmark for future measurements of the evolution of nonperturbative jet observables to validate QCD MC predictions and tune their free model parameters.

Chapter 2

Color flow

Due to the confining nature of the strong force, directly measuring the QCD interactions between quarks and gluons is not possible. The strength and direction of the strong force depends on the color charge of the particles involved. To a good approximation, the radiation pattern in QCD can be described through a color-connection picture, which consists of color strings connecting quarks and gluons of one color to quarks and gluons of the corresponding anti-color. An important question is whether there is evidence of these color connections (*color flow*) in the observable objects: color-neutral hadrons and the jets they form. The study of energy distributions inside and between jets in various topologies has a long history, dating back to the discovery of gluons in three-jet events at PETRA [241–244]. Color connections are still a poorly constrained QCD effect, which motivates the dedicated study presented in this chapter. If well understood, experiments can exploit color flow to aid Standard Model measurements and searches for physics beyond the SM.

One of the challenges in studying color flow is the selection of a final state with a known color composition. Color-singlet W bosons from $t\bar{t}$ events provide an excellent testing ground because these bosons have a known initial (colorless) state and such events can be selected with high purity. As a test that the color flow can be extracted from the observable final state, the data are compared to models with simulated W bosons that are color-charged or color-neutral.

2.1 Introduction

Information about the color connections of partons participating in the hard-scatter is embedded in the observable final state jets. This has been demonstrated by studying the energy distribution inside and between jets in events of various topologies. The first such measurement was by the JADE collaboration in 3-jet events at PETRA [245]. The JADE collaboration reported that the third leading jet in tri-jet events had a rather different shape than the leading or subleading jet in the same events or the two leading jets in dijet events. Among other properties, it was found that these third (gluon-like) jets had a broader distribution of energy and particle multiplicity as a function of distance from the jet axis compared to the other (quark-like) jets. Comparison with the models of the time suggested that this observation was in support of fragmentation along the color axes of the initiating partons. There are now a variety of three-jet studies aimed at investigating this phenomena of *color coherence* performed at PETRA [245, 246], PEP [247–250], Tevatron [251, 252], and LHC [253].

With center of mass energies large enough to produce on-shell W and Z bosons, the highest energy experiments since LEP have studied color flow directly from the decay of heavy color singlets. For instance, L3 [254] and DELPHI [255] studied hadronic diboson WW events in which the energy density between jets associated with the same W decay compared to the density between jets from different W decays was found to be sensitive to the color flow of models used to describe the data. Additional studies at LEP [256, 257] and the Tevatron [258] have used the known initial color state of the electroweak bosons to constrain models of color flow in hadronic final states.

All of the studies described thus far have used either the distribution of energy within a jet or the distribution of energy between jets as sensitive observables to constrain models of color flow. The combination of the orientation and distribution of intrajet and interjet radiation can provide additional discriminating power. First defined in Ref. [140], the *jet pull* is a kinematic variable built from momentum-weighted radial moments of jet constituents (*jet substructure*) combined with information from the relative orientations of jets in the event (*jet superstructure*) that was designed to

be sensitive to the color flow between the initiating partons of jets. Since its inception, jet pull has been suggested as a discriminating variable to isolate color singlets such as Higgs bosons from color octets (e.g. gluons) [140, 259] and has been used for this purpose experimentally in a variety of searches for the Higgs boson [260–262]. However, there has never been significant evidence from the data that this observable is directly sensitive to color flow.

The first experimental measurement of color connection using jet pull was performed in $t\bar{t}$ events with one lepton in the final state at $D\bar{O}$ [263]. Such events provide a relatively pure sample of hadronically decaying W bosons. By fitting the data with MC templates constructed from the jet pull distribution, exotic color flow models can be constrained directly. However, color flow is subtle and there was not sufficient precision at $D\bar{O}$ to observe significant differences between the singlet and octet models.

This chapter describes the first measurement to definitively show that the jet pull angle can differentiate color singlet and color octet dijet resonances¹. The jet pull angle is studied in $t\bar{t}$ at the LHC, where the $t\bar{t}$ cross section and integrated luminosity are much higher than at the Tevatron. In addition, improved analysis techniques have increased the precision of the measurement. Furthermore, the jet pull angle distribution is unfolded to correct for distortions from the detector resolution and finite acceptance in order to make the measurement available for MC tuning or testing models of color flow beyond the Standard Model.

This chapter is organized as follows. The remainder of Sec 2.1 describes color flow in the context of QCD and introduces the jet pull angle. Details about the analysis design, including the simulation, object reconstruction, and event selection are described in Sec. 2.2. The properties of the reconstruction and resolution of the jet pull are in Sec. 2.3 in preparation for unfolding the pull angle distribution, described in Sec. 2.4. A detailed description of the systematic uncertainties is documented in Sec. 2.5 and the unfolded results are given in Sec. 2.6. The chapter ends with some concluding remarks in Sec. 2.7.

¹The jet pull reconstruction studies and precision measurement presented here are published in Ref. [264] and Ref. [265], respectively. The measurement benefited from fruitful discussions and technical help from T. Neep, K. Joshi, M. Swiatlowski, Y. Peters, D. L. Mateos, and M. Schwartz.

2.1.1 Color flow in QCD

Color flow has important implications for all stages of jet formation. At the beginning of jet development, the leading effect is due to *color coherence*. A heuristic explanation [22, 266] for this effect is that soft (long wavelength) gluons cannot resolve individual partons that are close in angle. The same effect is true for photons. In electrodynamics, the rate of radiated photons in $\gamma \rightarrow e^+e^- \rightarrow e^+e^-\gamma$ will be suppressed outside of the e^+e^- opening angle. The soft photons cannot resolve the individual electrons; instead they are sensitive only to the sum, which is neutral. Symbolically, suppose that the positron has momentum $\mathbf{p}$ and the soft photon has momentum $z\mathbf{p}$ with $z \ll 1$, as in Fig. 2.1. The mass of the virtual positron is

$$m_{e^+} \sim 2zp(1-z)p(1-\cos\theta) \sim 2zp^2\theta_{e\gamma}^2. \quad (2.1)$$

By the uncertainty principle, the virtual e^+ can persist for a time $\Delta E \Delta t \sim 1^2$:

$$\frac{1}{\Delta t} \sim \sqrt{(m_{e^+})^2 + p^2} \sim p\sqrt{1 + 2z\theta^2} \sim pz\theta^2. \quad (2.2)$$

In this time, the electron and the positron have traveled a distance $\sim \theta_{ee}/pz\theta^2$. The wavelength of the soft photon in the direction away from the positron is $\lambda \sim 1/p_T^\gamma \sim 1/pz\theta$. A soft photons can resolve the separation between the e^- and e^+ if

$$\frac{1}{pz\theta} < \frac{\theta_{ee}}{pz\theta^2}, \quad (2.3)$$

which is the same as $\theta_{e\gamma} < \theta_{ee}$. In QCD, the impact is similar, only that the initial gluon in $g \rightarrow q\bar{q} \rightarrow q\bar{q}g$ is colored. For a parton with no color charge splitting into two quarks with opposite color, the impact of color coherence is the same as in the QED case. Large angle soft gluon radiation is suppressed because the gluons cannot

²This is a re-writing of the familiar relation $\Delta x \Delta p \gtrsim \hbar/2$. In natural units, one unit of angular momentum is $\hbar$ and the factor of 2 is absorbed in the $\sim$ sign.




Figure 2.1: An illustration of coherence. Photons emitted at large angles $\theta_{e\gamma}$ are not able to resolve the e^+e^- pair, leading to a suppression of radiation outside of θ_{ee} .

individually resolve the two quarks and instead are sensitive only to the color of the initial parton, which is zero. As a consequence of color coherence, the radiation pattern for jets initiated by two quarks originating from color singlets is enhanced in the interjet region relative to jets initiated by color triplet or color octets. Additionally, the radiation pattern for jets initiated two quarks resulting from a color singlets is suppressed in the intrajet region relative to jets initiated by color triplet or color octets. These radiation patterns are demonstrated in Fig. 2.2. The same high p_T Higgs boson undergoes fragmentation many times. By fixing the hard-scatter parton, the images in Fig. 2.2 show only the impact of the parton shower and hadronization on the distribution of radiation inside the jet. The two nodes in the image correspond to the initial quark locations. The radiation around the two nodes is enhanced in the right plot (octet) with respect to the left plot (singlet). Figure 2.3 shows a quantitative comparison between the two radiation patterns, where the enhancement between the nodes is also apparent for the singlet.

Color flow also has an impact on hadronization. While there is no first-principle perturbative description of hadronization, there are a variety of phenomenological QCD-inspired models. For example, in the popular string model, color connected quarks are bound by a linear confining potential that breaks into hadrons when the potential energy in the ‘string’ is sufficiently large. The resulting production of hadrons is enhanced between color connected partons.

The discussion so far has been in the limit $N_c = \infty$. For a finite number of colors, there are small effects due to *color reconnection*. These effects are suppressed by $1/N_c^2$, which is comparable to α_s .




Figure 2.2: The same high energy parton undergoes fragmentation many times (*re-showered*). Each time the parton fragments, the stable hadrons in the event are clustered into a $R = 1.0$ jet and trimmed using k_t subjets with $R_{\text{sub}} = 0.3$ which are removed using the parameter $f_{\text{cut}} = 0.05$. The high energy parton is a color singlet (left) or color octet (right) Higgs boson h with $m_h = 125$ GeV and $p_T = 500$ GeV. To ensure no other significant radiation in the event, momentum is conserved by balancing the h against a $Z(\rightarrow \nu\bar{\nu})$. The histograms are the average jet image (see Sec. 4.3.2) over all re-showers with p_T intensity and the L^2 norm.




Figure 2.3: The average difference between the left and right distributions in Fig. 2.2.

2.1.2 Jet Pull

One observable predicted to contain information about the color representation of a dijet resonance like the W , Z , or Higgs boson, is the *jet pull vector* [140]. The pull vector for a given jet J with transverse momentum, p_T^J , is defined as

$$\vec{v}_p^J = \sum_{i \in J} \frac{p_T^i |r_i|}{p_T^J} \vec{r}_i. \quad (2.4)$$

The sum in Eq. (2.4) runs over jet constituents with transverse momentum p_T^i and location $\vec{r}_i = (\Delta y_i, \Delta \phi_i)$, defined as the vector difference between the constituent and the jet axis (y_J, ϕ_J) in rapidity (y) - azimuthal angle (ϕ) space. Given the pull vector for jet J_1 , the angle formed between this pull vector and the vector connecting J_1 and another jet J_2 , $\vec{r}_{J_2}^{J_1} = (y_{J_2} - y_{J_1}, \phi_{J_2} - \phi_{J_1})$, is expected to be sensitive to the underlying color connections between the jets. This is shown graphically in Fig. 2.4, and the angle is called the *pull angle*, denoted $\theta_P(J_1, J_2)$. Symbolically:

$$\cos \theta_P(J_1, J_2) = (\vec{r}_{J_2}^{J_1} \cdot \vec{v}_p^{J_1}) / (|\vec{r}_{J_2}^{J_1}| |\vec{v}_p^{J_1}|). \quad (2.5)$$

The pull angle is symmetric around zero when it takes values between $-\pi$ and π and so henceforth $\theta_P(J_1, J_2)$ refers to the magnitude of the angle in $(\Delta y, \Delta \phi)$ space with $0 < \theta_P \leq \pi$. For jets originating from color-connected quarks, $\theta_P \sim 0$ since the radiation is predicted to fall mostly between the two jets. In other cases, θ_P need not be small, so the angle should be useful for determining color connections.

Due to the angular-weighting in Eq. 2.4, the contribution of large angle radiation is enhanced with respect to collinear radiation. For jets of size $R = 0.4$, this means that a jet constituent at the edge of a jet with transverse momentum p_0 contributes just as much to the jet pull vector as radiation in the jet core ($\Delta R \lesssim 0.02$) with $p_T \sim p_0(0.4/0.02)^2 \sim 400p_0$. This large radiation is sensitive to color flow, but it is also sensitive to the diffuse uncorrelated radiation in the event due to pileup and the underlying event. Therefore, the the momentum weighting in Eq. 2.4 is critical to suppress these contributions and formally render the jet pull infrared safe. The jet




Figure 2.4: The construction of the jet pull angle for jet J_1 with respect to J_2 .

pull is also collinear safe: for a particle $P_0 \rightarrow P_1 P_2$ collinearly, $\vec{r}_{P_0} = \vec{r}_{P_1} = \vec{r}_{P_2}$ and $\mathbf{p}_{T,P_0} = \mathbf{p}_{T,P_1} + \mathbf{p}_{T,P_2}$ so $\mathbf{p}_{T,P_1} |\mathbf{r}_{P_1}| \vec{r}_{P_1} + \mathbf{p}_{T,P_2} |\mathbf{r}_{P_2}| \vec{r}_{P_2} = \mathbf{p}_{T,P_0} |\mathbf{r}_{P_0}| \vec{r}_{P_0}$.

Any jet constituent can be used to construct the jet pull angle, several of which are used in this chapter. For detector-level jets, the natural constituent choice is the calorimeter-cell cluster. This results in the *calorimeter pull* angle. The analogous quantity for particle-level jets is the *all particles pull* angle using all jet constituents to compute Eq. 2.4. Charged particle tracks are not used explicitly in jet reconstruction, but offer superior angular resolution to the calorimeter-cell clusters. The *track pull* angle is built using tracks that are ghost-associated to the calorimeter jet. An analogous *charged-particles* pull angle is constructed for particle-level jets using only the electrically charged constituents. The jet axis definition is an implicit input to Eq. 2.4. The nominal jet four-vector is used for the particle-level pull angle axis. A discussion of the axis for detector-level pull angle is postponed until Sec. 2.3.6.1.

There are advantages and disadvantages for the track pull and likewise for the calorimeter pull. Tracks have a better angular precision than calorimeter clusters and so the charged pull is more precisely measured than the all particles pull angle. However, by using only charged particles, one is less sensitive to the underlying color flow and may also increase sensitivity to certain modeling uncertainties. Both constituent inputs are studied in parallel.

2.2 Analysis Design

The measurement presented in this chapter demonstrates the ability to extract the color charge of a dijet resonance using the jet pull variable. As part of this study, various performance aspects of the jet pull are studied in order to increase the precision of the measurement as well as to improve the tagging capabilities of the jet pull angle. Events enriched in top quark pair production provide a clean environment and a copious source of hadronically decaying W bosons, the model dijet resonance. The data are unfolded to correct for resolution and acceptance effects. These unfolded data are compared with various fragmentation models and can be used in the future to constrain models of jet formation. In order to quantify the sensitivity to the color charge of a dijet resonance, W bosons with an octet color charge are simulated and compared with the unfolded data.

The dataset and simulated samples are summarized in Sec. 2.2.1, along with a detailed discussion of the simulation for exotic color flow. Section 2.2.2 describes the object reconstruction and event selection.

2.2.1 Dataset and Simulation

The Monte Carlo simulation is similar to setup discussed in Sec 1.3.3.1. Table 2.1 contains a summary of the simulation setups used for each SM process. POWHEG-BOX+PYTHIA 6 is used for the nominal setup to model the $t\bar{t}$ process. Top quark pair production is derived from the data by subtracting all other background processes, but the $t\bar{t}$ simulation is used to construct the response matrix for unfolding the background-subtracted data. The additional $t\bar{t}$ samples listed in Table 2.1 are used as alternative models to assess systematic uncertainties.

Aside from the W +jets background and multijets backgrounds, all MC samples are normalized to their theoretical cross-sections, calculated to at least next-to-leading order (NLO) precision in QCD [267–272]. For the purpose of comparison between data and the SM prediction before unfolding, $t\bar{t}$ events are normalized to a cross-section of 253 ± 15 pb, calculated at next-to-next-to-leading order (NNLO) in QCD including next-to-next-to-leading logarithmic (NNLL) soft gluon terms [218],

Process	Generator	Type	Version	PDF	Tune
$t\bar{t}$	POWHEG-BOX [89,200,201] +PYTHIA 6 [96]	NLO ME + PS	- 6.426.2	CT10 [193,194] CTEQ6L1 [198]	- Perguia2011c [212]
Single top (t -, s -, and Wt -channels)	POWHEG-BOX +PYTHIA 6	NLO ME + PS	6.426.2	CT10(4f) CTEQ6L1	DR scheme (Wt) [273] Perguia2011c
WW, WZ, ZZ	SHERPA [3]	LO multi-leg ME + PS	1.4.1	CT10	Default
W/Z +jets	ALPGEN [222] +PYTHIA 6	LO multi-leg ME + PS	2.1.4 6.426.2	CTEQ6L1 CTEQ6L1	- Perguia2011c
$t\bar{t}^\dagger$	POWHEG-BOX +HERWIG [210] +JIMMY [211]	NLO ME + PS (MPI)	- 6.520.2 4.31	CT10 CT10 -	- AUET2 [274] -
$t\bar{t}^\dagger$	MC@NLO [88,275] +HERWIG +JIMMY	NLO ME + PS (MPI)	4.06 6.520.2 4.31	CT10 CT10 -	- AUET2 -

Table 2.1: Monte Carlo samples used in this analysis. The abbreviations ME, PS, PDF, MPI, LO and NLO respectively stand for matrix element, parton shower, parton distribution function, multiple parton interactions, leading order and next-to-leading order in QCD. Tune refers to the used set of tunable MC parameters. Those samples marked with a $\dagger$ are used as alternative $t\bar{t}$ samples to evaluate uncertainties due to the modeling of $t\bar{t}$ events. The 4-flavor scheme (4f) is used for CT10 only for the t -channel single top production.

assuming a top-quark mass of 172.5 GeV. The W +jets and multijet normalizations are described in Sec. 2.2.1.1.

Generated events are processed with a full ATLAS detector and trigger simulation [110] based on GEANT4 [108] and reconstructed using the same software as the experimental data. The effects of pileup are modelled by adding to the generated hard-scatter events multiple minimum-bias events simulated with PYTHIA 8.160 [189], the A2 set of tuned MC parameters (tune) [190] and the MSTW2008LO Parton Distribution Function (PDF) set [220]. The distribution of the number of interactions is then weighted to reflect the pileup distribution in the data.

2.2.1.1 Data-driven backgrounds

The event selection is described in Sec. 2.2.2, but the W +jets and QCD multijet backgrounds are normalized using general data-driven techniques [276]. In order for an event featuring only quarks and gluons to pass the event selection, a non-prompt lepton from a semi-leptonic heavy quark decay or a hadron faking a lepton must be reconstructed as a signal lepton. This background is estimated from data with the *matrix method*. Two data samples – tight and loose – are defined based on their observed lepton isolation, where all tight leptons are also loose leptons. Tight leptons are the ones used in the event selection, described in Sec. 2.2.2. The number of events passing the loose or tight selection can be decomposed as $N^{\text{loose}} = N_{\text{fake}}^{\text{loose}} + N_{\text{real}}^{\text{loose}}$ and $N^{\text{tight}} = N_{\text{fake}}^{\text{tight}} + N_{\text{real}}^{\text{tight}}$. Let ϵ_X be the probability for a $X \in \{\text{real}, \text{fake}\}$ event to pass the tight selection given that it passes the loose selection. Then, since all events that pass the tight selection also passed the loose selection, $N^{\text{tight}} = \epsilon_{\text{fake}} N_{\text{fake}}^{\text{loose}} + \epsilon_{\text{real}} N_{\text{real}}^{\text{loose}}$. Solving for number of fake events, $N_{\text{fake}}^{\text{tight}}$:

$$N_{\text{fake}}^{\text{tight}} = \frac{\epsilon_{\text{fake}}}{\epsilon_{\text{real}} - \epsilon_{\text{fake}}} \cdot (\epsilon_{\text{real}} N^{\text{loose}} - N^{\text{tight}}). \quad (2.6)$$

The probabilities ϵ_{real} and ϵ_{tight} are measured with auxiliary event selections that are enriched in real leptons ($Z \rightarrow l^+ l^-$) or fake leptons (low m_T and low E_T^{miss} ; see Sec. 2.2.2). Additional details can be found in Ref. [277]. Estimates from the matrix method are often reliable when the number of events is large, but there are known experimental and theoretical challenges³. In this chapter, the multijet background is small and the event selection is inclusive enough that these challenges can be ignored.

The W +jets process is normalized by exploiting the asymmetry of W^+ and W^- events produced at the LHC due to the charge asymmetric initial state. There are about 30% more W^+ events than W^- events and this ratio slightly increases with jet multiplicity beyond $n_{\text{jet}} > 0$ and the uncertainty on the theoretical uncertainty on

³For a discussion of some of the theoretical challenges, see Ref. [278]; these are of the same flavor as for the CR method, described in Sec. 2.2. Alternative methods similar to the matrix method exist (see e.g. the fake-factor method in Ref. [279]), but they have their own set of challenges.

the ratio is $\mathcal{O}(\%)$ [280]. Since the efficiency for passing the event selection and the distribution of the pull angle are nearly independent of the lepton charge, the total number of W +jets events in a given bin can be estimated by

$$N_{W^+} + N_{W^-} = \frac{N_{W^+}^{\text{MC}} + N_{W^-}^{\text{MC}}}{N_{W^+}^{\text{MC}} - N_{W^-}^{\text{MC}}} (N_{W^+}^{\text{Data}} - N_{W^-}^{\text{Data}}), \quad (2.7)$$

where other backgrounds do not contribute to the parenthetical term as they are charge symmetric. One large source of uncertainty on the W +jets yield is the heavy-flavor composition, as b -jets will be part of the event selection. While b -quarks are generated from $t\bar{t}$ production at tree-level, they are generated by higher order processes in W boson production. The charge asymmetry method is extended to include residual scale factors per W +jets flavor subprocess through charge-independent corrections for W + bb/cc , W + c , and W +light. The W + cc is grouped with W + bb and not with W + c because the two processes have a different charge asymmetry. See Ref. [281] for more detail. The overall scale factors are 1.3 ± 0.03 for the W + bb/cc component, 0.74 ± 0.04 for the W + c component and 0.96 ± 0.02 for the light component.

The charge asymmetry method is experimentally and theoretically robust, but the main drawback is the large sample size required for an acceptable precision. Figure 2.5 shows that almost 1000 events are required to reach a statistical uncertainty of 10%. In this chapter, the event selection is inclusive enough and the W +jets background is sufficiently small that this is unimportant. However, the search presented in Part III that probes extreme regions of phase space must use a different technique.




Figure 2.5: The total statistical uncertainty and the charge-asymmetry statistical uncertainty of a pure sample of W +jets events with 30% more W^+ than W^- .

2.2.1.2 Simulating Exotic Colorflow

To test the sensitivity of the jet pull angle to the singlet nature of the W boson, a simulated $t\bar{t}$ is generated with a color-octet W boson. Using the partons produced with POWHEG-BOX recorded in the Les Houches Accord format [282], the color flow is inverted such that one of the W decay daughters shares a color line with the b -quark and the other shares a line with the top quark, as demonstrated schematically in Fig. 2.6. This sample is referred to as color *flipped* in the rest of this chapter.




Figure 2.6: Diagram illustrating the color connections for the nominal sample with a color-singlet W (left) and the flipped sample with a color-octet W (right).

A sample event from a typical $t\bar{t}$ LHE file is shown in the top panel of Fig. 2.7. The LHE contains only the hard scatter information, including the incoming protons, the partons participating in the scattering, the outgoing top quarks, and the immediate top quark decay products. Some events in an LHE file from POWHEG-BOX will also contain an additional parton from ISR from real emission in the NLO matrix element. The second and third columns in Fig. 2.7 give some particle identification information, the fourth-sixth columns describe where in the event record the particles belong and the rest of the columns list the particles quantum numbers. The most relevant for this section is the column for the ‘color’ quantum number. All of the MC generators use a large N_c limit and assign a new color-anti color whenever a color charge is created; in this case the color is arbitrarily called 501. The first of the color columns is the color and the second column is the anti-color. Color cannot be created or destroyed so the total color must be the same at the end of the hard scatter (status 23) as it was in the beginning (status -21). The most relevant numbers in the example in

Fig. 2.7 are the color and anti-color of the quark and anti-quark decay products of the W^- : 502. Note that the W^- and the $\bar{t}$ are *removed from the event record*. This is one of the drawbacks from the color-flipping procedure⁴: the top quarks and W bosons participating in the color flipping must be removed⁵, otherwise the shower generator will identify an unphysical color flow (as W bosons are singlets in the SM). Just before the W is removed, its decay products are identified and their color strings are flipped with the $\bar{b}$ in the event. This is seen in the lower panel of Fig. 2.7, where the quark from the W has color 502 while the anti-quark has color 503 (the anti-color from the original W is with the $\bar{b}$). It is then critical that the particle numbers and mothers are re-aligned since the showering models all have many internal consistency checks. Some care must be taken for generators that allow for b quarks in the PDF. This is solved by checking that the parent of the b or $\bar{b}$ is not the proton.

The remainder of this section shows various distributions at particle-level with the nominal and flipped $t\bar{t}$ simulated samples. Both samples use PYTHIA 6 with the same settings for the fragmentation. To begin, Fig. 2.8 shows the invariant mass of the leading two non b -tagged jets. Like the pull vector, the dijet mass is sensitive to relatively soft wide angle radiation. The dijet mass squared is $m_{jj}^2 \approx \sum e_i e_j \theta^2$, where e_i is the energy of jet i and θ is the angle between jets i and j . Because of this angular weighting, the dijet mass distribution is distorted by inverting the color flow. In particular, the dijet invariant mass is slightly larger for the flipped sample due to the enhancement of radiation around and a reduction of radiation between the two jets with respect to the nominal sample. Therefore, the jet mass distribution could be used to constrain the color flow, along with the pull angle. However, this is not used in the subsequent analysis. This is because if a new resonance were discovered, its mass would not be known a priori and therefore would not contain any useful

⁴An alternative method is to directly simulate a color octet using a UFO in MadGraph. This setup was found to give qualitatively similar results for the pull angle distribution as the flipped sample. However, as it is LO, there are significant kinematic differences with compared to POWHEG-BOX.

⁵PYTHIA will produce radiation differently if the top quark is absent from the LHE file. This is because the virtuality of the top quark is preserved if present and otherwise, the radiation off of the dipole formed from the out-going b and its color-partner in the initial state can produce much larger radiation (as the invariant mass is easily $\mathcal{O}(\text{TeV})$). This may account for some of the changes in the mass distributions presented in this section. Thank you to Stefan Prestel for the explanation.

----- PYTHIA Event Listing (hard process) -----											
no	id	name	status	mothers	daughters	colors	px	py	pz		
0	90	(system)	-11	0	0	0	0	0.000	0.000	0.000	8000.000
1	2212	(p+)	-12	0	0	3	0	0.000	0.000	4000.000	4000.000
2	2212	(p+)	-12	0	0	4	0	0.000	0.000	-4000.000	4000.000
3	21	(g)	-21	1	0	5	8	504	503	0.000	94.198
4	21	(g)	-21	2	0	5	8	501	504	0.000	-466.748
5	6	(t)	-22	3	4	9	10	501	0	81.617	-59.344
6	-5	bbar	23	3	4	0	0	0	503	-0.916	-37.444
7	3	s	23	3	4	0	0	502	0	-54.108	10.412
8	-4	cbar	23	3	4	0	0	0	502	-26.593	86.377
9	24	(W+)	-22	5	0	11	12	0	0	86.539	-59.228
10	5	b	23	5	0	0	0	501	0	-4.922	-0.116
11	-13	mu+	23	9	0	0	0	0	0	33.200	0.577
12	14	nu mu	23	9	0	0	0	0	0	53.339	-59.805
Charge sum:				0.000	Momentum sum:		-0.000	0.000	-372.550	560.947	

----- PYTHIA Event Listing (hard process) -----											
no	id	name	status	mothers	daughters	colors	px	py	pz		
0	90	(system)	-11	0	0	0	0	0.000	0.000	0.000	8000.000
1	2212	(p+)	-12	0	0	3	0	0.000	0.000	4000.000	4000.000
2	2212	(p+)	-12	0	0	4	0	0.000	0.000	-4000.000	4000.000
3	21	(g)	-21	1	0	5	8	504	503	0.000	94.198
4	21	(g)	-21	2	0	5	8	501	504	0.000	-466.748
5	6	(t)	-22	3	4	9	10	501	0	81.617	-59.344
6	-5	bbar	23	3	4	0	0	0	503	-0.916	-37.444
7	3	s	23	3	4	0	0	502	0	-54.108	10.412
8	-4	cbar	23	3	4	0	0	0	503	-26.593	86.377
9	24	(W+)	-22	5	0	11	12	0	0	86.539	-59.228
10	5	b	23	5	0	0	0	501	0	-4.922	-0.116
11	-13	mu+	23	9	0	0	0	0	0	33.200	0.577
12	14	nu mu	23	9	0	0	0	0	0	53.339	-59.805
Charge sum:				0.000	Momentum sum:		-0.000	0.000	-372.550	560.947	

Figure 2.7: One LHE event for a nominal (top) and inverted (bottom) color flow.

information about the color flow. Furthermore, when using the jet pull for tagging jets originating from bosons with a known mass, it is likely that a mass requirement will already be applied. It is important to ensure that there is more information about the boson color flow aside from the jet mass distribution. This is tested by re-weighting the m_{jj} and ΔR (between the W daughter jets) distributions, as in the right plot of Fig. 2.8. Since they are determined mostly by the hard-scatter, the other kinematic properties of the jets are similar between the simulations. The individual η values of the non b -tagged jets are nearly identical, but there are small differences at low p_T for the momenta, shown in Fig. 2.10. This is due in part to the higher loss for octet (gluon) jets in the transition from partons to jets compared with triplets (quarks) [119]. The most important distribution is θ_P , shown in Fig. 2.11. There is a significant difference between the flipped and nominal samples, which is largely invariant to the mass reweighting. This is also true for $|v_P^J|$, which will be revisited in Sec. 2.4.1. For comparison, the charged-particles versions of θ_P and $|v_P^J|$ are shown in Fig. 2.13.




Figure 2.8: The invariant mass of the leading non b -jets with (right) and without (left) a reweighting to the invariant mass spectrum.




Figure 2.9: The ΔR between the non b -tagged jets with (right) and without (left) a reweighting to the invariant mass and ΔR spectra.




Figure 2.10: The p_T of the leading non b -tagged jet with (right) and without (left) a reweighting to the invariant mass and ΔR spectra.




Figure 2.11: The pull angle with (right) and without (left) a reweighting to the invariant mass and ΔR spectra.




Figure 2.12: The pull vector magnitude with (right) and without (left) a reweighting to the invariant mass and ΔR spectra.




Figure 2.13: The charged particle pull angle (left) and the charged pull vector magnitude (right). No reweighting is applied to the invariant mass spectrum.

2.2.2 Object Reconstruction and Event Selection

Jet reconstruction and charged particle track association is performed using the algorithms documented in Sec. 1.2.2.1. Clusters and jets are corrected to point toward the primary vertex, as motivated and detailed in Sec. 2.3.6.1. After this correction, the four-vector sum of clusters (treated as massless) is used for the all-particles pull angle jet axis implicit in Eq. 2.4. The axis for the charged-particles pull angle is constructed using the four-momentum sum of all the associated tracks treated as massless.

Aside from the W boson daughter jet selection, the event selection is identical to the one described in Sec 1.3.3.1. In particular, events are selected using single isolated electron and muon triggers and required to have exactly one reconstructed electron or muon with $p_T > 25$ GeV and $|\eta| < 2.5$. Furthermore, to enrich the selection in $t\bar{t}$ events, $E_T^{\text{miss}} > 20$ GeV and $E_T^{\text{miss}} + m_T > 60$ GeV. Events must have ≥ 4 jets with $p_T > 25$ GeV. At least two of these jets must be tagged using a 70% target efficiency. At least two jets must not be b -tagged; of these, the two leading- p_T jets with $|\eta| < 2.1$ are labelled as the jets from the hadronically decaying W boson, J_i with $p_T^{J_1} > p_T^{J_2}$. The b -tagged jets and the jets selected for the pull angle calculation are required to have $|\eta| < 2.1$ so that all constituents are within coverage of the inner detector used for tracking. The W daughter jet selection is described in more detail in Sec. 2.2.3.1. The event selection produces a sample that is expected to contain approximately 90% $t\bar{t}$ events. Table 2.2 shows the predicted composition compared to the data yield.

2.2.3 Particle-level Event Selection

Particle-level objects and a particle-level event selection are constructed to be as close as possible to the detector-level objects. The particle-level objects are the target when unfolding the data in order to make direct comparisons with various theoretical models and these objects are also used to study the reconstruction performance. Particle-level jets are constructed the same way as for the jet charge measurement, described in Sec. 1.2.2.2. The particle-level inputs to the all-particles pull angle are all of the charged and neutral particles clustered within particle-level jets. Only the charged

Process	Number of Events
$t\bar{t}$	95400 ± 14000
Wt-channel single top	2730 ± 600
s- and t-channel single top	150 ± 10
W+jets	3710 ± 120
Z+jets	560 ± 270
Dibosons	190 ± 40
Multijets	2500 ± 910
Total SM	105000 ± 14000
Data	102987

Table 2.2: Estimated composition of the selected event sample. The uncertainties are the sum in quadrature of the statistical uncertainties and either the uncertainties of the normalisation method (for the data driven W+jet and multi-jet estimates) or the uncertainties of the cross-section estimates.

particles clustered within the particle-level jets are used for the charged-particles pull angle. Particle-level electrons, muons, photons, and neutrinos are only considered if their parent in the MC ancestry is not a hadron or a tau which came from a hadron decay. Electrons and muons are *dressed* with photons by defined the lepton 4-vector as the sum of the electron or muon particle 4-vector and the sum of all photon 4-vectors within $\Delta R < 0.1$. Dressed leptons are a better approximation to the measured leptons than bare leptons because (nearly) collinear radiation cannot be resolved in the detector. The particle-level E_T^{miss} is the magnitude of the vector sum over all particle-level neutrinos. Note that particle-level electrons and photons assigned to electrons or muons through dressing are not used for jet clustering. A particle-level jet is **b**-tagged if a B hadron from the MC event record with $p_T > 5$ GeV is ghost-associated to the jet. Additional information about the particle-level object definitions can be found in Ref. [283].

The particle-level event selection is analogous to the detector-level selection described in Sec. 2.2.2 with detector-level objects replaced with particle-level objects. Exactly one electron or muon and at least four jets are required, each with $p_T > 25$

GeV and $|\eta| < 2.5$. The particle-level $E_T^{\text{miss}} > 20$ GeV and the sum of $E_T^{\text{miss}} + m_T > 60$ GeV. At least two of the selected jets are required to be identified as **b**-jets using the same definition as that found in Ref. [283]. As with the detector-level calculation of the pull angle, the two leading- p_T particle-level non **b**-jets with $|\eta| < 2.1$ are labelled as the jets from the hadronically decaying W boson. About 80% (70%) of the time, the (sub)leading particle-level jet is within $\Delta R < 0.4$ of the detector-level jet. Since this is not 100%, there is non-negligible contribution to the unfolding from combinatorics in addition to per-object resolutions (in Sec. 1.2.2.2, the leading jet contamination is $\sim 4\%$).

2.2.3.1 W Boson Candidate Selection

The W boson candidate is built from the two highest p_T jets that are not b -tagged. A common alternative W boson identification technique is to use the non b -tagged jets whose invariant mass is closest to the W boson mass. The reason for not using the invariant mass constraint is because the mass can bias the colorflow, as shown in Sec. 2.2.1.2. This section explores the identification efficiency of the the baseline method and the alternative mass-based method. As a result of fragmentation, it is not possible to uniquely associate jets with quarks and so one must *define* a metric for assessing the fidelity of the hadronic W boson reconstruction. A common scheme is to use ΔR between the selected jets and the W boson in simulation. However, such a scheme is not useful when the W boson is produced at low p_T and also removes most of the information about which jets contain the majority of the W boson energy. Let T be the set of all truth particles in a simulated event and used for particle-level jet clustering (stable hadrons). Define $\mathcal{I}(i) = 1$ if particle $i \in T$ is a descendent from a W boson and 0 otherwise. The function $\mathcal{I}$ is well-defined in a leading log parton shower where particle histories are recoverable from the succession of $1 \rightarrow 2$ splittings. Since the W boson is a color singlet, when no kinematic requirements are placed on the particles entering jet clustering, $\sum_{i \in T} \mathcal{I}(i) \mathbf{e}_i = \mathbf{e}_W$, where $\mathbf{e}_i$ is the energy of particle i and $\mathbf{e}_W$ is the energy of the hadronically decaying W boson. A useful metric for comparing W boson reconstruction algorithms is the fraction of the W boson energy contained in the selected jets, $f_W^{\text{jet}} = \sum_{i \in \text{jet}} \mathcal{I}(i) \mathbf{e}_i / \mathbf{e}_W$. Another useful quantity is the fraction of a jet's energy originated from the W boson, $f_{\text{jet}}^W = \sum_{i \in \text{jet}} \mathcal{I}(i) \mathbf{e}_i / \sum_{i \in \text{jet}} \mathbf{e}_i$. Ideally, $\sum_{i=1}^2 f_W^{\text{jet},i} \approx 1$ and each of the two jets is mostly built from W boson radiation, $f_{\text{jet},i}^W \approx 1$ for $i = 1, 2$. The left plot of Fig. 2.14 shows the distribution of $\sum_{i=1}^2 f_W^{\text{jet},i}$ for three methods: the baseline method, the alternative method, and the *best* method in which the two jets are chosen with the highest $f_W^{\text{jet},i}$. Events are simulated using POWHEG+PYTHIA 8 with a simple particle-level event selection that mimics the detector-level selection discussed in Sec. 2.2.2. In particular, jets are clustered with the anti- k_t $R = 0.4$ algorithm implemented in FASTJET using all stable particles from PYTHIA that are not leptons as input. Jets are considered if they have $p_T > 25$ GeV and are tagged as b -jets if there is a B -hadron from the event record within $\Delta R < 0.3$




Figure 2.14: Left: The fraction of the W boson energy carried by the jets selected with three methods: Best, Baseline (‘Base.’), and Alternative (‘Alt.’). The baseline method uses the leading non- b tagged jets, the alternative method uses an invariant mass constraint, and the best method uses truth information; see the text for details. Right: The difference in fractions between the various methods.

of the jet axis. Events are required to have at least two b -tagged jets and at least two non b -tagged jets. Figure 2.14 has several interesting features, including the spike at zero from selected jets resulting from ISR or other sources of non W jets and the fact that in the best case, the fraction of the W boson energy contained in the selected jets is often much less than unity. As expected, baseline method has a lower fraction than the alternative method. The right plot of Fig. 2.14 shows the difference between $\sum_{i=1}^2 f_W^{\text{jet},i}$ for the various methods. The spike at zero corresponds to cases in which the methods select the same jets. About 70% of the time, the baseline and best methods are the same and about 80% of the time, the best and alternative methods are identical. The alternative is typically better than the baseline, but not always, as indicated by the tail of the red histogram at positive values of the difference. Since the difference between the alternative and baseline methods is small, and the alternative method has a potential for bias, the baseline method is used exclusively for the rest of the chapter.

Figure 2.15 shows how the W energy fraction differs between the leading and




Figure 2.15: The difference in the W boson energy fractions between the leading (in p_T) and subleading jets, $f_{\text{jet},i}^W(\text{lead}) - f_{\text{jet},i}^W(\text{sublead})$.

sub-leading W candidate daughter jets. As expected, the higher p_T jet has more of the W boson energy on average compared to the lower p_T jet. Many of the two-jet pairs have a roughly symmetric fraction of the W boson energy, but the width of the distribution in Fig. 2.15 is broad compared to the range. While it is possible to form two jet pull angles $\theta_P(J_1, J_2)$ and $\theta_P(J_2, J_1)$, only the former is measured in part because it contains more information (radiation) from the hadronic W boson decay⁶.

More visualizations of the distribution of the W boson energy inside the selected jets are shown in Fig. 2.16. The top plots of Fig. 2.16 show the two-dimensional distribution of f_W^{jet} and f_{jet}^W for the leading (left) and sub-leading (right) jets. When some fraction of the jet energy is from the W boson, the fraction is nearly 100%. Consistent with Fig. 2.15, the top right plot of Fig. 2.16 is shifted to the left with respect to the top left plot by construction. The lower left plot in Fig. 2.16 shows that many events are along the diagonal, where most of the W energy is captured by the two selected jets. However, there is a large spread in the bulk where more than two jets are need to capture the full W energy. The lines at 0 correspond to ISR jets

⁶As will be discussed in Sec. 2.3.5.1, the two pull angles are largely uncorrelated so a statistical combination would improve the measurement. However, the systematic uncertainties are fully correlated and at present, the measurement is limited by systematic and not statistical uncertainty.




Figure 2.16: Top: The distribution of f_W^{jet} (Jet energy fraction of W) and f_{jet}^W (W energy fraction of jet) for the leading (left) and sub-leading (right) jets. Bottom: the distribution between $f_W^{\text{jet},1}$ and $f_W^{\text{jet},2}$ (left) and between $f_{\text{jet},1}^W$ and $f_{\text{jet},2}^W$ (right).

which have nothing to do with the W boson. The lower right plot in Fig. 2.16 shows that the energy fraction of a selected jet is either dominated by the W energy or is nearly zero. The fraction of events in which at least one jet has a W energy fraction above 80% is 90%, while the fraction of events in which the leading jet has at least 80% of its energy from the W is 62%.

In conclusion, the W boson often produces more than two jets and one must be careful when assessing the performance of any matching scheme.

2.2.3.2 Comparisons Between Data and Simulation

This section briefly describes the modeling of important kinematic distributions related to the event selection described in Sec. 2.2.2. Figure 2.17 shows the individual p_T of the selected W daughter jets and Fig. 2.18 shows the dijet p_T . All of these quantities have a slight slope in the data to MC ratio, which is discussed in more detail in Sec. 2.5.6.4. The η distributions are shown in Fig. 2.19 and the angular distance between the jets is shown in the left plot of Fig. 2.20. As the p_T of the hadronically decaying W bosons is generally $\lesssim 200$ GeV, the two selected jets are generally $\Delta R \gtrsim 1$ (see Chapter 4). The invariant mass of the two selected W daughter jets is shown in the right plot of Fig. 2.20. As expected, m_{jj} is peaked near m_W , though there is a broad tail from combinatorics and initial and final state radiation.




Figure 2.17: Left (right): The leading (sub-leading) non b -tagged jet p_T . The uncertainty band includes the detector-related experimental uncertainties described in Sec. 2.5. The final bin includes overflow.




Figure 2.18: The distribution of the dijet p_T constructed from the leading two non b -tagged jets (hadronically decaying W boson candidate). The uncertainty band includes the detector-related experimental uncertainties described in Sec. 2.5. The final bin includes overflow.




Figure 2.19: Left (right): The leading (sub-leading) non b -tagged jet η . The uncertainty band includes the detector-related experimental uncertainties described in Sec. 2.5. The final bin includes overflow.




Figure 2.20: Left (right): The ΔR (invariant mass) distribution of the leading to non b -tagged jets. The uncertainty band includes the detector-related experimental uncertainties described in Sec. 2.5. The final bin includes overflow.

2.3 Jet Pull Reconstruction Performance

2.3.1 Detector effects

This section uses the particle-level jets described in Sec. 2.2.3. Detector-level jets are matched to particle-level jets using a $\Delta R < 0.3$ criteria in order to understand how the detector response distorts the particle-level distributions. The output of the event selection in Sec. 2.2.3 is a set of four jets labeled B_1, B_2, J_1 and J_2 for every event. Since the jet pull angle $\theta_P(X, Y)$ requires two jets X and Y as input, there are 12 possible jet pull angles. In general $\theta_P(X, Y) \neq \theta_P(Y, X)$ since the former uses the substructure properties of X while the latter uses the substructure properties of Y . Figures 2.23(a)-2.23(f) show the pull angle distributions⁷ for all cases that involve the W daughter jets and the leading b -jet B_1 . The particle-level distributions are consistent with the corresponding particle level studies in the literature, where a peak at zero corresponds to jets which are ‘color-connected’ (e.g. the daughters of the color singlet W boson) and a uniform distribution corresponds to jets without such a connection [140].

Even though the particle-level distributions in Figures 2.23(c)-2.23(f) are nearly flat, all of the reconstructed shapes are non-uniform. However, there are clear trends: the track pull has a peak at $\pi/2$ and the calorimeter pull is peaked at zero⁸. Therefore, to understand the detector response for the jet pull in $t\bar{t}$, it suffices to study the truth to reconstructed jet pull angle detector response in Fig. 2.23(a) and Fig. 2.23(d) which are representative of the possible shapes and distortions in Fig. 2.23. To minimize the dependence on the physics processes creating the peak at zero in Fig. 2.23(a), most of the discussion in this section will be focused on Fig. 2.23(d) where any departure from a uniform distribution provides insight into detector effects.

⁷The calorimeter jet axis is used for the detector-level pull angles. One conclusion of Sec. 2.3 will be that instead the constituent axis should be used (see Sec. 2.3.6.1). As a result, the pull angle distributions in Sec. 2.4.1 and subsequent sections look qualitatively different than the ones shown here.

⁸An exception is Fig. 2.23(f) for which the peaks are slightly shifted. This is due to the dependence of the pull angle on the jet p_T ; with a higher p_T threshold, Fig. 2.23(f) resembles Fig. 2.23(d).




Figure 2.21: The jet pull angle $\theta_P(X, Y)$ distribution for various choices of X and Y for particle-level jets and also for detector-level jets matched to the particle-level jets.

2.3.2 Jet Pull Angle Response

The transition between particle-level and detector-level distributions is characterized by the *jet pull angle response*, $R(\theta_P)$ – the difference between the detector-level jet pull angle and the particle-level jet pull angle. The calorimeter/all-particles pull angle is calculated from clusters for detector-level jets and all constituents for particle-level jets. The track/charged-particles pull angle uses tracks ghost-associated to the jet for reconstructed jets and charged constituents for particle-level jets. The resolution of the jet pull angle is significantly different depending on the type of constituent used in the definition. Figure 2.22 shows the inclusive jet pull angle response for both the track/charged particle and calorimeter/all particles pull angles. It is evident from the different widths of the two sets of distributions in Fig. 2.22 that the track pull angle is measured more precisely than the calorimeter pull angle. In terms of the RMS of the jet pull angle response, this corresponds to about a 20% improved resolution of the track pull angle over the calorimeter pull angle resolution. The numbers in Fig. 2.22 also indicate small biases in the jet pull angle distributions. These are expected from Fig. 2.23, which show asymmetric shape deformations between the particle-level and detector-level distributions.

In order to fully understand the transition in shapes between particle-level and detector-level in Fig. 2.23 more information is needed beyond the inclusive jet pull angle response from Fig. 2.22. There are three sources contributing to the resolution of the jet pull angle⁹ $\theta_P(X, Y)$ response: the jet constituent angular resolution and momentum resolution with respect to X , the angular resolution of X , and the angular resolution of Y . For both the all-particles and charged-particles pull angles, all angles are computed with respect to the calorimeter (or all-particles) jet axes, independent of the constituents used in the calculation of the jet pull angle. The considerations so far have treated all the resolutions inclusively. It is difficult to systematically remove the resolution from the jet constituents, but it is straightforward to study the effect

⁹For the track-based pull, the definition also introduces some resolution. For instance, K_s decays and photon conversions that occur before/inside the pixel detector contribute to reconstructed tracks, but are not in the list of stable MC charged particles. Also, the $p_T > 500$ track threshold is not applied to the MC particles. All three of these effects have been studied and found to have a very small impact on the resolution and a negligible impact on the pull angle distribution shape.




Figure 2.22: The distribution of the jet pull angle response, $R(\theta_P)$, for both $\theta_P(J_1, J_2)$ and $\theta_P(J_1, B_1)$ as well as for the calorimeter pull angle and the track pull angle. Statistical uncertainties on the mean and RMS are an order of magnitude less than the values shown.

of the jet angular resolution on the jet pull angle.

One measure of the jet angular resolution is σ^{match} : the ΔR between reconstructed jets and matched particle-level jets¹⁰. Figure 2.23 shows the impact setting $\sigma^{\text{match}} = 0$ by systematically replacing detector-level jet axes with the corresponding matched particle-level jet axes. For both the calorimeter and track pull angles $\theta_P(J_1, B_1)$, setting $\sigma^{\text{match}} = 0$ of the b -jet has essentially no influence on the jet pull angle distribution due to the large lever-arm spanned by the vector connecting B_1 and J_1 . However, setting $\sigma^{\text{match}} = 0$ of the J_1 axis has a dramatic impact on the pull distribution shape. For the calorimeter pull, setting $\sigma^{\text{match}} = 0$ of the J_1 axis shifts the peak of the distribution to $\pi/2$ instead of at 0. Since the track angular resolution is much better than the calorimeter cluster angular resolution, the track pull angle resolution is dominated by the calorimeter jet angular resolution. By setting $\sigma^{\text{match}} =$

¹⁰There are at least two contributions to σ^{match} : 1) the angular distortions in momentum when particles become calorimeter clusters and 2) the set of particles associated with the measured calorimeter clusters may not be the same as the particles in the matched truth jet. The former effect can be studied by systematically smearing the truth jet axis and this dominates σ^{match} . The impact of increased distortions of the truth axis is discussed in the context of Fig. 2.24.

0, the pull angle response RMS decreases and the right plot of Fig. 2.23 shows that the jet pull angle distribution is nearly the same as the truth distribution.




Figure 2.23: The $\theta_p(J_1, B_1)$ distribution for the calorimeter (left) and track (right) pull angle after replacing detector-level jet axes by particle-level jet axes.

The right plot of Fig. 2.23 suggests a simple model for building intuition for the peak at $\pi/2$. Consider a *pseudo MC* model with N massless particles generated randomly from the decay of a single scalar particle whose mass and boost are tuned so that the lab frame p_T of the sum of the decay products is specified and such that all the decay products fall within $\Delta R < R$ of the ‘jet’ axis, defined by the vector sum of all the decay products. The dashed line in figure 2.24 shows the jet pull angle distribution for such a model in which two such jets are generated randomly, and with $N = 10$, $p_T = 80$ GeV, and $R = 0.4$. As expected for the undistorted distribution, the jet pull angle is uniform on $[0, \pi)$. To model the resolution, the constituents are fixed and the jet axis is smeared according to a bivariate normal distribution with zero correlation and $\sigma_\phi^{\text{nom}}, \sigma_y^{\text{nom}}$ taken from the ATLAS detector simulation: $\sigma_\phi^{\text{nom}} \approx 0.025 \approx \sigma_y^{\text{nom}}/1.5$. The resolution used in the simulation is given by $\sigma_\phi = r \times \sigma_\phi^{\text{nom}}, \sigma_y = r \times a \times \sigma_y^{\text{nom}}$, where r is a multiplicative factor and a is an asymmetry. The left plot of Fig. 2.24 for $r = 1$ shows that the peak at $\pi/2$ is a prediction of this simple model. By tuning the model parameters, one learns that this

feature can be explained if the resolution in $\mathbf{y}$ and resolution in ϕ are not the same in ATLAS; the right plot of Fig. 2.24 does not peak at $\pi/2$. The peak at $\pi/2$ comes from two facts: (1) in $(\Delta\mathbf{y}, \Delta\phi)$, the pull vector tends to be stretched towards the $\pm\Delta\mathbf{y}$ axis and (2) the distribution of $\Delta\mathbf{y}(\mathbf{J}_1, \mathbf{B}_1)$ is peaked at zero and thus in $(\Delta\mathbf{y}, \Delta\phi)$ coordinates, $\mathbf{B}_1$ lies on the $\Delta\phi$ axis. As r is increased so that the asymmetry is no longer relevant, the peak at $\pi/2$ disappears in all cases. In fact, it is possible to use these observations to *measure* the jet angular resolution with the jet pull angle. Figure 2.25 shows a χ^2 fit between 20 bins of the pull angle from simulation and templates formed from the toy MC. The minimum χ^2 is at $\alpha = 1.5, r = 1$ as desired, though the fit is much more sensitive to α than to r ¹¹.

A similar model can be created for the calorimeter pull angle, but the interpretation is less straight-forward. In particular, using the same pseudo MC model for the particle-level selection, the calorimeter pull angle resolution can be modeled by smearing all particles and then additionally recomputing the jet axis, since the cluster angular resolution need not be small compared to the jet angular resolution as was the case for tracks. Such a model can generically predict peaks at $0, \pi$ and with angular resolution asymmetry, $\pi/2$, but since there is not a one-to-one matching between particles and clusters, it is not possible to map these models onto a realistic description of the detector.

¹¹As a result, after the origin correction (see Sec. 2.3.6.1) this method loses precision.




Figure 2.24: The jet pull angle constructed in a pseudo MC with various resolution settings for the jet constituents.




Figure 2.25: A χ^2 -like scan between the simulated charged-particles pull angle distribution and templates using toy MC that are with varying asymmetry a and resolution r . The minimum χ^2 is for $a = 1.5$ and $r = 1.0$.

2.3.3 Jet Kinematics and the Jet Pull Angle Response

Unlike other jet substructure variables, the jet pull angle depends not only on the orientation of constituents within a jet, but also the placement of jets within an event, hence the term *jet superstructure*. Thus, even at particle-level, the jet pull angle can depend on the relative orientations of jets in a given event. The right plot of Fig. 2.20 shows the relationship between the particle-level jet pull angle $\theta_P(J_1, J_2)$ and the relative distance between jets, $\Delta R(J_1, J_2)$. The particle-level distribution shows a strong dependence on ΔR , with smaller values of ΔR corresponding to a larger peak at zero. In fact, it is mostly through ΔR that the particle-level distribution of $\theta_P(J_1, J_2)$ depends on the p_T of J_1 , as described below. The particle-level pull angle distributions that involve one of the b -jets are nearly independent of ΔR (and p_T). The right plot of Fig. 2.20 shows the RMS of the jet pull angle response as a function of $\Delta R(J_1, B_1)$, which is used because there is no ΔR dependence at particle-level.




Figure 2.26: Relationship between the θ_P (response) and ΔR between the jets.

The left plot of Fig. 2.27 shows the relationship between the jet pull angle $\theta_P(J_1, J_2)$ and the p_T of J_1 . There seems to be a clear relationship between $p_T^{J_1}$ and $\theta_P(J_1, J_2)$. However, this is inconsistent with the truth distributions in Figures 2.23(a) and 2.23(b); these distributions are nearly identical and yet the underlying p_T distribution for Fig. 2.23(b) must be softer than that of Fig. 2.23(a). The resolution is that the pull

angle distribution depends on p_T only through ΔR . At high J_1 p_T , J_1 and J_2 have smaller angular separation since the boost of the W boson in this case is larger (see Chapter 4). The right plot of Fig. 2.27 shows the distribution of $\theta_P(J_1, J_2)$ in bins of the p_T of J_1 for a fixed $\Delta R(J_1, J_2)$. The p_T dependence compared to the left plot of Fig. 2.27 is significantly reduced.




Figure 2.27: The jet pull angle distribution in three bins of p_T of J_1 .

Even though the jet pull angle is relatively independent of p_T , the response RMS does scale with p_T . Figure 2.28 shows the RMS of the jet pull angle response as a function of the p_T of the leading W daughter jet. The RMS improves with increasing p_T as the relative jet energy resolution improves with energy.




Figure 2.28: The RMS of the jet pull angle response as a function of the jet p_T . The ΔR between jets is restricted between 1.4 and 1.6 in the right plot. The track-based pull angle response is nearly flat by $p_T \sim 100 \text{ GeV}$.

2.3.4 Relationship Between the Jet Pull Angle Response and Jet Constituents

As the pull vector is determined from the constituents inside a jet, the jet pull angle response could depend on the number and orientation of the constituents of J_1 . There are many substructure variables which capture various properties of the orientation of constituents within a jet. One such property is the pull vector magnitude,

$$|v_p(J)| = \left| \sum_{i \in J} \frac{p_T^i |r_i|}{p_T^J} \vec{r}_i \right|. \quad (2.8)$$

In dedicated phenomenological studies, it was shown that the pull magnitude is not useful in discriminating octet from singlet color states [259]. However, this section will show that it is a useful handle on the jet pull angle resolution. The jet pull vector magnitude can be considered a radial moment, with the radial distance $(\Delta R)^2$ from the jet axis weighted by the fractional constituent p_T . The distribution of the magnitude for J_1 is shown in the left plot of Fig. 2.29. Events with a reconstructed pull vector magnitude of zero for the track pull, corresponding to cases in which there are no tracks ghost-associated to the jet, are not shown.

The right plot of Fig. 2.29 shows the relationship between the jet pull angle response RMS and the pull vector magnitude. A small pull vector magnitude corresponds to a worse resolution, in some cases because of a small lever arm. Since the pull vector magnitude can be measured, the right plot of Fig. 2.29 suggests that it can be used as an in-situ tool for improving precision. The left plot of Figure 2.30 shows the RMS of the pull response as a function of the efficiency for a threshold requirement on the pull vector magnitude. For instance, one can achieve a $\sim 10\%$ reduction in the RMS of the jet pull angle response while maintaining a 90% selection efficiency.

One undesirable property of the pull vector magnitude in terms of constraining the resolution is that it is anti-correlated with the jet p_T as shown in the right plot of Fig. 2.30. As the jet becomes more collimated, the constituents have a smaller ΔR




Figure 2.29: The pull vector magnitude (left) and the relationship (right) between the jet pull angle response and $|\mathbf{v}_P(J_1)|$ in a particular bin of jet p_T . For the track (charged particle) pull magnitude in the left plot, at least one track (charged particle) is required.




Figure 2.30: The RMS of the jet pull angle response as a function of the fraction of events that pass a cut on the pull vector magnitude (left) and the p_T dependence of the average pull vector magnitude (right).

with respect to the jet axis and so the pull vector magnitude decreases. Accordingly, an optimal threshold on the pull vector magnitude would be p_T dependent.

Another substructure observable that is correlated with the jet pull angle response

is the number of constituents. The pull angle resolution decreases with the number of constituents at low constituent multiplicity as shown in Fig. 2.31. The calorimeter pull angle and the track pull angle each require at least one cluster or track, respectively.




Figure 2.31: The jet pull angle response as a function of the number of jet constituents for J_1 .

2.3.5 Jet Pull Angle and Event Properties

2.3.5.1 Jet Labeling in $t\bar{t}$ Events

For a given jet pull angle $\theta_p(X, Y)$, there is the complimentary angle $\theta_p(Y, X)$ which uses different substructure information. Figure 2.32 shows that this information is largely uncorrelated. Furthermore, it is apparent from Figures 2.23 that there is a relationship between the shapes of the jet pull angle distributions and the assignment of the jets in the $t\bar{t}$ topology. For example, one can investigate the frequency with which the b -tag and dijet invariant mass assignment of J_1, J_2, B_1 and B_2 described in Sec. 3 aligns with the observed property that $\theta_p(J_1, J_2)$ and $\theta_p(J_2, J_1)$ tend to be smaller than $\theta_p(J_i, B_j)$, $\theta_p(B_i, J_j)$ or $\theta_p(B_i, B_j)$.

An event is called *matched* if $\theta_p(J_a, J_b) < \theta_p(J_a, B_1)$ for any $a, b \in \{1, 2\}$. Conversely, if $\theta_p(J_a, J_b) \geq \theta_p(J_a, B_1)$, an event is called *un-matched*. Figure 2.33 shows the tradeoff between matched and un-matched event efficiencies for a threshold on the

jet pull angle using particle-level jets¹². In other words, consider the $\theta_p(J_a, J_b)$ distribution as ‘signal’ and the $\theta_p(J_a, B_1)$ distribution as ‘background’. Then, Fig. 2.33 shows the relationship between signal and background efficiency as a function of the threshold on θ_p . Also plotted in Fig. 2.33 is the combined performance curve from both variables ($\theta_p(X, Y)$ and $\theta_p(Y, X)$), which is significantly better than either curve separately. In absolute units, the overall discrimination is poor – pull is not intended to be used as a stand-alone tagger. Since the jet pull angles with **b**-jets are independent of $\Delta R(X, Y)$ but $\theta_p(J_1, J_2)$ becomes more pronounced at smaller ΔR , there is a slight improvement in the efficiency curve, which is shown in the right plot of Fig 2.33.




Figure 2.32: Pairwise all-particles pull angle correlations using particle-level jets.

2.3.5.2 Pileup

An important event property from the point of view of the pull angle RMS is μ - the average number of additional **pp** interactions per bunch crossing at the LHC. The dependence of the RMS of the jet pull angle response is shown as a function of μ in Fig. 2.34. The RMS of the jet pull angle response is only weakly dependent on the pileup activity. For example, a linear fit to the data in Fig. 2.34 results in a slope of

¹²Particle-level jets are used here to illustrate the maximal achievable performance in the absence of selection biases and detector resolution effects.




Figure 2.33: Un-matched event (treated as a background) rejection versus the matched (treated as a signal) efficiency. The optimal combination is constructed from the full joint likelihood.

about $(1.6 \pm 0.1) \times 10^{-3}$ rad/interaction for the calorimeter pull angle response RMS and $(1.5 \pm 0.1) \times 10^{-3}$ rad/interaction for the track pull angle response RMS in the range $50 \text{ GeV} < p_T^J < 60 \text{ GeV}$. This trend does not vary greatly with p_T .




Figure 2.34: The RMS of the jet pull angle response as a function of μ for $50 \text{ GeV} < p_T^J < 60 \text{ GeV}$.

2.3.6 Comparisons between Data and Simulation

The purpose of this section is to qualitatively compare the pull vector in simulation with data. A quantitative comparison that disentangles detector-level and particle-level effects through unfolding is in Sec. 2.4. The MC is normalized by area to the data in all the following distributions. The uncertainty bands on the data/MC ratios include the experimental uncertainty on the tracking efficiency, the jet energy scale and the jet energy resolution in addition to a $\pm 6\%$ relative cross-section uncertainty on the $t\bar{t}$ component [213–218]. For the pull angle, the average uncertainty across all bins is plotted to remove fluctuations due to the small dependence of the pull angle on the jet energy scale and resolution uncertainties. Uncertainties on the cluster energy scale and angular resolution are not included in this section.

The pull vector magnitude is shown in Figure 2.35 for both MC@NLO+HERWIG and POWHEG-BOX+PYTHIA 6. For the track pull angle, at least two tracks are required in order to remove the portion of the resolution curve in Fig. 2.31 where the response RMS decreases at low constituent multiplicity. Both the calorimeter- and track-based distributions are within 10% of the data over nearly the entire range. There seems to be a minor overall slope in the data/MC ratio for the calorimeter-based pull angle that is due in part to the modeling of the angular resolution (see Sec. 2.3.6.1). Interestingly, the PYTHIA and HERWIG¹³ mis-modeling at low pull vector magnitude are in opposite directions of the data; this is likely due to the mis-modeling of the track multiplicity, as discussed in Chapter 3.

The distribution of the jet pull angle in the data is shown in Fig. 2.36 for $\theta_p(J_1, J_2)$. The resolution features at $\pi/2$ for the track-based pull angle and at zero for the calorimeter-based pull angle are both present and well described. The bias toward zero in the particle-level distribution (Fig. 2.23) that is also present in the particle-level selection (Fig. 2.23(a)) is reduced in Fig. 2.36 due to a selection bias: in a given event, the particle-level and detector-level assignment of jet labels can differ. This selection bias decreases with the increasing p_T of the jets, as is seen in the right plot of Fig. 2.36, where the peak at zero for the track pull dominates the resolution peak

¹³Sec. 2.5.6.1 shows that the ME generator is unimportant for the pull angle distribution.




Figure 2.35: The pull vector magnitude for both calorimeter pull and track pull. For the track pull, at least two tracks are required. Uncertainty bands include uncertainties on the jet energy scale and uncertainty as well as on the $t\bar{t}$ component of the MC. An uncertainty on the tracking efficiency is added for the track pull. No uncertainty is included for individual calorimeter clusters or for jet angular resolutions.

at $\pi/2$ in the MC. The size of the peak at zero also increases with p_T as discussed in Sec. 2.3.3. Figure 2.37 shows the jet pull angle distribution between the leading W daughter jet and the leading b -jet. Based on the studies summarized in Fig. 2.24, the slight parabolic trend in the track-based pull angle ratio in the left plot of Fig. 2.37 (and the left plot of Fig. 2.36) suggests that the scale or asymmetry parameter of the jet angular resolution may be over-estimated, though quantifying this statement is beyond the scope of this section.




Figure 2.36: The distribution of the jet pull angle $\theta_P(J_1, J_2)$ for both calorimeter cluster constituents and track constituents in both data and MC. The left plot has a 25 GeV requirement for the jets while the right plot has a tight threshold placed on the p_T of the dijet system.




Figure 2.37: Same as Fig. 2.36 only with $\theta_P(J_1, B_1)$ instead of $\theta_P(J_1, J_2)$.

2.3.6.1 Cluster and Jet Origin Corrections

Part of the jet calibration procedure is to correct the η of jets so that the detector-level pseudo-rapidity is an unbiased measurement of the corresponding particle-level quantity (see Sec. 4.3). During the LHC shutdown between Runs 1 and 2, the jet calibration procedure was improved to reduce the resolution in the η direction by correcting η event-by-event such that the jet axis is radially aligned with the primary vertex (defined by $\sum_{\text{tracks}} p_T^2$). A beamspot with finite size smears out the η resolution because of event-by-event distortions in the η value from a primary vertex that is not at the geometric center of ATLAS. Figure 2.38 illustrates the geometry of this distortion. The beamspot is $\mathcal{O}(10)$ cm in the z direction and the calorimeter is about 1 m away from the primary interaction. For $\eta = 0.5$, this is a correction of $\Delta\eta \sim 0.4$ for a primary vertex of $z \sim 50$ cm and $\Delta\eta \sim 0.1$ for $z \sim 10$ cm¹⁴.




Figure 2.38: A schematic diagram of the origin correction. The quantity θ^{detector} is the angle in the zx plane measured with respect to the geometric center of the detector, whereas the ‘true’ θ^{physics} is offset and based from the primary vertex (PV).

The origin correction significantly improves the jet η resolution. Figure 2.39 shows the η and ϕ angular resolution before and after the origin correction. The correction has no effect on the ϕ resolution, but reduces the width of the η resolution by a factor

¹⁴One can write $\theta_{\text{detector}} = 2 \arctan(e^{-\eta_{\text{detector}}})$ and then the z distance in detector coordinates is $z_{\text{detector}} \approx 1 \text{ m} / \tan(\theta_{\text{detector}})$. The ‘physics’ position $z_{\text{physics}} = z_{\text{detector}} - z_{\text{PV}}$. The physics angle is then approximately $\theta_{\text{physics}} \approx \arctan(1 \text{ m} / z_{\text{physics}})$.




Figure 2.39: The η (left) and ϕ (right) jet angular response before and after the jet origin correction. The response is defined as the RMS of the $\Delta\eta$ or $\Delta\phi$ between the reconstructed jet and the $\Delta R < 0.4$ matched particle-level jets.

of two. Since the pull angle resolution significantly depends on the jet axis resolution, the origin correction could significantly improve the pull angle resolution. However, it is not sufficient to correct the jet axis *without correcting the cluster positions as well*. The jet origin correction is performed *after* jet clustering so the jet constituents are unchanged. For jet substructure variables that depend on the jet axis (such as the jet pull), this introduces a mis-match that can be avoided by coherently origin correcting the constituent calorimeter-cell clusters. The cluster origin correction is performed as follows. Let $\mathbf{R}_i$ be the calorimeter-cell energy weighted center of the calorimeter-cell cluster i in detector coordinates. The transverse radius is defined by $R_{T,i} = R_i / \cosh(\eta_i)$. The new η position of the cluster i is

$$\begin{aligned} \eta_{\text{physics}} &= \text{asinh} \left(\frac{z_{\text{physics}}}{R_{T,i}} \right) = \text{asinh} \left(\frac{1}{R_{T,i}} (z_{\text{detector}} - z_{\text{PV}}) \right) \\ &= \text{asinh} \left(\sinh(\eta_{\text{detector}}) - \frac{z_{\text{PV}}}{R_{T,i}} \right). \end{aligned} \quad (2.9)$$

In order to preserve the total energy, the cluster transverse momentum becomes $\mathbf{p}_{T,\text{physics}} = \mathbf{p}_{T,\text{detector}} \cosh(\eta_{\text{detector}}) / \cosh(\eta_{\text{physics}})$. Figure 2.40 shows the improvement in the pull angle resolution from coherently applying the cluster origin correction. The reduction in the width of the pull angle response distribution ($\sim 5\%$) is modest, but there is a significant improvement in the modeling of the pull vector. Figure 2.41 compares the simulation to the data before and after coherently applying the origin correction. When neither the jet or cluster axes are origin corrected, the pull vector is well-modeled (Fig. 2.35) even if the resolution with respect to the particle-level quantity is worse because both axes are at the same angular ‘scale’. This is also true after both axes are corrected (left plot of Fig. 2.41). However, if only the jet axis is corrected, then the pull vector is maximally sensitive to the modeling of the beamspot because the cluster locations are distributed about the jet axis (which is independent of the PV) according to the width of the beamspot. The mis-modeling without coherent origin corrections is shown in the right plot of Fig. 2.41. Even though the uncertainty band is large, there is a clear systematic trend in the data/MC ratio and the distribution itself is stretched to higher values due to the offset between axes. Henceforth, both axes are coherently corrected¹⁵. The left plot of Fig. 2.42 confirms that the origin corrections improve over the uncorrected case. Most of the improvement in the resolution is from the jet axis correction and as already noted by Fig. 2.40, the resolution improvement from the cluster origin correction is smaller. However, the pull angle distribution qualitatively changes after each step of the correction, as shown by the right plot of Fig. 2.42¹⁶. Interestingly, when the jet axis resolution is significantly reduced with the jet origin correction, the pull angle distribution resembles the uncorrected track-based pull angle (i.e. a resolution peak at $\pi/2$).

Tracks are already ‘origin corrected’ by construction and so no further correction

¹⁵The actual axis used for the jet pull angle is the four-vector sum of the origin-corrected calorimeter-cell clusters. This is the nearly the same as the origin-corrected jet axis, but the cluster-based systematic uncertainties described in Sec. 2.5.2 are allowed to coherently vary the axis location.

¹⁶Note the slightly different definition due to the absolute value, $|\ast|$. This decreases the response width as events with e.g. pull vectors of π and $-\pi$ have a response of zero. However, this strategy is used for the unfolding described in Sec. 2.4 in order to reduce the number of bins. A scheme with an addition bin to account for the extreme migrations does not significantly reduce uncertainties.




Figure 2.40: The pull angle response before and after applying the cluster origin correction. The jet origin correction is applied in both cases.




Figure 2.41: The distribution of the all-particles pull vector magnitude with both the jet axis and the cluster axes origin corrected (left) and only the jet axis corrected (right).




Figure 2.42: Left: All-particles pull angle response. Right: the all-particles pull angle distribution after various origin corrections.

is required beyond the jet origin correction. However, it is possible to further improve the performance by using the *track-axis* formed from the four-vector sum of the tracks instead of the origin corrected axis. Figure 2.43 shows the axis angular response for the calorimeter jet axis and the track-axis. They have a similar resolution, but by using the track-axis, the charged-particles pull angle is nearly incentive to the calorimeter angular resolution. The pull angle response is shown in the left plot of Fig. 2.44 for the various jet axis definitions. As the origin corrected jet axis has a similar resolution to the track axis, the pull angle resolution is similar for these two choices of axis and both are improved with respect to the starting axis. As with the all-particles pull angle, the reduction in the angular resolution qualitatively changes the pull angle distribution shape (right plot of Fig. 2.44). As expected from Sec. 2.3.2, the reduction in the axis resolution removes the resolution peak at $\pi/2$; now the track-based pull angle between the two W daughter jets peaks at zero as is also the case for the all-particles pull angle. In all subsequent studies, the track axis is used for the charged-particles pull vector. The pull modeling of the track-based pull vector magnitude is shown in Fig. 2.45. The magnitude is generally shifted toward lower




Figure 2.43: The ϕ (left) and η (right) angular response for various choices of the jet axis. The ‘baseline’ configuration uses the calorimeter jet axis without the origin correction. The ‘calo axis’ uses the origin correction and the ‘track axis’ and ‘truth track’ use the four-vector sum of tracks. In all cases except the last one, the particle-level reference object is the full particle-level jet axis while in the last case, the reference is the four-vector sum of the charged particles only. The number in parenthesis is the RMS.

values than the all-particles pull vector due to the smaller constituent multiplicity.

As a summary, the particle-level, detector-level, and response for the all-particles and charged-particles pull angles are shown in Fig. 2.46. With the various axis modifications described above, all the detector-level distributions peak at zero just like the particle-level distributions. In addition to the SM pull angle distributions, Fig. 2.46 also shows the flipped W bosons for which the W decay products are not color connected. The pull angle distribution is more uniform for the octet than for the singlet; the remainder of Chapter 2 is aimed at studying how well these distributions can be distinguished with the ATLAS data. Figure 2.47 shows the all-particles and charged-particles pull angles in data at detector-level for all the axis modifications described in this section. Removing distortions from detector effects for a direct comparison with the particle-level models is described in Sec. 2.4.




Figure 2.44: Left: Charged-particles pull angle response. Right: the charged-particles pull angle distribution after various origin corrections.




Figure 2.45: The charged-particles pull vector magnitude using the track four-vector sum for the jet axis.




Figure 2.46: The all-particles (left), charged-particles (middle), and pull angle response (right) in simulation using the nominal color singlet $t\bar{t}$ model and additionally with the color octet model (left and middle only).




Figure 2.47: The detector-level all-particles (left) and charged-particles pull (right) angles in data and in simulation. The uncertainty band includes only the experimental uncertainties on the inputs to the event selection and the jet pull calculation (see Sec. 2.5). A large part of the uncertainty displayed here affects the overall normalization and is correlated between the individual bins. This component of the uncertainty is cancelled in the unfolded measurement of the unit-normalized pull angle distribution (see Sec. 2.4).

2.4 Unfolding

The rest of Chapter 2 is dedicated to the measurement of the pull angle by correcting for detector effects through unfolding.

2.4.1 Unfolding Parameters

In addition to the number of iterations in the Iterative Bayesian (IB) unfolding algorithm, various aspects of the unfolding setup are optimized to reduce the total uncertainty. The optimization is performed separately for the all-particles and charged-particles pull angles because the resolution is significantly worse for the former variable. Three settings were scanned in the optimization procedure:

Number of bins The closer the response matrix is to the identity matrix, the less dependent the results will be on the unfolding procedure and in particular on the prior in the IB algorithm. Generally, it is desirable for the diagonal elements of the response matrix to satisfy $\text{Pr}(\text{bin } i_{\text{truth}} \rightarrow \text{bin } i_{\text{reco}}) \gtrsim 50\%$. Since the pull angle resolution is comparable to the allowable range $\theta_p \in [0, \pi]$, it is expected that only a few, $\pi/\sigma(\theta_p^{\text{reco}} - \theta_p^{\text{true}}) \sim 3$, bins will be possible. More bins for the charged-particles pull angle are expected due to the superior resolution compared to the all-particles pull angle (see Fig. 2.46).

Number of iterations Increasing the number of iterations in the IB method reduces the dependence on the prior, chosen to be the particle-level spectrum in simulation. However, after a certain number of iterations the results saturate. The point at which the results do not change with more iterations (saturation) depends on the resolution. Figure 2.48 uses a simple calculation using a Toy MC to show how the saturation point depends on the resolution. For the pull angle $\sigma/\text{Range} \sim 1/3$, a saturation occurs at ~ 25 iterations. One does not necessarily want to use the number of iterations corresponding to the saturation point; increasing the number of iterations usually reduces the dependence on the truth spectrum in the MC used to construct the response matrix, but the cost is a larger statistical uncertainty.

Pull Magnitude cut The studies in Sec. 2.3.4 showed that $\sigma(\theta_p^{\text{reco}} - \theta_p^{\text{true}})$ depends on p_T and on the magnitude of the pull vector. The tradeoffs for a magnitude requirement are a reduction in statistics and a potential increase in model dependence, as the jet pull angle magnitude contains information about color flow.




Figure 2.48: The number of iterations required for the unfolded results to saturate using a Toy MC. The unfolding is said to saturate if the unfolded bin contents do not change by more than 0.01% between iterations. In this model, $\theta_p \sim \text{Uniform}(0, \text{Range})$ and $\theta_p^{\text{reco}} - \theta_p^{\text{true}} \sim \mathcal{N}(0, \sigma^2)$, where $\text{Range} = \pi$ and the smearing is done modulo π .

The parameters described above are interrelated and so the optimization must be performed simultaneously. The number of equal sized bins was scanned between 3 and 9, the number of iterations was scanned between 1 and 6 (1 and 14 for all-particles) and the pull vector magnitude was scanned between 0 (no requirement) and 3.5×10^{-3} in steps of 5×10^{-4} leading to 336 (784) configurations for the charged-particles (all-particles) pull angles. For each configuration, the data statistical uncertainty was combined with the dominant systematic uncertainties, including the color flow model,

fragmentation model, and the data-driven non-closure¹⁷. All of these uncertainties require running the unfolding algorithm with the given parameters at least once and are described in Sec. 2.5.

Using the bin-averaged uncertainty as a metric and allowing for some slight post-hoc modifications with unequal bin sizes, the parameters for the unfolding are as follows:

All-particles : 3 bins with ranges $[0, 0.275, 0.6375, 1.0] \times \pi$, 15 iterations, and no pull vector magnitude requirement.

Charged-particles : 4 bins with ranges $[0, 0.2, 0.5, 0.8, 1.0] \times \pi$, 3 iterations, and no pull vector magnitude requirement.

The optimization procedure suggested that a small requirement on the pull vector magnitude for both the all-particles and charged-particles pull angles could reduce the overall uncertainty. However, since such a gain is below 1% in the bin-averaged uncertainty and would introduce a new source of model dependence, the requirement is not used for the final configuration. Figure 2.49 shows the impact of the number of iterations on the uncertainty in each bin using only the color flow model and statistical uncertainties for illustration. For both the all-particles and charged-particles pull angles, the statistical uncertainty increases monotonically with the number of iterations. The combined uncertainties for all bins have a minimum in the plotted range except the second bin of the all-particles pull angle, due to the fact that the distributions are normalized before computing the uncertainties. For nearly all iteration choices, the color flow uncertainty is larger than the statistical uncertainty. Unlike the statistical uncertainty, the color flow uncertainty decreases with the number of iterations and then increases again as the unfolding overcorrects the simulation.

The remainder of this section describes in more detail the interplay between the resolution and the number of iterations required to reduce the model-dependence uncertainties. A first observation is that since the jet pull angle is a bounded variable, there is an induced correlation between the pull angle response and the pull angle

¹⁷The optimization was performed without data, so the statistical uncertainty is based on the expected yields from the simulation, and the non-closure uncertainty used the data only indirectly.




Figure 2.49: The sum in quadrature of the data statistical uncertainty and the color flow model uncertainty (See Sec. 2.5) as a function of the number of iterations in the IB unfolding method for the three bins of the all-particles pull angle (left) and for the four bins of the charged-particles pull angle (right). For the first bin, the open markers show how the total uncertainty is broken down into the two components.

itself. This is illustrated schematically in Fig. 2.50. If the (normalized by π) particle-level jet pull angle is 0.5, then the difference between the detector-level and the particle-level values can be at most 0.5^{18} . However, if the particle-level pull angle is 0 or 1, then the difference between the detector-level and the particle-level values can be as large as 1. This correlation is important because it means the response matrix depends on the particle-level pull angle distribution, which is most relevant for the color flow model uncertainty (which by construction has quite a different pull angle distribution).

A Toy MC is constructed to quantify this dependence. The particle-level spectrum is constructed as a one parameter family of distributions with varying peak heights at zero, emulating the important difference in the pull angle distribution between

¹⁸One important subtlety is about when the absolute value is taken when computing the response and the pull angle. A pull angle of π and a pull angle of $-\pi$ have the same probability under a given color flow model, but experimentally, $\pi = \theta_p^{\text{true}} \rightarrow \theta_p^{\text{reco}} = -\pi$ is a (maximal) mis-measurement. One way around this is to introduce another bin in the response matrix to account for negative value. This was tested and did not improve the uncertainty because increasing the number of bins resulted in lower transition probabilities in the response matrix.




Figure 2.50: A schematic diagram of the jet pull angle response which illustrates that the pull angle resolution is strongly correlated with the pull angle itself. For example, if you know that the difference between the particle-level and detector-level (normalized by π) pull angles is greater than 0.5 (blue shaded box on the right), then the truth pull angle must have been greater than 0.5.

the singlet and octet color flow models. Angles are generated uniformly at random between 0 and 2π and are then smeared with a Gaussian (modulo π) that has mean zero and standard deviation σ . The ‘measurement’ is performed with the absolute value of the angle divided by π so that the range is between 0 and 1. Each event is then re-weighted such that the truth spectrum probability distribution function is a right triangle with base length X and height set by normalization. As $X \rightarrow \infty$, the distribution between 0 and 1 is uniform and as $X \rightarrow 0$, the distribution is a δ -function at 0. Figure 2.51 shows the distribution corresponding to various values of X . Figure 2.52 shows the results of unfolding the measured (i.e. smeared) toy data. The z -axis is the bin normalized fractional uncertainty, defined as the difference between the truth distribution and the unfolded toy data. The toy truth and toy data both have $X = 5$ while the response matrix has a variable X_{MC} value. The uncertainty increases as X_{MC} moves away from 5 and the size of this uncertainty is bigger for larger angle smearing σ . The difference between the left and right plots in Fig. 2.52 shows that the size of the uncertainty can be mitigated by increasing the

number of iterations. For reference, the σ for the all-particles pull angle is $\sigma \sim 0.35$ and $\sigma \sim 0.28$ ¹⁹ for the charged-particles pull angle. In addition to the difference in resolutions, the absolute difference between the singlet and octet charged-particles pull angle distributions is smaller than for the all-particles pull angle because some of the discriminating information is lost in the neutral radiation.




Figure 2.51: The truth distribution for the toy experiment described in Sec. 2.4.1. The distributions are indexed by the base of a triangle which varies between $X = 1$ and 5. When the base length is 5, the distribution is close to uniform and when it is 1, the distribution is strongly peaked at zero.

¹⁹The standard deviation does not fully capture the large differences between the resolutions - see Fig. 2.46.




Figure 2.52: The bin averaged fractional uncertainty from comparing the truth and unfolded toy data ($X = 5$) using different response matrices (with $X = X_{MC}$). The value of σ is used in both the toy data and the response matrix. The left plot uses three iterations while the right plot uses 15 iterations.

2.4.2 Correction Factors

With the unfolding setup fixed, the next step in the unfolding procedure is to subtract non- $t\bar{t}$ processes from the data and apply correction factors. Background estimates (described in Sec. 2.2.1) are subtracted bin-by-bin in the pull angle distribution. Even though the expected background composition is about 10% of the total yield, the background pull angle distributions are nearly independent of the pull angle and therefore this correction has nearly no impact on the *normalized* pull angle distribution. Figure 2.53 shows the background composition as a function of the pull angle. The background distributions vary by less than 3% across bins.




Figure 2.53: The background composition as a function of the all-particles pull angle (left) and the all-particles pull angle (right).

One background that requires careful consideration is the single top Wt process. These events have a hadronically decaying W boson and the pull angle distribution would change depending on the color charge of the W boson. The nominal procedure is to subtract this component as if it were *background*, even though it is expected to behave as the *signal*. To assess the impact of this choice, single top Wt events were replaced by the nominal $t\bar{t}$ events, but scaled to the single top normalization. This ensemble was compared with an analogous one in which the Wt contribution is replaced with the color octet version of the nominal $t\bar{t}$ sample. The difference

in the unfolded result between these two setups across all bins is much less than the statistical uncertainty and therefore is ignored for the rest of the analysis (see Sec. 2.5.5 for more detail).

After subtracting the non- $t\bar{t}$ backgrounds, the data are corrected to account for events which may pass the detector-level selection but not the particle-level selection. Unlike for the jet charge measurement, the color flow measurement has a non-trivial event selection with requirements on many reconstructed objects. Due to the resolution and (in)efficiencies of these objects, there are a large fraction of events that pass one of the particle-level and detector-level selections, but not both. Figure 2.54 shows how the ratio of the number of events passing both the particle-level and detector-level event selections to the number of events passing only the detector-level event selection (*fake factor* - see Sec. 1.4.1.2) depends on the pull angle. The fake factor is about 70% and is largely independent of the pull angle. For the same reason as for the background subtraction, this small dependence on the pull angle means that the impact of mis-modeling in the fake factor is suppressed.




Figure 2.54: The pull-angle dependence of the fake factors for the all-particles pull angle (left) and for the charged-particles pull angle (right).

After the unfolding with the response matrix, discussed in the next section, *inefficiency factors* are applied to account for events in simulation that pass the particle-level selection but not the detector-level selection. Figure 2.55 shows the inefficiency factors as a function of the all-particles and charged-particles pull angles. Due to falling p_T spectra, the inefficiency factors are much smaller than the fake factors. However, similar to the fake factors, the inefficiency factors are nearly independent of the pull angle and therefore they have little impact on the final measurement.




Figure 2.55: The pull-angle dependence of the inefficiency factors for the all-particles pull angle (left) and for the charged-particles pull angle (right).

The acceptance for electron events is different than for muon events, which is reflected in the difference in fake and inefficiency factors shown in Fig. 2.56. The fraction of events in simulation that pass the particle-level electron channel and the detector-level muon channel selections (or vice versa) is less than $5 \times 10^{-3}\%$ and is ignored for constructing Fig. 2.56. The fake factors are nearly identical between the two channels while the inefficiency factor is approximately 20% higher for the muon channel. This is because the lepton contribution to the fake factor is due mostly to the mis-identification rate (very small) while the lepton contribution to the inefficiency factor is the particle identification efficiency.




Figure 2.56: The ratio of the fake (inefficiency) factor for muon events to electron events on the left (right).

2.4.3 Response Matrix

Figure 2.57 shows the nominal response matrix, constructed from POWHEG-BOX+PYTHIA 6. Despite the larger bin size, the diagonal entries for the all-particles pull angle are lower than the diagonal entries for the all-particles pull angle. The binning is chosen roughly so that the diagonal entries are $\gtrsim 50\%$. Due to the broad resolution, the migration probabilities are significant; as discussed in Sec. 2.4.1, this will have important implications for the theoretical modeling uncertainties described in Sec. 2.5.6.

Even though the inefficiency factors are slightly different between the electron and muon channels (Sec. 2.4.2), the response matrices are nearly identical. Figure 2.58 quantifies the difference in the response matrix between the two channels. Within the simulation statistical uncertainty, they are identical ($\chi^2/\text{NDF} \approx 0.3$). For all subsequent analysis, the two channels are pooled before unfolding with the response matrix.




Figure 2.57: The response matrix for the all-particles pull angle on the left and for the charged-particles pull angle on the right. The z-axis is the probability for an event to be reconstructed in a detector-level bin given that it started in a fixed particle-level bin on the horizontal axis (i.e. the columns are normalized to unity).




Figure 2.58: The difference between the response matrices for the muon and electron channels divided by the simulation statistical uncertainty.

2.5 Systematic uncertainties

The sources of uncertainty can be classified into two categories: experimental uncertainties and theoretical modelling uncertainties. In the first category, some uncertainties impact the pull angle directly and the others impact only the acceptance. As for the jet charge, systematic uncertainties are estimated by varying an aspect of the unfolding procedure, such as the response matrix. The following sections discuss each uncertainty and Sec. 2.5.9 contains a summary.

2.5.1 Tracking

The method for evaluating uncertainties related to the track reconstruction are detailed in Sec. 1.5.3. Unlike for the jet charge, the tracking uncertainties for the jet pull measurement are sub-dominant and so simple but conservative estimates were chosen while many of the detailed prescriptions in Sec. 1.5.3 were under development. For example, the tracking reconstruction efficiency systematic uncertainty is estimated without the final Run I ID material uncertainty constraint and thus tracks are randomly dropped with larger probabilities than are used for the jet charge measurement [226]. The probability in the region $2.3 < |\eta| < 2.5$ is 7%, $1.9 < |\eta| < 2.3$ corresponds to 4%, $1.3 < |\eta| < 1.9$ is 3%, and $0. < |\eta| < 1.3$ is 2%. These uncertainties do not explicitly take into account the modeling of the efficiency of the explicit track $\chi^2/\text{NDF} < 3$ requirement. However, the impact of any mis-modeling is subdominant to the already large uncertainties (see Sec. 1.5.3.2) and a comparison of simulation with data of the χ^2/NDF distribution in Fig. 2.59 confirms that there is no significant mis-modeling.

Most of the jets have $p_T < 400$ GeV where the impact of hit merging is insignificant. Conservatively, $\sim 50\%$ of the loss (see Sec. 1.5.3.3) is used to determine the rate of dropping tracks based on the jet p_T for estimating the uncertainty for reconstructing tracks inside high p_T jets (distinct from the inclusive efficiency described above). Between 400 and 500 GeV, 0.08% of tracks are randomly removed, for jets between 500 and 600 GeV, 0.8% are removed, between 600 and 800 GeV 1.9% are removed and 3.7% are removed for $p_T > 800$ GeV. The impact of a mis-modeling in the track

p_T resolution is conservatively estimated based on early Run I $Z \rightarrow \mu\mu$ studies by smearing track momenta randomly by 10% [226]. The tracking uncertainties only impact the charged-particle pull angle measurement. Table 2.5 quantifies the impact of the tracking uncertainties on the measured pull angle distribution. In all bins, the tracking uncertainties are significantly smaller than the data statistical uncertainty.




Figure 2.59: The χ^2/NDF distribution of tracks before the $\chi^2/\text{NDF} < 3$ requirement.

Source	Charged-particle θ_P Bin (uncertainty in [%])			
	1	2	3	4
Isolated Efficiency	0.17	0.14	0.05	0.11
Tracking Inside Jets	0.02	0.01	<0.01	0.03
Momentum Resolution	0.06	0.02	0.06	0.03
Total tracking uncertainty	0.18	0.14	0.08	0.12
Statistical uncertainty	0.68	0.47	0.48	0.74

Table 2.3: A summary of the tracking systematic uncertainty and their impact on the charged-particle pull angle measurement. Values are given in percent. For comparison, the data statistical uncertainty is the last line.

2.5.2 Calorimeter Cell Clusters

Uncertainties on the reconstruction of calorimeter cell clusters are estimated using comparisons between tracks and clusters in data and in simulation. Earlier versions of these uncertainties based on 2011 data were used in various jet property measurements in early Run 1 [284–286]. The cluster energy scale and angular resolution uncertainties described in Sec. 2.5.2.2 and 2.5.2.3, respectively, are derived for the first time based on the 2012 dataset. Table 2.5 quantifies the impact of the cluster uncertainties on the measured pull angle distribution.

Source	All-particle θ_P Bin (uncertainty in [%])		
	1	2	3
Reconstruction Efficiency	0.34	0.05	0.28
Energy Scale (Option 1)	$+0.28$ -0.22	$+0.02$ -0.66	$+0.26$ -0.50
Energy Scale (Option 2)	0.04	0.18	0.24
Angular Resolution	0.28	0.05	0.34
Total cluster uncertainty	0.52	0.66	0.67
Statistical uncertainty	1.14	0.58	1.19

Table 2.4: A summary of the cluster systematic uncertainty and their impact on the all-particle pull angle measurement. See Sec. 2.5.2.2 for an explanation of the two options for the cluster energy scale uncertainty. Values are given in percent. For comparison, the data statistical uncertainty is the last line.

The calorimeter cell cluster uncertainties described in this section do not fully take into account collective effects on the jet pull angle. In analogy to the jet energy energy scale uncertainty, it is possible that uncertainties on the jet pull angle from all of the input cluster measurements treated simultaneously may be different than the individual cluster-level approach given in this section. Developing general *bottom-up* cluster-based uncertainties for general jet substructure moments is an area of active research. Some studies addressing isolation and collective effects are addressed in Sec. 2.5.2.2 for the cluster energy scale.

2.5.2.1 Cluster Reconstruction Efficiency

Due to the material in and around the ID before the calorimeter, particles may have significant material interaction before reaching the calorimeter that prevent the seeding of calorimeter cell clusters. Calorimeter cell clusters require seed cells that exceed the noise threshold - if a particle interacts with the material and produces many spread out low energy secondary particles, there may not be sufficient localized energy to seed a cluster. The rate at which particles do not seed a cluster is studied with tracks that do not match a calorimeter cell cluster within $\Delta R < 0.2$ in studies of single hadron response at $\sqrt{s} = 900$ GeV [232]. By construction, this technique also includes in the reconstruction efficiency the rate at which particles scatter by a large angle after the ID, but this is a small effect for the choice of $\Delta R < 0.2$. Figure 2.61 shows a comparison between the rate of unmatched tracks in data and simulation as a function of the track momentum. The rate $P(E = 0) \approx \exp(-2E/\text{GeV})$. To conservatively estimate the uncertainty from mis-modeling the reconstruction efficiency, clusters with $E < 2.5$ GeV are randomly dropped 25% of $P(E = 0)$.




Figure 2.60: The probability for not finding a calorimeter cell cluster matched to a track as a function of the track momentum. Reproduced from Ref. [232].

2.5.2.2 Cluster Energy Scale

The cluster energy scale (CES) uncertainty is estimated using the E/p measurement based on the 2012 dataset [287]. Tracks are extrapolated to the various layers of the calorimeter and matched to clusters. Differences in the ratio of the track momentum and the LCW cluster energy between data and simulation are used as an estimate of the uncertainty. The ratio between data and MC is bounded by the following function:

$$f_{\pm}(p|\alpha, \beta) = 1 \pm \alpha \times \left(1 + \frac{\beta \text{ MeV}}{p}\right), \quad (2.10)$$

where $\alpha(\eta)$ and $\beta(\eta)$ are two dimensionless η -dependent functions and p is the track momentum. Figure 2.61 show the data and MC used to estimate α and β in two bins of η and Table 2.5 summarizes the values over all seven $|\eta|$ bins spanning $0 < |\eta| < 2.3$.

Coefficient	$ \eta $ bin lower edge						
	0	0.6	1.1	1.4	1.5	1.8	1.9
α	0.05	0.05	0.07	0.07	0.07	0.04	0.04
β	500	500	500	0	500	0	500

Table 2.5: A summary of the α and β coefficients in Eq. 2.10 used to bound the differences between data and simulation.

To estimate the impact of the CES uncertainty, the cluster energies inside the jet (after jet-finding) are scaled using the function f . The E/p measurement subtracts out the impact of neutral particles and so is directly applicable only to charged particle induced clusters. However, given the conservative nature of the prescription described below and that the CES uncertainty is subdominant, the same CES uncertainty is applied to all clusters.

Taking into account the correlations between the CES uncertainty is non-trivial and so two approaches are used, with the more larger one retained per bin of the pull angle.

1. For the ‘up’ (‘down’) uncertainty, multiply the four-vector of all clusters inside the jet by $f_+(\mathbf{p}|\alpha, \beta)$ ($f_-(\mathbf{p}|\alpha, \beta)$). The shift will be coherent for all clusters, but the actually scaling will change based on $\mathbf{p}$ and η .
2. There is only one uncertainty: multiply the four-vector of each cluster by a random number with mean one and standard deviation $f_+(\mathbf{p}|\alpha, \beta) - 1$. Generate the random numbers for this procedure in strips of η (with bin size half that of the E/p measurement bins) to allow for some coherence, but still mostly emulating local fluctuations.




Figure 2.61: The average LCW E/p for $0 < |\eta| < 0.6$ (left) and $0.6 < |\eta| < 1.1$ (right), using the same analysis framework as in Ref. [287] (but LCW is not in the public note - thank you B. Axen for the inputs). The blue band in the ratio shows the estimated uncertainty used for the cluster energy scale uncertainty.

By construction, the clusters used in the E/p measurement are isolated. However, the clusters inside the jets used for the jet pull angle measurement can be non-isolated. Figure 2.62 shows the distribution of the cluster energy inside jets in various bins of the cluster isolation (f_{iso} in [82]). The cluster isolation measures the sampling layer energy-weighted fraction of non-clustered neighbor cells on the outer perimeter of a topocluster. An isolation of 1 indicates that the clusters are isolated and an isolation of 0 indicates that the cluster is non-isolated. There is no evidence for significant isolation-dependent energy mis-modelling.




Figure 2.62: The distribution of cluster energy in three bins of the cluster isolation, with less isolated clusters on the left and more isolated cluster on the right. See the text for the definition of the cluster isolation.

2.5.2.3 Cluster Angular Resolution

Early Run 1 measurements of jet properties at $\sqrt{s} = 7$ TeV based on constituent clusters used a cluster angular resolution uncertainty based on differences between data and simulation in the $\Delta\phi$ and $\Delta\eta$ between isolated tracks and clusters. Cluster positions were smeared by 5 mrad independently in η and ϕ to account for potential mis-modeling [288]. Similar studies are shown in this section, based on the full $\sqrt{s} = 8$ TeV dataset that is about a factor of 4 larger than the $\sqrt{s} = 7$ dataset.

A $Z \rightarrow \mu\mu$ ($p_T^Z > 30$ GeV) event selection is chosen so that a significant fraction of clusters are isolated (no jet requirement). Tracks are selected which have a maximum of one cluster within $\Delta R < 0.15$ around their position extrapolated to the second layer of the calorimeter, excluding the muon tracks and with no cluster requirements other than $E > 0$. Fig. 2.63 shows the distribution of $\Delta R(\text{track, calo})$ for such tracks in the barrel ($|\eta| < 0.6$) and Fig. 2.64 shows the same distribution in the endcap ($2 < |\eta| < 2.4$). In all plots, there are clearly two peaks. The second peak is an artifact of the requirement that there be no additional clusters within $\Delta R < 0.15$. To study the impact of single particles, further analysis is only performed on cases in which $\Delta R(\text{track, calo}) < 0.075$ to remove the second peak. The momentum dependence of the $\Delta\phi$ and $\Delta\eta$ between tracks and clusters is tabulated in Fig. 2.65 and 2.66. Differences between the data and simulation are generally $\lesssim 1$ mrad. These differences are significantly smaller than the ones reported in the $\sqrt{s} = 7$ TeV analysis (by a factor of 5 in the endcap and 50 in the barrel). One reason is the restriction to the first peak and thus effectively suppressing the contribution from neutral particles. A version of Fig. 2.65 including the second peak results in resolutions similar to the early Run 1 numbers. However, due to the lack of additional studies to probe the full impact of neutral particles in clusters, a (likely) conservative 5 mrad smearing is also adapted at $\sqrt{s} = 8$ TeV.




Figure 2.63: Left (Right): The ΔR between isolated low (high) momentum tracks and clusters in $Z \rightarrow \mu\mu$ events in the barrel of the detector. Inputs from C. Young.




Figure 2.64: Left (Right): The ΔR between isolated low (high) momentum tracks and clusters in $Z \rightarrow \mu\mu$ events in the endcap of the detector. Inputs from C. Young.




Figure 2.65: Left (Right): The RMS of the $\Delta\eta$ ($\Delta\phi$) between isolated single particle tracks and clusters for tracks extrapolated to the second layer of the calorimeter in the barrel of the detector.




Figure 2.66: Left (Right): The RMS of the $\Delta\eta$ ($\Delta\phi$) between isolated single particle tracks and clusters for tracks extrapolated to the second layer of the calorimeter in the endcap of the detector.

2.5.3 Jet Angular Resolution

As demonstrated in Sec. 2.3.6.1 in the context of the origin correction, the jet pull angle is sensitive to the choice of the jet axis and therefore on the modeling of the jet axis angular resolution (with respect to the particle-level jet axis). One method to estimate the jet angular resolution (JAR) uncertainty is to use the angular displacement between calorimeter jets and track jets. Section 2.5.3.1 below describes this method in detail, but it is not used as the baseline JAR uncertainty for two reasons:

1. The uncertainty in the jet angular resolution should be very correlated with the uncertainties on the clusters. The uncertainty computed with the track jet difference would treat these uncertainties as fully uncorrelated.
2. The uncertainties determined with the track jet method are significantly larger than those determined from propagating cluster uncertainties (which is the baseline method), due at least in part to limited MC statistics in the measurement. Figure 2.70 compares the track-jet method of Sec. 2.5.3.1 with the JAR induced from the cluster uncertainties (baseline prescription).




Figure 2.67: A comparison between the track-jet and cluster-induced JAR uncertainties (baseline) for the jet η (left) and the jet ϕ (right).

2.5.3.1 In-situ method using track jets

Track jets are constructed from tracks using the same jet clustering algorithm as for calorimeter jets. The in-situ JAR uncertainty presented in this section utilizes the excellent angular resolution of these track jet. Let $\Delta\phi = \phi_{\text{track jet}} - \phi_{\text{calo jet}}$ and $\Delta\eta = \eta_{\text{track jet}} - \eta_{\text{calo jet}}$. The resolution of $\Delta\mathbf{x}$ for $\mathbf{x} \in \{\phi, \eta\}$ is given by

$$\sigma_{\Delta\mathbf{x}} \sim \sqrt{\sigma_{\mathbf{x}_{\text{track jet}}}^2 + \sigma_{\mathbf{x}_{\text{calo jet}}}^2}, \quad (2.11)$$

where the resolution for the calorimeter jet $\sigma_{\mathbf{x}_{\text{calo jet}}}$ is with respect to the corresponding particle-level jet with both charged and neutral particles while the resolution of the track jet $\sigma_{\mathbf{x}_{\text{track jet}}}$ is with respect to the corresponding particle level jet with only charged particles. The resolution of the track jets with respect to particle-level jets using both charged and neutral particles is even worse than the calorimeter jet angular resolution due to charge-to-neutral ratio fluctuations that are large compared to the detector-resolution. Standard error propagation on Eq. 2.11 gives an equation involving the uncertainty on the resolution of $\mathbf{x}$, $\sigma_{\sigma_{\mathbf{x}}}$:

$$\sigma_{\sigma_{\Delta\mathbf{x}}}^2 \sigma_{\Delta\mathbf{x}}^2 \sim \sigma_{\sigma_{\mathbf{x}_{\text{track jet}}}}^2 \sigma_{\mathbf{x}_{\text{track jet}}}^2 + \sigma_{\sigma_{\mathbf{x}_{\text{calo jet}}}}^2 \sigma_{\mathbf{x}_{\text{calo jet}}}^2. \quad (2.12)$$

Compared to the calorimeter angular resolution uncertainty, the track jet angular resolution uncertainty should be second order. Dropping the corresponding terms in Eq. 2.12 and solving for $\sigma_{\sigma_{\mathbf{x}}}$ gives an estimate for the uncertainty on the resolution of $\sigma_{\mathbf{x}}$:

$$\sigma_{\sigma_{\mathbf{x}_{\text{calo jet}}}}(\mathbf{p}_T, \eta) \sim \frac{\sigma_{\Delta\phi}(\mathbf{p}_T, \eta)}{\sigma_{\mathbf{x}_{\text{calo jet}}}(\mathbf{p}_T, \eta)} \times \sigma_{\sigma_{\Delta\mathbf{x}}}(\mathbf{p}_T, \eta). \quad (2.13)$$

The track jet method uses differences between data and simulation in the quantity $\sigma_{\sigma_{\Delta\mathbf{x}}}$ to estimate the uncertainty on $\sigma_{\sigma_{\mathbf{x}_{\text{calo jet}}}}$ via a scaling by $\sigma_{\Delta\phi}/\sigma_{\mathbf{x}_{\text{calo jet}}}$ that is

determined from simulation. The practical implementation of the JAR uncertainty would be to smear the ϕ and η of each jet by a Gaussian with mean zero and standard deviation s_x that solves the following equation ($\sigma_x \rightarrow \sigma_x + \sigma_{\sigma_x}$):

$$\sigma_x^2 + s_x^2 = (\sigma_x + \sigma_{\sigma_x})^2 \implies s = \sqrt{2\sigma\sigma_\sigma + \sigma_\sigma^2}. \quad (2.14)$$

Figures 2.68 and 2.69 show the jet p_T and jet η dependence, respectively, of the jet ϕ and η resolutions in simulation. The resolution decreases with jet p_T , dropping below two mrad at about 100 GeV, and is stable for central $|\eta|$, degrading at high $|\eta|$ due to the worse calorimeter granularity.




Figure 2.68: The p_T dependence of the jet angular resolution for η (left) and ϕ (right) for the leading non b -tagged jets in $t\bar{t}$ events. The error bars reflect the statistical uncertainty and the dashed line is a fit to $a + b/p_T$.

The next input to Eq. 2.13 is the resolution $\sigma_{\Delta x}$ and the corresponding uncertainty, $\sigma_{\sigma_{\Delta x}}$. Figures 2.70 and 2.71 show the p_T and η dependence of the track jet - calorimeter jet angular resolution. As expected, the resolution is larger than the corresponding distributions in Fig. 2.68 and 2.69. The uncertainty bands in Fig. 2.70 and 2.71 are the result of various simulation variations, including changes in the fragmentation model and comparisons of the amount of inner detector material. The data (not shown) are consistent with the simulation within these large $\sim 10\%$ uncertainties.




Figure 2.69: The η dependence of the jet angular resolution for η (left) and ϕ (right) for the leading non b -tagged jets in $t\bar{t}$ events.




Figure 2.70: The jet p_T dependence of the resolution of the $\Delta\eta$ (left) and $\Delta\phi$ (right) between track jets and calorimeter jets in simulated dijet events. The error band is described in the text. A dashed line is a fit to $a + b/p_T$. Inputs from F. Guescini.

Figure 2.72 shows the total fractional angular resolution and the smearing amount from Eq. 2.14. Due to the origin correction, the η and ϕ resolutions are comparable and the uncertainty on η is even smaller than for ϕ . However, the uncertainties are large - 15%-20% around $p_T \sim 50$ GeV. A significant contribution to this uncertainty




Figure 2.71: The same as Fig. 2.70, only with η instead of ϕ . The dashed line is a fit to $a + b\eta^2$. Inputs from F. Guescini.

is the limited MC sample size²⁰. A combination of larger simulations sets and reduced modeling systematic uncertainties will allow this technique to be a competitive validation in the future.




Figure 2.72: The relative jet angular resolution for ϕ and η (left) and the amount by which the jet angle should be smeared to conservatively cover the uncertainty in the angular resolution (right).

²⁰This is not apparent from the nearly smooth error band in Fig. 2.70 and 2.71 because the fluctuations from many variations are summed together.

2.5.4 Other Experimental Uncertainties

As the pull vector definition uses the calorimeter jet $\mathbf{p}_T$, both the all-particles and charged-particles pull angle are affected by the uncertainty in the jet energy scale [121, 289] and resolution [233]. However, changes in the jet energy scale and resolution do not impact the pull *angle*, but do impact the results via the acceptance due to $\mathbf{p}_T$ thresholds (evaluated in the same way as in Sec. 1.5.2). Similarly, uncertainties in the lepton energy scale, trigger efficiency, E_T^{miss} resolution and $\mathbf{b}$ -tagging efficiencies [224, 290–292] indirectly affect the results through changes in acceptance.

2.5.5 Background Processes

Other (minor) sources of uncertainty on the acceptance, which impact the measurement through the background subtraction, include those related to the luminosity [70], the multijet estimation, and the normalisation and heavy flavour content of the W +jets background [281]. The luminosity uncertainty of $\pm 2.8\%$ only affects those backgrounds that are estimated directly from simulation, including the single top, diboson, and Z +jets processes. Like the $t\bar{t}$ signal, the single production of a top quark in association with a W boson also can have one leptonically decaying W boson and one hadronically decaying W boson (Wt). An uncertainty on the quantum interference of the NLO Wt process with leading order $t\bar{t}$ process is estimated by comparing the DR and DS overlap removal schemes [273] (more detail in Sec. 6.2.2). Additionally, there is an uncertainty related to the flipped model. If the W boson radiation follows an octet pattern, then the contribution from Wt will be more like the flipped model than the SM $t\bar{t}$. Since the Wt is subtracted along with the other minor backgrounds, this could bias the measurement. However, the Wt is only about 3% of the total background composition and the difference between the singlet and octet radiation pattern is $\mathcal{O}(\%)$. A flipped Wt sample is not generated, but the impact can be estimated by replacing the Wt background with the flipped $t\bar{t}$ scaled to the predicted Wt yield. Figure 2.73 shows that such a conservative uncertainty would be much smaller than the statistical uncertainty (already subdominant to the uncertainties in Sec. 2.5.6) and is thus not considered for the remainder of the analysis.




Figure 2.73: The fractional difference in the unfolded result when the Wt is replaced with the flipped $t\bar{t}$ scaled to the Wt yield compared with the data statistical uncertainty for the all particles pull angle (left) and the charged particles pull angle (right).

2.5.6 Top Quark Pair Production Modeling

2.5.6.1 ME Generator and Fragmentation Model

As expected, differences in the the pull angle distribution for a fixed fragmentation model and variable ME generator are small compared to the reverse setup. This is demonstrated by Fig. 2.74 and 2.75. There are percent-level differences in the pull angle distribution between PYTHIA 6 and HERWIG.

2.5.6.2 Color Reconnection

While the color reconnection in the parton shower is also due to the exchange of color charge, it is expected to not have a large impact on the jet pull distribution as the pull vector should be set by the color flow at the hard scatter. This is supported by early studies in Ref. [293]. Figure 2.76 compares the nominal PYTHIA 6 tune (P2011C) with the 2012 Perugia lowCR tune [212] (as well as a tune for higher MPI). The lowCR Perugia tune differs from the nominal tune in the method and strength for calculating the reconnection probability for colored partons in the PS. The two




Figure 2.74: Particle-level distributions comparing PYTHIA 6 and HERWIG for the all particles pull angle (left) and the charged particles pull angle (right).

parameters which differ are MSTP(95) (probability calculation method) and MSTP(78) (strength of the connection). In the lowCR tune, the probability for a string piece to preserve its original connection is given by

$$P_{\text{keep}} = (1 - \zeta \times \text{MSTP}(78))^{n_{\text{int}}},$$

where n_{int} is the number of parton-parton interactions in the current event. The parameter $\zeta^{-1} = 1 + \text{MSTP}(77)^2 \times \langle p_T \rangle^2$ is a way to make this p_T dependent. In all the Perugia tunes, $\text{MSTP}(77)=1$. The probability in the nominal tune is given by

$$P_{\text{keep}} = (1 - \zeta \times \text{MSTP}(78))^{\langle n_s \rangle(y_1, y_2)},$$

where this tries to be more ‘local’ with the function $\langle n_s \rangle(y_1, y_2)$ that counts the number of string pieces (not counting the ones under consideration) between the rapidity endpoints of the piece under consideration y_1 and y_2 . The loCR tune is set to be consistent with the minimum bias data, with as low a CR setup as possible.




Figure 2.75: Particle-level distributions comparing PYTHIA 6 and HERWIG for the all particles pull vector magnitude (left) and the charged particles pull vector magnitude (right).

Figure 2.76 shows that the impact from varying the CR tune is very small at truth level, $\lesssim 1\%$. A similar trend is observed for the flipped model.

2.5.6.3 Initial and Final State Radiation

Figure 2.77 shows the impact on the pull angle distribution due to variations in the ISR/FSR modeling from varying the radiation simulated with ACERMC 3.8 [294] constrained by Ref. [295]. The ISR/FSR could impact the pull angle either directly by introducing more radiation around the two selected jets or indirectly by changing the event kinematics or by changing the number of jets in the event (and thus impact the combinatorics of which jets are selected).




Figure 2.76: Truth level distributions comparing the nominal and low CR tunes of Pythia for the all particles pull angle (left) and the charged particles pull angle (right).




Figure 2.77: The ISR/FSR variations at truth level for the all particle pull (left) and the charged particles pull (right). Electron and muon channels combined.

2.5.6.4 Top p_T

It is well known [281] that the POWHEG-BOX + PYTHIA 6 $t\bar{t}$ simulation at $\sqrt{s} = 8$ TeV with $h_{\text{damp}} = \infty$ (see Ref. [296]) has a slight mis-modeling at high p_T , which is seen clearly in the slope in the ratio plot in Figures 2.17 and 2.18. The $t\bar{t}$ p_T spectrum only enters the measurement indirectly through the unfolding, since the pull angle distribution and the pull angle resolution depend on the jet p_T spectrum. The following procedure is used to assess the impact of this mis-modeling:

- (1) Extract a joint distribution of particle-level pull angle distribution and the particle-level leading jet p_T from the nominal $t\bar{t}$ simulation.
- (2) Generate random points from (1) and smear the angle according to the Fig 2.28²¹. From this step we get a nominal response matrix and a nominal detector-level distribution.
- (3) Generate random points from (1) but re-weight the p_T spectrum (via assign event weights) so that it matches matches the data from Fig. 2.17. This produces a shifted response matrix.
- (4) Unfold the nominal detector-level distribution with the shifted response matrix and compare to the nominal.

For this test, five equally spaced bins between 0 and π are used for the pull angle distribution. The relative change in each bin after doing the comparison in step (4) is shown in Table 2.6. The changes are negligibly small and are ignored for the remainder of the analysis.

2.5.6.5 Color flow Model

For the purpose of comparing the unfolded data with the flipped color model, it is necessary to take into account any potential biases the model has on the unfolding. One way to estimate this uncertainty is to take the difference in the unfolded result

²¹The pull angle in Fig 2.28 is not exactly the same as the one here because of the origin correction. For the purpose of this test, the differences are sub-dominant.

Bin Number	All Particles	Charged Particles
1	0.01%	0.13%
2	-0.07%	-0.08%
3	-0.08%	-0.04%
4	0.05%	0.006%
5	0.06%	-0.02%

Table 2.6: The impact on the unfolded jet pull angle distribution from re-weighting the jet p_T spectrum to match the data.

when using the nominal versus flipped model for the response matrix. This procedure is excessively conservative, because the model dependence is already covered by the non-closure uncertainty described in Sec. 2.5.8. Therefore, for any result other than a comparison between the unfolded data and the flipped model, this uncertainty should not be included. Since the pull angle distribution for the flipped sample is significantly different than the nominal one, Sec. 2.4.1 suggests that this uncertainty may not be small. This is confirmed by Fig. 2.78. As desired, the color flow model uncertainty is significantly smaller than the color flow model difference, but is still $\sim 1\%$ in some bins.

2.5.6.6 Other

Other sources of uncertainty include the choice of factorization and renormalization scale in the ME calculation and the PDF [219]. As observed in Sec. 2.5.6.1, these variations have little impact on the color flow. Additionally, varying the top quark mass by ± 1 GeV has a negligible impact on this measurement.

2.5.7 Correction Factors

Uncertainties in the correction factors from Sec. 2.4.2 are accounted for as part of all other uncertainties described thus far. The fake and inefficiency factors are modified in addition to the response matrix for all the sample variations. As an example,




Figure 2.78: The fractional difference between the nominal particle-level distribution and the nominal detector-level distribution unfolded with a response matrix constructed from the flipped sample (*unfolding difference*). The dashed line shows the fractional difference at particle-level between the two models. All distributions are normalized to unity before computing fractional differences.

Fig. 2.79 (2.80) shows the variation in the fake (inefficiency) factor for the various ME and fragmentation models considered in Sec. 2.5.6.1. The correction factors are largely independent of the pull angle, and the largest uncertainty is on the overall acceptance from the fragmentation model. PYTHIA and HERWIG predict $\sim 3\%$ differences in the fake factors and $\sim 15\%$ in the inefficiency factors. Since the unfolded distributions are normalized to unity for the final result, overall differences in acceptance from the unfolding have no effect on the measurement.

2.5.8 Non-closure

The non-closure uncertainty uses the same data-driven technique that is described in detail in Sec. 1.5.4. As the detector-level simulation agrees well with the data (Fig. 2.40), the amount of reweighting is minimal.




Figure 2.79: The fake factors for the all-particles pull angle (left) and the charged-particles pull angle (right). The (1) and (8) in the legend refers to the nominal and flipped sample, respectively.




Figure 2.80: The inefficiency factors for the all-particles pull angle (left) and the charged-particles pull angle (right). The (1) and (8) in the legend refers to the nominal and flipped sample, respectively.

2.5.9 Summary

The systematic uncertainties are estimated by unfolding the data with varied response matrices or by subtracting varied background predictions from the data. Table 2.7 summarizes the various sources of systematic and statistical uncertainty for both the all-particles and charged-particles pull angles. The modeling uncertainties dominate and the total uncertainty is about 3% for the all-particles pull angle and about 2% for the charged-particles pull angle.

Uncertainty [%]	$\theta_p^{\text{all}} [\text{rad}]/\pi$			$\theta_p^{\text{charged}} [\text{rad}]/\pi$			
	0.0 - 0.275	0.275 - 0.6375	0.6375 - 1.0	0.0 - 0.2	0.2 - 0.5	0.5 - 0.8	0.8-1.0
$t\bar{t}$ NLO generator	1.61	0.50	1.00	0.94	0.17	0.05	1.47
Fragmentation Model	1.61	0.98	0.48	0.52	0.31	0.46	0.56
ISR/FSR	1.18	0.61	0.47	0.22	0.04	0.00	0.34
Color reconnection	0.54	0.37	0.92	0.40	0.29	0.16	0.23
MPI	0.20	0.13	0.04	0.59	0.32	0.41	0.42
Color model	1.22	1.70	0.69	1.12	0.18	0.52	0.25
Non-closure	0.47	0.06	0.38	0.61	0.58	0.32	1.19
JES	0.43	0.18	0.49	0.22	0.15	0.16	0.00
JER	0.27	0.01	0.26	0.03	0.12	0.17	0.49
Clusters	0.03	0.06	0.04	N/A			
Tracks		N/A		0.04	0.02	0.05	0.00
Other	0.38	0.13	0.45	0.20	0.15	0.14	0.00
Stats.	1.12	0.63	1.12	0.68	0.51	0.54	0.77
Total	3.20	2.26	2.16	1.97	1.00	1.07	2.26

Table 2.7: Uncertainties in each bin of the all-particle pull angle. The “Other” category includes uncertainties due to the non- $t\bar{t}$ backgrounds.

Figure 2.81 shows the full systematic uncertainty covariance matrix for the experimental and background normalization uncertainties²². The covariance matrix is the sum of the matrices from individual sources of uncertainty. A partition of these uncertainties into four components is displayed in Fig. 2.82. By construction, the cluster

²²There is some ambiguity on the sign of the off-diagonal terms, especially for the modeling uncertainties for which there is no well-defined notion of ‘shift up/down’. For the experimental uncertainties, a natural choice is to take the (signed) uncertainty as nominal - shifted. The covariance matrix for the modeling uncertainties is omitted here, but is revisited in Sec. 2.6.

and tracking uncertainties are only relevant for the all-particles or charged-particles pull angles, but not both. The significant correlation in the first bin of Fig. 2.81 is from the jet energy resolution, shown in the bottom left matrix in Fig. 2.82. The per-bin uncertainties are dominated by the diagonal components.

As observed in Fig. 2.81, the systematic uncertainties induce correlations between bins of the same observable and between the all-particles and charged-particles pull angles. Correlations between the variables are also present from coherent jet-by-jet statistical fluctuations and correlations between bins of the same variable are induced from the unfolding and from normalization. The top right plot of Fig. 2.81 shows that the all-particles and charged-particles pull angles are largely uncorrelated but there is a positive association, as expected ($\rho = 0.23$). The realization of this correlation in the binning used for the measurement is shown in the bottom right plot of Fig. 2.81. By construction, the bins of the same variable are independent of each other and correlations are measured by the off-diagonal blocks. There is a general positive correlation across all off-diagonal bins because the total yield between the two variables is correlated. This general correlation is removed in the bottom left plot of Fig. 2.81 by normalizing per variable. For example, the residual correlation between the first bins of both variables is about 15%. The normalization also induces a significant negative correlation between bins of the same variable due to the small number of total bin: if one bin content fluctuates up, the others have to be lower if the total integral is the same. The amount of the normalization-induced correlation scales with the bin width. After unfolding and normalization, there is still a significant negative correlation between bins of the same variable, but the magnitudes have changed. This is in part due to the large correlation between neighboring bins induced by the fact that the pull angle resolution is not small compared to the range, π .




Figure 2.81: The systematic uncertainty covariance matrix for all experimental and background normalization uncertainties. If the matrix were proportional to the identity matrix, then the systematic uncertainty in bin i of the all-particles pull angle would be $\sqrt{\Sigma_{ii}}$ and in bin i of the charged-particles pull angle would be $\sqrt{\Sigma_{i+3,i+3}}$ for Σ a matrix representing the plot above. The matrix Σ is the sum of the matrices from all individual sources of uncertainty.




Figure 2.82: The same as Fig. 2.81, but broken into four categories: cluster and tracking (top left), background normalization (top right), jet energy resolution (bottom left), and jet energy scale (bottom right).




Figure 2.83: Representations of correlations between and within bins of the all-particles and charged-particles pull angles induced from the unfolding (top left), normalization (bottom left), and jet-by-jet statistical fluctuations (top and bottom right). The unfolding correlation includes normalization. These matrices are generated by running the unfolding procedure on the data 10,000 times, each time generating *pseudo-data* $X_{ij} \sim \text{Poisson}(\lambda_{ij})$, where λ_{ij} is the number of events measured in bin (i, j) . For each pseudo-dataset, the backgrounds are subtracted before (normalization, unfolding, and) computing the correlation.

2.6 Results

The unfolded data are shown in Fig. 2.84 for both the all-particles and charged-particles pull angles. The data generally fall between the SM color flow and the flipped model, though the agreement with the SM is significantly better. Most of the difference between the models is in the first bin and due to the normalization, there is little spread in the distributions for the second bin of both variables. The flipped model is about 2.3σ away from the data in the first bin, while it is about 3.3σ in the first bin for the charged-particles pull angle. It is possible to quantify the compatibility using all bins by computing the probability distribution of the log likelihood ratio²³. Assuming the distribution of the pull angle follows a multivariate Gaussian distribution, the log likelihood ratio is (up to constants) simply the difference in χ^2 :

$$\log(\mathbf{p}_{\text{SM}}/\mathbf{p}_{\text{flipped}})(\vec{\mathbf{x}}) = \sum_{i=1}^4 \frac{(\mathbf{x}_i - \mathbf{x}_{i,\text{SM}})^2}{\sigma_i^2} - \sum_{i=1}^4 \frac{(\mathbf{x}_i - \mathbf{x}_{i,\text{flipped}})^2}{\sigma_i^2}, \quad (2.15)$$

where σ_i is the uncertainty on bin i and $\mathbf{x}_{i,M}$ is the i^{th} bin content under model M . The probability distribution of Eq. 2.15 can be evaluated numerically, taking into account correlation between bins, by generating pseudo-data from the measurement covariance matrix. As mentioned in Sec. 2.5.9, the covariance matrix is well-defined for the experimental systematic and statistical uncertainties but is not well-defined for the (dominant) modeling uncertainties. Despite this, one can estimate the mutlibin sensitivity by selecting a convention; in this case, the sign of the covariance matrix is from the varied sample prediction subtracted from the nominal prediction. Figure 2.85 shows the probability distribution for Eq. 2.15 (charged-particles pull angle) under both the SM and flipped hypotheses using ten million pseudo-experiments. The test statistic for the data is indicated by an arrow. The data is inconsistent with the flipped model at about 4σ (observed) while the nominal MC is inconsistent with the flipped model at about 5σ (expected), for $\sigma = \Phi^{-1}(1 - \text{p-value})$, for the Gaussian cumulative

²³The likelihood ratio test is the most *powerful* by the Neyman-Pearson lemma. See Sec. 7.1 for details. Since the logarithm is a monotonic function, the log-likelihood ratio test is also the most powerful.

distribution function Φ . A similar exercise with the all-particles pull angles results in lower significances such that the full combination of the two variables is dominated by the charged-particles pull angle significance.




Figure 2.84: Left (Right): the all-particles (charged-particles) pull angle distribution for the unfolded data and three particle-level simulations. The orange inner band on the data represents the statistical uncertainty while the yellow band is the sum in quadrature of the statistical and systematic uncertainty. Final version of this plot is from T. Neep.

Interestingly, the SM color flow from POWHEG-BOX+PYTHIA 6 is closer to the data than POWHEG-BOX+HERWIG in Fig. 2.84. With only the first bin, the HERWIG model is about 2σ away from the SM²⁴ while the PYTHIA 6 model is $\lesssim 1\sigma$ away. Both the HERWIG and PYTHIA 6 model predictions are further away from the flipped model than the data is from the flipped model. The unfolded data and particle-level analysis code are publicly available [33, 297] for further interpretation and can provide useful information for the tuning and model development of color flow.

²⁴Using only the charged-particles pull angle and ignoring the explicit color flow model uncertainty - see Sec. 2.5.6.5.




Figure 2.85: The distribution of the log-likelihood ratio test statistic under the SM and flipped color flow models. To generate the distributions, the full covariance matrix is resampled ten million times. In order to impose unitarity of the pseudo-data, only the first n bins are sampled and the $(n - 1)^{\text{th}}$ is fixed by normalization. Additionally removing this last bin from the log-likelihood ratio has little impact on the approximate significances quoted in the text.

2.7 Summary

The analysis presented in this chapter describes a measurement of the orientation of radiation from jets identified as originating from a W boson in $t\bar{t}$ events. The measurement uses 20.3 fb^{-1} of $\sqrt{s} = 8 \text{ TeV}$ pp collision data recorded by the ATLAS detector at the LHC. To quantify the distribution of energy inside one jet relative to another, the distribution of the *jet pull angle* is extracted from the data using information from both the ATLAS calorimeter and tracking detectors. The jet pull angle is found to correctly characterize the W boson as a color singlet, with data disfavoring an alternative color-octet model at greater than 3σ . This illustrates the potential to use the jet pull angle in future SM measurements and BSM searches. The jet pull angle measurement is presented as a normalized fiducial $t\bar{t}$ differential cross-section, allowing the results to be used to constrain implementations of color connection.

Chapter 3

Constituent Multiplicity

As discussed in earlier chapters, quarks and gluons produced in high-energy particle collisions hadronize before they can be observed directly. However, the properties of the resulting jets depend on the type of parton which initiated them. One jet observable sensitive to the quark or gluon nature is the number of charged particles inside the jet. Due to their larger color-charge under the strong force, gluon-initiated jets contain on average more particles than quark-initiated jets and the average (charged) particle multiplicity inside jets increases with jet energy [22]. These properties were used recently at the Large Hadron Collider (LHC) to differentiate between jets originating from a quark or a gluon [121, 167, 298–300]. These studies have found significant differences in the charged-particle multiplicity between the available simulations and data. Improved modelling based on measurements of the number of charged particles inside jets is thus crucial for future studies.

This chapter presents a measurement of the average charged-particle multiplicity inside jets as a function of the jet transverse momentum in dijet events in pp collisions at $\sqrt{s} = 8$ TeV with the ATLAS detector¹. The measurement of the charged-particle multiplicity inside jets has a long history from the SPS [302–304], PETRA [305, 306], PEP [307–310], TRISTAN [311], CESR [312], LEP [313–324], and the Tevatron [325]. At the LHC, both ATLAS [236, 326] and CMS [327] have measured the charged-particle multiplicity inside jets at $\sqrt{s} = 7$ TeV. One ATLAS result used jets that are

¹This analysis has been published in Ref. [301].

reconstructed with tracks and have transverse momentum less than 40 GeV. A second ATLAS analysis [236] has measured charged particles inside jets with transverse momenta spanning the range from 50 to 500 GeV with approximately constant 3–4% uncertainties. The CMS measurement spans jet transverse momenta between 50 and 800 GeV with 5–10% uncertainties in the bins of highest transverse momentum. The analysis presented here uses the full $\sqrt{s} = 8$ TeV ATLAS dataset, which allows for a significant improvement in the precision at high transverse momentum up to and beyond 1.5 TeV.

This chapter is organized as follows. Section 3.1 describes in more detail the motivation for a measurement of the charged particle multiplicity, including some theoretical considerations from QCD. The setup of the analysis, the corrections to remove detector distortions, and the systematic uncertainties, which are similar to the techniques used for the jet charge measurement (Chapter 1), are discussed in Sec. 3.2, Sec. 3.3, and Sec. 3.4, respectively. The results are in Sec. 3.5, both inclusive and exclusive in jet type. Section 3.6 ends the chapter with a summary and outlook.

3.1 Motivation

Despite being a basic jet quantity, the constituent multiplicity is non-trivial to describe precisely in perturbative QCD due to its sensitivity to very soft energy scales. Section 3.1.1 describes lowest order and state-of-the-art calculations that attempt to recover perturbative predictions for multiplicity. In addition to providing a basic probe of QCD at the highest energies, the constituent multiplicity is an important discriminant between quark and gluon initiated jets, as motivated in Sec. 3.1.2.

3.1.1 QCD Predictions for Multiplicity

The average particle multiplicity inside a jet was calculated in Sec. 1.1.3:

$$\langle n_p(E) \rangle = \sum_h \int_0^1 dz D_p^h(z, E), \quad (3.1)$$

where D_p^h is the fragmentation function describing the probability to find a hadron h with energy fraction z of the parton p . One could try to compute the p_T dependence of $\langle n_p \rangle$ using similar techniques as for the jet charge for which the multiplicity is related to the $\kappa \rightarrow 0$ limit. However, $\tilde{P}_{q \leftarrow q}(\kappa \rightarrow 0) \rightarrow \infty$ and so Eq. 1.11 cannot be used. One way to make a sensible lowest order prediction for the multiplicity is to include the suppression of large angle soft radiation due to color coherence (see Sec. 2.1.1). Color coherence can be incorporated into the DGLAP equation by using $t = E\theta$ instead of μ as the evolution variable, where E is the parton energy and θ is the opening angle of the radiation. Evolution to smaller values of θ is the *angular ordering* scheme. In this case, the equation governing the scale-dependence of $\tilde{D}$ is given by

$$t \frac{\partial}{\partial t} D_p^h(z, t) = \sum_{p'} \int_z^1 \frac{dz'}{z'} \frac{\alpha_s P_{p' \leftarrow p}(z')}{\pi} D_{p'}^h\left(\frac{z}{z'}, z' t\right), \quad (3.2)$$

which is identical to Eq. 1.9 except that the last term has an explicit dependence on

the integrand in its second coordinate. The corresponding equation for the Mellin moment of $\tilde{D}$ is

$$\begin{aligned}
 t \frac{\partial}{\partial t} \tilde{D}_p^h(\kappa, t) &= \frac{\alpha_s}{\pi} \sum_{p'} \int_0^1 dz z^\kappa \int_z^1 \frac{dz'}{z'} P_{p' \leftarrow p}(z') D_{p'}^h\left(\frac{z}{z'}, z' t\right) \\
 &\stackrel{x=z/z'}{=} \frac{\alpha_s}{\pi} \sum_{p'} \int_0^1 dz' (z')^\kappa P_{p' \leftarrow p}(z') \int_0^1 dx x^\kappa D_{p'}^h(x, z' t) \\
 &= \frac{\alpha_s}{\pi} \sum_{p'} \int_0^1 dz' (z')^\kappa P_{p' \leftarrow p}(z') \tilde{D}_{p'}^h(\kappa, z' t)
 \end{aligned} \tag{3.3}$$

which does not fully factor like Eq. 1.10. Nonetheless, one can try a solution of the same form that solves Eq. 1.10: $\tilde{D}(\kappa, t) \propto t^{\gamma(\kappa)}$ (γ is called the anomalous dimension). With this ansatz, Eq. 3.3 becomes

$$\gamma(\kappa) = \frac{\alpha_s}{\pi} \sum_{p'} \int_0^1 dz' (z')^{\kappa+\gamma(\kappa)} P_{p' \leftarrow p}(z'). \tag{3.4}$$

The most relevant regime is $z' \ll 1$, where the integral of the splitting function diverges using the original ordering scheme in Sec. 1.1.3. In this regime, $P_{p' \leftarrow p}(z') \approx \frac{2C}{\pi} \frac{1}{z} \delta_{pg}$, where $C = C_F$ for quarks and $C = C_A$ for gluons. Therefore,

$$\gamma(\kappa) \approx \frac{2\alpha_s C}{\pi} \int_0^1 dz' (z')^{\kappa+\gamma(\kappa)-1} \tag{3.5}$$

$$= \frac{2\alpha_s C}{\pi} \frac{1}{\kappa + \gamma(\kappa)}, \tag{3.6}$$

which is readily solved for γ :

$$\gamma = -\frac{\kappa}{2} + \sqrt{\frac{\kappa^2}{4} + \frac{2\alpha_s C}{\pi}}. \tag{3.7}$$

As desired, Eq. 3.7 is finite as $\kappa \rightarrow 0$. The difference with the solution in Sec. 1.1.3 is that Eq. 3.7 is the start of a series that is in powers of *the square root of* α_s . This is not the Taylor series of any function and thus the convergence of the series is not governed in the usual way for a perturbative series in α_s . This $\sqrt{\alpha_s}$ behavior has been observed and catalogued for a variety of related variables [328, 329] (*Sudakov safe*). For comparison, one could expand Eq. 3.7 in α_s and compare with the energy ordered calculation from earlier, using the gluon splitting function instead of the quark one:

$$\gamma_{\text{angular ordered}} = \text{finite} + \frac{2\alpha_s C}{\pi\kappa} + \mathcal{O}(\alpha^2) \quad (3.8)$$

$$\gamma_{\text{energy-ordered}} = \tilde{p}_{g \leftarrow p} = \frac{2C\alpha_s}{\pi} \int_0^1 \frac{dz}{z} = \frac{2\alpha_s C}{\pi\kappa} + \mathcal{O}(\alpha_s^2). \quad (3.9)$$

Equation 3.8 shows the importance of the $\sqrt{\alpha_s}$ expansion to recover a finite prediction, which is not achievable with any finite α_s expansion. Inserting the anomalous dimension from Eq. 3.7 into the ansatz $\tilde{D}(\kappa, t) \propto t^{\gamma(\kappa)}$ results in:

$$\langle n_p(E) \rangle \propto p_T^\gamma = \exp \left(\sqrt{\frac{2\alpha_s C}{\pi}} \log(p_T/\Lambda) \right) \sim \exp \left(\sqrt{C \log(p_T/\Lambda)} \right), \quad (3.10)$$

where the last line uses² $\alpha_s(p_T) \sim 1/\log(p_T/\Lambda)$. For quark and gluon jets, the first gluon emission is proportional to C_F and C_A , respectively. However, the subsequent parton shower is dominated in the $z \rightarrow 0$ limit by the gluon splitting function $g \rightarrow gg$ because the conversion of gluons back into quarks, $g \rightarrow q\bar{q}$, is suppressed by a factor of α_s . Therefore,

$$\langle n_p(E) \rangle \propto C_i \exp \left(\sqrt{C_A \log(p_T/\Lambda)} \right), \quad (3.11)$$

where $i = F$ for $p = \text{quark}$ and $i = A$ for $p = \text{gluon}$. The main features of Eq. 3.11

²This can be properly derived by including the running α_s in the ansatz for $\tilde{D} \propto \exp(\gamma \log(t)) \propto \exp(\int_{t_0}^t \gamma(\alpha_s(t')) dt'/t')$. See e.g. Sec. 6.1 in Ref. [22] for details.

are that the multiplicity increases with $\mathbf{p}_T$ and is larger for gluon jets than for quark jets. At lowest order, the ratio of quark to gluon multiplicity is a constant C_F/C_A .

The calculation of the anomalous dimension γ can be systematically improved as a series in $\sqrt{\alpha_s}$ despite the lack of control in α_s . Currently, the most precise calculation in this context is at next-to-next-to-next-to-leading-order (N³LO) using pQCD [330, 331]:

$$\langle n_g(\mathbf{y}) \rangle \propto \exp \left(f_{\text{LO}} \sqrt{\mathbf{y}} + f_{\text{NLO}} \log(\mathbf{y}) + f_{\text{N}^2\text{LO}}(\mathbf{y}) \frac{1}{\sqrt{\mathbf{y}}} + f_{\text{N}^3\text{LO}}(\mathbf{y}) \frac{1}{\mathbf{y}} \right) \quad (3.12)$$

$$\langle n_q(\mathbf{y}) \rangle = \frac{\langle n_g(\mathbf{y}) \rangle}{r_0(1 - r_1\gamma_0 - r_2\gamma_0^2 - r_3\gamma_0^3)}, \quad (3.13)$$

where $\mathbf{y} = \log(\mathbf{t}/\Lambda)$ and

$$\begin{aligned} f_{\text{LO}} &= 2C \\ f_{\text{NLO}} &= -a_1 C^2 \\ f_{\text{N}^2\text{LO}}(\mathbf{y}) &= C \left(2a_2 C^2 + \frac{\beta_1}{\beta_0^2} (\log(2\mathbf{y}) + 1) \right) \\ f_{\text{N}^3\text{LO}}(\mathbf{y}) &= C^2 \left(a_3 C^2 - \frac{a_1 \beta_1}{\beta_0^2} (\log(2\mathbf{y}) + 1) \right). \end{aligned} \quad (3.14)$$

The values of r_i and a_i are in Table 3.1. At the N³LO, the ratio of the quark and gluon jet multiplicities does vary with $\mathbf{p}_T$, though the overall C_F/C_A scaling is preserved.

Coefficient	Order			
	0	1	2	3
quark-gluon ratio r	$C_A/C_F = 2.25$	0.198	0.510	-0.041
gluon $\mathbf{p}_T$ scaling a	–	0.314	-0.301	0.112

Table 3.1: The coefficients of various parts of the N³LO prediction (Ref. [330, 331]) for the quark and gluon constituent multiplicity distributions from Eq. 3.13 and 3.14.

3.1.2 Quark and Gluon Tagging

In addition to using the charged particle multiplicity in jet to directly test the predictions of QCD, constituent multiplicity is a ubiquitous feature for discriminating quark jets from gluon jets. For example, the jet energy response in the ATLAS calorimeter differs between quark and gluon jets. The most recent jet calibration procedure, described in Sec 4.3, uses a residual correction based on the number of tracks inside jets. For a fixed jet energy, the jet energy response is lower for higher constituent multiplicity jets due to the non-linear calorimeter response. The impact of this residual calibration is shown in Fig. 3.1. The bottom panels show that the constituent multiplicity increases with jet p_T and the response decreases with the number of tracks. As a result of the residual calibration, the jet energy scale is less dependent on the composition of quark and gluon jets for a particular analysis since the distribution of the number of tracks inside jets significantly differs between quark and gluon jets.

Many SM processes and new physics scenarios of interest are produced with predominately quark jets. For example, top quark and W bosons decaying hadronically produce mostly quark jets and cascade decays of SUSY squarks or gluinos can result in high multiplicity quark jet final states. Therefore, it is desirable to have a tool that can differentiate quark jets from gluon jets. A dedicated performance study using early Run I data showed that n_{track} as well as the p_T - and ΔR -weighted sum of tracks (track width) inside jets are good variables for this task. However, these track-based variables have different distributions in data and simulation. As a result, the tagger performance in simulation is optimistic. Figure 3.2 shows the 2D likelihood ratio used for the tagger. There are qualitative differences between the two distributions, in particular the large likelihood in the lower left corner in simulation that is not as significant in data. The implication of this study is that quark/gluon tagging is significantly mis-modeled and one likely source³ is the modeling of jet fragmentation. An improved model of the number of particles inside jets is crucial for improved descriptions of quark/gluon tagging in the future.

³This analysis did not assess the systematic uncertainties related to the modeling of the detector response, which could account for some of the differences between data and simulation.




Figure 3.1: The response after applying an inclusive jet energy calibration (left) and after the residual correction (right). The lower panel is the distribution of the number of tracks inside jets for three jet p_T ranges. Reproduced from Ref. [125].

Despite the known mis-modeling of the n_{track} distribution, both ATLAS and CMS have used n_{track} -based quark/gluon tagging to search for new physics, taking care to assess the impact of potential sources of systematic bias. One prominent example is the ATLAS search for all-hadronic diboson resonances. Figure 3.3 shows the final dijet invariant mass spectrum before and after applying a requirement on the number of tracks. The signal to background ratio for the 2 TeV W' model increases because the W and Z bosons from the W' decay predominately produce quark jets. Interestingly, the excess in data also increases with this requirement.

The particle multiplicity inside jets is a powerful tool for probing the high energy behavior of QCD as well as for distinguishing quark jets from gluon jets to improve the significance of other SM measurements and searches for new physics beyond the SM. The remainder of this chapter describes a measurement of the charged particle multiplicity inside jets, exploiting both aspects of this tool.




Figure 3.2: The two-dimensional likelihood ratio for the track multiplicity (vertical axis) and track width (horizontal axis) quark/gluon tagger in data (left) and simulation (right). Reproduced from Ref. [167].




Figure 3.3: The dijet invariant mass spectrum for the full jet selection without the n_{track} requirement (left) and with all requirements (right). Reproduced from Ref. [298].

3.2 Analysis Design

In a fixed p_T bin, number of charged particles inside jets is nearly identical to the jet charge with the momentum-weighting factor $\kappa = 0$. As a result, most of the framework for the jet charge measurement can be re-used to perform the measurement of the p_T -dependence of $\langle n_{\text{track}} \rangle$. In particular, events are selected using single jet triggers and required to have at least two jets with $p_T > 50$ GeV that are well-balanced in p_T . One new p_T bin is added at $p_T > 1.5$ TeV and the measurement is performed for three track p_T thresholds (500 MeV, 2 GeV, and 5 GeV) in order to investigate the sensitivity of the modeling to the softness of the radiation. Figure 3.4 shows the track multiplicity ($p_T^{\text{track}} > 500$ MeV) in three jet p_T bins. As expected, the average n_{track} increases with jet p_T . Interestingly, the data distributions are largely between the predictions from PYTHIA 8 with the AU2 tune and HERWIG++ 2.63 with the EE3 tune.




Figure 3.4: The number of reconstructed tracks associated with a jet in three example jet p_T ranges for data and for PYTHIA 8 and HERWIG++ predictions. The data points have statistical uncertainties which in all bins are smaller than the marker size.

The simulation samples are the same as for the jet charge measurement except for three new particle-level models representing the latest underlying event tunes of PYTHIA 8 and HERWIG++. The differences between these models and the older ones will be discussed in the context of the unfolded results in Sec. 3.5. The details of the samples used are shown in Table 3.2.

ME Generator	PDF	Tune
PYTHIA 8.175 [189]	CT10 [194]	AU2 [190]
PYTHIA 8.186	NNPDF2.3 [32]	Monash [332]
PYTHIA 8.186	NNPDF2.3	A14 [333]
HERWIG++ 2.6.3 [195, 196]	CTEQ6L1 [334]	UE-EE3 [197]
HERWIG++ 2.7.1 [335]	CTEQ6L1	UE-EE5 [336]
PYTHIA 6.428 [96]	CTEQ6L1	P2012 [212]
PYTHIA 6.428	CTEQ6L1	P2012RadLo [212]
PYTHIA 6.428	CTEQ6L1	P2012RadHi [212]

Table 3.2: Monte Carlo samples used for measuring and studying the charged particle multiplicity inside jets.

In analogy to the procedure for the jet charge, the distribution of the jet $\mathbf{p}_T$ and $\mathbf{n}_{\text{track}}$ are discretized into a two-dimensional histogram. This histogram is *unfolded* to remove detector distortions. The average $\mathbf{n}_{\text{track}}$ is computed in each $\mathbf{p}_T$ bin, which is compared to a variety of particle-level models. In addition to studying the inclusive modeling of the $\mathbf{n}_{\text{track}}$ distribution, a novel technique is employed to extract the average charged particle multiplicity separately for quark and gluon jets. As discussed in the context of the jet charge measurement, the distribution of the jet type depends on rapidity. The more forward jet in dijet events is more likely to be the quark jet because the higher longitudinal momentum indicates a higher momentum fraction of the colliding proton. However, the scale of the shower is largely unaffected by the longitudinal momentum and therefore for a fixed jet $\mathbf{p}_T$, the difference in $\langle \mathbf{n}_{\text{charge}} \rangle$ between the more forward and the more central jet is due to the difference in the quark/gluon composition. Figure 3.5 shows the gluon jet fraction of the selected




Figure 3.5: The simulated fraction of jets originating from gluons as a function of jet p_T for the more forward jet (down triangle), the more central jet (up triangle), and the difference between these two fractions (circle). The fractions are derived from PYTHIA 8 with the CT10 PDF set and the error bars represent the PDF and matrix element uncertainties, further discussed in Sect. 4.1.2.3.4. The uncertainties on the fraction difference are computed from propagating the uncertainties on the more forward and more central fractions, treating as fully correlated.

jets in simulation as a function of the jet p_T . The fraction of gluon jets decreases monotonically as a function of jet p_T due to the higher fraction of momentum carried on average by quarks in the proton. However, the *difference* in the fractions between the more forward and more central jet peaks around $p_T \sim 350 \text{ GeV}$ and goes to zero at low and high jet p_T . Given the quark and gluon fractions $f_{q,g}^{f,c}$ with $f = \text{more forward}$, $c = \text{more central}$, $q = \text{quark}$, $g = \text{gluon}$ and $f_q + f_g = 1$, the average charged-particle

multiplicity for quark- and gluon-initiated jets is extracted by solving the system of equations in Eq. 3.15.

$$\begin{aligned}\langle n_{\text{charged}}^f \rangle &= f_q^f \langle n_{\text{charged}}^q \rangle + f_g^f \langle n_{\text{charged}}^g \rangle \\ \langle n_{\text{charged}}^c \rangle &= f_q^c \langle n_{\text{charged}}^q \rangle + f_g^c \langle n_{\text{charged}}^g \rangle.\end{aligned}\tag{3.15}$$

Figure 3.6 shows the closure of the extraction method based on Eq. 3.15. The filled circles show the more forward and more central $\langle n_{\text{charge}} \rangle$ as a function of the jet p_T , which are nearly identical at low and high p_T and are maximally different around $p_T \sim 350$ GeV as expected based on the discussion above and Fig. 3.5. The $\langle n_{\text{charge}} \rangle$ for the more forward gluons (quark) and the more central gluons (quarks) are identical with each other (upper ratio) and with the extracted $\langle n_{\text{charge}} \rangle$ gluon (quark) distribution (lower ratio) within MC statistical uncertainty. The small non-closure at low and high p_T is due in part to the effective number of MC events in those regions is very small due to the negligible difference between the more forward and the more central jet $\langle n_{\text{charge}} \rangle$. This method has several benefits compared to similar techniques for extracting quark and gluon jet properties. First, because only one sample is used for the entire extraction (as opposed to using e.g. γ +jets and dijets), the sample dependent differences between quark and gluon jets are suppressed⁴. Second, as the same events are used for the more forward and the more central jet, many of the experimental uncertainties cancel. This is also true of the PDF uncertainties: the absolute quark and gluon fractions have a bigger uncertainty than the rapidity-dependent differences between the quark and gluon fractions.

⁴This advantage only holds before comparing to quark and gluon jets from a different topology.




Figure 3.6: The p_T dependence of $\langle n_{\text{charged}}^f \rangle$, $\langle n_{\text{charged}}^c \rangle$, $\langle n_{\text{charged}}^g \rangle$, and $\langle n_{\text{charged}}^q \rangle$ in simulation for PYTHIA 8 AU2 as well as the extracted values of $\langle n_{\text{charged}}^g \rangle$, and $\langle n_{\text{charged}}^q \rangle$ using the procedure described in the text.

3.3 Unfolding

The procedure for removing detector distortions is the same procedure as was used for the jet charge measurement. In particular, the measurement is carried out within a fiducial volume matching the experimental selection to avoid extrapolation into unmeasured kinematic regions that have additional model dependence and related uncertainties. The particle level definitions, described in Sec. 1.2.2.2, are constructed to be as close as possible to the corresponding measured objects. For the jet charge measurement, the charged particle p_T threshold was irrelevant because of the p_T -weighting factor κ . However, the charged particle multiplicity is maximally infrared-sensitive and so it is crucial to specify a particle-level p_T threshold on the charged particles. In this case, the same threshold (500 MeV, 2 GeV, or 5 GeV) that is used for tracks is used for charged particles. The unfolding is performed over 11 bins in jet p_T : [0.5,1), [1,2), [2,3), [3,4), [4,5), [5,6), [6,8), [8,10), [10,12), [12,15), and [15, ∞) $\times$ 100 GeV. For the jet charge, there was no natural binning and the choice of bin size was chosen based on the resolution of the distribution. For the charged particle multiplicity, there is a natural bin size: one track. Figure 3.7 shows that it is important to use this binning scheme. For any coarser binning, there is a bias in the average charged particle multiplicity introduced when recovering the mean from the full distribution. In principle, one can correct for this bias, but since it is easy to remove and the unfolding can handle the large number of bins, a one-track-per-bin scheme is used. Another aspect of the binning is the total range. Figure 3.8 shows the fraction of events with more than 60 charged particles. Even in the highest p_T bin, this fraction is below 0.1%, so 60 bins is a conservative range for the measurement. Therefore there are 61 (including no charged particles/tracks) $\times$ 11 = 671 total bins in the measurement.

Figure 3.9 shows the *fake* and *inefficiency* factors as a function of the bin number $i = 1, \dots, 671$ that are applied in simulation before the response matrix can be used to perform the unfolding. There are some structures that are similar to the analogous figure for the jet charge measurement (Sec. 1.4.1.2), such as the generally decreasing correction as a function of jet p_T . However, the within p_T -bin structure is new -




Figure 3.7: The impact of binning the n_{charged} distribution at particle-level. The notation 1-1 means that there is one bin per charged particle multiplicity.




Figure 3.8: The fraction of events with more than 60 charged particles.

for the jet charge the fake and inefficiency factors are largely independent of the jet charge. These structures are due to the jet calibration - the jets in this measurement do not have the residual track-based correction from the global sequential calibration.

The response is lower for jets with a large number of tracks and so there are cases where an event does not pass the jet p_T symmetry requirement at particle-level but does at detector-level as a result of the lower response, leading to the cyclic dips in the right plot of Fig 3.9⁵.




Figure 3.9: For each bin of the combined jet p_T and n_{track} distribution, the inefficiency factors (left) and the fake factors (right) for the more forward jet.

After the correction factors are applied, the two-dimensional distribution of the n_{track} and jet p_T is unfolded using the same iterative Bayesian (IB) technique as for the jet charge measurement. The number of iterations, trading off unfolding bias with statistical fluctuations, is chosen by studying the unfolding bias when unfolding pseudo-data derived from HERWIG++ using a prior distribution and a response matrix derived from PYTHIA. Figure 3.10 shows the bias induced from a variety of iteration choices. The improvement from increasing the number of iterations beyond three is marginal, but to be consistent with the jet charge measurement, four iterations are used for all subsequent results.

The response matrix connects the prior to the posterior distribution in each step of the IB method. Figure 3.11 shows the nominal response matrix from PYTHIA 8.

⁵In principle, the left plot in Fig. 3.9 cannot exceed unity; it appears to do so in a few bins due to rounding errors.




Figure 3.10: The fractional bias induced when unfolding PYTHIA 8 simulation with a HERWIG++ response matrix for various numbers of iterations in the IB technique.

The matrix is nearly diagonal with several structures due to the nature of the binning. In particular, the nearly diagonal stripe in the left plot of Fig. 3.11 corresponds to events that were in the same particle- and detector-level p_T bin. The strip below the diagonal is more prominent than the one above the diagonal because given the jet p_T distribution is steeply falling and thus migrating to lower detector-level jet p_T values is more likely than higher jet p_T values. This effect diminishes as the size of the p_T bin goes to zero. The right plot in Fig. 3.11 shows the response matrix over the 61 n_{charged} bins averaged over all jet p_T bins. The matrix spreads away from the diagonal at high n_{track} due in part to the binomial effect⁶ and there is a bias that the unfolding needs to correct: the average detector-level $n_{\text{track}} < \text{particle-level } n_{\text{charged}}$. This bias increases with jet p_T , as shown in Fig. 3.12. The lower panel of Fig. 3.12 shows the average detector-level n_{track} divided by the particle-level n_{track} in each particle-level n_{track} bin. This offset is bigger for the higher jet p_T bin because of the increased loss of tracks due to hit merging.

⁶If every track is lost with probability p , then for n charged particles, the average number of reconstructed tracks is np and the standard deviation is $\sqrt{np(1-p)}$.




Figure 3.11: The total response matrix (left) and the response matrix for the 61 n_{track} bins averaged over the 11 jet p_T bins. The z-axis is truncated at 1%.




Figure 3.12: The response matrix for the 61 n_{track} bins in the ninth jet p_T bin (1.0 TeV < p_T < 1.2 TeV) on the left and the last jet p_T bin (p_T > 1.5 TeV) on the right. The lower panel shows the average detector-level n_{track} divided by the particle-level n_{track} in each particle-level n_{track} bin. The z-axis is truncated at 1%.

An overview of the unfolding is shown in Fig. 3.13. The top left plot in Fig. 3.13 shows the jet p_T dependence of n_{track} before unfolding for the three track p_T thresholds. As observed earlier, the PYTHIA 8 sample with the AU2 over-predicts the number of tracks inside jets. The relative over-prediction does not largely vary for the three track p_T bin. The top right plot of Fig. 3.13 shows the particle-level and detector-level simulations to illustrate the size of the bias corrected by the unfolding. This is quantified in the bottom right plot of Fig. 3.13, which is the ratio of the solid and dotted lines in the top right plot of Fig. 3.13. Note that this is for illustration purposes - the actual corrections are done over the 671 jet p_T and n_{track} bins and not to the $\langle n_{\text{track}} \rangle$ itself. The unfolded data with statistical uncertainty determined by bootstrapping are shown in the bottom left plot of Fig. 3.13. The next step to determine the quality of the modeling from simulation is to assess sources of systematic uncertainty.




Figure 3.13: The jet p_T dependence of (a) the average reconstructed track multiplicity for uncorrected data and detector-level simulation, (b) the average reconstructed track multiplicity for the detector-level simulation and the average charged-particle multiplicity for the particle-level simulation, (c) the average charged-particle multiplicity for the unfolded data and the particle-level simulation, and (d) the average charged-particle multiplicity divided by the average reconstructed track multiplicity in simulation. For the data, only statistical uncertainties are included in the error bars (which are smaller than the markers for most bins).

3.4 Systematic uncertainties

All stages of the charged-particle multiplicity measurement are sensitive to sources of potential bias. The method for evaluating the uncertainties is the same as for the jet charge measurement, as described below.

3.4.1 Response Matrix

For events that pass both the detector-level and particle-level fiducial selections, the response matrix describes migrations between bins when moving between the detector level and the particle level. The response matrix is taken from simulation and various experimental uncertainties in the charged-particle multiplicity and jet p_T spectra result in uncertainties in the matrix. These uncertainties can be divided into two classes: those impacting the calorimeter-based jet p_T and those impacting track reconstruction inside jets. The dominant uncertainty at high jet p_T is due to the loss of charged-particle tracks in the jet core due to track merging. This charged energy loss uncertainty is estimated using the data/MC differences in the ratio of the track-based jet p_T to the calorimeter-based jet p_T as was also done for the jet charge. More charged energy is lost in the data than in the MC and thus this uncertainty is one-sided. There are other tracking uncertainties in the track momentum scale and resolution, the track reconstruction efficiency, and the rate of tracks formed from random combinations of hits (fake tracks). The uncertainties related to the calorimeter-based jet are sub-dominant (except in the lowest p_T bins) and are due to the uncertainty in the jet energy scale and the jet energy resolution.

3.4.1.1 Charged-energy loss in the dense core of jets

The uncertainty on the charged-energy loss in the dense core of jets is estimated from the modeling of $\sum p_T^{\text{track}}/p_T^{\text{calo jet}}$, as described in Sec 1.5.3.3. The prescription for the uncertainty is to drop tracks randomly with the following probability: $\text{Pr}(\text{drop track } i) = \alpha p_{Ti}^n$, where n is some non-negative integer. The value of α is fixed by requiring the data and MC to agree on the average $\sum p_T^{\text{track}}/p_T^{\text{calo jet}}$. The only free parameter of the prescription is the power n of the track p_T used to model

the uncertainty so one must find the power that has the biggest impact on the observable. In the jet charge measurement, it was found that a very large power was conservative because at high p_T , the largest contribution to the jet charge comes from the highest p_T tracks. However, for n_{track} , a low power is conservative because all tracks are treated equally, independent of their momentum (as long as the p_T is large enough to pass the threshold). This is illustrated in figure 3.14, where the impact of dropping tracks randomly with $\text{Pr}(\text{drop track } i) = \alpha p_{Ti}^n$ for various values of n . Therefore, a power of 0 (i.e. all tracks are dropped with equal probability, regardless of their p_T) is adapted. Note that even for a power of zero the uncertainty increases with p_T because the data/MC difference in $\sum p_T^{\text{track}}/p_T^{\text{calo jet}}$ increases with p_T . The uncertainty is slightly larger for $n > 0$ for a p_T threshold of 5 GeV versus 0.5 GeV, but by construction the uncertainty is independent of the threshold when $n = 0$.




Figure 3.14: The impact of the charged energy loss uncertainty on the average number of charged particles in each p_T bin. The vertical axis is the relative uncertainty on the average unfolded n_{charged} . For each n , tracks are dropped randomly with a probability given by $\text{Pr}(\text{drop track } i) = \alpha p_{Ti}^n$, where α is fixed such that the MC is the same as the data in Fig. 1.61. Below 400 GeV, nuclear interactions dominate the track reconstruction efficiency uncertainty.

3.4.2 Correction Factors

Fake and inefficiency factors are derived from simulation to account for the fraction of events that pass either the detector-level or particle-level fiducial selection, but not both. These factors are generally between 0.9 and 1.0 except in the first jet- p_T interval ($50 < p_T < 100$ GeV), where threshold effects cause the correction factors to take values down to 0.8 (see Fig. 3.9). Experimental uncertainties correlated with the detector-level selection acceptance, such as the jet energy scale uncertainty, result in uncertainties in these correction factors. Another source of uncertainty in the correction factors is the explicit dependence on the particle-level multiplicity and jet p_T spectrum. A comparison of particle-level models (PYTHIA and HERWIG++) is used to estimate the impact on the correction factors. As was also done for the jet charge, the nominal fake and inefficiency factors from PYTHIA 8 are re-weighted to those from HERWIG++ and the unfolding is performed with the nominal PYTHIA 8 response matrix. Figure 3.15 shows the impact of the bin-by-bin re-weighting on the two-dimensional jet p_T and n_{track} distributions as well as their impact on the unfolded $\langle n_{\text{track}} \rangle$ distribution. In general, these uncertainties are $< 0.1\%$. A similar set of plots with nearly the same conclusion is shown for the fake factor uncertainty in Fig. 3.16.




Figure 3.15: The bin-by-bin inefficiency factor fractional uncertainty before unfolding (left) and the inefficiency factor fractional uncertainty on $\langle n_{\text{track}} \rangle$ as a function of the jet p_T bin (right) for track $p_T > 500$ MeV.




Figure 3.16: The bin-by-bin fake factor fractional uncertainty before unfolding (left) and the fake factor fractional uncertainty on $\langle n_{\text{track}} \rangle$ as a function of the jet p_T bin (right) for track $p_T > 500$ MeV.

3.4.3 Unfolding Procedure

The same data-driven non-closure uncertainty technique used for the jet charge measurement is used for $\langle n_{\text{charge}} \rangle$. In particular, the particle-level spectrum is reweighted so that the simulated detector-level spectrum, from propagating the reweighted particle-level spectrum through the response matrix, has significantly improved agreement with the uncorrected data. The modified detector-level distribution is unfolded with the nominal response matrix and the difference between this and the reweighted particle-level spectrum is an indication of the bias due to the unfolding method (in particular, the choice of a prior distribution). The re-weighting factors are simply determined at detector-level and applied at particle-level. These factors are shown in the left plot of Fig. 3.17 and the improvement in the data/MC agreement induced from the particle-level re-weighting is shown in the middle plot of Fig. 3.17. The right plot of Fig. 3.17 illustrates that the particle-level re-weighting brings the $\langle n_{\text{track}} \rangle$ distribution into nearly 100% agreement with the data.




Figure 3.17: The distribution of the weights used to re-weight the MC distribution for the non-closure test (left) and the re-weighted truth distribution (labeled after) (middle) for the more forward jet. The $\langle n_{\text{track}} \rangle$ in data and simulation before and after the re-weighting is shown in the right plot.

The uncertainty on $\langle n_{\text{track}} \rangle$ due to the data-driven non-closure is shown in the left plot of Fig. 3.18 and compared with the raw data/MC difference. Except in the first jet p_T bin where the track multiplicity is spread out over only a small number of bins, the non-closure uncertainty ($\lesssim 1\%$) is much smaller than the raw data/MC

difference ($\sim 10\%$). As a comparison, the relative difference in $\langle n_{\text{track}} \rangle$ when unfolding PYTHIA 8 with HERWIG++ is shown in the right plot of Fig. 3.18. This difference is not used as an uncertainty as it would over-count the non-closure uncertainty, but it is reassuring that the approximate size of the uncertainty is comparable to the data-driven technique. Note that the impact of the difference between PYTHIA 8 and HERWIG++ in the fake/inefficiency factors is already accounted for in Sec. 3.4.2.




Figure 3.18: The non-closure uncertainty for $\langle n_{\text{charged}} \rangle$ compared with the raw data/MC difference in the same quantity (left) and the relative difference in $\langle n_{\text{charged}} \rangle$ when unfolding PYTHIA 8 with itself (= particle-level PYTHIA by construction) and with a HERWIG++ response matrix (right).

3.4.4 Summary

A summary of the systematic uncertainties can be found in Table 3.3 and visualized in Fig. 3.19. The relative size of the uncertainties are similar for the three charged particle p_T thresholds. Aside from the first jet p_T bins, the dominant uncertainties are due to the isolated track reconstruction efficiency and the reconstruction efficiency of tracks inside jets. The statistical and systematic uncertainties are comparable in size ($\sim 4\%$) in the highest jet p_T bin.




Figure 3.19: A visualization of the systematic uncertainties for the $\langle n_{\text{charge}} \rangle$ distribution.

Average n_{charged} Systematic Uncertainty [%]	Jet p_T Range [100 GeV]											
	[0.5,1]	[1,2]	[2,3]	[3,4]	[4,5]	[5,6]	[6,8]	[8,10]	[10,12]	[12,15]	[15,18]	
Response Matrix												
Total Jet Energy Scale	+1.9 -1.9	+0.7 -0.9	+0.6 -0.8	+0.8 -0.7	+0.7 -0.7	+0.6 -0.7	+0.6 -0.7	+0.6 -0.5	+0.4 -0.4	+0.3 -0.3	+0.8 -0.7	
Jet Energy Resolution	+0.6 -0.6	+0.1 -0.1	+0.1 -0.1	+0.1 -0.1	+0.1 -0.1	+0.1 -0.1	+0.1 -0.1	+0.2 -0.2	+0.1 -0.1	+0.1 -0.1	+0.2 -0.2	
Charged Energy Loss	+0.0 -0.0	+0.0 -0.0	+0.0 -0.0	+0.0 -0.0	+1.2 -0.0	+1.1 -0.0	+1.1 -0.0	+1.1 -0.0	+1.0 -0.0	+3.6 -0.0	+3.3 -0.0	
Other Tracking	+1.2 -0.0	+1.0 -0.0	+0.9 -0.0	+0.8 -0.0	+0.8 -0.0	+0.7 -0.0	+0.7 -0.0	+0.7 -0.0	+0.7 -0.0	+0.7 -0.0	+0.8 -0.0	
Correction Factors	+0.1 -0.1	+0.1 -0.1	+0.1 -0.1	+0.1 -0.1	+0.1 -0.1	+0.1 -0.1	+0.0 -0.0	+0.0 -0.0	+0.0 -0.0	+0.0 -0.0	+0.0 -0.0	
Unfolding Procedure	+6.4 -6.4	+3.4 -3.4	+0.6 -0.6	+0.8 -0.8	+0.6 -0.6	+0.4 -0.4	+0.4 -0.4	+0.2 -0.2	+0.2 -0.2	+0.2 -0.2	+0.1 -0.1	
Total Systematic	+6.8 -6.7	+3.6 -3.5	+1.2 -1.0	+1.4 -1.1	+1.7 -0.9	+1.5 -0.8	+1.5 -0.8	+1.4 -0.6	+1.3 -0.5	+3.7 -0.4	+3.5 -0.7	
Data Statistics	0.5	0.2	0.1	0.1	0.0	0.1	0.1	0.3	0.6	1.2	3.9	
Total Uncertainty	+6.8 -6.7	+3.6 -3.6	+1.2 -1.0	+1.4 -1.1	+1.7 -0.9	+1.5 -0.8	+1.5 -0.8	+1.4 -0.6	+1.5 -0.8	+3.9 -1.3	+5.2 -4.0	
Measured Value	7.87	9.87	12.19	13.54	14.59	15.41	16.28	17.41	18.25	18.71	20.78	

Table 3.3: A summary of all the systematic uncertainties and their impact on the n_{track} mean for $p_T^{\text{track}} > 0.5$ GeV and the more forward jet. Uncertainties are given in percent. The last row is the measured average charged particle multiplicity. A value of 0.0 is quoted if the uncertainty is below 0.05%.

3.5 Results

The unfolded average charged-particle multiplicity combining both the more forward and the more central jets is shown in Fig. 3.20 for $p_T^{\text{track}} > 500$ MeV and Fig 3.21 for $p_T^{\text{track}} > 2$ GeV and 5 GeV, compared with various model predictions. As was already observed for the reconstructed data in Fig. 1.14, the average charged-particle multiplicity in data falls between the predictions of PYTHIA 8 and HERWIG++, independently of the underlying-event tunes. The PYTHIA 8 predictions are generally higher than the data and this is more pronounced at higher jet p_T . The default ATLAS tune in Run 1 (AU2) performs similarly to the Monash tune, but the prediction with A14 (the ATLAS default for the analysis of Run 2 data) is significantly closer to the data. A previous ATLAS measurement [326] of charged-particle multiplicity inside jets was included in the tuning of A14, but the jets in that measurement have $p_T \lesssim 50$ GeV. One important difference between A14 and Monash is that the value of α_s governing the amount of final-state radiation is about 10% lower in A14 than in Monash. This parameter has a large impact on the average charged-particle multiplicity, which is shown by the PYTHIA 6 lines in Fig. 3.20 where the Perugia radHi and radLo tunes are significantly separated from the central P2012 tune. The α_s value that regulates final-state radiation is changed by factors of one half and two for these tunes with respect to the nominal Perugia 2012 tune. The recent (and Run 2 default) EE5 underlying-event tune for HERWIG++ improves the modelling of the average charged-particle multiplicity with respect to the EE3 tune (Run 1 default). The general differences between data and simulation are similar for the three track p_T thresholds, but the level of agreement is slightly better for higher thresholds.




Figure 3.20: The measured average charged-particle multiplicity as a function of the jet p_T , combining the more forward and the more central jets for $p_T^{\text{track}} > 0.5 \text{ GeV}$. The band around the data is the sum in quadrature of the statistical and systematic uncertainties. Error bars on the data points represent the statistical uncertainty (which are smaller than the markers for most bins).




Figure 3.21: The measured average charged-particle multiplicity as a function of the jet p_T , combining the more forward and the more central jets for $p_T^{\text{track}} > 2$ GeV (left) and $p_T^{\text{track}} > 5$ GeV (right). The band around the data is the sum in quadrature of the statistical and systematic uncertainties. Error bars on the data points represent the statistical uncertainty (which are smaller than the markers for most bins).

3.5.1 Quark and Gluon Multiplicity

As discussed in Sec. 3.2, the difference in the average charged-particle multiplicity between the more forward and the more central jet is sensitive to the difference between quark and gluon constituent multiplicities. Figure 3.22 shows that the difference is significant for $p_T \lesssim 1.1$ TeV. The shape is governed by the difference in the gluon fraction between the more forward and the more central jet⁷, which was shown in Fig. 3.5 to peak around $p_T \sim 350$ GeV. The systematic uncertainties are significantly smaller on the difference than on the pooled (more forward and more central combined) average n_{track} . For example, at the peak around ~ 350 GeV, the systematic uncertainty is about a factor of three smaller for the difference compared with the combination of the more forward and more central jets.

The average difference, combined with the gluon fraction, can be used to extract the average charged-particle multiplicity for quark- and gluon-initiated jets separately. The extracted p_T dependence of the average charged-particle multiplicities for quark- and gluon-initiated jets is shown in Fig. 3.23. PYTHIA 8 with the CT10 PDF set is used to determine the gluon fractions. The experimental uncertainties are propagated through Eq. 3.15 by recomputing the quark and gluon average charged-particle multiplicities for each variation accounting for a systematic uncertainty; the more forward and more central jet uncertainties are treated as being fully correlated. In addition to the experimental uncertainties, the error bands in Fig. 3.23 include uncertainties in the gluon fractions from both the PDF and matrix element (ME) uncertainties. The PDF uncertainty is determined using the CT10 eigenvector PDF sets and validated by comparing CT10 and NNPDF. The ME uncertainty is estimated by comparing the fractions $f_{q,g}^{f,c}$ from PYTHIA 8 and HERWIG++ after reweighting the PYTHIA 8 sample with CT10 to CTEQ6L1 to match the PDF used for HERWIG++. All PDF re-weighting is performed using LHAPDF6 [337]. The PDF and ME uncertainties are comparable in size to the total experimental uncertainty. As expected, the average

⁷However, the peak is not in exactly the same location because the multiplicity for quarks and gluons is not the same and depends on p_T : $\langle n^c - n^f \rangle = \langle n^c \rangle - \langle n^f \rangle = (f_g^c n_g + f_q^c n_q) - (f_g^f n_g + f_q^f n_q)$, where n is the charged particle multiplicity for quarks (q) or gluons (g) and for the more forward (f) or more central (c) jets.




Figure 3.22: The jet p_T dependence of the difference in the average charged-particle multiplicity ($p_T^{\text{track}} > 0.5 \text{ GeV}$) between the more forward and the more central jet. The band for the data is the sum in quadrature of the systematic and statistical uncertainties and the error bars on the data points represent the statistical uncertainty. Bands on the simulation include MC statistical uncertainty.

multiplicity increases with jet p_T for both the quark-initiated jets and gluon-initiated jets. Furthermore, the multiplicity is significantly higher for gluon-initiated jets than for quark-initiated jets. The average charged-particle multiplicity in PYTHIA 8 with the AU2 tune is higher than in the data for both the quark- and gluon-initiated jets.

In addition to predictions from leading-logarithm parton shower simulations, calculations of the scale dependence for the parton multiplicity inside jets have been performed in perturbative quantum chromodynamics (pQCD). Up to a non-perturbative factor that is constant for the jet p_T range considered in this analysis⁸, these calculations can be interpreted as a prediction for the scale dependence of $\langle n_{\text{charged}} \rangle$ for quark- and gluon-initiated jets. There are further caveats to the predictability of such a calculation since n_{charged} is not infrared safe or even Sudakov safe [328]. Therefore,

⁸This factor is found to be about 0.19 for gluon jets and 0.25 for quark-initiated jets.

$\langle n_{\text{charged}} \rangle$ Systematic Uncertainty	Jet p_T Range [100 GeV]									
	[0.5,1]	[1,2]	[2,3]	[3,4]	[4,5]	[5,6]	[6,8]	[8,10]	[10,12]	[12,15]
Total exp.	+0.44 -0.34	+0.29 -0.24	+0.15 -0.24	+0.24 -0.17	+0.21 -0.33	+0.37 -0.43	+0.48 -0.58	+1.01 -1.03	+2.20 -2.39	+6.09 -6.16
ME	0.04	0.06	0.05	0.12	0.14	0.16	0.06	0.01	0.05	0.22
PDF	+0.01 -0.01	+0.06 -0.05	+0.11 -0.10	+0.18 -0.19	+0.22 -0.27	+0.25 -0.34	+0.30 -0.48	+0.30 -0.60	+0.41 -1.01	+0.23 -0.81
PDF II*	0.03	0.09	0.00	0.04	0.01	0.10	0.33	0.84	1.76	1.69
Half Cone**	0.01	0.03	0.03	0.04	0.03	0.03	0.03	0.02	0.03	0.01
ME ID***	0.06	0.03	0.04	0.05	0.04	0.03	0.01	0.01	0.04	0.05

Table 3.4: A summary of the systematic uncertainties on the average charged multiplicity extraction for gluons. (*) NNPDF versus CT10, used only as a cross-check. (**) Using a cone size of $\Delta R < 0.2$ instead of the nominal 0.4 in the q/g identification. Used only as a cross-check. (***) Matching the jets with the outgoing partons in the ME to do the q/g ID. Used only as a cross-check. The uncertainties are in units of n_{charged} .

the formal accuracy of the series expansion in $\sqrt{\alpha_s}$ is unknown. Given these caveats, the next-to-next-to-next-to-leading-order ($N^3\text{LO}$) pQCD calculation [330, 331] is overlaid in Fig. 3.23 with renormalization scale $\mu = R p_T$ in the five-flavour scheme and $R = 0.4$. The theoretical error band is calculated by varying μ by a factor of two. The prediction cannot give the absolute scale, and therefore the curve is normalized to the data in the second p_T bin ($100 \text{ GeV} < p_T < 200 \text{ GeV}$) where the statistical uncertainty is small. The predicted scale dependence for gluon-initiated jets is consistent with the data within the uncertainty bands while the curve for quark-initiated jets is higher than the data by about one standard deviation.

$\langle n_{\text{charged}} \rangle$ Systematic Uncertainty	Jet p_T Range [100 GeV]									
	[0.5,1]	[1,2]	[2,3]	[3,4]	[4,5]	[5,6]	[6,8]	[8,10]	[10,12]	[12,15]
Total exp.	+0.82 -1.16	+0.36 -0.41	+0.26 -0.28	+0.22 -0.30	+0.25 -0.32	+0.30 -0.35	+0.32 -0.36	+0.41 -0.47	+0.69 -0.67	+1.42 -1.70
ME	0.06	0.23	0.19	0.23	0.22	0.25	0.26	0.22	0.23	0.16
PDF	+0.02 -0.02	+0.11 -0.10	+0.17 -0.16	+0.27 -0.24	+0.33 -0.27	+0.38 -0.28	+0.44 -0.30	+0.47 -0.28	+0.62 -0.33	+0.45 -0.21
PDF II*	0.04	0.01	0.17	0.23	0.17	0.10	0.01	0.21	0.44	0.39
Half Cone**	0.01	0.02	0.02	0.02	0.02	0.01	0.01	0.01	0.01	0.00
ME ID***	0.07	0.03	0.01	0.01	0.01	0.01	0.02	0.02	0.02	0.02

Table 3.5: A summary of the systematic uncertainties on the average charged multiplicity extraction for quarks. (*) NNPDF versus CT10, used only as a cross-check. (**) Using a cone size of $\Delta R < 0.2$ instead of the nominal 0.4 in the q/g identification. Used only as a cross-check. (***) Matching the jets with the outgoing partons in the ME to do the q/g ID. Used only as a cross-check. The uncertainties are in units of n_{charged} .




Figure 3.23: The jet p_T dependence of the average charged-particle multiplicity ($p_T^{\text{track}} > 0.5 \text{ GeV}$) for quark- and gluon-initiated jets, extracted with the gluon fractions from PYTHIA 8.175 with the CT10 PDF. In addition to the experimental uncertainties, the error bands include uncertainties in the gluon fractions from both the PDF and ME uncertainties. The MC statistical uncertainties on the open markers are smaller than the markers. The uncertainty band for the N³LO pQCD prediction is determined by varying the scale μ by a factor of two up and down. The markers are truncated at the penultimate p_T bin in the right because within statistical uncertainty, the more forward and more central jet constituent charged-particle multiplicities are consistent with each other in the last bin.

3.6 Summary

This chapter presents a measurement of the p_T dependence of the average jet charged-particle multiplicity in dijet events from 20.3 fb^{-1} of $\sqrt{s} = 8 \text{ TeV}$ pp collision data recorded by the ATLAS detector at the LHC. The measured charged-particle multiplicity distribution is unfolded to correct for the detector acceptance and resolution to facilitate direct comparison to particle-level models. Comparisons are made at particle level between the measured average charged-particle multiplicity and various models of jet formation. Significant differences are observed between the simulations using Run 1 tunes and the data, but the Run 2 tunes for both PYTHIA 8 and HERWIG++ significantly improve the modelling of the average n_{charge} . Furthermore, quark- and gluon-initiated jet constituent charged-particle multiplicities are extracted and compared with simulations and calculations. As expected, the extracted gluon-initiated jet constituent charged-particle multiplicity is higher than the corresponding quantity for quark-initiated jets and a calculation of the p_T -dependence accurately models the trend observed in the data. The particle-level spectra are available [\[33\]](#) for further interpretation and can serve as a benchmark for future measurements of the evolution of non-perturbative jet observables to validate MC predictions and tune their model parameters.

Chapter 4

Boson and Top Quark Jets

Processes involving the production and decay of W , Z , and H bosons as well as top quarks provide benchmarks for testing the Standard Model (SM), as well as probes of physics beyond the SM (BSM). Since the cross section for the direct strong production of events with multiple jets (QCD multijets) at the LHC is many orders of magnitude larger than for the production of electroweak bosons or top quarks, it is usually the case that leptonic decays must be used to reduce the overwhelming background. This is an unfortunate limitation because the hadronic branching ratios are larger than the leptonic ones¹ and in some BSM theories, new particles similar to the SM electroweak bosons or top quarks do not couple directly to leptons. However, when the momentum of a boson or top quark is comparable with its mass, the spatial proximity of the decay products allows for a new set of tools that can be used to distinguish between single jets from hadronic boson decays and jets originating from QCD multijet backgrounds. Some of these *jet substructure* tools have already been introduced in earlier sections including the jet charge in Sec. 1, jet pull in Sec. 2, and multiplicity in Sec. 3. The most powerful tool is related to another quantum property of jets - the *jet mass*. Before discussing this jet observable in detail, it is important to quantify the size of a jet needed to capture most of the decay products of a boosted boson or top quark.

¹There are more *active* lepton than active quark types (five quarks since $m_{\text{top}} > m_{W/Z/H}$ and six total leptons) but since the electroweak bosons are blind to color, there are many more quarks.




Figure 4.1: A schematic illustration of the setup described in the text to compute the p_T dependence of ΔR . The black dot represents the W boson and the arrows from the dot represent the momentum of the quark decay products.

To illustrate the scaling of the angular distance ΔR between decay products, consider a W boson with momentum directed along the x -axis in the lab frame with magnitude p_T^W and assume $W \rightarrow qq'$. In the W boson rest frame, the two quarks each have energy $m^W/2$ and are back-to-back due to the conservation of energy and momentum. The angular distance in the lab frame will be maximized when the quark momenta are along the y direction (the y direction is the same in both the lab and W boson frames). Figure 4.1 illustrates this setup. Ignoring the quark masses, the momentum in the y direction before and after the boost is $m^W/2$ (momenta orthogonal to the boost direction are unchanged) and the x momentum goes from 0 in the W boson rest frame to $\gamma\beta m^W/2$ in the lab frame. Therefore,

$$\Delta R = \phi \sim \frac{m^W}{\gamma\beta m^W/2} = \frac{2}{\gamma\beta} = \frac{2m^W}{p_T^W}, \quad (4.1)$$

where ϕ is the opening angle between the quarks in the lab frame, β is the speed of the W boson in the lab frame ($\beta = p/E$), $\gamma = \frac{1}{\sqrt{1-\beta^2}} = E/m$ is the usual relativistic enhancement factor and $\gamma\beta = \frac{E}{m} \frac{p}{E} = \frac{p}{m}$. The $\sim$ represents the small angle approximation. The full form is given by

$$\Delta R = 2 \arctan \left(\frac{1}{\gamma \beta} \right) = \frac{2m}{p_T^W} - \frac{2}{3} \left(\frac{m}{p_T^W} \right)^3 + \mathcal{O} \left(\frac{m^5}{p_T^5} \right). \quad (4.2)$$

Since the sub-leading term in Eq. 4.2 is negative and the original setup was chosen to maximize ΔR , in general $\Delta R \geq \frac{2m}{p_T}$ for a particle of mass m and transverse momentum p_T decaying into two massless particles. The full joint distribution of ΔR and p_T is shown in the right plot of Fig. 4.2, ignoring effects of particle spin. The $2m/p_T$ scaling is an excellent approximation for the W decay where the decay products are nearly massless, but there are clear deviations in the case of the top quark where $m_W/m_t \sim 1/2$ is not negligible. The right plot of Fig. 4.2 shows the fraction of events in which the decay products are within $\Delta R < 1$ of the parent boosted W boson or top quark in the production of $t\bar{t}$. At low top quark p_T , the decay products are isotropically distributed. As the top quark boost increases, the W decay products are close enough to be contained within a cone of size $\Delta R = 1$ and then for very large top quark p_T , the b -quark is also contained within the cone.




Figure 4.2: Left: the joint distribution of ΔR and p_T and Right: the fraction of events in which the decay products ($W \rightarrow qq'$ or $t \rightarrow bqq'$) are within $\Delta R < 1$ of the parent particle.

Individual anti- k_t $R=0.4$ jets are an appropriate description of the fragmentation

from the well separated daughter quark decay products of low p_T boson and top quarks. However, isolating these events from the multijet background is an insurmountable challenge. Consider the case of $W \rightarrow qq'$, identified from two jets. If an event has N jets, then there are $\binom{N}{2}$ permutations which could give the W boson daughter jets. All hadronic top quark pair events produce six jets at leading order and there are often several additional jets from initial and final state radiation. A powerful discriminant is the invariant mass of the two jets, m_{jj} . One could require $m_{jj} \sim m_W$ in order to pick the two jets, but then m_{jj} is a less useful discriminant because the background will be sculpted. This is possible because while $m_{jj} \sim m_W$ for the signal, m_{jj} is set by $\sqrt{\hat{s}}$ for the background, which is often near m_W by coincidence. These challenges are resolved at high boson or top quark p_T . In that case, the decay products are geometrically close together, so a large radius jet $R \sim 1$ is likely to capture all of the decay products. Of course, one could always pick R large enough to capture all the decay products of the boson or top quark, but $R \sim 1$ is sufficiently small that jets of this size do not contain significant radiation from other sources. The dijet invariant mass now becomes the large radius jet mass m_j , whose square is defined as the square of the sum of the jet constituent's four-vectors. Just as in the resolved case, $m_j \sim m_W$ for the signal. The power of jet mass is that for the multijet background, $m_j \sim \alpha_s R p_{T,J} \ll \sqrt{\hat{s}}$ (see Sec. 4.1.1). This chapter describes the properties of boson and top quark jets in detail. Section 4.1 focuses on the jet mass, including measurements of the calorimeter jet mass resolution and new alternative jet mass definitions. The jet mass is combined with other jet substructure variables in Sec. 4.2 to distinguish boosted hadronically decaying bosons of different types, a natural extension of isolating these boson jets from multijet backgrounds. The chapter ends in Sec. 4.3 with a new paradigm for studying the rich structure of boson and top quark jets in the context of *machine learning*. State-of-the-art classification techniques are adapted to high energy physics for reconstructing and classifying boosted boson and top quark jets.

4.1 Jet Mass

When a jet is sufficiently large to contain most of the energy from a hadronically decaying boosted boson or top quark, the mass of a jet is approximately the boson or top quark mass. However, the particle-level and detector-level mass resolutions are both significant. At particle-level, the mass is obscured due to finite radius effects and sources of diffuse uncorrelated radiation. The detector-level radiation is affected by both the calorimeter-cell energy and angular resolution. Various techniques for improving the jet mass resolution and measuring its reconstruction properties using data-driven techniques will be discussed in this section. First, section 4.1.1 describes the mechanism by which generic quark and gluon jets acquire mass. Experimental techniques for calibrating the jet mass are discussed in Sec. 4.1.2. Alternative jet mass definitions are investigated in Sec. 4.1.3 and 4.1.4. Concluding remarks and future outlook are provided in Sec. 4.1.5.

4.1.1 The Mass of Quark and Gluon Jets

While the mass of on-shell quarks and gluons is negligibly small compared with the mass of electroweak boson and top quark jets, this is not always true for the mass of a generic QCD jet. Quark and gluon jets acquire significant mass as a result of (relatively) hard or wide angle gluon radiation. Many properties of the jet mass distribution can be well-described within the context of perturbation theory. The jet mass distribution at a hadron collider is known to approximate next-to-next-to-leading logarithmic order (NNLL_p) in the absence of non-global logarithms (extra-jet radiation that re-emits back into the jet) [338] and to even higher order (N³LL+NLO) for hemisphere mass at lepton colliders [339]. This section uses the lowest order results for illustration.

Since the quark masses are small compared with the energy scales relevant at the LHC, the QCD Lagrangian is approximately scale invariant. Consider a quark or gluon of energy E that radiates a gluon with energy e at an angle θ relative to the initial parton direction as depicted in Fig. 4.3. Define the energy fraction $z = e/E$. This section will consider the soft ($z \ll 1$) and collinear ($\theta \ll 1$) region of phase




Figure 4.3: A schematic diagram for the emission of one soft and collinear gluon off of a quark. In this section, $z \ll 1$ and $\theta \ll 1$.

space. As a result of the approximate scale invariance of the Lagrangian, one expects that the probability distribution of z is the approximately the same on all decades. More generally, for fixed $0 < a < b < 1$, $\Pr(a < z < b) = \Pr(ac < z < cb)$ for all $0 < c < 1/b$. In particular, taking the derivative shows that the probability distribution function of z , f_z , has the property $f_z(z) = cf_z(cz)$. As a result, for all $k = cz$, $f(k)k = f(cz)cz = f(z)z$. Therefore, $f(z)z$ must be constant. Let $y = \ln(z)$. Then, the probability density of y , $f_y(y)$, has the form

$$f_y(y) = f_x(x(y)) \left| \frac{\partial x(y)}{\partial y} \right| = f_x(x(y)) e^{\ln(x(y))} = f(x(y)) x(y) = \text{constant}, \quad (4.3)$$

and therefore the logarithm of z is uniformly distributed from $-\infty$ to 0 . The same argument applies for the logarithm of θ . To compute the differential cross section, the phase space is multiplied by the coupling factor α_s/π and the Altarelli-Parisi splitting functions [36], which were briefly introduced in Sec. 1.1. The relevant functions are

$$P_{gq} = C_F \frac{1 + (1-z)^2}{z} \quad (4.4)$$

$$P_{gg} = 2C_A \left[\frac{z}{(1-z)_+} + \frac{1-z}{z} + z(1-z) \right] + \frac{1}{2} \beta_0 \delta(1-z), \quad (4.5)$$

where $\beta_0 = (11C_A - 4n_f T_F)/3$ is the leading order coefficient of the QCD β -function, $T_F = 1/2$, $C_A = 4/3$ and $C_F = 3$ are the quark and gluon color factors, and n_f is the number of active quark flavors, which is five. The function $g(x)_+$ is defined by:

$$\int_0^1 dx f(x) g(x)_+ = \int_0^1 dx (f(x) - f(1)) g(x), \quad (4.6)$$

for some function $f(x)$. These functions already include the singular behavior $\frac{1}{z}$ and for small z , these functions reduce to $P \approx 2C/z$. A useful space² for describing the probability distribution of the soft and collinear gluon emissions is shown in Fig. 4.4. Since $\log(z)$ is uniform on $-\infty$ to 0, $\log(1/z)$ is uniform on 0 (hard) to ∞ (soft). Likewise, $\log(R/\theta)$ is uniform from 0 (wide angle) to ∞ (collinear). The invariant mass m of two massless particles with energies $(1-z)E$ and zE is given by $m^2 = z(1-z)E^2(1 - \cos \theta) \approx zE^2\theta^2$. In Fig. 4.4, contours of constant invariant mass squared are thus straight lines given by

$$\log\left(\frac{1}{z}\right) = -2\log \rho - 2\log\left(\frac{R}{\theta}\right), \quad (4.7)$$

where $\rho = m/(ER)$. At leading order (one real emission), the probability for the jet to have mass squared less than some fixed value m_0^2 is $\Pr(m^2 \leq m_0^2) = 1 - \frac{2\alpha_s}{\pi} \mathbb{L}$, where $\mathbb{L} = \log^2 \rho$ is the area of the blue triangle in Fig. 4.4. Therefore, the leading order probability distribution for m^2 is given by

$$f_{\text{LO}}(m^2) = -\frac{2\alpha_s C}{\pi} \frac{\partial \mathbb{L}}{\partial m^2} = -\frac{\alpha_s C}{2\pi} \frac{\partial}{\partial m^2} \log^2 \rho^2 = -\frac{\alpha_s C}{\pi m^2} \log\left(\frac{m^2}{E^2 R^2}\right). \quad (4.8)$$

Changing the variables to m gives $f_{\text{LO}}(m) = f_{\text{LO}}(m^2) \frac{\partial m^2}{\partial m} = 2m f_{\text{LO}}(m^2)$, which is

²This calculation is based on similar discussions in Ref. [329, 340, 341].

$$f_{\text{LO}}(\mathbf{m}) = -\frac{4\alpha_s C}{\pi \mathbf{m}} \log\left(\frac{\mathbf{m}}{\mathbf{E}R}\right). \quad (4.9)$$

The leading order distribution of the mass is not useful because it diverges too quickly³ as $\mathbf{m} \rightarrow 0$ (so $\int_0^{\mathbf{E}R} f_{\text{LO}}(\mathbf{m}) d\mathbf{m} = \infty$). Therefore, a different approximation is needed in order to make a sensible prediction of the jet mass distribution. For the leading order calculation, the soft and collinear regions of phase space are unregulated for one emission. However, the probability of many significant emissions is non-negligible and therefore another possibility is to consider all possible single gluon emissions. The initial quark or gluon is treated as a final state object that can radiate an arbitrary number of gluons (the *eikonal* approximation) with $z \ll 1$ for each emission. The leading order calculation showed that each emission has the form $\alpha_s \log^2 \rho$ - this approximation is therefore a *leading logarithm* approximation in which all double-logarithms $(\alpha_s \log^2 \rho)^n$ are summed to all orders. The beginning of the calculation is the same as for the leading order one - the emission with the highest $z\theta^2$ in Fig. 4.4 will set the jet mass. Therefore for a fixed $\mathbf{m}_0$ and $\mathbf{n}$ emissions, one is interested in the probability that all emissions have $z\theta^2 < \mathbf{m}_0^2$. To compute this probability, divide the blue triangle in Fig. 4.4 into $\mathbf{N}$ little boxes of equal area $\mathbf{a} = \mathbf{\Delta} / \mathbf{N}$. The size of the boxes is chosen so that the probability of multiple emissions within the box is small. In this leading logarithm approximation, all emissions are assumed independent of each other. Therefore,

³The divergence at zero can be regulated by considering the virtual corrections, which contribute at exactly $\mathbf{m} = 0$ by construction (if there is no second particle, then the jet mass is zero). However, the leading logarithm approach is still more useful for understanding the full distribution of the jet mass, especially at low jet mass.

$$\Pr(\text{no emissions in } \mathbb{L}) = \prod_{N \text{ boxes}} \Pr(\text{no emission in the box}) \quad (4.10)$$

$$= \prod_{N \text{ boxes}} (1 - \Pr(\text{emission in box})) \quad (4.11)$$

$$= \prod_{N \text{ boxes}} \left(1 - \frac{2\alpha_s C_a}{\pi}\right) \quad (4.12)$$

$$= \left(1 - \frac{2\alpha_s C_{\mathbb{L}}}{\pi N}\right)^N \stackrel{N \rightarrow \infty}{\equiv} e^{-\frac{2\alpha_s C}{\pi} \mathbb{L}} \quad (4.13)$$

The derivative Eq. 4.10 gives the probability distribution of the jet mass $f_{LL}(m^2)$:

$$f_{LL}(m) = 2m f_{LL}(m^2) = 2m \frac{\partial}{\partial m^2} \Pr(\text{no emissions in } \mathbb{L}) \quad (4.14)$$

$$= -2m \frac{\alpha_s C}{\pi m^2} \ln \left(\frac{m^2}{E^2 R^2} \right) \exp \left(-\frac{\alpha_s C}{2\pi} \log^2 \left(\frac{m^2}{E^2 R^2} \right) \right) \quad (4.15)$$

$$= -\frac{4\alpha_s C}{\pi m} \ln \left(\frac{m}{ER} \right) \exp \left(-\frac{2\alpha_s C}{\pi} \log^2 \left(\frac{m}{ER} \right) \right), \quad (4.16)$$

which is finite (actually zero) as $m \rightarrow 0$. The exponential suppression factor in Eq. 4.14 is called a *Sudakov factor*. The left plot of Fig. 4.5 shows the distribution in Eq. 4.14 plotted for $(E, R) = (200 \text{ GeV}, 1)$ and $(E, R) = (400 \text{ GeV}, 0.4)$ separately for quark and gluon jets. In general, the quark jet mass distribution is shifted to lower values of the jet mass. The energies and radii are chosen to approximately correspond to $2m/E = R$ for a W boson in order to illustrate how the mass distribution compares to $m_W \approx 80 \text{ GeV}$. Since $f_{LL}(m)$ is bounded and has compact support, it has finite moments. The average jet mass is given by

$$\langle m \rangle = \alpha_s R p_T C \left(4 \int_0^1 d\rho \log(\rho) \exp \left(-\frac{2\alpha_s C}{\pi} \log^2 \rho \right) \right). \quad (4.17)$$

The expression in parenthesis in Eq. 4.17 is an $\mathcal{O}(1)$ number that is approximately

0.9 for quark jets and 0.7 for gluon jets. The right plot of Fig 4.5 shows the average jet mass as a function of jet p_T , compared with the electroweak boson and top quark masses. Especially for quark jets, which dominate at high p_T (see Sec. 3.2), the average mass is significantly less than the mass of the boosted ‘signal’ objects. Interestingly, at some high p_T the average QCD jet mass will be the same and even higher than the mass of electroweak bosons and top quarks. For this reason, analyses using ultra-boosted bosons and top quarks would use a ceiling requirement on the jet mass instead of a lower mass threshold.




Figure 4.4: A schematic diagram of the $\log(1/z)$ versus $\log(R/\theta)$ plane in which the probability for the emission of a gluon is approximately uniform. The dashed lines show contours of constant m^2 , which increase from upper right to lower left. The blue triangle corresponds to mass m_0^2 .

Figure 4.6 shows how the average jet mass depends on the jet p_T for QCD jets in the early Run 2 data compared with simulation. Jets are clustered with a radius $R = 1.0$ and trimmed (see Sec. 4.1.2 for details). As expected, the average jet mass increases monotonically with p_T . The exact shape in Fig. 4.6 deviates from linear because (a) the composition of quarks and gluons changes as a function of p_T and (b) the jets are trimmed and so the *effective area* of the jet depends on p_T . The jet mass in the simulation is generated in the parton shower implemented in PYTHIA 8 which is based on the leading logarithm approximation, but includes additional effects such as a running α_s and the full LO quark and gluon splitting functions. The next sections describe how the jet mass is reconstructed in practice.




Figure 4.5: Left: The leading logarithm distribution of the jet mass for $(E, R) = (200 \text{ GeV}, 1)$ and $(E, R) = (400 \text{ GeV}, 0.4)$ separately for quark and gluon jets. The quark distribution is suppressed at zero, but it increases for finite mass much faster than the gluon distribution. Right: The average jet mass as a function of jet p_T for the same four settings as the left plot. Horizontal lines indicate the W boson, Higgs boson and top quark masses.




Figure 4.6: The p_T dependence of the average leading anti- k_t $R = 1.0$ trimmed with $f_{\text{cut}} = 0.05$ and $R_{\text{sub}} = 0.2$ jet mass. See Sec. 4.1.2 for details.

4.1.2 Reconstructing the Calorimeter Jet Mass

Reconstructing the jet mass is an experimental challenge because it requires a precise measurement of both the energy and location of particles inside a jet. This property of the jet mass is illustrated in Fig. 4.7 for a hadronically decaying boosted W boson in a MC model. Particles carrying a small fraction of the jet's $\mathbf{p}_T$ can contribute just as much to the mass as particles carrying a large fraction of the total momentum. Furthermore, at a hadron collider there is no conservation law that can be used for an in-situ study of the jet mass response. For the jet momentum, conservation in the plane transverse the beam is a powerful constraint that has no analogue for mass as $\sqrt{\hat{s}}$ is unknown. Even at a high energy electron-positron collider, for cases of interest for tagging, the jet mass is typically much smaller than the jet energy and therefore a constraint on the total energy is not useful.




Figure 4.7: The impact of a jet constituent on that jet's mass as a function of ΔR between the constituent and the jet axis and the $\mathbf{p}_T$ fraction carried by the constituent. W bosons are generated with $\mathbf{p}_T = 200$ GeV in a random direction and decay via a scalar two-body phase space into quarks which subsequently decay into 10 massless particles each also with the scalar phase space. The 'jet' is the four-vector sum of all these particles, which has mass 80 GeV and $\mathbf{p}_T = 200$ GeV.

In Run 1 of the LHC, the most used definition of the jet mass takes as input calibrated calorimeter-cell clusters. Most of this section will be focused on the mass of *large-radius* jets clustered with the anti- k_t algorithm using $R = 1.0$ and groomed with the trimming procedure [227] to reduce the sensitivity of the jet mass to contamination from pileup and the underlying event. These sources of diffuse energy are detrimental to the jet mass resolution because low-energy wide-angle radiation can have a big impact on the jet mass as illustrated by Fig. 4.7. For trimming, the jet constituents are re-clustered with the k_t algorithm⁴ using $R = R_{\text{sub}}$ and then the constituents of the resulting *subjects* with $p_T^{\text{subject}} < f_{\text{cut}} \times p_T^{\text{jet}}$ are removed. Note that this requirement is applied before any pileup mitigation and therefore the trimming becomes harsher for higher levels pileup. This is solved naturally by the re-clustering algorithm, described in Sec. 4.1.3. As a result of an extensive campaign [342, 343] to optimize R_{sub} and f_{cut} , the values $f_{\text{cut}} = 0.05$ and $R_{\text{sub}} = 0.3$ (0.2) are used in Run 1 (Run 2). The smaller R_{sub} value improves the jet mass resolution at high p_T where the jet constituents of a resonance with fixed mass are closer together. Figure 4.8 shows an event display in data illustrating the impact of trimming. Two high p_T well-isolated jets are nearly back-to-back in the transverse plane ($\Delta\phi \bmod \pi \approx 0$). The isolated anti- k_t jets have a circular catchment area whereas the k_t subjects have irregular areas whose sum is much less than the ungroomed jet area. Trimming has a small effect on the jet p_T , but a non-trivial impact on the jet masses. For example, the lower left jet loses less than 2% of its p_T after trimming while the jet mass is reduced by over 10%.

A jet-level calibration is applied to account for the residual detector response. This correction is first applied to the jet energy and then to the jet mass. In particular, the calibrated jet mass m of a jet J reconstructed with η_J is given by

$$m = c_{\text{JMS}} \left(c_{\text{JES}} \left(\sum_{i \in J} E_i, \eta_J \right), \eta_J \right) \times \sqrt{\left(\sum_{i \in J} E_i \right)^2 - \left(\sum_{i \in J} \vec{p}_i \right)^2}, \quad (4.18)$$

⁴The k_t instead of anti- k_t algorithm is used for subjects because it results in a more balanced distribution of energy - see Ref. [227] for more details.




Figure 4.8: An event display of a dijet event in the early Run 2 data. The gray and colored filled circles correspond to calorimeter cell clusters where the circle radius is proportional to the log of the cluster energy. Colored circles represent the constituents of the trimmed jets. The blue-gray circles are the ungroomed anti- k_t $R=1.0$ jets. The dark gray circles are the constituent clusters removed by trimming. The remaining gray area underneath the colored circles is the ghost area of the trimmed jet.

where E_i is the LCW calibrated energy of cluster i . Each cluster is treated as massless with three-momentum $\vec{p}_i = (E_i/\cosh\eta)(\cos\phi_i, \sin\phi_i, \sinh\eta)$. The calibration functions c_{JMS} and c_{JES} are for the jet mass and jet energy scales determined using numerical inversion. When generic QCD jets are used to derive the calibration, one needs to also control for the jet size. Early Run 2 calibrations therefore use m/p_T as one of the inputs to c_{JMS} . After this jet p_T - and jet mass-dependent calibration, the average reconstructed jet mass is the same as the particle-level jet mass in simulation for quark and gluon jets: the calibration *closes*. The response depends on the quark/gluon nature of the jets, so the calibration is only guaranteed to close in

a sample of events with the same composition as the one used to derive the calibration. Also, as a result of the dependence of the response on jet substructure, the calibration may not exactly close for boosted $W/Z/H$ boson or top quark jets. This is not necessarily a problem for jet tagging, but it can be mitigated by controlling for jet substructure in the calibration or performing the calibration on signal jets. Alternative jet mass definitions are described in Sec. 4.1.3 and 4.1.4.

Large radius jet 4-vector reconstruction performance is quantified by properties of the response (R)⁵: the ratio of the reconstructed jet mass to the jet mass of the corresponding particle-level jet. When distinguishing boosted hadronic resonance jets from generic quark and gluon jets, the most important property of R is its width. Since the distribution of R is not Gaussian, there is no universally accepted definition of the width. Figure 4.9 shows the distribution of R for boosted hadronically decaying W and Z bosons in four boson p_T ranges from 200 GeV up to 2 TeV. For illustration, two different fits are performed and overlaid on the input distributions. The first fit is an iterative χ^2 fit to a Gaussian that uses the histogram mean and standard deviation as seeds and then subsequently uses the fitted mean and standard deviation to set the fit range. By focusing on a $\pm 1\sigma$ interval about the mean, the fit captures the *core* of the distribution of R . Another way to isolate the core and down-weight the heavy tails is to fit a double-Gaussian. In Fig. 4.9 the core Gaussian is shown in red while the tail Gaussian is shown in blue. Both normal distributions are constrained to have the same mean and the core Gaussian is seeded with half the histogram standard deviation while the tail Gaussian is seeded with twice the histogram standard deviation. Similar fits are performed for the jet p_T response in Fig. 4.10. The p_T dependence of the core Gaussian resolution, along with the histogram standard deviation and 68% median-centered quantile are shown in Fig. 4.11 for both the jet mass and jet p_T response. For all definitions the measure of spread is divided by a measure of the distribution center, which is either the (fitted) mean or the median. The two fitting methods designed to isolate the core of the response distribution give similar results for both quantities (and for the inter-quantile range). However, the standard deviation is significantly

⁵Not to be confused with the jet radius, which is a constant. Unfortunately, the use of the symbol R is standard for both quantities.

larger, indicating the presence of non-negligible heavy tails.




Figure 4.9: The distribution of the jet mass response (R) for boosted hadronically decaying W and Z bosons in four boson p_T ranges from 200 GeV up to 2 TeV. See the text for a description of the various fitting methods.




Figure 4.10: The distribution of the jet p_T response (R) for boosted hadronically decaying W and Z bosons in four boson p_T ranges from 200 GeV up to 2 TeV. See the text for a description of the various fitting methods.




Figure 4.11: A summary of the p_T -dependence of the jet mass (left) and jet p_T (right) resolution. See the text for a description of the various fitting methods.

As expected from a calorimeter measurement, the jet p_T resolution monotonically decreases with p_T . At low p_T , this is also true for the jet mass resolution, but there is a turning point at about 500 GeV where the jet mass resolution degrades with p_T . This is because for a fixed mass, the particles inside a jet get closer together with increasing p_T . Due to the finite granularity of the detector, small angular separations cannot be resolved which reduces the jet mass scale and degrades the jet mass resolution. When the subjets inside the large-radius jet are well-separated (at low p_T), the mass and p_T resolutions are similar in magnitude because the mass resolution is mostly due to the energy resolution of the isolated subjets.

With an optimized definition of the jet mass and a calibration to remove most of the detector response, the key challenge is to determine the closure of the four-vector calibration and the jet mass and p_T resolutions in data. Section 4.1.2.1 introduces the most widely used technique for determining the closure of the large-radius jet p_T and jet mass calibration in data - the track-jet method. After a brief introduction to a bottom-up method in Sec. 4.1.2.2, the remainder of this section (Sec. 4.1.2.3) focuses on a new technique based on fitting resonance peaks.

4.1.2.1 Track-jet Method

The baseline method for measuring the closure of the calibration in data uses track jets. Tracks are clustered into jets using the same algorithm as for the calorimeter jets. These large-radius track jets are geometrically matched to calorimeter jets and their jet mass provides an independent measurement of the particle-level jet mass. Track-jets are particularly useful because the typical difference between the reconstructed track jet mass and the jet mass from the particle-level jet using only charged particles is small compared to the calorimeter jet mass resolution. However, the resolution of the track-jet mass with respect to the full particle-level jet mass is not small compared to the calorimeter jet mass resolution due to the large fluctuations in the charge-to-neutral ratio of particles inside the jet. Therefore, it is not possible to perform a measurement of the absolute closure of the jet mass calibration using track jets. Instead, the closure in data is studied relative to the closure in simulation. Define r_{track} as the ratio of the calorimeter jet mass to the matched track jet mass. Then,

$$r_{\text{track}} = \frac{m^{\text{calorimeter}}}{m^{\text{particle}}} \times \frac{m^{\text{particle}}}{m^{\text{charged-particle}}} \times \frac{m^{\text{charged-particle}}}{m^{\text{track}}}, \quad (4.19)$$

where the first term is the jet mass response (R), the second term is the inverse of the charged ratio of the jet (f_Q^{-1}) and the last term is the inverse of the track-jet resolution with respect to the particle-level jet mass using only charged particles (R_T^{-1}). The statement that track jets are well measured means that $\sigma(R_T) \ll \sigma(R)$. Due to isospin, $\langle f_Q \rangle \sim 2/3$. However, $\langle f_Q^{-1} \rangle > 3/2$, as shown by the left plot of Fig. 4.12. The right plot of Fig. 4.12 shows that $\langle f_Q \rangle$ is nearly independent of p_T , a fact that was used in Sec. 1.5.3.3 to determine tracking uncertainties for the jet charge. Interestingly, there is a slight difference between the the ratio based on the mass and the one based on p_T due to subtle differences in jet fragmentation to charged and neutral particles. The standard deviation of the f_Q distribution is also nearly independent of p_T and is approximately 0.2 which corresponds to $\langle f_Q^{-1} \rangle \sim 2$ based on the left plot of Fig. 4.12. The value $\sigma(f_Q)$ is smaller than one would expect if all particles carry an equal fraction of the jet's energy (see Fig. 4.12), but is not negligible

compared to $\langle f_Q \rangle$.




Figure 4.12: Left: The dependence of $\langle f_Q^{-1} \rangle$ on the standard deviation of the charged-to-neutral fluctuations. The charged fraction f_Q is modeled as a log normal distribution with mean $2/3$ and variable standard deviation σ . The horizontal dashed line is at $2/3$. The vertical dashed line is at $\sigma = 0.47 \approx \sqrt{\frac{2}{3}(1 - \frac{2}{3})}$, which is what one expects if each particle inside a jet carries an equal fraction of the jet energy and has a probability $2/3$ of being charged. Right: the p_T dependence of $\langle f_Q \rangle$ for both jet mass and jet p_T . The band is the standard deviation of the f_Q distribution using p_T .

When the terms on the righthand side of Eq. 4.19 are independent, $\langle r_{\text{track}} \rangle \propto \langle R \rangle$. In this case, an estimate for the relative uncertainty in the non-closure is given by

$$1 - \frac{\langle R^{\text{data}} \rangle}{\langle R^{\text{MC}} \rangle} = 1 - \frac{k^{\text{MC}}}{k^{\text{data}}} \times \frac{\langle r_{\text{track}}^{\text{data}} \rangle}{\langle r_{\text{track}}^{\text{MC}} \rangle}, \quad (4.20)$$

where k is the constant of proportionality between $\langle r_{\text{track}} \rangle$ and $\langle R \rangle$. This measurement is limited by data statistics at high jet p_T and elsewhere by uncertainties in $k^{\text{MC}}/k^{\text{data}}$ due to systematic uncertainties on the reconstruction of charged particle tracks and the modeling of jet fragmentation.

In principle, the track-jet method can also be used to determine the relative jet mass resolution, $\sigma(R)^{\text{data}}/\sigma(R)^{\text{MC}}$. Further assuming that R^2 and $f_Q^{-2} \times R_T^{-2}$ are independent, one can write $\langle r_{\text{track}} \rangle = k_1 \langle R \rangle$ and $\langle r_{\text{track}}^2 \rangle = k_2 \langle R^2 \rangle$ for some constants k_1, k_2

that are in general different between data and simulation. Then,

$$\sigma^2(\mathbf{R}) = \frac{\langle \mathbf{r}_{\text{track}}^2 \rangle}{k_2} - \frac{\langle \mathbf{r}_{\text{track}} \rangle^2}{k_1^2}. \quad (4.21)$$

This procedure is not applied in practice because it depends explicitly on the value of k_i , whereas for $\langle \mathbf{R} \rangle$, only the relative $k^{\text{MC}}/k^{\text{data}}$ are required.

Figures 4.13 and 4.14 show the distribution of $\mathbf{r}_{\text{track}}$ and $\langle \mathbf{r}_{\text{track}} \rangle(p_T)$ for generic quark and gluon jets and W boson-like jets in both data and simulation. For both sets of jets, the $\mathbf{r}_{\text{track}}$ distribution peaks just below 2 and the predicted $\langle \mathbf{r}_{\text{track}} \rangle$ is within $\lesssim 5\%$ of the data. The populations of jets in Fig. 4.13 and 4.14 are a representative set for applications of jet mass. Generic QCD jets are used to calibrate jets and set the most precise uncertainties using $\mathbf{r}_{\text{track}}$ because of their abundance. The jet mass is mostly used for tagging boosted hadronically decaying bosons and top quarks. There are no indications for significant biases in this procedure based on the $\mathbf{r}_{\text{track}}$ method. The total uncertainty from the track-jet method is about 5%, independent of jet p_T [122].




Figure 4.13: The (a) ratio of the calorimeter jet mass to track jet mass and (b) average value of this ratio as a function of the calorimeter jet mass for leading anti- k_t $R = 1.0$ trimmed with $f_{cut} = 0.05$ and $R_{sub} = 0.2$ jets. Calorimeter jets which contain only a single cluster have a mass of zero. MC is normalized to the number of events observed in data. The last bin includes overflow events.




Figure 4.14: (a) The distribution of r_{track} in the data for semi-leptonic $t\bar{t}$ events with the selected jet in the range $50 \text{ GeV} < m^{\text{jet}} < 120 \text{ GeV}$. (b) The median of the r_{track} distribution as a function of the jet p_T . (c) The inter-quantile range as a measure of the width. See Sec. 4.2 for a description of the event selection.

4.1.2.1.1 Interpreting r_{track} Uncertainties

While the track-jet method is simple to apply, there is an important caveat when interpreting the results. In general, f_Q and R are **not** independent. The calorimeter response is different between charged and neutral particles. The LCW corrects the difference on average, but the classification of individual clusters as EM or hadronic has a non-zero error and the finite calorimeter energy resolution can be non-negligible. Figure 4.15 shows the joint distribution of R and $f_Q^{-1} \times R_T$. The linear correlation between these two variables is small, but a small correlation is not sufficient for the average of two random variables X and Y to factorize: $\langle XY \rangle = \langle X \rangle \langle Y \rangle$.




Figure 4.15: The joint distribution of the response (R) and the product of the inverse charged ratio f_Q^{-1} and the inverse track mass response R_T^{-1} using simulated W boson jets with $p_T^{\text{truth}} > 200 \text{ GeV}$ from POWHEG-BOX+PYTHIA 6 $t\bar{t}$ production.

Figure 4.16 shows the gap between $\langle R \rangle \langle f_Q^{-1} R_T^{-1} \rangle$ and $\langle r_{\text{track}} \rangle = \langle R f_Q^{-1} R_T^{-1} \rangle$. Defining $\mathbf{c}$ to be the size of the gap, one can try to estimate if the r_{track} -based uncertainties are at least conservative, i.e. is $|\sigma| < |\Delta|$ for

$$\sigma = 1 - \frac{\langle R^{\text{data}} \rangle}{\langle R^{\text{MC}} \rangle} \quad \text{and} \quad \Delta = 1 - \frac{\langle r_{\text{track}}^{\text{data}} \rangle}{\langle r_{\text{track}}^{\text{MC}} \rangle}. \quad (4.22)$$

In other words, σ is the ‘true’ uncertainty and Δ is the uncertainty one estimates using the track-jet method. Assuming that the difference in $\mathbf{c}$ and $\langle f_Q^{-1} R_T^{-1} \rangle$ between data and simulation is small compared to $\epsilon = \mathbf{c} / \langle r_{\text{track}}^{\text{MC}} \rangle$, one can compute $\sigma = \Delta(1 - \epsilon) + \mathcal{O}(\epsilon^2)$. Since $\epsilon > 0$, $|\sigma| < |\Delta|$, as desired. If $\Delta \sim 5\%$ and $\epsilon \sim 30\%$, then the size of the bias could be 1-2%. This bias is currently not accounted for when applying the track-jet method, which can be justified if a precision uncertainty is not the goal.




Figure 4.16: Left: The gap between $\langle R \rangle \langle f_Q^{-1} R_T^{-1} \rangle$ and $\langle r_{\text{track}} \rangle = \langle R f_Q^{-1} R_T^{-1} \rangle$.

4.1.2.2 Bottom-up Method

Up to⁶ $p = 350$ GeV, the response of individual particles is well-measured using isolated tracks and test-beam experiments. The idea of the bottom-up method is to model the (average) jet response as the sum of the average jet response for all the constituent particles inside a jet. Symbolically,

$$p_{\text{jet, reco}}^\mu = \sum_{i \in \text{particle-level jet}} \langle R_i \rangle p_i^\mu, \quad (4.23)$$

where p_i^μ is the four-vector and $\langle R_i \rangle$ is the average calorimeter energy response of particle i . From $p_{\text{jet, reco}}^\mu$, one can compute the jet p_T or jet mass response for a given jet by dividing by the appropriate property of the particle-level jet. Figure 4.17 shows the average p_T and mass response using this bottom-up approach for boosted hadronically decaying W bosons. Only particles with energy above 500 MeV are included. For electrons and photons, $\langle R_i \rangle = 1$ and since muons do not usually deposit significant energy in the calorimeter, $\langle R_i \rangle = 0$. The ratio of charged-particle calorimeter energy to track p_T (E/p) is used for charged hadrons up to $p_T = 20$ GeV, after which test-beam data [344] is used up until $p_T = 350$ GeV. The test-beam response ranges from 0.65 at $E < 35$ to 0.78 for $p_T > 125$ GeV (for central η). For $p_T > 350$ GeV, $\langle R_i \rangle$ is not constrained by data; in Fig. 4.17 it is set to the highest value from the test-beam: 0.78. Protons and pions with $E < 10$ GeV can be identified using the amount of energy deposited as a function of distance traversed in the detector (dE/dx) and so their response values are individually computed while all other charged hadrons use generic values. For nearly collinear constituents (and mostly uncorrelated energy fluctuations), Eq. 4.23 is a good approximation for the jet p_T and therefore the bottom-up p_T response in Fig. 4.17 well-models the full response. However, jet mass is the result of significant angular splittings and so Eq. 4.23 is not a good approximation. It is therefore not surprising that the two models diverge at high p_T in the right plot of Fig. 4.17.

⁶This section includes technical input from Z. Marshall.




Figure 4.17: The average jet p_T (left) and mass (right) response as a function of p_T using simulated detector-level jets and propagating single particle responses via Eq. 4.23. The points marked data use E/p as measured with the ATLAS detector. Error bars on the points indicate the MC statistical uncertainty on the mean.

The jet p_T and mass scale uncertainty in the bottom-up approach arises due to uncertainty in the values of $\langle R_i \rangle$. This approach may be a useful technique for extending the jet mass and p_T scale (uncertainty) to high values beyond the capability of in-situ studies, though there seems to be some challenges for $p_T \gtrsim 1$ TeV. Furthermore, it can (approximately) describe some aspects of correlations between variables. However, it is not able to fully describe correlations and does not model fluctuations about the mean.

4.1.2.3 Resonance Method

Known resonance decays provide a standard reference for in-situ calibration and uncertainty studies. For example, Z boson, J/ψ , and Υ decays are used to measure the scale and resolution of the response function for muons [131], electrons and photons [128], and tau leptons [134]. However, these techniques are not directly applicable to hadronic resonance decays because the parton shower and jet clustering introduce a non-trivial distortion of the resonance's Breit-Wigner mass line-shape. This is illustrated by the difference between the black and red curves in the left plot of Fig. 4.18. The particle-level distribution depends on perturbative properties of the parton shower as well as non-perturbative effects such as hadronization and the underlying event. A measurement of the resonance peak will probe the convolution of these particle-level effects and the detector response. An extraction of the jet mass scale and resolution from a hadronic resonance requires the particle-level spectrum as input and therefore the precision can be limited by the corresponding modeling uncertainties. The right plot of Fig. 4.18 shows the impact of varying α_s in the parton shower on the particle-level jet mass spectrum - this uncertainty directly limits the precision of the resonance method. It is not possible to obtain a pure sample of hadronically decaying Z bosons at a hadron collider. However, it is possible to select events enriched in hadronic W boson decays from $t\bar{t}$ events where the second W boson is used to tag the event through its leptonic decay. Low p_T hadronic W boson decays have been used as a validation of the light quark jet energy scale in early Run 1 [121]. The precision of this measurement was limited by the modeling of the parton shower.

The measurement presented here⁷ is the first full estimate of the jet mass scale and resolution on the jet mass directly from boosted W boson jets. After a brief description of the event selection and simulation in Sec. 4.1.2.3.1, Sec. 4.1.2.3.2 describes a new technique for extracting the jet mass scale and resolution from resonance decays called the *forward-folding* method. The impact of systematic uncertainties in the measurement are given in Sec. 4.1.2.3.3 and the Run 1 result is summarized in Sec. 4.1.2.3.4. Finally, Sec. 4.1.2.3.6 contains improvements and extensions of the

⁷The $\sqrt{s} = 8$ TeV analysis presented in this section has been published in Ref. [345] and includes technical inputs from J. Veatch.

methods as well as results with the early Run 2 data.




Figure 4.18: Left: the mass distribution of the W boson (parton-level) and boosted and isolated W boson jets at particle-level and the same jet at detector-level. The jets are required to have $p_T > 200$ GeV. See Sec. 2.2.1 for details of the simulation. Right: the particle-level jet mass spectrum for candidate boosted W boson jets from $t\bar{t}$ events using the default Perugia 2012 PYTHIA 6 tune or the radHi/radLo variations. The one-loop running of the strong coupling-constant is $\alpha_s \propto 1/\ln(Q^2/\Lambda^2)$ with $\Lambda = 0.26$ in the default Perugia 2012 tune. This value is double or halved in the radHi/radLo variations.

4.1.2.3.1 Dataset and Event Reconstruction

The dataset and MC simulations are the same as those used for the color flow measurement, described in Sec. 2.2.1. In particular, the data are collected with single electron and muon triggers from the entire 2012 dataset corresponding to an integrated luminosity of 20.3 fb^{-1} . POWHEG-BOX+PYTHIA 6 is used for modeling the nominal $t\bar{t}$ sample. The definitions of reconstructed objects, aside from the addition of large-radius jets, are also the same as for the color flow measurement - see Sec. 2.2.2. The only exception is the isolation of electrons. Just as the size of W boson and top jets decreases with p_T , the leptons from W decays are closer, on average, to the b -jets originating from the same parent top quark with increasing p_T . A relative isolation

based on a shrinking cone is straightforward to apply for muons and recovers the efficiency at high p_T . A non-trivial complication for electrons is that they deposit most of their energy in the calorimeter which can be clustered with the radiation from the b -quark to form a single jet. Therefore, to recover efficiency at high top quark p_T , the electron energy deposit in the calorimeter is removed from the closest jet with $\Delta R < 0.4$ before applying a similar relative isolation procedure.

The event and object selections are based on the ATLAS search for $t\bar{t}$ resonances [346] and are summarized here for completeness. Candidate reconstructed $t\bar{t}$ events are chosen by requiring an electron or a muon with $p_T > 25$ GeV and $|\eta| < 2.5$, as well as a missing transverse momentum $E_T^{\text{miss}} > 20$ GeV. Events are rejected if there is not exactly one electron or muon. In addition, the sum of the E_T^{miss} and the transverse mass of the W boson, reconstructed from the lepton momentum and $\vec{p}_T^{\text{miss}}$, is required to be greater than 60 GeV. Events must have at least one b -tagged jet (at the 70% efficiency working point) and have at least one large-radius trimmed jet with $p_T > 200$ GeV and $|\eta| < 2$. Furthermore, there must be a small-radius jet with $p_T > 25$ GeV, and $\Delta R < 1.5$ to the selected lepton (targeting the decay chain $t \rightarrow bW(\rightarrow \ell\nu)$). The candidate W jet used for the measurement is selected as the leading large-radius trimmed jet with $\Delta\phi > 1.5$ from the lepton $\Delta R > 1.2$ from the small-radius jet that is matched to the lepton. To further ensure that the selected jet contains only the decay products of a W boson, it is required to have $\Delta R > 1.0$ to the nearest b -tagged small-radius jet. The jet mass and jet p_T distributions after the above event selections are shown in Fig. 4.19 and 4.20. There are about 35,000 events in data that pass the full selection; about 10,000 events that have $p_T > 300$ GeV; about 3,000 that have $p_T > 400$ GeV, and just over 1000 events with $p_T > 500$ GeV. The purity of events is about 65% over the entire mass range and about 80% for jet masses above 65 GeV.

4.1.2.3.2 Extracting the Jet Mass Scale and Resolution

There are two sets of considerations when constructing a procedure for extracting the jet mass response from the measured spectrum. First, the particle-level spectrum




Figure 4.19: The jet mass distribution for events passing the selection described in Sec. 4.1.2.3.1 for electron events (left) and muon events (right).




Figure 4.20: The jet p_T distribution for events passing the selection described in Sec. 4.1.2.3.1 for electron events (left) and muon events (right).

and/or response function can be parameterized or a non-parametric procedure can be used. In general, parametric forms are useful if the underlying function is known and also to make interpretation easier. In this case, approximations to the particle-level mass distribution are known (See Sec. 4.1.1), but these analytic forms do not capture non-perturbative effects such as hadronization and the underlying event that

are accounted for in the state-of-the-art MC generators. Therefore, the analysis presented in this section uses the non-parametric particle-level mass distribution from the simulation. The resolution plots in Sec. 4.1.2 showed that the jet mass response function is non-Gaussian and depends on the jet p_T . One possibility is to find a functional form (such as the double-Gaussian) for the non-Gaussian function and then to let the parameters depend on p_T . The benefit of using a parametric form is that the fitted parameters can be directly interpreted as the scale and resolution of the response function. However, even though the functions in Sec. 4.1.2 worked well to describe the jet mass response, there were still deviations from the empirical distributions from simulation and in general there is a dependence on the jet mass itself in addition to the jet p_T . Therefore, the method presented in this section is fully non-parametric, using the response function from the simulation. This response function is stretched and shifted so that when the particle-level mass distribution is *forward-folded*, it best matches with the data. In order to take into account the jet mass and jet p_T dependence of the response function, the response distribution from the simulation is binned in the jet p_T and the jet mass. More details about the forward-folding technique are described below.

The second consideration for constructing the extraction procedure is the particle-level event selection (fiducial volume). Ideally, the detector-level event selection would result in a sufficient pure selection. For generic jets from $t\bar{t}$ events, this is mostly true for jet masses near m_W . However, there are a large fraction of events that originate from top quark pair production, but the selected jet is not a fully contained W boson jet. This is evident from the large contribution to the particle-level jet mass spectrum in Fig. 4.19 for jet masses far from m_W . One approach is to subtract the non-resonant $t\bar{t}$ and non- $t\bar{t}$ backgrounds from the detector-level distribution prior to fitting the mass response (See the *subtraction method* in Ref. [345]). This makes the definition of the measurement conceptually cleaner, but introduces significant sources of model dependence. For example, the background jet mass distribution is not uniform under the W boson peak (it is falling) and therefore an uncertainty in the normalization of the backgrounds results in an uncertainty on the measured jet mass distribution shape. Additionally, the jet mass response of the background must

be taken as an input which can bias the measurement of the response in the signal. For these reasons, the measurement presented in this section is defined only by its detector-level selection - no components are subtracted prior to the measurement.

From the above considerations, a forward-folding method is used to extract the relative differences in the jet mass response between data and simulation. Let $\mathbf{R}(\mathbf{m}_{\text{true}}, \mathbf{p}_{\text{T}}^{\text{reco}})$ be the distribution of the jet mass response for given values of the particle-level jet mass $\mathbf{m}_{\text{true}}$ and the reconstructed (fully calibrated) jet transverse momentum $\mathbf{p}_{\text{T}}^{\text{reco}}$. In general, $\mathbf{R}$ is non-Gaussian and the full non-parametric form is taken from simulation, as well as the distribution of $\mathbf{m}_{\text{true}}$. For random variable ρ with $\rho \sim \mathbf{R}(\mathbf{m}_{\text{true}}, \mathbf{p}_{\text{T}}^{\text{reco}})$ and fixed $0 < r, s < \infty$ (to be determined below), define the new random variable $\mathbf{m}_{\text{folded}}$ as follows:

$$\mathbf{m}_{\text{folded}} | \mathbf{m}_{\text{true}}, \mathbf{p}_{\text{T}}^{\text{reco}}, r, s = [s\rho + (\rho - \langle \rho \rangle)(r - s)] \mathbf{m}_{\text{true}} \quad (4.24)$$

$$= (r\mathbf{m}_{\text{true}})\rho + (s - r)\langle \rho \rangle \mathbf{m}_{\text{true}} \quad (4.25)$$

The functional form of the transformation in Eq. 4.24 is chosen so that the distribution of $\mathbf{m}_{\text{folded}} | \mathbf{m}_{\text{true}}, \mathbf{p}_{\text{T}}^{\text{reco}}$ is the same as the distribution of $\mathbf{m}_{\text{reco}} | \mathbf{m}_{\text{true}}, \mathbf{p}_{\text{T}}^{\text{reco}}$ but with the average response scaled by s and the standard deviation of the response scaled by r . Symbolically:

$$\left\langle \frac{\mathbf{m}_{\text{folded}}}{\mathbf{m}_{\text{true}}} \middle| \mathbf{m}_{\text{true}}, \mathbf{p}_{\text{T}}^{\text{reco}} \right\rangle = s \langle \rho \rangle \quad (4.26)$$

$$\sigma \left(\frac{\mathbf{m}_{\text{folded}}}{\mathbf{m}_{\text{true}}} \middle| \mathbf{m}_{\text{true}}, \mathbf{p}_{\text{T}}^{\text{reco}} \right) = r \sigma(\rho), \quad (4.27)$$

where $\sigma(X)$ is the standard deviation of the random variable X . Eq. 4.26 follows by inspection of Eq. 4.24 because the second term has mean zero and Eq. 4.27 is evident from Eq. 4.25 because the second term is not random. For $r = s = 1$, $\mathbf{m}_{\text{folded}} | \mathbf{m}_{\text{true}}, \mathbf{p}_{\text{T}}^{\text{reco}}$ and $\mathbf{m}_{\text{reco}} | \mathbf{m}_{\text{true}}, \mathbf{p}_{\text{T}}^{\text{reco}}$ have the same distribution. The values of s and r are chosen such that the distribution of $\mathbf{m}_{\text{folded}}$ best matches the data. (Un)folding methods usually need to correct for migrations between the particle-level

and detector-level selections, but this is not necessary because the event selection is on the *reconstructed* jet $\mathbf{p}_T$ while the fitting is performed on the jet mass⁸ The fit to the detector-level jet mass distribution is performed by minimizing a χ^2 per degrees of freedom:

$$\mathbf{r}, \mathbf{s} = \operatorname{argmin}_{\mathbf{r}', \mathbf{s}'} \frac{1}{n-1} \left(\sum_{i=1}^n \frac{h_i(\mathbf{m}_{\text{folded}}|\mathbf{r}', \mathbf{s}') - h_i(\mathbf{m}_{\text{data}})}{\sigma_{i, \mathbf{m}_{\text{reco}}}^2 + \sigma_{i, \mathbf{m}_{\text{data}}}^2} \right)^2, \quad (4.28)$$

where $h_i(\cdot)$ is the content of a histogram of the variable $\cdot$ with n bins, $\sigma_{i, \mathbf{m}_{\text{data}}} = \sqrt{h_i(\mathbf{m}_{\text{data}})}$, and $\sigma_{i, \mathbf{m}_{\text{reco}}}$ is the MC statistical uncertainty in bin i . Since n is fixed, the normalization factor is only needed for the visualizations that appear later and do not impact the fitted values of $\mathbf{r}$ and $\mathbf{s}$. In order to render the fit insensitive to overall changes in the normalization, each template is normalized to have the same integral as the data. Since the multijet background is derived directly from the data, it is added unchanged to the MC-derived templates for each value of $\mathbf{r}$ and $\mathbf{s}$. In order to maximize the sensitivity to W boson-like jets, the fit is only performed in the mass range $50 \text{ GeV} < \mathbf{m}_{\text{reco}} < 120 \text{ GeV}$.

There are a variety of methods that could be used to generate $h_i(\mathbf{m}_{\text{folded}}|\mathbf{r}', \mathbf{s}')$. The most straight-forward but computationally intensive method would be to generate enough MC events to numerically estimate the full $\mathbf{p}_T^{\text{reco}}$ and $\mathbf{m}_{\text{true}}$ dependence of $\mathbf{R}$, sample events from $f(\mathbf{m}_{\text{true}}, \mathbf{p}_T^{\text{reco}})$, and then subsequently sample events from $f(\mathbf{m}_{\text{folded}}|\mathbf{r}', \mathbf{s}', \mathbf{m}_{\text{true}}, \mathbf{p}_T^{\text{reco}})$. This method requires a large number of sampled events per $\mathbf{r}'$ and $\mathbf{s}'$ in addition to a one-time cost of generating enough simulated events to populate the bins of the three-dimensional template for $\mathbf{R}$. The rest of the section uses a much simpler method that takes advantage of the fact that every event that is used for the estimation of the particle-level jet mass distribution also has a detector-level jet mass value. Each particle level jet mass $\mathbf{m}_{\text{true}}$ is transformed in the following way:

⁸The jet $\mathbf{p}_T$ spectrum in simulation well reproduces the data (Fig. 4.20) and any residual differences could be removed by re-weighting.

$$\mathbf{m}_{\text{true}} \mapsto \mathbf{m}'_{\text{true}}|\mathbf{r}, \mathbf{s} = \mathbf{s}\mathbf{m}_{\text{reco}} + (\mathbf{m}_{\text{reco}} - \langle \mathbf{m}_{\text{reco}}|\mathbf{m}_{\text{true}}, \mathbf{p}_T^{\text{reco}} \rangle)(\mathbf{r} - \mathbf{s}), \quad (4.29)$$

where $\mathbf{m}_{\text{reco}}$ is the detector-level jet mass from *the same simulated event* as $\mathbf{m}_{\text{true}}$. By construction, the distribution of $\mathbf{m}'_{\text{true}}|\mathbf{r}, \mathbf{s}$ is the same as $\mathbf{m}_{\text{reco}}|\mathbf{r}$, but does not require the distribution of $\mathbf{R}$ to be estimated directly, since it is built in automatically to the relationship between $\mathbf{m}_{\text{true}}$ and $\mathbf{m}_{\text{reco}}$. The only additional input that is required is a two-dimensional template for $\langle \mathbf{m}_{\text{reco}}|\mathbf{m}_{\text{true}}, \mathbf{p}_T^{\text{reco}} \rangle$, an example of which is shown in the left plot of Fig. 4.21. With the transformation in Eq. 4.29, the closure of the method is trivial: the χ^2 in Eq. 4.28 is exactly zero when $\mathbf{r} = \mathbf{s} = 1$ no matter how many MC events are available. This closure is illustrated in the right plot of Fig. 4.21.




Figure 4.21: Left: dependence of the average response $\langle \mathbf{m}_{\text{reco}}/\mathbf{m}_{\text{true}}|\mathbf{m}_{\text{true}}, \mathbf{p}_T^{\text{reco}} \rangle$ as a function of $\mathbf{p}_T^{\text{reco}}$ and $\mathbf{m}_{\text{true}}$. Right: the fitted distribution when using the detector-level simulation in place of the data. By construction, the fitted distribution is identical to the input detector-level distribution ($\mathbf{r} = \mathbf{s} = 1$).

4.1.2.3.3 Systematic Uncertainties

There are two sources of systematic uncertainty in the extraction of the relative jet mass scale and jet mass resolution when using the forward-folding method. First,

there is a theoretical modeling uncertainty on the particle-level jet mass distribution. Second, there are sources of theoretical or experimental uncertainties that impact the mass response. The method is not sensitive to sources of uncertainty that only change the overall normalization. Uncertainties are estimated by varying the simulation and then re-fitting the data. The difference between the nominal fitted values of r, s and the variation fits is used as the systematic uncertainty. Sources of theoretical modeling uncertainty include the NLO matching scheme, fragmentation, and initial- and final-state radiation (ISR/FSR). Any impact on the NLO matching scheme is estimated by replacing the nominal POWHEG-BOX+PYTHIA 6 $t\bar{t}$ sample with alternative samples generated with POWHEG-BOX+HERWIG and MC@NLO+HERWIG (all other processes remain unchanged). The fragmentation uncertainty uses a comparison between POWHEG-BOX+PYTHIA 6 with POWHEG-BOX+HERWIG and the ISR/FSR uncertainty is estimated by comparing two variations of POWHEG-BOX+PYTHIA 6 with different Perugia 2012 tunes⁹. The background modeling and experimental uncertainties have a much smaller impact on the relative jet mass scale and resolutions compared with the theoretical modeling uncertainties. The uncertainty on the W +jets background is the most relevant near the W mass peak, which is due to the statistical uncertainty on the charge asymmetry method used to derive the normalization [281].

4.1.2.3.4 Results

Figure 4.22 shows the χ^2 minimization for the relative jet mass scale and the relative jet mass resolution. Each value on the curve is the χ^2 per degree of freedom when fitting either the simulation or the data with a template from the simulation using a resolution function whose scale or resolution is shifted or stretched by the value indicated on the horizontal axis. As the fit is performed simultaneously for the relative scale and resolution, the curves in Fig. 4.22 are the value of the χ^2 per degree of freedom at a given relative jet mass scale or jet mass resolution minimized

⁹The factorization/renormalization scales and the h_{damp} parameter are simultaneously varied, but are expected to have a smaller impact on the jet mass (response). See Fig. 4.18 for details.

over the relative jet mass resolution or jet mass scale, respectively. By construction, the minimum χ^2 for the simulation fit to itself is zero at a relative scale of one. The relative jet mass scale is 1.001 and the relative jet mass resolution is 0.96. The range of the vertical axes relative to the range of the horizontal axes indicates that there is significantly more sensitivity to the relative scale than the relative resolution. The fit improves the χ^2 per degree of freedom by about 10%. A comparison between the template from simulation with the fitted parameters and the data is shown in Fig. 4.23. The dashed line is the particle-level jet mass spectrum that is smeared to detector-level before comparing with the data. The solid line shows the detector-simulation before fitting the relative jet mass scale and resolution and the dotted red line shows the post-fit distribution. There is only a small decrease in the χ^2/NDF from the fit, so the two distributions are similar. In the ratio plot, the band is the statistical uncertainty from the data while the black and red points are the pre- and post-fit ratios of the simulation with the data.




Figure 4.22: The value of the χ^2 per degree of freedom (NDF) at a given relative jet mass scale (left) or jet mass resolution (right) minimized over the variable not shown.

One- and two- σ statistical uncertainty ellipses are shown in Fig. 4.24. The statistical uncertainty is computed using the bootstrap technique: N pseudo-datasets are generated by (re)sampling from the data with replacement. Each pseudo-dataset i is




Figure 4.23: A comparison between the post- and pre-fit simulation and the data. The vertical dotted lines indicate the fit range.




Figure 4.24: The one (solid) and two (dashed) σ uncertainty ellipses for the relative jet mass scale and the relative jet mass resolution.

then fit with the nominal simulation and the measured values r_i and s_i are recorded. The circular marker in Fig. 4.24 represents $(\langle s_i \rangle, \langle r_i \rangle)$ while the star indicates the result of fitting the simulation to itself, which is at the point (1,1) by construction. Appendix C.1 describes how uncertainty ellipses are constructed. The mean and standard deviation of the joint distribution are estimated using the sample mean and standard deviation over the ensemble of pseudo-datasets. The values C for the 1σ and 2σ ellipses are computed by solving

$$\frac{1}{2} \int_0^C dx e^{-\frac{x}{2}} = \frac{1}{\sqrt{2\pi}} \int_{-Z}^Z dx e^{-\frac{x^2}{2}}, \quad (4.30)$$

where $Z = 1$ for the 1σ contour and $Z = 2$ for the 2σ contour. These results are combined with the systematic uncertainties in Table 4.1. The amount of (ISR and) FSR is the dominant uncertainty for extracting the relative JMR because the fitted resolution width compensates for changes in the width of the particle-level mass distribution (see Fig. 4.18). The total systematic uncertainty is about 2% for the relative jet mass scale and about 19% for the relative jet mass resolution.

Source of Uncertainty	Jet Mass Scale (s)	Jet Mass Resolution (r)
NLO matching	0.017	0.08
Fragmentation	0.018	0.05
ISR/FSR	0.004	0.15
Jet Energy Scale	0.002	0.03
Jet Energy Resolution	0.001	0.03
b -tagging	< 0.001	0.01
MC Normalization	0.001	0.01
Total Systematic Uncertainty	0.024	0.18
Data Statistical Uncertainty	0.004	0.05
Value	1.001	0.96

Table 4.1: A summary of the measured relative jet mass scale and jet mass resolution using both the subtraction and the forward folding methods. Uncertainties are given as a fraction of the nominal. The jet energy scale, ISR/FSR, and MC Normalization background uncertainties are treated as asymmetric but the maximum of the two variations are reported in this table.

4.1.2.3.5 Intermediate Conclusions

This section has reported a measurement of the relative jet mass scale and jet mass resolution using a sample enriched in boosted hadronically decaying W bosons from $t\bar{t}$ events in the $\sqrt{s} = 8$ TeV data collected by the ATLAS detector. A new method called *forward folding* uses non-parameteric shapes for both the particle-level distribution and the response function, derived from the simulation. The relative jet mass scale and jet mass resolution are compatible with unity within the statistical uncertainties at 0.4% for the jet mass scale and 5% for the jet mass resolution. This measurement can be used in the future to set a systematic uncertainty on the jet mass scale and the jet mass resolution for BSM searches. The scale and resolution are compatible with unity, but the uncertainty on the measurement should be used as uncertainties for analyses that use the large-radius jet mass to identify jets are resulting from boosted heavy particle decays. However, there are important caveats to this measurement:

Topology Dependence The jet mass scale and the jet mass resolution presented in this measurement use boosted hadronically decaying W bosons from $t\bar{t}$ events. The relative mass scale and resolution may depend on the jet p_T , the jet mass, the number of subjets within the jet, close-by radiation, and the presence of heavy flavor decays inside the jet.

Particle-level Input The measured hadronically decaying W boson resonance peak contains information about the convolution of the particle-level spectrum and the resolution function. In this measurement, the particle-level spectrum is taken as input to extract the resolution function. Therefore, the relative scale and resolution presented here are not applicable as uncertainties for precision measurements of the particle-level spectrum.

More, higher energy data will be available in Run 2 that will allow for many of the above challenges to be addressed as a new frontier is opened for new physics searches and precision measurements at high energies.

4.1.2.3.6 Improvements and Prospects with 13 TeV data

The 3.2 fb^{-1} of $\sqrt{s} = 13 \text{ TeV}$ data collected in 2015 are used in this section for a preliminary measurement of the relative jet mass scale and resolution, as well as for the introduction of a new technique to measure the jet p_T scale and resolution with forward-folding¹⁰. Even though the integrated luminosity is significantly lower with the early Run 2 dataset compared with Run 1, the increase in the inclusive $t\bar{t}$ cross-section coupled with a further increase at high p_T makes the total number of $t\bar{t}$ events roughly comparable. In particular, the inclusive cross-section increases by about a factor of 3.5 and Fig. 4.25 shows that there is another factor of ~ 2 at high p_T . The ratio of the number of top quarks pairs is therefore about $20.3/(3.2 \times 3.5 \times 2) \sim 1$.




Figure 4.25: The probability distribution (left) and the integral of the probability distribution above a given threshold (strictly speaking, this is one minus the cumulative distribution) on the top quark p_T for $\sqrt{s} = 8$ and 13 TeV. The increase in the inclusive $t\bar{t}$ cross-section between these two center-of-mass energies is about 3.5.

¹⁰The results in this section are published in Ref. [347] and include input from D. Melini, N. Norjoharuddeen, and M. Vos.

Figure 4.26 shows the jet mass distribution at $\sqrt{s} = 13$ TeV. The forward-folding method can be applied to any jet mass definition, such as the track-assisted jet mass (Sec. 4.1.4) shown in the bottom plots of Fig. 4.26. A scan in the relative jet mass scale and jet mass resolution produce the χ^2 curves in Fig. 4.27 that are analogous to the $\sqrt{s} = 8$ TeV curves from Fig. 4.22. As was the case at $\sqrt{s} = 8$ TeV, there is more sensitivity to the relative JMS than the JMR, evident from the width near the minimum χ^2 .

A comparison of the Run 1 and early Run 2 measured JMS and JMR is shown in the left plot of Fig. 4.28. Even though the number of top quark pairs is comparable between the two datasets, the statistical uncertainty ellipse is slightly larger for the $\sqrt{s} = 13$ TeV measurement as only the muon channel is used. As indicated in the caption of Fig. 4.26, the jet mass definition is slightly different between the two datasets: R_{sub} is 0.2 instead of 0.3 in order to be able to resolve the subjets of ultra boosted $W/Z/H$ bosons and top quarks. Despite this difference, the relative resolutions are statistically comparable with each other with a significance slightly above 1σ . The right plot of Fig. 4.28 shows the relative JMS and JMR for the two mass reconstruction algorithms in Fig. 4.26. The values are similar within the statistical uncertainties, but as expected, the systematic uncertainty is smaller for the track-assisted jet mass (more details in Sec. 4.1.4).

An innovation of the early Run 2 analysis is the extension of the forward-folding technique to measure the relative jet p_T scale and resolution. The resolution of any quantity can be measured with the forward-folding technique so long as one can identify a detector-level distribution that depends strongly on that resolution. The two top quarks in $t\bar{t}$ production tend to be produced with a similar p_T . Therefore, the leptonic top quark can be used as a proxy for the hadronically decaying top quark and thus $p_{\text{jet}}/p_T^{\text{lep top}}$ should be sensitive to the jet p_T scale and resolution of the hadronically decaying top quark. The full leptonic top quark p_T can be reconstructed from the two-vector sum of the lepton momentum, the $\vec{p}_T^{\text{miss}}$, and the transverse momentum of the nearby jet (see Sec. 4.1.2.3.1). One disadvantage of using the full leptonic top p_T is that it depends on calorimeter quantities ($\vec{p}_T^{\text{miss}}$ and the jet p_T). One can reduce the calorimeter-dependence by using either just the two-vector sum




Figure 4.26: The jet mass computed directly from calorimeter-cell clusters (top) and with the track-assisting procedure (bottom) for a lower p_T (left) and a higher p_T (right). The track-assisted jet mass is defined in Sec. 4.1.4. Note that the ATLAS default trimming parameters for Run 2 are slightly different than Run 1: R_{sub} is now 0.2 instead of 0.3 in order to improve the resolution at high p_T . The event selection is identical to the one at $\sqrt{s} = 8$ TeV, but only the muon channel is used in this section (negligible multijet contribution). The bands include detector-level jet and particle-level modeling systematic uncertainties. For illustration, the top left plot shows also the particle-level distribution. Note that all distributions are normalized to the data integral.




Figure 4.27: The value of the χ^2 per degree of freedom (NDF) at a given relative jet mass scale (left) or jet mass resolution (right) minimized over the variable not shown.




Figure 4.28: Left: The statistical (solid) and total (dashed) 1σ uncertainty ellipses for the relative jet mass scale and the relative jet mass resolution at $\sqrt{s} = 8$ TeV (red) and $\sqrt{s} = 13$ TeV (black). Right: the relative scale and resolutions for the two mass reconstruction algorithms in Fig. 4.26.

of the jet p_T and the lepton p_T or just the lepton p_T . The tradeoff for the reduced calorimeter-dependence is the reduced sensitivity to the large-radius jet p_T resolution.

This is illustrated by the ‘peakiness’ of the three distribution in Fig. 4.29. All three ratios show a peak near one and so would shift if the jet p_T scale were mis-modeled. However, the peak is sharpest with the full leptonic top, which suggests that the statistical uncertainty on the JER will be smallest when using this quantity.




Figure 4.29: The ratio of the large-radius jet p_T to the lepton p_T (left), the two-vector sum of the lepton and close-by jet transverse momentum (middle), and the full leptonically decaying top quark using the lepton, the close-by jet, and the $\vec{p}_T^{\text{miss}}$ (right).

A χ^2 fit using the forward-folding method for each of the three observables from Fig. 4.29 is shown in Fig. 4.30 for the JES and in Fig. 4.31 for the JER. The particle-level quantity entering the forward-folding is $p_T^{\text{particle-level jet}}/p_T^{\text{ref}}$, where p_T^{ref} is one of the *detector-level* quantities from Fig. 4.29. As expected from the sharpness of the ratio distributions in Fig. 4.29, the ratio with the full leptonic top quark candidate has the deepest χ^2 for the JES. The χ^2 distribution near the minimum in Fig. 4.31 is rather flat: there is not much sensitivity to the JER due to the large width of the ratio distributions (large relative to the JER itself). However, there is currently no in-situ constraint on the large-radius jet JER and so even a crude uncertainty is an important step forward.

Table 4.2 summarizes the fitted values of the relative JES, JER, JMS, and JMR using the early Run 2 dataset for the mass and p_T . The three jet p_T scale determinations are not all statistically consistent with each other, but there are significant systematic uncertainties that are not fully correlated between the methods. Theoretical modeling uncertainties, which are an important ingredient for the particle-level




Figure 4.30: The value of the χ^2 per degree of freedom at a given relative jet p_T scale minimized over the jet p_T resolution for three different reference objects: the lepton p_T (left), the two-vector sum of the lepton and close-by jet transverse momentum (middle), and the full leptonically decaying top quark (right).




Figure 4.31: The value of the χ^2 per degree of freedom at a given relative jet p_T resolution minimized over the jet p_T scale for three different reference objects: the lepton p_T (left), the two-vector sum of the lepton and close-by jet transverse momentum (middle), and the full leptonically decaying top quark (right).

templates dominate over the JES/JMS uncertainty¹¹ in most cases. The more calorimeter information used to extract the p_T scale and resolution, the larger the experimental uncertainty. The two methods that do not use the missing momentum have a similar precision for both the JES and JER.

These early measurements of the relative jet mass and jet p_T scale and resolution are already dominated by systematic uncertainties. In the future, it is therefore

¹¹The large-R JES (JMR) is not included as an uncertainty when the JES (JMR) is measured.

crucial to perform this measurement differentially in p_T and jet substructure (such as n_{subjets}) which will mitigate (some) of the modeling dependence.

Quantity		Value	Stat. Uncert	Modeling	Jets	Total Syst.
m^{calo}	$s_{\text{data}}^{\text{MC}}$	0.984	0.6 %	1.7 %	1.6 %	2.3 %
m^{calo}	$r_{\text{data}}^{\text{MC}}$	1.047	6.6 %	18.1 %	7.0 %	19.4 %
m^{TA}	$s_{\text{data}}^{\text{MC}}$	0.981	1.1 %	2.4 %	4.8 %	5.3 %
m^{TA}	$r_{\text{data}}^{\text{MC}}$	1.036	6.1 %	14.6 %	5.0 %	15.5 %
$p_{T,\text{jet}}/p_T^{\text{lep}}$	$s_{\text{data}}^{\text{MC}}$	1.011	0.7 %	1.3 %	0.4 %	1.3 %
$p_{T,\text{jet}}/p_T^{\text{lep}}$	$r_{\text{data}}^{\text{MC}}$	0.945	4.1 %	6.8 %	2.7 %	7.3 %
$p_{T,\text{jet}}/p_T^{\text{lep} + \text{b-jet}}$	$s_{\text{data}}^{\text{MC}}$	0.985	0.4 %	0.7 %	1.2 %	1.4 %
$p_{T,\text{jet}}/p_T^{\text{lep} + \text{b-jet}}$	$r_{\text{data}}^{\text{MC}}$	0.903	6.1 %	5.5 %	4.7 %	7.2 %
$p_{T,\text{jet}}/p_T^{\text{lep top}}$	$s_{\text{data}}^{\text{MC}}$	0.987	0.2 %	0.3 %	2.1 %	2.1 %
$p_{T,\text{jet}}/p_T^{\text{lep top}}$	$r_{\text{data}}^{\text{MC}}$	1.024	3.1 %	6.2 %	6.0 %	8.6 %

Table 4.2: Summary of the systematic uncertainties for the relative jet mass or energy scales ($s_{\text{data}}^{\text{MC}}$) and resolutions ($r_{\text{data}}^{\text{MC}}$). The first column states which observable is used to extract the relative jet mass (first four rows) or jet energy (rows 5-10) scale and resolutions.

4.1.3 Re-clustered Jet Mass

As discussed¹² in Sec. 4, the angular separation between decay products of a massive particle $\mathcal{P}$, such as a W or Z boson, scales as $2\mathbf{m}_{\mathcal{P}}/\mathbf{p}_T^{\mathcal{P}}$. This suggests that the radius parameter R of jet clustering algorithms aimed at collecting the hadronic decay products of $\mathcal{P}$ should be process dependent and scale with the momentum under consideration. However, at the LHC, most analyses use one global value of R fixed ahead of time. In ATLAS, this value is $R = 1.0$ for large-radius jets and $R = 0.4$ for small-radius jets. The reason for a fixed jet radius is that every jet configuration, which includes the algorithm, radius, and grooming parameters, must be calibrated to account for unmeasured energy deposits and other experimental effects [225, 349], even though the inputs to jet clustering are themselves calibrated. The calibration of inputs provides a partial calibration to the jet, but jet energy and mass scale corrections provide a *full* calibration by also correcting for particles that were missed, merged, or below noise thresholds, energy loss in un-instrumented regions of the calorimeter, and additionally takes into account correlations between particles. The dependence on these additional calibrations thus makes it desirable to reconsider the current jet clustering paradigm in favor of a modular structure that allows for a much broader class of algorithms and radius parameters to be selected by analyses.

One solution is to introduce a new angular scale $r < R$, such that jets of radius r can be the inputs to the clustering algorithm of large radius R jets¹³. If chosen appropriately, the fully calibrated small radius jets can make the calibration of the re-clustered large radius jets automatic. Furthermore, with no additional calibration needed, any large radius R , any clustering algorithm, and many grooming strategies can be simultaneously implemented in an analysis. Using optimal parameters can, for instance, significantly improve the discovery potential of searches for new physics [119]. In particular, every kinematic region of every analysis for every data-taking condition (e.g. level of pileup) can be individually optimized in order

¹²The phenomenological studies presented in this section are published in Ref. [348] and include technical inputs from M. Swiatlowski and P. Nef.

¹³Similar ideas have been proposed in the past such as variable R jets [350]. While these methods address the variability of R , they do not address the concerns about calibrations and uncertainties.

to maximize the sensitivity to particular physics scenarios. Another benefit is that the uncertainties on the re-clustered p_T and mass are also automatic consequences of propagating the corresponding uncertainties computed for small radius jets. In this way, the re-clustered jet mass can be viewed as any other kinematic variable, such as di-, tri-, or multi-jet invariant masses that are ubiquitous in measurements and searches for new physics.

The idea of re-clustering small radius jets is not new. These objects first appeared in an ATLAS search for supersymmetry in the multijet final state [351] and more recently in an ATLAS search for direct stop quark pair production in the all hadronic final state [352]. There are also related techniques which group small radius jets together to form pseudo-jets [353] or mega-jets [354–356]. This section introduces a new way of thinking about re-clustering. Instead of viewing jet grouping as a high-level analysis technique, the idea is to consider re-clustered jets as if they were any other jet collection clustered directly from low-level objects. This is a significant paradigm shift because there is an entire class of techniques for using and improving large-radius jets. For example, re-clustered jets can be groomed and their substructure can be useful for tagging.

This section is organized as follows. Section 4.1.3.1 introduces the technical details and benefits of re-clustering small-radius jets. Sections 4.1.3.2 and 4.1.3.3 describe performance studies at particle-level for the jet mass and other jet substructure, with a particular emphasis on pileup. The dependence of jet tagging on re-clustering parameters with the full ATLAS detector simulation is studied in Sec. 4.1.3.4 and the performance of re-clustered jet mass reconstruction relative to standard large-radius jets is discussed in Sec. 4.1.3.5. One of the key assumptions of re-clustering is that the impact of close-by jets on the jet energy scale response is small or at least well-modeled by the simulation. In-situ track jet methods are used to investigate the impact of close-by jets in Sec. 4.1.3.6. Section 4.1.3.7 provides an overview and outlook for re-clustering.

4.1.3.1 Re-clustering Jets

The inputs of jet clustering algorithm are typically stable particles (Monte Carlo truth studies), topological clusters (ATLAS), or particle flow objects (CMS). Re-clustered large radius R jets take as input the output of the small radius r jet clustering. Small radius jets have been calibrated with r as small as 0.2 [357] and there are no indications of sizable mis-modelling of close-by effects in the jet response for the standard $R = 0.4$ jets [225] (see Sec. 4.1.3.6 for more detail). In general, the algorithm used to cluster the small radius jets can be different than the algorithm used for re-clustering the entire event. Fig. 4.96 shows a simple example of an event clustered with anti- k_t $R = 1.0$ and with anti- k_t $R = 1.0$ re-clustered $r = 0.3$ anti- k_t jets. Unlike the inputs of clustering which are e.g. measured in a calorimeter and can be reconstructed and individually calibrated with very low energy, small radius jets can only be reliably fully calibrated for $\gtrsim 15$ GeV [225, 349], where the actual threshold may depend on r . This minimum p_T threshold acts as an effective grooming for the re-clustered jets (RC). This is seen clearly in Fig. 4.96, where the blue large radius jet has many constituents far away from the jet axis (which have low p_T) and are not part of the re-clustered jet. One could choose a more aggressive threshold to, for instance, remove the impact of additional pp collisions (i.e. pileup) on the jets. A more dynamic grooming scheme, named *re-clustered jet trimming* in analogy to large radius jet trimming [227], sets the p_T cut on the small radius jets based on the large radius jet p_T (calculated before any small- r jets are removed). Specifically, for re-clustered and trimmed jets (RT), the grooming removes any small radius jet constituent j of a large R re-clustered jet J if $p_T^j < f_{\text{cut}} \times p_T^J$. The parameter f_{cut} can be optimized for a particular kinematic selection and event topology. Other grooming schemes are possible, but beyond the scope of this section¹⁴.

Due to the increased catchment area of large radius jets over small radius jets, they are more susceptible to contributions from pileup. Just as there are pileup correction techniques for large radius jets and their subjets, one can benefit from pileup corrections to the small radius jet inputs that propagate to re-clustered jets. In

¹⁴Jet grooming procedures applied to jets-as-inputs have been studied in the past (see for instance Ref. [358]); these and other algorithms can be adopted to the re-clustering paradigm.




Figure 4.32: An example event which has been clustered using the anti- k_t $R = 1.0$ (left) and with anti- k_t $R = 1.0$ re-clustered $r = 0.3$ anti- k_t jets (right). The shaded regions show the jet area determined by clustering ghost particles. Only large radius jets with $p_T > 50$ GeV are shown and small radius jets are required to have $p_T > 15$ GeV. As with groomed large radius jets, re-clustered jets can have non-circular shapes.

particular, one can remove jets from pileup interactions with techniques like JVT [124] or pileup jet identification [359] and can correct the remaining jets with methods like the four-vector jet areas subtraction. Another way to mitigate the impact of pileup is to correct jet constituents before clustering [360–362], which is similar to the $r \rightarrow 0$ limit. However, applying a jet constituent pileup correction still requires an overall calibration and an intermediate scale $r \sim 0.1 - 0.5$ is one possibility.

In the growing field of jet substructure, there are many jet observables which depend explicitly on the jet constituents, not just the jet four-vector. These techniques are still applicable for re-clustered jets. One possibility is to compute substructure observables using the small-radius jet constituents inside the re-clustered jet. This approach should be similar to the jet substructure of a traditional groomed large-radius jet. An alternative *bottom-up* approach to jet substructure is to use the radius r jets directly as the inputs to jet substructure. The advantages and limitations of

bottom-up substructure are described in Section 4.1.3.3.

There are other technical benefits to re-clustering. For example, re-clustering can be much faster than traditional jet clustering. Jet clustering is an order $N \log N$ operation [363] – as the number of jet inputs in an event approaches 500 or more at high pileup conditions, jet clustering can take a significant amount of the full event reconstruction time. This is particularly relevant if one wants to scan the jet clustering parameters. However, there are typically 10 or fewer jets above the calibrated p_T threshold in any given event (though this obviously depends on the R size and threshold). With these typical numbers, creating a re-clustered jet is about 100 times faster than clustering a large- R jet directly. This kind of computational speed-up can allow analysis end-users – and not just large, central productions – to produce their own large- R jets, allowing for more creativity in exploring the optimal jet algorithms and parameters for analyses.

4.1.3.2 Particle-level Jet Mass Performance

Three processes are generated using PYTHIA 8.170 [96, 189] at $\sqrt{s} = 14$ TeV for studying the efficacy of re-clustered jets. Hadronic W boson and top quarks are used for studying hard 2- and 3-prong type jets. To simulate high p_T hadronic W decays, W' bosons are generated which decay exclusively into a W and Z boson which subsequently decay in quarks and leptons, respectively. The p_T scale of the hadronically decaying W is set by the mass of the W' which is tuned to 800 GeV for this study so that the $p_T^W \lesssim 400$ GeV. In this p_T^W range, not all of the decay products of the W are expected to merge into a small radius jet of $r \lesssim 0.4$, but should merge within a cone of $R = 1.0$. A sample enriched in 3-prong type jets is generated with $Z' \rightarrow t\bar{t}$, with $m_{Z'} = 1.0$ TeV, so that $p_T^t \gtrsim 350$ GeV. To study the tradeoff between signal and background jet identification, QCD dijets are generated with a p_T spectrum similar to the relevant signal process. Pileup is modeled by overlaying additional independently generated minimum-bias interactions with each signal event. The number of pileup interactions is between LHC Run 1 conditions, $n_{PU} = 20$, and the conditions toward the end of the LHC Run 2, $n_{PU} = 80$.

Jet are re-clustered using FASTJET [203] 3.0.3. While the large radius jets can

be defined using any set of parameters, the studies in this section use a fixed large jet algorithm: anti- k_t algorithm with $R = 1.0$. The reference jets are trimmed using $R_{\text{sub}} = 0.3$ k_t subjects with a p_T fraction threshold of $f_{\text{cut}} = 0.1$. Unlike the procedure used by most analyses, all momenta are pileup corrected using the jet areas technique prior to grooming. This is the natural setup for re-clustering and makes the optimal grooming parameters independent of n_{PU} .

Re-clustering is investigated with a series of schemes for the small radius jets: anti- k_t radius parameters in $\{0.2, 0.3, 0.4\}$ grooming $f_{\text{cut}} = 0.1$ and 0.2 (with a $p_T = 15$ GeV threshold). This list is not exhaustive, but encompasses a relevant set of parameters. Radii below $r = 0.2$ are not considered due to experimental limitations from calorimeter granularity and theoretical considerations from non-trivial non-perturbative effects. All small radius jets are required to have $p_T > 15$ GeV.

As the jet mass is the mostly widely used large-radius jet observable, it is used to benchmark various re-clustering schemes. The jet mass performance is quantified by the average jet mass $\langle m \rangle$, the standard deviation of the jet mass distribution (σ), and the dependance of these quantities with the amount of pileup. The averages and deviations are computed over a fixed mass range: 60-100 GeV. Another useful metric is the efficiency of a $60 < m_{\text{jet}}/\text{GeV} < 100$ requirement. Figures 4.33 and 4.34 compare RC with two settings of RT. In the region near the W mass peak, re-clustered trimming with $f_{\text{cut}} = 0.2$ performs the best in terms of the mass distribution standard deviation in the W mass window. However, there is a sizable peak at low mass where too many jets have been cut out by the aggressive trimming parameter. The fixed cut of 15 GeV is too low, especially at very high pileup where the large high mass tail is much bigger for RC than for RT. The re-clustered trimming using anti- k_t with the same f_{cut} as the trimming has very similar performance, though the peak position is slightly higher. Figure 4.34 shows the performance metrics as a function of NPV for the various grooming schemes. The average mass for RT is very stable, whereas there is a slight slope for RC. The mass resolution for RC is slightly worse than for RT, but the efficiency of RC is better because it avoids the peak at low masses well below the W boson mass.

The re-clustered jet mass distribution for several small radius jet sizes is shown




Figure 4.33: Various re-clustered grooming parameters for anti- k_t $r = 0.3$ jets for NPV = 20 on the left and NPV = 80 on the right.




Figure 4.34: Mean, mass resolution, and mass window efficiency of the mass distribution as a function of the number of additional vertices for various re-clustered jet grooming schemes.

in Figure 4.35 and the performance metrics are quantified in Fig. 4.36. For all three considered values of r , the minimum p_T cut is 15 GeV. In practice, this could be optimized, since smaller radius jets may be calibrated at smaller values of p_T . An alternative approach is to use *iterative re-clustering* by re-clustering $r = 0.2$ into $r' = 0.4$ and then into $R = 1.0$ to further increase the flexibility of the jet algorithms (also this reduces the effective jet area and so the resulting jets would be less susceptible

to pileup¹⁵). The right plot of Figure 4.35 and the top right plot of Fig. 4.36 show the $r = 0.2$ setting as resulting in the most peaked mass distribution.




Figure 4.35: Various small radii for a fixed algorithm of anti- k_t for NPV = 20 on the left and NPV = 80 on the right.




Figure 4.36: Mean, mass resolution, and mass window efficiency of the mass distribution as a function of the number of additional vertices for various small jet radii.

¹⁵If viewed as a uniform noise in the calorimeter, the contribution of pileup to a given jet scales proportionally to its area. However, there are local fluctuations that complicate this picture.

4.1.3.3 Re-clustered Jet Substructure

One natural possibility for computing jet substructure observables for re-clustered jets is to use the constituents of the small-radius jets inside the large-radius jet. An alternative method is to use the radius r jet momenta directly. For example, consider the k_t splitting scale¹⁶ $\sqrt{\bar{d}_{n,n+1}}$, which is sensitive to hard $(n+1)$ -prong structure in a jet. One can use directly the radius r jets inside a radius R re-clustered jet to compute $\sqrt{\bar{d}_{n,n+1}}$. If there are only two radius r jets, then $\sqrt{\bar{d}_{12}}$ is simply the k_t distance between the radius r jets. The advantage of this approach is that there is a natural prescription for calibrations and systematic uncertainties. The jet energy scale calibration and its uncertainties directly translate into the calibration of the bottom-up substructure variables. Furthermore, in this approach one knows how the substructure variable calibrations and uncertainties are correlated with the re-clustered jet calibrations and uncertainties. This information is available for the first time with this bottom-up procedure.

Figure 4.37 compares bottom-up and top-down jet substructure variables in classifying $Z' \rightarrow t\bar{t}$ and QCD multijet events. For the chosen parameters, the two techniques have comparable performance. The main drawback of bottom-up substructure is that the relative efficacy depends on p_T (and r). When $r \gtrsim m/p_T$, or equivalently, when there are not many radius r jets inside the radius R jet, the experimental gains from bottom-up substructure are diminished. For instance, if there is only one radius r jet, then $\sqrt{\bar{d}_{12}} = 0$. Thus, in certain kinematic regimes, bottom-up substructure may provide a powerful alternative to standard methods, but in other regimes a more dedicated analysis is required to understand correlations in calibrations and uncertainties (when jet substructure observables are built from the jet constituents).

¹⁶Computed by re-clustering a jet's constituents using the k_t algorithm and then considering the distance metric of the last n un-clusterings.




Figure 4.37: The performance of a bottom-up approach to jet substructure where the radius r jets are inputs to substructure variables. Solid lines show the performance curves for large radius trimmed jets ($R_{\text{sub}} = 0.3, f_{\text{cut}} = 0.05$) and the dashed lines show the analogous re-clustered variable. Random tagger denotes a classifier which picks signal and background with equal probability. The variable N_{sub} is the number of re-clustered subjets. The curves are determined by placing threshold requirements on the variable likelihoods.

4.1.3.4 Detector-level Jet Tagging

The studies¹⁷ in Sec. 4.1.3.2 are an important first step to quantifying the dependence of re-clustered jet mass performance on one of the most important experimental conditions, pileup. This section expands upon the study by using the full ATLAS detector-simulation to investigate how the full detector-resolution impacts the re-clustered jet mass performance. Re-clustering parameters are varied and the resulting jets are compared with traditional large-radius jets clustered directly from calorimeter-cell clusters. The relevant re-clustering parameters are f_{cut} , r , R and $p_{\text{T}}^{\text{cut}}$ (the small radius jet p_{T} threshold). Two metrics used for comparing algorithms are the *window size*, which is the size of the mass interval which contains at least 68% of the signal, and *window efficiency*, which is the fraction of background events which fall in the 68% window. PYTHIA $W' \rightarrow WZ$ events are signal and PYTHIA QCD dijets are background. The signal p_{T} spectrum is re-weighted to match that of the background¹⁸. The study is decomposed into three p_{T} ranges: $p_{\text{T}} \in [200 - 350]$, $[350, 500]$, and $[500 - 1000]$ GeV. The radius r jets used for re-clustering are pileup corrected but not calibrated. Pruned [364, 365] C/A jets with $R = 0.8$ are used as a benchmark as they perform well across a wide range of phase space [166].

4.1.3.4.1 Low p_{T} : 200 GeV - 350 GeV

Figure 4.38 compares anti- k_{t} $R = 1.0$ trimmed jets with $f_{\text{cut}} = 0.05$ and $R_{\text{sub}} = r$ with the analogous RT jets in the range $200 \text{ GeV} < p_{\text{T}} < 350 \text{ GeV}$. The main difference between the RT and traditional large-radius jets is that the small-radius jets for the former are anti- k_{t} while the k_{t} algorithm is used for subjet finding for the latter. As the W decay products are well-resolved by $r = 0.2, 0.3$, and $r = 0.4$ small-radius jets, the mass window is about the same size for all three algorithms. However, the small-radius jets with a larger size have a worse rejection (higher efficiency) for

¹⁷The results presented in this section include input from M. Solt. In particular, Solt made the final versions of the plots comparing the various algorithms.

¹⁸Instead of re-weighting for each jet algorithm, the p_{T} spectrum is weighted based on the leading C/A particle-level ungroomed jet with $R = 1.2$. This algorithm was chosen because in the signal, it has a high efficiency for capturing all of the W boson energy.

the background because the large small-radius jets tend to have a higher p_T and so a second (or third) background jet can survive the trimming and significantly increase the mass.




Figure 4.38: Left: The efficiency for signal and background jets to be reconstructed in the smallest mass window containing 68% of the signal. Right: the size of the window from the left plot. Both plots compare large radius trimmed jets with $R = 1.0$, $f_{\text{cut}} = 0.05$ and $R_{\text{sub}} = 0.05$ with re-clustered jets that have analogous parameters.

At low p_T , the choice of R is particularly important as the boson boost is not yet high enough to capture all of the decay products into one small radius jet. Therefore, re-clustered jets need to have multiple constituents in order to have a mass compatible with the W boson mass. Figure 4.39 shows the small-radius jet constituent multiplicity for several choices of R .

Figure 4.39 shows the p_T spectrum of the re-clustered jets in the lowest p_T range along with the number of small radius jets for $r = 0.2$ and several values of R . The distributions for $R \gtrsim 0.8$ are similar and collectively are significantly different than the $R = 0.6$ case. The loss of constituents degrades the mass-tagging performance, as illustrated in Fig. 4.40. The jets with $R \gtrsim 1.0$ have a significantly higher background rejection than re-clustered jets with $R < 1.0$. For $R = 0.6$, this is explained by the fact that many re-clustered jets have only one small-radius jet constituent for which

the jet mass distribution is nearly identical to the mass distribution for the QCD jet background. The $R = 0.8$ point in the left plot of Fig. 4.40 seems inconsistent with Fig. 4.39, which suggests that there is only a small difference in the constituent multiplicity between $R = 0.8$ and $R = 1.0$. This is explained by the bimodal mass distribution in the left plot of Fig. 4.41. Re-clustered jets with one constituent have a low-mass Sudakov peak while jets with multiple constituents have a jet mass near m_W . When the low mass peak contains more than $(100 - 68)\%$ of the distribution, the 68% interval is split across the two peaks, which is why the $R = 0.8$ point has such a larger window in the right plot of Fig. 4.40. In contrast, the $R = 0.6$ point has at least 68% in the low mass peak alone while the $R > 0.8$ points have at least that much probability in the high mass peak.




Figure 4.39: The small-radius jet constituent multiplicity for several choices of R and a fixed $r = 0.2$.

For a fixed value of R , the f_{cut} is also an important parameter, as it determines the amount of radiation that is preserved in the clustering. Figure 4.42 shows that so long as f_{cut} is large enough to remove unwanted radiation ($f_{\text{cut}} \sim 0.04$) and is small enough to preserve the hard structure ($f_{\text{cut}} \sim 0.1$), the performance metrics do not depend strongly on f_{cut} . Outside of these regimes, there is a strong dependence as unwanted radiation is preserved or desired radiation is removed.




Figure 4.40: Left: The efficiency for signal and background jets to be reconstructed in the smallest mass window containing 68% of the signal. Right: the size of the window from the left plot. In both plots, $r = 0.2$ and $f_{\text{cut}} = 0.05$.




Figure 4.41: The jet mass distribution for the signal (background) for various R values. In both plots, $r = 0.2$ and $f_{\text{cut}} = 0.05$.




Figure 4.42: Left: The efficiency for signal and background jets to be reconstructed in the smallest mass window containing 68% of the signal. Right: the size of the window from the left plot. In both plots, $r = 0.2$ and $R = 1.0$.

4.1.3.4.2 High p_T : 500 GeV - 1000 GeV

The trimming is harsher at high p_T so the r -dependence of the QCD rejection is slightly lower in Fig. 4.45 compared with Fig. 4.38. As a result, the background rejection for re-clustered jets is nearly independent of r and large-radius trimmed jets have only a small dependence. In this high p_T regime, most $r = 0.4$ jets have only one constituent (right plot of Fig. 4.44) while $r = 0.2$ jets still usually have at least two (left plot of Fig. 4.44). As a result, the width of the mass peak is largely insensitive to f_{cut} as long as it is not too large that it removes the second jet ($r = 0.2$) or too low that it lets in extraneous radiation ($r = 0.4$). In the range $0.04 \lesssim f_{\text{cut}} \lesssim 0.1$, the mass peak has approximately the same size for both algorithms in the signal. This is a similar range as for the low p_T bin, though the mass window size itself is significantly smaller for higher p_T bosons. The conclusion from this section is that there is a small preference for smaller radii at high p_T , but there is not nearly as much sensitivity as for lower p_T for the values of r and R . At low p_T , it is non-trivial

to pick an appropriate R (which can be compensated to some extent by varying r). From a practical point of view, there is a large incentive for using a smaller radius at high p_T because the jet mass from two constituents is mostly due to the p_T of those constituents, which is much better constrained than the mass of a single jet. The response of small radius jet mass is revisited in the context of close-by jets in Sec. 4.1.3.6 and for isolated high p_T jets in Sec. 4.1.4.




Figure 4.43: Left: Window efficiency for various r values. Right: Size of the smallest 68% window. Both plots compare large radius trimmed jets with $R = 1.0$, $f_{\text{cut}} = 0.05$ and $R_{\text{sub}} = 0.05$ with re-clustered jets that have analogous parameters.




Figure 4.44: The small-radius jet constituent multiplicity for several choices of f_{cut} and a fixed $r = 0.2$ (left) and $r = 0.4$ (right).




Figure 4.45: The size of the smallest 68% window as a function of f_{cut} for $r = 0.2$ (left) and $r = 0.4$ (right).

4.1.3.5 Re-clustered Jet Resolution

This section augments the signal and background comparisons from the previous section with a study of the four-vector resolution of signal jets. As they are both built directly or indirectly from calorimeter-cell clusters, one may expect the resolutions of re-clustered jets and large-radius trimmed jets to be similar. Figure 4.46 shows both the jet p_T and mass response. For both algorithms, $R = 1.0$ and $f_{\text{cut}} = 0.05$. The large-radius trimmed jets use k_t $R_{\text{sub}} = 0.3$ subjets and the re-clustered jets use anti- k_t $R = 0.4$ small-radius jets¹⁹. This setting will now be default for the rest of the chapter. For both re-clustered and large-radius trimmed jets, the particle-level reference is defined by running the same algorithm over detector-stable particles. The core of the response distributions are nearly identical (even slightly better for re-clustered jets in the case of the mass), but there are heavy tails for re-clustered jets. This can be explained by cases in which the re-clustering procedure picks a different number of jets at detector-level and particle-level. For nearly symmetric W boson decays, this asymmetry naturally introduces responses as big as 2 or as small as $1/2$ for p_T and even smaller for mass ($m^{\text{small-radius}} \ll m_W$). One way to quantify this effect is to introduce a new jet collection called *re-clustered trimmed truth*, which is constructed by matching each detector-level constituent of a re-clustered jet with a small-radius particle-level jet and then replacing the detector-level jet four-vector by the particle-level jet properties. When the detector-level and particle-level algorithms choose the same jets, the response of re-clustered trimmed truth jets is identically zero. The heavy tails of the red dashed lines in Fig. 4.46 show how the tails are explained by this mis-match in definition at particle-level and detector-level. One way to remove the impact of the mis-match is to use an algorithm-independent reference object. Figure 4.47 is the analogue to Fig. 4.46, but using the W boson four-vector as a reference for both re-clustered and large-radius trimmed jets. The tails as well as the core of both distributions are similar. One slight disadvantage of this method is that the resolution is now convolved with a non-negligible particle-level resolution from fragmentation that can hide differences between the algorithms in the tails. A

¹⁹Smaller radius jets are not (yet) calibrated and understood at the same level as $R = 0.4$ jets; the previous section does suggest that smaller would be better for the future.

direct way of comparing the resolutions of the two methods is to directly compare the four-vectors jet-by-jet. The ratio of the re-clustered $\mathbf{p}_T$ to the large-radius trimmed jet $\mathbf{p}_T$ is plotted as a function of boson $\mathbf{p}_T$ in the left plot of Fig. 4.48. For $\mathbf{p}_T \gtrsim 200$ GeV where one $R \sim 1$ jet is expected to capture most of the boson decay products, the ratio is strongly peaked at one. The two jets share most of the same regions of the calorimeter and so the response fluctuations shown in the middle and right plots of Fig. 4.48 are highly correlated²⁰. One potential difference in the calibration between re-clustered jets and large-radius trimmed jets is the treatment of close-by hadronic activity. Large-radius jets are calibrated as one unit, integrating over the distribution of energy inside the jet. However, the calibration of re-clustered jets is factorized, first calibrating small-radius jets. Figure 4.49 shows that the simulation does not predict a significant ΔR dependence of the re-clustered jet $\mathbf{p}_T$ scale. The next section describes a measurement of close-by effects on the jet mass using the $\sqrt{s} = 8$ TeV data.

²⁰In fact, one could use one of the collections to calibrate or establish uncertainties for the other.




Figure 4.46: The p_T (left) and mass (right) response of large radius trimmed jets and re-clustered jets for $200 \text{ GeV} < p_T^V < 400 \text{ GeV}$. In this p_T bin, over 80% of the re-clustered jets have at least two constituents.




Figure 4.47: The p_T response, with the p_T^V in the denominator instead of the truth jet p_T , of large radius trimmed jets and re-clustered jets for various bins of boson p_T .




Figure 4.48: Left: the ratio of the re-clustered jet p_T to the large radius trimmed jet p_T as a function of the boson p_T . Middle (right): the joint distribution of the large-radius jet response and the re-clustered jet mass response in a low (high) p_T^V bin.




Figure 4.49: The p_T response as a function of the ΔR between the re-clustered jet constituents in events where the re-clustered jet has exactly two constituents and $200 \text{ GeV} < p_T^{\text{re-clustered jet}} < 400 \text{ GeV}$.

4.1.3.6 Close-by Effects

When two jets are physically close in $\Delta\mathbf{R}$, their four-vector response can be different with respect to isolated jets. For instance, the energy deposits from particles in one particle-level jet may end up clustered into a different reconstructed jet. In these cases, the energy response will be high for one jet and low for the other. Even at particle-level, the presence of close-by jets can skew jet shapes due to the properties of the jet clustering algorithm, or changes in the physical distribution of energy due to color flow between jets. The jet response depends on these jet shapes and so this is another source of bias in the response of non-isolated jets. Non-isolated jets are common in high multiplicity final states but are always present in re-clustered jets with more than one small-radius jet constituent. Jet calibrations and their corresponding systematic uncertainties are derived in simulation from jets that are well-isolated. These calibrations are applied to all jets, regardless of other close-by hadronic activity. Studies at $\sqrt{s} = 7$ TeV suggested that 2%-5% shifts in the jet energy scale are possible due to the presence of close-by jets [289]. However, detailed studies with the larger $\sqrt{s} = 8$ TeV dataset demonstrated that these shifts are well-modeled by the simulation and therefore no additional uncertainty is applied to the jet $\mathbf{p}_T$ to account for close-by hadronic activity [122]. The right plot of Fig 4.2 shows that it is not sufficient for the jet $\mathbf{p}_T$ to be well-modeled - there are important kinematic regimes where small-radius jets with a significant mass are in close proximity to other jets. For example, this occurs in boosted top quark jets with the W boson decay products are merged inside one small-radius jet distinct from a close-by b -quark jet. Close-by effects on the jet mass response have never studied, but are critical for jet tagging in dense environments. This section presents the first such measurement using an extension of the track-jet method (see Sec. 4.1.2.1) called the *triple-ratio technique*, described in Sec. 4.1.3.6.1. Large-radius jets in $t\bar{t}$ events are used to study the jet-area dependence of close-by shifts in the jet mass response because there are not enough small-radius jets with a significant mass and close-by activity. These results are presented in Sec. 4.1.3.6.2.

4.1.3.6.1 Triple Ratio Technique

The first step in studying the dependence of the jet mass response on close-by jet activity in the data is to quantify the level of nearby radiation. One possibility is the quantity f_{closeby} , defined as

$$f_{\text{closeby}} = \sum_j \frac{\vec{p}_j \cdot \vec{p}}{|\vec{p}|^2}, \quad (4.31)$$

where the sum runs over all jets above a p_T threshold inside a cone of radius $\Delta R < X$ with respect to the probe jet. Less isolated jets have a larger value of f_{closeby} . There are a few ways to naturally extend the definition of the small-radius jet f_{closeby} in Eq. 4.31 to large radius jets. The most obvious definition is to take all small-radius jets that are in some annulus of the jet axis, as is done with small radius jets. A disadvantage of this definition is that the trimmed jet area can be rather small so that small-radius jets with $\Delta R \sim 1$ can already be quite far away from the jet center. An alternative definition uses calorimeter-cell clusters instead of jets in Eq. 4.31. The list of clusters could be inclusive or exclude those clusters dropped in the trimming process. Figure 4.50 compares three definitions of f_{closeby} for large-radius jets. The definition based on small-radius jet (clusters) uses $1 < \Delta R < 2$ ($\Delta R < 2$). By construction, the clusters removed from trimming are relatively soft, so the cluster-based definitions are highly correlated. If there are no close-by clusters, there will not be close-by jets, but there can also be no close-by jets above 25 GeV but many close-by clusters. For the rest of this section, the cluster-based definition, excluding trimmed clusters, is used for making comparisons with data.

In order to probe the impact of close-by effects on $R = 0.4$ jets, the jet area dependence is studied in a sample of large-radius trimmed jets with a range of sizes. Figure 4.51 shows the jet m/p_T and jet area regions considered in this analysis in addition to the distribution of the cluster f_{closeby} . The f_{closeby} distribution does not depend strongly on the jet area, but does decrease with p_T due to the close-by b-jet.




Figure 4.50: The joint distributions of f_{closeby} computed with nearby ($1 < \Delta R < 2$) small-radius jets or all non-jet clusters within $\Delta R < 2$ (left) and two cluster-based definition that either include or exclude clusters dropped during trimming (right).




Figure 4.51: Left: the jet m/p_T and trimmed jet area regions used in this analysis. Right: the distribution of the cluster f_{closeby} in all the nine bins of jet m/p_T and trimmed jet area, evenly spaced between 3×3 bins in the range 0.2 and 0.5 in m/p_T and in the range 0.2π and 0.5π in the trimmed jet area.

The impact of close-by jets is quantified by comparing the response (using $r_{\text{track}} = m_{\text{calo}}/m_{\text{tracks}}$) of isolated jets with non-isolated jets: $R = r_{\text{track}}^{\text{non-iso}}/r_{\text{track}}^{\text{iso}}$. A third ratio (r_{track} is itself a ratio) is formed to compare data and simulation: $R^{\text{data}}/R^{\text{MC}}$.

4.1.3.6.2 In-situ Close-by Results for the Jet Mass

The ratio the median²¹ $r_{\text{track}} = m_{\text{from calo}}/m_{\text{from tracks}}$ as a function of the jet m/p_T and the trimmed jet area between low $f_{\text{closeby}} < 0.6$ and $f_{\text{closeby}} > 0.6$ in data and MC are presented in Fig. 4.52. The value $f_{\text{closeby}} = 0.6$ is chosen to separate isolated and non-isolated jets because it is approximately the median of the f_{closeby} distribution. A small negative trend in m/p_T is predicted by the simulation in the right plot of Fig. 4.52, but this is not in the data distribution. Instead, there is a small trend in the opposite direction. This is quantified by the triple ratio in Fig. 4.53 that is the ratio of the left and right plots from Fig. 4.52. Except for low m/p_T and large jet area, the triple ratio is statistical consistent with unity, suggesting that no additional uncertainty is required for re-clustered jets due to the modeling of close-by jets. There may be a $\lesssim 2\% - 5\%$ bias for large jet areas, but this is not the focus of this section. Figure 4.52 does not include any systematic uncertainties, but many of the experimental sources of bias cancel in one of the three ratios. For example, Fig. 4.54 shows the impact of varying the JMS up and down within its uncertainty. The resulting change in the triple ratio is $\lesssim 1\%$ in all bins.




Figure 4.52: The ratio of r_{track} between isolated (lose f_{closeby}) and non-isolated jets (high f_{closeby}) as a function of the jet area and m/p_T in data (left) and simulation (right). The smaller number in the data plot is the statistical uncertainty.

²¹The median is less sensitive to outliers than the mean.




Figure 4.53: The ratio of the left and right plots from Fig. 4.52. The smaller number in each bin indicates the data statistical uncertainty.




Figure 4.54: The same as Fig. 4.53, but with the JMS varied down (left) or up (right).

4.1.3.7 Summary and Outlook

Re-clustering is a modular paradigm for large-radius jet clustering that introduces analysis flexibility and a natural scheme for estimating systematic uncertainties. This flexibility can increase the discovery potential of the LHC as the large-radius jet parameters can be individually optimized for each analysis. One last appealing property of re-clustering is that it provides a continuous bridge between the low and high p_T regimes. At low p_T , re-clustered jets tend to have several constituents and most of the mass information is from the p_T of the small-radius jets. At high p_T , a re-clustered jet is identical to a single small-radius jet (see Fig. 4.55). Calibrations (and uncertainties) of small-radius jet mass is not well-constrained, but this is also true for traditional ‘large-radius jets’ at high p_T with small groomed area. The only difference is that the division between the two cases is made explicit in the re-clustering paradigm. Where re-clustering reduces to a single jet, one can use the *track-assisted* jet mass introduced in the next section to obtain calibrations and uncertainties.




Figure 4.55: The distribution of the ratio of the leading large-radius jet mass to small-radius jet mass as a function of W boson p_T .

4.1.4 Track-assisted Jet Mass

The main challenge of re-clustering is at high p_T where the re-clustered large-radius jets have only one small radius jet constituent. The mass of small-radius jets is less constrained than large-radius jets from the data due to the lack of ultra high p_T W boson and top quark jets. Additionally, the mass resolution for both small- and large-radius jets degrades at high p_T , as discussed in Sec. 4.1.2, as the distance between particles approaches the cluster angular resolution and ultimately the detector granularity. One strategy to mitigate this degradation in the resolution at high p_T is to use information from charged particle tracks *as part of the jet mass reconstruction*. The track momentum resolution also degrades with p_T , but the angular resolution is significantly superior to the calorimeter angular resolution. Track and calorimeter information are already combined as part of particle-flow techniques in CMS [366], but these procedures suffer at high p_T from the inability to accurately match tracks and clusters, especially in ATLAS where the magnetic field is weaker than in CMS by a factor of two. The idea in this section is to use track-based properties of (sub)jets without attempting to match tracks with individual calorimeter clusters.

The left plot Fig. 4.56 shows the average number of tracks and clusters inside an $R = 0.4$ boosted boson jet as a function of the jet p_T . The particle multiplicity for a W boson jet should be nearly independent of p_T because the starting scale for the parton shower is set by the quark p_T in the W boson center-of-mass frame, which is independent of the boost. The number of tracks is nearly constant up to $p_T \sim 2$ TeV and then drops by about 1 track over the next 1 TeV. In contrast, the number of calorimeter clusters decreases significantly with jet p_T as the particles become more collimated with increasing boost. This is illustrated with an event display of a particular high p_T W boson jet in the right plot of Fig. 4.56. The $p_T \sim 3.5$ TeV W boson jet has a particle level mass of about 80 GeV and a calorimeter mass of about 150 GeV. This large mass is due in part to the very soft radiation at the periphery of the jet (possibly from pileup), while the true mass is mostly contained within the core $\Delta R \sim 2m_W/p_T \sim 0.05$. There are only eight reconstructed calorimeter clusters while there are 20 reconstructed tracks.

While tracks are measured precisely, the jet mass built only from tracks is not




Figure 4.56: Left: the average multiplicity of tracks and clusters in $R = 0.4$ boosted W and Z boson jets as a function of the particle-level jet p_T . Right: an event display of the clusters inside one particular high p_T boosted boson jet. The TA mass is the track-assisted jet mass (see text for details).

directly useful. There are significant charged-to-neutral fluctuations that induce a resolution with respect to the particle-level jet mass constructed from all constituents. This resolution is significantly larger than the calorimeter mass resolution. Figure 4.57 shows the calorimeter-only and track-only jet mass distributions. The uncalibrated track-only mass has a much lower average value than the calorimeter jet mass due to the missed neutral energy. When accounting for this average, the width of the track-mass is substantially broader than the width of the calorimeter mass distribution.

A minimal, but powerful way to improve the track-only mass is to apply a jet-by-jet correction for the charged-to-neutral ratio to form the *track-assisted jet mass*:

$$m_{\text{track-assisted}} = m_{\text{track}} \times \frac{p_{T,\text{calo}}}{p_{T,\text{track}}}, \quad (4.32)$$

where $p_{T,\text{track}}$ is the 4-vector sum of the tracks associated to a (trimmed) jet. The track-assisted jet mass for the example in Fig. 4.56 is closer to the particle-level jet mass, in part because of the larger number of tracks, improved angular resolution, and reduced sensitivity to pileup. The reduced sensitivity to the charged-to-neutral




Figure 4.57: The jet mass distribution of reconstructed boosted hadronically decaying W and Z boson jets with $p_T > 1.5 \text{ TeV}$. A jet collection is calibrated if the average value of $m_{\text{reco jet}}/m_{\text{truth jet}}$ is unity.

fluctuations compared with the track-only mass leads to a resonance peak in Fig. 4.57 that has a sharper maximum near the boson mass, comparable to the calorimeter jet mass.

A procedure for correcting the jet mass as in Eq. 4.32 was first proposed using hadronic calorimetry to correct electromagnetic-only measurements [367, 368]. The extension to charged particle tracks was introduced in the context of top-quark jet tagging [369] using the HEPTopTagger algorithm [370, 371]. Since that time, there have been phenomenological studies using track-assisted jet mass²² for ultra boosted ($p_T \gtrsim \mathcal{O}(10) \text{ TeV}$) top quark and boson jets [372, 373]. This remainder of this section is the first experimental study of the track-assisted jet mass, including a discussion of calibrations and systematic uncertainties.

Figure 4.58 shows the jet mass distribution for boosted W and Z boson jets clustered using the $R = 0.4$ anti- k_t algorithm. For jets with $600 < p_T < 800 \text{ GeV}$, the calorimeter jet mass peak is sharper than for the track-assisted jet mass; the

²²The phenomenological studies have not given a name to the quantity to Eq. 4.32, so it is defined here as the track-assisted jet mass.

inter-quantile range divided by the median (the quantile analogue to the coefficient of variation) is about 40% larger for the track-assisted jet mass. However, at higher momenta, $p_T > 2$ TeV, the peaks have nearly the same resolution. Figure 4.59 quantifies the p_T dependence of the inter-quantile range for the jet mass response distribution. The truth jet mass in the definition of the response is the same for both the track-assisted and calorimeter jet mass. The calorimeter jet mass response distribution is broader for the track-assisted jet mass for $p_T \lesssim 1.7$ TeV after which the track-assisted jet mass resolution is smaller than the calorimeter jet mass resolution. These differences are not confined to the core of the response distribution as the trends are quantitatively the same for both the $\pm 10\%$, $\pm 20\%$, and $\pm 30\%$ quantiles centered around the median.




Figure 4.58: The track-assisted and calorimeter jet mass distribution for boosted W and Z boson jets clustered using the $R = 0.4$ anti- k_t algorithm for $600 < p_T < 800$ GeV (left) and $p_T > 2$ TeV (right).

In addition to improving the jet mass resolution at high jet p_T , there are many experimental benefits to the track-assisted jet mass compared with the traditional calorimeter jet mass. First of all, if the jet p_T in Eq. 4.32 is already calibrated, the track-assisted jet mass is also nearly calibrated without any extra effort. This is observed in Fig. 4.58: the masses are constructed with uncalibrated EM-scale jet mass, but EM+JES p_T ; the track-assisted jet mass peak is around m_W while the average




Figure 4.59: The p_T dependence of the inter-quantile range divided by the median of the track-assisted and calorimeter jet mass distributions. The 10%, 20%, and 30% inter-quantile ranges are centered around the median.

calorimeter jet mass is much larger. Additionally, if the jet p_T is corrected for pileup, the track-assisted jet mass is also corrected for pileup, since pileup tracks can be removed from m_{track} . Most importantly, the in situ momentum balancing techniques used to calibrate and estimate systematic uncertainties for the calorimeter jet p_T apply directly to the track-assisted jet mass. In particular, the jet mass scale and resolution uncertainty are the convolution of calorimeter p_T and tracking uncertainties. Unlike the uncertainty on the calorimeter jet mass, these two components can be well-estimated to the highest accessible jet momenta. This is true for both small and large radius jets. For example, Fig. 4.60 shows how the small radius jet energy resolution uncertainty translates into an uncertainty on the track-assisted jet mass resolution. By construction, the resolution uncertainty on the response is nearly the same for the p_T and the track-assisted jet mass. For small radius jets, the resolution is measured precisely and so the uncertainty is $\mathcal{O}(1\%)$. For comparison, the difference between PYTHIA and HERWIG++ is the same order of magnitude. This last quantity gives a sense of the fragmentation uncertainty, entering through the tracking

component of the track-assisted jet mass.




Figure 4.60: The calorimeter jet momentum resolution uncertainty (circle), the jet p_T resolution-induced track-assisted jet mass resolution uncertainty (up triangle), and the difference in the track-assisted jet mass resolution when comparing PYTHIA and HERWIG++. The resolution uncertainty is defined as $\sqrt{v^2 - n^2}$, where n is the nominal response resolution and v is the response resolution from the alternative sample. The alternative sample for the first two cases is created by smearing the jet energy within the resolution uncertainty and the alternative sample for the last case is the HERWIG++ simulation.

Since calorimeter jet mass is not used in the construction of the track-assisted jet mass, one may consider combining the two mass definitions to achieve even better performance. Even though fluctuations in the calorimeter induce correlations between the jet mass and jet p_T response, the left plot of Fig. 4.61 shows that the correlation between the track-assisted jet mass response and the calorimeter jet mass response is negligible. It is therefore a good approximation to treat the track-assisted jet mass X and the calorimeter jet mass Y as independent when determining the optimal combination. For a fixed truth mass, $X \sim \mathcal{N}(1, \sigma_1)$ and $Y \sim \mathcal{N}(1, \sigma_2)$. In this application, X is the calibrated track-assisted jet mass response and Y is the calibrated

calorimeter jet mass response. Let²³ $Z = \alpha X + \beta Y$. Assuming X and Y are calibrated, Z is also calibrated if $\alpha + \beta = 1$. The variance of Z is $\sigma^2(Z) \approx \alpha^2 \sigma^2(X) + \beta^2 \sigma^2(Y)$. The first order conditions $\partial_\alpha \sigma^2(Z) = 0$ and $\partial_\beta \sigma^2(Z) = 0$ result in the minimum variance unbiased estimator of the particle-level mass $\hat{Z}$: $\alpha \propto 1/\sigma_1^2, \beta \propto 1/\sigma_2^2$. The performance of the optimal combination of the calorimeter and track-assisted jet mass is shown in Fig. 4.61. For $p_T \lesssim 1$ TeV, the calorimeter jet mass resolution is significantly better than that of the track-assisted jet mass and so the improvement from the combination is negligible. However, for $p_T \gtrsim 1$ TeV, there is a full $\approx 40\%$ ($\approx \sqrt{2} - 1$) improvement in the resolution. Distributions of the track-assisted jet mass, the calorimeter jet mass, and the optimal combination are shown in a few jet p_T bins in Fig. 4.62.




Figure 4.61: Left: the joint distribution of the track-assisted jet mass response and the calorimeter jet mass response. Right: the p_T dependence of the normalized jet mass resolution for the track-assisted jet mass, the calorimeter jet mass, and the optimal combination of the two mass definitions.

The track-assisted jet mass is a promising technique for jet mass at high p_T . At moderate and low p_T , there may be modifications of the algorithm to recover performance. In particular, since the charge-to-neutral ratio has large local fluctuations, subjet corrections may be able to improve the resolution at lower p_T when the

²³The linear combination is optimal over a wider class of functions, but this is beyond the scope of this section.




Figure 4.62: Distributions of the track-assisted jet mass, the calorimeter jet mass, and the optimal combination are shown in a few jet p_T bins.

calorimeter can reliably resolve a jet's substructure [369].

4.1.5 Conclusions and Future Outlook

Section 4.1 has presented several new techniques for improving jet four-vector reconstruction and measuring the performance in data. At low p_T , re-clustering offers a powerful and flexible scheme for tailoring jet clustering that naturally reduces to traditional small-radius jets at high p_T . The mass of single small-radius jets is not well-constrained due to the lack of a pure sample of massive small radius jets in the data, but one can use the track-assisted mass at high p_T to improve the performance and retain a natural scheme for calibrations and uncertainties. As more data at higher p_T is collected in Run 2, the in-situ resonance method can be used as an independent measurement of the uncertainties for both re-clustering and track-assistance across a wide range of jet p_T . This toolkit will hopefully be able to improve the sensitivity to boosted bosons and top quarks for a wide range of searches in Run 2 and beyond. Part III will already show that re-clustering is a useful tool for extending the sensitivity to electroweak scale SUSY using the early $\sqrt{s} = 13$ TeV data.

The next section presents a new application of the traditional large-radius trimmed mass to distinguish different types of boson jets.

4.2 Boson Type Tagger

4.2.1 Introduction

Jet substructure techniques²⁴ developed to distinguish hadronically decaying W and Z bosons from QCD multijet background processes have become increasingly sophisticated. A recent review is given in Ref. [293]. Both ATLAS [375] and CMS [166] have performed detailed comparisons of the various tagging variables and jet-grooming techniques with the overall conclusion that large QCD multijet suppression factors²⁵ are possible while maintaining acceptable levels of boson tagging efficiency. Given a W/Z -boson tagger, a natural next step is to distinguish between boson types, e.g. W -boson jets from Z -boson jets.

There are several important possible applications of a boson-type tagger at the LHC. First, a type tagger could enhance the SM physics program with W and Z bosons in the final state. Measurements of this kind include the determination of the cross sections for V +jets, WV , and $t\bar{t} + V$. Another important use of a boson-type tagger is in searches for flavor-changing neutral currents (FCNC). Due to the Glashow–Iliopoulos–Maiani (GIM) mechanism [376], FCNC processes in the SM are highly suppressed. Many models of new physics predict large enhancements to such processes. Both ATLAS and CMS have performed searches for FCNC [377, 378] of the form $t \rightarrow Zq$ in the leptonic channels, but these could be extended by utilizing the hadronic Z decays as well. FCNC process mediated by a leptophobic Z' may be detected only via hadronic type-tagging methods. A third use of a boson-type tagger is to categorize the properties of new physics, if discovered at the LHC. For instance, if a new boson were discovered as a hadronic resonance, a boson-type tagger could potentially distinguish a $W'(\rightarrow qq)$ from a $Z'(\rightarrow qq)$ (where mass alone may not be useful). This is especially relevant for leptophobic new bosons, which could not be distinguished using leptonic decays.

Labelling jets as originating from a W or Z boson is less ambiguous than quark/gluon labelling. A W boson can radiate a Z boson, just like a quark can radiate a gluon,

²⁴The tool presented in this section has been published in Ref. [374].

²⁵ $\mathcal{O}(1\%)$ QCD multijet efficiency at 50% signal efficiency.

but this is heavily suppressed for the former and not for the latter. The radiation pattern of jets from W - and Z -bosons is less topology dependent because it is largely independent of the other radiation in the event as W and Z bosons are color singlets. Aside from the production cross section and subtle differences in differential decay distributions, the only features that distinguish between W and Z bosons are their mass, charge, and branching ratios. Experimentally, this means that the only variables that are useful in discriminating between hadronic decays of W and Z bosons are those which are sensitive to these properties. The three variables used in the analysis presented here are *jet mass*, sensitive to the boson mass, *jet charge*, sensitive to the boson charge, and a *b-tagging* discriminant which is sensitive to the heavy-flavor decay branching fractions of the bosons. The application of a boson-type tagger in practice will be accompanied by the prior use of a boson tagger (to reject QCD multijet processes). The type-tagger variables are largely independent of typical boson-tagger discriminants like n -subjettiness [379], which rely on the two-prong hard structure of both the W and Z decays.

Due to the large QCD backgrounds with experimental signatures similar to hadronic electroweak boson production, isolating W and Z bosons at the LHC is challenging. However, at lepton machines, electroweak boson production is often dominant and can be a background for many other interesting processes (see e.g. Ref. [380]). Jet tagging in boosted topologies has matured considerably since LEP and so some of the techniques presented here may be applicable to future high energy lepton machines.

This section introduces a new jet tagging method to distinguish between hadronically decay W and Z bosons at the LHC, and documents its performance with the ATLAS detector at $\sqrt{s} = 8$ TeV. The section is organized as follows. Section 4.2.2 describes the simulated datasets used in constructing and evaluating the boson-type tagger. Following a discussion of the differences between the properties of W and Z bosons in Sec. 4.2.3, Sec. 4.2.4 defines the three discriminating variables. The construction and performance of the tagger are detailed in Sec. 4.2.5 and the sensitivity to systematic uncertainties is described in Sec. 4.2.6. The input variables are studied in a dataset enriched in boosted W bosons in Sec. 4.2.7. The section ends with a discussion of possible uses of the tagger in Sec. 4.2.8 and conclusions in Sec. 4.2.9.

4.2.2 Datasets

Two sets of Monte Carlo (MC) simulations are generated, one to study the tagger's W versus Z performance and the other to compare the tagger inputs for W bosons with the data. For the tagger performance, it is useful to have a source of isolated high p_T W and Z bosons. One physics process that produces such final states is the production of a hypothetical W' boson. Predicted in models of new physics with an $SU(2)$ gauge group, the W' is analogous to the SM W boson. For this analysis, a 100% branching ratio $W' \rightarrow WZ$ is used to generate events with simultaneously boosted W and Z bosons; the p_T of the SM bosons is set by the mass of the W' boson. Unfortunately, there is no evidence in real data for W' bosons and it is not possible to measure the tagger performance directly in the data due to the lack of a pure sample of boosted, hadronically decaying Z bosons. However, the modelling of the tagger inputs can be studied using hadronically decaying W bosons from $t\bar{t}$ events in the data. The simulation and event selection used for the modelling studies are identical to those from Sec. 4.1.2.

A simulated sample of W' bosons is generated with PYTHIA 8 using the leading-order parton distribution function set (PDF) MSTW2008 [193, 194] and the AU2 [190] set of tunable parameters (tune) for the underlying event. The baseline samples use PYTHIA for the $2 \rightarrow 2$ matrix element calculation, as well as p_T -ordered parton showers [98] and the Lund string model [101] for hadronization. Additional samples are produced with HERWIG++ [195], which uses angular ordering of the parton showers [99], a cluster model for hadronization [102], as well as the EE3 [197] underlying-event tune. In order to remove artifacts in the p_T distributions of the W and Z bosons due to the generation of W' particles with discrete masses, the p_T^V spectra are re-weighted to be uniform in the range $200 \text{ GeV} < p_T^V < 400 \text{ GeV}$. As is discussed in Sec. 4.2.1, for $p_T > 200 \text{ GeV}$, a jet with large radius is expected to capture most of the W or Z boson decay products. The range is truncated to $p_T < 400 \text{ GeV}$ because hadronically decaying W bosons can be probed with data in this p_T range; there are too few events in the 8 TeV dataset for $p_T > 400 \text{ GeV}$. Figure 4.63 shows the boson p_T spectrum before any re-weighting. The shape of the distribution is set by the available range of W' masses that spans a few hundreds of GeV to 4 TeV. Spikes are due to kinematic

jacobian peaks from individual W' masses. Since the W and Z mass difference is small compared to the W' masses, the shapes of the W and Z boson p_T spectrum are nearly identical.

The W' events are processed with a full simulation of the ATLAS detector [110] based on the **Geant4** [108] toolkit, and reconstructed using the same software as for the experimental data. The average number of additional pp collisions per bunch crossing (pileup interactions) was 20.7 over the full 2012 run. The effects of pileup are modelled by adding multiple minimum-bias events, which are simulated with Pythia 8.160 [189], to the generated hard-scatter events. The distribution of the number of interactions is then weighted to reflect the pileup distribution in the 2012 data.




Figure 4.63: The p_T spectrum of the simulated W and Z bosons from $W' \rightarrow WZ$ decays before applying any p_T re-weighting.

4.2.3 Distinguishing a Z boson from a W boson

Decays of W or Z bosons are characterized by the boson's mass and coupling to fermions. The mass difference between the W and Z boson is about 10 GeV and if produced from a hard scatter or the decay of a heavy enough resonance, both bosons are produced nearly on-shell since the width $\Gamma_V = 2.1$ (2.5) GeV is much less than the mass $m_V = 80.4$ (91.2) GeV for W (Z) bosons [27]. The Breit–Wigner resonance curves for W and Z bosons are shown in Fig. 4.64(a). The separation between the curves is a theoretical limit on how well mass-sensitive variables can distinguish between W and Z bosons. For hadronic boson decays, the mass peaks measured with jets are broader. This is because the jet-clustering algorithm for final-state hadrons loses particles at large angles to the jet axis and includes extra particles from the underlying event and pileup.

The generic coupling of a boson V to fermions is given by $g_V \gamma_\mu [c_V - c_A \gamma_5]$, where g_V is a boson-dependent overall coupling strength, and c_V and c_A are the vector and axial-vector couplings, respectively. The W boson couples only to left-handed fermions so $c_V = c_A = 1$ with $g_W \propto k N_C G_F m_W^3 |V_{ij}|^2$, where G_F is the Fermi coupling constant, V_{ij} is a Cabibbo–Kobayashi–Maskawa (CKM) matrix element [26, 381], k represents higher-order corrections, and $N_C = 3$ for the three colours of quarks and $N_C = 1$ for leptons. The CKM matrix is nearly diagonal so $W^+ \rightarrow u\bar{d}$ and $W^+ \rightarrow c\bar{s}$ are the dominant decay modes. Small off-diagonal elements contribute to the other possible decay modes, and the overall branching ratios assuming hadronic decay are approximately 50% for $W \rightarrow cX$ and 50% for $W \rightarrow \text{light-quark pairs}$. The W boson has electric charge ± 1 in units of the electron charge, so by conservation of charge, its decay products have the same net charge. The scalar sum of the charge of all the final-state hadrons originating from a W boson decay is not infrared safe (directly sensitive to the non-zero detection threshold), so there are limits to the performance of charge tagging dictated by the energy threshold placed on charged particles in the event reconstruction.

In contrast to W boson decays, Z bosons decay to both the left- and right-handed fermions. The partial width for $Z \rightarrow f\bar{f}$ is proportional to $k N_C G_F m_Z^3 [c_V^2 + c_A^2]$. The factors c_V and c_A are slightly different for up- and down-type fermions. The $b\bar{b}$

branching ratio is 22%, the $c\bar{c}$ branching ratio is 17% and the sum of the remaining branching ratios is 61%, assuming a hadronic decay. W boson decays to b -quarks are highly suppressed by the small CKM matrix elements V_{cb} and V_{ub} , so that identifying b -hadron decays associated with a hadronically decaying boson is a powerful discriminating tool. Branching ratios are plotted in Fig. 4.64d for Z decays to light quarks, c -quarks, and b -quarks, and in Fig. 4.64(c) for the W boson decays to light quarks and c -quarks.

Since the coupling structure is not identical for W and Z bosons, the total decay rates differ, and the angular distributions of the decay products also differ slightly. However, even at parton level without any combinatoric noise, the differences in the angular distributions are subtle. This is illustrated in Fig. 4.64(b) for transversely polarized W and Z bosons (details can be found in Appendix B.1). The angular distributions are identical for the two bosons for longitudinal polarization because the distributions for right- and left-handed fermions is the same. The relative contribution of left- and right-handed components for the Z decays depends on the quark flavor; for up-type quarks the relative contribution from right-handed fermions is 15% while it is only 3% for down-type quarks. In $t\bar{t}$ decays, the fraction of longitudinally polarized W bosons (ignoring the b -quark mass) is $m_t^2/(m_t^2 + 2m_W^2) \sim 0.7$. In contrast, the boson is mostly transversely polarized in inclusive V +jets events. See Appendix B.2 for a derivation of these polarization properties. Any discrimination shown in Fig. 4.64(b) is diluted by combinatorics (including distinguishing q from $\bar{q}$ jets), non-perturbative effects, and detector reconstruction, so angular distributions are not considered further in this section²⁶.

²⁶Boson polarization does have an impact on the jet mass distribution and thus on distinguishing boson jets from QCD jets [166]. However, the impact on distinguishing W jets from Z jets is highly suppressed because there are only (small) differences when the bosons are transversely polarized. Polarizations would be important only if the W and Z were predominately produced with different polarizations, which does not happen in e.g. V +jets, $t\bar{t}$, or $W' \rightarrow WZ$ events.




Figure 4.64: (a) Breit–Wigner resonances for the W (red) and Z (blue) bosons, (b) angular distribution of the decay products of transversely polarized W/Z bosons with respect to the spin direction in the boson rest frame, (c) hadronic branching fractions of the W^+ boson, and (d) of the Z boson. In (c) and (d), *light* stands for decay modes not involving c and b quarks.

4.2.4 Definitions of reconstructed objects

Jets are formed from clusters using two different jet algorithms. Small-radius jets are built with the anti- k_t algorithm with jet radius parameter $R = 0.4$ and large-radius jets are formed using the anti- k_t algorithm with $R = 1.0$ and then trimmed using k_t $R = 0.3$ subjets with $f_{\text{cut}} = 0.05$. Since the W and Z boson masses differ by about 10 GeV, the jet mass can be used to discriminate between these two particles. The distributions of the boson masses and jet masses for hadronically decaying W and Z bosons are shown in Fig. 4.65. The particle-level (‘truth’) jet mass is constructed from stable particles in the MC simulation ($c\tau > 10$ mm), excluding neutrinos and muons, clustered with the same jet algorithm as for calorimeter-cell clusters. The QCD multijet processes that govern the formation of stable particles from the W and Z decay products create a broad distribution of jet masses even without taking into account detector resolution. Constructing the jet mass from calorimeter-cell clusters further broadens the distribution. The jet-mass resolution (physical $\oplus$ detector) is large compared to the natural width of the W and Z bosons and comparable to the difference in their masses. For example, the standard deviation of $p_T^{\text{reco jet}}/p_T^{\text{truth jet}}$ is approximately 10%. The jet-mass variable nevertheless has some discriminating power. Figure 4.66 shows that as long as $p_T \gtrsim 200$ GeV, the jet mass distributions are relatively independent of p_T for isolated W and Z bosons.

The momentum and electric charge of particles traversing the detector contain information about the charge of their parent boson. To suppress the impact of pileup, tracks are required to originate from the primary collision vertex, which is defined as the vertex with the largest $\sum p_T^2$ computed from associated tracks. Additionally, tracks must satisfy a very loose quality criterion for the track fit χ^2 per degree of freedom, which must be less than three. Tracks are associated with jets using ghost association [120]. The charge of tracks associated with a jet is sensitive to the charge of the initiating parton. In order to minimize the fluctuations due to low- p_T particles, the *jet charge* is calculated using a p_T -weighting scheme (see Chapter 1 for details):




Figure 4.65: The boson mass at generator level (top left), the ‘truth’ jet mass at particle-level after parton fragmentation (top right), and the detector-level jet mass distributions (bottom) for W and Z boson jets separately (left) and for an even admixture of the two jet types (right). The parton-level plot has a different vertical scale than the other plots and also has no p_T requirement.

$$Q_J = \frac{1}{(p_{T,J})^\kappa} \sum_{i \in \mathbf{Tracks}} q_i \times (p_T^i)^\kappa, \quad (4.33)$$

where $\mathbf{Tracks}$ is the set of tracks with $p_T > 500 \text{ MeV}$ associated with jet J , q_i is the charge (in units of the electron charge) determined from the curvature of track i with




Figure 4.66: The p_T dependence of the jet mass distribution. The points are the mean value in a given boson p_T bin and the shaded region is the standard deviation for the Z boson distribution.

associated p_T^i , κ is a free parameter, and $p_{T,J}$ is the transverse momentum of the jet measured in the calorimeter. The calorimeter energy is used in the denominator to determine p_T instead of the sum of track momenta to account for the contribution from neutral particles. Dedicated studies have shown that $\kappa = 0.5$ is generally best for determining the charge of partons from the jets they produce (see Sec. 1.3). The distributions of the jet charge for jets initiated by W^+ , W^- and Z bosons are shown in Fig. 4.67. There is an observable separation between positive and negative W bosons, though the width of the jet charge distribution is larger than the separation of means. Figure 4.66 shows that the standard deviation is about twice as large as the separation between the jet charge means of W^+ boson jets and Z jets. As with Fig. 4.66, the jet charge distribution is relatively stable with $p_T \gtrsim 200$ GeV, though there is a small increase in the standard deviation due to the degradation in tracking performance at high p_T .

The expected charge composition of a W sample is process dependent. There are more W^+ than W^- bosons in inclusive W' production because of the initial charge asymmetry of quarks in the proton resulting in more W'^+ than W'^- . The

discrimination between Z bosons and a near even mixture of $W^\pm$ is greatly diminished with respect to e.g. Z versus W^+ . In that case charge sensitive variables are not very useful for the tagger and so all results are shown also without such variables. In a variety of physics processes, the charge of the hadronically decaying W boson is known from other information in the event. For example, in searches for FCNC effects in $t\bar{t}$ events with one leptonically decaying W boson, the charge of the lepton is opposite to the charge of the hadronically decaying W boson. Henceforth, only W^+ bosons are used for constructing the boson-type tagger; the results are the same for W^- bosons.




Figure 4.67: The jet charge distribution for jets originating from $W^\pm$ and Z bosons in simulated W' decays. Each distribution is normalized to unity. The parameter κ controls the p_T -weighting of the tracks in the jet charge sum.

The tracks from charged particles can be used further to identify the decays of certain heavy-flavor quarks inside jets due to the long b -hadron lifetime. This is useful for boson-type tagging because the Z boson couples to $b\bar{b}$ while decays of the W boson to b -quarks are highly suppressed and can be neglected. ATLAS has commissioned a b -tagging algorithm called MV1 (defined in Ref. [75]) which combines information about track impact-parameter significance with the explicit reconstruction of displaced b - and c -hadron decay vertices. The MV1 distribution is shown in Fig. 4.69 for the leading and sub-leading small-radius jets matched to the leading




Figure 4.68: The p_T dependence of the jet charge distribution. The points are the mean value in a given boson p_T bin and the shaded region is the standard deviation for the Z boson distribution.

large-radius jet. The boson-type tagger presented here uses multiple bins of the MV1 distribution simultaneously. Five bins of MV1 are defined by b -tag efficiencies (probability to tag a b -quark jet as such) of 0%–50%, 50%–60%, 60%–70%, 70%–80%, and 80%–100% as determined in simulated $t\bar{t}$ events. A lower b -tag efficiency leads to higher light-quark jet rejection. The five b -tagging efficiency bins are exclusive and MV1 is constructed as a likelihood with values mostly between zero and one (one means more like a b -jet). For example, a 100% b -tagging efficiency corresponds to a threshold of $MV1 > 0$ and an 80% b -tagging efficiency corresponds to a threshold value of $MV1 > z$ for $z \ll 1$. The 80%–100% b -tag efficiency bin then corresponds to jets with an MV1 value between 0 and z . Constructed in this way, the fraction of true b -jets inside an efficiency bin $x\%$ – $y\%$ should be $(y - x)\%$.

Small-radius jets are matched to a large-radius jet by geometric matching²⁷ ($\Delta R < 1.0$). Of all such small-radius jets, the two leading ones are considered. There are thus 30 possible bins of combined MV1 when considering the leading and sub-leading

²⁷In the definition of jets, R is the characteristic size in (y, ϕ) and the rapidity y is used in the jet clustering procedure, whereas geometrical matching between reconstructed objects is performed using $(\Delta R)^2 = (\Delta \phi)^2 + (\Delta \eta)^2$, where η is the pseudorapidity.

matched small-radius jet. The number of bins is 25 from the 5×5 efficiency-binned MV1 distributions in addition to five more for the case in which there is no second small-radius jet matched to the large-radius jet. The distribution for the efficiency-binned MV1 variable for the leading and sub-leading matched small-radius jets is shown for W and Z bosons in Fig. 4.70. The flavor of a small-radius jet is defined as the type of the highest energy parton from the parton shower record within $\Delta R < 0.4$. As expected, a clear factorization is seen in Fig. 4.70 – the MV1 value depends on the flavor of the jet and not the process that created it. This means that, for example, c -jets from W decays have the same MV1 distribution as c -jets from Z decays. Jets originating from b -hadron decays tend to have a larger value of MV1, which means they fall in a lower efficiency bin. Jets not originating from b - or c -decays are called light jets and are strongly peaked in the most efficient bin of MV1. There is always one small-radius jet matched to the large-radius jet, but about 20% of the time there is no sub-leading jet with $p_T > 25$ GeV matched to the large-radius jets. These jets are all predicted to originate from light-quark decays of the W and Z bosons.




Figure 4.69: The distribution of the MV1 discriminant for the leading (left) and subleading (right) small-radius jets matched to the leading large-radius jet. The spike at -1 in the right plot corresponds to cases in which there is not a second small-radius jet. The other features in the distribution correspond to transitions in the dominant input algorithms to MV1 [75].

Figures 4.71 illustrates why there is a slightly different fraction of W events that




Figure 4.70: The efficiency-binned MV1 distribution for small-radius jets associated with large-radius jets resulting from W and Z boson decays. The left (right) plot shows the leading (sub-leading) small-radius jet MV1 distribution. The bins correspond to exclusive regions of b -jet efficiency. As such, the bin content of the black line (b -tagging for b -jets) should be proportional to the size of the efficiency window: about 50% for the rightmost bin, 10% for the three middle bins and 20% for the second bin.

have no second matched small radius jet compared to Z events. The W and Z transverse momentum spectrum are identical, so the boson mass difference has no effect on the spectrum at the low end, i.e. the probability for the subleading jet to be below threshold is independent of the boson. However, since $m_W < m_Z$, the angular separation between the two boson decay products is slightly smaller for W bosons and thus at high p_T , the W daughter jets merge into a single small radius jet earlier than for Z jets.

The p_T -dependence of the matched b -jet multiplicity is shown in Fig. 4.72. As with the jet mass and jet charge, there is a clear turn on for $p_T \sim 200$ GeV. However, there is a second feature at $p_T \gtrsim 400$ GeV when the two small-radius jets begin to merge into a single small-radius b -tagged jet.




Figure 4.71: Left: The fraction of events with $\Delta R(q, q) < 0.4$ in a toy MC simulation as a function of the simulated boson mass. Right: The p_T spectrum of the softer of the two decay products for a simulated (scalar) boson with mass 1 TeV when one of the decay products is below 25 GeV or the two decay products are within $\Delta R < 0.4$.




Figure 4.72: The p_T dependence of the number of small radius b -tagged jets (70% working point) associated to the large radius jet. The points are the mean value in a given boson p_T bin and the shaded region is the standard deviation for the Z boson distribution.

4.2.5 Tagger performance

The optimal multivariate tagger combining jet mass, jet charge, and the MV1 of matched small-radius jets is constructed from a three-dimensional (3D) likelihood ratio. For N bins each of jet mass and jet charge, as well as 30 combined MV1 bins, the 3D likelihood ratio would have $30 \times N^2$ total bins. Populating all of these bins with sufficient MC events to produce templates for the likelihood ratio requires an unreasonable amount of computing resources, especially for the high-efficiency bins of combined MV1. Estimating the 3D likelihood as the product of the 1D marginal distributions, where all variables but the one under consideration are integrated out, is a poor approximation for jet mass and combined MV1 due to the correlation induced by the presence of semileptonic b -decays, which shift the jet mass to lower values due to the presence of unmeasured neutrinos²⁸. It is still possible to use a simple product by noting that all three tagger inputs are independent when the flavor of the decaying boson has been determined. Thus, for each possible boson decay channel, templates are built for the jet mass, the jet charge, and the efficiency-binned MV1 distributions. For a particular decay flavor, the joint distribution is then the product of the individual distributions. Summing over all hadronic decay channels then gives the full distribution. To ease notation, the efficiency-binned MV1 is denoted $B = (B_{\text{lead}}, B_{\text{sub-lead}})$. The distribution for B_{lead} ($B_{\text{sub-lead}}$) is shown in the left (right) plot in Fig. 4.70. Symbolically, for decay flavor channel $\mathcal{F}$, mass M , charge Q , and efficiency-binned MV1 B , the likelihood is given by:

$$p(M, Q, B|V) = \sum_{\mathcal{F}} \text{Pr}(\mathcal{F}|V) p(M|\mathcal{F}, V) p(Q|\mathcal{F}, V) \text{Pr}(B|\mathcal{F}, V), \quad (4.34)$$

where²⁹ $V \in \{W, Z\}$ and the sum is over $\mathcal{F} = bb, cc, cs, cd$ and light-quark pairs.

²⁸The muons from semileptonic decays are added back to the jet using a four-momentum sum. Adding back the muon has a negligible impact on the inclusive mass distribution due to the semileptonic branching ratios and lepton identification requirements. For details about the muon reconstruction and selection, see Sec. 4.2.7 (the only difference here is that the isolation is not applied). Figure 4.73 shows the impact of this muon correction on the jet mass.

²⁹The symbol p denotes a probability density whereas Pr denotes a discrete probability distribution.

The distribution of B is well approximated as the product of the distributions for B_{lead} and $B_{\text{sub-lead}}$ when the flavors of the leading and sub-leading jets are known. This is exploited for hadronically decaying W bosons and for the light-quark flavor decays of Z bosons to construct templates for B that have a sufficient number of simulated events for large values of B , i.e. $\Pr(B|\mathcal{F}, V) = \Pr(B_{\text{lead}}|\mathcal{F}, V) \Pr(B_{\text{sub-lead}}|\mathcal{F}, V)$. Figure 4.74 shows that factorization holds within the statistical uncertainty of the simulation. The unit-normalized templates for B are shown in Fig. 4.70 and the unit-normalized templates $p(M|\mathcal{F}, V)$ and $p(Q|\mathcal{F}, V)$ are shown in Fig. 4.75. For a given boson type, the jet-charge template is nearly independent of the flavor. However, there is a dependence of the jet mass on the (heavy) flavor of the boson decay products. The independence of the jet mass and jet charge distributions is demonstrated in Fig. 4.76.




Figure 4.73: The distribution of the jet mass for Z boson jets with 0, 1, or 2 matched muons before (left) and after (right) adding the muon to the jet four-vector before computing the mass.

The likelihood function is constructed by taking the ratio of the probability distribution functions $p(M, Q, B|V)$, for $V \in \{W, Z\}$, determined from the templates in Eq. 4.34. Every bin i of the 3D histogram that approximates $p(M, Q, B|V)$ is assigned a pair of numbers $(i, s_i/b_i)$ where s_i is the overall fraction of the signal (Z or




Figure 4.74: Left: The relationship between the leading and sub-leading MV1 distributions. The distributions are normalized per bin of the leading small-radius jet (horizontal axis). The two MV1 values are independent if the distribution in each bin along the vertical axis does not change as a function of the horizontal axis (which is true within the MC statistical uncertainties). Right: A validation of the templates for the combined binned MV1 B (template is dashed).

W) in bin i and b_i is the fraction of the overall background (the other boson flavor) in bin i . Bins are then sorted from largest to smallest s_i/b_i , with $f(i)$ defining a map from the old bin index to the new, sorted one. There are then two 1D histograms: for the signal, bin j has bin content $s_{f^{-1}(j)}$ and for the background, bin j has bin content $b_{f^{-1}(j)}$. The optimal tagging procedure is then to set a threshold on the new 1D histograms. The full likelihood ratio of the combined tagger is shown in Fig. 4.77 where the thresholds required for 90%, 50%, and 10% Z -boson tagging efficiency are marked with shaded regions.

Curves displaying the tagging performance for all possible subsets of $\{M, Q, B\}$ are shown in Fig. 4.78. There are 30 possible values for B , which are therefore represented by discrete points. The jet mass is the best performing single variable for medium to high Z -boson efficiencies, with visible improvement for $M+B$ and $M+Q$. There is a significant gain from combining all three variables for Z -boson tagging efficiency above about 20%. Below 20%, the combined tagger is dominated by B where the




Figure 4.75: (a) The jet mass $p(M|\mathcal{F}, V)$ and (b) jet charge $p(Q|\mathcal{F}, V)$ templates conditioned on the flavor $\mathcal{F}$ of the boson V decay for jets with $200 \text{ GeV} < p_T < 400 \text{ GeV}$. The solid lines are for Z boson decays and the dashed lines are for W boson decays.




Figure 4.76: The joint distribution of the jet mass and jet charge for W boson jets (left) and Z boson jets (right). The peak of the distribution along the jet mass axis is shifted toward higher values for the Z due to its higher mass. The linear correlation is less than 1% in both cases and the two distributions are nearly independent.

$Z \rightarrow b\bar{b}$ branching fraction no longer limits Z -boson tagging efficiency. For Z -boson efficiencies of about 50%, one can achieve W^+ rejection factors ($1/\epsilon_{W^+}$) of 3.3 by using Q or B alone and about 5.0 using mass alone. For Z efficiencies of $\epsilon_Z = 90\%$, 50%, and 10%, W^+ rejection factors of 1.7, 8.3, and 1000, respectively, can be achieved with the combined tagger. Although most applications of boson-type tagging will target Z bosons as the signal while rejecting W bosons as background, the likelihood constructed in Fig. 4.77 can also be used to optimally distinguish W^+ bosons from Z bosons. The corresponding performance curves are shown in Fig. 4.79. The locations of the b -tagging points are all now shifted to high efficiency with respect to Fig. 4.78 because, for W^+ tagging, one wants to operate in the high-efficiency b -tagging bins (whereas the opposite is optimal for Z tagging). At an efficiency of $\epsilon_{W^+} = 50\%$, a Z -boson rejection factor of $1/\epsilon_Z \approx 6.7$ can be achieved.




Figure 4.77: The full likelihood ratio for the tagger formed from jet mass, jet charge, and a small-radius jet b -tagging discriminant. The black histogram shows the likelihood ratio for Z bosons and the red histogram is the likelihood ratio for W^+ bosons. The shaded areas show the region of the likelihood ratio corresponding to 90%, 50%, and 10% working points of the Z -boson tagging efficiency.




Figure 4.78: The tradeoff between Z efficiency and (a) $1 - (W^+ \text{ efficiency})$ (b) or $1/(W^+ \text{ efficiency})$ on (a) a linear scale and (b) a logarithmic scale. Each curve is constructed by placing thresholds on the likelihood constructed from the inputs indicated in the legend. Since the b-tagging discriminant is binned in efficiency, there are only discrete operating points for the tagger built only from B.




Figure 4.79: The tradeoff between W^+ efficiency and (a) $1 - (Z \text{ efficiency})$ or (b) $1/(Z \text{ efficiency})$ on (a) a linear scale and (b) a logarithmic scale. Each curve is constructed by placing thresholds on the likelihood constructed from the inputs indicated in the legend. Since the b-tagging discriminant is binned in efficiency, there are only discrete operating points for the tagger built only from B.

4.2.6 Systematic uncertainties

The performance curves in Figures 4.78 and 4.79 are based on the nominal modelling parameters of the ATLAS simulation. Additional studies show how the curves change due to the systematic uncertainties on the inputs to the likelihood function. Sources of experimental uncertainty include the calibrations of the large- and small-radius jet four-momenta, the b -tagging (which incorporates e.g. impact parameter modelling), and the modelling of track reconstruction.

The uncertainty on the scale of the large-radius jet mass calibration is estimated using the double ratio in data and MC simulation of calorimeter jet mass to track jet mass [343]. Tracks associated with a jet are well measured and provide an independent observable correlated with the jet energy. Uncertainties on the jet-mass resolution can have a non-negligible impact on the performance of the tagger. The jet-mass resolution uncertainty is determined from the difference in the widths of the boosted W boson jet-mass peak in semileptonic $t\bar{t}$ simulated and measured data events [343] and also from varying the simulation according to its systematic uncertainties [382]. The resolution is about 5 GeV in the Gaussian core of the mass spectrum and its uncertainty is about 20%. The impact of the jet-mass scale and resolution uncertainties on the boson-type tagger built using only the jet mass is shown in Fig. 4.81 for two nominal working points of 50% and 90% Z -boson tagging efficiency. Both the likelihood map f from Sec. 4.2.5 and the threshold value are fixed. Inputs to the tagger are shifted by their uncertainties and the 1D histograms described above are re-populated. The efficiencies for W and Z bosons are recomputed and shown as markers in Fig. 4.81(a). Coherent shifts of the jet masses (JMS) for W and Z bosons result in movement along the nominal performance curve corresponding to $\pm 10\%$ changes in the efficiency. However, there are also shifts away from the nominal curve because the optimal jet-mass cut is not a simple threshold. Variation of the jet-mass resolution (JMR) preserves the scale and so the movement is nearly perpendicular to the original performance curve, at the $\lesssim 5\%$ level, because of the increased overlap in the Z and W mass distributions³⁰. Shifts along the nominal curve optimally use

³⁰Although such shifts retain optimal use of the tagger (highest rejection for a fixed efficiency), they can degrade the quality of e.g. a cross-section measurement.

the input variables (albeit at different efficiencies), while shifts away from the nominal curve are a degradation in the performance. The impact of the fragmentation is estimated by using input variables from HERWIG but with the likelihood map from PYTHIA. PYTHIA and HERWIG have similar W/Z efficiencies at both the 50% and 90% benchmark points.

The systematic uncertainty on the efficiency of the tracking reconstruction is estimated by removing tracks associated with jets using an η -dependent probability [226]. The probability in the region $2.3 < |\eta| < 2.5$ is 7%; it is 4% for $1.9 < |\eta| < 2.3$, 3% for $1.3 < |\eta| < 1.9$, and 2% for $0 < |\eta| < 1.3$. These probabilities are known to be conservative in the most central η bins. There is also an uncertainty on the modelling of track merging for high- p_T jets, but the loss is expected to be negligible for jets with $p_T < 400$ GeV. Differences in the modelling of fragmentation can affect the expected performance for all the input variables, especially for the track-dependent observables. The impact of various uncertainties on the boson-type tagger built using only the jet charge is shown in Fig. 4.81(b). Since W and Z boson decays produce on average many tracks (see Sec. 4.2.7), removing a small number of them does not have a big impact on the jet-charge tagger as a result of the p_T -weighting in the jet charge sum.

The efficiency to b -tag jets of various flavors (b , c , and light) is measured in data using $t\bar{t}$ events, jets with identified charm hadrons, and multijet events [75]. The differences between data and MC simulation are typically a few percent and are applied as independent correction factors on a per-jet basis. The uncertainties on these scale factor measurements are used as estimates of the systematic uncertainty on the b -tagging. The sources of uncertainty are decomposed into many uncorrelated components (24 for b -jets, 16 for c -jets, and 48 for light-flavor jets) and the impact on the rejection is added in quadrature for a fixed value of ϵ_{signal} . The b -tagging of matched small-radius jets is also affected by uncertainties on the jet-energy scale and resolution. These quantities are varied within their uncertainties and if the shifted jet has $p_T < 25$ GeV, its MV1 value is not considered. The impact of various uncertainties on the boson-type tagger built using only the b -tagging discriminant for a 10% nominal Z efficiency is shown in Fig. 4.82. At this efficiency, the full

boson-type tagger is dominated by the **b**-tagging inputs, as seen in Fig. 4.78. The scale factor uncertainty for **b**-jets has no impact on the W efficiency (no real **b**-jets), but there is approximately a 10% uncertainty on the Z efficiency. The uncertainties on the jet-energy scale for small-radius jets are relevant only because of the 25 GeV p_T threshold. Since all of the large-radius jets are required to have $p_T > 200$ GeV, the threshold is relevant only in the rare case that one of the W daughters is nearly anti-parallel in the W rest frame to the direction of the W boost vector.

The **b**-tagging scale factors are only determined up to $p_T = 300$ GeV and then are extrapolated up to $p_T \gtrsim 500$ GeV using simulation. The left plot of Fig. 4.80 shows that the fraction of **b**-tagged jets with $p_T > 300$ GeV is negligible in the large-radius jet p_T range $200 < p_T < 400$ GeV. In principle, the (mis)modeling could depend on the ΔR between the **b**-jets as the scale factors are extracted for isolated **b**-jets. However, the studies in Ref. [383] based on $g \rightarrow b\bar{b}$ suggest that the systematic uncertainties are small.




Figure 4.80: The p_T spectrum of the small-radius **b**-jets and the ΔR between them when there are two.

The impact of the uncertainties on the jet-mass scale and resolution on the boson-type tagger built using all of the inputs (jet mass, jet charge, and **b**-tagging) is shown in Fig. 4.83(a). At very low Z -boson tagging efficiency, the tagger is dominated by **b**-tagging, so Fig. 4.82 is a good representation of the uncertainty on the full tagger's

performance. For higher efficiencies, the tagger is dominated by the jet mass, although the jet charge and b -tagging discriminant significantly improve the performance. The uncertainty on the full tagger's performance at the 50% and 90% Z -boson tagging efficiency benchmark points is due mostly to the uncertainty on the jet mass, which is why these uncertainties are shown in Fig. 4.83.




Figure 4.81: The impact of selected systematic uncertainties on benchmark working points of the boson-type tagger. (a) a jet-mass-only tagger, for 50% (left) and 90% Z efficiency benchmarks. (b) a jet-charge-only tagger, for 50% (left) and 90% Z efficiency benchmarks. The point marked HERWIG uses the alternative shower and hadronization model for the simulation, with the likelihood template from PYTHIA. See the text for an explanation of the notation in the legend.




Figure 4.82: The impact of selected systematic uncertainties on benchmark working points of a b-tagging-only tagger at a 10% Z efficiency benchmark. The b-tagging discriminant is binned, so there are only discrete operating points. The point marked **HERWIG** uses the alternative shower and hadronization model for the simulation, with the likelihood template from **PYTHIA**. The b-tagging scale factor uncertainties are determined separately for b-, c-, and light-quark jets. Variations are added in quadrature for each ‘truth’ jet flavor. There is no contribution from the b-jet scale factor uncertainties on the W rejection because there are no ‘truth’ b-jets. Conversely, the c- and light-jet scale factor uncertainties do not impact the Z bosons because at this low efficiency, all the selected Z bosons decay into $b\bar{b}$.




Figure 4.83: The impact of uncertainties on the jet-mass scale and resolution for 50% (a) and 90% (b) Z efficiency working points of the full boson-type tagger. The point marked HERWIG uses the alternative shower and hadronization model for the simulation, with the likelihood template from PYTHIA.

4.2.7 Validation of tagging variables using data

The tagger cannot be fully tested with data because it is not possible to isolate a pure sample of hadronically decaying Z bosons in pp collisions. However, the modelling of the variables used to design the tagger can be studied with a relatively pure and copious sample of hadronically decaying W bosons in $t\bar{t}$ events which can be tagged by the leptonic decay of the other W boson in the event (semileptonic $t\bar{t}$ events). Single-lepton triggers are used to reject most of the events from QCD multijet background processes. Candidate reconstructed $t\bar{t}$ events are chosen by requiring an electron or a muon with $p_T > 25$ GeV and $|\eta| < 2.5$, as well as a missing transverse momentum $E_T^{\text{miss}} > 20$ GeV. The electrons and muons are required to satisfy a series of quality criteria, including isolation³¹. Events are rejected if there is not exactly one electron or muon. In addition, the sum of the E_T^{miss} and the transverse mass³² of the W boson, reconstructed from the lepton and E_T^{miss} , is required to be greater than 60 GeV. Events must have at least one b -tagged jet (at the 70% efficiency working point) and have at least one large-radius trimmed jet with $p_T > 200$ GeV and $|\eta| < 2$. Furthermore, there must be a small-radius jet with $p_T > 25$ GeV, and $\Delta R < 1.5$ to the selected lepton (targeting the decay chain $t \rightarrow bW(\rightarrow \ell\nu)$). The other W boson candidate is selected as the leading large-radius trimmed jet with $\Delta R > 1.5$ from the small-radius jet that is matched to the lepton. The W +jets and multijet backgrounds are estimated from the data using the charge asymmetry and matrix methods, respectively [281]. The other backgrounds are estimated directly from MC simulation. Although the resulting event selection is expected to have a high $t\bar{t}$ purity (about 75%), the events cannot be compared directly to the isolated W bosons from the simulated W' boson decays. This is because there are several effects that make the typical large-radius jet in semileptonic $t\bar{t}$ events different from isolated W and Z boson jets in typical W' boson events:

1. The event selection is based on the reconstructed jet p_T , so even if $p_T^{\text{jet}} \gtrsim$

³¹Leptons are considered isolated if they are well separated from jets ($\Delta R > 0.4$) and the track/calorimeter energy within a small cone, centred on the lepton direction but excluding the lepton itself, is below a fixed relative value.

³²The transverse mass, m_T , is defined as $m_T^2 = 2p_T^{\text{lep}}E_T^{\text{miss}}(1 - \cos(\Delta\phi))$, where $\Delta\phi$ is the azimuthal angle between the lepton and the direction of the missing transverse momentum.

200 GeV for an $R = 1.0$ jet, the true hadronically decaying W boson in the event may have $p_T^W < 200$ GeV and thus the W boson decay products might not be collimated within $\Delta R < 1$.

2. There are more (close-by) jets in semileptonic $t\bar{t}$ events than in W' boson events. Jets not originating from the W boson can form the leading large-radius jet, or the b -jet from the same top-quark as the hadronically decaying W bosons can merge with the W boson decay products to form a large-radius jet.




Figure 4.84: The distribution of p_T^{jet}/p_T^W (left) and $\Delta R(\text{jet}, W)$ (right) for W' and $t\bar{t}$ events. The large tails in $t\bar{t}$ events are due to cases in which more than the W boson hadronic decay products are merged inside the large-radius jet.

The variables p_T^{jet}/p_T^W and $\Delta R(\text{jet}, W)$, for the W boson from the MC ‘truth’ record and the selected large-radius jet, are used to classify the various $t\bar{t}$ event sub-topologies. Events are labelled as having a **Boosted W** if $|p_T^{\text{jet}}/p_T^W - 1| < 0.1$ and $\Delta R(\text{jet}, W) < 0.1$. These numbers are based on the distribution for isolated W and Z bosons from the W' simulation. Figures 4.84 shows the distributions of p_T^{jet}/p_T^W and $\Delta R(\text{jet}, W)$ in both the W' sample and in the $t\bar{t}$ simulation. If the b -quark from the top-quark decay has an angular distance $\Delta R < 1.0$ from the selected large-radius jet, this jet is labelled as **b-contaminated**. All other $t\bar{t}$ events, including events where both W bosons decay into leptons, are labelled as **Other**. The p_T spectrum of the jets from




Figure 4.85: The joint distribution of p_T^{jet}/p_T^W and $\Delta R(\text{jet}, W)$ for $t\bar{t}$ events (left) and W' events (right). Dashed lines indicate the selection for the **Boosted W** category.

the classified events is shown in Fig. 4.86. In Fig. 4.86 and subsequent figures, systematic uncertainties on the simulation include the jet p_T and jet mass uncertainties described in Sec. 4.2.6, but exclude tracking uncertainties, which are sub-dominant. Events are vetoed if the selected large-radius jet has $p_T > 400$ GeV or if the ΔR between the selected large-radius jet and a tagged b -jet is less than 1.0. This suppresses the b -contaminated $t\bar{t}$ events. The effectiveness of the $t\bar{t}$ event classification is most easily seen from the jet mass distribution, shown in Fig. 4.87(a). The mass of the boosted W bosons from $t\bar{t}$ events is peaked around m_W , as is a small contribution from the hadronically decaying W bosons in single-top events in the Wt channel. There is no peak at m_t in the b -contaminated spectrum because of the b -jet veto, but there is a small non-resonant contribution below the top-quark mass, due to events in which one W daughter is matched with the b -jet. This is akin to the b -jet+lepton invariant mass used in other circumstances to measure top-quark properties and naturally has a scale around 150 GeV [384]. The low-mass peak in W +jets and the ‘other’ $t\bar{t}$ events is due to the Sudakov peak from QCD jets, the location of which scales with $R \times p_T$. The dependence on p_T of the W -peak position in Fig. 4.87(a) is shown in Fig. 4.87(b). Events with the leading jet in a window around the W mass, $50 \text{ GeV} < m^{\text{jet}} < 120$ GeV are selected and the median of the mass distribution is plotted in Fig. 4.87(b)

as a function of the jet p_T . The similar trend for the simulation and the data shows that the combination of the reconstructed jet-mass scale and ‘truth’ jet-mass scale is well modelled. To quantify the spread in the jet mass peak, various inter-quantile ranges are shown as a function of p_T in Fig. 4.87(c). The inter-quantile range of size $0\% < X < 50\%$ is defined as the difference between the $50\% + X\%$ quantile and the $50\% - X\%$ quantile, and is a measure of the spread in the distribution. The width of the boosted- W mass peak is well modelled within the statistical precision of the 2012 data sample.




Figure 4.86: The p_T distribution of the selected large-radius jets. The uncertainty band includes all the experimental uncertainties on the jet p_T and jet mass described in Sec. 4.2.6.

The modelling of boosted W bosons can also be studied using the jet-mass scale measured from tracks. Defining the variable r_{track} as the ratio of the jet mass determined from tracks to the jet mass determined from the calorimeter, the jet mass scale uncertainty is related to the difference from unity of the ratio of $\langle r_{\text{track}} \rangle$ in data to $\langle r_{\text{track}} \rangle$ in MC simulation. The mass scale uncertainty is calculated using the procedure described above, but with r_{track}^{-1} . If the jet consists only of pions, the natural scale for r_{track} is $2/3$, although there are significant physics and detector effects that introduce a large spread of values. The distribution of r_{track} in the $t\bar{t}$ -enriched event

sample with the same p_T and b -jet veto requirements as in Fig. 4.87 is shown in Fig. 4.88(a). Unlike the raw jet-mass distribution, the r_{track} distribution is similar for all of the sub-processes, as expected. The scale and spread of the r_{track} distribution are quantified in figures 4.88(b) and 4.88(c) using the p_T dependence of the median and inter-quantile ranges. Previous studies have indicated that the track multiplicity, n_{track} , in quark and gluon jets is not well modelled, especially for gluon jets, where n_{track} is lower in the data with respect to PYTHIA (see Chapter 3). The distribution of the track multiplicity for large- R jets in the $t\bar{t}$ -enriched event sample is shown in Fig. 4.89. The boosted W events are peaked at slightly lower values of the number of associated tracks compared to the quark/gluon jets from the other processes. The (charged) particle multiplicity increases for generic quark and gluon jets as a function of jet energy. However, the mass-scale of the jets produced from W boson decays is set by m_W so that in the absence of detector reconstruction effects, the track multiplicity distribution should be largely p_T independent. The p_T dependence of the track multiplicity is shown in Fig. 4.89(b) and 4.89(c) in the form of the median and the inter-quantile ranges. The median does increase because of the large non- W component as well as the finite detector acceptance for charged particles from the boosted W boson decay. The width is well modelled within the statistical precision of the data. However, there is disagreement for the median. Previous studies (including Ref. [167]) suggest that this is due to fragmentation modelling and not the modelling of the detector response.

The p_T -weighted distribution of the track charges defines the jet charge, which is shown in Fig. 4.90(a). The charge of the lepton from the leptonic W boson decay determines the expected charge of the hadronically decaying W boson candidate, allowing for a tag-and-probe study of the capability of charge tagging in hadronic W boson decays [164]. The jet charge for boosted W bosons for positively (negatively) charged leptons is clearly shifted to the left (right) of zero. There is also some separation between positive and negative W boson decays when the selected large-radius jet does not satisfy the criteria for being a boosted W boson. This is because the jet still contains some of the W boson decay products, and the jet charge is correlated with the charge of the W boson. The difference between the inclusive and boosted

W -boson jets is clearer in the p_T dependence plot of the median jet charge shown in Fig. 4.90(b). The medians of the distributions for boosted W jets are nearly twice as far apart as the medians for inclusive jets. However, in both cases the spread is less than the width of the distribution, shown as the inter-quantile range (inter-quantile range with $X = 25\%$) in Fig. 4.90(c). Even though there is some small disagreement for the median number of tracks, the p_T -weighted sum defining the jet charge is reasonably well modelled.

The remaining input to the boson tagger is the b -tagging discriminant for the matched small-radius jets. The efficiency-binned MV1 distributions are shown in Fig. 4.91(a) and 4.91(b) with the same selection criteria as for the previous figures, except that the b -jet veto is removed. The contamination due to the b -jet from the top-quark decay complicates a direct study of the MV1 distribution for boosted W jets; contamination from the b -quark decay products is seen clearly in the MV1 distribution at lower values of the efficiency. Most of the boosted W jets are in the highest efficiency bin because they have no real b -hadron decay.

Overall, the simulation models all three input variables well.




Figure 4.87: (a) The jet-mass distribution of the selected jets in semi-leptonic $t\bar{t}$ events. (b) The median of the mass distribution as a function of the jet p_T for events with the selected jet in the range $50 \text{ GeV} < m^{\text{jet}} < 120 \text{ GeV}$. This includes the contributions from events which are not classified as Boosted W. (c) For the same events as in (b), the inter-quantile range as a measure of spread. The quantiles are centred at the median. The uncertainty band includes all the experimental uncertainties on the jet p_T and jet mass described in Sec. 4.2.6. The inter-quantile range of size $0\% < X < 50\%$ is defined as the difference between the $50\% + X\%$ quantile and the $50\% - X\%$ quantile. Statistical uncertainty bars are included on the data points but are smaller than the markers in many bins.




Figure 4.88: (a) The distribution of r_{track} in the data for semi-leptonic $t\bar{t}$ events with the selected jet in the range $50 \text{ GeV} < m^{\text{jet}} < 120 \text{ GeV}$. (b) The median of the r_{track} distribution as a function of the jet p_T . This includes the contributions from events that are not classified as Boosted W. (c) The inter-quantile range as a measure of the width. The quantiles are centred at the median. The uncertainty band includes all the experimental uncertainties on the jet p_T and jet mass described in Sec. 4.2.6. The inter-quantile range of size $0\% < X < 50\%$ is defined as the difference between the $50\% + X\%$ quantile and the $50\% - X\%$ quantile. Statistical uncertainty bars are included on the data points but are smaller than the markers in many bins.




Figure 4.89: (a) The distribution of the number of tracks associated with the selected large-radius jet in the semi-leptonic $t\bar{t}$ data for events with the selected jet in the range $50 \text{ GeV} < m^{\text{jet}} < 120 \text{ GeV}$. (b) The median of the distribution of the number of tracks as a function of the jet p_T . This includes the contributions from events that are not classified as Boosted W. (c) The inter-quantile range as a measure of the width. The quantiles are centred at the median. The uncertainty band includes all the experimental uncertainties on the jet p_T and jet mass described in Sec. 4.2.6. The inter-quantile range of size $0\% < X < 50\%$ is defined as the difference between the $50\% + X\%$ quantile and the $50\% - X\%$ quantile. Statistical uncertainty bars are included on the data points but are smaller than the markers in many bins.




Figure 4.90: (a) The distribution of the jet charge in the data for semi-leptonic $t\bar{t}$ events with the selected jet in the range $50 \text{ GeV} < m^{\text{jet}} < 120 \text{ GeV}$. The ratio uses the positive lepton charge. (b) The median of the jet charge distribution as a function of the jet p_T . This includes the contributions from events that are not classified as Boosted W (except for the blue triangles, for which only the Boosted W is included). (c) The inter-quartile range as a measure of the width. The quantiles are centred at the median. The uncertainty band includes all the experimental uncertainties on the jet p_T and jet mass described in Sec. 4.2.6. The inter-quartile range is defined as the difference between the 75% quantile and the 25% quantile. Statistical uncertainty bars are included on the data points but are smaller than the markers in many bins.




Figure 4.91: The efficiency-binned MV1 distribution for the (a) leading and (b) sub-leading matched small-radius in semi-leptonic $t\bar{t}$ events. If there is no second small-radius jet with $p_T > 25$ GeV and $\Delta R < 1$ to the selected large-radius jet axis, the event is put in the ‘No jet’ category in (b). The uncertainty band includes all the experimental uncertainties on the jet p_T and jet mass and those related to the b -tagging described in Sec. 4.2.6. Statistical uncertainty bars are included on the data points but are smaller than the markers in many bins.

4.2.8 Outlook

The simulation studies of the boson-type tagger presented in Sec. 4.2.5 show that for bosons with $200 \text{ GeV} < p_T < 400 \text{ GeV}$, it is possible to achieve Z-boson efficiencies of $\epsilon_Z = 90\%$, 50% , and 10% with W^+ boson rejections of 1.7, 8.3 and 1000, respectively. Putting this into context, with $R(\epsilon_Z)$ defined as the lowest possible W-boson tagging efficiency at a fixed Z-boson tagging efficiency:

- The WZ/WW cross-section ratio is $\sim 20\%$ [385]. At the 50% type-tagger working point, one can change the ratio of events to

$$\frac{50\%}{R(50\%)} \times \frac{\sigma(WZ)}{\sigma(WW)} = \frac{50\%}{12\%} \times \frac{\sigma(WZ)}{\sigma(WW)} = \frac{50}{12} \times 20\% \approx 83\%, \quad (4.35)$$

with the possibility for a high-purity extraction of the WZ cross section in the semileptonic channel ($\ell\nu q\bar{q}$).

- Diboson resonances are predicted by many models of physics beyond the Standard Model. The all-hadronic channel provides a significantly higher yield than the leptonic channels. At the 90% type-tagger working point, one can distinguish ZZ from WZ with a likelihood ratio of $0.9^2/(0.9 \times 0.6) \sim 1.5$.
- At the 10% type-tagger working point, a leptophobic flavor-changing neutral current (with decays like in the SM) with a branching ratio of 1% would have the same number of events as the $t \rightarrow bW$ decay in $t\bar{t}$ production.³³

$$\frac{10\%}{R(10\%)} \times \frac{\Gamma(t \rightarrow Zc)}{\Gamma(t \rightarrow Wb)} = \frac{10\%}{0.1\%} \times \frac{\Gamma(t \rightarrow Zc)}{\Gamma(t \rightarrow Wb)} = 100 \times 1\% = 100\%. \quad (4.36)$$

Only the range $200 \text{ GeV} < p_T < 400 \text{ GeV}$ was studied thus far due to the availability of W bosons in the data. MC simulation suggest that the separation between

³³Up to impurities due to the high-occupancy $t\bar{t}$ environment.

W bosons and Z bosons from jet mass and jet charge is still powerful up to and beyond 1 TeV. The information from b -tagging degrades around 400 GeV as the two decay products from the boson become too close to resolve as two separate jets³⁴.

4.2.9 Conclusions

A tagger for distinguishing hadronically decaying boosted Z bosons from W bosons using the ATLAS detector has been presented. It will most likely be used after a boson tagger has rejected most QCD multijet events³⁵. Three discriminating variables are chosen which are sensitive to the differences in boson mass, charge, and branching ratios to specific quark flavors: large-radius jet mass, large-radius jet charge, and an associated small-radius jet b -tagging discriminant. For moderate and high Z -boson tagging efficiencies, the jet mass is the most discriminating of the three variables, but there is significant improvement in discrimination when combining all three inputs into a single tagger. At low Z -boson efficiencies, smaller than the $Z \rightarrow b\bar{b}$ branching ratio, the b -tagging discriminant is the most useful for rejecting W bosons. The full tagger is largely unaffected by many systematic uncertainties on the inputs, with the exception of the uncertainties on the jet-mass scale and resolution. While it is not possible to measure the tagger efficiencies directly in data due to the lack of a pure sample of boosted, hadronically decaying Z bosons, modelling of the likelihood function using hadronically decaying W bosons has been studied in the data. Overall, the simulation agrees well with the 20.3 fb^{-1} of $\sqrt{s} = 8 \text{ TeV}$ pp data collected at the LHC.

³⁴Smaller radius (track) jets can be used to recover the efficiency in this regime [386].

³⁵See Fig. 4.92 and 4.93 for a demonstration the (near-)independence of the jet mass and jet charge with a standard boson-versus-QCD tagging variable, 2-subjettiness [379].




Figure 4.92: The joint distribution of the jet mass and 2-subjettiness for W boson jets (left) and Z boson jets (right). The peak of the distribution along the jet mass axis is shifted toward higher values for the Z due to its higher mass. The linear correlation is less than $\pm 10\%$ in both cases and the two distributions are nearly independent.




Figure 4.93: The joint distribution of the jet charge and 2-subjettiness for W boson jets (left) and Z boson jets (right). The linear correlation is less than $\pm 2\%$ in both cases and the two distributions are nearly independent.

4.3 Machine Learning for Jet Tagging

Machine learning is a generic term to describe procedures for identifying and classifying structure within a dataset. As such, most analysis techniques can be described as a form of machine learning. However, there is a deeper connection between machine learning and jet physics: the fundamental object of study only exists in the context of machine learning. A jet is defined by a clustering algorithm, which is an example of an *unsupervised machine learning* technique. Unlike the output of most clustering procedures, jets have a *physical meaning*. The earlier sections in this chapter have shown that the quantum properties of jets can be calculated, calibrated, and measured with high precision. Even though there is an extensive literature on clustering techniques, the most commonly used jet algorithms were established within the high energy physics community. This is because the physical meaning of a jet only makes sense if the defining algorithm satisfies particular properties such as infrared and collinear safety. The first half of this section (Sec. 4.3.1) is dedicated to bridging this gap by minimally modifying one of the most common unsupervised learning techniques for use in jet physics. A new jet algorithm called *fuzzy jets* uses *mixture modeling* to cluster jets and is demonstrated on events with the hadronic decays of boosted boson and top quarks. The parameters of the learned fuzzy jets contain information about the quantum properties of jets, which can be used for jet tagging.

Jet tagging is an example of *supervised learning*. High energy physics is a unique setting for supervised learning because it is possible to generate arbitrarily large high fidelity simulation datasets that are *labeled* (have a known type or origin). This chapter has introduced many jet substructure observables useful for separating jets initiated by different partons or particles. The optimal tagger is one that uses the likelihood ratio based on the full radiation pattern within the jet. In practice, it is not possible to compute the full likelihood. However, many sophisticated supervised learning techniques have been designed to be close approximations to the likelihood and can achieve near-optimal performance. Section 4.3.2 will demonstrate how state-of-the-art techniques borrowed from computer vision can improve the performance of jet tagging by thinking of the jet radiation pattern as an image. Machine learning is

a tool to guide but not replace physical intuition. Therefore, one of the main focuses of Sec. 4.3.2 is to visualize what the machine learning algorithms are learning from the radiation pattern in jets.

There is an ever-growing machine learning literature that will aid physics analyses at the LHC to fully exploit the data. This section ends with a brief discussion of prospects for the future in Sec. 4.3.3.

4.3.1 Fuzzy Jets

The purpose of this section is to introduce a new paradigm for jet clustering, called *fuzzy jets*³⁶, based on probabilistic mixture modeling and to demonstrate its use in boosted topologies. Section 4.3.1.1 introduces the statistical concept of a mixture model and describes the necessary modification to make the procedure IRC safe (see Sec. 4.3). Section 4.3.1.2 gives one efficient method for clustering fuzzy jets based on the Expectation-Maximization (EM) algorithm. Section 4.3.1.3.4 contains several examples comparing fuzzy jets with sequential recombination and Sec. 4.3.1.4 describes how one might mitigate the impact of overlapping proton-proton collisions (pileup). Conclusions are presented in Sec. 4.3.1.5 with some summary remarks and outlook for the future.

4.3.1.1 Mixture Model Jets

Mixture models [388] are a statistical tool for clustering which postulate a particular class of probability densities for the data to be clustered. Generically, for grouping n m -dimensional data points into k clusters, the mixture model density is

$$p(x_1, \dots, x_m | \pi, \theta) = \prod_{i=1}^n \left(\sum_{j=1}^k \pi_j f(x_i | \theta_j) \right), \quad (4.37)$$

³⁶The ideas presented in this section are published in Ref. [387]. Many of the studies presented in this section were performed by Conrad Stansbury. In particular, Stansbury made the final versions of Fig. 4.96-4.102, 4.105, and 4.106.

where π_j is the unknown weight of cluster j such that $\sum_j \pi_j = 1$ and $f(\mathbf{x}_i|\theta_j)$ is a probability density on $\mathbf{n}$ -dimensions with unknown parameters θ_j to be learned from the data. A common choice for f is the normal density Φ with $\theta_j = (\mu_j, \Sigma_j)$ for μ_j the $\mathbf{m}$ -dimensional mean and Σ_j the $\mathbf{m} \times \mathbf{m}$ covariance matrix. In the mixture model paradigm, the θ_j are the cluster properties; in the Gaussian case, μ_j is the location of cluster j and Σ_j describes its shape in the $\mathbf{m}$ -dimensional space. When clustering with a finite mixture, the number of clusters k must be specified ahead of time³⁷, which is dual to the usual use of sequential recombination³⁸ in which k is learned and the size of jets is specified ahead of time. The standard objective in mixture modeling is to select the parameters θ_j which maximize the likelihood (Eq. 4.37) of the observed dataset. Figure 4.94 illustrates what the learned event density might look like for $k = 3$ and Gaussian $f = \Phi$ in $\mathbf{m} = 2$ dimensions.

An equivalent way of approaching mixture modeling is to view Eq. 4.37 as the density used to generate the data. In other words, view the data as having been drawn randomly from the density specified in Eq. 4.37, with the following setup:

1. Throw $\mathbf{n}$ independent and identical k -sided dice with probability π_j to land on side $j = 1, \dots, k$ and label the outcomes $\lambda_1, \dots, \lambda_n$.
2. Independent of the others, data point $i \in \{1, \dots, \mathbf{n}\}$ is drawn randomly from $f(\cdot | \theta_{\lambda_i})$.

Once θ and π are learned by minimizing Eq. 4.37, one can compute $q_{ij} = \Pr(\lambda_i = j | \mathbf{x}_i)$, the posterior probability that $\mathbf{x}_i$ was generated by $f(\cdot | \theta_j)$ or, intuitively, the posterior probability that $\mathbf{x}_i$ belongs to cluster j . The q_{ij} are the *soft assignments* of particles i to jet j and will play an important role in Sec. 4.3.1.2 when showing how to maximize the likelihood in Eq. 4.37. In particular, $q_{ij} = \pi_j f(\mathbf{x}_i|\theta_j) / \sum_j \pi_j f(\mathbf{x}_i|\theta_j)$.

³⁷There is a wealth of literature on the subject of choosing k , for a survey of methods, see [389]. The likelihood monotonically increases with k ; as alternatives to maximum likelihood, one can for instance look for kinks in the likelihood as a function of k [390].

³⁸It is similar to the exclusive form of the k_t sequential recombination scheme [391]. The exclusive nature of the algorithm (and the minimization procedure used to find the jets) is similar to the XCone algorithm [392, 393].




Figure 4.94: An example of the learned per-particle probability density specified in Eq. 4.37 with $k = 3$ and Gaussian $f = \Phi$ in $m = 2$ dimensions. One cluster is associated with each component density $\Phi_i = \Phi(\cdot \mid \mu_i, \Sigma_i)$, where the dot $\cdot$ is a placeholder for the function argument.

Jets produced with mixture modeling are called *fuzzy jets* because of the soft memberships - every particle can belong to every jet with some probability³⁹. This can be seen explicitly in Fig. 4.94 where the densities of all three clusters are everywhere nonzero, so $q_{ij} > 0$ for all j . The idea of probabilistic membership was recently studied in the context of the Q-jets algorithm [395] in which the same event is interpreted many times by injecting randomness into the clustering procedure. Unlike Q-jets, fuzzy jets allocates the soft membership functions deterministically throughout the clustering procedure. However, like Q-jets, there is an ambiguity in how to assign kinematic properties to the clustered jets. Fuzzy jets are defined by their shape (and location), not their constituents. This is in contrast to anti- k_t jets, which are defined by their constituents without an explicit shape determined from the clustering procedure. One simple assignment scheme is to define the momentum of a jet j as

³⁹Soft assignments for jets during clustering was studied in the context of the “optimal jet finder” [394] which maximizes a function of the soft assignments.

$$p_{\text{jet } j} = \sum_{i=1}^m p_i \begin{cases} 1 & j = \text{argmax}_k q_{ik} \\ 0 & \text{else} \end{cases}. \quad (4.38)$$

This procedure assigns every particle to its most probable associated jet and will be known as the hard maximum likelihood (HML) scheme, but is not the only possible assignment algorithm. The dual problem in sequential recombination is the jet area, which must be defined [396], whereas the jet kinematics are the ‘natural’ coordinates.

For the remainder of the section, the likelihood in Eq. 4.37 is specialized to the case of clustering particles into jets at a collider like the LHC. Consider a mixture model in two dimensions⁴⁰ with $\mathbf{x}_i = \rho_i$. The resulting mixture model (MM) jets are inherently not IR safe: particle $\mathbf{p}_T$ does not appear in the likelihood and therefore arbitrarily low energy particles can influence the clustering procedure. Therefore, the log likelihood is slightly modified:

$$\log \mathcal{L}(\{\mathbf{p}_{T,i}, \rho_i\}|\theta) = \sum_{i=1}^m p_{T,i}^\alpha \log \left(\sum_{j=1}^k \pi_j f(\rho_i|\theta_j) \right), \quad (4.39)$$

where α is a weighting factor. Equation 4.39 is the log of Eq. 4.37 with the term $p_{T,i}^\alpha$ inserted in the outer sum. For $\alpha > 0$, the resulting *modified* mixture model (mMM) jets are IR safe, and when $\alpha = 1$, the jets are C safe. Therefore, for $\alpha = 1$, the jets are IRC safe. Different choices of component densities f in Eq. 4.39 give rise to different IRC safe MM jet algorithms. Several possibilities for f have been studied, but for the remainder of this section uses a (wrapped) Gaussian⁴¹ $f = \Phi$. The resulting fuzzy jets are called modified Gaussian Mixture Model jets (mGMM) and are parameterized by the locations μ_j , the covariance matrices Σ_i , and the cluster

⁴⁰One must take care in selecting a class of densities appropriate for the angular quantity ϕ . For more details on the wrapped Gaussian distribution and motivation for its use in this context, see Appendix C.2.

⁴¹When f is a circular step function, the algorithm is related to the Snowmass iterative cone algorithm [397] via the ‘Snowmass Potential’ [398].

weights π_j . For initialization, $\pi_j = 1/k$ and $\Sigma_j = I$.

Since practical procedures for maximizing the modified likelihood in Eq. 4.39 may converge to stationary points that are not globally optimal, the output of a fuzzy jet algorithm will depend on an initial setting of the cluster parameters θ and π . One simple procedure, used exclusively for the rest of the section, is to seed fuzzy jets based on the output of a sequential recombination jet algorithm. This guarantees an IRC safe initial condition and therefore the entire procedure is IRC safe.

4.3.1.2 Clustering Fuzzy Jets: the EM Algorithm

One iterative procedure for maximizing the mixture model likelihood in Eq. 4.37 is the *Expectation-Maximization* (EM) algorithm [399–401]. After initializing the cluster locations and prior density π , the following two steps are repeated:

Expectation Given the current values of θ_j , compute the fuzzy membership probabilities $q_{ij} = \pi_j \Phi(\vec{\rho}_i | \mu_j, \Sigma_j) / \sum_{j'} \pi_{j'} \Phi(\vec{\rho}_i | \mu_{j'}, \Sigma_{j'})$.

Maximization Given q_{ij} , maximize the *expected modified complete log likelihood* over the parameters π, μ, Σ .

The expected modified complete log likelihood has the form

$$\sum_{i=1}^N \sum_{j=1}^k p_{Ti}^\alpha (q_{ij} \log \Phi(\vec{\rho}_i; \vec{\mu}_j, \Sigma_j) + q_{ij} \log \pi_j). \quad (4.40)$$

Note that the expected modified complete log likelihood is not the same as the expected modified log likelihood, shown in Eq. 4.39. They differ in that the complete log likelihood has the second sum outside the logarithm while Eq. 4.39 has the sum inside the logarithm. The power of the EM algorithm is that maximizing the complete log likelihood results in iteration scheme that monotonically improves the original log likelihood. This desirable property of the EM algorithm is still true when $\alpha > 0$; for a proof, see Appendix C.3. Many choices for f have closed form maxima for the M step; in the Gaussian $f = \Phi$ case outlined above, the updates are given by

$$\mu_j^* = \sum_{i=1}^n \tilde{q}_{ij} x_i \quad \Sigma_j^* = \sum_{i=1}^n \tilde{q}_{ij} (x_i - \mu_j)(x_i - \mu_j)^T \quad \pi_j^* = \frac{1}{\sum_{i=1}^n p_{\tilde{\pi}}^\alpha} \sum_{i=1}^n p_{\tilde{\pi}}^\alpha \tilde{q}_{ij}, \quad (4.41)$$

where $\tilde{q}_{ij} = q_{ij} p_{\tilde{\pi}}^\alpha / \sum_{l=1}^n q_{il} p_{\tilde{\pi}}^\alpha$. The well-known k-means clustering algorithm [402] can be recovered as the limit of expectation-maximization in a Gaussian mixture model with $\Sigma = \sigma^2 I, \sigma^2 \rightarrow 0$. Figure 4.95 illustrates GMM clustering using the EM algorithm with $k = 2$ clusters. The EM algorithm readily accommodates constraints on the model parameters. One constraint for simplicity that is used throughout the rest of the section is $\Sigma_j = \sigma_j^2 I$ for all j , which requires the curves of constant likelihood in (y, ϕ) to be circular. The learned value of σ_j will be useful for distinguishing jets originating from different physics processes. Note that since the modified complete log likelihood is IRC safe, the EM algorithm does not break the IRC safety of the original log likelihood.




Figure 4.95: An illustration of the EM algorithm for $k = 2$. The circles represent data points, the triangles represent the estimated cluster locations μ_j , and the ellipsoids are equidensity contours describing the shapes Σ_j of the learned cluster distributions. In the E-step, bluer colors correspond to higher value of $p_{i, \text{blue jet}}$.

4.3.1.3 Comparisons with Sequential Recombination and Jet Tagging

This section describes some numerical comparisons between sequential recombination and fuzzy jets. Section 4.3.1.3.1 summarizes the simulation details with some first event displays showing both fuzzy and sequential recombination jets. These two approaches to jet clustering are studied over an ensemble of events in Sec. 4.3.1.3.2. A third subsection, Sec. 4.3.1.3.3, illustrates that fuzzy jets captures new information about the hadronic final state, and in the fourth section, Sec 4.3.1.3.4, it is demonstrated that this new information can be used to classify the jet type.

4.3.1.3.1 Details of the Simulation

Simulated W' , Z' , and QCD multijet events are generated using the same setup as in Sec. 4.1.3.2. Large-radius $R = 1.0$ anti- k_t trimmed jets with $k_t R = 0.3$ subjets groomed with $f_{\text{cut}} = 0.05$ are used as a benchmark. These jets are also used to seed the fuzzy jets using a threshold of 5 GeV⁴². The choice of the parameters for the anti- k_t jet seeds is akin to the radius parameter R in the usual sequential recombination paradigm in that they can have a significant impact on the clustered jet properties. In complete analogy to the choice of R , the choice of seed jet parameters will depend on the targeted final state and the initial event conditions (e.g. pileup).

The EM algorithm for fuzzy jet clustering is terminated when the per iteration increase in the log likelihood is less than 10^{-6} for five consecutive iterations, or when a maximum of 100 iterations is reached. In practice most events converge after a much smaller number of iterations than this bound, with only a small fraction of events stopping for lack of convergence, and then only in high pileup scenarios ($n_{\text{PU}} > 80$).

To model the discretization and finite acceptance of a real detector, a calorimeter of towers with size 0.1×0.1 in (y, ϕ) extends out to $y = 5.0$. The total momentum of the simulated particles incident upon a particular cell are added as scalars and the four-vector $\mathbf{p}_j$ of any particular tower j is given by

⁴²This low threshold guarantees that there are enough seed jets around to capture the radiation from the underlying event. Another strategy could be to use the *event jet* (see Sec. 4.3.1.4) even when there is no pileup.

$$\mathbf{p}_j = \sum_{i \text{ incident on } j} E_i (\cos \phi_j / \cosh y_j, \sin \phi_j / \cosh y_j, \sinh y_j / \cosh y_j, 1). \quad (4.42)$$

Without any corrections, fuzzy jets are significantly sensitive to pileup (see Sec. 4.3.1.4). One simple way to mitigate this sensitivity is to use a local pileup mitigation technique such as charged-hadron-subtraction, by which charged pileup particles (identified by their primary vertex) are subtracted from towers within the acceptance of the tracker $|\eta| < 2.5$.




Figure 4.96: A $Z' \rightarrow t\bar{t}$ event clustered with fuzzy jets (left) or anti- k_t $R = 1$ (right) without pileup ($n_{PU} = 0$). The circles indicate the 1σ contour (fuzzy jets) or the ungroomed jet area (anti- k_t). The small filled colored circles are the particles, with the color and size indicating their energy. The top quark locations from the generator-record are indicated by red stars. In the left plot, anti- k_t jet locations are shown with gray crosses where the long tail points towards the mGMM jet for which it was a seed. In the right plot, darker discs correspond to higher p_T jets. The highest p_T fuzzy jet has a black 1σ contour while all others are shown in gray.

A representative $Z' \rightarrow t\bar{t}$ event is shown in Figure 4.96. In contrast to the anti- k_t jets, fuzzy jets vary widely in radial size. The jets centered around the top quark locations did not move far from their anti- k_t seed jets, though the final size is much smaller than one. The lower p_T fuzzy jets moved a long distance from the seed jet location and are bigger than 1 in order to accommodate the diffuse radiation in the event. Another new feature of fuzzy jets compared to anti- k_t jets is that they can overlap with each other. Overlapping mGMM jets are an expression of structure inadequately captured with a single Gaussian shape. The ability to learn features at different scales in the same event without relying on a size parameter like the anti- k_t radius parameter can give mGMM fuzzy jets additional descriptive power over anti- k_t and other traditional jet algorithms.

4.3.1.3.2 Kinematic Properties of Fuzzy Jets

Due to the p_T weighting in the event likelihood, the hard mGMM jets (under HML) have a similar location and total energy as the leading anti- k_t jets. This is demonstrated by Fig. 4.97, which shows that the p_T spectrum of the leading mGMM jet is nearly identical to spectrum for the leading anti- k_t jet. The mGMM algorithm differs from the anti- k_t algorithm in the size and structure of clustered jets. One important variable sensitive to the distribution of energy within a jet is the jet mass. Figure 4.98 shows the jet mass distribution for the same jets as in Fig. 4.97, still using the HML scheme. Even though the two algorithms learn a similar core, the mass distributions are significantly different. Both mass distributions show clear peaks near the W boson and top quark masses, but the size and shape of the peaks differs by algorithm. The W mass peak is higher using fuzzy jets for both the W' and Z' processes. In Z' events, the fuzzy jets tend to resolve the three-prong structure of top quark jets into two (often overlapping) fuzzy jets. One of these jets captures the hadronic W decay while the other corresponds to the b -jet. The low mass peak for the W' in the right plot of Fig. 4.98 occurs when the fuzzy jets decompose a single boosted W jet into two jets that each have a QCD-like jet mass. The trend toward lower masses is also observed for the leading jet in QCD multijets.




Figure 4.97: The jet p_T for the leading anti- k_t jet (left) and leading fuzzy jet under the HML particle assignment scheme (right). All the processes are re-weighted so that the anti- k_t p_T spectra are the same.




Figure 4.98: The jet mass for the leading anti- k_t (left) and leading fuzzy jet under the HML particle assignment scheme (right), in an anti- k_t leading jet p_T window of 350 to 450 GeV. All the processes are re-weighted so that the anti- k_t p_T distributions are the same. The dashed white lines mark $m_W \approx 80$ GeV and $m_{\text{top}} \approx 175$ GeV.

4.3.1.3.3 New Information from Fuzzy Jets

The properties Σ of a fuzzy jet can be useful in distinguishing jets resulting from different physics processes. In the simplest realization of mGMM jets already described above, $\Sigma = \sigma^2 \mathbf{I}$, where σ is a measure of the size of the core of a jet. Although σ is a simple variable to construct from the wealth of data available after clustering with the mGMM algorithm, it captures at least some of the schematic differences in the likelihood for $Z' \rightarrow t\bar{t}$ and $W' \rightarrow WZ$ relative to a QCD multijet background. The left plot of Fig. 4.99 also shows the distribution of σ over all fuzzy jets. The generic jet is nearly independent of the hard-scatter process and tends to be much larger than the usual small-radius jet size ($R = 0.4$). Fuzzy jets capturing the highest p_T structure in the event tend to be small (as the structure tend to be small), but the rest of the diffuse radiation in the event requires large fuzzy jets spread out over the detector. The distribution for the leading fuzzy jet σ is shown in the right plot of Fig. 4.99. The distribution for the sub-leading jet in signal events is qualitatively similar to the leading jet and is largely uncorrelated. In background events, the sub-leading jet is systematically wider than the leading jet. As expected from the $2m/p_T$ scaling⁴³ of the jet size, the right plot of Fig. 4.99 shows that top quark jets have a larger σ than W jets which have a larger σ than generic quark and gluon jets. However, Fig. 4.100 shows that σ is not 100% correlated with m/p_T ; the next section will show that σ provides additional information for jet tagging beyond m/p_T . Note that part of the new information in σ is resulting from the clustering procedure itself and not just the definition of the observable. For example, computing σ from the constituents of an anti- k_t jet (i.e. running fuzzy jets on these constituents with $k = 1$) would result in

$$\sigma^2 = \frac{\sum_{i=1}^n p_{T,i} \Delta R^2}{\sum_{i=1}^n p_{T,i}}, \quad (4.43)$$

which is nearly the same as m/p_T .

⁴³At leading order, there is an exact relationship between σ and m/p_T - See Appendix C.4.




Figure 4.99: The learned value of σ for all fuzzy jets (left) and for the highest p_T jet under the HML scheme (right).




Figure 4.100: The joint distribution of the leading fuzzy jet σ and the leading anti- k_t jet m/p_T for Z' events (left) and QCD multijet events (right). The Pearson correlation coefficient is shown in the bottom right of both plots.

4.3.1.3.4 Fuzzy Jets for Tagging

Many properties of events clustered with fuzzy jets may be useful for jet tagging, but for a brief illustration, Fig. 4.101 shows the performance of a tagger based on σ . The σ -based tagger is significantly better than the random tagger, providing a rejection of ~ 40 at a signal efficiency of 50% for both top quark event tagging and W boson event tagging. The word ‘event’ is used as a reminder that even though Z' events produce boosted top quarks, the fuzzy or anti- k_t jet may only contain the W-boson decay products (see Fig. 4.98). A relevant benchmark variable is the anti- k_t jet m/p_T , which is similarly useful and in contains similar information. The likelihood (i.e. optimal) combination of σ and m/p_T is significantly better than σ or m/p_T alone, indicating the the information in σ that is uncorrected with m/p_T from Fig. 4.100 adds useful discriminating information.




Figure 4.101: The tradeoff between signal efficiency versus QCD multijet rejection ($=1/\text{efficiency}$) when the signal process is $Z' \rightarrow t\bar{t}$ (left) or $W' \rightarrow WZ$ (right). The *random* tagger line the curve signal efficiency = background efficiency.

4.3.1.4 Underlying Event and Pileup

One of the interesting features of fuzzy jets is that for densities Φ with infinite support (such as the Gaussian), the area over which particles can belong to that jet is infinite. This is in sharp contrast to anti- k_t jets for which the area is bounded by πR^2 . Hard-scatter anti- k_t jets are unaffected by relatively soft nearby jets. However, if there are not enough fuzzy jets to capture the diffuse soft radiation in an event, the jets that would otherwise capture the hard-scatter energy must become larger. This is illustrated in Fig. 4.102 where both top-quark jets are significantly larger at $n_{\text{PU}} = 40$.




Figure 4.102: The same Z' event with $n_{\text{PU}} = 0$ (left) and $n_{\text{PU}} = 40$ (right). The grid lines show the 0.1×0.1 tower size and the vertical dashed lines show the range over which the charged pileup energy is subtracted from each tower. The top quark locations from the generator-record are indicated by red stars and anti- k_t jet locations are shown with gray crosses where the long tail points towards the mGMM jet for which it was a seed. The fuzzy jets themselves are represented by their 1σ contour.

One can force the fuzzy jets algorithm to focus on the hard-scatter by using $\alpha > 1$, but at the cost of losing IRC safety. Another possibility is to increase the number of seed jets. A third possibility is to artificially add a jet to the event likelihood that has a uniform constant density over the entire detector. This *event jet* can absorb the




Figure 4.103: The same schematic $k = 3$ per-particle probability density from Fig. 4.94, but with a constant likelihood added to represent the event jet.

diffuse soft radiation and allow the other fuzzy jets to ‘focus’ on the hard-scatter. The idea of an event jet is illustrated schematically in Fig. 4.103. A constant density can provide probability to soft particles far from the hard-scatter jets, which can stabilize the size of the hard-scatter fuzzy jets. Quantitatively, the algorithm is modified with $q_{ij} \rightarrow \frac{q_{ij}}{\gamma + \sum_k p_{ik}}$, where γ is the event jet weight. In principle, the algorithm could learn γ , but since it should scale linearly with the median pileup density ρ , one could reduce the algorithm complexity by fixing $\gamma = \kappa\rho$. A value of $\kappa \sim 0.3$ was found to be optimal over a wide range of processes. Under the HML scheme, a particle is assigned to the event jet if $\max_k p_{ik} < \gamma$. Studies indicate that when the event-jet is coupled with a simple tower-level pileup subtraction scheme, the resulting properties of the leading fuzzy jet are robust against pileup. Many complex constituent-based pileup subtraction schemes exist (see e.g. Ref. [360–362]); one simple procedure used here for illustration is $p_T \mapsto \max(p_{T,\text{uncorrected}} - \rho A, 0)$, where $A = 0.1^2$. The same event from Fig. 4.102 is shown with the above pileup corrections in Fig. 4.104. As desired, the the two leading jets corresponding to the top quarks are nearly the same size for $n_{PU} = 0$ and 40. The sub-leading jets shift as the soft radiation is balanced between them and the event jet. The stability of σ is quantified in Fig. 4.105 where both the mean and standard deviation of the σ distribution are nearly independent

of n_{PU} . Note that the standard deviation of the σ distribution *decreases* at high n_{PU} as all jets are large and nearly all the same size.




Figure 4.104: The same events as in Fig. 4.102, but with event-jet and tower-based pileup subtraction applied.




Figure 4.105: The mean and standard deviation of the σ distribution in Z' events as a function of n_{PU} with and without the event jet and tower-based pileup corrections.

4.3.1.5 Conclusions

The modified mixture model algorithms provide a new way of looking at whole event structure. In contrast to the usual uses of hierarchical-agglomerative algorithms like anti- k_t , the number of seeds is fixed ahead of time and their properties are learned during the clustering process. The learned parameters provide a new set of handles for distinguishing jets of different types. Even simple variables constructed out of the learned parameters of a mixture of isotropic Gaussian jets, like σ , offer complementary information for tagging W boson and top quark jets. Even though the variable σ is sensitive to pileup, small modifications to the fuzzy jets algorithm can mitigate the impact of pileup.

Fuzzy jets provide a new paradigm for jet clustering in high energy physics. These IRC safe likelihood-based clustering schemes set the stage for many possibilities for future studies related to jet tagging, probabilistic clustering, and pileup suppression. Figure 4.106 is the first step to bridge the gap between new machine-learning motivated unsupervised learning algorithms and analysis at the LHC: a first glimpse at the fuzzy jet σ with the $\sqrt{s} = 8$ TeV ATLAS data.




Figure 4.106: The distribution of the leading jet σ using the same event selection as in Sec. 4.2.7. Only the muon channel is included (negligible QCD multijets contribution).

4.3.2 Jet Images

One of the most complex and important supervised learning tasks is facial recognition. The field of computer vision has developed sophisticated tools for performing this task with ever-increasing gains in performance. The ATLAS calorimeter is analogous to a digital image: it is a scalar field⁴⁴ in two discrete dimensions where the pixels are calorimeter cells and the intensity is the measured energy. By using image representations of jets (*jet images* [403]), the entire set of computer vision techniques can be directly applied to jet tagging. Linear-discriminant based tagging with jet images can provide a similar performance as a simple tagger based on jet observables motivated by physical intuition [403]. With shallow neural networks, the jet images-based tagging performance can out-perform simple jet observables [404]. This section⁴⁵ investigates the use of deep neural networks (DNN) that are the state-of-the-art algorithms in the field of computer vision [406–408]. In addition to studying the performance of these algorithms, the focus is on exploring where the networks have extracted discriminating information. To begin, Sec. 4.3.2.1 is a detailed description of image pre-processing and the physical impact of each step. Section 4.3.2.2 briefly summarizes describes the neural network architectures explored in Sec. 4.3.2.3. The section ends with conclusions in Sec. 4.3.2.4.

4.3.2.1 Pre-processing and the Symmetries of Spacetime

The setup from Sec. 4.1.3.2 is used to simulate boosted W bosons and QCD multijets and the detector discretization from Sec. 4.3.1.3.1 is used to pixelate the energies. In practice, the detector and jet image granularities can be different, but are set equal here for simplicity. Large-radius trimmed jets are clustered with $R = 1.0$ with k_t $R = 0.3$ subjets groomed with $f_{\text{cut}} = 0.05$. Trimming mitigates the contribution

⁴⁴A natural extension of these methods is to use vector fields incorporating information from calorimeter segmentation (like RGB images) or even tracks. These provide interesting challenges as the granularity would vary by component.

⁴⁵The ideas presented in this section are published in Ref. [405]. Many of the studies presented in this section were performed by Luke de Oliveira. In particular, de Oliveira developed the network architectures and ran the training. In addition, M. Kagan helped setup some of the technical framework for the studies.

from pileup; a detailed investigation into the performance of the neural network for $n_{\text{PU}} > 0$ is beyond the scope of Sec. 4.3.2.

Three key jet features for distinguishing between W jets and QCD jets are the *jet mass*, *n -subjettiness*⁴⁶ and the ΔR between subjects of the trimmed jet. These observables are used for benchmarking the performance of the neural network in Sec. 4.3.2.3. The distributions of these three discriminating variables are shown in Fig. 4.107. The transverse momentum is also a useful observable for distinguishing signal from background. However, in practice different techniques may be optimized for individual p_T bins because most of the input variables have a strong particle-level and/or detector-level p_T dependence. To prevent the neural network from learning the jet p_T as a useful discriminant, the momentum spectrum is re-weighted so that the signal has the same p_T distribution as the background.




Figure 4.107: The distributions of the jet mass (left), 2-subjettiness (τ_{21}) (middle) and the ΔR between subjects (right) for signal (blue) and background (red) jets.

A *jet image* is formed by taking the constituents of a jet and discretizing its energy into pixels in (η, ϕ) . In this section, the detector and jet images have the same granularity, so the first step in forming the image is to draw a box of a fixed size (25×25) around the jet area.

In order for the machine learning algorithms to most efficiently learn discriminating features between signal and background and to not learn the symmetries of space-time, the jet images are pre-processed. This procedure can greatly improve

⁴⁶Defined using the winner-takes-all axis that increases the robustness to pileup [409].

performance and reduce the required size of the sample used for testing. The pre-processing procedure happens in four steps: translation, rotation, re-pixelation, and inversion. To begin, the jet images are translated so that the leading subjet is at $(\eta, \phi) = (0, 0)$. Translations in ϕ are rotations around the z -axis and so the pixel intensity is unchanged by this operation. On the other hand, translations in η are *Lorentz boosts* along the z -axis, which do not preserve the pixel intensity. A proper translation in η would modify the intensity. One simple modification of the jet image to circumvent this change is to replace the pixel intensity E_i with the transverse energy $p_{T,i} = E_i / \cosh(\eta_i)$. This new definition of intensity is invariant under translations in η and is used exclusively for the rest of this section.

The second step of pre-processing is to rotate the images around the center of the jet. If a jet has a second subjet, then the rotation is performed so that the second subjet is at $-\pi/2$. If no second subjet exists, then the jet image is rotated so that the first principle component of the pixel intensity distribution is aligned along the vertical axis. Unless the rotation is by an integer multiple of $\pi/4$, the rotated grid will not line up with the original grid. Therefore, the energy in the rotated grid must be re-distributed amongst the pixels of the original image grid. A cubic spline interpolation is used in this case - see Ref. [403] for details. The last step is a parity flip so that the right side of the jet image has the highest sum pixel intensity.

Figure 4.108 shows the average jet image for W boson jets and QCD jets before and after the rotation, re-pixelation, and parity flip steps of the pre-processing. The more pronounced second-subjet can already be observed in the left plots of Fig. 4.108, where there is a clear annulus for the signal W jets which is nearly absent for the background QCD jets. However, after the rotation, the second core of energy is well isolated and localized in the images. The spread of energy around the leading subjet is more diffuse for the QCD background which consists largely of gluon jets that have an octet radiation pattern. This is compared to the singlet nature of the W jets where the radiation is mostly restricted to the region between the two hard cores (see Chapter 2).

One standard pre-processing step that is often additionally applied in computer vision algorithms is normalization. A common normalization scheme is the L^2 norm




Figure 4.108: The average jet image for signal W jets (top) and background QCD jets (bottom) before (left) and after (right) applying the rotation, re-pixelation, and inversion steps of the pre-processing. The average is taken over images of jets with $240 \text{ GeV} < p_T < 260 \text{ GeV}$ and $65 \text{ GeV} < \text{mass} < 95 \text{ GeV}$.

such that $\sum I_i^2 = 1$ where I_i is the intensity of pixel i . This is particularly useful for the jet images where pixel intensities can span many orders of magnitude, and when there is large pixel intensity variations between images. In this study, the jet transverse momenta are all around 250 GeV, but this can be spread amongst many pixels or concentrated in only a few. The L^2 norm helps mitigate the spread and thus makes training easier for the machine learning algorithms. However, normalization can distort the information contained within the jet image. Some observables, such as the ΔR between subjets, is invariant under all of the pre-processing steps as well as normalization. However, consider the *image mass*,

$$m_I^2 = \sum_{i < j} E_i E_j (1 - \cos(\theta_{ij})), \quad (4.44)$$

where $E_i = I_i / \cosh(\eta_i)$ for pixel intensity I_i and θ_{ij} is the angle between massless four-vectors with η and ϕ at the i and j pixel centers. The image mass is not invariant under all pre-processing steps but does encode key information to identify highly boosted bosons that would ideally be preserved by the pre-processing steps. As discussed earlier, with the proper choice of pixel intensity, translations preserve the image mass since it is a Lorentz invariant quantity. However, the rotation pre-processing step does not preserve the image mass. To understand this effect, consider two four-vectors: $p^\mu = (1, 0, 0, 1)$ and $q^\mu = (0, 1, 0, 1)$. The invariant mass of these vectors is $\sqrt{2}$. The vector p^μ is at the center of the jet image coordinates and the vector q^μ is located at $\pi/2$ degrees. If the image is rotated around the jet axis so that the vector q^μ is at 0 degrees, akin to rotating the jet image so that the sub-leading subjet goes from $\pi/2$ to 0, then p^μ is unchanged but $q^\mu \rightarrow (1, 0, \sinh(1), \cosh(1))$. The new invariant mass of q^μ and p^μ is about 1, which is reduced from its original value of $\sqrt{2}$. The parity inversion pre-processing step does not impact the image mass, but a L^2 normalization does modify the image mass. The easiest way to see this is to take a series of images with exactly the same image mass but variable L^2 norm. The map $I_i \mapsto I_i / \sum_j I_j^2$ modifies the mass by $m_I \mapsto m_I / \sum_j I_j^2$ and so the variation in the normalizations induces a smearing in the jet-image mass distribution.

The impact of the various stages of pre-processing on the image mass are illustrated in Fig. 4.109. The finite segmentation of the simulated detector slightly degrades the jet mass resolution, but the translation and parity inversion (flip) have no impact, by construction, on the jet mass. The rotation that will have the biggest potential impact on the image mass is when the rotation angle is $\pi/2$ (maximally changing η and ϕ), which does lead to a small change in the mass distribution. A translation in η that uses energy as the intensity instead of $\mathbf{p}_T$ (referred to as the *naive translation*) and the L^2 normalization scheme both significantly broaden the mass distribution. One way to quantify the amount of information in the jet mass that is lost by various pre-processing steps is shown in the Receiver Operator Characteristic (ROC) curve of Fig. 4.110. Information about the mass is lost when the ability to use the mass to differentiate signal and background is diminished. The naive translation and the L^2 normalization schemes are significantly worse than the other image mass curves which are themselves similar.




Figure 4.109: The distribution of the image mass after various states of pre-processing for signal jets (left) and background jets (right). The naive translation and the L^2 normalization image masses are both multiplied by constants so that the centers of the distribution are roughly in the same location as for the other distributions.




Figure 4.110: The tradeoff between W boson (signal) jet efficiency and inverse QCD (background) efficiency for various pre-processing algorithms applied to the jet (images).

4.3.2.2 Network Architecture

One of the most successful architectures for modern computer vision is the convolution neural network (CNN or Convnet). A detailed description of the CNN, its components, and related ideas, is beyond the scope of this section. Traditional (shallow) neural networks are now standard tools (and likely common knowledge), but Ref. [410] is a thorough textbook on *deep* neural networks and Ref. [411] is a review with many references to current research. The basic feature of a CNN that distinguishes it from a regular (*fully connected*) network is that each node of the output layer is connected to only a small number of nodes (=pixels for the first layer) from the input layer. The connection from an output node to an input node is the result of a discrete convolution of a filter with a patch of the input. Convolutional networks work well for detecting features, wherever they may be in the image. However, unlike images of natural or human-made scenery, jet images are very sparse. Figure 4.111 shows the distribution of the occupancy. Typically only 5-10% of pixels are non-zero and lack edges or other obvious features. Tests with different filter sizes found that

an usually large filter of 11×11 was optimal. This size is just big enough to capture the only clear jet-by-jet feature: a core separated from a second node of radiation. To complement the convolutional network, a fully connected network based on the MaxOut activation function [407] is used for comparisons in the next section. For a complete description of the sequence of activation functions, non-linearities (rectified linear except at the last layer, where a sigmoid is used), and down-samplings see Sec. 4 in Ref. [405].




Figure 4.111: Left: The distribution of the fraction of pixels (occupancy) that have a nonzero entry (blue) or at least 1% of the scalar sum of the pixel intensities from all pixels (red). Right: A typical single W jet image.

4.3.2.3 Performance and Visualization

Figure 4.112 shows the W tagging performance of the DNNs compared with the benchmark physically-motivated observables. Both the CNN and MaxOut networks out-perform the single benchmarks and their pairwise combinations. For example, at a signal efficiency of 30%, the best DNN has a 60% larger rejection than the likelihood combination of mass and τ_{21} . The fully connected network outperforms the CNN and interestingly the CNN with normalized input images outperforms the CNN with unnormalized images. Section 4.3.2.1 showed that normalization washes out information about the jet mass, which is the first indication (more below) that the network(s) are not fully learning information about the jet mass.




Figure 4.112: Left: ROC curves for individual physically motivated features as well as three deep neural network discriminants. Right: the DNNs are compared with pairwise combinations of the physically motivated benchmarks.

One way to test if a neural network has learned the discriminating information in a benchmark observable is to assess the performance of a combination of the variable with the DNN output. The combinations of jet mass, τ_{21} and ΔR with the MaxOut network are shown in Fig. 4.113. Combining ΔR or τ_{21} with the DNN output does not improve the performance while there is a significant improvement for the mass+DNN combination. One common feature of ΔR and τ_{21} is that they are *scale-invariant*, i.e.

scaling the jet image by a constant amount (as in normalization) does not change their values. They encode strictly geometric information about the radiation pattern within the jet. In contrast, the jet mass depends on both geometric and scale information. Figure 4.113 may indicate that scale information is not well-learned by the network. Corresponding curves for the CNN show the same qualitative features as Fig. 4.113.




Figure 4.113: ROC curves that combined the DNN outputs with physics motivated features MaxOut architecture.

One way to visualize the trends from Fig. 4.114 is to consider how the DNN sculpts the distribution of the benchmark observables in background events, i.e. $p(\mathbf{x}|\text{DNN})$. If the background distribution of $\mathbf{x}$ for signal-like DNN output is not the same as the signal distribution of $\mathbf{x}$, then there is more information in $\mathbf{x}$ than is contained in the network output. The left plot of Fig. 4.114 shows that the peak of the τ_{21} distribution is essentially a non-linear function of the DNN output. For low values of the network output, the τ_{21} distribution is peaked at high (background-like) values while high DNN output morphs the distribution to be peaked at low (signal-like) values. A similar trend is observed for ΔR . When the network output is small (background-like), the distribution of ΔR is nearly uniform. However, for high DNN output (signal-like), the ΔR distribution is peaked around 0.6 (set by m and p_T) just like the signal distribution in Fig. 4.107. Even though Fig. 4.113 indicates that

not all of the information about the jet mass is learned by the network, the DNN does appropriately sculpt the background distribution for the extreme DNN outputs. When the DNN output is close to one or close to zero, the jet mass distribution is peaked at m_W or 65 (i.e. steeply falling) GeV, respectively. However, the distribution at intermediate values of the DNN is broad much broader than either extreme.




Figure 4.114: The distribution of τ_{21} (left), ΔR (middle) and the jet mass conditioned on the CNN output for background jets. Distributions for the MaxOut network look qualitatively the same.

The remainder of this section is dedicated to probing what about the radiation pattern the DNN has learned beyond e.g. τ_{21} and the jet mass. A first step is to study *what information could be learned* by the network by considering the typical signal and background jet images in a small window of τ_{21} and jet mass. Figure 4.115 shows the average jet image in three windows of τ_{21} for a fixed small window of the jet mass and jet p_T . As expected, in these small windows the signal and background distributions look nearly identical: at low τ_{21} the jets have two distinct cores of energy and at high τ_{21} there are no longer two clear subjects. The subtle differences between the top and bottom rows of Fig. 4.115 are magnified by taking the image differences, shown in Fig. 4.116. In the window with $\tau_{21} \in [0.19, 0.21]$, there are five features: a localized blue patch in the bottom center, a localized red patch just above that, a red diffuse region between the red patch and the center and then a blue dot just left of center surrounded by a red shell to the right. Each of these have a physics meaning: the lower two localized patches give information about the orientation of

the second subjet (ΔR) which is slightly wider for the QCD jets that need a wider angle to satisfy the mass requirement. The red diffuse region just above the localized patches is likely an indication of color flow: the W bosons are color singlets compared to the color octet gluon jet background, and thus one expects the radiation pattern to be mostly between the two subjets for the W . One can draw similar conclusions for all the features in each of the plots in Figure 4.116.




Figure 4.115: $W' \rightarrow WZ$ (top) and QCD (bottom) average jet-images in three small windows of τ_{21} : $[0.19, 0.21]$ (left), $[0.39, 0.41]$ (middle), and $[0.59, 0.61]$ (right). In all cases, jet mass is restricted to be between 79 GeV and 81 GeV and the jet p_T is required to be in the interval $[250, 260]$ GeV.

Figure 4.117 is one way of visualizing if the information available in Fig. 4.116 is learned by the network. Each pixel shows the linear correlation with the network output. The DNN is output is a non-linear function of the inputs, but the distribution of the correlation contains non-linear spatial information about where discrimination information is contained in the jet radiation pattern. Many of the same features from




Figure 4.116: The average difference between $W' \rightarrow WZ$ jet-images in same small windows of τ_{21} as Fig. 4.115. Red (blue) colors are more signal- (background)-like.

Fig. 4.116 appear in these correlation images. In particular, the radiation between the subjets does seem to be strongly correlated with the DNN - an indication that color flow information is playing a role in the DNN performance.




Figure 4.117: Pearson correlation coefficient for pixel intensity and the CNN output for $W' \rightarrow WZ$ and QCD (combined) in the same small windows of τ_{21} as Fig. 4.115.

Chapter 2 introduced the jet pull variable and showed that it is sensitive to color flow. The jet pull angle θ_P is adapted here for large-radius jets by using subjets instead of resolved jets. If the leading subjet is labeled J and the subleading subjet is labeled j , then there are two pull angles that may contain useful discriminating information related to color flow: $\theta_P(J, j)$ and $\theta_P(j, J)$. The former pull angle contains substructure information about J and the latter angle uses the substructure of j . Figure 4.118 compares the performance of θ_P -based taggers with the other benchmark observables

and the DNNs. In the left plot of Fig. 4.118, the jet mass and τ_{21} are restricted to a small range as in the previous figures. By construction, the jet mass and τ_{21} observable have little discriminating information. The DNNs are significantly better than the random tagger, but much worse than the inclusive performance from Fig. 4.112 (jet mass and τ_{21} are important inputs to the DNN). The jet pull-based taggers perform significantly better than the random tagger, but are significantly worse than the DNNs. A similar trend is true for the right plot of Fig. 4.118. Instead of restricting the phase space, the event weights have been applied in the right plot of Fig. 4.118 that make the joint distribution of jet mass and τ_{21} identical (uniform) for both the signal and background. All events are used, but by construction the jet mass and τ_{21} do not contain any useful discriminating information. The performance of the θ_P -based taggers are a significant fraction of the DNN-based tagger performance. However, when the DNN is trained with the weighted applied, it significantly out-performs the pull angles. This suggest that there is possibly more color flow information in the DNN that is not captured by θ_P and also shows that a significant fraction of the DNN ‘memory’ is dedicated to learning about τ_{21} and mass.




Figure 4.118: ROC curves including the jet pull angle θ_P in a restricted phase space (left) and using an inclusive event selection, but with events weighted so that the joint jet mass and τ_{21} distributions are identical between signal and background.

4.3.2.4 Outlook and Conclusions

Jet Images are a powerful paradigm for visualizing and classifying jets. This section has shown that when applied to jet images, deep neural networks outperform several known and highly discriminating physically-motivated jet observables such as the jet mass and n -subjettiness, τ_{21} . A series of tests have shown that some of these benchmark jet features are learned by the network, but others are not. In particular, the networks are able to effectively learn geometric information about the radiation pattern, but not scale information as captured in part by the jet mass. It is an important next step to develop techniques that allow the networks to also learn mass-like features. The visualization studies in re-weighted or redacted regions of phase space show that some of the residual information learned by the network can be attributed to the differences in color flow between the signal and background. Chapter 2 introduced the jet pull variable and demonstrated that it is sensitive to color flow. In this section, the jet pull has been adapted to large-radius jets using subjets. While the jet pull angle does carry some discriminating power beyond the jet mass and τ_{21} , it does not contain enough information to fully explain the DNN performance. Further studies of the visualizations may help to identify a simple feature like the jet pull angle that captures all or most of the color flow information learned by the networks.

The methods presented in this section have built a new link between high energy physics and computer vision. State-of-the-art classification techniques applied to jet images shows that there is a great potential to improve the performance of tagging algorithms using the extensive machine learning literature. In addition to improving the sensitivity of BSM searches, these new techniques may ultimately be able to improve the physical understanding of jets and their complex radiation pattern.

4.3.3 Conclusions and Future Outlook


Sections 4.3.1 and 4.3.2 have shown two successful applications of adapting machine learning techniques to jet physics. Domain specific knowledge (IRC safety, the symmetries of spacetime, etc.) have played an important role in specializing these techniques to jet physics and understanding what they have learned. Three main conclusions from these studies:

1. State-of-the-art machine learning techniques can significantly improve upon the performance of traditional techniques motivated directly by physical intuition.
2. Representing the data in new ways can expand physical intuition by highlighting properties that are not readily captured by current methods.
3. Most importantly: advanced machine learning techniques are tools to guide but not replace physical intuition. An algorithm is most useful if the performance gains can be physically understood and independently validated.

There are many interesting directions to take this work in the future. The extensive machine learning literature offers numerous possibilities for studying more complex tagging and reconstruction tasks such as low level tracking/calorimeter-cell clustering and calibration, full event tagging, pileup discrimination, and combining multiple detector elements into a single (multi-‘color’) jet image. With the large amount of high energy data to be collected in Run 2 and beyond, it will be important to study these techniques in-situ in order to develop calibrations and systematic uncertainties. Then, advanced machine learning techniques can be fully utilized to increase the sensitivity of LHC searches and measurement, including the study of rare and or subtle aspects of the SM and beyond.

Part III

The Search for a Light Stop Squark



Display of a candidate boosted top quark pair production event from proton-proton collisions recorded by ATLAS with LHC stable beams at a collision energy of 13 TeV. The red line shows the path of a muon with transverse momentum around 50 GeV through the detector. The dashed line shows the direction of the missing transverse momentum, which has a magnitude of about 470 GeV. The green and yellow bars indicate energy deposits in the liquid argon and scintillating-tile calorimeters, from these deposits 4 small-radius ($R = 0.4$) jets are identified with transverse momenta between 70 and 300 GeV. Three of these small-radius jets are re-clustered into the leading large-radius ($R = 1.0$) jet (not shown explicitly) with a transverse momentum of about 600 GeV and a jet mass of about 180 GeV, near the top quark mass. One of these three jets in addition to the fourth jet above 70 GeV are identified as having originated from b-quarks. Tracks reconstructed from hits in the inner tracking detector are shown as arcs curving in the solenoidal magnetic field.

The studies in Part II showed that quarks and gluons contain a wealth of information about the structure of the SM; in Part III, they will serve as a window to beyond the SM. In particular, the top quark holds a special place in the SM. With a near-unity Yukawa coupling, the top quark is the most massive known elementary particle and has the strongest coupling to the Higgs boson. As such, many extensions of the SM predict new particles that couple strongly or even exclusively to top quarks. One of the most compelling such theories is Supersymmetry (SUSY). Theoretically elegant and practical, SUSY is a powerful paradigm for explaining some of the issues with the SM discussed in Sec. 1.2. At the core of weak-scale SUSY is a light top squark (stop), the supersymmetric partner to the top quark. If sufficiently light, the stop will be copiously produced at the LHC and can result in experimentally rich final states, often via top quarks. Under mild assumptions, the top quarks from stop decays will always be accompanied by a stable weakly interacting SUSY particle that escapes detection. This results in a $t\bar{t} + E_T^{\text{miss}}$ topology that will be main focus of Part III.

Chapter 1 begins Part III with an introduction and motivation for SUSY and in particular for a relatively light stop. This introduction is slightly nontraditional, beginning with a purely theoretical motivation instead of the usual practical one associated with the ‘hierarchy problem’, which is discussed in Sec. 1.3. This order is chosen to stress that SUSY is a logical model-building extension of the SM, despite being broken below the electroweak scale. The hierarchy problem and the dark matter relic density (the ‘WIMP’ miracle) motivate the close proximity of the SUSY breaking scale with the electroweak scale. Light stops are a generic prediction of SUSY models that naturally solve the hierarchy problem and are produced in association with the lightest supersymmetric particle (LSP) that is a dark matter candidate.

Stop pair production produces an experimentally complex and challenging final state. Chapter 2 provides an overview of the analysis strategy that involves both simulation-based and data-driven techniques to estimate and validate background predictions. A variety of event selections are constructed to target a wide range of phenomenological signatures that are discussed in Sec. 2.1. The search presented in Part III spans all of Run 1 of the LHC with $\sqrt{s} = 8$ TeV and the beginning of Run 2 at

$\sqrt{s} = 13$ TeV. The analysis strategy has evolved over time, increasing in sophistication and sensitivity. The focus will be on the state-of-the-art, but the early methods are also discussed in order to show the origin of the enhanced sensitivity at each stage. The search targets the one lepton final state of stop pair production. Leptons are precisely measured with high efficiency and provide useful handles for differentiating signal from backgrounds. In particular, the generic production of quark and gluon jets is highly suppressed by requiring at least one reconstructed lepton. The two-lepton final state offers a particularly clean environment for searching for stop pair production, but the branching ratio is significantly smaller than the zero- and one-lepton final states.

An extensive toolkit of discriminating variables is constructed specifically for the $t\bar{t} + E_T^{\text{miss}}$ topology in the one lepton final state. Many of the variables utilize the missing momentum vector combined with kinematic properties of the other reconstructed objects. Chapter 3 describes all of the variables in detail, including new techniques that are used in this search for the first time. Due to its similarity to the signal signature, SM top quark pair production is one of the most important background processes. However, $t\bar{t}$ events with a single lepton can be reduced to a negligible level based on kinematic endpoints (Sec. 3.2.1.1). One of the dominant residual backgrounds is the pair or single production of top quarks resulting in final states with two real leptons. The construction of variables that can effectively suppress dilepton $t\bar{t}$ events will be a large focus of Chapter 3.

The discriminating variables from Chapter 3 are combined to form signal-sensitive event selections called *signal regions*. Chapter 4 describes the construction of the signal regions, including the optimization procedure for maximizing the sensitivity to stop pair production. The kinematic properties of the stop decay products depend on the mass of the stop as well as the mass difference between the stop and the LSP. Higher stop masses and wider mass gaps give rise to harder energy spectra. However, the stop cross section decreases with mass leading to a tradeoff between acceptance and absolute yield. Compressed spectra are challenging because the signature is relatively similar to SM top quark pair production. For low stop masses, the cross-section is sufficiently high to take advantage of subtle differences in the shapes of

distributions to increase the sensitivity when the phase space for the LSP is restricted.

In order to reduce the dependence on simulation and the sensitivity to mis-modeling the data, a series of background-enriched event selections are constructed to normalize background yields. These *control regions* are designed to be as close as possible to the signal regions in order to reduce the required simulation-based extrapolation in phase space to the signal region. Chapter 5 documents the background estimation, including the construction of the control regions. Chapter 6 describes a complete study of potential sources experimental and theoretical bias on the background estimates. Many sources of uncertainty are reduced by normalizing the predicted yield in the control region to the observed data. In addition, the total number of predicted events in the signal region is sufficiently small that the data statistical uncertainty dominates any residual systematic uncertainty. Nonetheless, there are some signal regions for which the systematic uncertainty is significant and plays a major role in setting the sensitivity of the search.

After combining the background predictions with the observations in the signal regions, there is no significant evidence for electroweak scale stops. As a result, limits are calculated to set bounds on the excluded models. Chapter 7 documents these limits and discusses future directions for the search as well as a retrospective analysis of all Run 1 LHC SUSY searches.

As a result of the search presented in Part III⁴⁷, simple stop models with m_{stop} up to almost 800 GeV for a wide range of LSP masses are excluded. This puts a severe constraint on electroweak scale SUSY and many other models that predict light top quark partners. There are always loopholes and the search will continue to push the limits up to and beyond $m_{\text{stop}} = 1$ TeV as well as fill in gaps at lower masses where more complicated models can allow stops to evade the current limits. This rich program has a strong foundation in the tools and techniques developed in Part III and will hopefully result in uncovering new aspects of the SM or discovering new particles in the (near) future.

⁴⁷The ATLAS search results from Part III are published in Ref. [412–415] and include technical input and many useful discussions with the entire ATLAS stop one-lepton analysis team, including M. Barisonzi, J. Montejo Berlingen, D. Boerner, T. Eifert, J. Gramling, A. Henrichs, J. Kuechler, P. Pani, S. Pataia, K. Rosbach, S. Strandberg, M. Ughetto, X. Wang, A. Yiming, and K. Yoshihara.

Chapter 1

Introduction and Motivation

The most elegant construction of a model is to be as extensive as possible while respecting all known symmetries of nature. In classical mechanics, this leads to the familiar $\mathcal{L} \propto m\mathbf{v}^2$ and in the SM requires and forbids certain terms in the Lagrangian. Supersymmetry is no exception. The symmetry group of the SM can be written as $\mathcal{S} \simeq \mathfrak{P} \times \text{SU}(3) \times \text{SU}(2) \times \text{U}(1)$, where $\mathfrak{P}$ is the Poincaré group that encodes the symmetries of spacetime and the second part of $\mathcal{S}$ is the internal symmetry group of the SM. The famous ‘no-go’ theorem from Coleman and Mandula [416] in 1967 showed that this structure is maximal: there is no non-trivial (direct product) way to mix the spacetime symmetry group with the internal symmetry group in 3+1 dimensions and retain non-zero scattering amplitudes. However, there is a unique [417] loophole - the symmetry group can be extended if the generators are not *bosonic*. This is quantified in Sec. 1.1 with construction of the super Poincaré group leading to SUSY as a symmetry of super spacetime. The effect of SUSY is to relate bosons and fermions and as such predicts many new particles. No SUSY partners have been observed and so SUSY must be broken below (at least) the electroweak scale, as explained by Sec. 1.2. A motivation for electroweak scale SUSY breaking is in Sec. 1.3, where SUSY is shown to be an elegant solution to the hierarchy problem and provides a natural candidate for dark matter. The minimal supersymmetric SM (MSSM) is introduced in Sec. 1.4 along with the properties of a light stop.

1.1 Superspace as an Extension of Spacetime

To illustrate¹ the construction of superspace and the action of SUSY without a heavy burden of notation, this section uses a $1+1$ dimensional model that retains most of the key features of the full $3+1$ setting². To begin, consider the $1+1$ representations of the usual spacetime. Let $\eta = \text{diag}(1, -1)$ be the $1+1$ dimensional metric. As in $3+1$ dimensions, a Lorentz transformation is a linear map which preserves the Minkowski distance. If $\mathbf{v}$ is a two-vector (the $1+1$ analogue of a four-vector), then the Minkowski distance is given by $\mathbf{v}^\top \eta \mathbf{v}$, where $\mathbf{v}^\top$ denotes transpose of $\mathbf{v}$ and the product represents matrix multiplication. Then, a Lorentz transformation Λ is a matrix such that $\mathbf{v}^\top \Lambda^\top \eta \Lambda \mathbf{v} = \mathbf{v}^\top \eta \mathbf{v}$ for all two-vectors $\mathbf{v}$. Therefore, Λ is characterized by $\Lambda^\top \eta \Lambda = \eta$. Without loss of generality, let $\Lambda_{12} = -\sinh(\mathbf{y})$ for $\mathbf{y} \in \mathbb{R}$ (the rapidity). Then, $\Lambda^\top \eta \Lambda = \eta$ results in three equations:

$$\begin{aligned}\Lambda_{11}^2 - \sinh(\mathbf{y})^2 &= 1 \\ \Lambda_{12}^2 - \Lambda_{22}^2 &= 1 \\ \Lambda_{11}\Lambda_{21} - \Lambda_{22}\sinh(\mathbf{y}) &= 0.\end{aligned}\tag{1.1}$$

Solving these equations leads to the general form of a Lorentz transformation:

$$\Lambda(\mathbf{y}) = \begin{pmatrix} \cosh(\mathbf{y}) & -\sinh(\mathbf{y}) \\ -\sinh(\mathbf{y}) & \cosh(\mathbf{y}) \end{pmatrix} = \exp(i\mathbf{y}J), \text{ where } J = \begin{pmatrix} 0 & i \\ i & 0 \end{pmatrix}, \tag{1.2}$$

A general *Poincaré transformation* is a combination of a Lorentz transformation and a translation in spacetime: $\mathbf{v} \mapsto \Lambda \mathbf{v} + \mathbf{w}$, for a two-vector $\mathbf{w}$.

¹Part of this introduction is based on the Part III essay in Ref. [418].

²Some aspects of this model appear while studying superstring theory [419], although in that context it is an auxiliary device. This section will take an approach which resembles aspects of previous work on $1+1$ super QED [420–422], but with a slightly different (concrete) angle. The complications of multidimensional representations can obscure the physical intuition and simplistic motivation for supersymmetry. The construction in $1+1$ removes many complications, such as those associated with the properties of spinors [423].

The Poincaré transformation can be represented by 3×3 matrix multiplication:

$$\begin{pmatrix} \Lambda & w \\ 0 & 1 \end{pmatrix} \quad \text{acting on} \quad \begin{pmatrix} v \\ 1 \end{pmatrix}. \quad (1.3)$$

The representation in Eq. 1.3 allows for an easy computation of the group laws of the Lie group $\mathfrak{P}$ of Poincaré transformations. The matrix J embedded in the 3×3 matrix is one of the generators. The full set of generators are given from the Taylor series expansion around the identity matrix:

$$M = \begin{pmatrix} 0 & i & 0 \\ i & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix} \quad E = \begin{pmatrix} 0 & 0 & -i \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix} \quad P = \begin{pmatrix} 0 & 0 & 0 \\ 0 & 0 & -i \\ 0 & 0 & 0 \end{pmatrix}. \quad (1.4)$$

Simple matrix multiplication with these explicit representations, shows that the defining commutation relations of the Poincaré algebra are $[M, E] = iP$ and $[M, P] = iE$ (boosts do not commute with translations). In 1+1 dimensions, there is no ‘spin,’ but one can construct the analogy of a spinor representation of the Lorentz subgroup $\mathfrak{L}$ of the Poincaré group. In higher dimensions, $\mathfrak{L}$ has multiple generators with non-trivial commutation relations. However, in the lower-dimensional case, the Lorentz group is Abelian and as such all irreducible representations are one-dimensional. Define the lower dimensional analogues of the gamma matrices [419]:

$$\gamma^0 = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix} \quad \gamma^1 = \begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix} \quad (1.5)$$

The matrices in Eq. 1.5 satisfy the Clifford algebra $\{\gamma^\mu, \gamma^\nu\} = 2\eta^{\mu\nu}I_2$, where I_2 is the 2×2 identity matrix. A representation of the Lorentz group is $K = \frac{i}{4}[\gamma^0, \gamma^1] = \frac{i}{2}\eta$

which is *similar* to J . Define a *Dirac spinor* ψ as a two component object which transforms as $\psi \mapsto e^{iyK}\psi$, where $e^{iyK} = \text{diag}(\exp(-y/2), \exp(y/2))$. In order to construct a Lagrangian out of Dirac spinors, they need to be combined to form Lorentz invariant quantities. As in the 3+1 case, $\psi^\dagger\psi$ does not work since $\psi^\dagger\psi \mapsto \psi^\dagger e^{2iyK}\psi$. Instead, let $\bar{\psi} = \psi^\dagger\gamma^0$, then $\bar{\psi}\psi \mapsto \psi^\dagger e^{iyK}\gamma^0 e^{iyK}\psi = \bar{\psi}\psi$. Simple matrix multiplication shows that $\bar{\psi}\gamma^\mu\psi$ is a Lorentz vector, i.e. transforms by $\exp(iyJ)$. This leads to the Lorentz invariant Dirac Lagrangian $\mathcal{L} = \bar{\psi}(i\gamma^\mu\partial_\mu - m)\psi$.

As in the 3+1 case, the Dirac spinors are not irreducible representations of the Lorentz group. In the 1+1 case this is evident because all irreducible representations of an Abelian group are one dimensional. This is also clear because under the action of the Lorentz group, the two components of ψ transform independently, as e^{iyK} is diagonal. Let $\psi = (\psi_L \ \psi_R)^T$ where the ψ_L and ψ_R are called *Weyl spinors* and transform as $\psi_L \mapsto e^{-y/2}\psi_L$ and $\psi_R \mapsto e^{y/2}\psi_R$. A curiosity of 1+1 dimensions is that one can choose ψ_L and ψ_R to be purely real and thus are *Majorana-Weyl* spinors [423].

Now, Minkowski space is extended to include two new *Grassman-valued* degrees of freedom, θ_1 and θ_2 . Unlike *bosonic* degrees of freedom (regular commuting numbers), the Grassman-valued degrees of freedom *anti-commute* with themselves and each other. In particular, this means that $\theta_i^2 = 0$. Furthermore, these new coordinates do not transform as a vector. Instead, they transform as Weyl spinors and when combined into θ , transform as a Dirac spinor. The resulting space is known as *superspace*. In addition to extending the space of coordinates, one can extend the group of transformations to include translations in the spinorial degrees of freedom. In general, let a spinorial translation, $\theta \mapsto \theta + \epsilon$, also affect the vector components of the superspace coordinate. If the effect is required to be linear, then for $\mathbf{a}, \mathbf{b} \in \mathbb{R}$, the most general form of a spinor coordinate translation is $(x^\mu, \theta_i) \mapsto (x^\mu + \bar{\epsilon}\gamma^\mu\theta, \theta + \epsilon)$ where $\bar{\epsilon} = \epsilon^T\gamma^0$. To see this, note that the only way to combine ϵ_i with one of x, t, θ_i and form a vector-like object is $\epsilon_i\theta_i$, which is both a commuting number and transforms as $\exp(\pm y)$. Thus, in combinations of $\pm\epsilon_i\theta_i$, one may hope to get the correct transformation law of x^μ .

A general transformation then has the form

$$\begin{pmatrix} t \\ x \end{pmatrix} \mapsto \begin{pmatrix} t' \\ x' \end{pmatrix} = \begin{pmatrix} t + a\epsilon_1\theta_1 + d\epsilon_2\theta_2 \\ x + a'\epsilon_1\theta_1 + d'\epsilon_2\theta_2 \end{pmatrix}. \quad (1.6)$$

One can transform $(x^\mu)'$ using the vector law (lefthand side of Eq. 1.7) and compare to the transformation of the summands (righthand side of Eq. 1.7).

$$\Lambda(\phi) \begin{pmatrix} t' - t \\ x' - x \end{pmatrix} = \begin{pmatrix} ae^{-2\phi/2}\epsilon_1\theta_1 + de^{2\phi/2}\epsilon_2\theta_2 \\ a'e^{-2\phi/2}\epsilon_1\theta_1 + d'\epsilon_2e^{2\phi/2}\theta_2 \end{pmatrix} \quad (1.7)$$

Equating terms results in $d = -d'$, $a = a'$. After renaming constants³, this becomes

$$\begin{pmatrix} t' \\ x' \end{pmatrix} = \begin{pmatrix} t + \epsilon_1\theta_1 + \epsilon_2\theta_2 \\ x + \epsilon_1\theta_1 - \epsilon_2\theta_2 \end{pmatrix} = x^\mu + \bar{\epsilon}\gamma^\mu\theta. \quad (1.8)$$

Note that this is not the most general transformation one could make. For example, a non-linear transformation of the form $x^\mu \mapsto (1 + \epsilon_1\theta_2)x^\mu$ is valid. In addition, one could try to generalize the affect of a vector translation on the spinorial coordinates, but there is no nontrivial linear transformation.

Combining the form of a spinor coordinate transformation with the action of the Poincaré group, one then can construct the full group of isometries on superspace, called the super-Poincaré group. The three dimensional matrix representation from Eq. 1.3 of the Poincaré group can be extended to a five dimensional representation of the super-Poincaré group:

³This is a subtle point. The spinors are Weyl-Majorana and so are real. Thus, one would not be able to absorb imaginary constants and so a priori cannot be ruled out. If one wants $\{Q_1, Q_1\}$ to be real, then a factor of i is required.

$$\begin{pmatrix} \cosh \phi & -\sinh \phi & \exp(-\phi/2)\epsilon_1 & \exp(\phi/2)\epsilon_2 & w_1 \\ -\sinh \phi & \cosh \phi & \exp(-\phi/2)\epsilon_1 & -\exp(\phi/2)\epsilon_2 & w_2 \\ 0 & 0 & \exp(-\phi/2) & 0 & \epsilon_1 \\ 0 & 0 & 0 & \exp(\phi/2) & \epsilon_2 \\ 0 & 0 & 0 & 0 & 1 \end{pmatrix} \text{ acting on } \begin{pmatrix} v \\ \theta \\ 1 \end{pmatrix},$$

which gives rise to the generators of the super-Poincaré group algebra:

$$M = \begin{pmatrix} 0 & i & 0 & 0 & 0 \\ i & 0 & 0 & 0 & 0 \\ 0 & 0 & i/2 & 0 & 0 \\ 0 & 0 & 0 & -i/2 & 0 \\ 0 & 0 & 0 & 0 & 0 \end{pmatrix} \quad E = \begin{pmatrix} 0 & 0 & 0 & 0 & -i \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \end{pmatrix} \quad P = \begin{pmatrix} 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & -i \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \end{pmatrix}$$

$$Q_1 = \begin{pmatrix} 0 & 0 & -i & 0 & 0 \\ 0 & 0 & -i & 0 & 0 \\ 0 & 0 & 0 & 0 & -i \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \end{pmatrix} \quad Q_2 = \begin{pmatrix} 0 & 0 & 0 & -i & 0 \\ 0 & 0 & 0 & i & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & -i \\ 0 & 0 & 0 & 0 & 0 \end{pmatrix}.$$

The explicit form of the generators M, E, P, Q_1 and Q_2 allows for easy computation of the defining relations of the super-Poincaré graded algebra. In particular, the SUSY translations commute with the space-time translations $[P, Q_i] = [E, Q_i] = 0$ and anticommute with each other $\{Q_1, Q_2\} = 0$. Furthermore, SUSY translation operators have a spinor Lorentz structure: $[M, Q_1] = \frac{i}{2}Q_1$ and $[M, Q_2] = -\frac{i}{2}Q_2$. The most important relation is the self anticommutation of the SUSY translations, which yield $\{Q_1, Q_1\} = -2i(E + P)$ and $\{Q_2, Q_2\} = -2i(E - P)$. Heuristically, these self anticommutation relations say that a SUSY translation is the ‘square root’ of spacetime translations.

As a quantum field theory, the fundamental objects in the Standard Model (SM) are the fermionic and bosonic quantum fields. Likewise, in SUSY, quantum fields are the objects governed by the equations of motion. The only difference is that the fields in the SM are maps from Minkowski space, while in SUSY, fields are maps from

superspace. The latter are called *superfields*. A scalar superfield Ψ is a map from superspace into $\mathbb{C}$ which is invariant under a super-Poincaré transformation⁴. Since $\theta_i^2 = 0$, a Taylor expansion of a generic scalar superfield is of the form

$$\Psi(x^\mu, \theta_i) = \phi(x^\mu) + \bar{\theta}\psi(x^\mu) + \theta_1\theta_2 F(x^\mu), \quad (1.9)$$

where ϕ is a scalar, F is a pseudoscalar, and ψ is a Dirac spinor. The field F is a pseudoscalar because under parity, a Dirac spinor S transforms into $\gamma^0 S$, which in the chosen basis means that the two components of θ are interchanged (for a nice discussion, see 8.10 in [21] or 3.6 in [20]). Thus, $\theta_1\theta_2 \mapsto \theta_2\theta_1 = -\theta_1\theta_2$ and since Ψ is a scalar, $\Psi \mapsto \Psi$ under parity. To compensate, $F \mapsto -F$. Under a SUSY translation of $\epsilon^T Q$ on the coordinates of superspace, $\Psi(x^\mu, \theta, \omega) \mapsto \Psi'(x^\mu, \theta, \omega) = \Psi(x^\mu - \bar{\epsilon}\gamma^\mu\theta, \theta - \epsilon)$, where ϵ is a Dirac spinor with components ϵ_1 and ϵ_2 . The minus sign in the expression for Ψ' comes from the fact that a SUSY translation has been defined to shift the coordinates forward and thus the field must compensate by evaluation at a shifted backward location in superspace. A simple computation shows that one can express Ψ in terms of an operator action on Ψ in the following way:

$$\Psi'(x^\mu, \theta, \omega) = \exp(-\bar{\epsilon}\gamma^\mu\theta\partial_\mu - \epsilon_i\partial_{\theta_i})\Psi(x^\mu, \theta), \quad (1.10)$$

with the standard notation $\partial_\mu = (\partial_t, \vec{\nabla})$ and $x^\mu = (t, \vec{x})$ so coordinates are initially given raised while derivatives are all positive when lowered⁵. Taylor expanding the expression for Ψ' gives the form of an infinitesimal SUSY translation along the θ direction:

$$\delta\Psi \equiv \Psi'(x^\mu, \theta, \omega) - \Psi(x^\mu, \theta, \omega) = (-\bar{\epsilon}\gamma^\mu\theta\partial_\mu - \epsilon_i\partial_{\theta_i})\Psi(x^\mu, \theta). \quad (1.11)$$

⁴In this one-dimensional case, the field is real-valued.

⁵Following the convention of [20].

Let

$$\mathcal{Q}_i = i(\gamma^0 \gamma^\mu)_{ij} \theta_j \partial_\mu + i \partial_{\theta_i}, \quad (1.12)$$

so that $\delta\Psi = i\epsilon^\top \mathcal{Q}\Psi$. The factors of i come from the desire to have (anti)Hermitian operators, as is done for the familiar construction of $i\mathcal{P}_\mu = -\partial_\mu$. Similarly, define $i\mathcal{M} = x\partial_t + t\partial_x + \frac{1}{2}\theta_1\partial_{\theta_1} - \frac{1}{2}\theta_2\partial_{\theta_2}$. With these identifications, $\mathcal{Q}$, $\mathcal{P}$, and $\mathcal{M}$ form a representation of the super-Poincaré group through their action on superfields. This can be shown by computing the (anti)commutation relations of the various operators. The momentum operators have the expected Lorentz structure: $[\mathcal{M}, \mathcal{P}_0] = i\mathcal{P}_1$ and $[\mathcal{M}, \mathcal{P}_1] = i\mathcal{P}_0$ and the index on the operators $\mathcal{Q}_i$ is indeed a spinor index, since $[\mathcal{M}, \mathcal{Q}_1] = \frac{i}{2}\mathcal{Q}_1$ and $[\mathcal{M}, \mathcal{Q}_2] = -\frac{i}{2}\mathcal{Q}_2$. As in the case with the matrix representation, $[\mathcal{P}_\mu, \mathcal{Q}_i] = \{\mathcal{Q}_1, \mathcal{Q}_2\} = 0$. The only difference in the defining algebra of the operator versus matrix representation is the self anticommutation relations of the SUSY translations: $\{\mathcal{Q}_1, \mathcal{Q}_1\} = 2i(\mathcal{P}_0 + \mathcal{P}_1)$ and $\{\mathcal{Q}_2, \mathcal{Q}_2\} = 2i(\mathcal{P}_0 - \mathcal{P}_1)$, which differ by a relative minus sign. This sign comes from the fact that the fields compensate for a coordinate change in the opposite way that the coordinates themselves shift and thus the SUSY algebra defined by the operators is not identical to the algebra we encountered earlier.

With the form of a SUSY translation in Equation 1.12 one can compute the changes in the component fields of Ψ as in $\delta\Psi = \delta\phi + \bar{\theta}\delta\psi + \theta_1\theta_2\delta F$. The transformations are

$$\begin{aligned} \delta\phi &= -\bar{\epsilon}\psi \\ \delta\psi &= \gamma^\mu \epsilon \partial_\mu \phi + \gamma^5 \epsilon \\ \delta F &= -\partial_\mu \bar{\psi} \gamma^5 \gamma^\mu \epsilon, \end{aligned} \quad (1.13)$$

where $\gamma^5 \equiv \gamma^0 \gamma^1$. The key feature of Eq. 1.13 and the main result of this section is that the boson ϕ transforms into the spinor ψ and the spinor transforms into a (translated) boson. In this way, SUSY is a symmetry relating bosons and fermions

by transforming one into the other.

1.2 Broken Supersymmetry

Still working in 1+1 dimensions, let Ψ be an irreducible super-Poincaré *multiplet* containing some bosonic and fermionic degrees of freedom. Massive Poincaré multiplets are identified by their mass and spin. This means that a single multiplet can contain only bosonic degrees of freedom or fermionic degrees of freedom, but not both. Consider the operator N_F which is defined by $N_F|\text{boson}\rangle = |\text{boson}\rangle$ and $N_F|\text{fermion}\rangle = -|\text{fermion}\rangle$. This definition is chosen such that the operator trace $\text{Tr}(N_F) = \sum_{\text{boson in } \Psi} \langle \text{boson} | N_F | \text{boson} \rangle + \sum_{\text{fermion in } \Psi} \langle \text{fermion} | N_F | \text{fermion} \rangle$ is simply the number of bosonic degrees of freedom minus the number fermionic degrees of freedom. The operator N_F anti-commutes with Q_i :

$$\begin{aligned} (N_F Q_i + Q_i N_F) |\text{boson/fermion}\rangle \\ = N_F |\text{fermon/boson}\rangle + (+/-) Q_i |\text{boson/fermion}\rangle \\ = (-/+) + (+/-) = 0. \end{aligned} \tag{1.14}$$

Since the trace is linear and has the cyclic property $\text{Tr}(Q_i N_F Q_i) = \text{Tr}(N_F Q_i Q_i) = \text{Tr}(-Q_i N_F Q_i)$ and therefore, this quantity is zero. However, $\text{Tr}(N_F \{Q_i, Q_i\}) = 2(E - P)\text{Tr}(N_F)$. Thus, $\text{Tr} = 0$ and the number of fermions and bosons must be the same in the multiplet. Just as for the Poincaré group, in the super-Poincaré group, $E^2 - P^2$ commutes with all the generators and so the mass is still characterizes a multiplet. This means that in SUSY, every boson has a *superpartner* fermion with the same mass and vice versa. Even though this was derived in 1 + 1 dimensions, it holds for 3 + 1 as well.

While the construction in Sec. 1.1 is elegant, it cannot be true - no superpartners of the SM particles have been observed. Therefore, if is a real symmetry of nature, SUSY must be broken below the energy scales currently accessible to experiments. The next section describes a strong motivation for the SUSY breaking scale to be close to the electroweak energy scale.

1.3 The Hierarchy Problem and Weak-Scale SUSY

One of the fundamental limitations of the SM is that it does not describe gravity. This should not be relevant for physics at the electroweak scale, where the strength of classical gravity is negligible compared with the other forces. However, at energies near the *Planck scale*⁶ $\Lambda \sim 10^{19}$ GeV, gravity will be comparable in strength to the other forces at which point there must be significant contribution from physical laws beyond the SM. The electroweak scale and the Planck scale are theoretically connected by quantum corrections to particle properties. For example, the input mass parameter for a particle in the Lagrangian receives corrections from next-to-leading-order effects encoded by Feynman diagrams like the one shown in Fig. 1.1 for the Higgs boson mass. If the SM is valid up to Λ , then the correction from Fig. 1.1 has the form

$$\delta m \sim -\frac{m_f^2}{v^2} \int_0^\Lambda \frac{d^4 k}{(2\pi)^4} \frac{\not{k} \not{k}}{k^2 k^2} \sim -\frac{m_f^2 \Lambda^2}{v^2}, \quad (1.15)$$

for the vacuum expectation value v and where the minus sign is for the closed fermion loop and results from the difference between Fermi versus Dirac statistics. Similar calculations show that other particles are also sensitive to this *cutoff* scale Λ , but not all quadratically (for fermions, it is $\log(\Lambda)$). For particle masses near the electroweak scale, this seems like an enormous cancellation of $\mathcal{O}(\Lambda/v)$ effects. However, for fermions and gauge bosons, the impact of the corrections is naturally suppressed by symmetry. Gauge invariance ensures that the mass of the photon, gluon, Z and W bosons before electroweak symmetry breaking is exactly zero. Corrections for massless fermions are zero by chiral symmetry⁷. Therefore, the corrections for fermions with a small mass must go to zero as the mass goes to zero. This accounts for all the SM particles except the Higgs boson, which has no symmetry to suppress quantum corrections to the bare mass. This is further complicated because all of the SM masses are tied to the Higgs boson mass after electroweak symmetry breaking. The apparent

⁶This is the energy scale E at which the gravitational potential energy from two objects with mass E/c^2 separated by a distance r is the same as a photon with wavelength r , i.e. $GE^2/rc^2 = \hbar c/r$. At this energy scale, gravitational effects are not small compared to quantum mechanical effects.

⁷Invariance under independent transformations of left- and right-handed fermions.

large cancellation giving rise to the physical Higgs boson mass is called the *hierarchy problem*. Before proceeding, it should be stressed that the hierarchy problem is a formal/aesthetic problem and not a logical inconsistency in the theory. This is in contrast to related problems arising earlier in the history of particle physics such as the non-renormalizability of the Fermi theory of the weak force which had a cutoff at the electroweak scale. However, the hierarchy problem is intriguing/suggestive and continues to be one of the core drivers of model building in high energy physics research.




Figure 1.1: A one-loop diagram contributing to the Higgs boson self energy at next-to-leading order.

One elegant method for eliminating the hierarchy problem is to protect the Higgs boson mass using similar strategies as for the fermions or Gauge bosons. Gauge symmetries do not directly help because there still needs to be a mechanism for generating a non-zero mass. Suppose that there was a new fermion which shared a mass parameter with the Higgs boson. Chiral symmetry would protect the fermion from receiving large quantum corrections to its mass and therefore would indirectly suppress corrections for the Higgs boson. Such a theory was introduced in Sec. 1.1: Supersymmetry. Under SUSY, there is a fermion partner to the Higgs boson (called the *Higgsino*) which is in the same multiplet as the Higgs boson with exactly the same mass. Under exact SUSY, the Higgs boson mass is protected by chiral symmetry.

However, as discussed in Sec. 1.2, SUSY is not exact. There are many mechanisms for breaking SUSY below the electroweak scale such that the SUSY partners to the SM particles are heavier than experimental limits. Before discussing models with SUSY breaking (see Sec. 1.4), consider the largest contributions to Eq. 1.15. Since the

correction scales with the Yukawa coupling of the fermions, the dominant contribution is from top quark loops. The corresponding largest corrections from SUSY therefore needs to come from stop loops, shown in Fig. 1.2. Therefore, if broken SUSY is to provide a solution to the Hierarchy problem, the stop must be relatively light. This can be quantified by limiting the amount of *fine-tuning* [424, 425] required for a SUSY model to reproduce the observed SM spectrum at the electroweak scale. The SM has a large amount of fine-tuning because the input Higgs boson mass parameter in the Lagrangian and the quantum corrections to the Higgs boson mass, each $\mathcal{O}(10^{19})$ GeV, must cancel at one part in 10^{17} to produce the measured 125 GeV Higgs boson mass. There is no unique way to quantify fine-tuning, but there is general consensus that $\mathcal{O}(1\%)$ tuning (suitably defined) requires $m_{\text{stop}} \lesssim 1$ TeV⁸.




Figure 1.2: SUSY NLO corrections to the Higgs boson self energy. While the right diagram is topologically the same as the leading NLO fermion diagram in Fig. 1.1, it is suppressed with respect to the left diagram due to the two scalar propagators ($\sim \int d^4k/k^4$) compared with one ($\sim \int d^4k/k^2$). There are also two powers of the coupling constant for the right diagram, but the Yukawa coupling is nearly one for the top/stop.

There is other indirect evidence for weak scale SUSY in addition to solving the Hierarchy problem. For example, any of the neutral SUSY particles could make a natural dark matter candidate due to the *Weakly Interacting Massive Particle (WIMP) miracle*. Define $\Omega_\chi = \rho_\chi/\rho_{\text{critical}}$ as the normalized mass density of dark matter particle χ . The critical density ρ_{critical} is derived from the Friedmann metric for a flat, homogeneous, and isotropic universe and is given by $\rho_{\text{critical}} = 3H^2/(8\pi G) \sim 10(\text{GeV}/c^2)/\text{m}^3$,

⁸Early references include Ref. [426, 427] and this is an area of active research - see for instance Ref. [428–433]

where $H \sim 100(\text{km/s})/\text{Mpc}$ [434] is the Hubble constant. The equation of motion for the number density of a dark matter particle χ is given by the Boltzman equation:

$$\dot{n}_\chi + 3n_\chi \dot{a}/a = -\langle\sigma v\rangle (n_\chi^2 - n_{\chi,\text{equilibrium}}^2), \quad (1.16)$$

where the dot denotes a derivative with respect to time and a is the scale factor of the universe ($H = \dot{a}/a$). The lefthand side of Eq. 1.16 is the equation for an expanding universe with constant mass; the factor of 3 simply results in the usual $n_\chi \propto 1/a^3$. The righthand side of Eq. 1.16 accounts for creation and annihilation of χ where $\langle\sigma v\rangle$ is the thermally averaged annihilation cross section multiplied by the relative speed. In the early universe when the temperature was very high, $k_B T \gg m_\chi c^2$, pairs of dark matter particles were constantly being created and destroyed. When the temperature dropped below m_χ , the dark matter particles no longer annihilated and so the density was fixed at $n_{\chi,\text{equilibrium}}$, giving rise to the relic density Ω_χ observed today. The solution to the equilibrium number density from Eq. 1.16 is given by [435]

$$n_\chi \sim s_0 10^{-8} \left[\left(\frac{m_\chi}{\text{GeV}} \right) \left(\frac{\langle\sigma v\rangle}{10^{-27} \text{cm}^3/\text{s}} \right) \right]^{-1}, \quad (1.17)$$

where s_0 is the current entropy density of the universe. For the highly relativistic particles contributing to the entropy density, the only dimensionful number is the temperature T (by definition, the particle masses are irrelevant) and so $s \propto T^3$. The exact form is $s = 2\pi^2 g(T) T^3/45$, where $g(T)$ is the number of effective degrees of freedom⁹. The current temperature of the universe is about 3 K at which basically only photons and neutrinos contribute to $g(T) \sim 4$. In units of cm^3 (with units in which $k_b = 1$),

$$s_0 \sim 2T^3 \sim 2(3\text{K})^3 \times \left(\frac{1\text{eV}}{10^4\text{K}} \right) \times \left(\frac{1}{2 \times 10^{-7}\text{m}} \right) \sim 4000\text{cm}^{-3}. \quad (1.18)$$

⁹See Ref. [436] for a pedagogical explanation.

Using $\Omega_\chi = m_\chi n_\chi / \rho_{\text{critical}}$,

$$\Omega_\chi h^2 \sim (4 \times 10^{-27} \text{cm}^3 \text{s}^{-1} / \langle \sigma v \rangle) . \quad (1.19)$$

The total dark matter relic density has been measured to be $\Omega h^2 \sim 0.1$ [27]. The “WIMP miracle” is that the cross-section for a weak-scale interaction is about $\alpha_{\text{weak}}^2 / m_{\text{weak}}^2$ and for $\alpha_{\text{weak}} \sim 0.01$ and $m_{\text{weak}} \sim 100$ GeV, Eq. 1.19 is the same order of magnitude as the measurement.

Two other related sources of indirect motivation for electroweak SUSY are grand unification and successful electroweak symmetry breaking. An intriguing curiosity of a minimal SUSY extension of the SM (see Sec. 1.4) is that the three gauge group coupling constants seem to be equal to each other at a high energy and this *grand unification* (GUT) scale is near the Plank scale. As described in Sec. 1.1, the running of the coupling constant g is described by solutions to the Callan-Symanzik equation:

$$\frac{dg}{d \log(Q/M)} = \beta(g), \quad (1.20)$$

where M is a fixed energy scale (such as m_Z) and Q is the running energy scale. At leading order, $\beta(g) = b_0 g^3 / (4\pi)^2$, with $b_0 = \sum_{\text{fields } f} \kappa_f C(G, r_f)$, where¹⁰ $f \in \{\text{gauge, Weyl fermion, scalar}\}$, $\kappa_f = -11/3$ for gauge fields, $2/3$ for Weyl fermions, and $1/3$ for scalars. The factors $C(G, r_f)$ depend on the gauge group G as well as the representation of the field r_f . In the adjoint representation, $C(G) = C_2(G)$, the quadratic Casimir operator of group that is N for $SU(N)$, $N > 1$ and 0 for $U(1)$. In the fundamental representation of $SU(N)$, $C(G) = \frac{1}{2}$ and for $U(1)$, $C(G) = Y^2$, where Y is the weak hypercharge. For $SU(3)$, the Higgs does not contribute and there is no

¹⁰Somewhat surprisingly, the full derivation of these factors is not usually presented all at once in the main QFT texts. Most advanced QFT students will have derived the equations for QED and QCD, but there is a small jump to the general $U(1)$ from QED (also to include complex scalars) - see for instance Chapter 66 in Ref. [437]. With some careful thought, the inclusion of complex scalars in the non-Abelian case can be extracted using the results of the background field method presented in Chapter 16.6 in Ref. [20].

distinction between left and right handed fields, so the Weyl fermion combine to give the familiar equation

$$b_0 = -\frac{11}{3}N_c + \frac{2}{3}n_f, \quad (1.21)$$

where N_c is the number of colors and n_f is the number of quarks. For the three gauge couplings of the SM, $g_1 = e/\cos(\theta_W)$, $g_2 = e/\sin(\theta_W)$, g_3 corresponding to the gauge groups $U(1)$, $SU(2)$, and $SU(3)$, the three leading order β functions are

$$\begin{aligned} b_0^1 &= \frac{2}{3} \times (2Y_{E_L}^2 + Y_{e_R}^2 + 2N_c Y_{Q_L}^2 + N_c Y_{u_R}^2 + N_c Y_{d_R}^2) \times 3 + \frac{1}{3} \times 2Y_H^2 = \frac{41}{6} \\ b_0^2 &= -\frac{11}{3} \times 2 + \frac{2}{3} \times \frac{1}{2} \times 12 + \frac{1}{3} \times \frac{1}{2} = -\frac{19}{6} \\ b_0^3 &= -\frac{11}{3} \times 3 + \frac{2}{3} \times \frac{1}{2} \times 12 = -7, \end{aligned} \quad (1.22)$$

where the terms in red are the κ factors, the terms in blue are the $C(G, r)$ factors and the remaining terms count the number of fields. For instance, there are 12 Weyl quarks (= 6 Dirac fermion quarks) and 12 total left-handed fields contributing to b_0^2 (each quark type contributes three times, one for each color). The weak-hypercharges are $Y_{E_L} = -\frac{1}{2}$, $Y_{e_R} = 1$, $Y_{Q_L} = \frac{1}{6}$, $Y_{u_R} = -\frac{2}{3}$, and $Y_{d_R} = \frac{1}{3}$, where $Y = Q - T_3$ (T_3 is the weak isospin). For one copy of SUSY added to the SM, the gauge bosons have fermionic partners in the adjoint representation and the fermions have complex scalar partners in the fundamental representation. Therefore, instead of $-11/3C$ in Eq. 1.22, in SUSY the contribution is $(-11/3 + 2/3)C = -3C$. Likewise for the fermions (and the complex scalars), instead of $2/3C$ (or $1/3$), in SUSY the contribution is $(2/3 + 1/3)C = C$. Therefore, the leading order β functions become

$$\begin{aligned}
b_0^1 &= 1 \times (2Y_{E_L}^2 + Y_{e_R}^2 + 2N_c Y_{Q_L}^2 + N_c Y_{u_R}^2 + N_c Y_{d_R}^2) \times 3 + 1 \times 2Y_H^2 \times 2 = 11 \\
b_0^2 &= -3 \times 2 + 1 \times \frac{1}{2} \times 12 + 1 \times \frac{1}{2} \times 2 = 1 \\
b_0^3 &= -3 \times 3 + 1 \times \frac{1}{2} \times 12 = -3,
\end{aligned} \tag{1.23}$$

where the extra factor of two for the Higgs fields is due to a second Higgs doublet that is required in the minimal SUSY extension of the SM (see Sec. 1.4). Figure 1.3 shows the one-loop running of the three gauge group constants, $\alpha_i = g_i^2/4\pi$. Conveniently, for $t = \log(Q/M)$,

$$\frac{d\alpha^{-1}(t)}{dt} = -\frac{1}{\alpha^2} \frac{d\alpha}{dt} = \frac{g}{2\pi\alpha^2} \frac{dg}{dt} = \frac{b_0 g^4}{2\pi\alpha^2 (4\pi)^2} = -\frac{b_0}{2\pi}, \tag{1.24}$$

i.e. the inverse coupling depends linearly on $\log(Q/M)$. The $U(1)$ coupling in Fig. 1.3 is scaled by $\sqrt{5/3}$ as predicted by grand unified theories¹¹ such as the $SU(5)$ theory of Georgi and Glashow [439]. The PDG values of $\alpha_{EM}^{-1}(m_Z) = 127.916 \pm 0.015$, $\sin^2(\theta_W)(m_Z) = 0.23116 \pm 0.00013$, and $\alpha_s(m_Z) = 0.1184 \pm 0.0007$ are used as the initial condition (the error bands are too small to see) [440]. In the SM, the three couplings do not unify at any scale, but amazingly for the MSSM, there is a point around $Q = 10^{16}$ GeV where all three couplings are the same within uncertainties. There are some changes to this picture by including higher order corrections, but the prospect of unification is unchanged.

Related to supersymmetric grand unification is the successful breaking of electroweak symmetry. In order for the Higgs potential to have a local minimum and thus a positive vacuum expectation value, the Higgs boson mass squared in the Lagrangian must be negative. In grand unified SUSY theories where the Higgs mass is set to a positive value at the grand unified scale, the renormalization group flow down to the electroweak scale drives down the Higgs boson mass squared and over a large

¹¹Chapter 8.3 in Ref. [438] has a simple explanation of this factor and a more detailed approach can be found in e.g. Chapter 97 of Ref. [437].




Figure 1.3: The one-loop renormalization group evolution of the inverse couplings as a function of energy. The $U(1)$ coupling is scaled by $\sqrt{5/3}$, which is the quantity that unifies with the others in grand unified theories.

range of parameter space is negative due to the large top quark Yukawa coupling¹². Therefore, GUT SUSY can *explain* why electroweak symmetry is broken [441, 442]. Another impact of the large top Yukawa coupling for RGE flow in SUSY grand unified theories is that the stop is often the lightest squark near the electroweak scale, even if all the scalar masses are unified at the GUT scale (see Sec. 1.4).

The next section describes a complete model of weak-scale SUSY called the *Minimal Supersymmetric Standard Model* (MSSM) and will be the default SUSY model discussed for the remainder of Part III.

¹²In the minimal SUSY extension of the SM, there are two Higgs boson doublets, one associated with up type quarks and one associated with down type quarks (see Sec. 1.4). As the top quark is up-type, the associated Higgs mass squared is the one usually driven negative.

1.4 The Minimal Supersymmetric Standard Model

There are many ways to extend the SM with SUSY. For example, one could augment spacetime with multiple copies of the fermionic dimensions ($N > 1$ SUSY) or add additional SUSY multiplets beyond those that match the SM fields (for one additional scalar, this is the NMSSM). However, the focus of this section and much of the SUSY literature is the *minimal* SUSY extension to the SM (MSSM) that has one chiral multiplet¹³ for each SM fermion and one vector multiplet for each gauge boson. One new multiplet in the MSSM with respect to the SM is a second Higgs field. The main point of SUSY was to add a fermionic partner to the Higgs so that the mass would be protected by chiral symmetry. However, all electroweakly interacting fermions contribute to the Feynman diagrams in Fig. 1.4 which generate an anomaly in the SM¹⁴: if this diagram does not exactly vanish, then $U(1)$ symmetry is violated beyond leading order in perturbation theory. The matrix element from Fig. 1.4 is proportional to Y^3 , which amazingly sums to zero in the SM:

$$\begin{aligned} \mathcal{M} &\propto \sum_{\text{left-handed}} Y^3 - \sum_{\text{right-handed}} Y^3 \\ &= 2Y_{E_L}^3 - Y_{e_R}^3 + 2N_c Y_{Q_L}^3 - N_c Y_{u_R}^3 - N_c Y_{d_R}^3 = 0 \end{aligned} \quad (1.25)$$

Since the superpartners of the SM particles all have the same hypercharge, the sum in Eq. 1.25 remains zero. When only one fermionic Higgs partner is added, the anomaly will not vanish. This is solved by simply adding a second Higgs field with *opposite* hypercharge. Table 1.1 summarizes the complete field content of the MSSM. There are 17 chiral supermultiplets, each containing one fermion and one complex scalar, and three vector supermultiplets, each containing a vector boson and a fermion. The unbroken MSSM Lagrangian is the same (in form and number of parameters) as the SM case except for the Higgs and the lepton/baryon number violating sectors. The

¹³The procedure for constructing a scalar superfield introduced in Sec. 1.1, namely Taylor expanding a field with certain transformation properties can be generalized to form chiral and vector superfield - see e.g. Chapter 4 in Ref. [443].

¹⁴See for instance Chapter 20.2 in Ref. [20].




Figure 1.4: The Feynman diagram that is the source of the chiral anomaly in the SM. Any electroweakly interacting fermion contributes to the loop.

Higgs part of the Lagrangian is given by

$$\mathcal{L}_{\text{Higgs}}^{\text{MSSM}} = (y_U)_{ij} Q_i H_u u_j^c + (y_D)_{ij} Q_i H_d d_j^c + (y_L)_{ij} L_i H_d e_j^c + \mu H_u H_d, \quad (1.26)$$

where the group indices are suppressed, y_U, y_D , and y_L are the SM Yukawa mass matrices and μ is a new term that is allowed by $SU(2)$ symmetry. Under electroweak symmetry breaking, the $\mathcal{L}_{\text{Higgs}}^{\text{MSSM}}$ behaves similarly to the SM case (ignoring the μ term¹⁵), except masses are generated separately for up-type quarks by H_u and for down-type quarks and leptons by H_d ¹⁶. In addition to the Higgs sector, there are a set of terms allowed by all of the internal symmetries, but explicitly violate low energy effective symmetries of the SM (lepton and baryon number conservation):

$$\mathcal{L}_{\text{RPV}}^{\text{MSSM}} = \lambda_{ijk} L_i L_j e_k^c + \lambda'_{ijk} L_i Q_j d_k^c + \lambda''_{ijk} u_i^c d_j^c d_k^c + \kappa_i L_i H_u, \quad (1.27)$$

where the λ and κ terms are new dimensionless parameters. While a priori there is no symmetry which forbids $\mathcal{L}_{\text{RPV}}^{\text{MSSM}}$, it has significant phenomenological consequences. Most importantly, if the $\lambda_i \neq 0$ ¹⁷, the proton could rapidly decay even though the

¹⁵The μ term actually introduces a fine-tuning problem - its value is arbitrary, yet needs to be near the electroweak scale. There is a large literature on this subject - see any papers which cite the earliest ideas: Ref. [444, 445].

¹⁶See e.g. Chapter 10 in Ref. [438] or Sec. 8.1 in Ref. [443] for slightly more information and Ref. [446] for extensive details.

¹⁷Technically, two of the λ need to be nonzero for proton decay. However, there are other constraints and issues of naturalness if only one of the $\lambda \neq 0$.

experimental lifetime is greater than 10^{33} years [447] (see e.g. Sec. 6.2 in Ref. [443]). One process contributing to proton decay is illustrated in Fig. 1.5. The standard assumption to remove $\mathcal{L}_{\text{RPV}}^{\text{MSSM}}$ is to impose a new $\mathcal{Z}_2$ symmetry called R-parity [448] (RPV = R-parity violation) under which the SM particles are neutral and the SUSY partners are charged. Symbolically, the R-charge of a product of fields $F_1 F_2 \cdots F_n$ is given by

$$R(F_1 F_2 \cdots F_n) = (-1)^{\sum_i 3B_i + L_i + 2s_i}, \quad (1.28)$$

where B_i , L_i , and s_i are the baryon number, lepton number, and spin of field F_i . Requiring conservation of R-charge has many important phenomenological consequences. First of all, there is no baryon or lepton number violation at tree level in the MSSM. Second, at collider experiments with a SM-only initial state, SUSY particles must be produced in pairs. The lightest SUSY particle (LSP) must be stable because it cannot decay into only SM particles. When combined with the WIMP miracle, this last property makes the LSP an attractive dark matter candidate particle. For the remainder of Part III, R-parity is assumed conserved.




Figure 1.5: The Feynman diagram illustrating proton decay with RPV couplings.

There are many ways to break SUSY in the MSSM, but they all involve additional model assumptions. To avoid making specific model assumptions, consider the Lagrangian of the MSSM augmented with terms that explicitly violate SUSY, $\mathcal{L}_{\text{SUSY}}^{\text{MSSM}}$. Terms are only allowed if they do not *reintroduce* the hierarchy problem and preserve all other symmetries (*softly broken SUSY*). Since the quadratic divergences giving rise to the hierarchy problem are associated with the dimensionless Yukawa couplings

Field	SM component	SM spin	SUSY partner	Partner spin	U(1)	SU(2)	SU(3)	Comment
Q_i	$(u_L \ d_L)$	1/2	$(\tilde{u}_L \ \tilde{d}_L)$	0	$\frac{1}{6}$	2	3	3 generations
u_i^c	u_R^c	1/2	$\tilde{u}_R^\dagger$	0	$\frac{2}{3}$	1	$\bar{\mathbf{3}}$	3 generations
d_i^c	d_R^c	1/2	$\tilde{d}_R^\dagger$	0	$-\frac{1}{3}$	1	$\bar{\mathbf{3}}$	3 generations
L_i	$(e_L \ \nu_L)$	1/2	$(\tilde{e}_L \ \tilde{\nu}_L)$	0	$\frac{1}{2}$	2	1	3 generations
e_i^c	e_R^c	1/2	$\tilde{e}_R^\dagger$	0	-1	1	1	3 generations
H_u	$(H_u^+ \ H_u^0)$	0	$(\tilde{H}_u^+ \ \tilde{H}_u^0)$	1/2	$\frac{1}{2}$	2	1	
H_d	$(H_d^0 \ H_d^-)$	0	$(\tilde{H}_d^0 \ \tilde{H}_d^-)$	1/2	$-\frac{1}{2}$	2	1	
B	B	1	$\tilde{B}$	1/2	0	1	1	
W	$W^\pm, W^0$	1	$\tilde{W}^\pm, \tilde{W}^0$	1/2	0	3	1	
G	g	1	$\tilde{g}$	1/2	0	1	8	

Table 1.1: A summary of the MSSM field content in terms of electroweak eigenstates before symmetry breaking. Even though the SUSY partners of the left- and right-handed fermions are scalars, they still carry the L or R subscript to emphasize their relationship to the SM particles. The superpartners of the fermions are called *sfermions* (squarks and sleptons) and the superpartners of the bosons are called *bosinos* (bino, wino, gluino, and higgsino). The right-handed field are specified in terms of the charge conjugate of left-handed fields because chiral superfields only contain left-handed fermions.

($y \propto m/v$ as in Eq. 1.15), the Hierarchy problem can be avoided by omitting dimensionless interactions in $\mathcal{L}_{\text{SUSY}}^{\text{MSSM}}$. Fermion mass terms are not allowed by SU(2), but mass terms for the complex scalars and fermionic partners of the gauge bosons are allowed¹⁸. The full soft SUSY breaking Lagrangian is given by

$$\mathcal{L}_{\text{SUSY}}^{\text{MSSM}} = \frac{1}{2} M_i G_i^2 + m_{\Phi, ij}^2 \tilde{\Phi}_i^\dagger \tilde{\Phi}_j + m_{\tilde{\Phi}, ij}^2 \tilde{\Phi}_{R, i}^\dagger \tilde{\Phi}_{R, j} + A_{ijk} \tilde{\Phi}_i \tilde{\Phi}_j \tilde{\Phi}_k + \text{h.c.}, \quad (1.29)$$

where $G_i \in \{\tilde{g}, \tilde{W}, \tilde{B}\}$, $\tilde{\Phi} \in \{\tilde{Q}, \tilde{L}, \tilde{H}\}$ for e.g. $\tilde{Q} = (\tilde{u}_L \ \tilde{d}_L)$, $\tilde{\Phi} \in \{\tilde{u}, \tilde{d}, \tilde{e}\}$, and all group indices are suppressed. The mass terms in Eq. 1.29 allow the SUSY partners to have

¹⁸For a longer explanation, see Chapter 9.2 in Ref. [438].

a mass much higher than and unrelated to the SM fermions and bosons which acquire a mass through EWSB. In particular, the partners of the left- and right-handed SM fields have different soft SUSY masses and therefore can have significantly different masses. In total, the full softly broken MSSM (from now on, *this* will be called the MSSM) has 105 new parameters with respect to the SM [449]. Many of these terms are highly constrained by current experiments. For example, the off-diagonal terms in the mass matrix induce large neutral flavor changing processes ruled out by flavor physics experiments. However, the MSSM still has an enormous parameter space. There is a vast literature of SUSY models that make various predictions for the relationships between parameters. One well-studied set of models is the Constrained Minimal Supersymmetric Standard Model (CMSSM) [448, 450–453] in which the particle masses in addition to the gauge couplings unify at a GUT scale. In particular, at the GUT scale the scalar supersymmetric particles have the same mass m_0 , the gauge fermion supersymmetric particles have the mass $M_{1/2}$ and the trilinear scalar couplings are given by a new parameter A_0 multiplied by the corresponding Standard Model Yukawa matrices. The only other required input to fully specify the full MSSM is the ratio of the Higgs' vacuum expectation values $\tan \beta$ and the sign of the Higgsino mass term $\text{sign}(\mu)$. The value of μ is set by requiring the calculated Z^0 mass is equal to the measured value. Thus, the CMSSM has only five more parameters than the SM, far fewer than the full MSSM. The SUSY particle spectrum at any given scale is then determined by solving the RGEs with boundary conditions at the three scales: GUT, SUSY breaking¹⁹, and electroweak. The standard is fixed point iteration [455–458]. Figure 1.6 shows an example calculation, running the CMSSM GUT scale parameters down to the electroweak scale. By construction, the gauge boson masses are equal to $m_{1/2} = 500$ GeV at the GUT scale, which is just beyond 10^{16} GeV. Successful EWSB is a *prediction* of this model and the stop is generally lighter than the other sfermions at the electroweak scale. The RGEs for the soft SUSY breaking masses M_i , $i = 1, 2, 3$ are similar to the equations for the gauge couplings discussed earlier (Eq. 1.24). In particular, at leading order [443],

¹⁹This intermediate scale is used because the radiative corrections associated with EWSB are smallest [454].

$$\frac{dM_i}{dt} = -\frac{b_i}{2\pi}\alpha_i M_i, \quad (1.30)$$

where b_i was defined in Eq. 1.23 for the MSSM. Amazingly,

$$\frac{d}{dt} \left(\frac{M_i}{\alpha_i} \right) = \frac{1}{\alpha_i} \frac{dM_i}{dt} + M_i \frac{d\alpha_i^{-1}}{dt} = 0, \quad (1.31)$$

which means that this ratio does not run with scale, at one-loop. In the CMSSM and in any other model where the gauge masses unify at the GUT scale, this gives a concrete prediction for the mass hierarchy in the MSSM. Using the input parameters at m_Z from earlier,

$$\frac{M_1}{M_2} = \frac{\alpha_1}{\alpha_2} = \frac{5 \sin^2 \theta_W}{3 \cos^2 \theta_W} \sim 0.5 \quad (1.32)$$

$$\frac{M_3}{M_2} = \frac{\alpha_s \sin^2 \theta_W}{\alpha_{EM}} \sim 3.5, \quad (1.33)$$

where the factor of 5/3 is assuming that the gauge unification happens with the GUT scaling mentioned in Sec. 1.3. This gives the famous ratio $M_3 : M_2 : M_1 \sim 7 : 2 : 1$ at the electroweak scale and the expectation that the gluino mass is higher than the mass of the electroweak superpartners.

Specifying parameters at the GUT scale and at the electroweak scale is a powerful technique for reducing the number of input parameters, but it also creates challenges. In particular, the five parameters of the CMSSM are actually not enough to uniquely define an electroweak scale SUSY spectrum - there can be multiple solutions to the RGE equations [459, 460]. Figure 1.7 illustrates the presence of these multiple solutions, some of which have significantly different phenomenology. This loophole may allow CMSSM-type models to evade current limits, though all of the known extra spectra have similar masses to the previous spectra and so inclusive search results should be largely unaffected.

The CMSSM was the main set of models used by experiments at LEPP, the Tevatron, and the early part of LHC Run 1 for designing and interpreting experimental searches. However, it has largely fallen out of favor because it is *too constrained* and many physical parameters such as the measured Higgs mass are not predicted correctly.




Figure 1.6: An example CMSSM model specified with sfermion mass $m_0 = 400$ GeV, gaugino mass $m_{1/2} = 500$ GeV, zero trilinear couplings, $\tan(\beta) > 0$ and a positive μ . The running of the masses is calculated at NLO using SOFTSUSY 3.7.1 [455].

The opposite extreme to the CMSSM is an approach where parameters are only specified near the electroweak scale. One class of such models is called the phenomenological MSSM (pMSSM) [461, 462], as it reduces the number of MSSM parameters by imposing reasonable phenomenological constraints. In particular, by requiring pMSSM models to not introduce non-SM sources of CP violation, lack flavor changing neutral currents, and have degenerate first and second generations, the total number of parameter is reduced to 19. Various groups have performed scans in (subsets of) this 19 parameter space to identify regions of the pMSSM that are also consistent with SM measurements and SUSY searches. For example, Fig. 1.8 shows one part of a pMSSM model with a light stop that is not ruled out by the direct stop searches, but




Figure 1.7: Left: the number of electroweak scale spectra consistent with the CMSSM parameters as a function of m_0 and $M_{1/2}$ for fixed $A_0 = 0, \tan\beta < 0$ and $\mu < 0$. Right: the difference in select phenomenological parameters along the strip of two solutions from the left plot just below where A^0 is tachyonic. Sparticle masses are nearly identical between the two spectra, but stop and sbottom branching ratios vary by more than 10% and the predicted dark matter relic density differs by more than 100% for $M_{1/2} \sim 840$ GeV. See Ref. [460] for more detail.

is excluded by searches with a more inclusive scope due to the complexity of the final state. While there is an inherent bias in these scans due to the choice of parameter priors, they are useful for identifying a class of ‘realistic’ models.

The 19 parameters of the pMSSM is still too large for most practical purposes. Currently, the most popular approach is to focus on specific topologies or *simplified models* and *ignore* the rest of the spectrum, assuming it is largely decoupled or at least factorized from the process of interest [468–470]. Simplified models are useful for organizing searches based on experimental signatures instead of unobservable theoretical parameters. In addition, searches based on simplified models can easily be reinterpreted in any model that has a simplified model-like component. To construct simplified models, it is useful to recast the MSSM fields in terms of the mass eigenstates instead of the weak eigenbasis (though this is not particular to simplified models). After electroweak symmetry breaking, the mass matrix for the neutral




Figure 1.8: One model from a recent ATLAS summary [463] of SUSY searches interpreted in the context of a random scan in the pMSSM [462, 464–466]. This particular model is not excluded by direct stop searches but is ruled out by other searches with a broader scope. One reason the direct searches do not rule out this model, and one powerful use of the pMSSM, is the model complexity: there are multiple light neutralinos/charginos with cascade decays. This plot was created with PySLHA [467]. All particles not shown are heavier than 700 GeV.

electroweak superpartners is given at leading order by (e.g. Sec. 8.2 in Ref. [443]):

$$M_N = \begin{pmatrix} M_1 & 0 & -c_\beta s_W m_Z & s_\beta s_W m_Z \\ 0 & M_2 & c_\beta c_W m_Z & -s_\beta c_W m_Z \\ -c_\beta s_W m_Z & c_\beta c_W m_Z & 0 & -\mu \\ s_\beta s_W m_Z & -s_\beta c_W m_Z & -\mu & 0 \end{pmatrix}, \quad (1.34)$$

where $s_x = \sin(x)$ and $c_x = \cos(x)$. The diagonalization of Eq. 1.34 results in the mass matrix for the four *neutralinos* $\tilde{\chi}_i^0$, $i = 1, \dots, 4$ with the convention $m_{\tilde{\chi}_i^0} < m_{\tilde{\chi}_{i+1}^0}$. The lightest neutralino is an excellent dark matter candidate as it is stable if it is the LSP

(assumed henceforth) and only interacts via the weak force²⁰. When $m_Z \ll |\mu \pm M_i|$, $i = 1, 2$, the neutralinos are nearly pure *binos*, *winos*, and *higgsinos*. In such a case, one may expect the LSP to be mostly bino-like or higgsino-like (assuming $M_1 < M_2$ as in the GUT-inspired scenario). Similarly, there is a two-by-two matrix for the charged electroweak superpartners that forms the two electrically positive and two negative *charginos* $\tilde{\chi}_i^\pm$, $i = 1, 2$. All of the scalar sfermions can also mix to form the mass eigenstates. The most important is the stop mass matrix²¹:

$$M_t^2 = \begin{pmatrix} m_{\tilde{Q},33}^2 + m_{\text{top}}^2 + \left(\frac{1}{2} - \frac{2}{3}s_W^2\right) m_Z^2 c_{2\beta} & m_{\text{top}}(A_{\tilde{t}_R \tilde{Q}_3 \tilde{H}_u} - \mu \cot(\beta)) \\ m_{\text{top}}(A_{\tilde{t}_R \tilde{Q}_3 \tilde{H}_u} - \mu \cot(\beta)) & m_{\tilde{q},33}^2 + m_{\text{top}}^2 + \frac{2}{3}s_W^2 m_Z^2 c_{2\beta} \end{pmatrix}. \quad (1.35)$$

The stop mixing angle θ_t is defined as the angle of the rotation matrix required to diagonalize Eq. 1.35. After diagonalizing the fields, the two stop mass eigenstates are called $\tilde{t}_1$ and $\tilde{t}_2$ with $m_{\tilde{t}_1} < m_{\tilde{t}_2}$. In the literature, $X_t = A_{\tilde{t}_R \tilde{Q}_3 \tilde{H}_u} - \mu \cot(\beta)$ is often called the effective mixing parameter, as it controls the amount of mixing between the weak eigenstates in Eq. 1.35.

The main motivation of electroweak scale SUSY was the cancellation of quantum corrections to the Higgs boson mass. In the MSSM, the lightest Higgs boson mass is not a free parameter; at tree level, it is given by

$$m_h^2 = \frac{1}{2} \left(m_{A^0}^2 + m_Z^2 - \sqrt{(m_{A^0}^2 - m_Z^2)^2 + 4m_Z^2 m_{A^0}^2 s_{2\beta}^2} \right), \quad (1.36)$$

where A^0 is the pseudoscalar Higgs boson generated after electroweak symmetry breaking by the scalar part of the Higgs field and has mass $m_{A^0}^2 = 2\mu^2 + m_{H_u}^2 + m_{H_d}^2$.

²⁰See Ref. [471] for an argument why any of the electrically or color charged particles would not make good dark matter candidates. The relic abundance of SUSY LSP dark matter depends on the field content of the lightest neutralino - mostly higgsino and wino LSP dark matter tends to overproduce and mostly bino LSP tends to underproduce the measured density (see e.g. the review Ref. [435] and references therein).

²¹In principle, there can be mixing between the sfermion families, but this is assumed negligible due to the often unacceptable flavor changing neutral currents.

It seems that in SUSY, the Higgs mass is actually *too well regulated*:

$$m_h^2 \leq \lim_{m_{A^0} \rightarrow \infty} m_h^2(m_{A^0}) = m_Z^2(1 - s_{2\beta}^2) = m_Z^2 \cos^2(2\beta) \leq m_Z^2. \quad (1.37)$$

If the tree-level calculation was (close) to the full answer, the MSSM would be ruled out by the fact that $m_h \approx 125 \text{ GeV} > m_Z$. Fortunately, the corrections to the Higgs boson mass are not small. The dominante correction comes from loops with stops²²:

$$\Delta(m_h^2) = \frac{3}{4\pi^2} c_\beta^2 y_t^2 m_t^2 \left[\ln \left(\frac{m_{\tilde{t}_1} m_{\tilde{t}_2}}{m_t^2} \right) + \frac{\Delta_{\text{mix}}}{m_t^2} \right] + \mathcal{O} \left(\frac{m_t^2 m_Z}{m_{A^0}}, \frac{m_t^2 y_t^4}{(4\pi^2)^2} \right), \quad (1.38)$$

where the term Δ_{mix} goes to zero when $\theta_{\tilde{t}} \rightarrow 0$:

$$\Delta_{\text{mix}} = c_\beta^2 s_\beta^2 (m_{\tilde{t}_2}^2 - m_{\tilde{t}_1}^2) \ln \left(\frac{m_{\tilde{t}_2}^2}{m_{\tilde{t}_1}^2} \right) + \frac{c_\beta^4 s_\beta^4}{m_t^2} \left[(m_{\tilde{t}_2}^2 - m_{\tilde{t}_1}^2)^2 - \frac{1}{2} (m_{\tilde{t}_2}^4 - m_{\tilde{t}_1}^4) \ln \left(\frac{m_{\tilde{t}_2}^2}{m_{\tilde{t}_1}^2} \right) \right]. \quad (1.39)$$

Even with the correction in Eq. 1.38, it is difficult to accommodate a 125 GeV Higgs boson. Figure 1.9 shows the maximum value of Eq. 1.38 without stop mixing and with mixing that maximizes the Higgs mass correction. In order for the Higgs mass to be heavy enough, one or both of the stops have to be relatively heavy ($m \gtrsim 1 \text{ TeV}$). In order for one of the stops to be light ($m \lesssim 1 \text{ TeV}$), there must be significant stop mixing.

In addition to setting the mass hierarchy, stop mixing has a significant impact on stop decay. The stop can decay via a neutral current $\tilde{t} \rightarrow t + \tilde{\chi}^0$ or charged current $\tilde{t} \rightarrow b + \tilde{\chi}^\pm$. If the lightest chargino is heavier than $\tilde{t}_1$, then the neutral current decay dominates unless $m_{\tilde{t}} - m_{\tilde{\chi}_1^0} \ll m_t, m_W$ so that loop-induced processes can compete with the off-shell top/ W boson in the decay. Since the coupling between the stop

²²Due to its importance, there is an extensive literature on this calculation, see e.g. Ref. [472] for a review. This equation is based on Eq. 8.1.24 in Ref. [443], which is a nicer version of Eq. 62 in Ref. [472].




Figure 1.9: The maximum Higgs mass $\sqrt{m_Z^2 c_{2\beta}^2 + \Delta(m_h^2)}$ from Eq. 1.39 for no stop mixing $\theta_{\tilde{t}} = 0$ (left) and mixing that maximizes the correction (right).

and the bino is via the weak hypercharge, by Table 1.1, the coupling is stronger to $\tilde{t}_R$ than to $\tilde{t}_L$. Therefore, $\mathcal{BR}(\tilde{t}_R \rightarrow t + \tilde{\chi}_1^0)/\mathcal{BR}(\tilde{t}_L \rightarrow t + \tilde{\chi}_1^0) > 1$ for a mostly bino LSP. In fact, if the chargino is mostly a wino and $\theta_{\tilde{t}} \sim \pi/2$ (i.e. $\tilde{t}_1 \sim \tilde{t}_R$), then the $\mathcal{BR}(\tilde{t}_1 \rightarrow t + \tilde{\chi}_1^0) \approx 100\%$ regardless of the mass of $\tilde{\chi}_1^\pm$ since the superpartner of the right handed top does not couple to $\tilde{W}^\pm$. In general, the partial widths of the two processes are determined by $\theta_{\tilde{t}}$, $m_{\tilde{t}}$, $m_{\tilde{\chi}_1^\pm}$, $m_{\tilde{\chi}_1^0}$, and the neutralino/chargino mixing matrices which depend on μ , M_1 , M_2 , and $\tan(\beta)$ - see Ref. [473] for a full set of formulae at leading order²³. Due to its general importance and unique final state, the reminder of Part III focuses exclusively on $\tilde{t}_1 \rightarrow t + \tilde{\chi}_1^0$.

Given that the stop decays via $\tilde{t}_1 \rightarrow t + \tilde{\chi}_1^0$, the stop and neutralino mixing parameters determine the polarization of the top quark. Standard Model $t\bar{t}$ production results in unpolarized top quarks, i.e. equal numbers of left- and right-handed quarks. However, the production of top quarks via stops can result in significant asymmetry.

²³The formulae include the branching partial widths to the other three neutralinos and the higher mass charginos. A derivation in the case of the neutral current decay can be found in e.g. Chapter 12.1 of Ref. [438]. Reference [474] presents a clear discussion of the dependence of the $\mathcal{BR}(\tilde{t}_1 \rightarrow t + \tilde{\chi}_1^0)$ on μ and M_1 , including the case of additional neutralinos.

The interaction vertex is proportional to [475]:

$$\tilde{t}_1 \tilde{\chi}_1^0 (\cos(\theta_{\text{eff}}) P_L + \sin(\theta_{\text{eff}}) P_R) t, \quad (1.40)$$

where P_L and P_R are the usual spin projection operators $\frac{1}{2}(1 \pm \gamma^5)$ and the effective mixing angle is given by

$$\tan \theta_{\text{eff}} = \frac{Y_t N_{14} \cos(\theta_{\tilde{t}}) - \frac{2\sqrt{2}}{3} g_1 N_{11} \sin(\theta_{\tilde{t}})}{\sqrt{2} \left(\frac{g_2}{2} N_{12} + \frac{g_1}{6} N_{11} \right) \cos(\theta_{\tilde{t}}) + Y_t N_{14} \sin(\theta_{\tilde{t}})}. \quad (1.41)$$

The parameter $Y_t = y_t / \sin \beta$, where y_t is the SM top quark Yukawa coupling and the matrix N diagonalizes the mass matrix M_N , $N^\dagger M_N N^{-1} = \text{Diag}(\mathbf{m}_{\tilde{\chi}_i^0})$. In the case $\mathbf{m}_Z \ll |\mu \pm M_i|$, $i = 1, 2$, $N_{11} \approx 1$ and $N_{1j} \approx 0$ for $j > 0$; then, $-\tan \theta_{\text{eff}} \sim \tan(\theta_{\tilde{t}})/6$. The factor of six is due to the asymmetric coupling of the bino to $\tilde{t}_L$ and $\tilde{t}_R$ [474]. Note that the effective mixing angle depends on both the stop mixing matrix **and** the neutralino mixing matrix. Changes in the top quark polarization result in different energy spectra of the final state objects, leading to changes in the efficiency for a given event selection [474, 476]. The phenomenology of stop decay is discussed in more detail in Sec. 2.1.

1.5 Related Models

Before describing the analysis strategy for searching for stops, it is important to note that the $t\bar{t} + E_T^{\text{miss}}$ signature is an important property of many extensions of the SM. Another natural source of $t\bar{t} + E_T^{\text{miss}}$ within SUSY is the pair production of gluinos where each gluino decays $\tilde{g} \rightarrow t\bar{t}$ (*gluino mediated stop* or GMS). In a natural SUSY spectrum, the gluino should not be too much heavier than the stop because the stop mass receives large quantum corrections from the gluino just as the Higgs receives large contributions from the stop (see e.g. Ref. [433]). When the stop is significantly heavier than the LSP, GMS models can have fantastic signatures including many top or b quarks. However, when $m_{\text{stop}} \sim m_{\text{LSP}}$ as might be needed to regulate the amount of dark matter [477–480], the stop decay products can be too soft to measure and therefore the total signature is $t\bar{t} + E_T^{\text{miss}} + \text{soft}$, as illustrated by Fig. 1.10. The properties of these models are discussed more detail in Sec. 7.4.2.1.

There is also a wide range of non-SUSY models that produce $t\bar{t}$ in association with weakly interacting particles. For example, new particles with both lepton and baryon number (*leptoquarks* [481]) could decay to a top quark and a neutrino. These third generation leptoquarks have been recently proposed as a model to explain the $\bar{B} \rightarrow D^* \tau \bar{\nu}$ excess [482]. Scalar leptoquark production is identical to stop pair production, but there can be differences in kinematic distributions of the decay products due to the spin configurations of the final state objects. Vector leptoquarks have a significantly higher cross section due to the extra spin states. Another model that results in an increased cross-section is the case of vector-like quarks [483] T' that are fermions with right-handed charged current interactions. These spin 1/2 particles often occur in theories where the Higgs is not fundamental (as a solution to the hierarchy problem) as in the little Higgs [484–486], top-color assisted technicolour [487], composite Higgs [488–495] models. When the $T' \rightarrow tZ$ and $Z \rightarrow \nu\nu$, the final state is similar to the pair production of stops. Figure 1.10 also shows diagrams for the leptoquark and the vector-like quark. There are also dark matter models with non-resonant production giving rise to $t\bar{t}\chi\chi$ for dark matter particle χ [496]. The coupling to mass is a strategy to avoid large flavor changing neutral currents [497].




Figure 1.10: Schematic diagrams illustrating models with similar signatures as stop pair production (top left). Top right: the gluino mediated stop process with nearly mass-degenerate stop and neutralino such that the stop decay products are too soft to be reconstructed. Bottom left: pair production of vector-like quarks T' decaying into Z bosons that decay into neutrinos that are undetected just like neutralinos. Bottom right: leptoquarks carry both lepton and baryon numbers and can decay to a top quark and a neutrino. All of these models are constrained by the search presented in Part III.

Chapter 2

Analysis Strategy

While stop pair production shares many similarities with other searches for new particles, it also requires a dedicated approach. For example, high mass stops produce many high p_T jets and a large E_T^{miss} , but without explicitly targeting final states with top quarks, there is a significant loss in sensitivity. The first searches for the $t\bar{t} + E_T^{\text{miss}}$ topology were performed by the CDF collaboration at the Tevatron at $\sqrt{s} \approx 2$ TeV using the one-lepton [498] and all-hadronic final states [499]. However, the stop pair production cross-section is too low for any model to be excluded with 95% confidence (fermionic top quark partners were excluded up to about 400 GeV). The first stop search to be sensitive to $t\bar{t} + E_T^{\text{miss}}$ was an early $\sqrt{s} = 7$ TeV result by ATLAS in the one-lepton channel [500] using 1 fb^{-1} of data that excludes stops with massless LSP up to about $m_{\text{stop}} \sim 280$ GeV. Using the full $\sqrt{s} = 7$ TeV dataset, ATLAS was able to exclude simplified stop models with stop masses between 230 GeV and 440 GeV for massless LSPs, and top squark masses around 400 GeV are excluded for LSP masses up to 125 GeV [501]. These early analyses focused on applying standard tools to relatively low stop mass models. The remaining natural parameter space is complex and requires a series of dedicated techniques to effectively suppress and estimate backgrounds. Chapter 2 begins with an introduction to stop phenomenology for $m_{\text{stop}} \lesssim 1$ TeV (Sec. 2.1). The background estimation paradigm, called the *control region method* is described in Sec. 2.2 and the technical setup of the analysis is documented in Sec. 2.3.

2.1 Phenomenology

In the approximation that the stop sector decouples from the rest of the MSSM, the cross section for stop pair production depends only on the stop mass. Like any other non-resonant process with $m \ll \sqrt{s}$, the stop production cross section falls off rapidly as a function of m_{stop} , with a $1/m^2$ matrix element suppression compounded with a significant PDF suppression. Figure 2.1 shows the pair-production cross-section at $\sqrt{s} = 8$ and $\sqrt{s} = 13$ TeV as a function of mass from $100 \text{ GeV} < m_{\text{stop}} < 2 \text{ TeV}$. For $m_{\text{stop}} \gtrsim 250 \text{ GeV}$, $\sigma(m_{\text{stop}}) \sim 1/m_{\text{stop}}^6$. For $m_{\text{stop}} \sim m_{\text{top}}$, the stop cross section is about 15% of the $t\bar{t}$ cross section due to the additional spin states available for the spin 1/2 top quark. Around $m_{\text{stop}} \sim 600 \text{ GeV}$, the $t\bar{t}$ cross section is comparable to the irreducible SM $t\bar{t} + Z(\rightarrow \nu\bar{\nu})$ background. The high mass stop pair production cross section increases more than $t\bar{t}$ between $\sqrt{s} = 8$ and $\sqrt{s} = 13$ TeV due to the relatively larger gain in parton luminosity at high momentum fraction. However, event selections targeting stop production will enhance the high $m_{t\bar{t}}, p_{T,t\bar{t}}$ tails, which for the same reason also get a larger increase with energy than the inclusive production. This is illustrated in the right plot of Fig. 2.1, which shows that even though the inclusive $t\bar{t}$ cross section increases by a factor of about three, after requiring $E_T^{\text{miss}} > 300 \text{ GeV}$ the increase is by more than a factor of six.

In addition to the stop mass, the other relevant mass scale is the neutralino mass, which sets how much phase space is available to the stop decay products for a fixed stop mass. Fig. 2.2 shows the average invariant mass of the stop decay products for the decay $\tilde{t} \rightarrow b f f' \tilde{\chi}_1^0$. Events are generated with a four-body phase-space and then re-weighted via

$$\frac{1}{(m_{b f f'}^2 - m_t^2)^2 + m_t^2 \Gamma_t^2} \frac{1}{(m_{f f'}^2 - m_W^2)^2 + m_W^2 \Gamma_W^2}, \quad (2.1)$$

where $m_t = 175 \text{ GeV}$, $\Gamma_t = 1.3 \text{ GeV}$, $m_W = 80 \text{ GeV}$ and $\Gamma_t = 2.5 \text{ GeV}$. The three stripes correspond to the on-shell ($m_{\text{stop}} > m_{\text{top}} + m_{\text{LSP}}$), three-body ($m_W + m_{\text{LSP}} < m_{\text{stop}} < m_{\text{top}} + m_{\text{LSP}}$), and four-body ($m_{\text{stop}} < m_W + m_{\text{LSP}}$) regions of parameter space. Away from the on-shell region, the decay through a virtual top quark competes




Figure 2.1: Left: The stop pair production cross section at $\sqrt{s} = 8$ TeV and $\sqrt{s} = 13$ TeV as a function of the stop mass compared with important SM background processes $t\bar{t}$ and $t\bar{t} + Z(\rightarrow \nu\bar{\nu})$. The SUSY cross section is calculated at NLO+NLL [502, 503] while the $t\bar{t}$ cross section has NNLO+NNLL accuracy using TOP++2.0 [218] with the PDF4LHC prescription [504]. The $t\bar{t} + Z$ cross section is computed at NLO from Ref. [505]. Right: the $t\bar{t}$ cross section as a function of the particle-level E_T^{miss} .

with the loop suppressed flavor changing neutral current process $\tilde{t} \rightarrow c\tilde{\chi}_1^0$ (could also be $\tilde{t} \rightarrow u\tilde{\chi}_1^0$). Up until the CDF Tevatron Run II searches in Ref. [498, 499], this was the only decay channel used for searching for flavor neutral stop decays. The LEP experiments ruled out stops in this decay channel with $m_{\text{stop}} \lesssim 100$ GeV [506–509] and the Tevatron experiments excluded these models for $m_{\text{stop}} \lesssim m_t$ and $m_{\text{LSP}} \lesssim m_{\text{stop}} - 40$ GeV [510, 511]. At the LHC, the most powerful search strategies for this topology involve dedicated charm-jet tagging techniques [512] and the associated production of stops with initial state radiation (ISR) jets [513], excluding m_{stop} up to 300 GeV. Traditionally, the $\tilde{t} \rightarrow c\tilde{\chi}_1^0$ decay mode was only considered in the four-body region of Fig. 2.2. However, the tradeoff between the two processes can be relevant all the way until the boundary of the on-shell top decay [514]. The off-shell regions of parameter space are briefly discussed in subsequent sections, but the remainder

of Part III will be focused on the on-shell regime, which gives rise to the $t\bar{t} + E_T^{\text{miss}}$ signature¹ from the simplified model shown in Fig. 2.3.




Figure 2.2: The average invariant mass of the b-quark and SM fermions from the decay $\tilde{t} \rightarrow b f f' \tilde{\chi}_0^1$ as a function of the stop mass and neutralino mass.

The kinematic properties of a stop decay event are determined by the momentum of the top quark and the neutralino, as illustrated by Fig. 2.4. Due to a steeply falling PDF, high mass stops are produced nearly at rest in the lab frame and so the magnitude of the top quark and neutralino momentum are nearly same in this frame. For a given stop mass M and neutralino mass m , this momentum is given by

$$p(M, m) = \sqrt{\frac{(M^2 - (m_{\text{top}} - m)^2)(M^2 - (m_{\text{top}} + m)^2)}{4M^2}}. \quad (2.2)$$

Figure 2.5 shows the top quark momentum as a function of the stop mass and neutralino mass using Eq. 2.2. Over a large region of parameter space, the relative acceptance (using the top momentum as a proxy) is relatively constant and near 100%. For $m_{\text{LSP}} \sim \frac{1}{2}m_{\text{stop}} + 100$ GeV, there is a sharp transition where the acceptance drops

¹The transition regions have finite width, so care is required when considering models in those regions where the phenomenology is rapidly changing. The boundary region is discussed in the context of the search results in Sec. 7.




Figure 2.3: A leading order Feynman diagram for the simplified stop model considered in Part III. The pair production of stops is followed by the subsequent decay to top quarks and neutralinos. The top quark decays nearly 100% of the time into a W boson and a b -quark. One of the W bosons decays leptonically and the other decays hadronically (into quarks).

to zero at the kinematic boundary $m_{\text{stop}} = m_{\text{top}} + m_{\text{LSP}}$. The distributions of kinematic variables will be significantly different in this transition region compared with the ‘bulk’ and therefore multiple event selections are required to maintain sensitivity across the full parameter space. Changes in acceptance are combined with the falling cross section in Fig. 2.1. Since the level curves of Fig. 2.1 represent lines of constant signal yield, the sensitivity to stop models plotted in this plane should have the same form for a fixed signal region.

There is one other parameter needed to determine the kinematic properties of the stop decay products: the effective coupling to left- and right-handed top quarks θ_{eff} from Eq. 1.41. Before quantifying the impact of this effective coupling, Fig. 2.7 illustrates how it can impact the phenomenology. For simplicity, suppose that all decays happen along a line. The b -quark mass and the neutrino masses are negligibly small and so they always have left-handed helicity (= chirality for massless particles). Focusing on the left diagram of Fig. 2.7, if the top quark is produced with a positive helicity (blue), then the fixed helicity of the b -quark requires the W to




Figure 2.4: For the kinematically tightest event selection from the Run 1 stop search in the one-lepton channel (see Sec. 4), the acceptance is plotted as a function of the top quark momentum given by Eq. 2.2. The acceptance is defined as the fraction of simulated signal events that pass a particle-level version of the analysis (no detector simulation). The ratio of the particle-level acceptance to the acceptance using the full detector simulation are all within a few percent of one in the relevant region and so can be safely ignored. In the region beyond 600 GeV, the acceptance is well described by a straight line $\epsilon = 0.027 \frac{\%}{\text{GeV}} \times p - 4.3\%$. Every point with the same color corresponds to a model with the same stop mass. The spread in the top/neutralino (LSP) momentum is due to the spread in the neutralino masses.

be longitudinally polarized. In this case, there is no preferred momentum orientation (along the line) for the W decay products. However, if the top quark is produced with negative helicity, then only one configuration of W boson decay products is allowed: the charged lepton must be going to the right in the W boson rest frame. This means that in the lab (stop) frame, the charged lepton will tend to have a softer momentum spectrum than the neutrino. Following the same logic in the right diagram of Fig. 2.7 shows that for a (helicity) right-handed top, the charged lepton tends to be anti-parallel to the W boson momentum but parallel to the top quark momentum. The result is a slightly harder charged lepton momentum spectrum in the lab frame relative to the neutrino. There are two other diagrams for all possible orientations of




Figure 2.5: The top (and neutralino) momentum given by Eq. 2.2 as a function of the stop mass and neutralino mass. The momentum is normalized to one at $m_{\text{LSP}} = 0$ for a direct comparison of the acceptance different stop masses.




Figure 2.6: A combination of the acceptance from Fig. 2.5 and the cross section from Fig. 2.1. The z-axis is normalized to unity at $(m_{\text{stop}}, m_{\text{LSP}}) = (400, 0)$.

the momenta, but they are mirror images of the ones in Fig. 2.7 and give the same conclusions. The overall conclusion is that (helicity) left-handed top quarks result in a softer charged lepton momentum spectrum while (helicity) right-handed top quarks result in a harder charged lepton momentum spectrum.




Figure 2.7: Diagrams indicating the various spin configurations of the stop decay products. Single black arrow indicate the direction of the momentum in the rest frame of the particle above the origin of the arrows and double-lined colored arrows indicate the production of the spin along the momentum direction. There are two other possible collinear diagrams per panel where the top begins moving to the right, but the conclusions are the same as for these two.

Chiral left-handed top quarks tend to result in helicity left-handed top quarks and vice versa [474]. For a bino LSP, chiral right-handed stop² result in chiral right-handed top quarks and vice versa since $U(1)$ connects particles of like-chirality. In contrast, since the higgs coupling is like the mass terms in the Lagrangian which couples left to right chiral states, for a higgsino LSP, chiral right-handed stops result in chiral left-handed top quarks. The impact on the distribution of the lepton kinematics can

²The stop is spin 0, so this is short-hand for the superpartner of the chiral right-handed top quark. The helicity is with respect to the lab frame; for massive particles, one can always find a frame in which the sign of the helicity is reversed.

be quantified using the effective mixing angle from Eq. 1.41 [475]:

$$\frac{dN}{d \cos \theta_l} \propto E_{\tilde{\chi}_1^0}^{\text{top frame}} + 2 \sin(\theta_{\text{eff}}) m_{\tilde{\chi}_1^0} + p_{\tilde{\chi}_1^0}^{\text{top frame}} \cos(2\theta_{\text{eff}}) \cos \theta_l, \quad (2.3)$$

where θ_l is the angle between the lepton and the neutralino (from the same stop decay) in the top quark rest frame, $E_{\tilde{\chi}_1^0}^{\text{top frame}}$ is the energy of the neutralino in the top quark rest frame and $p_{\tilde{\chi}_1^0}^{\text{top frame}}$ is the momentum of the neutralino in this frame. Equation 2.3 shows that the behavior of the lepton angle in the top quark rest frame is determined by the quantity $\cos(2\theta_{\text{eff}})$. Figure 2.8 uses Eq. 1.41 to show how $\cos(2\theta_{\text{eff}})$ depends on the stop mixing angle $\theta_{\tilde{t}}$ for a pure bino LSP ($N_{11} = 1$, $N_{1i} = 0$, $i = 2, 3, 4$) and a pure higgsino LSP ($N_{14} = 1$, $N_{1i} = 0$, $i = 1, 2, 3$). As expected from the previous description³; for a purely chiral stop (i.e. $\tilde{t}_1 = \tilde{t}_L$ or $\tilde{t}_1 = \tilde{t}_R$), the top chirality will be the same as the stop for a bino LSP and the opposite for a higgsino LSP.




Figure 2.8: The dependence of $\cos(\theta_{\text{eff}})$ on the stop mixing angle for a pure bino (black dashed) and a pure higgsino (red) LSP.

The impact of the effective mixing angle θ_{eff} on the lepton p_T spectrum is shown in Fig. 2.9. As expected from the simple description around Fig. 2.7, the charged lepton p_T spectrum is harder in the case of mostly right-handed top quarks ($\theta_{\text{eff}} = \pi/2$)

³Note that the graphs in this plot are -1 times the ones appearing in Ref. [475]. See the text for the explanation why the form given here agrees with the expectation of the coupling structure.

than for mostly left-handed top quarks ($\theta_{\text{eff}} = 0$). One way to quantify the impact of the change in the p_T spectrum is to evaluate the efficiency for a fixed lepton p_T requirement as a function of stop mass. Figure 2.10 shows the ratio of efficiencies between the mostly right- and mostly left-handed top quark configurations for a $p_T > 25$ GeV threshold on the charged lepton momentum⁴. Over most of the $(m_{\text{stop}}, m_{\text{LSP}})$ plane, the change in acceptance is about 25%. Near the $m_{\text{stop}} \sim m_{\text{LSP}}$ diagonal, the impact of the polarization is less because the coefficient of the $\cos(2\theta_{\text{eff}})$ term in Eq. 2.3 is suppressed by the reduced phase space. In the limit $m_{\text{stop}} \rightarrow m_{\text{LSP}}$, $dN/d\cos\theta$ is constant, independent of θ_{eff} . Figure 2.11 combines information from Fig. 2.10 with the stop cross section in Fig. 2.1 to estimate the reduction in the expected sensitivity for the extreme values of θ_{eff} . For a stop with $\theta_{\text{eff}} = \pi/2$ and $m_{\text{stop}} = 500$ GeV, the number of predicted events is comparable to a stop with $\theta_{\text{eff}} = 0$ and $m_{\text{stop}} \sim 480$ GeV, resulting in a predicted loss in sensitivity of about 20 GeV. The actual impact in the limit will be discussed in Sec. 7.4.2 with the results.




Figure 2.9: The charged lepton p_T spectrum in stop decays for three values of θ_{eff} for $(m_{\text{stop}}, m_{\text{LSP}}) = (500, 0)$.

⁴The top quark polarization impacts the other decay products as well, but the effect is largest for the lepton p_T because it is further down the decay chain compared to the b -quark and additional sources of jets can mitigate the impact from the hadronically decaying W boson.




Figure 2.10: The ratio of the efficiency for a $p_T^{\text{charged lepton}} > 25$ GeV requirement between the mostly left-handed top quark ($\theta_{\text{eff}} = 0$) and the mostly right-handed top quark ($\theta_{\text{eff}} = \pi/2$).




Figure 2.11: For a fixed stop mass $m_{\text{stop,R}}$ (LSP is massless), plotted is $m_{\text{stop,R}} - m_{\text{stop,L}}$ for $\sigma(m_{\text{stop,L}}) \times \epsilon_L = \sigma(m_{\text{stop,R}}) \times \epsilon_R$. The value $m_{\text{stop,R}}$ is the stop mass for $\theta_{\text{eff}} = \pi/2$ and $m_{\text{stop,L}}$ is the stop mass for $\theta_{\text{eff}} = 0$, so the plot shows how much the stop mass needs to be reduced to have the same predicted number of events with a most left-handed top quark configuration compared to a given mostly right-handed top quark configuration.

For high mass $m_{\text{stop}} \gtrsim 200$ GeV and m_{LSP} sufficiently far away the diagonal $m_{\text{stop}} = m_{\text{top}} + m_{\text{LSP}}$, the cross section $\times$ acceptance map in Fig. 2.6 coupled with the top polarization completely categorizes the properties of the signal necessary for the search. For these models, the strategy is to develop event selections estimated to have a high purity and a high yield of stop events for a particular *benchmark model*. The reach of each selection will then be determined by the discussion above and the choice of benchmark models is set by the goal to cover a wide range of the $(m_{\text{stop}}, m_{\text{LSP}})$ parameter space. The stop signal is characterized by a large $\vec{p}_{\text{T}}^{\text{miss}}$ from the neutralinos, one isolated charged lepton and at least four high p_{T} jets resulting from the tree-level top quarks shown in Fig. 2.3. Two of these jets are expected to originate from b -quarks. The construction of discriminating variable and their use in event selections are described in Chapters 3 and 4, respectively.

When m_{LSP} is near $m_{\text{stop}} - m_{\text{top}}$, the neutralinos and top quarks have very little momentum in the stop rest frame and so the signal $E_{\text{T}}^{\text{miss}}$ and lepton/jet momentum spectra start to approach the distributions of the $t\bar{t}$ background. For relatively low m_{stop} , the signal cross section is sufficiently large so that one can take advantage of subtle differences in the shapes of various kinematic distributions. Section 2.2.2 describes how the single event selection paradigm can be modified to incorporate shape information. Additional strategies are possible in this regime, and are likely needed for the sensitivity to cross the $m_{\text{LSP}} = m_{\text{stop}} - m_{\text{top}}$ limit. These include requiring additional high p_{T} (ISR) jets to boost the neutralino momentum [515–518] and forward jets in a vector-boson-fusion (VBF) topology [519].

For $m_{\text{stop}} \lesssim 200$ GeV, the stop cross section is so high that top quark properties measurements are sensitive to the presence of a light stop. Recent studies have suggested that the $t\bar{t}$ cross section [520, 521] and spin-correlation measurements between the leptons from top quark decays [522] are sensitive to stops. Exploiting precision measurements of the $t\bar{t}$ cross-section makes use of the current NNLO+NNLL accuracy that reduces the theoretical uncertainty in $\sigma_{t\bar{t}}$ to about 5%, which is sensitive to the $\mathcal{O}(10\%)$ contribution of a degenerate stop. The angular distribution between the two leptons from top quarks encodes information about the production mechanism

and correlations between the top quark spins. Like top quarks, stops are also produced mostly via gluon-gluon fusion at the LHC. However, because stops are scalar particles, there is no direct correlation between the spins of the resulting top quarks. Both ATLAS and CMS have performed measurements to constrain light stop models using the total cross section [520, 523, 524] and spin correlations [522], ruling out stop models with m_{stop} between 150 GeV and 190 GeV. A non-negligible ($\sim 30\%$) contribution to the sensitivity for the spin correlation measurement is due to the constraint from the $t\bar{t}$ cross-section.

Despite these innovative efforts to constrain degenerate $m_{\text{stop}} \sim m_{\text{top}}$ case, there is an important challenge with this regime that can obscure the results. Measurements which exploit the cross section could be effected by a bias in the top quark mass measurement due the presence of a light stop. In particular, since $\sigma_{t\bar{t}}$ increases with decreasing top quark mass, a negative shift in the measured top quark mass would increase the predicted $t\bar{t}$ cross-section and could hide the additional contribution to the measured cross-section from direct stop pair production. Figure 2.12 illustrates the invariant mass distribution of the three jets associated to the hadronically decaying (off-shell) top quark in $pp \rightarrow t\bar{t}$ and $pp \rightarrow \tilde{t}_1 \tilde{t}_1^*$ where the other (off-shell) top quark decays into $W^\pm \rightarrow l^\pm \nu$. The invariant mass distribution is sensitive to the top quark mass and thus can be used to extract the mass from data. When m_{stop} is just below m_{top} , the distribution of m_{jjj} is biased toward lower values. This is also true even when m_{stop} is just above m_{top} because the top quark Breit-Wigner is skewed to lower values (m_{top} must be less than m_{stop}). The presence of these stop events could bias the top quark mass measurement to low values, if the *calibration curve* is derived using only $t\bar{t}$ simulation (right plot of Fig. 2.12). In particular, the measured values of the top quark mass and $t\bar{t}$ cross section are given by

$$\begin{aligned}
m_t^{\text{measured}} &= \frac{\langle m_{\text{jjj}} \rangle_{t\bar{t}} \times \sigma_{t\bar{t}}(m_{\tilde{t}_1}) \times \epsilon + \langle m_{\text{jjj}} \rangle_{t\bar{t}}(m_t) \times \sigma_{t\bar{t}}(m_t)}{c_1(\sigma_{t\bar{t}}(m_{\tilde{t}_1}) \times \epsilon + \sigma_{t\bar{t}}(m_t))} - \frac{c_0}{c_1} \\
\sigma_{t\bar{t}}^{\text{measured}} &= \sigma_{t\bar{t}}(m_{\tilde{t}_1}) \times \epsilon + \sigma_{t\bar{t}}(m_t),
\end{aligned} \tag{2.4}$$

where ϵ is the ratio of the SUSY acceptance to the $t\bar{t}$ acceptance and c_0, c_1 are the

m_t^{true}	m_t^{measured}		True $\sigma_{t\bar{t}}(m_t^{\text{true}})$		True $\sigma_{t\bar{t}}(m_t^{\text{measured}})$		True $\sigma_{t\bar{t}}$		Measured $\sigma_{t\bar{t}}$	
	LHC8	Tevatron	LHC8	Tevatron	LHC8	Tevatron	LHC8	Tevatron	LHC8	Tevatron
170	168.6	169.0	271.1	8.0	279.0	8.1	42.6	0.87	295.4	8.5
172.5	170.8	171.3	251.7	7.3	264.4	7.6	42.6	0.87	276.0	7.8
175	172.9	173.5	233.8	6.8	249.7	7.2	42.6	0.87	258.1	7.3

Table 2.1: Bias in the measured top quark mass and $t\bar{t}$ cross-section due to the presence of a light stop ($m_{\bar{t}} = 170$ GeV) that decays via the three-body process. All masses are in GeV and all cross-sections are in pb. The measured top quark mass is biased low from the true mass which results in the true cross-section at the measured top mass, true $\sigma_{t\bar{t}}(m_t^{\text{measured}})$ to be higher than the true cross-section at the true mass, true $\sigma_{t\bar{t}}(m_t^{\text{true}})$. The former quantity is what would be predicted under the SM-only hypothesis in the presence of the 170 GeV stop. The measured $\sigma_{t\bar{t}}$ is the sum of true $\sigma_{t\bar{t}}(m_t^{\text{true}})$ and true $\sigma_{t\bar{t}}$, corrected for the lower acceptance for the three-body decay.

slope and intercept from the calibration curve in Fig. 2.12, respectively. For example, a stop with $m_{\bar{t}} \sim 170$ GeV that decays via an off-shell top quark together with a true top quark mass of about 175 GeV would cause a bias in the top quark mass that makes it compatible with the measurements with Run 1 of the LHC (LHC8). As a consequence, the predicted $t\bar{t}$ cross-section would be over-estimated by about 16 pb which in turn would make it much harder to find the stop with a cross-section of about 43 pb (which is further reduced to about 60% since the acceptance is lower than for $t\bar{t}$). The cross-section over-estimation increases with the true top quark mass, while the compatibility of the measured top quark mass with the LHC8 decreases when going beyond about 175 GeV. Figure 2.13 and Table 2.1 summarize how the change in the measured mass could hide such a *sneaky stop*⁵. The impact for $m_{\text{stop}} > m_{\text{top}}$ is greatly reduced because the top quarks have nearly the same mass distribution as SM $t\bar{t}$ production.

Due to the relatively lower stop cross section, the shift in the measured top quark mass is predicted to be smaller at the Tevatron compared to the LHC. Figure 2.14 shows the size of the shift as a function of the true top quark mass between the two colliders. Interestingly, there is a small tension between measured top quark mass values between the Tevatron and LHC experiments that is in the correct direction

⁵The analysis in this section has been published in Ref. [525].




Figure 2.12: Left: Unit normalized distributions of the m_{jjj} variable for $t\bar{t}$ with $m_t = 172.5 \text{ GeV}$, and for $\tilde{t}\bar{\tilde{t}}$ pair production with a two-body $\tilde{t} \rightarrow tN$ decay with $m_{\tilde{t}} = 175 \text{ GeV}$ (and $m_t = 172.5 \text{ GeV}$), and a three-body decay $\tilde{t} \rightarrow bWN$ for $m_{\tilde{t}} = 170 \text{ GeV}$. The neutralino is assumed massless. Jets are assigned to the hadronically decaying top quark by minimizing $\chi^2 = (m_{j_1 j_2 b_1} - m_{b_2 l\nu})^2 / (20 \text{ GeV})^2 + (m_{j_1 j_2} - m_W)^2 / (10 \text{ GeV})^2$ for $\{j_i\}$ the set of jets not identified as originating from a b -quark. Even though the neutrino p_z is unmeasured, it can be inferred by solving $m_{l\nu} = m_W$. The simulations are performed using HERWIG++ 2.7 [195, 335] and analyzed using the RIVET 1.8.2 framework [297] with FASTJET 3.0.6 [203] for clustering anti- k_t jets with $R = 0.4$ [117]. Right: Calibration curve that relates the measured value $\langle m_{jjj} \rangle$ to the (MC) top quark mass, m_{top} in $t\bar{t}$ events. See Ref. [525] for more detail.

predicted by a light stop. However, all of the most precise measurements use a calibration scheme involving the *Monte Carlo* mass, which is related to a well-defined QFT top quark mass only within ambiguities of $\mathcal{O}(\Lambda_{\text{QCD}})$ and the relation may depend on $\sqrt{s}$; see e.g. ref. [527]. Since a wide range of simulation schemes where various MC mass definitions are used, this tension is not a significant indication of deviations from the SM.

The mass measurement based on $\langle m_{jjj} \rangle$ is highly simplified from the state-of-the-art. However, prompted by the above argument, a detailed study using the most precise ATLAS top quark mass measurement technique [528] indicates that the impact




Figure 2.13: Summary of the effects leading to the sneaky stop scenario: the shifts in the measured $t\bar{t}$ cross-section and measured top quark mass. The solid line corresponds to an unbiased measurement of the $t\bar{t}$ cross-section as a function of the top quark mass. The dot-dashed line is what would be measured in the presence of a $\tilde{t} \rightarrow b W N$ with $m_{\tilde{t}} = 170$ GeV for an unbiased top quark mass measurement. However, under the SM+SUSY hypothesis the top quark mass measurement would be *biased* which translates into what would actually be observed shown in the dashed line. For all three lines, the band reflects the $\sim 5 - 6\%$ theory uncertainty on the cross-section. For comparison, the measured top quark mass and $t\bar{t}$ cross-section are shown from recent CMS [526] and ATLAS [520] results.




Figure 2.14: The measured top quark mass as a function of the true top quark mass. The bias in the measurement arises from the presence of a light $\tilde{t}$ with $m_{\tilde{t}} = 170$ GeV and decaying via the three-body process (left) or with $m_{\tilde{t}} = 175$ GeV and decaying via the two-body process.

on the stop limits could be as much as 5 GeV [523]. As the LHC accumulates more data at $\sqrt{s} = 13$ TeV and systematic uncertainties are reduced, cross section and other top quark properties should continue to be exploited to ensure that no stop is hiding around $m_{\text{stop}} \lesssim m_{\text{top}}$.

The remainder of Part III will focus on the direct search for a light stop in the one lepton + four jets + missing momentum final state using the *control region method*, described in Sec. 2.2.

2.2 The Control Region Method

In order to identify stop events among an overwhelming SM background, key variables are identified for which the probability distribution significantly differs between signal and background. Figure 2.16 illustrates how these variables are used to estimate, validate, and test the background predictions. Each background process is separately estimated, but for the sake of simplicity, suppose that there is one type of SM background and one powerful variable V . Examples of V appear in Sec. 3, but typically V is associated with an energy scale in the event and the likelihood $p_S(V)/p_B(V)$ monotonically increases as a function of V . A *signal region* (SR) is an interval of V predicted to have low background and high signal yield. As the likelihood is often monotonically increasing, these regions usually take the form $[v_0, \infty)$ for some fixed v_0 . The goal is compare the number of predicted signal events to the number of predicted background events in the SR. If the Poisson fluctuations in the background are (much) larger than the predicted signal yield, then this is a hopeless exercise. For this reason, the signal region is usually defined by $v_0 \gg 1$ (with the appropriate units, often GeV) where $p_S(V)/p_B(V) \gtrsim 1$. The shapes $p_S(V)$ and $p_B(V)$ are obtained from MC simulation. The number of predicted events in the signal region⁶ is then given by $N_B \int_{\text{SR}} p_S(V) dV$ and $N_S \int_{\text{SR}} p_S(V) dV$ for the background and signal, respectively. The factor N_S is given by $\mathcal{L}_{\text{int}} \times \sigma \times \epsilon$, where $\mathcal{L}_{\text{int}}$ is the integrated luminosity, σ is the cross section calculated for the signal, and ϵ is the efficiency of all event selections prior to the selection on V (also estimated from simulation). In contrast, N_B is defined such that $N_B \int_{\text{CR}} p_B(V) dV$ agrees with the observed data in a *control region* (CR), which is an interval of V where the signal is expected to be negligible compared with the background. By normalizing the background in the control region, the predicted number of events at low values of V is ‘correct’ (see Sec. 2.2.3 for caveats) and the only uncertainty is due to extrapolating this prediction to the SR using $p_B(V)$. The control region should be kinematically close to the SR in order to reduce any uncertainty from this extrapolation, but should also be loose enough so that the total number of events in the control region allows for a relatively precise

⁶The background events in the signal region are also used for the final result, but have little influence due to the small total yield. See Sec. 7.1 for details.

measurement of N_B . Often the requirement for higher event yields in the CR results in a significant gap between the CR and the SR. Part of the region in between where $p_S(V)/p_B(V)$ is still small can be used to validate the CR prediction, albeit with a significant uncertainty. Such a region is called a *validation region*.

In practice, many variables are combined to form the signal and control regions and there are a variety of background processes. Section 2.2.1 presents an overview of the signal regions defined by a single set of kinematic requirements (*single bin SR*) and Sec. 2.2.2 describes the setup when there are multiple SR bins used at the same time. The control region method is a powerful tool for exploring regions of phase space that are not well constrained by previous measurements. However, it does have some limitations (Sec. 2.2.3) which are important to understand in the exploration of the TeV scale and beyond.




Figure 2.15: A schematic diagram to illustrate the control region method. The red distribution represents $N_B p_B(V)$ and the blue is $N_S p_S(V)$ (see the text for details). The left and right distributions are identical, only with a logarithmic scale for the vertical axis in the right plot. Control, validation, and signal regions are denoted by CR, VR, and SR, respectively. The symbol S is shorthand for $N_S p_S(V)$ and B represents $N_B p_B(V)$.

2.2.1 Single Bin Signal Regions

In the most basic and widely used form of the control region method, there is one SR and multiple control regions that constrain various background processes. All but one of the signal regions for Part III have this structure. Figure 2.16 illustrates the setup using the early $\sqrt{s} = 13$ TeV analysis as an example. A variety of kinematic variables are used to define a signal region. Two of these variables, in this case called m_T (Sec. 3.2.1) and am_{T2} (Sec. 3.2.1.2), are changed to form control regions that are disjoint from the SR and to each other. In addition to the kinematic requirements, the number of b -jets is a powerful tool for building control regions. Requirements on other kinematic variables may also be loosened in order to increase the CR statistics, but the background composition is determined by a few key variables. All of these control regions are described in more detail in Chapter 5. In Fig. 2.16, there are four control regions (TCR, WCR, STCR, TZCR) and in between these regions and the SR are three validation regions (TVR, WVR, and WVR-tail). Not all control regions have validation regions and some regions (e.g. WCR) can have multiple validation regions to probe different aspects of the extrapolation from the CR to the SR.




Figure 2.16: The single bin control and validation region setup for the $\sqrt{s} = 13$ TeV analysis. All regions share a common selection and then are further distinguished by requirements on m_T , am_{T2} and the number of b -jets. The prefixes stand for ST = single top, T = $t\bar{t}$, W = W +jets, and TZ = $t\bar{t} + Z$. The exact definitions of the regions are given in Sec. 5.

2.2.2 Multibin (Shape Fit) Signal Region

One way to increase the sensitivity of a signal region is to split it into multiple bins that have different p_s/p_b . Multiple bins increase the sensitivity because it provides a finer scale for the likelihood and effectively gives a higher weight to events where the likelihood is higher. Figure 2.17 quantifies this statement with an example; compared with the one-bin setup, the two-bin setup has a lower probability for rejecting the SM when there is SUSY for a fixed probability to reject the SM when it is in fact true. The gain is bigger when the difference in the likelihoods between bins is bigger.




Figure 2.17: The tradeoff between Type II (do not reject SM when SUSY is true) and Type I errors (reject SM when SM is true) for three scenarios. In all cases, the number of background events is 200 in one bin and 100 in a second bin. The total number of signal events is 30. In the left plot, both bins have 10% signal; in the middle plot the fraction is twice as high in the first bin; in the right plot, the fraction is four times as high in the first bin compared with the second bin. The errors are computed by scanning threshold requirements on the (log) likelihood ratio distribution. For the two-bin case, the likelihood ratio is a product of the individual bin likelihood ratios. Only statistical uncertainties are included in the likelihood. The vertical dashed line is at 5%.

For more intuition, consider a two-bin setup with mean background event yields of B_1 and B_2 and signal yields S_1 and S_2 . The optimal test procedure is based on the likelihood ratio $p_{s+b}(x)/p_b$ (see Sec. 7.1). Suppose that $S_1 = 0$. Ideally, one would remove the first bin, as it contains no useful discriminating information for the signal. If the two bins are lumped together, the first bin dilutes the power of the two bins

together. However, if the two bins are split and the likelihood is a product over the two bins, then the first bin automatically does not contribute ($p_{s_1+b_1}(x_1)/p_{b_1}(x_1) = 1$). This is an extreme case, but it illustrates the main point. In principle, the optimal procedure is to weight every event by its log likelihood ratio (i.e. put each event in its own bin⁷), but this makes it difficult to validate the modeling of the weights and so the focus here is on a small number of bins. Binning is not used (yet) for the SR setup introduced in Sec. 2.2.1 because of the explicit model dependence through p_s .

The shape fit region still uses the control region method, but in a more integrated way than for the single bin regions. Control regions and signal regions are simply bins of a multibin SR where the p_s/p_b is very low in the CR-like bins and high in the SR-like bins. This setup is illustrated in Fig. 2.18. The selection (Sec. 4.4) and fit procedure (Sec. 7.1) are described in later sections. Just like Fig. 2.16, there are two key kinematic variables in addition to b -tagging information to control the purity of various background processes. The expected p_s/p_b is about 20% in some of the most signal-sensitive bins (upper-right of Fig. 2.16) and about 10% in others. In the CR-like regions, the likelihood is less than 1%. A global normalization of the various data-driven backgrounds is possible but puts a stringent constraint on the modeling as a function of the variables defining the bins in Fig. 2.18. Since the background composition changes most strongly with the m_T variable and not E_T^{miss} , each column in Fig. 2.18 is separately normalized using the CR method. This also mitigates any potential mis-modeling of the E_T^{miss} trigger turn-on for $E_T^{\text{miss}} \sim 100$ GeV. By construction, the fitted background prediction is nearly identical to the data yield in the CR-like bins. The implications the differences between data and simulation in the SR-like regions have for light stops are discussed in Sec. 7.4.2.

⁷To see that these are equivalent, consider a case where there are only two possible values of p_s/p_b . Label the bins 1 and 2 and then the log likelihood for the two bins is (up to a constant) $x_0 \log(1 + s_1/b_1) + x_1 \log(1 + s_2/b_2)$, where x_i is the number of observed events, s_i is the mean number of signal events, and b_i is the mean number of background events in bin i . Instead, suppose each event in bin i is weighted by $\log(1 + s_i/b_i)$ so that the total number of ‘measured events’ is $x = x_0 \log(1 + s_1/b_1) + x_1 \log(1 + s_2/b_2)$. Then, the log likelihood is (up to a constant) $x \log(1 + s/b)$ for $s = s_1 + s_2$ and $b = b_1 + b_2$. Since $\log(1 + s/b)$ is a constant across bins, the likelihood with the weighted setup is a monotonic function of the binned likelihood. Therefore, they result in the same statistical power for a fixed signal model. However, this result may not (exactly) hold with a different test statistic.

2.2.3 Drawbacks of the CR Method and Alternatives

The main disadvantage of the control region method is the assumption that the shape $p_B(V)$ from Sec. 2.2 is known. Differential distributions are usually known with less precision than the total cross-section and tails of distributions are known to be sensitive to higher order effects (and in some cases, non-perturbative modeling). Therefore, a thorough investigation of potential sources of systematic bias in extrapolating from the control region to the signal region is presented in Sec. 6.

Another, more subtle disadvantage of the standard control region method is that it can be very sensitive to statistical fluctuations when the number of events in the control region is small. If there is an under-fluctuation in the data, then the predicted background yield in the signal region will be too small. This is partially accounted for in the statistical uncertainty from the data in the control region, but the *central value* will be biased. Figure 2.19 illustrates the source of bias by showing the probability of obtaining 2σ evidence for SUSY when there is only background. The probability should be 5%. It is not exactly so even when the number of events in the control region is infinite because the number of observed events in the signal region can only take discrete values. The most striking feature of Fig. 2.19 is that the probability for a 2σ excess is almost a factor of two higher for $N[\text{control region}] = 10$ compared with $N[\text{control region}] \rightarrow \infty$. Figure 2.20 is another way to view the problem, but in the case where there is SUSY. Suppose there would be a 3σ excess if the true expected number of events in the signal region were known (i.e. an infinite number of events in the control region). Fig. 2.20 shows the probability for this value to drop to less than 3σ when the number of events in the control region is finite. For example, the probability for a 3σ to drop to a 2σ is about 10% when the true expected number of events is 10 in both the signal and control regions. This has important implications for discovery as the threshold for ‘evidence’ is usually set at 3σ , while 2σ fluctuations are largely ignored. These biases are mostly mitigated when $N[\text{control region}] \gtrsim \mathcal{O}(100)$, as is mostly true for the regions constructed in Sec. 5.

One simple modification of the standard control region method is to adapt it to a Bayesian framework. The main problem arises because the standard method puts too much emphasis on the observed data in the CR. A Bayesian approach would be:




Figure 2.19: The probability for a 2σ excess when there is only background as a function of the number of expected events in the control region and in the signal region using the standard control region method. A 2σ excess is defined as a case when the probability for the observed number of events in the signal region to exceed the number of predicted events in the signal region to be less than 5%. Only statistical uncertainties are included in this calculation.




Figure 2.20: The probability for a 3σ excess (with ∞ events in the control region) to be measured as a $< 3\sigma$ excess.

$$p(N_B|N_{CR}^{\text{observed}}) \propto \text{Poisson}(N_{CR}^{\text{observed}}|N_B)p(N_B), \quad (2.5)$$

where $p(N_B)$ is a prior distribution. One reasonable prior is the Gamma distribution, which is conjugate for the Poisson. The posterior mean for a Gamma prior with parameters α and β is

$$\langle N_B | N_{CR}^{\text{observed}} \rangle = \left(\frac{1}{1 + \beta} \right) N_{CR}^{\text{observed}} + \left(\frac{\beta}{1 + \beta} \right) \frac{\alpha}{\beta}, \quad (2.6)$$

which is a linear superposition of the prior mean (α/β) and the observed number of events from one observation. The parameter β plays the role of the number of effective events ‘observed’ prior to seeing any data. If there are auxiliary measurements that can be used to constrain α/β , then the number of events in such a region could be used to set β . Another possibility is to use the estimated systematic uncertainty on the number of events from the raw simulation. For example, a systematic uncertainty of 30% would correspond to an auxiliary measurement of ~ 10 events, since $1/\sqrt{10} \sim 30\%$. Therefore, one could pick $\beta = 10$ and then $\alpha = \beta \times N_B^{\text{MC}}$. This procedure is illustrated in Fig. 2.21. Compared to Fig. 2.19, the probability for observing an excess when there is no SUSY (often called Type 1 error) is significantly reduced.

The Bayesian framework also allows a natural method for incorporating systematic uncertainties into the control region method. The usual procedure for estimating the systematic uncertainty is to consider alternative plausible simulations and compare the differences in the predicted events in the signal region using the various models. Usually, one simulation is taken as nominal and the differences with respect to the other models are taken as Gaussian uncertainties on the expected number of events in the signal region. One way to avoid this ad-hoc approach is to use a hierarchical Bayes model in which there is a latent variable z that describes which simulation is the best description of the data. There could be many acceptable models and the final prediction is estimated by integrating out z . The posterior variance would also give a sense of the uncertainty (in fact, any measure of spread based off of the posterior could be used). This model is illustrated graphically in Fig. 2.22. Additionally, it may be possible to combine this with an empirical Bayes approach in which many searches are simultaneously used to provide a prior for z . If no other information is




Figure 2.21: The probability for a 2σ excess under the null (only known particles) hypothesis as a function of the number of expected events in the control region and in the signal region using the Bayesian control region method described in the text. A 2σ excess is defined as a case when the probability for the observed number of events in the signal region to exceed the number of predicted events in the signal region to be less than 5%. A systematic uncertainty of 30% is assumed when choosing β .

available, a discrete uniform random variable on $\{1, \dots, n\}$ could be used. Figure 2.23 illustrates the posterior under such a scheme with three plausible simulations, using the uniform prior for z . As expected, the distribution is slightly broader when the sample variations are included.




Figure 2.22: A graphical representation of the hierarchical Bayes model described in the text. For model M_j , the predicted background yield (N_B^{MC}) is denoted λ_i^j .




Figure 2.23: Sample posteriors using the hierarchical Bayes model described in the text. The nominal simulation predicts $(N_{\text{CR}}^{\text{raw}}, N_{\text{SR}}^{\text{raw}}) = (10, 5)$ and the two alternative models have $(N_{\text{CR}}^{\text{raw}}, N_{\text{SR}}^{\text{raw}}) = (10.5, 5.5)$ and $(N_{\text{CR}}^{\text{raw}}, N_{\text{SR}}^{\text{raw}}) = (9.5, 4.5)$. The number of observed events in the control region is 11 and the systematic uncertainty, used to construct β is 30%.

One could additionally generalize the above procedures when there are multiple control regions or when a control region has multiple bins. In that case, the simulation will be properly down-weighted to account for the new information from the data. These Bayesian methods are promising ways to incorporate prior information from simulations and auxiliary measurements, but require more investigation before they can be properly incorporated into an analysis.

2.3 Dataset and Monte Carlo Samples

The $\sqrt{s} = 8$ TeV data from Run 1 were collected between March and December 2012 resulting in an integrated luminosity of 20.3 fb^{-1} and the $\sqrt{s} = 13$ TeV data from the early part of Run 2 where collected between October and December 2015 leading to $\int \mathcal{L} dt = 3.2 \text{ fb}^{-1}$. These data were recorded using a combination of single lepton and E_T^{miss} triggers. Figure 2.24 shows the efficiency for various trigger algorithms. Isolated lepton triggers require $p_T > 24$ GeV in addition to particle identification and isolation criteria. The efficiency drops at high stop mass due to the isolation failing when the jets from the same boosted top quark as the lepton are close-by. Inclusive lepton triggers have a higher p_T threshold of $p_T > 60, 36$ GeV for electrons and muons, respectively. The E_T^{miss} trigger is fully efficient for offline $E_T^{\text{miss}} \gtrsim 150\text{-}200$ GeV. Boosted W and top quark jets can fire the single large-radius jet trigger, though the efficiency is not competitive with the other triggers as it is not fully efficient until $p_T \gtrsim 400$ GeV.




Figure 2.24: Left: the trigger efficiency for various algorithms as a function of the stop mass. Right: The additional efficiency provided by various triggers beyond the efficiency already provided by the triggers to the left of it in the legend.

For $m_{\text{stop}} \gtrsim 500$ GeV and a massless LSP, the E_T^{miss} trigger is the most efficient, recording $\gtrsim 95\%$ of events. The isolated single lepton triggers are also highly efficient,

but only add about $\lesssim 2\%$ beyond the E_T^{miss} trigger for high mass stops. For lower stop masses, the single lepton triggers are an essential complement to the E_T^{miss} trigger. As such, the early $\sqrt{s} = 13$ TeV analysis uses only the E_T^{miss} trigger, but the full $\sqrt{s} = 8$ TeV analysis uses events that pass the E_T^{miss} trigger or either of the isolated or inclusive single lepton triggers.

The diverse background processes relevant for the stop search require a complete set of simulated SM samples. Several matrix element (ME) generators are combined with parton shower (PS) generators. Signal SUSY samples are generated at leading order (LO) with MG5_aMC v2 [505] (HERWIG++) at $\sqrt{s} = 13$ ($\sqrt{s} = 8$) TeV. All signal samples generated with MG5_aMC are interfaced with PYTHIA 8.186 [189]. The nominal stop mixing angle is given by $\theta_{\tilde{t}} = 1$ ($\tilde{t}_1$ is mostly $\tilde{t}_R$) and $N_{11} = 1$ (pure bino LSP). This corresponds to an effective mixing angle of $\theta_{\text{eff}} \approx 1.4$. Background samples use one of four setups:

- MG5_aMC or MadGraph [529] interfaced with PYTHIA or HERWIG++ using the CKKW-L [530, 531] or the MC@NLO method for matching a LO or next-to-leading-order (NLO) ME to the PS, respectively.
- POWHEG-BOX [201, 202, 532–534] interfaced to PYTHIA 6 [96], HERWIG+JIMMY, or HERWIG++ using the POWHEG method [89, 200] to match the NLO ME to the PS.
- SHERPA [3] using Comix [535] (LO+extra partons) and OpenLoops [536] (NLO) ME generators interfaced with the SHERPA parton shower [537]. Leading order samples with extra partons use the CKKW matching scheme.
- ACERMC 3.8 [294] interfaced with PYTHIA 6 for fragmentation.

The samples produced with MG5_aMC and POWHEG-BOX at $\sqrt{s} = 13$ TeV use EVTGEN v1.2.0 [105] for the modelling of b -hadron decays. Similarly, the generators TAUOLA [106] and PHOTOS [100] are used to model τ -lepton decays and QED radiative corrections. The simulation setup is summarized in Table 2.2 and more details can be found in Ref. [538–541] for $t\bar{t}$ and single top, W/Z +jets, dibosons, and

$t\bar{t} + W/Z$, respectively. Due to the sub-optimal $h_{\text{damp}} = \infty$ setting for the $t\bar{t}$ sample at $\sqrt{s} = 8$ TeV, a $p_T^{t\bar{t}}$ re-weighting is performed based on the dedicated early Run 1 measurement of this quantity [281]. Additional samples aside from those shown in Table 2.2 are used to assess theoretical modeling uncertainties and will be discussed in Section 6. Such samples are generated using one of the four setups listed above.

As in previous chapters, pileup is simulated by overlaying PYTHIA 8 minimum bias events on the samples listed above. The particle-level simulations are processed using either a full detector simulation [110] based on GEANT 4 [108] or a fast simulation [109] with a parameterized calorimeter response and GEANT 4 for all other processes. There is no significant difference between these setups for the event selections considered in Part III.

Process	ME Generator	ME PDF	Fragmentation	UE Tune	Cross-section Order
$t\bar{t}$	POWHEG-BOX	CT10	PYTHIA 6	P2012 (2011C)	NNLO+NNLL [213–218]
Single top	POWHEG-BOX	CT10	PYTHIA 6	P2012 (2011C)	NNLO+NNLL [268–270]
Single top (t-chan.)	ACERMC	CTEQ6L1	PYTHIA 6	P2011C	NNLO+NNLL [270]
W/Z +jets	SHERPA 2.1.1 (1.4.1)	CT10	SHERPA	Default	NNLO [267]
Diboson	SHERPA 2.1.1 (1.4.1)	CT10	SHERPA	Default	NLO [271, 272]
$t\bar{t} + W/Z$	MG5_AMC v2	NNPDF2.3	PYTHIA 8	A14	NLO [505]
	MadGraph 5	CTEQ6L1	PYTHIA 6	AUET2B	NLO [542, 543]
$t\bar{t} + \gamma$	MG5_AMC v2	CTEQ6L1	PYTHIA 8	A14	NLO [505]
	MADGRAPH 5	CTEQ6L1	PYTHIA 6	AUET2B	NLO [544]
SUSY Signal	MG5_AMC v2	NNPDF2.3	PYTHIA 8	A14	NLO+NLL [502]
	HERWIG++	CTEQ6L1	HERWIG++	EE3	NLO+NLL [503]

Table 2.2: Overview of the nominal simulated samples. The last row indicates the order and reference for the inclusive cross-section to which the (lower order) MC simulations are normalized. The blue indicates a setup at $\sqrt{s} = 8$ that differs from the one used at $\sqrt{s} = 13$ TeV. More information about the $t\bar{t} + \gamma$ generation can be found in Sec. 5.4.2.

Chapter 3

Object and Variable Definitions

The main difference between the stop search and a more inclusive search for squarks and gluinos is the particular stop pair production event topology resulting from high p_T top quarks. In the one-lepton channel, all reconstructable high p_T objects are utilized: (b-tagged) jets, electrons, muons, photons, hadronically decaying τ leptons, and E_T^{miss} . These objects are combined to form discriminating variables designed specifically for $t\bar{t} + E_T^{\text{miss}}$. Section 3.1 provides an overview of the object reconstruction, including the procedures for resolving ambiguities in object labeling (overlap removal). A detailed description of the discriminating variables used in the signal region optimization (Sec. 4) follows in Sec. 3.2.

3.1 Object Selection

Many of the objects used by the stop search were already introduced in Part I and Part II. This section concisely describes each object and highlights the differences between the $\sqrt{s} = 8$ analysis and the early Run 2 $\sqrt{s} = 13$ TeV search. The general strategy is to devise two sets of objects labeled *baseline* and *signal*, where the former are a subset of the latter. Baseline objects are used in the ambiguity solving (overlap removal) and for vetoing events with a second reconstructed electron or muon. Signal objects are used as inputs to the discriminating variables described in Sec. 3.2 and the final event selections. Table 3.1 presents an overview of the object definitions, with references to more detailed documentation. A quantitative comparison of the resolutions and reconstruction efficiencies is presented in Chapter 6. The paragraphs below briefly summarize the selections and notable changes between Runs 1 and 2.

The inputs to jet clustering at $\sqrt{s} = 8$ TeV are calorimeter cell clusters with the LCW calibration while at $\sqrt{s} = 13$ TeV, clusters directly at the electromagnetic scale are used for jet finding. The EM-scale jet energy resolution is worse than LCW, but at $\sqrt{s} = 13$ TeV, the global sequential calibration [125] (EM+JES) reduces these differences. An anti- k_t radius parameter of $R = 0.4$ is used for default jet clustering. Large-radius jets based on re-clustering (see Sec. 4.1.3) are used to identify boosted hadronically decaying top quarks and W bosons for high mass stops. These objects are discussed in more detail in Sec. 3.2.4. Due to the high energy nature of the target signal, the analysis is robust against the impact of pileup. Nonetheless, at $\sqrt{s} = 13$ TeV, the jet-vertex-tagger (JVT) [545] based on the tracks associated to the jet used for $p_T^{\text{jet}} < 50$ GeV to suppress spurious jets¹ from pileup interactions. The JVT is configured for a 92% hard-scatter jet efficiency, which corresponds to a pileup jet efficiency of about 1%. Fake jets may also be generated by non- pp collision processes such as calorimeter noise and beam-induced interactions with the imperfect beampipe vacuum. Such jets are readily identified by various quality criteria such as

¹For $p_T \sim 20$ GeV, about half the pileup jets are from random combinations of pileup interactions (stochastic pileup) while the other half are genuine quark and gluon jets (QCD pileup jets). The fraction of QCD pileup jets increase with p_T and the fraction of stochastic pileup jets increase with the number of pileup interactions.

the fraction of energy in the electromagnetic calorimeter compared with the fraction of the jet energy accounted for in reconstructed tracks [121, 546]. Since jets are used for many aspects of the event reconstruction, entire events are vetoed if any of these bad jets are identified. The resulting efficiency for hard-scatter events is higher than 99%.

Tracks and secondary/tertiary vertices associated with jets are also used to classify jets as resulting from b -quarks (b -tagged jets) using the MV1 (MV2c20) algorithm with 70% (77%) efficiency in simulated $t\bar{t}$ events [75, 547, 548] at $\sqrt{s} = 8$ (13) TeV. The ‘c20’ in the Run 2 version indicates that the background composition in the algorithm training had 20% charm-jets and 80% light-flavor jets. The 70% MV1 working point has a light-quark jets rejection ($=1/\text{efficiency}$) of about 140 and a charm-quark jet rejection of about 5. For the same b -quark jet efficiency, the MV2c20 algorithm improves the light-quark jet rejection by about a factor of 4 and the charm-quark jet rejection by about 50%. Part of this improvement came from the addition of the new pixel layer (IBL) and part from algorithmic improvements. Scale factors are applied in the simulation to correct for differences in the efficiency between data and simulation. The choice of the b -tagging working point was optimized for the search, as discussed in Sec. 4.

Jets and their associated tracks are additionally used to identify hadronically decaying τ leptons. The distribution of energy inside the jet is combined with tracking information to form a multivariate classifier separately for τ leptons decaying into one ($\sim 85\%$) or three charged pions ($\sim 15\%$) [134, 549, 550]. At $\sqrt{s} = 8$ TeV, a very loose working point was optimized that has a 3% efficiency for $t\bar{t}$ events without a hadronically decaying τ and 31% (37%) efficiency for $t\bar{t}$ events with a one- (three-) prong τ decay. This corresponds to a tight veto of 97% in events without a hadronically decaying τ . At $\sqrt{s} = 13$ TeV, tau jets are identified using the Loose identification algorithm [549, 550] which has a 60% and 50% efficiency for reconstructing one- and three-prong τ decays, respectively. In both Run 1 and Run 2, candidate hadronically decaying τ leptons are required to have no more than three tracks and if there three tracks, the sum of the track electric charges must be ± 1 . All signal region event selections in Part III require exactly one signal electron or muon; the τ candidate

must have opposite electric charge to these leptons, unless there are two tracks and the net charge is zero. Two track τ candidates are only permitted in the Run 1 analysis. There is no distinction between baseline and signal reconstructed hadronically decaying τ leptons.

One of the most important background processes to the stop search is the pair production of top quarks resulting in two charged leptons, where one is not identified as such. Therefore, it is advantageous to have the baseline lepton definition be as inclusive as possible to efficiently veto such events. The p_T threshold for baseline leptons is ≤ 10 GeV for both Run 1 and Run 2 and only loose quality criteria are imposed on the various track and electromagnetic calorimeter shower properties [131, 133, 551, 552]. For the same efficiency, the Run 2 electron identification has a $\sim 40\%$ larger background rejection due to the multivariate combination of reconstructed electron properties [551]. Most high p_T muons are measured by both the inner detector (ID) and the muon spectrometer (MS) resulting in combined muon candidates (CB). Muons beyond the ID acceptance are selected using muons reconstructed only with the MS (stand-alone, SA) and the efficiency is recovered for $|\eta| < 0.1$, where the MS is only partially instrumented due to calorimeter and ID services, by using ID tracks matched to either calorimeter energy deposits consistent with a minimum ionizing particle (calo-tagged, CT) or a track segment in the MS (segment-tagged, ST). The electrons or muons from stops are predicted to be significantly harder than in background processes and so the signal leptons are required to have $p_T > 25$ GeV. This high threshold is also useful for suppressing QCD multijet backgrounds and is required for the $\sqrt{s} = 8$ TeV analysis for the lepton trigger to be fully efficient. In addition, signal leptons must pass several quality criteria on their transverse (d_0) and longitudinal (z_0) impact parameters (IP). In particular, $|d_0| < 0.2$ mm and $z_0 < 1$ mm for muons at $\sqrt{s} = 8$ TeV and $d_0/\sigma_{d_0} < 5$ (3) and $z_0 \sin(\theta) < 0.5$ mm for electrons (muons) at $\sqrt{s} = 13$ TeV. Leptons from W boson decays are generally well-separated from other objects in the event, so other processes can be suppressed by imposing isolation criteria. These criteria are based on the scalar sum of the p_T from tracks (excluding the electron or muon track) within a cone around the lepton. When the top quarks from stop decays are produced with sufficient boost, the leptons are naturally close to the b -jet from the same top decay.

This is illustrated by Fig. 3.1, which shows the joint distribution of the $\Delta(\ell, \mathbf{b})$ and stop mass. To maintain efficiency for high stop masses, the isolation cone scales with the inverse of the lepton $\mathbf{p}_T$ (see Chapter 4).




Figure 3.1: The joint distribution of the ΔR between the lepton and $\mathbf{b}$ -quark from the same top quark decay for $\tilde{t} \rightarrow t\tilde{\chi}^0$. At $m_{\text{stop}} = 500$ GeV, about 2% (10%) of events have $\Delta R < 0.4$ (0.8) and at $m_{\text{stop}} = 800$ GeV, this fraction increases to $\sim 5\%$ (25%).

Explicitly reconstructed isolated photons are only considered if their $\mathbf{p}_T \gtrsim m_Z$ in order to mimic high $\mathbf{p}_T$ Z -bosons as described in Sec. 5.4.2. The actual threshold is chosen to be as high as possible while maintaining a sufficient event yield. If not explicitly reconstructed, isolated photons will be labeled as jets. At $\sqrt{s} = 13$ TeV, the threshold was slightly raised with respect to the Run 1 search due to the use of the photon trigger that is nearly 100% efficient² for $\mathbf{p}_T^\gamma = 125$ GeV. There are many sources of high $\mathbf{p}_T$ photons inside jets that are not useful for the $t\bar{t} + Z$ background estimation. The Run 1 analysis did not impose any explicit isolation requirements, but photons do participate in the overlap removal (see below). At $\sqrt{s} = 13$ TeV, in addition to the overlap removal, an explicit isolation is imposed on the calorimeter energy inside a ΔR cone around the photon. As with hadronically decaying τ leptons,

²The single lepton triggers are never fully efficient - see Sec. 2.3.

there is no difference between signal and baseline photons.

All of the objects described above are used to form the $\vec{p}_T^{\text{miss}}$. Jets, electrons, and muons that enter the calculation use the dedicated calibrations for those objects. Hadronically decaying τ leptons enter in the calculation as electrons or jets without the dedicated τ calibration. At $\sqrt{s} = 8$ TeV, calibrated photons with $p_T > 10$ GeV explicitly enter the $\vec{p}_T^{\text{miss}}$ calculation. For Run 2, only the high p_T photons for the $t\bar{t} + \gamma$ control region described in Sec. 5.4.2 enter the calculation as photons; otherwise photons are part of the jet collection. The main difference between Runs 1 and 2 is that the former uses a calorimeter-based soft-term [136, 137] for the unclustered energy while the latter uses a track-based soft-term [138, 139]. Information about soft neutral particles is lost in the track-based soft term, but tracks are largely insensitive to pileup. For the early Run 2 conditions, the insensitivity to pileup is the dominant effect and the E_T^{miss} with a track-based soft-term has a $\sim 10\%$ better resolution than the calorimeter-based soft-term.

Object	$p_T > [\text{GeV}]$	$ \eta <$	Identification	Isolation	Other
Baseline Jets	20	2.8	Looser* [121]	–	LCW
	20	none	BadLoose [546]	–	EM+GSC
Signal Jets	25	2.5	same as baseline	–	
	25	2.5	JVT@92% [545]	–	
b-tagged Jets	25	2.5	MV1@70% [75]	–	
	25	2.5	MV2c20@77% [547, 548]	–	
Baseline Electron	10	2.47	Loose [551]	none	
	7	2.47	VeryLooseLH [552]	none	
Signal Electron	25	2.47	Tight [551]	$p_T^{\Delta R < 0.2} / p_T^e < 0.1$	
	25	2.47	LooseLH [552]+IP	$p_T^{\Delta R < \max\{10 \text{ GeV}/p_T^e, 0.2\}} @99\%$	
Baseline Muon	10	2.4	CB+ST [131]	none	
	6	2.7	Loose CB+ST+CT+SA [133]	none	
Signal Muon	25	2.4	IP (see the text)	$p_T^{\Delta R < 0.2} < 1.8 \text{ GeV}$	
	25	2.7	IP (see the text)	$p_T^{\Delta R < \max\{10 \text{ GeV}/p_T^\mu, 0.3\}} @99\%$	
Photons	125	2.37	Tight [553, 554]	none (see OR)	
	100	2.37	Tight [553, 554]	$E_T^{\Delta R < 0.4} / < 0.022 p_T^\gamma + 2.45/\text{GeV}$	
Hadronic τ	15	2.47	Jet BDT@31%-37% [555, 556]	none	
	20	2.5	Loose [549, 550]	none	
E_T^{miss}	–	–	–	–	CST [136]
	–	–	–	–	TST [139]

Table 3.1: An overview of the objects definitions used for the stop search. The color blue indicates the criteria at $\sqrt{s} = 13 \text{ TeV}$ whereas black is for $\sqrt{s} = 8 \text{ TeV}$. See the text for details.

Baseline and signal jets and leptons considered for further use must pass an object ambiguity resolving algorithm. These algorithms are tailored for the stop search and optimized using particle-level labels in simulation to maintain a low mis-classification rate in order to increase the one lepton signal and reduce the significant two-lepton background. Tables 3.2 and 3.3 summarize these overlap removal (OR) procedures at $\sqrt{s} = 8$ TeV and $\sqrt{s} = 13$ TeV, respectively. For example, if an electron and a jet overlap with $\Delta R < 0.2$ and the jet is not **b**-tagged, then the object is interpreted as an electron and the overlapping jet is removed from the list of jets. Overlap between leptons and **b**-jets is treated differently than inclusive jets because semi-leptonic B decays can naturally produce overlapping leptons. At high stop mass, signal leptons can also naturally be close to jets. For this reason, at $\sqrt{s} = 13$ TeV, a p_T -dependent ΔR cone is used to remove leptons in favor of jets. Another change from Run 1 to Run 2 is the muon / non-**b** jet overlap condition which only applies if the jet has less than three associated tracks with $p_T > 500$ MeV or $p_T^{\text{muon}}/p_T^{\text{jet}} > 0.7$. The track requirement removes jets that are seeded by muon radiation and the p_T asymmetry requirement identifies jets that are unassociated with the muon, which is unlike the case for electrons that deposit most of their energy in the calorimeter and thus to the close-by jet. Due to the use of calo-tagged muons at $\sqrt{s} = 13$ TeV, there can be some overlap between electrons and muons near $|\eta| \approx 0$. In nearly all cases, the simulation predicts that these objects are due to true electrons. The overlap between electrons and photons has a significant improvement in the photon purity in the $t\bar{t} + \gamma$ control region (Sec. 5.4.2).

Object 1	e	μ	l	γ	τ
Object 2	j	j	j	j	l
$\Delta R <$	0.2	0.4	0.4	0.2	0.2
Condition	j not b-tagged		–	–	–
Resolution	e	μ	j	γ^*	l

Table 3.2: A summary of the procedure to resolve ambiguous object labels at $\sqrt{s} = 8$ TeV. The first two rows list the overlapping objects: electrons (e), muons (μ), electron or muon (l), jets (j), photons (γ), and hadronically decaying τ leptons (τ). The third and fourth rows give the overlap condition and the last row lists which label is given to the ambiguous object. The procedure is applied from left to right. (*) Jets are removed only in the $t\bar{t} + \gamma$ validation region (see Sec. 5.4.2).

Object 1	e	e	μ	l	γ	γ	τ
Object 2	μ	j	j	j	j	e	e
$\Delta R <$	0.1	0.2	0.2	$\min\left(0.4, 0.04 + \frac{10}{p_T^l}\right)$		0.2	0.1
Condition	CT μ j not b-tagged		j not b-tagged and $n_{\text{track}}^j < 3$ or $\frac{p_T^\mu}{p_T^j} > 0.7$		–	–	–
Resolution	e	e	μ	j	γ	e	e

Table 3.3: A summary of the procedure to resolve ambiguous object labels at $\sqrt{s} = 13$ TeV. The first two rows list the overlapping objects: electrons (e), muons (μ), electron or muon (l), jets (j), photons (γ), and hadronically decaying τ leptons (τ). The third and fourth rows give the overlap condition and the last row lists which label is given to the ambiguous object. The procedure is applied from left to right.

3.2 Discriminating Variables

The key to a powerful and robust search is the use of a relatively small number³ of highly discriminating variables. There are two strategies when developing variables. One possibility is to target particular aspects of the background that are absent in the signal (*veto*). A second tactic focuses on properties of the signal that are absent in the background (*tag*). This section explores a series of veto and tag variables, many of which are specifically designed for the stop search and used here for the first time. One of the key themes in the development of the variables is a focus to use *tailored variables* when possible. Many simple variables such as $H_T = \sum_i p_{T, \text{jet } i}$ and $m_{\text{eff}} = E_T^{\text{miss}} + H_T + p_T^{\text{lepton}}$ are generically useful for signatures with high multiplicity final states involving multiple missing particles. However, the price of simplicity is sub-optimality in particular situations such as the $t\bar{t} + E_T^{\text{miss}}$ topology. One of the most powerful variables is the transverse mass (m_T), which has been mentioned at several points in earlier chapters. Section 3.2.1 describes m_T in detail in order to demonstrate exactly how and why it is useful for the stop search. A generalization of m_T to cases when there are multiple missing particles is the m_{T2} family of observables (Sec. 3.2.1.2). Background events can be reconstructed with large m_T and m_{T2} when jets are sufficiently mis-measured. Section 3.2.2 described how resolution information can be incorporated into kinematic variables to suppress these events. Another veto variable that can be combined with kinematic information is hadronic tau identification (Sec. 3.2.3). A significant fraction of $t\bar{t}$ contain a hadronically decaying τ to pass harsh requirements on m_T . The section ends with a brief discussion of boosted top quark and W boson tagging techniques (Sec. 3.2.4). Some of the techniques already described in Chapter 4 are directly applicable to the stop search.

³An alternative paradigm is to process all available information using sophisticated machine learning techniques (see Sec. 4.3). This can be a powerful approach, but for *tail searches* such as this one, by construction there is little data available near the signal region to thoroughly validate such methods. Therefore, robust and powerful methods grounded in physical intuition are preferred.

3.2.1 Transverse Mass Variables

One of the most striking characteristics of top quarks and the targeted particles in theories beyond the SM is their large mass. If all particles from these heavy particle decays could be reconstructed and unambiguously identified, the invariant mass would be a powerful variable. However, due to neutrinos and neutralinos, a significant fraction of the resonance mass goes into undetected energy. At a hadron collider, the $\sqrt{s}$ is unknown and therefore only the sum of the transverse momentum of the undetected particles can be inferred. Transverse mass variables are modifications of the usual invariant mass to cases where there is at least one undetected particle and the total longitudinal momentum is unknown. Even though longitudinal information is missing, the transverse mass of the decay products of massive particles tends to be higher than for background processes. This fact was first used to discover the W boson at CERN by the UA1 [59] and UA2 [61] collaborations. The left plot of Fig. 3.2 shows the transverse mass spectrum for the first six W boson candidate events.




Figure 3.2: Left: The transverse mass distribution for the discovery of the W boson by the UA1 collaboration [59]. Right: the transverse mass distribution used by the CDF collaboration for one of the most precise determinations of the W boson mass [557].

The most important feature of transverse mass variables is that they are bounded by the resonance mass and tend to have a probability distribution that is concentrated near the kinematic endpoint. These features allowed the UA1 and UA2 collaborations to estimate the mass of the W boson and is still used today for the most precise determination of the W boson mass by the Tevatron collaborations [557, 558]. The

right plot of Fig. 3.2 shows an example from CDF [557] with $\mathcal{O}(10^6)$ events. This bounded property of the transverse mass variables helps to suppress the W boson and top quark backgrounds that have transverse mass well above other backgrounds, but well below the scale of new physics. Section 3.2.1.1 contains a detailed description of the transverse mass in topologies with one missing particle.




Figure 3.3: The distribution of m_{T2} in dilepton $t\bar{t}$ events using the leptons (b -tagged jets) as the visible particles in the left (right) plot. These distributions are used to measure the top quark, W boson, and neutrino masses. Published in Ref. [384].

When a decay chain has multiple missing particles, there is an ambiguity in the assignment of momentum to each particle because only the sum of their transverse momentum is measured by $\vec{p}_T^{\text{miss}}$. However, in this case there are still ways to bound the parent particle mass by computing the minimum mass consistent with the observed decay products. This is basis for a generalization of the transverse mass called m_{T2} and is described in detail in Sec. 3.2.1.2. Like the simple transverse mass described above, m_{T2} will have the property that for $t\bar{t}$ events it is relative large, but bounded well below the scale for signal. Figure 3.3 shows two m_{T2} variables in dileptonic $t\bar{t}$ events. The left (right) plot of Fig. 3.3 is bounded by m_W (m_{top}), a fact which is used to measure these masses just as the simple transverse mass is used to measure m_W in inclusive W +jets events.

3.2.1.1 Transverse Mass

Consider two particles with four-momentum $\mathbf{p}^\mu = (p_x, p_y, p_z, E)$ and $\mathbf{q}^\mu = (q_x, q_y, q_z, F)$. Define $\tilde{\mathbf{p}}^\mu = (p_x, p_y, 0, \sqrt{E^2 - p_z^2})$ and $\tilde{\mathbf{q}}^\mu = (q_x, q_y, 0, \sqrt{F^2 - q_z^2})$. The *transverse mass* is defined by $m_T^2 = (\tilde{\mathbf{p}}^\mu + \tilde{\mathbf{q}}^\mu)(\tilde{\mathbf{p}}_\mu + \tilde{\mathbf{q}}_\mu)$. By construction, the transverse mass is invariant under longitudinal boosts. If the particles have masses m_p and m_q and are the decay products of a two-body resonance decay with mass M then,

$$\begin{aligned}
 M^2 &= m_p^2 + m_q^2 + 2(EF - p_z q_z - p_T q_T) \\
 &= m_p^2 + m_q^2 + 2(E_T F_T \cosh(\Delta y) - \mathbf{p}_T \cdot \mathbf{q}_T) \\
 &\geq m_p^2 + m_q^2 + 2(E_T F_T - \mathbf{p}_T \cdot \mathbf{q}_T) \\
 &= \left(\sqrt{E^2 - p_z^2} + \sqrt{F^2 - q_z^2} \right)^2 - (p_x + q_x)^2 - (p_y + q_y)^2 \\
 &= m_T^2,
 \end{aligned} \tag{3.1}$$

where $E_T^2 = m_p^2 + p_T^2$ and Δy is the difference in rapidity of the two particles. The second line follows from trigonometry identities and tedious arithmetic to show that $p_z = E_T \sinh(y)$ ($p_z = p_T \sinh(\eta)$) and so $E = E_T \cosh(y)$ ($|p| = p_T \cosh(\eta)$). Suppose that the resonance is a W boson decaying into $W \rightarrow e\nu$. Since $m_e, m_\nu \ll m_W$, $E \approx |\vec{p}|$ and $F \approx |\vec{q}|$ is an excellent approximation. Therefore, $m_T^2 \approx 2p_{T,e}p_{T,\nu}(1 - \cos \theta_{e\nu})$, where $\theta_{e\nu}$ is the angle between $\vec{p}_{T,2}$ and $\vec{p}_{T,\nu}$. Since the neutrino is not directly detected, the definition of the transverse mass used in practice is

$$m_T^2 = 2E_T^{\text{miss}} p_T^{\text{lepton}} \left(1 - \cos \Delta\phi \left(\vec{E}_T^{\text{miss}}, \vec{p}_T^{\text{lepton}} \right) \right). \tag{3.2}$$

In the W boson rest frame, $p_{T,e} = p_{T,\nu} = \frac{m_W}{2} |\sin \theta|$, where θ is the angle of the electron-neutrino axis from the z -axis. With this formulation, it is clear that $m_T \leq m_W$ with equality only if $\theta = \pm\pi/2$. Since the differential volume element in spherical coordinates is $dV = \rho^2 \sin \theta d\rho d\theta d\phi = \rho^2 d\rho d(\cos \theta) d\phi$, an isotropic distribution for

the decay⁴ in the W rest frame results in the following probability distribution for $\mathbf{m}_T$:

$$\begin{aligned}
 f_{\mathbf{m}_T}(\mathbf{m}_T) &= \sum_{\theta \geq \pi/2, \theta < \pi/2} \left| \frac{\partial}{\partial \mathbf{m}_T} \mathbf{m}_T^{-1}(\mathbf{m}_T) \right| f_{\cos \theta}(\mathbf{m}_T^{-1}(\mathbf{m}_T)) \\
 &= \sum_{\theta \geq \pi/2, \theta < \pi/2} \frac{1}{\mathbf{m}_W} \frac{(\mathbf{m}_T/\mathbf{m}_W)}{\sqrt{1 - (\mathbf{m}_T/\mathbf{m}_W)^2}} \times \frac{1}{2} \\
 &= \frac{1}{\mathbf{m}_W} \frac{(\mathbf{m}_T/\mathbf{m}_W)}{\sqrt{1 - (\mathbf{m}_T/\mathbf{m}_W)^2}}
 \end{aligned} \tag{3.3}$$

where $\mathbf{m}_T(x) = \mathbf{m}_W \sqrt{1 - x^2}$ so $\mathbf{m}_T^{-1}(y) = \pm \sqrt{1 - (y/\mathbf{m}_W)^2}$, one for each branch: $\theta \geq \pi/2, \theta < \pi/2$. The probability distribution $f_{\cos \theta} = \frac{1}{2}$ because $\cos \theta$ is uniform on $[-1, 1]$. The sum in Eq. 3.3 is over the two values of $\cos \theta$ that result in the same $\mathbf{m}_T$ value. The most striking feature of Eq. 3.3 as a function is that it is monotonically increasing with its maximum at $\mathbf{m}_T = \mathbf{m}_W$. Since the term in the absolute value in the first line of Eq. 3.3 is the Jacobian of the variable transformation, this peak is called the *Jacobian peak*. The left plot of Fig. 3.4 shows the distribution of the transverse mass for a W boson produced at rest. By construction, the analytical formula derived in Eq. 3.3 is identical to the red filled histogram for simulated (scalar) $W \rightarrow e\nu$ events. For comparison, additional distributions are shown in Fig. 3.4 for cases in which there are additional neutrinos in the event. In these cases, Eq. 3.2 is used as the definition of the transverse mass, where the E_T^{miss} includes all non-reconstructed particle momenta. For leptonic τ decays $W \rightarrow \tau\nu_\tau, \tau \rightarrow e\nu_e\nu_\tau$, the endpoint of the $\mathbf{m}_T$ distribution is still $\mathbf{m}_W$, but the probability distribution is mostly concentrated at low values of the transverse mass. Since events with only one true lepton-neutrino pair have an $\mathbf{m}_T \leq \mathbf{m}_W$, most of the events with $\mathbf{m}_T > \mathbf{m}_W$ originate from events

⁴The scalar decay is used for illustration; in reality the W boson is a spin-1 particle. In the W boson rest frame, $dN/d\cos \theta \propto (1 \pm \cos \theta)^2$ for transversely polarized bosons (with spin ± 1) and $dN/d\cos \theta \propto \sin^2 \theta$ for longitudinally polarized bosons (for the spin axis along z). These correspond to $f_{\cos \theta}(x) = 3(1 \pm x)^2/8$ and $f_{\cos \theta}(x) = 3(1 - x^2)/4$, respectively. W bosons produced in top decays are mostly longitudinal whereas inclusive W boson production results in mostly transversely polarized bosons. See Appendix B.2.

with two true leptons, one of which is out of acceptance (lost) or reconstructed as a jet or is a hadronically decaying τ lepton (mis-id). When the second lepton is lost, its momentum is part of the E_T^{miss} while if it is mis-identified, its momentum is not directly⁵ part of the E_T^{miss} . Dilepton events are simulated by independently generating two W bosons. When these bosons are produced at rest, the momentum of the lost lepton cancels with its neutrino pair and so the m_T spectrum is identical to the single $W \rightarrow e\nu$ case. However, if the second lepton is mis-identified, then the m_T distribution can exceed m_W . The kinematic maximum is achieved when the decay axis of both $e\nu$ pairs are parallel and in the transverse plane. In that case, $m_T^{\text{max}} = \sqrt{2}m_W \approx 113$ GeV. Figure 3.5 shows the distribution of m_T in simulated events with a full spin-1 W boson. All cases have the Jacobian peak, but the tail of the distribution toward zero depends on the spin.

While the m_T distribution is invariant under longitudinal boosts, it is not invariant under transverse boosts. To illustrate the impact of a transverse boost, begin with a configuration close to the kinematic limit: $p^\mu = (\epsilon, m_W/2, 0, m_W/2) + \mathcal{O}(\epsilon^2/m_W)$ and $q^\mu = (-\epsilon, -m_W/2, 0, m_W/2) + \mathcal{O}(\epsilon^2/m_W)$. Under a boost along the x -axis with magnitude β , $p_x \rightarrow \gamma\epsilon + \gamma\beta \frac{m_W}{2}$. Therefore,

$$\begin{aligned} m_T^2 &= (p_T + q_T)^2 - (p_x + q_x)^2 - (p_y + q_y)^2 \\ &= m_W + \epsilon\gamma\beta m_W + \mathcal{O}\left(\frac{\epsilon^2}{m_W}, \beta^2 m_W^2\right) \\ &= m_W^2 \left(1 - \frac{\epsilon\beta}{m_W} + \mathcal{O}\left(\frac{\epsilon^2}{m_W^2}, \beta^2\right)\right). \end{aligned} \quad (3.4)$$

The interesting properties of Eq. 3.4 are that there is no impact of a boost if the maximum value is already achieved ($\epsilon = 0$) and if $\epsilon \neq 0$, a transverse boost *reduces* the transverse mass. As the W boson becomes more boosted, the momentum of the electron and neutrino increase, but the angle between them decreases. The later happens faster than the former, which flattens out the m_T distribution at high p_T .

⁵All the visible momenta are used to construct the E_T^{miss} , so these mis-identified leptons will contribute indirectly.




Figure 3.4: The distribution of the transverse mass for the isotropic decay of a W boson for $p_T^W = 0$ GeV on the left and $p_T^W = 200$ GeV on the right. The red histogram shows the m_T distribution for the decay $W \rightarrow e\nu$ and the blue histogram shows the distribution for $W \rightarrow \tau\nu$ where the τ decays to an electron and neutrinos. The black histograms correspond to the pair production of $W \rightarrow e\nu$ where the second lepton is either included (lost) or not included (mis-id) in the E_T^{miss} . The dotted line is the analytic formula derived in the text.

The right plot of Fig. 3.4 shows the m_T distribution for the same configurations as the left plot, but now with $p_T^W = 200$ GeV. The kinematic maximum for $W \rightarrow e\nu$ is still $m_T \leq m_W$, but the distribution has filled in at lower values of the transverse mass. For the pair production of W bosons, both have the same p_T but an arbitrary direction. Since the angle between the W bosons can be large, the transverse mass can significantly exceed m_W .

Figure 3.4 showed that if there are anomalous contributions to the E_T^{miss} , the transverse mass can exceed its natural kinematic maximum. Another way for $m_T > m_W$ in events with one reconstructed lepton is if there are additional genuine contributions to the E_T^{miss} . For example, the high energy undetected neutralinos in stop events can push m_T well beyond m_W . Figure 3.6 shows the m_T distribution for simulated signal $\tilde{t} \rightarrow t\tilde{\chi}^0$ events and dilepton $t\bar{t}$ events where one lepton is lost or mis-identified. The




Figure 3.5: The transverse mass distribution for a scalar W decay (spin 0) and a vector W decay (spin 1). Each distribution is generated by MG5_aMC 2.1.1. A Higgs boson with $m_H = m_W$ is used for the spin 0 line, which is identical to the lines from Fig. 3.4 using a phase-space only generator and the analytic formula.

stops are produced at rest, resulting in a kinematic maximum of $m_T \leq \sqrt{2}m_{\tilde{t}}$, which occurs when the neutralino-top axes are aligned and in the transverse plane. The left and right plots of Fig. 3.6 differ in the source of the top quark boost. In the left plot, the invariant mass of the $t\bar{t}$ system is zero, but its p_T is significant. Conversely, in the right plot of Fig. 3.6, the $t\bar{t}$ system is produced at rest, but with significant invariant mass. In both cases, the typical top quark p_T is half of the p_T or mass scale. However, the m_T distribution is significantly different between the $p_T^{\tilde{t}\bar{t}}$ and $m_{\tilde{t}\bar{t}}$ schemes. The $m_{\tilde{t}\bar{t}}$ case is most similar to Fig. 3.4 where the W bosons can have a significant angle between them, while in the $p_T^{\tilde{t}\bar{t}}$ case, the two W bosons are spatially close. The larger opening angle results in a larger m_T value. In practice, both $p_T^{\tilde{t}\bar{t}}$ and $m_{\tilde{t}\bar{t}}$ will be nonzero, but the former is more important for the stop search. This is because a four-jet event selection requires dilepton events to be produced with additional jets, as discussed in more detail in Sec. 5.




Figure 3.6: The distribution of the transverse mass for signal $\tilde{t} \rightarrow t\tilde{\chi}^0$ events and dilepton $t\bar{t}$ events where one lepton is lost or mis-identified. The signal histograms are identical in the left and right plots. In the left plot, the p_T of the $t\bar{t}$ system is boosted in an arbitrary direction with a fixed p_T . In the right plot, the $t\bar{t}$ pair is produced at rest, but with a large $m_{t\bar{t}}$ so that the top quarks have significant boost.

The right plot of Fig. 3.6 shows that there is a strong correlation between the E_T^{miss} and m_T in events where the two top quarks are independent. Figure 3.7 quantifies the correlation for the various event types discussed above. The correlation increases with mass for stop events and $t\bar{t}$ events with increasing $m_{t\bar{t}}$, reaching about 60%. In contrast, when there is a relationship between the direction of the two top quarks, as is the case when the entire $t\bar{t}$ system is boosted, there is little correlation between E_T^{miss} and m_T . These differences in correlation show that the m_T can add useful information beyond what is already contained in the E_T^{miss} . Note that the correlation is 100% for a single $W \rightarrow e\nu$ produced at rest because both the m_T and $p_{T,\nu} \sim \cos(\theta)$.

For all the reasons described above, the transverse mass is a powerful variable at suppressing the single lepton $t\bar{t}$ and W +jets backgrounds. Figure 3.8 shows the transverse mass distribution in the early Run 2 data for a selection of events enriched in the pair production of top quarks. In particular, events are required to have exactly one reconstructed signal lepton with no additional baseline leptons, four jets




Figure 3.7: The linear correlation between the E_T^{miss} and m_T as a function of the p_T or mass scale for various processes. The $p_T^{t\bar{t}}$ and $m_{t\bar{t}}$ are setup in the same way as for Fig. 3.20. The correlation is set to zero for unphysical parameter values such as $m_{t\bar{t}} < 2m_t$.

with $p_T > 25$ GeV and at least one b -tagged jet. The resulting events are predicted to have a $t\bar{t}$ purity of about 80%. As expected, most of the $t\bar{t}$ events have one real lepton, due to the second lepton veto. Therefore, there is a clear cutoff near the W boson mass. However, the right plot of Fig. 3.8 shows that at high values of m_T , events with two real leptons dominate, split between events with a second real electron or muon and events with a hadronically decaying τ . The single lepton background at high m_T is negligible. Figure 3.9 shows the analogous plots for events enriched in the single production of a W boson in association with jets. The event selection for Fig. 3.9 differs from the one used for Fig. 3.8 only by exchanging the b -jet requirement for a b -jet veto. The Jacobian peak is clearly present in both the log and linear scale plots of Fig. 3.9. For $m_T > m_W$, there is an enhancement of diboson events with a second real lepton, but this is not as significant as for the $t\bar{t}$ case. This is because the diboson to W +jets cross section ratio⁶ is $\mathcal{O}(1\%)$ [385, 559] while the dilepton branching ratio

⁶The ratio is even smaller when including the leptonic branching ratio for the second boson.

is about 25% of the single lepton $t\bar{t}$ branching ratio. Since the W +jets events with $m_T \gg m_W$ are well beyond the parton-level kinematic maximum, the shape of the distribution in the high m_T tail is determined mostly by resolution effects. Methods for incorporating resolution information into kinematic variables will be described in Sec. 3.2.2.




Figure 3.8: A comparison of data and simulation using a loose selection requiring exactly one signal lepton, four jets with $p_T > 25$ GeV and at least one b -tagged jet. The left and right plots differ only in the scaling of the vertical axis. The uncertainty band includes jet energy scale and resolution uncertainties (see Sec. 6).

The transverse mass constructed above assumes that the mass of the missing particle is known a priori. This is a valid assumption when the missing object is a SM particle such as the neutrino. However, in general, the full m_T as defined by Eq. 3.1 (and not Eq. 3.2) implicitly requires the input of a missing particle mass, otherwise $F^2 - q_z^2 = q_x^2 + q_y^2 + m_{\text{test}}^2$ cannot be computed from only transverse quantities. This will be most relevant for generalizing the transverse mass in Sec. 3.2.1.2, but one can already see the consequences of a non-trivial test mass for the case of a single W boson. Figure 3.12 shows the impact of choosing $m_{\text{test}} \neq m_{\text{true}}$. The minimum value of m_T is no longer zero - the transverse mass is bounded below by m_{test} . This happens when the transverse momentum of the decay products is zero. The most




Figure 3.9: A comparison of data and simulation using a loose selection requiring exactly one signal lepton, four jets with $p_T > 25$ GeV and at exactly no b -tagged jets. The left and right plots differ only in the scaling of the vertical axis. The uncertainty band includes jet energy scale and resolution uncertainties (see Sec. 6).

important change is that there is no longer a p_T -independent kinematic limit. In the context of $W \rightarrow e\nu$, the generalized transverse mass can be written

$$m_T^2 = m_{\text{test}}^2 + 2 \left(p_{T,l} \sqrt{m_{\text{test}}^2 + p_{T,\nu}^2} - p_{T,l} p_{T,\nu} \cos(\theta_{e\nu}) \right). \quad (3.5)$$

Eq. 3.5 is maximized when $\theta_{e\nu} = \pi$ and the electron and neutrino momentum are in the transverse plane. This was clear earlier in the context of a W boson at rest, but for a boosted W boson the implication is that the boost must be collinear with the electron-neutrino axis. A boost along the electron-neutrino axis will enhance the momentum of the electron relative to the neutrino or vice versa since in the W rest frame, the two are back-to-back. Since the neutrino momentum is added in quadrature with the test mass, the maximum value of m_T is achieved when the boost is parallel to the electron direction. If the boost has magnitude β :

$$\mathbf{p}_{T,l} \rightarrow \gamma \mathbf{p}_{T,l} + \beta \gamma \mathbf{p}_{T,l} = \gamma \mathbf{p}_{T,l} (1 + \beta) = \frac{m_W}{2} \sqrt{\frac{1 + \beta}{1 - \beta}}, \quad (3.6)$$

where without loss of generality, one can take the electron momentum to be aligned with the x-axis. A similar calculation shows that $\mathbf{p}_{T,\nu} \rightarrow \frac{m_W}{2} \sqrt{\frac{1 - \beta}{1 + \beta}}$. If $m_{\text{test}} = 0$, these two factors exactly cancel and the endpoint is invariant under transverse boosts, as was observed earlier. However, when $m_{\text{test}} > 0$, the two boost factors do not cancel and therefore the endpoint scales with the boost. This is a general feature of the transverse mass whenever $m_{\text{test}} \neq m_{\text{true}}$ and will be investigated further in Sec. 3.2.1.2.




Figure 3.10: Identical m_T distributions as in Fig. 3.4, except the neutrino mass parameter in the transverse mass calculation is set to 20 GeV (left) and 50 GeV (right)

3.2.1.2 Multiple Missing Particles: m_{T2}

Sec. 3.2.1.1 demonstrated that a simple threshold requirement $m_T \gg m_W$ is an effective procedure for suppressing the single lepton $t\bar{t}$ and W +jets backgrounds while maintaining high stop signal efficiency. One of the dominant residual backgrounds is $t\bar{t}$ with two real leptons where the second lepton is out of acceptance, fails the particle identification, or is a hadronically decaying τ . This section describes an extension of the transverse mass to cases where there are multiple missing particles in order to reduce the two lepton backgrounds with multiple neutrinos.

Figure 3.11 shows the generic setup: two particles P with the same mass m_P are pair produced and decay to visible particles V_1, V_2 and undetected particles C . The momenta of V_1 and V_2 are measured in the detector but only the sum of the transverse momenta of the C particles are inferred from momentum conservation. Define the variable m_{T2} ⁷ [560, 561] as


$$m_{T2}(m_{C_{1,\text{test}}}, m_{C_{2,\text{test}}}) = \min_{\substack{\vec{p}_{T,\text{test}}^{C_1} + \vec{p}_{T,\text{test}}^{C_2} = \vec{E}_T^{\text{miss}}}} \max_i \left\{ m_T \left(\vec{p}_T^{V_i}, \vec{p}_{T,\text{test}}^{C_i}, m_{V_i}, m_{C_{i,\text{test}}} \right) \right\}, \quad (3.7)$$

where $m_T(\vec{p}, \vec{q}, m_{\text{vis}}, m_{\text{test}})$ is the generalized transverse mass introduced in Eq. 3.1 and m_{test} is the assumed mass for the invisible particle C_i . The *test* in $\vec{p}_{T,\text{test}}^{C_i}$ and $m_{C_{i,\text{test}}}$ is to distinguish the dummy variable in the minimization from the true and unknown values of $\vec{p}_T^{C_i}$ and m_{C_i} . In a form similar to Eq. 3.5 but including mass effects for the visible particle, the generalized transverse mass takes the form:

$$m_T^2(\vec{p}, \vec{q}, m_{\text{vis}}, m_{\text{test}}) = m_{\text{vis}}^2 + m_{\text{test}}^2 + 2 \left(\sqrt{m_{\text{vis}}^2 + p^2} \sqrt{m_{\text{test}}^2 + q^2} - \vec{p} \cdot \vec{q} \right). \quad (3.8)$$

Like the simple transverse mass, m_{T2} is invariant under longitudinal boosts but not transverse boosts. When the masses $m_{C_{1,\text{test}}}, m_{C_{2,\text{test}}}$ are chosen correctly, $m_{T2} \leq$

⁷Pronounced ‘M-Tee-Two’ and in the literature is also referred to as the ‘stransverse mass’ or the ‘Cambridge m_{T2} ’.

Figure 3.11: A schematic diagram showing the generic m_{T2} setup.

m_P . This is because $m_{T2} \leq \max_i \left\{ m_T \left(\vec{p}_T^{V_i}, \vec{p}_T^{C_i}, m_{V_i}, m_{C_i} \right) \right\} \leq m_P$, as shown in Sec. 3.2.1.1. Figure 3.12 demonstrates the power of m_{T2} in events with two W bosons where one W decays into a hadronically decaying τ and the other W boson decays into an electron or muon (ℓ). Multiple neutrinos contribute to the missing momentum so the m_T of the ℓ and the $\vec{p}_T^{\text{miss}}$ can exceed the m_W bound as long as the bosons are not produced at rest. In contrast, the m_{T2} using $V_1 = \ell$, $V_2 = \tau$, $m_{C_1} = m_{C_2} = 0$ is kinematically bound by m_W (neglecting m_τ). Unlike the m_T case, the m_{T2} endpoint is not as saturated due to the minimization in the definition.

The visible objects V_1 and V_2 from Fig. 3.11 are often *composite systems of particles*. For example, V_i might be the combination of a b-jet and a lepton. In such cases, it is often true that $m_{V_1} \neq m_{V_2}$. When the composite systems are not even made of the same types of constituent particles, the presumed lost children may not be the same and in general $m_{C_1} \neq m_{C_2}$. Just like the generalized m_T from Sec. 3.2.1.1, the general m_{T2} variable is bounded from below by $\max\{m_{C_1} + m_{V_1}, m_{C_2} + m_{V_2}\}$. This is readily calculated by taking the derivative of each m_T branch with respect to the assigned missing particle momentum $p_{\alpha, \text{test}}^{C_i}$ for $\alpha \in \{x, y\}$, shown in Eq. 3.9. The second implication (2) in Eq. 3.9 is from summing and squaring the first implication (1) and the third implication (3) is the result of simplifying after inserting the second implication (2) back into the first one (1).




Figure 3.12: A comparison of m_T and m_{T2} for events with an identified hadronically decaying τ .

$$\begin{aligned}
 \frac{1}{2} \frac{\partial}{\partial p_{\alpha, \text{test}}^{C_i}} m_T^2(\vec{p}_T^{V_i}, \vec{p}_{T, \text{test}}^{C_i}, m_{V_i}, m_{C_i, \text{test}}) &= \frac{p_{\alpha, \text{test}}^{C_i} \sqrt{m_{V_i}^2 + (p_T^{V_i})^2}}{\sqrt{m_{C_i}^2 + (p_{T, \text{test}}^{C_i})^2}} - p_{\alpha}^{V_i} \\
 &\stackrel{(1)}{\Rightarrow} p_{\alpha, \text{test}}^{C_i} = p_{\alpha}^{V_i} \frac{\sqrt{m_{C_i}^2 + (p_{T, \text{test}}^{C_i})^2}}{\sqrt{m_{V_i}^2 + (p_T^{V_i})^2}} \stackrel{(2)}{\Rightarrow} p_{T, \text{test}}^{C_i} = \frac{m_{C_i}}{m_{V_i}} p_T^{V_i} \\
 &\stackrel{(3)}{\Rightarrow} p_{\alpha, \text{test}}^{C_i} = \frac{m_{C_i}}{m_{V_i}} p_{\alpha}^{V_i} \stackrel{(4)}{\Rightarrow} m_{T, \text{min}} = m_{V_i} + m_{C_i}.
 \end{aligned} \tag{3.9}$$

In addition to the flexibility to choose the particles composing V_1 and V_2 as well as the test masses m_{C_i} , there is an extensive literature on variations of m_{T2} and related variables that aim to solve the same multiple-missing-particle problem. See Ref. [562, 563] for a extensive reviews of the existing methods. One important variation is the *perpendicular m_{T2} variables*. Unlike the simple transverse mass, m_{T2} generally does not have a closed-form solution to the minimization in Eq. 3.8. Numerical techniques for computing m_{T2} are described in Sec. 3.2.1.2.2, but first Sec. 3.2.1.2.1 documents the perpendicular m_{T2} which does have a closed-form solution.

3.2.1.2.1 Perpendicular m_{T2}

Define the upstream transverse momentum as $\vec{p}_T^{\text{up}} = -\vec{p}_T^{V_1} - \vec{p}_T^{V_2} - \vec{p}_T^{\text{miss}}$, i.e. all transverse momentum aside from the momenta from V_1 and V_2 . The perpendicular momenta $\vec{p}_{T,\perp} = \vec{p}_T - (\hat{p}^{\text{up}} \cdot \vec{p}_T) \hat{p}^{\text{up}}$, where $\hat{p}^{\text{up}} = \vec{p}_T^{\text{up}} / |\vec{p}_T^{\text{up}}|$. In a topology like the one used to make Fig. 3.12 but with $t\bar{t}$ production (Fig. 3.13), the upstream momentum includes the b -jets directly from the top quark decay and any ISR radiated prior to the $t\bar{t}$ production. At the end of Sec. 3.2.1.1, Fig. 3.12 showed that when the test mass m_{C_i} is not equal to the mass of the missing object, the presence of nonzero upstream momentum makes the kinematic maximum $\vec{p}_T^{\text{up}}$ -dependent. By constructing a m_{T2} variable with only perpendicular momenta, $m_{T2,\perp}$ [564], the $\vec{p}_T^{\text{up}}$ -independence of the kinematic maximum is restored because $\vec{p}_{T,\perp}^{\text{up}} = \vec{0}$ by construction.

In addition to its $\vec{p}_T^{\text{up}}$ -independence, $m_{T2,\perp}$ is useful because it has an analytic formula for the event-by-event⁸ value. To illustrate how this works, consider a special but important case where $m_{C_i} = m_{V_i} = 0$. Since all perpendicular momenta lie along a line:

$$m_{T,\perp}^2(\vec{p}_{T,\perp}^{V_i}, \vec{p}_{T,\perp}^{C_i}) = \begin{cases} 4p_{T,\perp}^{V_i} p_{T,\perp}^{C_i} & \vec{p}_{T,\perp}^{V_i} \text{ is anti-parallel to } \vec{p}_{T,\perp}^{C_i} \\ 0 & \text{else} \end{cases}. \quad (3.10)$$

The calculation of m_{T2} is particularly simple because it is now a one-dimensional optimization problem. Without loss of generality, suppose that $p_{T,\perp}^{V_1} \geq p_{T,\perp}^{V_2}$. To ease the notation, let x be the signed projected test momentum in the minimization, $x = \vec{p}_{T,\perp}^{C_1} \cdot \hat{p}_{T,\perp}^{V_1}$ where $\hat{p}_{T,\perp}^{V_1} = \vec{p}_{T,\perp}^{V_1} / p_{T,\perp}^{V_1}$ is a unit vector pointing in the direction of the projected momentum for V_1 . Analogously, define p, q, y and ϵ as signed scalars representing the momenta of V_1, V_2, C_1 and $E_{T,\perp}$, i.e. $p = p_{T,\perp}^{V_1}$, $q = \vec{p}_{T,\perp}^{V_2} \cdot \hat{p}_{T,\perp}^{V_1}$, $y = \vec{p}_{T,\perp}^{C_2} \cdot \hat{p}_{T,\perp}^{V_1}$, and $\epsilon = \vec{E}_{T,\perp}^{\text{miss}} \cdot \hat{p}_{T,\perp}^{V_1}$. By momentum conservation, $p + q + \epsilon = 0$ and $x + y = \epsilon$. Therefore, $y = -p - q - x$. With this notation, Eq. 3.10 can

⁸In the general case, even though there is no analytic formula for the event-by-event quantity, there are general formulae for the kinematic maxima - see Ref. [565]. In general, this endpoint depends on p_T^{up} .

be re-written as $(\mathbf{m}_{T,\perp}^{V_1})^2 = 4|\mathbf{p}\mathbf{x}|(1 - \Theta(\mathbf{x}))$ and $(\mathbf{m}_{T,\perp}^{V_2})^2 = 4|\mathbf{q}\mathbf{y}|(1 - \Theta(\mathbf{y}))$ where $\Theta(\mathbf{x})$ is the Heavyside step function. The two branches of Eq. 3.10 are illustrated in the upper diagrams of Fig. 3.14. In the left diagram of Fig. 3.14, the two visible perpendicular momenta are on opposite sides of the upstream momentum so $\mathbf{p}\mathbf{q} < 0$. Since $|\mathbf{p}| > |\mathbf{q}|$, $\epsilon < 0$. When $\mathbf{x} > 0$, $\mathbf{m}_{T,\perp}^{V_1} = 0$ and for $\mathbf{x} < 0$, $(\mathbf{m}_{T,\perp}^{V_1})^2 = 4\mathbf{p}|\mathbf{x}|$. Similarly, $\mathbf{m}_{T,\perp}^{V_2} = 0$ when $\mathbf{y} > 0$, which occurs when $\mathbf{x} < -|\mathbf{p}| + |\mathbf{q}|$. For $\mathbf{x} > 0$, $\mathbf{m}_{T,\perp}^{V_2} = \mathbf{m}_{T,\perp}^{V_1} = \max_i\{\mathbf{m}_{T,\perp}^{V_i}\} = 0$, and therefore $\mathbf{m}_{T2,\perp} = 0$. This is illustrated graphically in the lower left graph in Fig. 3.14.

The second possibility is that $\mathbf{p}\mathbf{q} > 0$, as in the upper right diagram of Fig. 3.14. The curve for $(\mathbf{m}_{T,\perp}^{V_1})^2$ in the lower right graph of Fig. 3.14 is unchanged from the first case. However, now $\mathbf{y} < 0$ is required for $\mathbf{m}_{T,\perp}^{V_2}$ to be nonzero. This occurs when $\mathbf{x} > -|\mathbf{p} + \mathbf{q}|$. When $\mathbf{x}$ is large, $\max_i\{\mathbf{m}_{T,\perp}^{V_i}\} = \mathbf{m}_{T,\perp}^{V_2}$ and when $\mathbf{x}$ is much less than zero, $\max_i\{\mathbf{m}_{T,\perp}^{V_i}\} = \mathbf{m}_{T,\perp}^{V_1}$. As illustrated by lower right graph of Fig. 3.14, the global minimum occurs when $\mathbf{m}_{T,\perp}^{V_1} = \mathbf{m}_{T,\perp}^{V_2}$, which implies $\mathbf{p}\mathbf{x} = \mathbf{q}\mathbf{y} = \mathbf{q}(-\mathbf{p} - \mathbf{q} - \mathbf{x})$, or $\mathbf{x} = -\mathbf{q}$. Substituting $\mathbf{x} = -\mathbf{q}$ in the formula for $\mathbf{m}_{T,\perp}^{V_1}$ then gives $\mathbf{m}_{T2,\perp}^2 = 4\mathbf{p}\mathbf{q}$. A formula that covers both of the above cases is $\mathbf{m}_{T2,\perp}^2 = 2\mathbf{A}_{T,\perp}$ where $\mathbf{A}_{T,\perp} = (|\mathbf{p}\mathbf{q}| - \mathbf{p}\mathbf{q}) = (\mathbf{p}_{T,\perp}^{V_1}\mathbf{p}_{T,\perp}^{V_1} - \vec{\mathbf{p}}_{T,\perp}^{V_1} \cdot \vec{\mathbf{p}}_{T,\perp}^{V_2})$. A straightforward extension of the above argument to the case where the test masses are not zero, but are both equal to the same value $\mathbf{m}_C$ gives [564] $\mathbf{m}_{T2,\perp} = \sqrt{\frac{1}{2}\mathbf{A}_{T,\perp}} + \sqrt{\frac{1}{2}\mathbf{A}_{T,\perp} + \mathbf{m}_C^2}$. Additionally including nonzero visible particle masses has the form⁹ [566, 567]:

$$\mathbf{m}_{T2,\perp}^2 = \mathbf{m}_C^2 + \mathbf{A}_{T,\perp} + \sqrt{\left(1 + \frac{4\mathbf{m}_C^2}{2\mathbf{A}_{T,\perp} - \mathbf{m}_{V_1}^2 - \mathbf{m}_{V_2}^2}\right) (\mathbf{A}_{T,\perp}^2 - \mathbf{m}_{V_1}^2 \mathbf{m}_{V_2}^2)}, \quad (3.11)$$

where the generalized $\mathbf{A}_{T,\perp} = \frac{1}{2}(\mathbf{E}_{T,\perp}^{V_1}\mathbf{E}_{T,\perp}^{V_1} - \vec{\mathbf{p}}_{T,\perp}^{V_1} \cdot \vec{\mathbf{p}}_{T,\perp}^{V_2})$ for $(\mathbf{E}_{T,\perp}^{V_i})^2 = \mathbf{m}_{V_i}^2 + (\mathbf{p}_{T,\perp}^{V_i})^2$. The full formula, including the possibility for $\mathbf{m}_{C_1} \neq \mathbf{m}_{C_2}$ has also been computed [568]:

⁹Note that this formula also holds in the limit that the upstream momentum is zero. In that case, all $\perp$ quantities are replaced by the regular momenta.

$$m_{T2,\perp}^2 = \Sigma_C + A_{T,\perp} + \frac{\Delta_C \Delta_V}{A_{T,\perp} - \Sigma_V} \pm \sqrt{\left(1 + \frac{2\Sigma_C}{A_{T,\perp} - \Sigma_V} + \left(\frac{\Delta_C}{A_{T,\perp} - \Sigma_V}\right)^2\right) (A_{T,\perp}^2 - m_{V_1}^2 m_{V_2}^2)}, \quad (3.12)$$

where $\Sigma_C = \frac{1}{2}(m_{C_1}^2 + m_{C_2}^2)$, $\Delta_C = \frac{1}{2}(m_{C_2}^2 - m_{C_1}^2)$ and similarly for Σ_V, Δ_V and m_{V_i} .

In all cases, the event-by-event formula for $m_{T2,\perp}$ depends on the particle momenta only through the quantity $A_{T,\perp}$. When the visible particle masses are small compared to their momenta, this means that the distribution of $m_{T2,\perp}$ will have a δ -function corresponding to the case $A_{T,\perp} = 0$ that occurs whenever the visible particle momenta are on opposite sides of the upstream momentum. Furthermore, the additional projection compared with the original m_{T2} means that $m_{T2,\perp}$ tends to be lower for background and signal events. A full comparison of the discriminating power of $m_{T2,\perp}$ with m_{T2} is described in Sec. 3.2.1.2.5.




Figure 3.13: An illustration of the setup for the projected m_{T2} variable. The upstream momentum is the sum of all particles not associated with the visible particles V_i and child particles C_i . In this $t\bar{t}$ topology, the upstream momentum is due to the b -jets directly from the top quark decay and any FSR produced before the $t\bar{t}$ production.




Figure 3.14: Diagrams illustrating the two possible orientations of projected transverse momenta (top) and the corresponding $\mathbf{m}_{T,\perp}$ graphs for both branches of the decay (bottom). The circles represent a transverse cross-section of the detector - the beam is into and out of the page. In the left configuration, the visible transverse momenta are on opposite sides of the upstream momentum and the resulting $\mathbf{m}_{T2,\perp}$ value is equal to zero. In the right diagrams, the visible momenta are on the same side of the upstream momentum which allows for a finite $\mathbf{m}_{T2,\perp}$ value.

3.2.1.2.2 Numerical Methods

Aside from special cases such as the perpendicular $\mathbf{m}_{T2}$ in the previous section, there is no general analytic formula for the event-by-event value of $\mathbf{m}_{T2}$. The level sets of the $\mathbf{m}_T^2$ curves in Eq. 3.7 are conic sections in $x = \vec{p}_{x,\text{test}}^{C_i} \cdot \hat{x}$ and $y = \vec{p}_{y,\text{test}}^{C_i} \cdot \hat{y}$:

$$(\mathbf{m}_{V_i}^2 + (\mathbf{p}_y^{V_i})^2)x^2 + (\mathbf{m}_{V_i}^2 + (\mathbf{p}_x^{V_i})^2)y^2 - \mathbf{m}^2 \mathbf{p}_x^{V_i} x - \mathbf{m}^2 \mathbf{p}_y^{V_i} y - 2\mathbf{p}_x^{V_i} \mathbf{p}_y^{V_i} xy - M = 0, \quad (3.13)$$

where $\mathbf{m}^2 = \mathbf{m}_T^2 - \mathbf{m}_{C_i}^2 + \mathbf{m}_{V_i}^2$ and $M = \frac{1}{4}\mathbf{m}^4 - (\mathbf{E}_T^{V_i})^2(\mathbf{m}_{C_i}^2)$. The coefficients A, B, C of the x^2 , xy and y^2 terms in Eq. 3.13 satisfy $B - 4AC \leq 0$:

$$4(\mathbf{p}_x^{V_i})^2(\mathbf{p}_y^{V_i})^2 - 4(\mathbf{m}_{V_i}^2 + (\mathbf{p}_y^{V_i})^2)(\mathbf{m}_{V_i}^2 + (\mathbf{p}_x^{V_i})^2) = -4\mathbf{m}_{V_i}^2(\mathbf{m}_{V_i}^2 + (\mathbf{p}_T^{V_i})^2). \quad (3.14)$$

Therefore, the level sets are ellipses as long as $\mathbf{m}_{V_i} > 0$ and parabolas otherwise. Writing Eq. 3.13 as $f_i(x, y) = 0$, it is possible to plot both curves with one set of coordinates using the conservation of momentum constraint, $f_1(x, y) = 0$ and $f_2(\mathbf{E}_x^{\text{miss}} - x, \mathbf{E}_y^{\text{miss}} - y) = 0$. A vertical slice in $(x, \mathbf{m}_T)$ space that goes through the value of $\mathbf{m}_{T2}$ looks like either the left or right graph in Fig. 3.15. In the left graph of Fig. 3.15, the minimum over $\max\{\mathbf{m}_T\}$ occurs at the intersection of the two $\mathbf{m}_T$ curves (balanced) while in the right graph, the minimum of one $\mathbf{m}_T$ curve is above the other curve and is thus equal to $\mathbf{m}_{T2}$ (unbalanced). Without loss of generality, assume that $\mathbf{m}_{V_1} + \mathbf{m}_{C_1} > \mathbf{m}_{V_2} + \mathbf{m}_{C_2}$. From Eq. 3.9, the minimum of $\mathbf{m}_T^{V_1}$ occurs when $(x, y) = \frac{\mathbf{m}_{C_1}}{\mathbf{m}_{V_1}} \vec{p}_T^{V_1}$. Therefore, the condition for the unbalanced case is $\mathbf{m}_{V_1} + \mathbf{m}_{C_1} > \mathbf{m}_T(\vec{p}_T^{V_2}, \vec{E}_T^{\text{miss}} - \frac{\mathbf{m}_{C_1}}{\mathbf{m}_{V_1}} \vec{p}_T^{V_1}, \mathbf{m}_{V_2}, \mathbf{m}_{C_2})$. This case can only occur if $\mathbf{m}_{V_1} + \mathbf{m}_{C_1} \neq \mathbf{m}_{V_2} + \mathbf{m}_{C_2}$. Figure 3.16 demonstrates the probability for the unbalanced case as a function of $\mathbf{m}_{C_1}$ and $\mathbf{m}_{C_2}$ in $t\bar{t}$ events. The fraction of unbalanced events is increased by choosing V_1 and V_2 with significantly different masses ($\mathbf{m}_{b_l}$ and $\mathbf{m}_b$). The larger the difference between the $\mathbf{m}_T$ minima, the higher the probability for the unbalanced case.




Figure 3.15: Graphs illustrating the two possible configurations leading to the value of m_{T2} at the intersection of the two m_T conic sections (left) or at the minimum of one of the sections, if it is above the other curve (right).




Figure 3.16: The fraction of events for which the m_{T2} value is $m_{V1} + m_{C1}$ (unbalanced case) for $t\bar{t}$ events as a function of m_{C1} and m_{C2} . The visible object V_1 is the four-vector sum of the b -quark and lepton from the same top quark decay and V_2 is the b -quark from the other top quark decay. A white star indicates the choice for the $\mathbf{am}_{T2}$ variable described in Sec. 3.2.1.2.3.

In the balanced case, the value of $\mathbf{m}_{T2}$ is determined by computing the minimum value along the intersection of the $\mathbf{m}_T$ surfaces, as illustrated by Fig. 3.17. Equivalently, the value of $\mathbf{m}_{T2}$ is equal to the point at which the ellipses from the $\mathbf{m}_T^2$ level curves are tangent. This second condition, illustrated in Fig. 3.18 is used for quickly and accurately computing $\mathbf{m}_{T2}$. Rewriting $f_1(\mathbf{x}, \mathbf{y}) = 0$ and $f_2(\mathbf{E}_x^{\text{miss}} - \mathbf{x}, \mathbf{E}_y^{\text{miss}} - \mathbf{y}) = 0$ with $\mathbf{m}_T^{V_1} = \mathbf{m}_T^{V_2} = \mathbf{m}_T$:

$$\begin{aligned} \mathbf{a}_{y2}\mathbf{y}^2 + (\mathbf{a}_{xy}\mathbf{x} + \mathbf{a}_y(\mathbf{m}_T^2))\mathbf{y} + (\mathbf{a}_{x2}\mathbf{x}^2 + \mathbf{a}_x(\mathbf{m}_T^2)\mathbf{x} + \mathbf{a}_0(\mathbf{m}_T^4)) &= 0 \\ \mathbf{b}_{y2}\mathbf{y}^2 + (\mathbf{b}_{xy}\mathbf{x} + \mathbf{b}_y(\mathbf{m}_T^2))\mathbf{y} + (\mathbf{b}_{x2}\mathbf{x}^2 + \mathbf{b}_x(\mathbf{m}_T^2)\mathbf{x} + \mathbf{b}_0(\mathbf{m}_T^4)) &= 0, \end{aligned} \quad (3.15)$$

where the coefficients $\mathbf{a}_i$ and $\mathbf{b}_i$ are given by Eq. 3.13. Solving for $\mathbf{y}$ gives

$$\mathbf{y} = \tilde{\mathbf{a}}_1(\mathbf{x}, \mathbf{m}_T^2) \pm \sqrt{\tilde{\mathbf{a}}_2(\mathbf{x}^2, \mathbf{m}_T^4) + \tilde{\mathbf{a}}_3(\mathbf{x}^2, \mathbf{m}_T^4)}, \quad (3.16)$$

where the coefficients $\tilde{\mathbf{a}}$ are the usual solution to the quadratic equation from Eq. 3.15. A similar expression holds for the second branch but with $\mathbf{a} \leftrightarrow \mathbf{b}$. Where the two ellipses intersect, the values of $\mathbf{y}$ will be the same:

$$\tilde{\mathbf{a}}_1(\mathbf{x}, \mathbf{m}_T^2) \pm \sqrt{\tilde{\mathbf{a}}_2(\mathbf{x}^2, \mathbf{m}_T^4) + \tilde{\mathbf{a}}_3(\mathbf{x}^2, \mathbf{m}_T^4)} = \tilde{\mathbf{b}}_1(\mathbf{x}, \mathbf{m}_T^2) \pm \sqrt{\tilde{\mathbf{b}}_2(\mathbf{x}^2, \mathbf{m}_T^4) + \tilde{\mathbf{b}}_3(\mathbf{x}^2, \mathbf{m}_T^4)}. \quad (3.17)$$

Rearranging to remove the radicals:

$$\frac{1}{4}(\mathbf{c}_1(\mathbf{x}^2, \mathbf{m}_T^4) - \mathbf{c}_2(\mathbf{x}^2, \mathbf{m}_T^4))^2 = \tilde{\mathbf{a}}_2(\mathbf{x}^2, \mathbf{m}_T^4) + \tilde{\mathbf{a}}_3(\mathbf{x}^2, \mathbf{m}_T^4)\tilde{\mathbf{b}}_2(\mathbf{x}^2, \mathbf{m}_T^4) + \tilde{\mathbf{b}}_3(\mathbf{x}^2, \mathbf{m}_T^4), \quad (3.18)$$

where $\mathbf{c}_1(\mathbf{x}^2, \mathbf{m}_T^4) = (\tilde{\mathbf{a}}_1(\mathbf{x}, \mathbf{m}_T^2) - \tilde{\mathbf{b}}_1(\mathbf{x}, \mathbf{m}_T^2))^2$ and $\mathbf{c}_2(\mathbf{x}^2, \mathbf{m}_T^4) = \tilde{\mathbf{a}}_2(\mathbf{x}^2, \mathbf{m}_T^4) + \tilde{\mathbf{a}}_3(\mathbf{x}^2, \mathbf{m}_T^4) + \tilde{\mathbf{b}}_2(\mathbf{x}^2, \mathbf{m}_T^4) + \tilde{\mathbf{b}}_3(\mathbf{x}^2, \mathbf{m}_T^4)$. Equation 3.18 is a quartic equation of $\mathbf{x}$ for a fixed $\mathbf{m}_T$. From

Fig. 3.18, when $m_T < m_{T2}$, there are no intersection points and so Eq. 3.18 will have no real roots. In contrast, when $m_T > m_{T2}$, Fig. 3.18 shows that there are two points of intersection for the ellipses and so Eq. 3.18 will have two real roots. At exactly the value $m_T = m_{T2}$, Eq. 3.18 will have one repeated root. A polynomial has a repeated root if and only if its determinant is zero. The discriminant of a quartic polynomial is a sixth degree polynomial in the coefficients of the polynomial. However, the coefficients of Eq. 3.18 are quadratic functions of m_T^2 and therefore solving for m_{T2} using the discriminant requires finding the roots of a 12th degree polynomial. There is no general analytic formula for such a high degree polynomial so the roots must be computed numerically. Even though this is a well-defined procedure, it is not the usual way m_{T2} is computed because it would be relatively slow and possibly numerically unstable.




Figure 3.17: Three views of the m_T^2 surfaces from the same event, rotated so illustrate the structure of the intersection. The value of m_{T2}^2 is indicated in the middle graph by a white arrow.

The state-of-the-art numerical calculators for m_{T2} are based on the observation that it is quick and easy to check if two ellipses intersect. Then, m_{T2} is computed by iteratively bisecting an interval known to contain the point at which the ellipses are tangent. Bisection techniques are powerful because they achieve a precision of n decimal places with only $\log_2(10^n)$ bisections. The first bisection method [569] used the above observation that the number of real roots of Eq. 3.18 differs if m_T is above or below m_{T2} . A fast way to check the number of real roots of a polynomial




Figure 3.18: Level surfaces of the graphs from Fig. 3.17. As noted above, these curves are ellipses and the point at which they are tangent (indicated by a black star) corresponds to the value of $\mathbf{m}_{T2}$.

is to use the *Sturm sequence*, which is based on a few evaluations of a simple series of five polynomials (the original one, its derivative, and various divisors). An even faster method is based on the observation by C. Lester that it is easier and more robust to check if the area of two ellipses overlap than to check if their boundaries intersect [570]. The numerical procedure for the quick evaluation of overlapping conic sections is from Ref. [571]. In addition to the speed of evaluation, this new procedure is more robust compared with the Sturm sequence method because it removes the need for special cases when e.g. one of the visible particles is massless and the $\mathbf{m}_T$ level sets are parabolas instead of ellipses. Furthermore, before the availability of the calculator from Ref. [570], there was no specialty $\mathbf{m}_{T2}$ calculator for the case $\mathbf{m}_{C_1} \neq \mathbf{m}_{C_2}$ (Ref. [569] is only programed for the symmetric case). Therefore, at $\sqrt{s} = 8$ TeV $\mathbf{m}_{T2}$ was evaluated using a generic function minimizer (Migrad - see Sec. 4.2) initialized with $\mathbf{x} = \vec{\mathbf{p}}_{\mathbf{x}}^{\text{miss}} \cdot \hat{\mathbf{x}}/2$ and $\mathbf{y} = \vec{\mathbf{p}}_{\mathbf{y}}^{\text{miss}} \cdot \hat{\mathbf{y}}/2$. Each $\mathbf{m}_T$ surface is smooth, so the numerical minimization is robust except near the intersection of the surfaces where some instability is caused by the discontinuity in the first derivative.

The $\sqrt{s} = 13$ TeV analysis uses the dedicated calculator from Ref. [570] that is orders of magnitude faster than the generic approach for the same precision. Figure 3.19 shows that the bisection approach is just as good as the analytic formula in the case $p_T^{\text{up}} = 0$.




Figure 3.19: A demonstration of the numerical accuracy of the m_{T2} calculator based on the overlap of conic sections instead of the intersection of their boundaries. Plotted is the difference d between the numerical calculation and the analytic formula using the projected m_T from Sec. 3.2.1.2.1 in $t\bar{t}$ events for the variable am_{T2} described in Sec. 3.2.1.2.3. The scale of $m_T \sim \mathcal{O}(100 \text{ GeV})$. See Ref. [570] for more detail.

3.2.1.2.3 Tailoring m_{T2} for the stop search

With a high multiplicity final state, there are many choices for V_i and m_{C_i} in constructing an m_{T2} variable for the stop search. This section describes two particular m_{T2} variables that are designed to suppress dilepton $t\bar{t}$ events. After requiring $m_T > m_W$, Fig. 3.8 showed that the majority of surviving $t\bar{t}$ events have a second lepton that is either lost, mis-identified, or is a hadronically decaying τ lepton. First, consider the case in which the second lepton is undetected so that $\vec{p}_T^{\text{miss}} \approx \vec{p}_T^{\nu_1} + \vec{p}_T^{\nu_2} + \vec{p}_T^{\text{lost } \ell}$. One could construct an m_{T2} variable using the b -jets as the V_i and then grouping everything downstream of the b -quarks in the top decay chain into $\vec{p}_T^{C1} + \vec{p}_T^{C2}$. However, additional information is available by using asymmetric objects [568, 572] for the V_i . Following an idea in Ref. [573], the *asymmetric* m_{T2} ($\mathbf{a}m_{T2}$) is formed by letting V_1 be the b -jet from one top quark decay and setting V_2 to be the four-vector sum of the b -jet and lepton from the other top quark decay. As illustrated in Fig. 3.20, this means that the missing particle for the top branch is an entire W boson and on the bottom branch, only a neutrino. Therefore, $m_{C1} = m_W$ and $m_{C2} = m_\nu \approx 0$. With these choices, $\mathbf{a}m_{T2} \leq m_{\text{top}}$ for the background depicted in Fig. 3.20.




Figure 3.20: A schematic diagram of dileptonic $t\bar{t}$ decay where one of the charged leptons is lost. Lost particles are circled with a dashed line. For the $\mathbf{a}m_{T2}$ variable, the visible particle on the top (bottom) branch is the b -jet (sum of the b -jet and lepton). The missing particle in the top (bottom) branch is a W boson (neutrino).

An important practical complication for constructing $\mathbf{am}_{T2}$ is the selection of the $\mathbf{b}$ -jets and the pairing of the lepton with the $\mathbf{b}$ -jet from the same branch. The signal region event selections described in Sec. 4 only require one explicitly identified $\mathbf{b}$ -tagged jet. Section 3.2.1.2.5 will explore two methods for choosing the two $\mathbf{b}$ -tagged jet: the two jets with the highest $\mathbf{b}$ -tag discriminant weights or the two highest $\mathbf{p}_T$ jets. As the $\mathbf{b}$ -quarks appear higher in the top quark decay chain than the other tree-level quarks, it is expected that they have a higher $\mathbf{p}_T$ on average. Section 3.2.1.2.5 also considers two algorithms for matching $\mathbf{b}$ -jets with the lepton: take the closest in ΔR or compute both possibilities and set $\mathbf{am}_{T2}$ to be the minimum of the two $\mathbf{m}_{T2}$ values. Figure 3.21 shows the scalar parton-level distribution of $\mathbf{am}_{T2}$ for dileptonic $t\bar{t}$ and stop events. By construction, $\mathbf{am}_{T2} < \mathbf{m}_{\text{top}}$ when the second lepton is lost. In contrast, when the second lepton is measured but not identified as a lepton, there is a small tail of events with $\mathbf{am}_{T2} > \mathbf{m}_{\text{top}}$. This is also true when the second lepton is lost, but the $\mathbf{b}$ -jet / lepton pairing is performed with ΔR . For stop events, $\mathbf{am}_{T2}$ can greatly exceed $\mathbf{m}_{\text{top}}$. The peak around 150 GeV is due to the unbalanced case in which $\mathbf{am}_{T2} = \mathbf{m}_{\text{bl}}$.




Figure 3.21: The distribution of $\mathbf{am}_{T2}$ for dileptonic $t\bar{t}$ and stop events with various configurations as indicated in the legend. The $\mathbf{p}_T$ spectrum of the top quarks in $t\bar{t}$ events is chosen to be identical to the distribution for stop events.

The distribution of $\mathbf{am}_{T2}$ after a relatively loose event selection is shown in Fig. 3.22. All three $t\bar{t}$ components are significantly reduced for $\mathbf{am}_{T2} > m_{\text{top}}$. Interestingly, the other backgrounds have relatively uniform distributions over the plotted range; this observation is revisited in Sec. 5.3 to isolate single top events.




Figure 3.22: A comparison of data and simulation using a loose selection requiring exactly one signal lepton, four jets with $p_T > 25$ GeV, at least one b -tagged jet, and $m_T > 100$ GeV to enrich the dilepton contribution. The b -tag weight is used to select the two b -jet candidates and the minimum value over both b -lepton pairings is used to resolve the matching ambiguity. The uncertainty band includes jet energy scale and resolution uncertainties (see Sec. 6).

To target the case of dileptonic $t\bar{t}$, a second m_{T2} variable called m_{T2}^τ is constructed based on the topology illustrated in Fig. 3.23. The visible particle on one branch is an identified hadronic τ and on the other branch is the reconstructed electron or muon. For the lower branch, a single neutrino is the lost particle so $m_{C_2} = m_\nu$. If the full hadronic τ were stable, than there would be a single neutrino on the upper branch. However, the τ decays into a ν_τ in addition to hadrons so the mass of the missing object in the upper branch is not strictly m_ν . However, when $m_{C_1} = 0$, m_{T2} is still bounded by m_W and is therefore used in the construction of m_{T2}^τ . Figure 3.24 demonstrates that the m_W bound is still preserved with $m_{C_1} = 0$, but the kinematic maximum is not as saturated when the full τ energy is not measured.

The only combinatorial challenge for m_{T2}^τ is the selection of the τ candidate. Section 3.2.1.2.5 considers two possibilities: using the third highest p_T jet or the highest p_T jet that is not one of the two jets with the highest b -tagging weight. The possibility of using an explicit τ candidate is investigated in Sec. 3.2.3.




Figure 3.23: A dileptonic $t\bar{t}$ event with one hadronically decaying τ . Lost particles are circled with a dashed line. For the m_{T2}^τ variable, the visible particle on the top (bottom) branch is the hadronic τ (reconstructed lepton). The missing particle in the top (bottom) branch is the sum of τ neutrinos (e or μ neutrino).




Figure 3.24: The distribution of m_{T2}^τ in dilepton $t\bar{t}$ events with top quark p_T spectra that match that of a 800 GeV stop and $m_{LSP} = 0$. For the filled histogram, the τ is decayed using a scalar three-body phase space to mimic $\tau \rightarrow \nu_\tau \pi^0 \pi^\pm$. The ‘pions’ are added to form the visible τ and the ν_τ is added to the $\vec{p}_T^{\text{miss}}$.

3.2.1.2.4 Comparisons between transverse mass variables

Figure 3.25 presents an overview of transverse mass variables constructed to identify and suppress dileptonic $t\bar{t}$ events. The dashed lines indicate which subsystem of the $t\bar{t}$ decay is targeted by the variables with the same color code and the particles composing the V_i are circled. Section 3.2.1.2.3 introduced the $\mathbf{am}_{T2}$ and $\mathbf{m}_{T2}^\tau$ variables. The variable $\mathbf{m}_{bl}$ is the $\mathbf{b}$ -jet / lepton invariant mass and is identical to $\mathbf{am}_{T2}$ in the unbalanced case. One new $\mathbf{m}_{T2}$ variable, simply denoted $\mathbf{m}_{T2}$ in Fig. 3.25, uses two $\mathbf{b}$ -tagged jets as the visible particles, adds the reconstructed lepton to the $\vec{p}_T^{\text{miss}}$ and then $\mathbf{m}_{C_1} = \mathbf{m}_{C_2} = \mathbf{m}_W$. The *contransverse mass* [574], denoted $\mathbf{m}_{CT}$, is similar in spirit to $\mathbf{m}_{T2}$ and is defined by

$$\mathbf{m}_{CT}^2 = \mathbf{m}_{V_1}^2 + \mathbf{m}_{V_2}^2 + 2 \left(\sqrt{\mathbf{m}_{V_1}^2 + (\mathbf{p}_T^{V_1})^2} \sqrt{\mathbf{m}_{V_2}^2 + (\mathbf{p}_T^{V_2})^2} + \vec{p}_T^{V_1} \cdot \vec{p}_T^{V_2} \right). \quad (3.19)$$

Equation 3.19 is similar to the equation for $\mathbf{m}_T^2$, but with the missing particle replaced with the second visible particle and with the sign of the last term flipped (compare with Eq. 3.8). The contransverse mass is constructed so that its kinematic maximum is invariant under equal and opposite boosts of the particles V_1 and V_2 in analogy to the invariance of the transverse mass under coherent boosts of V_1 and V_2 in the same direction¹⁰. The endpoint of the $\mathbf{m}_{CT}$ distribution is a known combination of the parent, visible, and invisible particle masses, which makes it useful for discriminating signal events from background events. One advantage of $\mathbf{m}_{CT}$ over $\mathbf{m}_{T2}$ is that Eq. 3.19 is simple, without requiring any optimization. In the context of the stop search, $\mathbf{m}_{CT}$ is constructed with two $\mathbf{b}$ -jet candidates as the visible particles.

The above variables and their variations are quantitatively compared in terms of their ability to separate $t\bar{t}$ events from stop events. Each variable has several variations, including the $\mathbf{b}$ -jet and τ -jet identification algorithms for $(\mathbf{a})\mathbf{m}_{T2}^{(\tau)}$, the $\mathbf{b}$ -jet / lepton matching scheme for $\mathbf{am}_{T2}$ and $\mathbf{m}_{bl}$, and the projection perpendicular to

¹⁰As with the transverse mass, the contransverse mass is not invariant under these *contra-linear* boosts event-by-event, but the endpoint of the $\mathbf{m}_{CT}$ distribution is invariant.




Figure 3.25: A diagram of dileptonic the $t\bar{t}$ decay chain and kinematic variables targeting various components of the chain. The dotted lines highlight which aspects of the top quark decay are involved with the construction of the variables with the same color. Particle inputs to each variable are circled with the same colors.

$\vec{p}_T^{\text{up}}$ (or not). The $m_{\text{CT},\perp}$ [575] is defined analogously to the $m_{\text{T2},\perp}$ variables. Another method to reduce the $\vec{p}_T^{\text{up}}$ dependence of m_{CT} is to apply a *boost-correction* [576] based on the kinematic properties of the visible particles and the $\vec{p}_T^{\text{miss}}$. A metric to quantify the separation power is given by the overlap integral [577, 578]:

$$\langle S^2 \rangle = \frac{1}{2} \int \frac{(f_S(x) - f_B(x))^2}{f_S(x) + f_B(x)} dx, \quad (3.20)$$

where f_S and f_B are the probability distribution functions for a random variable X with signal and background processes, respectively¹¹. Figure 3.26 shows example distributions of m_T , E_T^{miss} , am_{T2} , and $am_{\text{T2},\perp}$ for the $t\bar{t}$ background and a stop signal with $(m_{\text{stop}}, m_{\text{LSP}}) = (500, 1)$. As expected, the m_T peak in $t\bar{t}$ events is around m_W with a long tail due to resolution effects and the dilepton contribution. The am_{T2} distribution has an endpoint near m_{top} for $t\bar{t}$ events while the $am_{\text{T},\perp}$ distribution is concentrated at low values for both the background and signal. As a result, the separation $\langle S^2 \rangle$ is significantly worse for $am_{\text{T2},\perp}$ compared with am_{T2} .

¹¹Note the similarity of the overlap integral with the χ^2 -divergence. In fact, the overlap integral is an f-divergence with $f(u) = (u-1)^2/(u+1)$ (for the χ^2 -divergence, $f(u) = (u-1)^2$). Despite this, the overlap integral has not been applied to e.g. signal processing outside of high energy physics.




Figure 3.26: The distributions of m_T (top left), E_T^{miss} (top right), am_{T2} (bottom left), and $am_{T2,\perp}$ (bottom right) for $t\bar{t}$ and stop events with $(m_{\text{stop}}, m_{\text{LSP}}) = (500, 1)$. The separation (sep) is defined in Eq. 3.20. The b -tagging weight (MV1) is used to select the b -jets for am_{T2} and a ΔR scheme matches the b -jet with the lepton.

Table 3.4 summarizes the separation power for all the variables mentioned above. In addition to the separation, the table also provides the correlation with the E_T^{miss} and m_T . These are two known powerful variables, so a low correlation is an important metric for deciding on the usefulness of a new technique. The unprojected am_{T2} and m_{T2}^{τ} variables have the highest separation power amongst the possible variables. Additionally, am_{T2} has only a modest correlation with m_T and E_T^{miss} . The variants have similar performance; for the sake of reducing the parameter space for later optimization, the (B, min) for am_{T2} and the B setting for m_{T2}^{τ} are chosen as default¹². Further comparisons between the transverse mass and other variables are described in Sec. 4.

¹²The ‘B’ method also has a significantly higher accuracy: the jet with the highest b -tag weight is nearly twice as likely as the highest p_T jet to be matched to a particle-level b -jet.

Variable	Variant	Separation	Corr. E_T^{miss}	Corr. m_T
E_T^{miss}	—	0.55	100%	49%
m_T	—	0.5	49%	100%
m_{CT}	P	0.19	53%	1%
	P, $\perp$	0.13	40%	1%
	P, BC	0.23	60%	3%
	B	0.06	29%	6%
	B, $\perp$	0.06	23%	4%
	B, BC	0.13	39%	9%
	P	0.24	63%	2%
m_{T2}	B	0.16	47%	8%
	P, $\perp$	0.13	40%	1%
	B, $\perp$	0.06	23%	4%
	P, min	0.33	65%	24%
	P, ΔR	0.31	61%	24%
	B, min	0.28	53%	30%
	B, ΔR	0.27	51%	29%
am_{T2}	P, $\perp$, ΔR	0.15	40%	14%
	B, $\perp$, ΔR	0.10	25%	18%
	B, $\perp$, min	0.08	26%	14%
	P	0.36	54%	66%
	B	0.40	63%	70%
	P, $\perp$	0.15	19%	34%
	B, $\perp$	0.14	24%	31%
m_{T2}^τ	B, min	0.02	0%	15%
	P, min	0.04	2%	4%
	B, ΔR	0.02	3%	11%
	P, ΔR	0.03	4%	7%
m_{bl}				

Table 3.4: The separation power and correlation with E_T^{miss} and m_T for a variety of variables described in the text. The background is $t\bar{t}$ and the signal is a stop model with $(m_{\text{stop}}, m_{\text{LSP}}) = (500, 1)$. The variant P means that p_T is used to pick the b -jets while B means that the b -tagging weight is used. The symbol $\perp$ denotes the perpendicular variant of the variable in the first column. A boost correction is indicated by the letters BC. The symbols ΔR and min represent the scheme for addressing the matching ambiguity between the lepton and the b -jet by using the closest pair or considering the minimum of both possible pairings.

3.2.1.2.5 Additional Considerations

This section briefly describes a few aspects of m_{T2} that are slightly out of, or beyond the scope of the rest of the chapter. In particular,

- The variable am_{T2} is also useful for stop decays other than $\tilde{t} \rightarrow t\tilde{\chi}^0$. For example, when $m_{\text{stop}} < m_{\text{top}}$, the stop can undergo a three-body $\tilde{t} \rightarrow bW\tilde{\chi}^0$ decay that has systematically *lower* am_{T2} values than the dileptonic $t\bar{t}$ background for which $am_{T2} \sim m_{\text{top}}$. This is illustrated in the left plot of Fig. 3.27. The unbalanced case is set by $\max\{m_{bl} + m_\nu, m_b + m_W\}$, which is $m_{bl} \lesssim \sqrt{m_{\text{top}}^2 - m_W^2}$ for $t\bar{t}$. In contrast, $m_{bl} < m_W$ for the signal and so am_{T2} tends to be close to the m_W lower bound. Therefore, an *upper threshold* on am_{T2} is a powerful discriminant for targeting three-body stop decays.

Another possibility is the flavor-changing decay $\tilde{t} \rightarrow b\tilde{\chi}^\pm$ (bC). Without a resonant top quark, the bl invariant mass for the bC decay is significantly higher than the $t\bar{t}$ case resulting in a larger separation between signal and background relative to the flavor-neutral decay. This is illustrated in the right plot of Fig. 3.27. The am_{T2} variable has successfully improved the sensitivity to both three-body and bC stop decays in the 3body and bCx signal regions of Ref. [414].

- One can always improve or create m_{T2} variables by incorporating more information. Section 3.2.2 will describe how to include resolution information and Sec. 3.2.3 will combine m_{T2} with an explicit hadronically decaying τ reconstruction. In a similar spirit, events with a second low quality lepton that does not pass the baseline criteria can be coupled with kinematic information via m_{T2} to create a high-efficiency veto for dilepton events. Often, low quality leptons are not well modeled, but large uncertainties are suppressed when the veto efficiency is close to unity. Figure 3.28 gives a concrete example where an m_{T2} variable is formed from the selected lepton and the next highest p_T muon. All events in Fig. 3.28 pass the second lepton veto, so the muon used in the m_{T2} calculation does not pass the requirements to be baseline. However, if the muon were truly




Figure 3.27: The distribution of am_{T2} for dileptonic $t\bar{t}$ where one lepton is lost and a three-body stop decay (left) with $(m_{\text{stop}}, m_{\text{LSP}}) = (200, 100)$ GeV and for a flavor-changing $b + \tilde{\chi}^\pm$ decay and flavor-neutral two-body decay $t + \tilde{\chi}^0$ (right) with $(m_{\text{stop}}, m_{\text{chargino}}, m_{\text{LSP}}) = (800, 300, 150)$ GeV.

from a W decay, $m_{T2}^{\text{lost}} \leq m_W$. Even in cases where there is not a lost muon, the scale of m_{T2}^{lost} is significantly less than that of the signal. Unfortunately, there are a significant fraction of signal events with $m_{T2}^{\text{lost}} \approx 0$, but it still may be useful to veto events in the first bin of Fig. 3.28. It may also be possible to improve the performance by adding slightly more quality criteria to the muon definition (but still below the baseline requirements).

- While all of the applications so far have been focused on vetoing the background, m_{T2} variables can also be used to directly tag the signal. The decay $\tilde{t} \rightarrow t\tilde{\chi}^0$ has the form of Fig. 3.11 with top quarks for V_i and neutralinos for C_i . Therefore, one could construct an $m_{T2}^{\tilde{t}}$ with $m_{C_i} = m_{\tilde{\chi}^0}$. The endpoint of such a variable would be $m_{T2}^{\tilde{t}} \leq m_{\tilde{t}}$ in stop events. If signal events saturate this bound and background events are relegated to lower values, then this variable could be useful. The left plot of Fig. 3.29 gives a concrete example of an $m_{T2}^{\tilde{t}}$ variable where one visible particle is the large-radius jet hadronic top quark candidate (see Sec. 3.2.4) and the b -lepton pair is the other visible particle.




Figure 3.28: The distribution of m_{T2} in diboson ($ZZ/WW \rightarrow l^+l^-\nu\bar{\nu}$) and stop events using the selected lepton as one visible particle and the next highest p_T muon for the second visible particle. For the stop model, $(m_{\text{stop}}, m_{\text{LSP}}) = (800, 1)$ GeV.

By construction, $m_{T2}^{\tilde{t}} \leq m_{\tilde{t}}$ and the average $m_{T2}^{\tilde{t}}$ is significantly larger for the signal than for the irreducible $t\bar{t} + Z(\rightarrow \nu\bar{\nu})$ background. The right plot of Fig. 3.29 shows that $\mathbf{a}m_{T2}$ has a similar separation, suggesting that the two variables may be related. The correlation between $\mathbf{a}m_{T2}$ and $m_{T2}^{\tilde{t}}$ is shown in Fig. 3.30 for both the background and the signal. There is a strong correlation in the signal, but little correlation in the background. One exception is at low m_{T2} values where the unbalanced case can result in both variables giving the same value. The $m_{T2}^{\tilde{t}}$ is promising tool for selecting stop events and suppressing events with a similar event topology; it will be interesting to expand upon these studies in the future.

Additional properties and uses of m_{T2} are discussed in subsequent chapters.




Figure 3.29: The distribution of $m_{T2}^{\tilde{t}}$ (left) and am_{T2} (right) for $t\bar{t} + Z(\rightarrow \nu\bar{\nu})$ and $t\bar{t}$ events with $(m_{\text{stop}}, m_{\text{LSP}}) = (800, 1)$ GeV. The left plot also includes the distributions using particle-level top quarks for the visible particles.




Figure 3.30: The joint distribution of $m_{T2}^{\tilde{t}}$ and am_{T2} for stop events (left) and $t\bar{t} + Z(\rightarrow \nu\bar{\nu})$ events (right).

3.2.2 Significance Variables

Event-by-event and object-by-object resolutions can be estimated from simulation and auxiliary measurements. These resolutions are often the source of background events with apparent signal-like kinematic properties. For example, Fig. 3.31 shows a schematic diagram illustrating a dijet event where one of the two jets has a significant mis-measurement in the direction transverse to the jet axis. As a result of jet angular resolution, an event with no real missing momentum from neutrinos or other weakly interacting particles can have a large apparent missing momentum. Jet-by-jet kinematic covariance matrices could be used to identify such events and rule out the mis-measurement as insignificant. The uncertainty on the jets and other objects can be used for all kinematic variables, in addition to the E_T^{miss} . However, this information is mostly unused in the construction of discriminating variables at the LHC. This section describes simple procedures for optimally combining kinematic variables with estimates of their resolution to form *significance variables*.




Figure 3.31: The green and red annuli represent the ATLAS calorimeters and the yellow arrows represent the measured jet directions. The direction of the right jet is mis-measured leading to an apparent E_T^{miss} represented by the dashed arrow. Such events could be identified if the jet kinematic covariance matrices show that the jets are statistically consistent with being back-to-back.

To concretely illustrate the idea¹³, consider a kinematic variable $\mathbf{m}$ for a particular process which has a kinematic maximum M in the absence of detector resolution. For example, $\mathbf{m}$ could be transverse momentum or the actual mass of some system of particles. The usual procedure for using $\mathbf{m}$ is to place a threshold $\mathbf{m}_{\text{threshold}}$ and then to count the number of events for which $\mathbf{m} > \mathbf{m}_{\text{threshold}}$. If this number significantly exceeds expectation, then one has evidence for new physics. However, one can do better than this by including more information such as event-by-event resolutions and the mass scale M . To construct an example, consider three variants of the variable $\mathbf{m}$:

- $\mathbf{m}^{\text{truth}}$: the value of $\mathbf{m}$ for a single realization of a given experiment in the absence of detector resolution.
- $\mathbf{m}^{\text{measured}}$: the measured value of $\mathbf{m}$ for a single realization of a given experiment.
- $\mathbf{m}^{\text{re-measured}}$: a random variable with probability distribution function given by the posterior distribution for the outcome of an experiment. This variable only makes sense in the context of conditioning on a measured value from a single experiment $\mathbf{m}^{\text{measured}}$. In special cases, $\mathbf{m}^{\text{measured}}$ may be the mean or median of the distribution of $\mathbf{m}^{\text{re-measured}}$, but this is not true in general due to asymmetric detector responses and a falling truth-level spectrum.

One quantity that captures resolution and scale information is the probability P_M that the measured value $\mathbf{m}^{\text{measured}}$ for a fixed event would exceed the scale M if re-measured¹⁴. Symbolically, this is

$$P_M = \Pr(\mathbf{m}^{\text{re-measured}} > M | \mathbf{m}^{\text{measured}}), \quad (3.21)$$

The probability P_M will depend on the probability distribution $p(\mathbf{m}^{(\text{re-})\text{measured}} | \mathbf{m}^{\text{true}})$, known as the resolution function. For many applications, the resolution function is

¹³The ideas presented in this section are published in Ref. [579] and include input from C. Lester.

¹⁴A related quantity, is $Q_M = \Pr(\mathbf{m}^{\text{truth}} > M | \mathbf{m}^{\text{measured}})$. One can show that P_M and Q_M do not induce the same ordering on events and therefore one may be better than the other for a particular application.

well-approximated by a Gaussian centered at the true value with standard deviation σ_m . If the true distribution is only slowly varying in a neighborhood $m^{\text{measured}} \pm \sigma_m$ such that $p(m^{\text{true}})$ is approximately constant, the value of P_M can be computed analytically as shown in Eq. 3.22.

$$\begin{aligned}
P_M &= \int_M^\infty p(m^{\text{re-measured}} | m^{\text{measured}}) dm^{\text{re-measured}} \\
&= \int_M^\infty \int_{-\infty}^\infty p(m^{\text{re-measured}} | m^{\text{measured}}, m^{\text{true}}) p(m^{\text{true}} | m^{\text{measured}}) dm^{\text{re-measured}} dm^{\text{true}} \\
&\propto \int_M^\infty \int_{-\infty}^\infty p(m^{\text{re-measured}} | m^{\text{true}}) p(m^{\text{measured}} | m^{\text{true}}) dm^{\text{re-measured}} dm^{\text{true}} \\
&\propto \int_M^\infty \exp\left(-\frac{(m^{\text{re-measured}} - m^{\text{observed}})^2}{4\sigma_m^2}\right) dm^{\text{re-measured}} \\
&= \frac{1}{2} \left(1 + \operatorname{erf}\left(\frac{m^{\text{observed}} - M}{2\sigma_m}\right)\right), \tag{3.22}
\end{aligned}$$

The second line in Eq. 3.22 is the law of total probability, the third line is from the fact that $m^{\text{re-measured}}$ is independent of m^{measured} given m^{true} , Bayes theorem, and the approximation that $p(m^{\text{true}})$ is approximately constant near m^{measured} . The fourth line in Eq. 3.22 is from completing the square and integrating out m^{true} . Since the erf function is monotonic and smooth, the complete behavior of P_M is determined by the quantity:

$$X_M \equiv \frac{m^{\text{observed}} - M}{\sigma_m}. \tag{3.23}$$

The only current use of a variable like X_M is the “ E_T^{miss} significance”. First constructed at DØ [580], the E_T^{miss} significance in its most complete form usually refers to the log of a likelihood ratio:

$$\log\left(\frac{p(E_T^{\text{miss}} = E_T^{\text{miss,measured}})}{p(E_T^{\text{miss}} = 0)}\right). \tag{3.24}$$

The purpose of E_T^{miss} significance is to differentiate events with real missing energy from invisible particles like neutrinos from those without (see Fig. 3.31), and it is constructed from the resolution functions of all the objects used to construct the E_T^{miss} itself. For Gaussian resolutions, the E_T^{miss} significance is a monotonic function of $(E_T^{\text{miss}})^2 / 2\sigma_{E_T^{\text{miss}}}^2$. The resolutions are well approximated by $\sigma_{E_T^{\text{miss}}} \propto \sqrt{H_T}$, the scalar sum of the visible $\mathbf{p}_T$ in the event [137, 581]. Therefore, an approximate E_T^{miss} significance may be written as a monotonic function of $(E_T^{\text{miss}})^2 / H_T$. The most widely used choice is $E_T^{\text{miss}} / \sqrt{H_T}$. Note that the approximate E_T^{miss} significance is a realization of X_M in which $M = 0$, the resolution function is Gaussian, and $\sigma \propto \sqrt{H_T}$.

Even though $E_T^{\text{miss}} / \sqrt{H_T}$ and E_T^{miss} are correlated, one can gain statistical power by considering $E_T^{\text{miss}} / \sqrt{H_T}$ in addition to or instead of E_T^{miss} itself. This has been shown in numerous analyses spanning a wide range of physics processes including Standard Model measurements and searches for SUSY. In addition to studying the general properties of significance variables, the next sections explore the potential gains from building significance variables for other kinematic variables.

3.2.2.1 Constructing Significance Variables

The optimal method for using event-by-event and object-by-object resolutions is to combine them in multidimensional likelihood with the kinematic variables themselves. By the Neyman-Pearson lemma [582], a threshold requirement on the likelihood is no worse than any other possibility, and thus optimal. However, it is often not possible or highly non-trivial to compute the likelihood combination. The E_T^{miss} significance example motivated the formation of a particular combination of the kinematic variable and its associated resolution into a single quantity. This quantity is equivalent to the significance variable X_M , which may contain all of the relevant discriminatory information. Ideally, it is a general trend that most of the relevant resolution information can be condensed into a single simple X_M -like variable. Fortunately, this will be true under certain conditions – principally those in which the signal and backgrounds are associated with different mass or energy scales. Before showing specific examples, it is important to note that while the X_M significance variables may capture most of the relevant resolution information, they may not always be optimal for every kinematic

variable. Any case in which resolutions are significantly non-Gaussian may require, for optimality, the use of a significance variable based on the full likelihood ratio. Nonetheless, X_M is simple to compute and contains information that is currently unused by most analyses.

3.2.2.2 Examples of Optimal Significance Variables

To begin, consider a simple model in which the variable $\mathbf{m}$ has a delta function distribution, $(1/N)d\mathbf{m}_i/dN = \delta(\mathbf{m} - M_i)$, where $i \in \{s, b\}$ (signal/background). For example, suppose that $\mathbf{m} = \mathbf{m}_T$ in a search with a resonance decaying into a lepton and a neutrino. Due to the Jacobian peak, most of the probability for $\mathbf{m}$ is near M_i , and so this simple model captures some important aspects of the analysis. Let the resolution functions of $\mathbf{m}$ be Gaussian with width σ . Then, the joint probability distribution of $\mathbf{m}$ and σ is given by

$$p_i(\mathbf{m}, \sigma) = g(\sigma) \frac{1}{\sqrt{2\pi}\sigma} \exp\left(-\frac{(\mathbf{m} - M_i)^2}{2\sigma^2}\right), \quad (3.25)$$

where $g(\sigma)$ is the distribution of σ . The optimal use of $\mathbf{m}$ and σ is to place a threshold on the ratio $p_s(\mathbf{m}, \sigma)/p_b(\mathbf{m}, \sigma)$. Dividing the probability functions from Eq. 3.25 and monotonically transforming them results in the following optimal significance variable

$$V_{\text{opt}}^{(\text{Gaussian})} = \frac{\mathbf{m} - (M_s + M_b)/2}{\sigma^2}. \quad (3.26)$$

This significance variable is similar to X_M (with $M = (M_s + M_b)/2$) and only differs in the use of the variance instead of the standard deviation of the resolution in the denominator. The simple Gaussian example shows that while simple and intuitive, X_M may not always be optimal. However, as long as $g(\sigma)$ is not too broad, the difference between σ and σ^2 in the denominator should be minimal, since linearizations about a characteristic scale σ_0 will give similar results:

$$\frac{x - M}{\sigma} = \frac{\sigma_0}{2} \frac{x - M}{\sigma^2} + \text{constant} + \mathcal{O}((\sigma - \sigma_0)^2)$$

Now, consider a variant of the previous example with an asymmetric resolution function defined by the Gumbel distribution [583]:

$$p_i(m) = \frac{1}{\beta} \exp\left(\frac{m - M_i}{\beta}\right) \exp\left(-\exp\left(\frac{m - M_i}{\beta}\right)\right). \quad (3.27)$$

The Gumbel is chosen because its Taylor series is the same as a Gaussian with parameters μ and σ^2 up to the third order term in $\frac{m - M_i}{\beta}$ with the identification $\sigma = \frac{e}{\sqrt{2\pi}}\beta$ and $\mu = M_i$. Figure 3.32 overlays a Gaussian on top of the Gumbel distribution with this identification scheme. The two distributions have the same core, but the tail of the Gumbel distribution is heavier on the left than the right, which represents the generic case in which events are more likely to have smeared from lower values due to falling priors. Taking the logarithm of the likelihood ratio results in the following optimal significance variable:

$$V_{\text{opt}}^{(\text{Gumbel})} = \exp\left(\frac{m - M_b}{\beta}\right) - \exp\left(\frac{m - M_s}{\beta}\right) + \frac{M_b - M_s}{\beta}. \quad (3.28)$$

Lines of constant $V_{\text{opt}}^{(\text{Gumbel})}$ are shown in Fig. 3.32. When m is small compared to M_s and M_b , the first two terms in Eq. 3.28 are highly suppressed relative to the constant third term and so the lines of constant $V_{\text{opt}}^{(\text{Gumbel})}$ are horizontal lines in Fig. 3.32 on the left. This region is uninteresting as usually the region of interest is $m > M_b$ since M_b is often less than M_s . For $M_b < m < M_s$, the first term in Eq. 3.28 dominates so lines of constant p_s/p_b are well approximated by lines of constant X_M with $M = M_b$. These are found in Fig. 3.32 as straight lines radiating from $(m = M_b, \beta = 0)$. Finally, in the region in which $m > M_s$ and β is small compared $M_s - M_b$, both exponentials are large and so the dominant part of (3.28) can be re-written as




Figure 3.32: The Gumbel distribution and an equivalent Gaussian distribution.

$$\exp\left(\frac{m - \bar{M}}{\beta}\right) \sinh\left(\frac{M_s - M_b}{\beta}\right), \quad (3.29)$$

where $\bar{M}$ is the average of M_s and M_b . Since the $\sinh$ term is relatively smaller and slowly varying, lines of constant likelihood ratio are thus just lines of constant X_M with $M = \bar{M}$, which are once again straight lines in Fig. 3.32. This simple example shows that an optimal use of m , σ_m , and M is well approximated by a threshold requirement on X_M even when the resolution function is realistically asymmetric.

3.2.2.3 Choosing the Separation Scale M

The construction in Sec. 3.2.2.2 shows that M can play a dynamic role in the definition of X_M . First of all, note that an analysis that uses a threshold requirement on X_M is truly a generalization of an analysis that uses a threshold on m . Let c be the threshold for the latter analysis, such that signal-like events are those with $m > c$. Then, the events chosen by $X_c > 0$ will be identical to those chosen with $m > c$ and therefore there is always a choice of M that reduces the significance variable-based analysis to one based only on the kinematic variables themselves. In particular, an optimal analysis based on X_M can be no worse than one based on m alone and will




Figure 3.33: Contours of constant $V_{\text{opt}}^{(\text{Gumbel})}$ in the $(\mathfrak{m}, \beta)$ plane for $M_b = 80$ and $M_s = 85$. The contours are drawn at powers of 2 from $-2^3, \dots, -2, 0, 2, \dots, 2^8$.

likely be better since X_M incorporates more information and has an additional degree of freedom (M).

The interpretation of M as the scale of Standard Model physics does not require that it be fixed ahead of time, since detector resolutions can distort the *reconstructed* scale away from the *true* scale. It is often the case that the distribution of σ itself is independent of the underlying process and thus not useful for distinguishing signal and background. Another way to visualize how M mixes with σ to add discriminating power on top of $\mathfrak{m}$ is to consider the ordering of events induced by X_M versus $\mathfrak{m}$. For example, suppose that there are only two events with $\mathfrak{m}$ values $\mathfrak{m}_1, \mathfrak{m}_2$ and resolutions σ_1 and σ_2 . The quantity which controls the ordering of X_M is $\Delta \equiv (\mathfrak{m}_2\sigma_1 - \mathfrak{m}_1\sigma_2)/(\sigma_1 - \sigma_2)$. When $\Delta < 0$ or infinite in magnitude, then $X_M^1 > X_M^2$ for all M . However, if $\Delta > 0$, then there is a critical M^* such that for $M < M^*$, $X_M^1 > X_M^2$ for $M > M^*$, $X_M^1 < X_M^2$. The value of M^* is Δ . For $N > 2$, the situation is more complicated, but the result is the same; different values of M can rearrange the distribution of events based on X_M from the distribution based on M . One can generalize the plots in Figure 3.34 for $N > 2$. Note that the distribution of points of intersection with the M axis forms the observed distribution of $\mathfrak{m}$.




Figure 3.34: The dependance of X_M on M for two events with $\Delta \equiv (m_2\sigma_1 - m_1\sigma_2)/(\sigma_1 - \sigma_2) > 0$ in the left plot and $\Delta \in [-\infty, 0) \cup \{\infty\}$ in the right plot.

Consider a kinematic variable $\mathbf{m}$ with zero resolution maximum $\tilde{m}$. Intuitively, it may seem that M should be equal to or greater than $\tilde{m}$. However, the value of M which maximizes the significance $\hat{s}(c) \equiv s/\sqrt{b}$, for c a threshold value on X_m , could be less than $\tilde{m}$. If σ is constant over all events, X_M induces the same ordering on events as $\mathbf{m}$ and so any value of M maximizes $\hat{s}$. As an example, recall the model in Eq. 3.25. If the distribution of σ_m is also a delta function, then X_M and $\mathbf{m}$ will give the same significance. Therefore, take a simple extension:

$$g(\sigma) = p\delta(\sigma - \sigma_1) + (1 - p)\delta(\sigma - \sigma_2), \quad (3.30)$$

where σ_i are two fixed values of σ and $p \in [0, 1]$. With this simple model, one can easily compute the distributions of $\mathbf{m}$, X_M and $\hat{s}$, as seen in Figure 3.35 for $\tilde{m} = 80$ for the background, $\tilde{m} = 90$ for the signal, $p = 1/2$ and ρ is the signal efficiency, defined by $\rho(c) = \int_c^\infty dx f(x)$ for $f(x)$ the signal probability density function and c a cut value. Furthermore, $\sigma_1 = 5$ and $\sigma_2 = 10$. In this setup, there is an $M < \tilde{m}$ which outperforms the significance at $M = \tilde{m}$. This is seen clearly in the second plot of the figure in which the low value of M can allow for X_M to distinguish between low and high resolution events for the signal. In the limit as $\tilde{m} - M > \sigma$, X_M will be able to distinguish the low and high resolution events, thus increasing $\hat{s}$. For $\tilde{m} - M \gg \sigma$,




Figure 3.35: These plots illustrate the distributions of m , X_M and $\hat{s}$ for a simple model in which m is always ‘on shell’ at 80 for the background and 90 for the signal. The resolutions can take one of two values with probability $1/2$, independent of the physics process. The optimal (V_{opt}) variable is based on the likelihood ratio detailed in Sec. 3.2.2.2 and Eq. 3.26.

the efficacy of X_M approaches the constant resolution case and so one cannot gain more by decreasing M .

Before proceeding, here are two further remarks about the above toy model. First, note that due to the simplicity, one can actually derive the optimal variable, described in Sec. 3.2.2.2 as $V_{\text{opt}} = (m - \bar{M})/\sigma^2$, where $\bar{M}$ is the average of the signal and background true m values. The distribution and significance of V_{opt} are shown alongside X_M in Fig. 3.35. One can see that while X_M does no better than the optimal variable, for the appropriate choice of M it can have essentially the same maximum significance. A second remark is that since the distribution of σ for signal and background is identical, the resolution alone cannot distinguish signal and background. Thus, the improved performance of X_M over m is due entirely to the event-by-event combination of m , M , and σ to capture resolution and kinematic properties of the reconstructed objects.

3.2.2.4 Empirical Examples

This section contains a few illustrative examples of significance variables using realistic physics processes but simplistic models for the detector resolution. Events are generated using PYTHIA 8.170 [96, 189] to cover three canonical searches that exploit endpoints in kinematic distributions. The resolution of the missing momentum is modeled as $\sigma_{E_{x,y}^{\text{miss}}} = 0.5\sqrt{\sum E_T}$, where $\sum E_T$ is the sum of all visible momentum and follows the measured spectra in dijets [137]. The distributions of X_M are computed by numerically propagating the underlying kinematic uncertainties.

A first example is the $W' \rightarrow \mu\nu$ search using the transverse mass of the muon and the neutrino as the main discriminant. In this search, the W mass is a natural choice for M in constructing X_M , where $m = m_T$. The W' boson is created with a mass of 100 GeV¹⁵ and the SM CKM matrix. The distributions of m_T , X_M and $\hat{s}$ are shown in Fig. 3.36. The various rows of Fig. 3.36 demonstrate the affect of the W width on the efficacy of X_M . For a vary narrow resonance background, X_M is much better than m_T , but as the width becomes large, the advantage decreases.

Another possible use of the m_T significance is in the standard $H \rightarrow \tau\tau$ search (measurement) [586, 587] where the di-tau system is the ‘visible particle’ in the calculation of the transverse mass. In the dilepton channel, the dominant background is Z boson production and so the natural value for M is 90 GeV. Figure 3.37 shows the distributions of m_T , X_M , and $\hat{s}$ for a 125 GeV Higgs. The optimal value of M was found to be less than 90, as indicated in the diagram. The $\hat{s}$ figure shows that there can be a significant improvement from X_M over m_T by about 20%.

A third illustrative example is the pair production of stops with $\tilde{t} \rightarrow t + \text{LSP}$ in the dilepton channel using m_{T2} . With the leptons as the visible particles in the definition of m_{T2} , this system once again has the feature that the resolution is mostly due to the missing momentum vector. With $t\bar{t}$ as the dominant background, the natural scale is $M = 80$ GeV. The m_{T2} distribution, m_{T2} significance, and $\hat{s}$ are shown in Fig. 3.38 for a compressed scenario of $m_{\text{stop}} = 350$ GeV and $m_{\text{LSP}} = 170$ GeV. The use of X_M improves the significance by about 30% over m_{T2} alone.

¹⁵Excluded by [584, 585], useful here for illustration only.




Figure 3.36: In each row, the left plot compares the transverse mass distribution for a Standard Model W and a W' with mass 100 GeV. The middle plot is the corresponding distributions of X_M with $M = 80$ GeV. The right plot shows the rejection $s\sqrt{b}$ as a function of the signal efficiency, in arbitrary units. The bands show the statistical uncertainty due to limited Monte Carlo statistics. The top row has a boson mass width of 0, the middle has a width of 20% of the natural width, and the bottom row has the full width of about 2 GeV [27]. We can see that for this fixed value of M , the performance of X_M is better than m_T for a narrow width and then worse at higher width. By construction, X_M cannot be worse than m_T and thus the optimal M in the last row must be different than 80. The inset plot shows X_M for $M = 100$, for which the performance of X_M and m_T is the same.




Figure 3.37: The left plot is the m_T distribution for dileptonic $Z \rightarrow \tau\tau$ and $H \rightarrow \tau\tau$ for a 125 GeV Higgs. The middle plot is the corresponding X_M curve with $M=70$ and the right plot is the rejection versus efficiency relationship.




Figure 3.38: The left plot is the m_{T2} distribution for dileptonic $t\bar{t}$ and $\tilde{t} \rightarrow t + \text{LSP}$ for a 350 GeV stop and 170 GeV LSP. The middle plot is the corresponding X_M curve with $M=80$ and the right plot is the rejection versus efficiency relationship.

3.2.2.5 Examples with Full Detector Simulation

The examples in Sec. 3.2.2.4 show that the additional information from X_M can be useful for improving the significance of bounded kinematic variables. This section re-focuses on the stop search in the one lepton channel which makes use of several variables with kinematic edges (see Sec. 3.2.1). Section 3.2.2.5.1 begins with the construction of approximate significance variables that continue to utilize the relationship $\sigma_{E_T^{\text{miss}}} \propto \sqrt{H_T}$. A more complex set of variables based on known jet-by-jet resolutions are built in Sec. 3.2.2.5.2.

3.2.2.5.1 Approximate m_T Significance

An approximate m_T significance is defined as $X_{m_T} = (m_T - M)/\sigma$, where σ is constructed from $\sqrt{H_T}$ and E_T^{miss} by linearly propagating uncertainties assuming no angular resolution. In the massless approximation, $m_T^2 = 2E_T^{\text{miss}} p_T^\ell (1 - \cos(\theta))$ where p_T^ℓ is the transverse momentum of the lepton and θ is the angle between $\vec{p}_T^\ell$ and $\vec{p}_T^{\text{miss}}$. Linear error propagation results in the following formula for σ :

$$\sigma \propto (p_T^\ell)^2 (1 - \cos(\theta))^2 \sigma_{E_T^{\text{miss}}}^2 + (\text{term proportional to } \theta \text{ resolution}). \quad (3.31)$$

Neglecting the θ resolution and modeling $\sigma_{E_T^{\text{miss}}}^2 \propto H_T$ results in $\sigma \propto m_T \sqrt{H_T}/E_T^{\text{miss}}$. Figure 3.39 compares the distributions of m_T and X_{m_T} (with $M = 100$ GeV) in simulations of $t\bar{t}$ and stop pair production. The m_T significance distribution for the background falls off below 0 while the peak in the signal is greater than zero. A quantitative comparison of the performance between m_T and X_{m_T} is in Fig. 3.41. An approximate statistical significance is given by $s/\sqrt{b}$, where s and b are the signal and background yield after a threshold requirement on m_T or X_{m_T} . The statistical significance of X_{m_T} is nowhere worse than m_T , even without a thorough optimization of M . At the peak of the statistical significance, around a signal efficiency of $\sim 1/3$, there is a $\sim 10\%$ improvement when incorporating the resolution information. The joint distribution of m_T and X_{m_T} in Fig. 3.42 shows that there is a strong relationship

between these two variables, as might be expected from the simple form of σ in Eq. 3.31. However, there is still a significant spread, which leads to the improvement in Fig. 3.41.




Figure 3.39: The distributions of m_T (left) and X_{m_T} (with $M = 100$ GeV) (right) for $t\bar{t}$ and stop pair production.




Figure 3.40: A comparison of m_T (left), σ (middle), and X_{m_T} (right) in data and simulation at $\sqrt{s} = 8$ TeV.

An important test for any new variable is the ability of the simulation to model the data. Figure 3.40 compares the numerator and denominator of X_{m_T} as well as the significance variable itself. All three quantities are well-modeled, suggesting that it is ready for use. However, the m_T significance constructed in this section is rather simple - a more sophisticated approach to significance variables that will be used for signal region optimization is discussed in the next section.




Figure 3.41: A comparison of the statistical significance after a threshold requirement on m_T and the (approximate) m_T significance, X_{m_T} . The background is $t\bar{t}$ and the signal is a stop model with $(m_{\text{stop}}, m_{\text{LSP}}) = (350, 150)$ GeV. This statistical significance is the ‘discovery significance’ because it quantifies the number of standard deviations a signal would be above the background-only noise.




Figure 3.42: The joint distribution of m_T and X_{m_T} with $M = 100$. When $m_T = 100$, $X_{m_T} = 0$ by construction.

3.2.2.5.2 $H_{T,\text{sig}}^{\text{miss}}$ Significance

The p_T and η dependence of the jet resolutions are well-understood in simulation and have been well-measured in data (see e.g. Ref. [122]). Parameterizations of the resolutions can be used to calculate resolutions for kinematic quantities that depend on jets event-by-event. Consider a quantity similar to the E_T^{miss} called the H_T^{miss} :

$$H_T^{\text{miss}} = \left| \sum_{\text{jets } j} \vec{p}_{T,j} + \vec{p}_T^\ell \right|, \quad (3.32)$$

where the sum runs over all signal jets and the momentum of the lepton $\vec{p}_T^\ell$. The symbol H is used instead of E to indicate that only the *hard-objects* are used to construct H_T^{miss} , whereas E_T^{miss} also includes energy not associated with signal jets and leptons. As expected, there is a strong correlation between the two definitions for high E_T^{miss} when the contribution from these softer energy sources is small. Figure 3.43 shows the distribution of E_T^{miss} conditioned on H_T^{miss} for $t\bar{t}$ events with $E_T^{\text{miss}} > 100$ GeV. For $H_T^{\text{miss}} \gtrsim 100$ GeV, there is a strong correlation with well over 50% of E_T^{miss} values within 15-30 GeV of the H_T^{miss} . The advantage of Eq. 3.32 is that the resolutions of the jets are known parametrically and so the resolution $\sigma_{H_T^{\text{miss}}}$ can be computed as

$$\sigma_{H_T^{\text{miss}}}^2 = \frac{1}{N} \sum_{i=1}^N \left(\sum_{\text{jets } j} \Sigma_i^j \vec{p}_{T,j} + \vec{p}_T^\ell \right)^2 - \left(\frac{1}{N} \sum_{i=1}^N \left| \sum_{\text{jets } j} \Sigma_i^j \vec{p}_{T,j} + \vec{p}_T^\ell \right| \right)^2, \quad (3.33)$$

where Σ_i^j is a diagonal two-by-two matrix with entries $1 + z_i^j$, for $z_i^j \sim \mathcal{N}(0, \sigma(p_{T,j}, \eta_j))$. To suppress the fluctuations in the calculation of $\sigma_{H_T^{\text{miss}}}^2$, N in Eq. 3.33 is chosen to be 1000. As expected, there is a strong relationship between the approximate E_T^{miss} resolution and the resolution computed with Eq. 3.33 (Fig. 3.44). By incorporating more local information about the resolution, $\sigma_{H_T^{\text{miss}}}$ should be a better approximation to the full significance. The H_T^{miss} significance, $H_{T,\text{sig}}^{\text{miss}} = (H_T^{\text{miss}} - M)/\sigma_{H_T^{\text{miss}}}$.




Figure 3.43: The distribution of E_T^{miss} conditioned on H_T^{miss} in bins of 20 GeV for each variable in $t\bar{t}$ events. All events have $E_T^{\text{miss}} > 100$ GeV.




Figure 3.44: The joint distribution of the approximate E_T^{miss} resolution, $\sqrt{H_T}$ and the H_T^{miss} resolution computed via Eq. 3.33.

A quantitative comparison of $H_{T,\text{sig}}^{\text{miss}}$ with the traditional $E_T^{\text{miss}}/\sqrt{H_T}$ using the separation power from Eq. 3.20 is shown in Fig. 3.46. The separation is largest for $M = 100$ GeV, with a $\sim 15\%$ improvement over $E_T^{\text{miss}}/\sqrt{H_T}$. One of the disadvantages of $E_T^{\text{miss}}/\sqrt{H_T}$ is that it is strongly correlated with E_T^{miss} . As a result of neglecting the soft energy and due to the scale shift M , $H_{T,\text{sig}}^{\text{miss}}$ can be less correlated with E_T^{miss} . This intuition is quantified in Fig. 3.47 which shows that the linear correlation is smaller for all considered values of M . Table 3.5 summarizes the information from Fig. 3.46 and Fig. 3.47 and shows that $H_{T,\text{sig}}^{\text{miss}}(M = 100)$ is strictly better than $E_T^{\text{miss}}/\sqrt{H_T}$ in the important metrics considered here and is therefore chosen as baseline for optimizations studies in later chapters. Furthermore, Fig. 3.45 indicates that this $H_{T,\text{sig}}^{\text{miss}}$ (the $M = 100$ GeV is henceforth dropped) is well-modeled by the simulation¹⁶.




Figure 3.45: A comparison of data and simulation using a loose selection requiring exactly one signal lepton, four jets with $p_T > 25$ GeV and at least one b -tagged jet. The left and right plots differ only in the scaling of the vertical axis. The uncertainty band includes jet energy scale and resolution uncertainties (see Sec. 6).

¹⁶In order to reduce data/MC differences in the measured jet resolutions, *the same* (simulation) resolution parameterizations are used for data and simulation.




Figure 3.46: The distributions of $H_{T,\text{sig}}^{\text{miss}}$ with $M = 50$ GeV (bottom left), $M = 100$ GeV (top left), and $M = 0$ GeV (top right) along with the distribution of $E_T^{\text{miss}} / \sqrt{H_T}$ (bottom right) for $t\bar{t}$ and stop events.




Figure 3.47: The joint distribution of $H_{T,\text{sig}}^{\text{miss}}$ or $E_T^{\text{miss}}/\sqrt{H_T}$ with E_T^{miss} . The linear correlation is indicated in each plot.

Variable	Separation (650,1)	Separation (350,50)	Correlation with E_T^{miss}
E_T^{miss}	0.59	0.22	1.00
$E_T^{\text{miss}}/\sqrt{H_T}$	0.50	0.18	0.71
$H_{T,\text{sig}}^{\text{miss}} (M = 100)$	0.58	0.21	0.68
$H_{T,\text{sig}}^{\text{miss}} (M = 0)$	0.50	0.18	0.50

Table 3.5: A summary of the separation power and correlation with E_T^{miss} for $H_{T,\text{sig}}^{\text{miss}}$ and $E_T^{\text{miss}}/\sqrt{H_T}$ based on Fig. 3.46 and Fig. 3.47.

3.2.3 Tau veto

Dilepton $t\bar{t}$ events where one of the two leptons is a τ that decays hadronically is a major background to the search because the extra neutrinos allow events to evade m_T and E_T^{miss} thresholds and the hadronic activity contributes an extra jet to meet n_{jet} requirements. Section 3.1 introduced the explicit hadronically decaying τ reconstruction algorithms used at both $\sqrt{s} = 8$ and $\sqrt{s} = 13$ TeV. This section explores how to create a powerful τ veto while maintaining a nearly 100% efficiency for events without a hadronically decaying τ . Figure 3.48 shows that most hadronically decaying τ leptons are reconstructed as a signal jet. Hadronically decaying τ leptons with $|\eta| < 2.5$ are not reconstructed as a signal jet about 15% of the time due to the $p_T > 25$ GeV threshold. Even if a τ lepton has $p_T > 25$ GeV, a significant fraction of its energy can be lost to unmeasured neutrinos. To begin, the next paragraph describes important properties of the $\sqrt{s} = 13$ TeV τ reconstruction efficiency.




Figure 3.48: The pseudorapidity (left) and p_T (right) distributions of hadronically decaying τ leptons from W decays in $t\bar{t}$ events decomposed by how the τ is reconstructed (if at all). A τ lepton in the simulation is *matched* to a jet if $\Delta R < 0.4$. If the hadronic decay products of the τ constitute $\leq 90\%$ of the jet p_T , the τ is labeled as *hiding in a jet*. Events are required to have exactly one signal lepton, at least four signal jets, at least one b -tagged jet, $E_T^{\text{miss}} > 200$ GeV, $m_T > 150$ GeV, and at least one large-radius jet with $p_T > 150$ GeV (see Sec. 3.2.4).

Three efficiency working points are optimized by combining track and calorimeter information such as the mass of the four-vector sum of tracks in the core of the jet [549, 550]. The points are called *loose*, *medium*, and *tight* and correspond to efficiencies for one- (three-)prong decays of about 60% (50%), 55% (40%), and 45% (30%), respectively. Figure 3.49 shows the probability for a hadronically decaying τ lepton from a W boson decay to be reconstructed and pass additional criteria. With the same event selection as Fig. 3.48, about 77% of τ leptons are reconstructed as signal jets. Since the jets used for the dedicated τ reconstruction have a threshold lower than for signal jets (20 GeV versus 25 GeV), the efficiency to reconstruct a τ as a ‘reco τ ’ is slightly higher than for signal jets by about 5%. The number of tracks inside a jet is a powerful τ discriminant because it is p_T -independent for τ jets and increases with p_T for quark and gluon jets (see Chapter 3). The third bin of Fig. 3.49 shows the efficiency for reconstructing τ leptons as signal jets with less than five tracks. This simple τ identification scheme has a similar efficiency to the dedicated τ reconstruction with exactly one or three tracks. Tracks for the dedicated algorithm are only chosen from the jet core, $\Delta R < 0.2$. Additionally requiring that the reco tau has opposite electric charge to the signal lepton reduces the efficiency by a few percent and a $p_T > 20$ GeV threshold further lowers the efficiency relatively by about 10%. The overall efficiency of also applying the loose, medium, or tight identification criteria is about 38%, 34%, or 28%, respectively.

The reconstruction efficiency in Fig. 3.49 does not include an important efficiency from combinatorics. A reconstructed object may match to a generator-level τ lepton with high efficiency, but if an event has many such objects, then the ambiguity lowers the efficiency for reconstructing the correct τ momentum. The left plot of Fig. 3.50 shows the probability that various objects are matched to the generator-level hadronically decaying τ lepton. The Run 1 scheme that uses the leading non b -tagged jet to form m_{T2}^τ (see Sec. 3.2.1.2.3) has a low (20%) efficiency for correctly selecting the τ lepton, while the n_{track} -based taggers have a much higher (40-50%) efficiency. Additionally, the right plot of Fig. 3.50 shows that the Run 1 scheme often chooses the wrong object as the leptonic τ , while the dedicated τ reconstruction algorithms almost never picks the wrong object. In other words, even though the leading non




Figure 3.49: The frequency for a hadronically decaying τ lepton to be reconstructed as one of a variety of objects, described by the labels on the horizontal axis.

b -tagged jet and the leading reco τ with a tight identification have similar efficiencies, the former is not the τ about 80% of the time while the latter is not the $\tau < 1\%$ of the time. A similar trend is true for events without a τ lepton from a W boson decay. Figure 3.51 shows the probability that a particular object is selected as a hadronic τ candidate when there is no particle-level τ lepton in the event. The Run 1 scheme was chosen so that every event has a m_{T2}^{τ} value; therefore it has a 100% probability of picking an object in Fig. 3.51. In contrast, loose, medium, or tight reco τ algorithms only have candidates in 7%, 6%, 4% of events, respectively.

The two main points from Fig. 3.49, 3.50, and 3.51 are that the ‘ τ ’ used for m_{T2}^{τ} does not usually correspond to a particle-level τ and the highest veto efficiency in stop events (no actual τ) using one of the dedicated reco τ algorithms is 93%-96% (rightmost bins of Fig. 3.51). One solution can improve both of these statistics: combining kinematic information from m_{T2}^{τ} with identification information from the reco τ algorithms. A new $m_{T2}^{\tau}(\text{ID})$ variable is formed by using a reco τ with a particular identification algorithm (ID) as the visible particle for m_{T2} . By construction, this variable can only be calculated a small fraction of the time. Instead of vetoing events if a reco τ exists, events are only vetoed if $m_{T2}^{\tau}(\text{ID}) \geq X$. When $X = 0$, then the veto




Figure 3.50: Given an event has a hadronically decaying τ , the frequency that a given τ identification technique selects the τ correctly (left) or incorrectly (right).




Figure 3.51: Given an event does not have a hadronically decaying τ , the frequency that a given τ identification algorithm produces a reconstructed τ candidate.

has the 93%-96% efficiency quotes above for stop events and a 62%-72% efficiency for background τ events (Fig. 3.49). However, as $X \rightarrow \infty$, the veto is 100% efficient for signal events. The goal is to optimize X and ID to achieve a $\sim 99\%$ efficiency for

signal events and the best possible rejection of τ events.

Figure 3.52 summarizes the efficiencies for all combinations of τ identification algorithms and m_{T2}^τ . The ΔR between the τ candidate and the leading large-radius jet also provides useful information for rejecting events with a hadronically decaying τ lepton. In single lepton $t\bar{t}$ (and stop) events, the (fake) reco τ is usually within the large-radius jet, while in events with a τ , there can be a large separation between the τ and the jet¹⁷. For all the combinations in Fig. 3.52, the threshold requirement on ΔR and m_{T2}^τ are optimized (where possible) so that the signal efficiency is 99%. The best combination is for a loose τ identification and a threshold requirement on $m_{T2}^\tau \gtrsim m_W \approx 80$ GeV. This combination is used for the $\sqrt{s} = 13$ GeV signal region optimization described in Chapter 4.




Figure 3.52: A summary figure with all the considered hadronic τ vetoes based on the above considerations. The (black) red line shows the efficiency for $t\bar{t}$ events with(out) a hadronically decaying τ lepton. The signal $(m_{\text{stop}}, m_{\text{LSP}}) = (800, 1)$ GeV is has mostly semi-leptonic $t\bar{t}$. The ΔR is between the hadronically decaying τ candidate and the leading large-radius jet. In addition to the loose, medium, and tight reco τ working points, a simple $n_{\text{track}} = n_{\text{trk}} < 5$ identification scheme is part of the comparison. The dashed line is at an efficiency of 90%.

¹⁷In particular when leptons are part of the jet clustering, the large-radius including these leptons will tend to be harder than one from hadronically decaying τ leptons due to the lost energy in neutrinos.

In addition to the improvement in the signal efficiency of the modified m_{T2}^τ veto, the new variable is significantly less correlated with m_T (see Table 3.4). The correlation is reduced by over a factor of 10 for the $t\bar{t}$ background and by a factor of about two in the signal. Figure 3.53 shows that the m_{T2}^τ distribution using a loose reco τ as one of the visible particles is relatively well-modeled and as expected, $m_{T2}^\tau \lesssim m_W$ for the background. For the signal, m_{T2}^τ often significantly exceeds m_W , with only a small peak at $m_{T2}^\tau = 0$ corresponding to the unbalanced case¹⁸.




Figure 3.53: The distribution of m_{T2}^τ with an event selection enriched in dilepton $t\bar{t}$ events with a hadronically decaying τ . Events are required to have at least one reconstructed reco τ passing the loose identification. In addition, events must have at least four jets with $p_T > 80, 50, 40, 25$ GeV, $E_T^{\text{miss}} > 200$ GeV and at least one b -tagged jet. In order to suppress semi-leptonic $t\bar{t}$ events, $m_T > 100$ GeV. See Sec. 5.1.1.1 for details.

¹⁸It is stated in Sec. 3.2.1.2.2 that the unbalanced case can only occur when $m_{C_1} + m_{V_1} \neq m_{C_2} + m_{V_2}$. However, when the $m_{V_i} = m_{C_i} = 0$, the m_T^2 surfaces allow for the minimum value to be reached even if $E_T^{\text{miss}} > 0$. This is related to Fig. 3.14 and is described in Ref. [588] in detail.

3.2.4 Hadronic Top Mass Reconstruction

Much of the focus in the previous chapters is centered on identifying leptonically decaying top quarks for suppressing the dilepton $t\bar{t}$ background. Another possibility is to target *hadronically* decaying top quarks that are present in the mostly semileptonic $t\bar{t}$ signal, but absent in the dileptonic $t\bar{t}$ background. Hadronically decaying top quarks produced with a small or moderate boost often result in three daughter jets¹⁹. The first top-tagging technique for the stop search was introduced in the $\sqrt{s} = 7$ GeV analysis [501]. A relatively unoptimized simple combination of jets tries to capture a hadronically decaying W boson matched with another jet to give the full top quark decay:

1. Let j_1 and j_2 be the two jets with $m_{j_1 j_2} > 60$ GeV closest in ΔR (W candidate). If no such jets exist, set $m_{\text{had top}} = 0$.
2. Take the signal jet j_3 with $m_{j_1 j_2 j_3} > 130$ GeV closest in ΔR to the dijet system $(j_1 + j_2)$. If no such jet exist, set $m_{\text{had top}} = 0$.
3. Define $m_{\text{had top}} = m_{j_1 j_2 j_3}$.

In signal events with a hadronically decaying top quark, it is expected that $m_{\text{had top}} \sim m_{\text{top}}$. The early $\sqrt{s} = 8$ TeV signal regions use a loose top-quark tag defined by $0 < m_{\text{had top}} < m_{\text{top}} + 30$ GeV = 205 GeV. Two main challenges with this tag are that (a) the top quark and W daughter jets need not be close when the top quark is produced with a small boost and (b) the signal efficiency for even the loose tag $0 < m_{\text{had top}} < 205$ GeV is only $\sim 50\%$ - 60% in the signal. In addition, there is more useful information in the event that can be used to improve the hadronic top quark mass candidate such as b -tag weights and jet resolutions. This additional information could be used to select the jets and dynamically vary the thresholds in the algorithm that should depend on the jet resolutions. Therefore, for the search based on the entire $\sqrt{s} = 8$ TeV dataset, a new hadronic top quark identification technique is

¹⁹This is an ill-defined notion, especially since the top quark is not colorless. See Sec. 2.2.3.1 for detail. In this context, the statement about the number of daughter jets is used heuristically and not quantitatively.

developed. The new variable exists for all events (100% efficiency for $m_{\text{top},\chi^2} > 0$) and is given by:

1. Let b_1, b_2 be the signal jets with the highest b -tagging weight (consistent with the choice used for m_{T2} in Sec. 3.2.1.2.3).
2. Consider all pairs j_1, j_2 of signal jets that are not b_1, b_2 .
3. Compute $\sigma_{m_{j_1 j_2}}^2 = m_{j_1 j_2}^2 (r_1^2 + r_2^2)$ and $\sigma_{m_{j_1 j_2 j_3}}^2 = m_{j_1 j_2 j_3}^2 (r_1^2 + r_2^2 + r_3^2)$, where r_i is the fractional energy uncertainty of jet i (same as for $H_{T,\text{sig}}^{\text{miss}}$ in Sec. 3.2.2.5.2).
4. Select j_1, j_2 and i to minimize the following:

$$\chi^2 = \frac{(m_{j_1 j_2 j_3} - m_{\text{top}})^2}{\sigma_{m_{j_1 j_2 j_3}}^2} + \frac{(m_{j_1 j_2} m_W)^2}{\sigma_{m_{j_1 j_2}}^2}. \quad (3.34)$$

Figure 3.54 compares the ‘simple’ and χ^2 -based approaches for the dileptonic $t\bar{t}$ background and a stop model with $(m_{\text{stop}}, m_{\text{LSP}}) = (600, 250)$ GeV. The top quark mass peak in the signal is sharper for the signal, but due to Eq. 3.34, the background also has a peak around m_{top} . The separation power (Eq. 3.20) is about 0.035 for the simple definition and 0.05 for χ^2 -definition. In addition, the signal ($t\bar{t}$) efficiency for a $m < 200$ GeV threshold increases from about 60% (47%) for the simple definition to about 69% (54%) for the χ^2 -definition.

When $m_{\text{stop}} \gtrsim 700$ GeV so that $p_T^t \gtrsim 350$ GeV, the top quark hadronic decay products are sufficiently collimated that a single large-radius jet can capture most of the energy. For the stop search at $\sqrt{s} = 13$ GeV that targets such high mass stops, re-clustered trimmed jets are used to form a hadronic top quark mass from m_{jet} . Chapter 4 describes boosted top quarks, large-radius jets, and re-clustering in detail. One of the benefits of re-clustering is that the jet algorithm parameters can be easily optimized for each signal region. Section 4 describes the parameter optimization in the case of the stop search. Large-radius jets are formed with signal small-radius jet inputs after the overlap removal. Figure 3.55 shows the modeling of the jet mass distribution in a selection enriched in semi-leptonic $t\bar{t}$ events. There are peaks at the W and top quark masses and most of the large-radius jets near m_W have

two constituent small-radius jets while most of the jets near m_{top} have at least three small-radius jet constituents.




Figure 3.54: A comparison of the two resolved hadronic top quark mass tagging techniques described in the text. The stop model is $(m_{\text{stop}}, m_{\text{LSP}}) = (600, 250)$ GeV.




Figure 3.55: The large-radius jet mass in a one-lepton $t\bar{t}$ selection requiring $\Delta\phi(\text{jet}_i, \vec{p}_T^{\text{miss}}) > 0.4$, at least four jets with $p_T > 80, 50, 40, 40$ GeV, $E_T^{\text{miss}} > 200$ GeV, $m_T > 30$ GeV, at least one b-jet, at least one large-radius jet with $R = 1.2$ with $p_T > 200$ GeV and $m_{\text{jet}} > 50$ GeV. Leptons are not included in the re-clustering and small-radius jets are dropped (trimmed) if they have $p_T < 5\%$ of the jet p_T . The jet mass distribution is decomposed by constituent multiplicity.

Chapter 4

Signal Regions

A series of increasingly complex event selections starting with the early $\sqrt{s} = 8$ TeV data and covering the beginning of the $\sqrt{s} = 13$ TeV data in Run 2 are sensitive to a broad region of the light stop parameter space. Figure 4.3 presents an overview of the expected limits in the $m_{\text{stop}}, m_{\text{LSP}}$ mass plane. This chapter covers the optimization of seven signal regions from three datasets. With 13 fb^{-1} of $\sqrt{s} = 8$ data, SR1, SR2, and SR3 extend the $\sqrt{s} = 7$ TeV limits at intermediate masses, to higher neutralino masses, and to higher stop masses respectively. With the full Run 1 $\sqrt{s} = 8$ TeV dataset (20.3 fb^{-1}) three additional signal regions further extend the limits in all three directions. The tN_diag signal region pushes the limit at low stop mass toward the challenging *diagonal* in the $m_{\text{stop}}, m_{\text{LSP}}$ mass plane where $m_{\text{stop}} \approx m_{\text{LSP}} + m_{\text{top}}$. Kinematically tighter regions tN_med and tN_high are analogues to SR2 and SR3 and extend the sensitivity at high stop and neutralino masses. Even though the dataset at $\sqrt{s} = 13$ TeV is significantly smaller than at $\sqrt{s} = 8$ TeV (only 3.2 fb^{-1}), the significant increase in the stop cross section coupled with new techniques allows the early Run 2 dataset to further expand the sensitivity to nearly $m_{\text{stop}} = 800 \text{ GeV}$.

Each signal region is individually optimized starting from a loose event selection (Sec. 4.1) using procedures described in 4.2. There are two classes of signal regions: single-bin and multi-bin regions. The single-bin regions are documented in Sec. 4.3 and the tN_diag multi-bin region optimization and final event selection is described in Sec. 4.4.




Figure 4.1: An overview of the various signal regions described in Chapter 4. The horizontal axis is the stop mass and the vertical axis is the neutralino mass; each point in this plane corresponds to a simplified stop model. Model cross-sections are set by the stop mass and decrease from left to right. The dashed line corresponds to the kinematic boundary above which stop decay to an on-shell top quark is forbidden. The various lines and shaded region are the expected exclusion limits using the statistical procedures documented in Sec. 7.1. Seven signal regions are optimized to ensure a broad sensitivity across the plane. Names of the signal regions are placed in the locations of parameter space where they add the most unique sensitivity.

4.1 Preselection

The starting point for the signal region optimization is a loose event selection (*preselection*) with many of the irrelevant backgrounds already suppressed. This preselection includes the trigger and isolated lepton requirements in addition to the second lepton veto. Events are further required to have at least four signal jets, at least one b -tagged jet, $E_T^{\text{miss}} > 100$ GeV, and $m_T > 30$ GeV. After the preselection, the Z +jets and QCD multijets backgrounds are negligible (see Chapter 5) and $t\bar{t}$ events dominate. Additional intermediate preselections are used to reduce the gap in phase space to the potential signal regions. For example, at $\sqrt{s} = 13$ TeV, the preselection used for the SR13 optimization additionally required $E_T^{\text{miss}} > 150$ GeV, $m_T > 100$ GeV, and $|\Delta\phi(\text{jet}_i, \vec{p}_T^{\text{miss}})| > 0.4$ for $i = 1, 2$ (to suppress mis-measured E_T^{miss}). Figure 4.2 shows the jet p_T spectra after this preselection. The pair production of top quarks is the dominant process and due to the relatively high E_T^{miss} requirement, the leading jets have a hard p_T spectrum. Additional distributions with the SR13 preselection appear in Sec. 4.3.

All signal region optimizations are performed prior to observing the data in signal-like regions of phase space. In order to avoid looking at data with signal sensitive event selections while still monitoring the data in looser event selections, the optimization is performed *blinded*. Prior to the finalization of the SR definitions, all data (and simulation) passing the preselection, $E_T^{\text{miss}} > 200$ GeV, and $m_T > 140$ (150) GeV at $\sqrt{s} = 8$ (13) TeV are removed from monitoring plots. Blinding does not effect simulation-only studies such as the optimization described in Sec. 4.2. None of the plots shown in subsequent sections have the blinding applied as all SR are now fixed.




Figure 4.2: The distribution of the leading (top left), second leading (top right), third leading (bottom left) and fourth leading (bottom right) jet p_T spectra after the preselection at $\sqrt{s} = 13$ TeV. All plots are normalized to unity. Note that the horizontal axes have been chosen per distribution so that the distributions fill out the entire plot.

4.2 Optimization Procedure

In principle, the best event selection using only threshold requirements on a series of N variables is obtained by scanning the entire $\sim \mathbb{R}^N$ space and computing the test statistic for each point, taking the selection that optimizes this statistic. This is never possible in practice for the following reasons:

Number of combinations If each dimension is discretized into ~ 10 intervals and there are ~ 10 variables, then the number of combinations is $\# \text{intervals}^{\# \text{variables}} \sim 10^{10}$. Ten intervals is already coarse; if instead there are ~ 50 intervals, then this number is $\sim 10^{16}$. Large regions of parameter space can be eliminated based on simple criteria (e.g. no signal events remain), but this is still an unfeasible number of combinations to check.

Test Statistic Evaluation The full test statistic (described in Sec. 7.1) for a given selection takes $\mathcal{O}(10 \text{ seconds})$ to compute. With 10^{10} configurations and 1000 full time batch nodes, the brute force optimization would require ~ 3 years.

For optimizing SR1, SR2, and SR3, the two challenges above were addressed by (1) only considering a small number of combinations and (2) using a simplified version of the test statistic. About 10,000 total combinations of threshold requirements on $\mathbf{am}_{T2}$, $\mathbf{m}_{T2}^{\tau}$, E_T^{miss} , $\mathbf{m}_T$, $E_T^{\text{miss}}/\sqrt{H_T}$, $\mathbf{p}_T^{\text{jet } 1}$, hadronic top mass, and the isolated track veto were studied using a brute force approach. The simplified metric for (2) is $s/\sqrt{b + (0.25 \times b)^2}$, for s signal events and b background events. This formula is a comparison of s signal events to the background uncertainty that is the sum in quadrature of a $\sqrt{b}$ Poisson uncertainty with a 25% background systematic uncertainty. If all yields could be treated as Gaussian, then this *significance* value would be the Z-score with $Z > 2$ corresponding to a 2σ sensitivity. This simple optimization procedure was able to quickly converge on signal regions that could extend beyond the $\sqrt{s} = 7$ TeV performance, even after adjusting for the increase in luminosity. For example, at $(\mathbf{m}_{\text{stop}}, \mathbf{m}_{\text{LSP}}) = (500, 0)$, the significance was increased by a factor of ~ 3 . Because the procedure was so simple, each point in the coarse $(\mathbf{m}_{\text{stop}}, \mathbf{m}_{\text{LSP}})$ plane was

separately optimized. Small ad-hoc adjustments of the optimized thresholds resulted in three distinct signal regions with broad sensitivity across the parameter space.

A more sophisticated approach was used for the analysis of the full $\sqrt{s} = 8$ TeV dataset which is a closer approximation to the optimal configuration described above. For a given benchmark model (one for each target signal region), a set of selections is chosen that minimizes the background composition for a fixed number of predicted signal events. This is repeated for a scan in the number of signal events between 5 and 10. For each signal efficiency, the final test statistic is computed and the point with the best value is selected. If the minimization step works successfully, then the selected signal region will be *globally optimal* because all (reasonable) test statistics will be improving¹ for decreasing background yield for a fixed signal efficiency. The minimization is performed using the Nelder-Mead simplex method [589] as implemented in the MINUIT [590] package. A simplex-based routine is used instead of the more standard Davidon-Fletcher-Powell switching method [591–593] (called MIGRAD in MINUIT) because the number of simulated events passing a selection is discrete so methods based on derivatives do not perform well. The objective function is:

$$f(\mathbf{b}, s, S) = \begin{cases} -s/b & s > S \\ -s/b + g(s - S) & s \leq S \end{cases}, \quad (4.1)$$

where $g(x)$ is a penalty function that forces s to be close to the target signal yield S . The optimization of the $\sqrt{s} = 8$ TeV data analysis used an exponential function $g(x) = \alpha \exp(\beta(x) - 1)$, which makes f continuous at $s = S$ and takes advantage of a large derivative to force $s \rightarrow S$. The values $\alpha = \beta = 1$ worked well. One disadvantage of the exponential function is that f is not smooth at $s = S$. Therefore, the optimization at $\sqrt{s} = 13$ TeV used the alternative function $g(x) = \gamma x^2$, which lead to better convergence because $\partial_x g(x)|_{x=0} = 0$. Values of $\gamma \sim 5 - 10$ resulted in

¹For a p-value, this means decreasing and for a significance, this means increasing. At $\sqrt{s} = 8$ TeV, the metric is the CL_s value (see Sec. 7.1) and at $\sqrt{s} = 13$ TeV, it is the ‘discovery significance’ described in the caption of Table 4.3. With no evidence for SUSY, the final result is the exclusion limit based on the CL_s in all cases.

relatively stable performance. Figure 4.3 illustrates the above procedure.




Figure 4.3: An example scan in the number of signal events. For each fixed S , the background yield is minimized and the CL_s value is computed (see Sec. 7.1 for details).

One additional component of the algorithm is the use of background weights. Background estimates that are at least partially data-driven are of higher quality than those constructed only from simulation. Therefore, it is desirable to have a background composition that is enriched in backgrounds that can be well-estimated and suppressed in backgrounds that cannot be predicted using some insight from the data. To achieve this in the minimization, the background yield $\mathbf{b}$ is constructed as a weighted sum over all of the background components: $\mathbf{b} = \sum \omega_i \mathbf{b}_i$, where $\mathbf{b}_i$ is the predicted background for process $i \in \{\bar{t}t, \bar{t}t + V, W + \text{jets}, \text{single top}, \text{dibosons}\}$ and ω_i is a fixed weight. By increasing the weight of one background relative to another, the algorithm can be steered toward suppressing a target background process. For the $\sqrt{s} = 8$ TeV analysis, $\omega_{\bar{t}t} = 1$, $\omega_{\bar{t}t+V} = 6$, $\omega_{W+\text{jets}} = 3$, $\omega_{\text{single top}} = 6$, and $\omega_{\text{dibosons}} = 6$. Even though there is a dedicated $W+\text{jets}$ CR for the $\sqrt{s} = 8$ TeV analysis, most of the $W+\text{jets}$ events in the SR are produced in association with heavy flavor jets, while the events in the control region are mostly light flavor; this is the reason for $\omega_{W+\text{jets}} = 3$. With data driven estimates for $\bar{t}t + V$ and single top as part of the $\sqrt{s} = 13$ TeV analysis, non-unity ω factors are no longer necessary².

²The diboson background is already negligible without imposing $\omega_{\text{dibosons}} > 1$.

4.3 Single Bin Regions

Six single bin regions are constructed for a broad coverage of sensitivity in the $(m_{\text{stop}}, m_{\text{LSP}})$ mass plane using the procedures described in the previous section. Table 4.1 provides an overview of the regions of parameter space targeted by each signal region. There are three phenomenologically distinct regions. At low stop mass, the signal cross section is relatively high ($\mathcal{O}(1\%)$ of $t\bar{t}$), but many kinematic distributions do not significantly differ from the dominate $t\bar{t}$ background (SR1). In contrast, at high stop mass, the cross section is very low, but many kinematic distributions are significantly different between the signal and the $t\bar{t}$ background (SR3, tNhigh, SR13). When additionally the LSP mass is large the difference between signal and background is reduced and requires a dedicated event selection to maintain sensitivity to this region of parameter space (SR2, tNmed). In order to increase the sensitivity for discovering SUSY with the early $\sqrt{s} = 13$ TeV data, a gluino mediated stop (GMS) model with a nearly degenerate stop and LSP was used for the optimization. As described in Sec. 1.5, such a model also results in $t\bar{t} + E_{\text{T}}^{\text{miss}}$ when the stop decay products are too soft to reconstruct. The gluino model $(m_{\tilde{g}}, m_{\tilde{t}}, m_{\text{LSP}}) = (1250, 750, 745)$ GeV was chosen to be kinematically equivalent (see Sec. 7.4.2.1.1 for details) to a model with stop pair production at $(m_{\tilde{t}}, m_{\text{LSP}}) = (800, 0)$ GeV. The only significant difference between the GMS and the direct stop models is that for a fixed mass, the former has a cross section that is about a factor of 50 more than the latter due to the additional spin and color states for the gluino.

Table 4.2 shows the defining selections for each region. A complete description of each variable can be found in Chapter 3.2. Horizontal lines in Table 4.2 group variables with a similar purpose. Even though the event selections were constructed in a mostly automated fashion, it is useful to examine their anatomy to understand why each value was chosen. *A well-motivated event selection is a robust event selection.* The logic for the various event selections is similar amongst the six event selections; for brevity SR13 is used as an example. Figure 4.4 shows the $E_{\text{T}}^{\text{miss}}$ distribution using the $\sqrt{s} = 13$ TeV preselection and after the full SR13 event selection before the $E_{\text{T}}^{\text{miss}}$ requirement. The LSPs can carry significant momentum and therefore the $E_{\text{T}}^{\text{miss}}$ is

one of the most powerful variables in any search for R-parity conserving SUSY. The pair of gluinos (or equivalently, a pair of 800 GeV stops) for the model in Fig. 4.4 will be produced nearly at rest in the lab frame and therefore $E_T^{\text{miss}} \lesssim 2 \times m_{\text{stop}}/2$. The factor of two is from the two LSPs and the factor of 1/2 is from the split in the stop mass between the top quark and the LSP boost. As indicated by Fig. 4.4, most of the stop events have significantly less E_T^{miss} than this bound because the orientation of the two LSPs relative to each other is random and the neutrino from the top quark decay can further reduce the total missing momentum when it has a large momentum component anti-parallel to the LSP directions. The optimal threshold value for SR13 is $E_T^{\text{miss}} \gtrsim 350$ GeV. Beyond that value, the reduction in signal outweighs the reduction in background. Note that the peak of the E_T^{miss} distribution is higher after the event selection than it is with only the preselection. This is due in part to the $H_{T,\text{sig}}^{\text{miss}}$ requirement.

Region	L_{int} [fb $^{-1}$]	$\sqrt{s}$ [GeV]	Stop Mass [GeV]	LSP Mass [GeV]
SR1	13	8	250	50
SR2	13	8	500	200
SR3	13	8	650	50
tNmed	20.3	8	550	200
tNhigh	20.3	8	650	1
SR13	3.2	13	800	1

Table 4.1: Benchmark signal models used for optimizing the single bin regions. The regions SR1-3 were optimized using the entire grid and so the chosen models are representative of the regions of parameter space that the three regions target. SR13 was optimized using a GMS model, but the equivalent stop model is the one in the table above.

A similar set of plots for the m_T are in Fig. 4.5. The peak of the m_T distribution for the signal is lower than the E_T^{miss} because the m_T is essentially the geometric average of the E_T^{miss} and the lepton p_T . As it is further down the decay chain than the neutralinos, the lepton p_T is expected to be softer and therefore brings down the geometric average.

Despite this, the $\mathbf{m}_T$ is still one of the most powerful variables, with a separation³ of about 14% in the signal region – the same as E_T^{mis} (at preselection, the separation is 35% for $\mathbf{m}_T$ and about 55% for E_T^{miss}). While all of the SM backgrounds are suppressed at high values of $\mathbf{m}_T$, W +jets events are reduced the most because they have no additional source of missing momentum to surpass the $\mathbf{m}_T \lesssim \mathbf{m}_W$ edge.

Variable	SR1	SR2	SR3	tNmed	tNhigh	SR13
Jet $\mathbf{p}_T > [\text{GeV}]$	80, 60, 40, 25			80, 60, 40, 25	100, 80, 40, 25	100, 80, 50, 25
$ \Delta\phi(\text{jet}_1, \vec{\mathbf{p}}_T^{\text{miss}}) >$	0.8	–	0.8	–	–	0.4
$ \Delta\phi(\text{jet}_2, \vec{\mathbf{p}}_T^{\text{miss}}) >$	0.8	0.8	0.8	0.8	–	0.4
$E_T^{\text{miss}} > [\text{GeV}]$	150	200	225	300	320	350
$H_{T,\text{sig}}^{\text{miss}} >$	–			12.5	12.5	20
$\mathbf{m}_T > [\text{GeV}]$	140 (*)	140	180	140	200	200
$\mathbf{am}_{T2} > [\text{GeV}]$	–	170	200	170	170	175
$\mathbf{m}_{T2}^\tau > [\text{GeV}]$	–	–	120	–	120	80
$\mathbf{m}_{\text{top},\chi^2} \in [\text{GeV}]$	[130,205]			[130,195]	[130,250]	[140, ∞]
$\Delta R(\mathbf{b}, \mathbf{l}) <$	–			–	3	2.5
$\Delta\phi(\vec{\mathbf{p}}_T^{\text{miss}}, 2^{\text{nd}} \text{ Large R jet}) >$	–			–	–	1

Table 4.2: A summary of the six single bin signal region event selections. Dashed lines indicate that there is no requirement on the given variable. Even though the same symbol might be used for all six regions, some of the variables have a different meaning across columns. For example, the jets used in the first five columns have the LCW calibration while the jets in the last column are at the EM scale but with the GS calibration applied. Furthermore, the choice of the τ candidate differs between the first five columns and the last column. The $\mathbf{m}_{\text{top},\chi^2}$ variable is an explicit tri-jet mass reconstruction in the first five columns (using jet resolution information in columns 4 and 5) and the large radius jet mass in the last column (see Sec. 3.2.4). (*) There is also an upper threshold of 250 GeV for this loose selection.

The $\mathbf{am}_{T2}$ distribution is shown in Fig. 4.6. At preselection, $\mathbf{am}_{T2}$ has a similar separation between signal and background as $\mathbf{m}_T$ ($\sim 35\%$), but after the rest of the event selection, it offers the most discriminating power with a separation of about $\sim 20\%$. In all signal regions that use $\mathbf{am}_{T2}$, the threshold value is around $\mathbf{m}_{\text{top}} \sim 175$ GeV. The large drop in the $t\bar{t}$ distribution at this point is visible in all of the plots

³Using the same heuristic metric as introduced in Sec. 3.2.1.2.5.





Figure 4.4: The distribution of the E_T^{miss} for the event preselection (top) and after all SR13 requirements except E_T^{miss} (bottom). Figures with all signal region requirement but the one displayed are called N-1 plots. Both the total background and signal yields are normalized to unity in the right plots. An arrow indicates the signal region requirement. The first number in parenthesis is the expected yield without applying any normalization factors and the second number is the raw event count in simulation (an indication of the statistical uncertainty). The Wt component includes all single top processes, but is dominated by the single production of a top quark in association with a W boson.

Figure 4.5: The same as Fig. 4.4, but for m_T instead of E_T^{miss} .

in Fig. 4.6. In contrast, single top events tend to be *above* m_{top} . This will motivate a data-driven technique to estimate this background in Sec. 5.3.1.




Figure 4.6: The same as Fig. 4.4, but for am_{T2} instead of E_T^{miss} .

A technique proposed in the literature to further suppress $t\bar{t}$ is a χ^2 compatibility test with the two-lepton $t\bar{t}$ background hypothesis [594]. For a missing leptonically decaying W boson W_2 and one neutrino from the identified leptonically decaying W boson ν_1 , there are 8 total unknowns: $p_{W_2}^\mu$ and $p_{\nu_1}^\mu$. Imposing $p_{W_2}^2 = m_W^2$, $p_{\nu_1}^2 = 0$, $p_{W_2x} + p_{\nu_1x} = p_x^{\text{miss}}$ and $p_{W_2y} + p_{\nu_1y} = p_y^{\text{miss}}$ reduces this to only four unknowns. One can choose these unknowns to be p_{W_2x} , p_{W_2y} , p_{W_2z} , and p_{ν_1z} . A χ^2 variable S is then the sum $(p_{W_1}^2 - m_W^2)^2/a_W^4 + \sum_{i=1}^2 (p_{t_i}^2 - m_t^2)^2/a_t^4 + (4m_t^2 - (\sum_i p_i)^2)/a_{\text{SM}}$, where

$\mathbf{p}_{t_i}$ is the sum of $\mathbf{p}_{W_i}$ and the four-vector of a $\mathbf{b}$ -tagged jet and the $\mathbf{a}_x$ are resolution parameters (see Ref. [594]). The *topness* variable is then given by $\log(\min S)$. The left plot of Fig. 4.7 shows a double-peak structure that separates the dilepton background with low values of S with the signal that has higher values of S . However, after the $\mathbf{am}_{T2}$ requirement (and other selections), shown in the right plot of Fig. 4.7, there is little additional discriminating power from topness.





Figure 4.7: The distribution of the $\mathbf{m}_T$ for the event preselection (top) and after all SR13 requirements except $\mathbf{m}_T$ (bottom). Both the total background and signal yields are normalized to unity in the right plots. An arrow indicates the signal region requirement.

The $H_{T,\text{sig}}^{\text{miss}}$ has a significant correlation with the E_T^{miss} , but Fig. 4.8 shows that it is still has significant separation power after preselection and in the signal region. Most events with the preselection have $H_{T,\text{sig}}^{\text{miss}} > 0$ because $H_T^{\text{miss}} \sim E_T^{\text{miss}}$ and $H_{T,\text{sig}}^{\text{miss}} > 0$ implies $H_T^{\text{miss}} > 100$ GeV. For the same reason, the peak of the $H_{T,\text{sig}}^{\text{miss}}$ distribution in the background shifts toward higher values in the SR due to the higher E_T^{miss} requirement.

Transverse mass variables are designed for cases where particles are lost or misidentified; in contrast the next two variables target explicit top reconstruction. The large-radius jet mass is shown in Fig. 4.10. As discussed in Chapter 4, when the top quark has sufficient boost in the lab frame, its decay products can be captured

Figure 4.8: The same as Fig. 4.4, but for $H_{T,\text{sig}}^{\text{miss}}$ instead of E_T^{miss} .

by a single large-radius jet. For $m_{\text{stop}} \sim 800$ GeV and a massless neutralino, $p_T^{\text{top}} \sim m_{\text{stop}}/2 \sim 400$ GeV. Re-clustered jets are used for this purpose, which allow for the large radius to be optimized per selection. Several radii were studied and the optimal value was found to be $R = 1.2$, which is consistent with the naive expectation that $R \sim 2m_t/p_T$. All signal jets with $p_T > 25$ GeV enter the re-clustering procedure and those with $p_T < 5\% \times p_T^{\text{large-radius}}$ are trimmed away. Re-clustering also allows for testing the inclusion (or not) of leptons in the re-clustering procedure. It was found that the sensitivity is higher when leptons are explicitly excluded from the re-clustering: both signal and background yields increase, but background increases more than the signal. In particular, the signal yield for the SR13 benchmark increases by about 20%, but the background increases by about 30%, with the biggest increase from $t\bar{t}$. Note that this exclusion of leptons is effectively an overlap removal procedure between large-radius jets and leptons. This overlap is trivial for re-clustered jets with a moderate boost as in this search; for large-radius jets clustered directly from calorimeter-cell clusters, the overlap is non-trivial due to significant energy deposits by electrons in the calorimeter. The SUSY signal shows a clear top quark mass peak in the left plot of Fig. 4.10 while the mostly dileptonic $t\bar{t}$ background has no resonant mass peak. In the signal region, the separation is reduced due to the correlation with other variables, but the top quark mass peak for the signal and $t\bar{t} + V$ are still separated from the dileptonic $t\bar{t}$, which is concentrated at lower values of the jet mass.

The leading large-radius jet is generally back-to-back in ϕ with the $\vec{p}_T^{\text{miss}}$. However, the subleading large-radius jet (if it exists) tends to be aligned with the $\vec{p}_T^{\text{miss}}$ in dileptonic $t\bar{t}$ events and back-to-back in signal events. This is because the E_T^{miss} from the neutrinos in $t\bar{t}$ events are generally close to at least the sub-leading large-radius jet formed in part by the lost or mis-identified second lepton. In contrast, in signal events, both the hadronic and leptonic top quark candidates are recoiling from the $\vec{p}_T^{\text{miss}}$ from the neutralinos. Figure 4.11 illustrates these properties of both the leading and sub-leading large-radius jets. About 25% of background events have a second signal large-radius jet ($p_T > 150$ GeV, $m_{\text{jet}} > 50$ GeV, and $|\eta| < 2.5$) while only about 10% of signal events have such a jet.

In addition to tagging the hadronically decaying top quark, it is possible to identify




Figure 4.9: The distribution of the large-radius $R = 1.2$ jet mass for the event preselection (top) and after all SR13 requirements except the jet mass (bottom). Both the total background and signal yields are normalized to unity in the right plots. An arrow indicates the signal region requirement.




Figure 4.10: The distribution of the large-radius $R = 1.2$ jet p_T after all SR13 event selections except the large-radius jet p_T . For large-radius jets with $p_T > 500 \text{ GeV}$, the trimming actively removes low p_T signal jets constituents.




Figure 4.11: Top: the $\Delta\phi$ between the leading large radius jet and the $\vec{p}_T^{\text{miss}}$. Bottom: the same as the top, but for the subleading large radius jet if one exists. Large radius jets are required to have $p_T > 150 \text{ GeV}$, $m > 50 \text{ GeV}$, and $|\eta| < 2.5$.

the boosted leptonically decaying top quark from the proximity of the lepton with a b -jet due to the same $2m/p_T$ scaling from above. Figure 4.12 shows the distribution of $\Delta R(b, \ell)$ using the highest p_T b -jet. This distance tends to be $\lesssim 1$ for the signal, but there is a heavy tail due to combinatorics. The $t\bar{t}$ background also has leptonically decaying boosted top quarks, but with less boost than for the signal. The mass of the b -jet and lepton pair also contains information about the top quark mass, but $m_{b\ell}$ can be naturally large for the background: for the correct pairing it has the same distribution as the signal ($m_{b\ell}$ is a Lorentz invariant) and for the incorrect $b\ell$ pairing, it is naturally large due to the significant distance between the b and the ℓ .




Figure 4.12: Same as Fig. 4.10, but with the ΔR between the highest p_T b -jet and the selected lepton instead of the jet mass.

The distribution of m_{T2}^τ is shown in Fig. 4.13. As expected, the background distribution has an endpoint near m_W while the signal is shifted toward much higher values. There are some geometric orientations (the *unbalanced configuration* - see Sec. 3.2.1.2.2) such that m_{T2}^τ is exactly zero. This results in a finite loss in acceptance for any positive threshold on m_{T2}^τ and is the source of the non-negligible signal yield in the first bin of the histograms in Fig. 4.13. Even though the $W+\text{jets}$ and $t\bar{t}+V$ processes are a significant contribution to the SR, they are largely absent from Fig. 4.13 because they do not usually have a second lepton, whereas $t\bar{t}$, Wt , and VV

events can use a hadronically decaying τ to exceed the m_T threshold. The number of $t\bar{t}$, Wt , and VV events that are removed by the m_{T2}^τ requirement (i.e. present in Fig. 4.13), $\sim (0.4, 0.2, 0.1)$ are comparable to the total yield of these backgrounds with the full signal region selection, $\sim (0.3, 0.1, 0.1)$ (see e.g. Fig. 4.10).




Figure 4.13: The normalized distribution of the m_{T2}^τ for the event preselection (left) and after all SR13 requirements except m_T (right). An arrow indicates the signal region requirement in the right plot.

Figure 4.15 shows the b -jet multiplicity at preselection and in the signal region. The b -tagging algorithm is configured to be 77% efficient for inclusive $t\bar{t}$ events. Table 4.3 shows the yields in SR13 for a variety of b -tagging working points, ranging from 60% to 85% efficiency. There is not a strong dependence on the working point for efficiencies below 77%, but the 85% efficiency is clearly worse than the others as proportionally more background events pass the ≥ 1 b -tagged jet requirement. This is largely due to the significant increase in the W +jets b -tagged jet multiplicity for tN high, which is similar to SR13. Between the 70% and 80% working points, there is a large increase in the charm-jet contribution to the one b -tagged jet bin.

Signal events in Fig. 4.15 are significantly more likely to have a second reconstructed b -jet compared to background events. However, requiring at least two b -tagged jets would remove too many signal events for such a threshold to be useful. In the future, one could exploit the asymmetry in the $n_{b\text{-jets}}$ distribution by using event weights or dividing up the single bin SR into (at least two) bins. Figure 4.16 illustrates the improvement one might expect from using event-weights. The optimal weights are signal-model dependent, but generic nearly optimal weights can be derived for a broad class of models.

Process	60%	70%	77%	85%
$t\bar{t}$ 1L	0.0	0.0	0.0	0.0
$t\bar{t}$ 1L1 τ	0.1	0.1	0.2	0.2
$t\bar{t}$ 2L	0.1	0.1	0.1	0.1
$t\bar{t}$ total	0.2	0.2	0.3	0.3
Single Top	0.1	0.1	0.1	0.1
W +jets	0.1	0.1	0.2	0.5
Dibosons	0.1	0.1	0.1	0.2
$t\bar{t} + V$	0.3	0.4	0.4	0.4
Total SM	0.8	1.0	1.1	1.6
Gtc (1250, 750)	5.0	5.1	5.4	5.7
Discovery σ (30% syst.)	3.1	3.0	3.0	2.7

Table 4.3: Yields for the SR13 defined in Table 4.2 but with various b -tagging working points. The last row is an approximate significance using the RooStat [595] `NumberCountingUtils` routine `BinomialExpZ`. A p -value is computed with a likelihood given by the product of a Poisson term for the statistical uncertainty and a Poisson term for the systematic uncertainty treated as a statistical uncertainty from an auxiliary measurement, i.e. a Poisson with mean $\tau = 1/(30\%)^2$ so that the fractional uncertainty of the auxiliary measurement is $1/\sqrt{\tau} = 30\%$. The conversion from p -value to σ is given by $\sigma = \Phi^{-1}(1 - p)$, for Φ the Gaussian cumulative distribution function.




Figure 4.14: The b -tagged jet multiplicity for W +jets events in tN_{high} broken down by jet flavor. The defining efficiency of the b -tagging working point is evaluated in inclusive $t\bar{t}$ events.




Figure 4.15: The distribution of the number of b -tagged jets at preselection (left) and in SR13 (right).




Figure 4.16: The median significance $= \Phi^{-1}(1 - \text{CL}_s)$ (see Sec. 7.1) as a function of the number of background and signal events, assuming zero systematic uncertainty. The numbers show the gain in significance when weighting each event by 2 if there are ≥ 2 **b**-jets. With this simple setup, 2 is nearly optimal, though it depends on **s** and **b** (see Sec. 2.2.2). The relative frequencies of **b**-tagged jets is taken from Fig. 4.15.

4.4 Compressed Signal Region

Adding multiple bins to a signal region significantly increases the complexity of the optimization procedure. To simplify the approach, the `tN_diag` signal region uses SR1 as a base. Two variables are loosened from their requirements in SR1 to define the bins of `tN_diag`; the tightest bin roughly corresponds to the SR1 event selection. Figure 4.17 shows the distribution of various kinematic variables with a one-lepton, four jets at $p_T > 80, 60, 40, 25$ GeV and $E_T^{\text{miss}} > 100$ GeV preselection for $t\bar{t}$ and $\tilde{t}_1\tilde{t}_1$ with $(m_{\text{stop}}, m_{\text{LSP}}) = (250, 50)$. It is clear that the m_T is the most discriminating variable given this preselection and is therefore chosen to define the shape fit. The E_T^{miss} is also a useful discriminating variable, which is partially hidden from Fig. 4.17 due to the $E_T^{\text{miss}} > 100$ GeV requirement. Figure 4.19 shows the E_T^{miss} , αm_{T2} and $m_{\text{had}}^{\text{top}}$ after an upper requirement on m_T . The signal and background distributions are nearly the same, but the E_T^{miss} shows the most difference and is therefore used as a second defining variable of the multibin signal region. Even though the likelihood ratio may not significantly change as a function of E_T^{miss} , the modeling, in particular for the trigger, may depend on E_T^{miss} . Therefore, the background normalization parameters and key systematic uncertainty nuisance parameters are assigned per E_T^{miss} bin.

The binning of the shape fit signal region is shown in Table 4.4. There are four bins of m_T and three bins of E_T^{miss} . Control regions are integrated into the signal region at low m_T and with an inverted b -tag requirement, as indicated by the last column of Table 4.4. The upper boundaries for the tightest m_T and E_T^{miss} bins were optimized using a coarse scan in $E_T^{\text{miss}} \in (150 \text{ GeV}, 175 \text{ GeV})$ and $m_T \in (140 \text{ GeV}, 160 \text{ GeV})$. Additionally, the E_T^{miss} significance requirement is scanned in the range $(3, 5, 8) \text{ GeV}^{1/2}$. For each combination of thresholds, pseudo-data from the simulation are fit using the procedure described in Sec. 7.1. This fit incorporates the most important experimental systematic uncertainties, including the jet energy scale and jet energy resolution uncertainties with one nuisance parameter each per E_T^{miss} bin. To demonstrate the power of the shape fit, the full multibin approach is compared with a single-bin region using only three bins from the $E_T^{\text{miss}} > 150$ GeV column: one for a SR, one for a $t\bar{t}$ CR, and one for a W +jets CR. This region should have comparable sensitivity to




Figure 4.17: The distribution of E_T^{miss} , m_T , $a m_{T2}$, and $m_{\text{had}}^{\text{top}}$ with a one-lepton, four jets at $p_T > 80, 60, 40, 25$ GeV and $E_T^{\text{miss}} > 100$ GeV preselection.

SR1. Table 4.5 shows the results of the scan. The most striking feature of Table 4.5 is that the multibin setup improves upon the single-bin regions by nearly an order of magnitude for all parameter values. The single bin regions are not able to exclude the benchmark model chosen for the optimization, but it is within reach of the multibin SR. Due to the per E_T^{miss} bin normalization/nuisance parameters, the level of profiling




Figure 4.18: The distribution of E_T^{miss} , $a m_{T2}$, and $m_{\text{had}}^{\text{top}}$ after the preselection from Fig. 4.17 with the additional requirement $m_T < 50$ GeV.

is minimal⁴. Therefore, the gain in sensitivity is mostly from the additional bins with various signal-to-background ratios, as desired. A high E_T^{miss} threshold of $8 \text{ GeV}^{1/2}$ is worse than the lower values in the scan, though there is not much difference between 3 and $5 \text{ GeV}^{1/2}$. Therefore, the SR1 requirement of $8 \text{ GeV}^{1/2}$ is loosened to $5 \text{ GeV}^{1/2}$ for tN_{diag} . More generally, the fit seems to perform best when the bin with the tightest selection has a relatively high signal yield (i.e. is relatively loose). For this reason, the $m_T = 140 \text{ GeV}$ and $E_T^{\text{miss}} = 150 \text{ GeV}$ thresholds are chosen for the tightest bin. The yields for this signal region are shown in Fig. 4.19⁵. For $m_T < 120 \text{ GeV}$, the bins have $\mathcal{O}(1000)$ events and for $m_T > 120 \text{ GeV}$, the bins have $\mathcal{O}(100)$ events. The signal-to-background ratio ranges between 10-20% in the tightest bins of the SR.

The tN_{shape} region described above was released as a preliminary result in the summer of 2013 [413]. The additional event selection beyond the m_T and E_T^{miss} requirements was optimized slightly more for the final result published in Ref. [414] and shown in Fig. 2.18 from Sec. 2.2.2. In particular, the tight τ veto from Sec. 3.1 effectively removes about 10% of the $t\bar{t}$ background with only a negligible impact on

⁴In the usual CR method approach, the fit is setup so that each CR has significantly more events than the SR so that in the full fit, the background normalization is simply set by the ratio of data to MC in the CR. However, the various bins of the shape fit region have a significant background yield and so many bins can play a role in normalizing the background. There are 7 total free parameters in the fit (one $t\bar{t}$ and one $W+\text{jets}$ parameter per E_T^{miss} bin and one overall signal normalization and 15 total bins. Breaking up the normalization parameters into E_T^{miss} bins significantly reduces the amount of profiling from an over-constrained fit. See Sec. 7.2 for more details.

⁵The background yields are *post-fit* - see Sec. 7.1 for details of the fit.

m_T window [GeV]	E_T^{miss} window [GeV]	b-tags	Comment
$60 < m_T < 90$	$100 < E_T^{\text{miss}} < 125$	$= 0$	W +jets enriched
$60 < m_T < 90$	$100 < E_T^{\text{miss}} < 125$	> 1	$t\bar{t}$ enriched
$90 < m_T < 120$	$100 < E_T^{\text{miss}} < 125$	> 1	
$120 < m_T < 140$	$100 < E_T^{\text{miss}} < 125$	> 1	
$m_T > 140$	$100 < E_T^{\text{miss}} < 125$	> 1	
$60 < m_T < 90$	$125 < E_T^{\text{miss}} < 150$	$= 0$	W +jets enriched
$60 < m_T < 90$	$125 < E_T^{\text{miss}} < 150$	> 1	$t\bar{t}$ enriched
$90 < m_T < 120$	$125 < E_T^{\text{miss}} < 150$	> 1	
$120 < m_T < 140$	$125 < E_T^{\text{miss}} < 150$	> 1	
$m_T > 140$	$125 < E_T^{\text{miss}} < 150$	> 1	
$60 < m_T < 90$	$E_T^{\text{miss}} > 150$	$= 0$	W +jets enriched
$60 < m_T < 90$	$E_T^{\text{miss}} > 150$	> 1	$t\bar{t}$ enriched
$90 < m_T < 120$	$E_T^{\text{miss}} > 150$	> 1	
$120 < m_T < 140$	$E_T^{\text{miss}} > 150$	> 1	
$m_T > 140$	$E_T^{\text{miss}} > 150$	> 1	

Table 4.4: The definition of tN_diag. In addition to the variables shown above, the signal region is defined by a four-jet requirement with $p_T > 80, 60, 40, 25$ GeV, $m_{\text{top}}^{\text{had}} \in [130, 205]$ GeV, $E_T^{\text{miss}}/\sqrt{H_T} > 5$ GeV^{1/2}, and $\Delta\phi(\text{jet}_i, \vec{p}_T^{\text{miss}}) > 0.8$ for $i = 1$ and 2.

the signal. Furthermore, the jet p_T requirements are loosened to $p_T > 60, 60, 40, 25$ GeV. The thresholds above have the same values as SR2 and SR3, whose benchmark models have much harder p_T spectra than the tN_diag benchmark model. One other small change is the addition of a $\Delta R(b, \ell) < 2.5$ requirement. The top quarks from the tN_diag benchmark model are not so boosted such that one of the b-tagged jets is always near the lepton, but this requirement is useful for suppressing dilepton $t\bar{t}$ where the two top quarks are back-to-back. These and other modifications were studied using a similar setup to the one described above and also included approximate theoretical modeling systematic uncertainties for the $t\bar{t}$ and W +jets

$E_T^{\text{miss}}/\sqrt{H_T}$ [GeV ^{1/2}]	m_T [GeV]	E_T^{miss} [GeV]	CL _s multibin	CL _s single bin	Notes
5	140	150	0.0151	0.146	0.0012*
5	140	175	0.0179	0.301	
5	160	150	0.0156	0.221	
5	160	175	0.0177	0.451	
3	140	150	0.0152	0.145	
3	140	175	0.0176	0.301	
3	160	150	0.0162	0.224	
3	160	175	0.0168	0.451	
8	140	150	0.036	0.149	Fit Failed
8	140	175	NaN	0.301	
8	160	150	0.0218	0.26	
8	160	175	NaN	0.454	

Table 4.5: CL_s values (see Sec. 7.1) computed with the shape-fit (multibin) and single bin setups described in the text. These values are approximately **p**-values for a hypothesis test (smaller values are better). The E_T^{miss} and m_T values are the thresholds for the tightest region of the shape fit and define the single bin region. The * denotes the CL_s value for the shape fit without any systematic uncertainties. As expected, systematic uncertainties have a big impact on the sensitivity. In two cases (marked ‘Fit Failed’) the multibin fit did not converge.

processes. Between the two selections, the signal-to-background ratio increased from about 15% in the tightest bin of the shape fit to about 20% (for the same integrated luminosity). Table 4.6 summarizes the final tN_{diag} event selection.




Figure 4.19: The yields for the preliminary version of tN_diag. An analogous table for the final version appears in Fig. 2.18. All background yields are after the post-fit - see Sec. 7.1 for details.

Variable	tN_diag (preliminary)	tN_diag (final)	Comment
Jet $p_T > [\text{GeV}]$	80, 60, 40, 25	60, 60, 40, 25	
$\Delta\phi(\text{jet}_i, \vec{p}_T^{\text{miss}}) >$	0.8	0.8	$i = 1, 2$
$E_T^{\text{miss}}/\sqrt{H_T} > \text{GeV}^{1/2}$	5	5	
$E_T^{\text{miss}} > [\text{GeV}]$	100	100	3 bins of the shape fit
$m_T > [\text{GeV}]$	60	60	4 bins of the shape fit
$m_{\text{had}}^{\text{top}} \in [130, 205] \text{ GeV}$			
$\Delta R(b, \ell)$	–	< 2.5	
τ -veto	–	tight	

Table 4.6: A summary of the multibin shape fit region, tN_diag. This region was released with a preliminary selection in Ref. [413] with a small change for the final result in Ref. [414]. Dashed lines indicate that there is no requirement on the given variable.

Chapter 5

Background Estimation

The strategy for estimating the SM background in each signal region is to use the control region method to predict the yield for each distinct subprocess. Control regions are constructed *for each signal region* to be as close as possible to the signal region phase space while maintaining a high yield and purity of the target background process. Processes which cannot be normalized using data-driven techniques are subdominant and are estimated using simulation. Figure 5.1 shows the background composition in each of the signal regions described in Chapter 4.

Due to its large cross section and phenomenological similarity to the signal, top quark pair production is one of the most important background processes in all signal regions (see Fig. 5.1). As described in Sec. 3.2.1.1, m_T is a powerful tool for suppressing single lepton $t\bar{t}$ and therefore the dominant $t\bar{t}$ background in the signal regions has two real leptons, where one is not identified as a lepton or is a hadronically decaying τ . Section 5.1 documents the construction of $t\bar{t}$ control regions. As a result of the many tools available for reducing the yield of dilepton $t\bar{t}$ events, the remaining background composition in the signal regions is diverse. In addition to $t\bar{t}$, the production of W bosons in association with many jets (Sec. 5.2), the single production of top quarks (Sec. 5.3), and the associated production of top quark pairs with a Z boson (Sec. 5.4) are also significant contributions to the SM background. As shown in Fig. 5.1 these backgrounds are increasingly relevant for higher target stop masses. All of the background processes mentioned thus far are integrated into the control region




Figure 5.1: Left: The yields in each signal region broken down by background type after the CR-only fit setup described in Sec. 7.1. Right: yields are normalized to unity to show the fractional background composition. The numbers in the right plot show the fraction of non- $t\bar{t}$ background. tNdiagxy is the $(x+1)^{\text{th}}$ E_T^{miss} bin and $(y+2)^{\text{th}}$ m_T bin of the shape fit region.

method for a data-driven estimate of the normalization; for single top and $t\bar{t} + Z$, this is the first time data-driven techniques are used in the $t\bar{t} + E_T^{\text{miss}}$ topology. The next most important background is the production of multiple electroweak bosons in association with jets. Section 5.5 describes the modeling of this background, which uses a simulation-based approach.

Other SM processes are negligibly small, due to a small cross-section or a low acceptance. The latter category includes the single production of Z bosons in association with jets and QCD multijets. Both of these processes require significant fake E_T^{miss} , and fake leptons, or both. The exception is $Z(\rightarrow \tau_{\text{had}} \tau_{\text{lep}}) + \text{jets}$, which has the same final state as $W(\rightarrow l\nu) + \text{jets}$. However, the former is suppressed by a factor of at least $\mathcal{BR}(Z \rightarrow \tau\tau) \times \mathcal{BR}(\tau_{\text{lep}}) \times \mathcal{BR}(\tau_{\text{had}}) \times \sigma_{Z+\text{jets}} / (\alpha_s \times \sigma_{W+\text{jets}} \times \mathcal{BR}(W \rightarrow l\nu)) \sim 2\%$ with respect to $W \rightarrow l\nu$; the ratio $\sigma_{Z+\text{jets}} / \sigma_{W+\text{jets}} \sim 0.1$ and the factor of α_s is due to the fact that hadronically decaying τ would be reconstructed as a jet and thus reduce the number of needed quark and gluon jets. A further suppression results from the

required mis-measurement of m_T , which is naturally low for τ events relative to direct $W \rightarrow \mu/e + \nu$ events (see Sec. 3.2.1.1) and thus needs to smear further to pass the high m_T threshold. The other possible Z +jets decays and estimates for their suppression factors are summarized in Table 5.1. In all cases, since W +jets is already a small background, the Z +jets is negligible and is henceforth ignored. Generic QCD multi-jet production has a large cross-section compared with W +jets, but requires both a fake lepton and fake E_T^{miss} . The estimates for the regions described in sections 2.2.2 and 4.1.2.3.1 showed that this background is already subdominant for an inclusive one-lepton $t\bar{t}$ event selection and Ref. [277] shows that it is negligible at high E_T^{miss} . Therefore, multijets are ignored for the remainder of the chapter.

Z Decay Mode	Lepton	Additional Jets	E_T^{miss}	Fraction of $W \rightarrow l\nu$
$\nu\bar{\nu}$	fake	4	correct	$10\% \times \epsilon_f \times 20\%/30\% \lesssim 0.1\%$
$e^+e^-/\mu^+\mu^-$	one lost	4	fake	$10\% \times \rho_f \times \epsilon_l \times 6.6\%/30\% \ll 0.1\%$
$\tau_{\text{lep}}\tau_{\text{lep}}$	one lost	4	mis-measured	$10\% \times \rho_m \times \epsilon_l \times 3.3\% \times 35\%^2/30\% \ll 0.1\%$
$\tau_{\text{lep}}\tau_{\text{had}}$	correct	3	mis-measured	$10\% \times \rho_m \times 3.3\% \times 35\% \times 65\%/(30\% \times \alpha_s) \lesssim 2\%$
$\tau_{\text{had}}\tau_{\text{had}}$	fake	2	mis-measured	$10\% \times \rho_m \times \epsilon_f \times 3.3\% \times 65\%^2/(30\% \times \alpha_s^2) \lesssim 0.1\%$
$q\bar{q}$	fake	2	fake	$10\% \times \rho_f \times \epsilon_f \times 70\%/(30\% \times \alpha_s^2) \ll 0.1\%$

Table 5.1: Estimates for the yield of Z +jets events relative to the $W \rightarrow l\nu$ yield. The number of additional jets only contributes to the last column when it differs from four, which is the necessary number of extra jets already needed by W +jets events to pass the event selection. If a tau decays leptonically, it is considered to be reconstructed as a jet, reducing the number of extra jets required. The rate of fake or non-prompt leptons $\epsilon_f \lesssim 1\%$ [277]. The probability for events with no real E_T^{miss} to be mis-reconstructed as events with large E_T^{miss} , ρ_f , is negligible because the resolution scales as $\sqrt{\sum E_T} \text{ GeV}^{1/2} \sim 20 \text{ GeV}$. Events with tau decays naturally have E_T^{miss} , but this needs to be severally mis-measured (with rate ρ_m) to pass at least the m_T requirements.

The chapter ends with an overview in Sec. 5.6 with all of the control region definitions and signal region yields. In addition, approximate scale factors for the data-driven background estimates are calculated as a function of key discriminating variables. Uncertainties associated with the background estimates are presented in Chapter 6.

5.1 Top Quark Pair Production

Top quark pair production in the lepton+jets final state has the same signature at leading order as the targeted signal: one lepton, missing momentum (from the neutrino), four jets (two b -jets). All of the single-bin signal regions use a strict m_T requirement that effectively eliminates the single lepton background, replacing it with dilepton $t\bar{t}$ processes. This technique is used in reverse to estimate the $t\bar{t}$ background in the signal region: the MC is normalized in a low m_T window enriched in single lepton $t\bar{t}$ events. Events that pass all signal requirements except have low m_T are kinematically similar to the signal region events, but have a small predicted signal contamination and a high single lepton $t\bar{t}$ event yield and purity. The disadvantage of the low m_T method is the required extrapolation over lepton multiplicity from the control region to the signal region. The cross section and event kinematics of a $t\bar{t}$ event are determined by $m_{t\bar{t}}$ and $p_{T,t\bar{t}}$. Therefore, to reduce the theoretical systematic uncertainties on the extrapolation from one lepton events in the CR to two lepton events in the signal region, it is important to ensure that $m_{t\bar{t}}$ and $p_{T,t\bar{t}}$ are as similar as possible in the CR and SR. For a fixed $m_{t\bar{t}}$ and $p_{T,t\bar{t}}$, the leading jets in one lepton $t\bar{t}$ events should have approximately the same distribution as the corresponding jets in two lepton $t\bar{t}$ events. However, the third and fourth jets in one lepton $t\bar{t}$ events are already present at leading order (from W boson decays) unlike in dilepton $t\bar{t}$ events. Section 5.1.1 explores the modeling of these subleading jets using an explicit dilepton event selection. With the same fixed top quark kinematic properties, the E_T^{miss} will be softer in one lepton $t\bar{t}$ events because there is only one neutrino from W boson decays¹. This suggests a lower E_T^{miss} threshold is appropriate for the $t\bar{t}$ control region compared with the signal region. A lower E_T^{miss} requirement is also useful because it can improve the $t\bar{t}$ purity. Figure 5.2 shows how the $t\bar{t}$ purity of the tNhigh control

¹How much softer depends on the reconstruction of the second lepton in dilepton events. If the second lepton is out of acceptance, then the difference between single lepton and dilepton events is larger than for events where the second lepton is within acceptance, but not identified as a lepton. At high E_T^{miss} , the former is largely irrelevant by construction since a low p_T lepton will not contribute significantly to the E_T^{miss} . A high p_T lepton that is too far forward to reconstruct as such that is also not reconstructed as a jet could contribute significant additional E_T^{miss} , but this is suppressed because the $|\eta|$ distribution is falling from 0.

region depends on the E_T^{miss} threshold. A value of 70% was chosen for the tNHigh TCR in order to increase the total event yield and also $t\bar{t}$ event purity. Similar studies for all of the signal regions produced customized $t\bar{t}$ control regions that are summarized in Table 5.2. The $H_{T,\text{sig}}^{\text{miss}}$ is varied by a similar amount as the E_T^{miss} when relevant because both variables scale the same way with additional real missing momentum. Additional kinematic requirements are relaxed for the tighter signal regions where the total $t\bar{t}$ event yield is too low from simply using the low m_T window. Due to the Jacobian peak (see Sec. 3.2.1.1), most of the single lepton events have $m_T \sim m_W$ and so the lower bound of the m_T window is set greater than zero to reduce non $t\bar{t}$ backgrounds. The upper edge of the m_T window is chosen to allow a gap between the control region and signal region for validation purposes, described in Sec. 7.2 in more detail.




Figure 5.2: Optimization of the $t\bar{t}$ CR associated with tNhigh. The left plot shows the control region composition for various fractional E_T^{miss} thresholds after changing the m_T window and removing the m_T requirements. The signal region requirement on E_T^{miss} is 320 GeV. The $H_{T,\text{sig}}^{\text{miss}}$ threshold is varied coherently with the E_T^{miss} requirement.

The predicted signal yield in all of the $t\bar{t}$ control regions is negligible for regions of $(m_{\text{stop}}, m_{\text{LSP}})$ near the corresponding benchmark models. This is illustrating explicitly for SR13 in Fig. 5.3, where the benchmark model is $(m_{\text{stop}}, m_{\text{LSP}}) = (800, 0)$.

Requirement	SR1	TCR1	SR2	TCR2	SR3	TCR3	tNmed	TCRmed	tNhigh	TCRhigh	SR13	TCR13
m_T [GeV]	[140, 250]	[60, 90]	> 140	[60, 90]	> 180	[60, 90]	> 140	[60, 90]	> 200	[60, 90]	200	[30, 90]
E_T^{miss} [GeV]	–	–	–	–	> 225	> 220	–	–	> 320	> 225	> 350	> 250
am_{T2} [GeV]	–	–	> 170	> 120	> 200	> 170	> 170	> 120	> 170	> 80	> 175	[100, 200]
m_{T2}^τ [GeV]	–	–	–	–	> 120	> 0	–	–	> 120	> 0	–	–
$H_{T,\text{sig}}^{\text{miss}}$	–	–	–	–	–	–	–	–	> 12.5	> 8.8	> 20	> 15
$\Delta R(b, l)$	–	–	–	–	–	–	–	–	–	–	< 2.5	< ∞
Total Yield	125	1661	9.6	169	4.3	195	13.0	159	5.0	359	1.3	102
$t\bar{t}$ Purity	83%	82%	56%	66%	44%	57%	50%	79%	39%	80%	25%	88%

Table 5.2: The definition of the $t\bar{t}$ control regions for each signal region presented in Chapter 4. Only the requirements that differ from the corresponding signal region are indicated in the table, with a ‘–’ if there is no change between the signal and control region. The lower two rows show the total background yield and the fraction of $t\bar{t}$ events in both the signal and control region using the CR-only fit, described in Sec. 7.1. The upper am_{T2} requirement in TCR13 is to ensure orthogonality from STCR13, described in Sec. 5.3.1.

The predicted yield for $m_{\text{stop}} \sim 800$ GeV is less than 0.2 events whereas the entire SM prediction for this region is about 100 events (see Table 5.2).




Figure 5.3: The number of signal events in TCR13 as a function of the stop mass and neutralino mass. The benchmark model for SR2 has a stop mass of 800 GeV and a neutralino mass of 0 GeV. The number of fitted SM events in TCR13 is 102.

5.1.1 Modeling Dilepton Events

As the dilepton $t\bar{t}$ background is estimated using a mostly single lepton $t\bar{t}$ event selection at low m_T , it is critical to validate the modeling of the jet-related variables in the high m_T tail. The leading order matrix element for dilepton events only has two out-going quarks and so at least two extra jets must originate from somewhere else in the simulation. For a hadronic origin of the jets, the two possibilities are the real emission from the NLO matrix element and extra radiation at leading logarithm from the parton shower. Another possibility is that the extra jets are mis-identified leptons, which nearly always occurs for hadronically decaying taus with sufficient p_T . This section presents two event selections for probing the modeling of the extra jets in dilepton $t\bar{t}$ events by explicitly requiring a second lepton. The events in Sec. 5.1.1.1 are required to have a reconstructed tau candidate in addition to an electron or muon to form a $1L1\tau$ validation region (VR) and the events in Sec. 5.1.1.2 have an explicitly reconstructed electron-muon pair. Note that these validation regions could be used for a direct estimation of the $t\bar{t}$ background yield in the signal region via the control region method in place of the low m_T region. However, due to the much lower cross-section for dilepton $t\bar{t}$ events, it is likely that the selection would need to be significantly looser than the one-lepton region and thus a larger phase-space extrapolation is required.




Figure 5.4: Feynman diagrams illustrating dilepton $t\bar{t}$ events passing the four-jet selection. In the left diagram, one of the electron or muon is not identified or reconstructed as a jet. In the right diagram, the tau is reconstructed as a jet.

5.1.1.1 Tau Validation Region

About half of the dilepton background has one hadronically decaying τ ($1L1\tau$). Many of the signal regions have hadronic τ vetos designed to reject such events, but τ reconstruction is not as clean as electron or muon identification and therefore many hadronic τ events still pass the full event selection. Events that would have been rejected due to the hadronic τ veto can be used to study the modeling of jets beyond those produced from the tree-level $t\bar{t}$ system. A τ validation region is therefore constructed with the event selection shown in Table 5.3. After requiring one reconstructed hadronically decaying τ in addition to the preselection, the validation region is still dominated by one lepton $t\bar{t}$ events with a fake τ . The one lepton component is suppressed by requiring $m_T > 100$ GeV, as illustrated in the left plot of Fig. 5.5. In total, there are about 100 SM events predicted in the validation region with over 70% $1L1\tau$ purity.

Requirement	Value
Preselection	Exactly one lepton
Reconstructed τ	> 0
n_{jets}	≥ 4
1 st jet p_T [GeV]	> 80
2 nd jet p_T [GeV]	> 50
3 rd jet p_T [GeV]	> 40
4 th jet p_T [GeV]	> 25
m_T [GeV]	> 100
E_T^{miss} [GeV]	> 200
$n_{\text{b-jets}}$	> 0

Table 5.3: An event selection requiring a hadronic τ candidate to study the modeling of $1L1\tau$ events.

The jet multiplicity in the τ validation region is shown in Fig. 5.6, beginning at four jets as required by all signal region selections. Frequently, one of the four jets is




Figure 5.5: Left: the m_T distribution in the $1L1\tau$ validation with all requirements except the m_T threshold, which is indicated by an arrow. Right: the am_{T2} distribution in the $1L1\tau$ validation region. Jet energy scale and resolution uncertainties are included in the error band. The last bin contains overflow.

the hadronically decaying τ itself as there is no τ -jet overlap removal. For this reason, it is slightly ‘easier’ for a $1L1\tau$ event to pass the event selection compared with a dilepton $t\bar{t}$ event with only electrons or muons. Formally, the POWHEG-BOX+PYTHIA 6 simulation is NLO accurate to the fourth jet (assuming one of the four is the hadronically decaying τ) and only leading logarithmically accurate for $n_{\text{jets}} > 4$. However, the agreement is significantly better than naively expected because the parton shower has been extensively tuned to collider data. While the $\chi^2/\text{NDF} < 1$, there is a small slope in the data/MC ratio. Inclusive measurements of the ‘extra’ jets in $t\bar{t}$ events find a slope in the opposite direction [596], suggesting that the potential trend in Fig. 5.6 is possibly insignificant.

The subleading jets are examined in more detail in Fig. 5.7. There is no unambiguous way to select jets that are not produced from the leading order $t\bar{t}$ decay, but one useful proxy is to consider non b -tagged jets. Figure 5.7 shows the p_T distribution of the leading non b -tagged jets in the $1L1\tau$ validation region. The data/MC ratio does not provide any significant evidence for a mis-modeling of these p_T spectra.




Figure 5.6: The number of signal jets in the 1L1 τ VR. Jet energy scale and resolution uncertainties are included in the error band. The last bin contains overflow.




Figure 5.7: The leading (left) and sub-leading (right) non b -tagged jets in the 1L1 τ validation region. When there are more than two b -tagged jets in the event, the plotted jets are those after the leading two b -tagged jets, ordered by p_T . Jet energy scale and resolution uncertainties are included in the error band. The last bin contains overflow.

5.1.1.2 Dilepton Validation Region

The other half of the dilepton $t\bar{t}$ events that pass the event selection have a second electron or muon that is not reconstructed as such. To study the modeling of the jets in this case, a minimal two lepton event selection is constructed and is summarized in Table 5.4. Low p_T unscaled dilepton triggers are available. However, these triggers do not add a significant number of events for the region of interest $E_T^{\text{miss}} > 200$ GeV where the E_T^{miss} trigger is fully efficient. An opposite flavor $e\mu$ selection is chosen to suppress Z +jets events². The predicted yield in the resulting validation region is about 400 events with over $> 85\%$ $t\bar{t}$ 2L purity.

Requirement	Value
Trigger	E_T^{miss}
n_e	$= 1$
n_μ	$= 1$
$q_e \times q_\mu$	< 0
n_{jets}	≥ 4
E_T^{miss} [GeV]	> 200
$n_{b\text{-jets}}$	> 0

Table 5.4: An event selection requiring two reconstructed signal leptons. The variable q_l denotes the charge of lepton l .

There is no unique way to define the m_T and am_{T2} variables in two-lepton events, but a way to probe the case where the second lepton is reconstructed but mis-identified is shown in Fig. 5.8. The lepton with the higher p_T is treated as the signal lepton and the second lepton is ignored. Even though it is not explicitly identified, this second lepton still contributes to the E_T^{miss} calculation. The second neutrino allows events to exceed the $m_T = m_W$ and the population of the tail is determined by $p_T^{t\bar{t}}$ and $m_{t\bar{t}}$ as in Sec. 3.2.1.1. In contrast, the am_{T2} distribution is mostly contained³ within

²A selection with ee or $\mu\mu$ would be possible with an additional requirement on $m_{l\bar{l}}$ to be away from m_Z .

³The main topology motivating the am_{T2} variable is when the second lepton is lost (not part of

$\mathbf{am}_{T2} \lesssim m_{\text{top}} \sim 175$ GeV which is the reason this variable is powerful at suppressing the dilepton $t\bar{t}$ background.

Figure 5.9 shows the jet multiplicity in the 2L validation region and is the analogue to Fig. 5.6 from the 1L1 τ region. All other plots in this section require $n_{\text{jets}} \geq 4$, but the modeling of the third jet is already interesting because only two jets are expected from the ME in dilepton events. There is a small slope in the data/MC ratio for $n_{\text{jets}} > 4$, but it is well within the systematic uncertainty from the jet energy scale and resolution even though it is in the same direction as dedicated studies [596]. The modeling of the momentum of the ‘extra’ jets is shown in Fig. 5.10. As in the 1L1 τ case, there is no unique way to identify such jets, but the leading non- b tagged jets are a good proxy. There is no significant evidence for mis-modeling the p_T spectra.




Figure 5.8: Left: the m_T distribution in the 2L validation region where the softer lepton is treated as measured but not reconstructed. Right: the $\mathbf{am}_{T2}$ distribution with the same lepton treatment as the left plot. Jet energy scale and resolution uncertainties are included in the error band. The last bin contains overflow.

the E_T^{miss} , but Sec. 3.2.1.2 shows that it is also useful when the lepton is only mis-identified.




Figure 5.9: The jet multiplicity in the 2L VR. Jet energy scale and resolution uncertainties are included in the error band. The last bin contains overflow.





Figure 5.10: The leading (left) and sub-leading (right) non b-tagged jets in the 2L validation region. When there are more than two b-tagged jets in the event, the plotted jets are those after the leading two b-tagged jets, ordered by p_T . Jet energy scale and resolution uncertainties are included in the error band. The last bin contains overflow.

5.2 W +jets

The inclusive W +jets cross-section is about 500 times higher than the inclusive $t\bar{t}$ cross-section. Accounting for the production of extra jets (four for W +jets and 2 for dilepton $t\bar{t}$) and the leptonic branching ratios, this factor reduces to approximately $500 \times \alpha^4 \times 30\% / (\alpha^2 \times 10\%) \sim 15$. However, unlike for $t\bar{t}$ events, there is no possibility for a second lepton to allow W +jets events to naturally exceed a stringent m_T threshold. Therefore, the W +jets events that pass the signal region event selections must have significant mis-measurement of the $\vec{p}_T^{\text{miss}}$. Section 5.2.2 investigates the modeling of W +jets events in the m_T tail. Control regions for W +jets are constructed analogously to the $t\bar{t}$ ones in Sec. 5.1 with one additional modification. In order to suppress $t\bar{t}$ events in the W +jets control region, events are required to have exactly no b -tagged jets. The b -jet veto is a powerful tool for removing $t\bar{t}$ events while maintaining a high yield and purity of W +jets events, as demonstrated in Fig. 5.11. For a $\sim 70\%$ b -tag working point, one expects that $t\bar{t}$ events fall in the 2 b -tag bin $\sim 0.7^2 \sim 50\%$ of the time, in the 1 b -tag bin $\sim 2 \times 0.3 \times 0.7 \sim 40\%$ of the time, and in the zero b -tag bin $\sim 10\%$ of the time. However, one significant drawback of this method is that most of the W +jets events in the signal region are associated with heavy flavor jets (e.g. $W + b\bar{b}$ and $W + c$) while the ones in the control region are nearly all from light flavor jets. The kinematic properties of the flavor extrapolation are discussed in Sec. 5.2.1 and the associated systematic uncertainties are documented in Sec. 6.2.4.

All of the W +jets control regions are recorded in Table 5.5, analogously to Table 5.2 for the $t\bar{t}$ control regions. The regions are nearly identical to the $t\bar{t}$ control regions with only the b -jet multiplicity inverted. One exception is WCR13, which has no upper bound on αm_{T2} as this is not needed to enforce orthogonality with STCR13, as described in Sec. 5.3.1. The W +jets background is under 20% of the total SM in all signal regions except SR3 where it is predicted to be less than 30%. Conversely, the signal contamination in the W +jets control regions is negligible as the stops receive a similar suppression to $t\bar{t}$ and already have a small cross section.

Figure 5.11: The distribution of the number of b -jets with the inclusive preselection.

Requirement	SR1	WCR1	SR2	WCR2	SR3	WCR3	tNmed	WCRmed	tNhigh	WCRhigh	SR213	WCR213
m_T [GeV]	[140, 250]	[60, 90]	> 140	[60, 90]	> 180	[60, 90]	> 140	[60, 90]	> 200	[60, 90]	200	[30, 90]
$n_{b\text{-jets}}$	> 0	< 1	> 0	< 1	> 0	< 1	> 0	< 1	> 0	< 1	> 0	< 1
E_T^{miss} [GeV]	–	–	–	–	> 225	> 220	–	–	> 320	> 225	> 350	> 250
am_{T2} [GeV]	–	–	> 170	> 120	> 200	> 170	> 170	> 120	> 170	> 80	> 175	> 100
m_{T2}^{τ} [GeV]	–	–	–	–	> 120	> 0	–	–	> 120	> 0	–	–
$H_{T,\text{sig}}^{\text{miss}}$	–	–	–	–	–	–	–	–	> 12.5	> 8.8	> 20	> 15
$\Delta R(b, l)$	–	–	–	–	–	–	–	–	–	–	< 2.5	< ∞
Total Yield	125	897%	9.6	147	4.3	169	13.0	161	5.0	482	1.3	135
$t\bar{t}$ Purity	9%	62%	17%	76%	28%	79%	16%	68%	18%	69%	12%	71%

Table 5.5: The definition of the W control regions for each signal region presented in Chapter 4. Only the requirements that differ from the corresponding signal region are indicated in the table, with a ‘–’ if there is no change between the signal and control region. All changes highlighted in red are different from the definition of the corresponding $t\bar{t}$ control region. The lower two rows show the total background yield and the fraction of W +jets events in both the signal and control region using the CR-only fit, described in Sec. 7.1.

5.2.1 Extrapolating in b -jet Multiplicity

The extrapolation from the W +jets control region to the signal region can be decomposed into two components: first a kinematic extrapolation across m_T and then a flavor extrapolation from 0 b -jets to > 0 b -jets. This section demonstrates that the kinematic extrapolation is similar for the b -veto and b -tag selection. It is not possible to isolate a pure sample of W +jets events in the data with at least one b -tagged jet due to the contamination from $t\bar{t}$, so W +jets simulation is used for this study. The top left and right plots in Fig. 5.12 build upon the preselection with additional jet requirements to be kinematically similar to SR13: $p_T > 100, 80, 50, 25$ GeV. There are significantly more events in the b -veto region than in the b -tag region, but the shape (in simulation) of the m_T distribution is nearly the same, as seen in the top middle plot of Fig. 5.12. There seems to be a small systematic feature just beyond m_W that could be due to the difference in light jet and b -jet energy resolutions, which could shift the location of the Jacobian edge. The ratio of the factors used to extrapolate from low m_T to high m_T (transfer factors, or TF) are largely independent of the SR m_T requirement, as illustrated by the bottom middle plot of Fig. 5.12. This is not strictly necessary for the method to work, but the similarity in transfer factors does make the extrapolation more robust.

5.2.2 Modeling the m_T Tail

While $t\bar{t}$ events enter the signal region mostly through mis-identification of leptons, W +jets events pass the signal region event selection mostly through resolution smearing. Therefore, it is crucial to validate the modeling of the high m_T tail for W +jets events. Since the signal is largely suppressed by a b -jet veto, there is little concern for signal contamination for nearly all of the 0 b -tag phase space. The kinematic region $90 \text{ GeV} < m_T < 120 \text{ GeV}$ is used to form a W +jets validation region in analogy to the $t\bar{t}$ validation regions and is discussed in Sec. 7.3. The region $m_T > 120 \text{ GeV}$ is investigated in this section, using SR13 as an example. To be as kinematically close to the signal region as possible, events are required to have four jets with $p_T > 100, 80, 50, 25 \text{ GeV}$ in addition to the preselection that includes $E_T^{\text{miss}} > 200 \text{ GeV}$.




Figure 5.12: An illustration of the extrapolation from the W +jets control region to the signal region decomposed into two steps. The top left plot is the m_T distribution in a preselection with a b -jet veto. The bottom left plot shows the transfer factor (TF) that relates the yield in the region $30 \text{ GeV} < m_T < 90 \text{ GeV}$ region to various regions with m_T greater than the value on the x-axis. The right column of plots are analogous to the first column, but with a b -jet requirement instead of a veto. The middle column plots are ratios of the left and right columns. The error bars represent statistical uncertainty.

Figures 5.13 and 5.14 show various kinematic distributions in the resulting WVR-tail validation region which has about 100 events with approximately 60% W +jets purity. Even though events in the WVR-tail have significantly mis-measured $\vec{p}_T^{\text{miss}}$, the MC is a good model within the statistical uncertainties.




Figure 5.13: The m_T (top left), E_T^{miss} (top right), am_{T2} (bottom left), and $H_{T,\text{sig}}^{\text{miss}}$ (bottom right) distributions in the WVR-tail validation region.




Figure 5.14: The leading (top left), subleading (top right), third leading (bottom left) and fourth leading (bottom right) p_T distributions in the WVR-tail validation region.

5.3 Single Top

Even though the required $\sqrt{s}$ is lower, the cross-section for the single production of top quarks is lower than for $t\bar{t}$ due to the involvement of electroweak couplings and/or b -quarks in the proton. Inclusively, the dominant single top production mechanism is the t -channel exchange of a W boson, illustrated by the middle diagram of Fig. 5.15. However, with only one out-going W boson (from the top quark) and a small number of out-going matrix-element quarks and gluons (two if the W decays leptonically in Fig. 5.15), the t -channel single top process is effectively suppressed by m_T and mild jet requirements. In contrast, the Wt -channel process has two W bosons, and similarly to top quark pair production, the doubly leptonically decaying W boson final state can effectively circumvent an m_T threshold. As with $t\bar{t}$, nearly all of the Wt events predicted to pass the SR requirements have two real leptons and so much of the discussion from Sec. 5.1 related to validating the modeling of extra jets directly applies to the single top process.




Figure 5.15: Feynman diagrams for the s -, t -, and Wt -channels on the left, middle, and right, respectively. Even though these are leading order in α_s , these are not necessarily the dominant diagrams because the b -quark PDF is highly suppressed due to the b -quark mass.

Several variables described in Sec. 3.2 are designed to select events with a resonant hadronically decaying W boson. Single top events with two leptonically decaying W bosons would fail this requirement, except the invariant mass of one of the leptons and one of the b -quarks is naturally much larger than the corresponding observable in top quark pair production. Figure 5.17 shows the invariant mass of the non-resonant W boson and b quark from Wt events. By construction, $m(bW) > m(W)$, but approximately 50% of events have $m(bW) > m_{\text{top}} \sim 175$ GeV. Requirements

that select relatively high $m(\mathbf{bl})$ such as the large-radius jet mass and $\mathbf{am}_{T2}$ will have higher efficiency for Wt events compared with top quark pair production. As a result, Wt is predicted to be a non-negligible background in many of the signal regions.




Figure 5.16: The distribution of the invariant mass of the W boson and b -quark from the non-top quark side. Diagrams with doubly resonant top quarks are explicitly removed. About 50% of the distribution has $m(Wb) > m_{\text{top}} \sim 175$ GeV.

The Wt -channel Feynman diagram in Fig. 5.15 only contains one out-going b -quark from the matrix element. However, POWHEG-BOX predicts that inclusively 40% of the events have a second out-going b -quark in the NLO ME using the 5-flavor scheme in which b -quarks are treated as constituents of the proton. Figure 5.17 shows that this fraction increases with E_T^{miss} . After the full event preselection, nearly all Wt events have a second b -jet at particle-level.

One key challenge with the Wt process is the interference between Wt simulated at NLO and LO $t\bar{t}$. Representative beyond leading order Wt Feynman diagrams with a second b -quark in the final state are shown in Fig. 5.18. There are some diagrams that overlap with $t\bar{t}$ when a Wb pair go on-shell. This is further discussed in Sec. 6.2.2 in the context of systematic uncertainties, but is an important motivation




Figure 5.17: The average number of b-jets, the probability for at least and for exactly one b-jet at particle-level from POWHEG-BOX+PYTHIA 6.

for constraining aspects of this process with data. Another motivation is that unlike $t\bar{t}$, the Wt process has only recently been observed [597, 598] and has essentially no constraints on the modeling of its kinematic properties. Section 5.3.1 describes a single-top CR used for the first time in a $t\bar{t} + E_T^{\text{miss}}$ search in the early $\sqrt{s} = 13$ TeV search. For the signal regions at $\sqrt{s} = 8$ TeV, the single top background is predicted directly from simulation.

5.3.1 A Data-driven Approach

The only difference between Wt and $t\bar{t}$ when there are two out-going b-quarks in the ME is the presence of one non-resonant Wb pair. Therefore, variables aimed at reconstructing the hadronic or leptonically decaying top quarks can (partially) separate Wt from $t\bar{t}$. One particularly powerful variable for this task is $a\mathbf{m}_{T2}$, for which $t\bar{t}$ events are significantly reduced beyond $a\mathbf{m}_{T2} \gtrsim m_{\text{top}}$. As mentioned in the previous section, many Wt events have a second b-jet at particle-level. Requiring two b-jets is crucial for obtaining a high Wt purity because of the contamination from W +jets events which are also not bounded by $a\mathbf{m}_{T2} \lesssim m_{\text{top}}$. Additional $t\bar{t}$ suppression is possible in the two b-jet selection when the ΔR between b-jets is required to be relatively




Figure 5.18: Next-to-leading-order (up to $\alpha_s^4 \alpha_w^2$) Feynman diagrams from MG5_AMC that have the same outgoing particles as $t\bar{t}$ at tree level. The top left diagram overlaps with $t\bar{t}$ when the two intermediate top quarks go on-shell. The top right diagram has no top quarks at all and the bottom diagrams contain top quarks but do not interfere with top quark pair production even when the intermediate top quark(s) go on-shell.

large. This is because one way for one-lepton $t\bar{t}$ events to exceed the $\mathbf{am}_{T2}$ endpoint is for a charm jet from the hadronically decaying W boson to be $\mathbf{b}$ -tagged with higher $\mathbf{b}$ -tagging weight than a second $\mathbf{b}$ -jet from the top quark decay. This is illustrated in Fig. 5.19 at parton level. For a given choice of $\mathbf{b}$ -jet, $\mathbf{am}_{T2} \sim \max(80, m(\mathbf{bl}))$. To account for combinatorics, the selected $\mathbf{am}_{T2}$ is minimized over both pairings of $\mathbf{b}$ -jets; therefore $\mathbf{am}_{T2} \sim \min(m(\mathbf{b}_1\mathbf{l}), m(\mathbf{b}_2\mathbf{l}))$. When the $\mathbf{b}$ -jet entering the $\mathbf{am}_{T2}$ calculation is on the same side as the lepton, then $m(\mathbf{bl})$ is bounded from kinematics by $\sqrt{m_{\text{top}}^2 - m_W^2} \sim 155 \text{ GeV}$ independent of the top quark p_T so this will generally be smaller than the invariant mass of the lepton and the charm-jet which increases with p_T^{top} . However, if both the (true) $\mathbf{b}$ -jet and the charm-jet are from the opposite

top quark from the lepton, the minimum can be much larger than m_{top} .




Figure 5.19: The value of am_{T2} from a simple scalar parton-level simulation. In the correct assignment, the two b -quarks from the top decay are used to calculate am_{T2} . When the charm quark from the W decay is used in place of one of the b -quarks, then the second b -quark used in the calculation can either be from the same top quark as the lepton or the opposite top quark. In all cases, the plotted value is the minimum over both pairings of the lepton and ‘ b -jets’.

Table 5.6 summarizes the event selection for the single top control region. The m_T window is larger than for the $t\bar{t}$ control region in order to increase statistics; there are not enough events for a validation region at high m_T which is what the region beyond $m_T = 90$ GeV is used for in the $t\bar{t}$ case. A few of the other requirements are also loosened from the $t\bar{t}$ case in order to increase statistics. With about 80 predicted events in the single top control region, the Wt purity is about 40%. About 3% of the single top events are due to other single top processes, dominated by t -channel production.

Figure 5.20 shows the am_{T2} distribution in the Wt control region with all selections applied aside from the am_{T2} requirement. There may be a small slope in the data/MC ratio for $am_{T2} \lesssim 200$ GeV, but in the control region there is no significant evidence for mis-modeling. The bottom ratio panel in Fig. 5.20 shows that the Wt purity increases by nearly two orders of magnitude between ~ 150 GeV and

Requirement	SR13	TCR	STCR
m_T [GeV]	> 200	[30, 90]	[30, 120]
$n_{b\text{-jets}}$	> 0	> 0	> 1
E_T^{miss} [GeV]	> 350	> 250	> 200
$H_{T,\text{sig}}^{\text{miss}}$	> 20	> 15	> 5
am_{T2}	> 175	[100, 200]	> 200
Large R jet mass [GeV]	> 140	> 140	–
$\Delta R(b, l)$	< 1.5	–	–
$\Delta R(b, b)$	–	–	< 1.2
$t\bar{t}$ 1L	0.0	72.0	30.8
$t\bar{t}$ (other)	0.3	12.4	4.7
Wt	0.1	3.8	29.9
Single top (other)	0.0	1.6	1.0
W +jets	0.2	7.9	11.9
VV	0.1	1.8	2.5
$t\bar{t} + V$	0.4	0.7	1.2
Wt purity	12%	4%	36%
SM (pre-fit)	1.1	100.3	82.0
Data	Sec. 7	102	71

Table 5.6: The definition of the single top control region for SR13 compared with the definitions of SR13 and the corresponding $t\bar{t}$ control region. Only the requirements that differ from the corresponding signal region are indicated in the table, with a ‘–’ if there is no requirement on the given variable. The lower panel indicates pre-fit predictions from the simulation compared with the data.

~ 300 GeV. The b -jet related selections are illustrated in Fig. 5.21. The lower ratio in the $n_{b\text{-jets}}$ clearly shows the significant improvement in the Wt purity by explicitly requiring a second b -jet. As motivated earlier, the $\Delta R(b_1, b_2)$ distribution peaks at low values for $t\bar{t}$ and the purity significantly increases for $\Delta R \gtrsim 1$. The overall MC prediction is a slightly above the measured data so the Wt contribution in the final

results is scaled down from the control region method.




Figure 5.20: The distribution of am_{T2} in the single top control region with all selections applied except the am_{T2} requirement. The left and right plots are identical except for the difference in scale.




Figure 5.21: The distribution of $n_{b\text{-jets}}$ and $\Delta R(b_1, b_2)$ in the single top control region with all selections applied except the ones shown in the plots.

5.4 Top Quark Pair Production with a Z Boson

The production of a Z boson in association with a top quark pair that decays into neutrinos is an irreducible background for the stop search with $\tilde{t} \rightarrow t + \tilde{\chi}^0$. Figure 5.22 shows two representative Feynman diagrams where the phenomenological similarity between the two processes is highlighted with dashed circles around the missing momentum from neutralinos/neutrinos. The cross-section for $t\bar{t} + Z(\rightarrow \nu\nu)$ is suppressed with respect to generic top quark pair production by $\sim \alpha_w \times \text{BR}(Z \rightarrow \nu\nu) \times P \sim 0.001P$ where $P \sim (m_{t\tilde{t}}/m_{t\tilde{t}Z})^2 \sim \mathcal{O}(0.1)$ phase space factor (see Fig. 2.1). Therefore, this process is only a significant background for the high stop mass search where the signal cross-section is comparably small. In this regime, the $t\bar{t} + Z$ is a dominant background even though its cross-section is only known with a 30% uncertainty from dedicated measurements [599–602]. The early $\sqrt{s} = 8$ TeV analysis used a simulation-only method to estimate the $t\bar{t} + Z$ background and the full $\sqrt{s} = 8$ TeV began using a data-driven technique to validate this estimate. The data-driven estimate became fully integrated into the control region method for the default estimate in the early $\sqrt{s} = 13$ TeV analysis. Section 5.4.1 briefly describes the simulation-only estimate and Sec. 5.4.2 details the data-driven technique using photons.




Figure 5.22: Feynman diagrams for stop pair production and decay (left) and top quark pair production in association with a Z boson that decays to neutrinos (right). The dashed circles show the dominant contribution to the E_T^{miss} . There are many other leading order Feynman diagrams for the $t\bar{t} + Z$ process, which are described in Sec. 5.4.2.

5.4.1 Estimation from Simulation

MadGraph 5 (MG5_aMC) is used to simulate $t\bar{t} + Z$ events at $\sqrt{s} = 8$ (13) TeV. The top row of Fig. 5.23 shows the distribution of the E_T^{miss} and m_T for $t\bar{t} + Z$, $t\bar{t}$ and signal events. The stop mass sets a natural scale for these kinematic variables, which are steeply falling for the SM processes. However, the bottom row of Fig. 5.23 illustrates the challenge with $t\bar{t} + Z$: it does not need a second lepton to pass the m_T requirement and therefore has a hadronically decaying top quark⁴. As a result, the m_{jet} distribution is nearly the same as for signal. Furthermore, the am_{T2} distribution is generally harder than the $t\bar{t}$ background due to the extra energy from the Z boson.




Figure 5.23: Distributions of (clockwise) E_T^{miss} , m_T , m_{jet} , and am_{T2} . The stop model has $(m_{\text{stop}}, m_{\text{LSP}}) = (800, 0)$. Events are required to have four jets with $p_T > 50, 50, 25, 25$ GeV, $E_T^{\text{miss}} > 300$ GeV, $m_T > 120$ GeV and at least one $R = 1.2$ large radius re-clustered jet with $p_T > 300$ GeV. The last bin contains overflow.

⁴In contrast to $t\bar{t}$ events without a Z boson that require a second lepton in order for $m_T > m_W$.

5.4.2 A Data-driven Method with Photons

5.4.2.1 Motivation and Overview

Since the process $t\bar{t} + Z(\rightarrow \nu\bar{\nu})$ is a significant and irreducible background, in particular for signal regions targeting high mass stops, it is desirable to constrain the normalization using data-driven techniques. However, it is not possible to isolate a pure sample of $t\bar{t} + Z$ events with sufficient statistics to set a useful constraint on the yield in the signal region. Leptonic Z decays allow for a pure sample, but the cross-section times branching ratio ($\sim 10\%$) is too small - similar to the pair production of 800 GeV stops (see Fig. 2.1). One possibility is to use a similar process: $t\bar{t} + \gamma$. Using photons to constrain Z production is a standard technique for estimating inclusive $Z(\rightarrow \nu\bar{\nu}) + \text{jets}$ background processes [603–607] that has been extensively studied theoretically [608–610] and experimentally [611], but has never before been studied or applied to $t\bar{t} + Z$. The main benefit of using photons is that they can be directly identified with high purity with no loss due to a small branching ratio to leptons. Even though the $t\bar{t} + \gamma$ process itself was only recently observed [612–614], the studies in this section show that a selection with high purity and sufficient statistics at high p_T can be constructed. Since the Feynman diagrams for $t\bar{t} + Z$ and $t\bar{t} + \gamma$ are nearly identical, the theoretical uncertainty in extrapolating from $t\bar{t} + \gamma$ to $t\bar{t} + Z$ is expected to be small. Section 5.4.2.2 quantifies this similarity at leading order.




Figure 5.24: Leading order Feynman diagrams for $t\bar{t} + Z$ (left) and $t\bar{t} + \gamma$ (right) with the boson produced in the final state. In addition to these diagrams, there are two additional sets of FSR diagrams with gluons in the initial state: one with the same setup as above and one with a t -channel exchange of top quarks fusing into the boson (see Fig. 5.22).

5.4.2.2 Relating $t\bar{t} + \gamma$ to $t\bar{t} + Z$ at Leading Order

The properties of $t\bar{t} + Z$ and $t\bar{t} + \gamma$ matrix elements are similar: the sets of Feynman diagrams are nearly identical. In addition to the final state radiation diagrams in Fig. 5.24, the other leading order diagrams with the boson radiated in the initial state are shown in Fig. 5.25. Since the gluon is not charged under the electroweak force, only $q\bar{q}$ initial states contribute to the ISR diagrams while both gluon-gluon and $q\bar{q}$ diagrams contribute at leading order to the FSR diagrams. The only diagrams which are different between $t\bar{t} + Z$ and $t\bar{t} + \gamma$ are the ones that begin at NLO (such as Fig. 5.26) due to the coupling of neutrinos to Z bosons that does not exist for photons.




Figure 5.25: Leading order Feynman diagrams for $t\bar{t} + Z$ (left) and $t\bar{t} + \gamma$ (right).

Even though the set of Feynman diagrams are basically identical for $t\bar{t} + Z$ and $t\bar{t} + \gamma$, the relative contributions are different because the Z boson couples stronger to down-type quarks and the photon couples stronger to up-type quarks. In particular, the photon couples with strength eQ_q which is $Q_q = 2/3$ for up-type quarks (including the top quark) and $Q_q = 1/3$ for down-type quarks. The Z boson coupling is different for left- and right-handed fermions with $e(T^3 - \sin^2 \theta_W Q_q)/(\cos \theta_W \sin \theta_W)$, where T^3 is weak isospin. For up-type quarks, the first term is $(\frac{1}{2} - \frac{2}{3} \sin^2 \theta_W)$ for left-handed quarks and $-\frac{2}{3} \sin^2 \theta_W$ for right-handed quarks. Likewise, for down-type quarks, the Z boson coupling strength is proportional to $(-\frac{1}{2} + \frac{1}{3} \sin^2 \theta_W)$ for left-handed quarks and $\frac{1}{3} \sin^2 \theta_W$ for right-handed quarks. The $t\bar{t} + Z$ to $t\bar{t} + \gamma$ cross-section ratio for a fixed quark type is given by




Figure 5.26: A lowest order (α_w^5) diagram that exists for $t\bar{t} + Z$ that has no $t\bar{t} + \gamma$ version (left) and the strong production ($\alpha_w \alpha_s^4$) analogue (right). The left diagram is suppressed with respect to the right one by $(\alpha_w/\alpha_s)^4 \sim 10^{-4}$.

$$R_q = \frac{\sigma_{t\bar{t}+Z}^{\text{via } q}}{\sigma_{t\bar{t}+\gamma}^{\text{via } q}} = \frac{\frac{1}{2} (\sigma_{t\bar{t}+Z}^{\text{via } q_L} + \sigma_{t\bar{t}+Z}^{\text{via } q_R})}{\sigma_{t\bar{t}+\gamma}^{\text{via } q}} = \frac{(\frac{1}{2} - |Q_q| \sin^2 \theta_W)^2 + (Q_q \sin^2 \theta_W)^2}{2Q_q^2 \cos^2 \theta_W \sin^2 \theta_W}, \quad (5.1)$$

where the factor of two in the denominator is from averaging over the initial state spins. For $\cos \theta_W = m_W/m_Z$ with $m_W \approx 80.385$ GeV and $m_Z \approx 91.1876$ GeV, $R_u \approx 0.945$ and $R_d \approx 4.851$. If the bosons were only produced via the FSR process, then the total cross-section ratio is expected to be R_u since the top quark is up-type. In contrast, if the bosons were only produced via the ISR processes, then there would be a tradeoff between R_u and R_d due to the mixture of $u\bar{u}$ and $d\bar{d}$ initial states. As the valence up quark PDF dominates at high p_T , asymptotically the initial state cross-section ratio should also approach R_u . The behavior in the cross-section ratio for the ISR processes is similar to the inclusive V +jets [610] case. However, the FSR processes dominate for $t\bar{t} + V$ ($\gtrsim 80\%$) for $p_T^V \gtrsim 100$ GeV.

Even if the couplings were identical between $t\bar{t} + Z$ and $t\bar{t} + \gamma$, there would still be a significant difference in the cross-sections due to the boson masses. The biggest impact of the large Z boson mass is at low p_T^V . Generically, electroweak radiation receives double Sudakov logarithm enhancements [20]:

$$d\sigma(p \rightarrow p' + \gamma/Z) \approx d\sigma(p \rightarrow p') \times \frac{\alpha}{\pi} \log\left(\frac{-q^2}{\mu^2}\right) \log\left(-\frac{q^2}{m^2}\right), \quad (5.2)$$

where the first logarithm is due to a soft singularity that is cutoff by the detectability scale μ for photons and the boson mass for $t\bar{t} + Z$ and the second logarithm is the collinear divergence that is cutoff by the mass of the emitting particle⁵. Since $\mu \ll m_Z$, there is a large enhancement of photon bremsstrahlung at low p_T . This is further complicated by the fact that this enhancement occurs for all electrically charged particles in the final state, independent of their relationship to the hardscatter process (see Sec. 5.4.2.3). Fortunately, the region where $t\bar{t} + Z$ is relevant corresponds to high p_T^V where differences due to the Z boson mass are less important. Figure 5.27 shows the conditional distribution of the Z p_T given the particle-level E_T^{miss} . For a selection requiring $E_T^{\text{miss}} > 300$ GeV, 2/3 of the Z bosons have $p_T > 300$ GeV and about 90% have $p_T > 200$ GeV.




Figure 5.27: The distribution of the Z p_T given the particle level E_T^{miss} in $t\bar{t} + Z(\rightarrow \nu\bar{\nu})$ events. All neutrinos contribute to the particle level E_T^{miss} .

⁵See Ref. [615] for a nice discussion on electroweak radiation in the ultra high $p_T \gtrsim 10$ TeV regime.

Figure 5.28 shows the cross-section ratio of $t\bar{t} + Z$ to $t\bar{t} + \gamma$ for parton-level calculations of the ISR processes and the FSR processes with gluon-gluon initial states. At low $p_T^V \ll m_Z$, the ratio is very small due to the large enhancement for photons that is cutoff for Z bosons (see Eq. 5.2). The kinematic differences between $t\bar{t} + Z$ and $t\bar{t} + \gamma$ are mostly eliminated as $p_T^V \gg m_Z$. As expected, the ratio for the ISR processes (red) lies between R_u and R_d and is closer to R_u due to the valence up quarks (as $p_T \rightarrow \sqrt{s}, R \rightarrow R_u$). This ratio slightly decreases with p_T^V as the fraction of up quarks increases with $\sqrt{s}$. However, this is a small effect, illustrated by Fig. 5.29. Over three orders of magnitude in p_T^V , the relative contribution from $u\bar{u}$ increases by only 5-10%. The most puzzling aspect of Fig. 5.28 in the context of the above discussion is the blue line. Since the top quark is an up-type quark, the above argument suggests that the blue line should be $R_u \sim 1$, about half of what is observed. The reason is subtle and is a new feature of the $t\bar{t} + Z$ that is not relevant for generic Z +jets. Due to its mass, the Z boson has three polarization states while the photon only has two (transverse) states. The bosons produced in inclusive Z +jets are mostly transverse and so the additional polarization state is irrelevant. However, just as W bosons from $t\bar{t}$ are mostly longitudinally polarized, the FSR Z bosons have a significant ($\sim 50\%$) longitudinal polarization. When only the transverse polarizations are considered⁶, the ratio is indeed close to R_u , as shown by the green line ratio in Fig. 5.28.

Figure 5.30 shows the cross-section ratio for all sub-processes using the simulation setup described in Sec. 2.3. The next section describes the identification of photons in the simulation in more detail, which in part accounts for some differences (larger photon contribution) between Fig. 5.30 and the expectations from Fig. 5.28. A lower ratio is expected at $\sqrt{s} = 8$ TeV compared with $\sqrt{s} = 13$ TeV in part because a fixed p_T samples a lower momentum fraction at $\sqrt{s} = 13$ TeV which moves the ISR process ratio away from R_u .

For high p_T bosons, Fig. 5.30 shows that the cross-section ratio is nearly unity.

⁶This is accomplished by observing all longitudinal helicity states DATA (NHEL(I, 2), I=1,5) /-1,-1,-1,-1, 0/ in `matrix1.f` and then skipping these states (2 in this case) in the loop DO I=1,NCOMB so that TS(I) remains zero. Thank you to Michael Peskin for the idea and Valentin Hirschi for the assistance in implementation.




Figure 5.28: The cross section ratio for $t\bar{t} + Z$ and $t\bar{t} + \gamma$ for various sub-processes described in the text. Unlike the simulation setups described in Sec. 2.3, the $t\bar{t} + Z$ and $t\bar{t} + \gamma$ use exactly the same setup: MG5_AMC 2.1.1 with PDF set CTEQ6L1. No extra partons are generated in the ME and photon radiation from top decay products is not included. The ISR only processes are generated with the syntax `generate p p > t t~ a / t t~`.

The power of the photon method is that the $Z(\rightarrow e^+e^-/\mu^+\mu^-)$ branching ratio is about 6% so the usable cross-section for the photon process is about 20 times larger than the Z process. As will be described in Sec. 5.4.2.5, photons can be identified and reconstructed with high efficiency and purity. However, there is a finite acceptance for photon reconstruction while the neutrinos from $t\bar{t} + Z(\rightarrow \nu\bar{\nu})$ can be anywhere in the detector. Especially at high p_T where the photons and Z bosons are mostly central, this is a subdominant effect to the others discussed above. For example, about 2.2% of photons have $|\eta| > 2.5$ at $p_T > 100$ GeV, 1.6% for $p_T > 200$ GeV and about 1.2% for $p_T > 300$ GeV.




Figure 5.29: Given a $t\bar{t} + V$ ISR event, this is the fraction of events originating from a $u\bar{u}$ initial state. The dotted lines show the $p_T = 0$ fractions.




Figure 5.30: The full cross-section ratio between $t\bar{t} + Z$ and $t\bar{t} + \gamma$ at 8 TeV (left) and 13 TeV (right).

5.4.2.3 Simulation and Matrix Element Photons

The discussion in Sec. 5.4.2.2 was focused on photons originating directly from the hard scatter process. However, there are two significant sources of additional photons at particle-level: radiation from charged particles from the top quark decays and the decays of neutral pions, $\pi^0 \rightarrow \gamma\gamma$. Figure 5.31 shows representative diagrams from MADGRAPH when the photon is radiated from one of the charged decay products from the top quark. As noted in Sec. 5.4.2.2, Z Bremsstrahlung is highly suppressed compared to photon radiation and so these photons are not directly useful for constraining the $t\bar{t} + Z$ cross-section. The fraction of photons from the charged decay products of the top quark decreases with photon p_T . Figure 5.32 shows the fraction of photons produced from the charged top quark decay products ($b, W^\pm, l$) as a function of the photon p_T . In agreement with the fractions reported by Ref. [544], photons from the charged top quark decay products dominate until about $p_T \gtrsim 60$ GeV and this fraction decreases to reach about 25-30% by $p_T \gtrsim 100$ GeV.




Figure 5.31: Representative diagrams from MADGRAPH where the photon originates from one of the charged decay products of the top quark.

High p_T photons are also abundantly produced from neutral pions. A majority of the hadrons from quark and gluon fragmentation are pions and about 1/3 of the pions are π^0 . When one of the π^0 inside a jet carries a large momentum fraction of the initiating quark or gluon, the photons from the π^0 decay can have significant p_T . Most of these photons can be separated from the hard-scatter photons because they are non-isolated from the remaining hadronic activity inside the jet. However, on occasion a real photon from a π^0 will be reconstructed as an isolated photon,




Figure 5.32: The fraction of photons radiated from the charged top quark decay products (Top Decay) versus the fraction of photons directly from the hard-scatter (Not from Top Decay) as a function of the photon p_T . Events are generated with MG5_aMC at leading order with the full process `generate p p > t t~ > l+ vl b b~ l- vl~ a`. Photons are categorized as originating from the charged decay products of the top quark by parsing the ancestry stored in the LHE output of MG5_aMC. The band is the MC statistical uncertainty.

which will artificially decrease the $t\bar{t} + Z$ to $t\bar{t} + \gamma$ cross-section ratio. Experimental tools for suppressing these photons are described in Sec. 5.4.2.5. The remainder of this section focuses on the labeling of photons in the simulation and the removal of overlap between different generators that cover the same regions of phase space.

Photons are generated at nearly every stage of event simulation. The MadGraph (or MG5_aMC) matrix elements include $t\bar{t} + \gamma$ (but not the POWHEG-BOX matrix elements), PYTHIA and PHOTOS ($\sqrt{s} = 8$ TeV only) add photons as ISR and FSR during fragmentation, and photons can be generated by GEANT4 during the interactions of particles with the detector. Furthermore, photons generated at one stage can be removed at another stage. For example, there is a small probability that photons from the ME can be converted to fermion pairs in PYTHIA and photon conversions in the detector are common. It is therefore crucial to specify a hierarchy in order to

avoid double-counting of photons. The highest preference is given to photons that originate from the ME generator. Therefore, ISR photons from PYTHIA or PHOTOS must be removed as they cover the same region of phase space. Table 5.33 shows an example event where PYTHIA adds a high p_T ISR photon to a $t\bar{t}$ event that needs to be removed as it is covered by the ME $t\bar{t} + \gamma$ sample. The composition of high p_T photons in the POWHEG-BOX $t\bar{t}$ sample is shown in Figure 5.34. Figure 5.35 illustrates how the labeling is performed. Most of the photons in particle level events with at least one photon with $p_T > 80$ GeV are from (asymmetric) neutral hadron decays. Only 3% originate from ISR and only 15% are radiated off of the top quark or its immediate decay products. These events are the ones that need to be removed.

----- PYTHIA Event Listing (complete event) -----											
no	id	name	status	mothers		daughters		colours		px	py
0	90	(system)	-11	0	0	0	0	0	0	0.000	0.000
1	2212	(p+)	-12	0	0	307	0	0	0	0.000	0.000
2	2212	(p+)	-12	0	0	308	0	0	0	0.000	0.000
3	21	(g)	-21	7	0	5	6	101	102	0.000	0.000
4	21	(g)	-21	8	8	5	6	102	103	0.000	0.000
5	6	(t)	-22	3	4	9	9	101	0	-4.567	-88.578
6	-6	(tbar)	-22	3	4	10	10	0	103	4.567	88.578
7	2	(u)	-41	12	12	11	3	101	0	0.000	-0.000
11	2	(u)	-43	7	0	16	16	102	0	175.819	-37.109
16	2	(u)	-44	11	11	22	22	102	0	184.181	-42.930
22	2	(u)	-44	16	16	30	30	102	0	183.933	-42.817
30	2	(u)	-52	22	22	38	38	102	0	182.014	-42.370
38	2	(u)	-44	30	30	64	64	102	0	182.669	-39.669
64	2	(u)	-44	38	38	104	105	102	0	182.666	-39.696
104	2	(u)	-51	64	0	128	128	125	0	181.192	-38.581
105	21	(g)	-51	64	0	122	122	102	125	1.494	-1.236
128	2	(u)	-52	104	104	143	144	125	0	167.269	-35.617
143	2	(u)	-51	128	0	177	177	136	0	159.955	-34.472
177	2	(u)	-44	143	143	232	233	136	0	160.057	-34.206
232	2	(u)	-51	177	0	265	266	136	0	25.298	-5.486
233	22	(gamma)	-51	177	0	350	350	0	0	134.759	-28.720
350	22	gamma	62	233	233	0	0	0	0	134.760	-28.362

Figure 5.33: An example (abridged) event record from PYTHIA showering a $t\bar{t}$ event in which a high p_T photon is added as ISR (red).

Table 5.7 shows the photon labeling hierarchy and the relative fractions of the various categories in the $t\bar{t} + \gamma$ validation region that is described in Sec. 5.4.2.5. After the basic event selection, most events are from the dedicated $t\bar{t} + \gamma$ sample with a matrix element photon with $p_T > 80$ GeV. The $\sim 18\%$ of events from POWHEG-BOX with a MEP with $p_T > 80$ GeV should be removed to avoid double-counting. The next biggest category of events have a particle-level photon with $p_T > 80$ GeV that originates from somewhere other than the ME. The dedicated $t\bar{t} + \gamma$ sample is

Photon Origin in Powheg+Photos




Figure 5.34: The composition of the leading photon origin in POWHEG-BOX+PYTHIA 6+PHOTOS $t\bar{t}$ events with at least one photon at particle-level with $p_T > 80$ GeV. The ME- $t\bar{t}$ system category includes photons radiated from the decay products of the top quarks.

generated with a $p_T > 15$ GeV photon filter in the ME and a $p_T > 80$ GeV photon filter after PYTHIA. While most events have a ME photon with $p_T > 80$ GeV, about 2% of events pass the PYTHIA filter due to a pion decay. These events should be removed in favor of the higher order corrections in the POWHEG-BOX sample. The remaining small fraction of events have no particle-level photon with $p_T > 80$ GeV and are the result of fake photons either from electrons or jets that have a large electromagnetic fraction.




Figure 5.35: Example particle ancestries for the categories shown in Fig. 5.34. Photons are identified with their PDG ID (22) and then the MC event ancestry is parsed to find the origin of the photon. Each line contains four items: PDG ID, particle name, HepMC status code [616], and location in the event record. Not all numbers are used for the location in the event record, but higher numbers do correspond to later in the simulation.

Condition	Label	VR Fraction
if $\exists$ MEP with $p_T > 80$ GeV,	$t\bar{t} + \gamma$	$\sim 80\%$
else if $\exists$ a truth γ with $p_T > 80$	$t\bar{t}$ +shower γ (with $t\bar{t}$)	$\sim 15\%$
else if $\Delta R(\text{reco } \gamma, \text{truth } e) < 0.1$	$t\bar{t} + \text{fake } (e \mapsto \gamma \text{ fake})$	$\sim 2\%$
else	$t\bar{t}$ +fake (jet $\mapsto \gamma$ fake)	$\sim 3\%$

Table 5.7: The composition of photon origins in the $t\bar{t} + \gamma$ and $t\bar{t}$ samples at $\sqrt{s} = 8$ TeV. The order of photon precedence is described in the first column. Except for the first row, all other rows are labeled $t\bar{t}$ and not $t\bar{t} + \gamma$. The last column quantifies the fraction of the various categories in the $t\bar{t} + \gamma$ validation region described in Sec. 5.4.2.5

5.4.2.4 Higher Order QCD Corrections

Next-to-leading-order QCD corrections⁷ for the $t\bar{t}Z$ and $t\bar{t}\gamma$ processes have been calculated and are non-negligible. Table 5.8 summarizes the values of the k-factors from the literature. Only the MG5_AMC collaboration has reported k-factors for both processes using the same setup, which is desirable to minimize theoretical uncertainties. Scale and PDF uncertainties in the individual k-factors are also not small - about 20% when reported. Some of the calculations include stable top quarks while others allow for radiation from the top quark decay products. The $\sqrt{s} = 8$ TeV version of the analysis used a k-factor of 1.9 for the $t\bar{t} + \gamma$ validation region following the procedure of Ref. [617] based on the calculation in Ref. [544]. However, a careful investigation of this k-factor reveals that it is likely over-estimated because it is based off of a result using a fixed order calculation with jet requirements and thus artificially increases the NLO cross-section via a higher acceptance. One other difference is that the top decay is correctly treated in the calculation for the case when the k-factor was 1.9. However, the authors state that this is likely not the cause of the higher k-factor, and when considering the fraction of events with high p_T photons from the top decay products (see Sec. 5.4.2.3), this cannot account for the difference with the MG5_AMC calculation. The inclusive k-factor from Ref. [544] is closer to 1.5, which agrees with the inclusive k-factor from the MG5_AMC collaboration. In the control region method, only the cross-section *ratio* between $t\bar{t} + \gamma$ and $t\bar{t} + Z$ is relevant (working only at high p_T^V). Since the diagrams are basically identical, one may expect that the QCD corrections are nearly the same for the two processes. This is supported by the k-factor ratio calculation in Ref. [505] shown in the second row of Table 5.8 and is further examined in this section. Additionally, this section explores the p_T dependence of the k-factor ratio. It is not known from the calculations presented in Table 5.8 if there is a significant p_T dependence to the k-factor or the ratio (the external studies only go to $p_T < 200$ GeV) of k-factors between $t\bar{t} + Z$ and $t\bar{t} + \gamma$.

Tables 5.9 and 5.10 show the result of fixed-order calculations of the LO and NLO

⁷This section benefited from many useful conversations with Till Eifert, Javier Montejo Berlingen, Josh McFayden, Stefan Hoche, and Lance Dixon.

$t\bar{t}\gamma$ and $t\bar{t}Z$ cross sections as a function of the boson p_T threshold for SHERPA+OPENLOOPS (5.9) and MG5_AMC (5.10) at $\sqrt{s} = 13$ TeV. In both cases, an isolation for the photon of $\Delta R = 0.4$ is used implemented by the Frixione cone with $n = 2$ and $\epsilon = 0.025$ [618]. The two calculations give similar results and show that the k-factor ratio is consistent with unity and independent of p_T within 10% over the range $100 \text{ GeV} < p_T^V < 600 \text{ GeV}$. One reason⁸ it might decrease is that at low boson p_T , the dominant contributions are gluon-gluon fusion where the boson comes from a top quark line, whereas at high boson p_T , the quark-quark annihilation dominates where the Z and γ come from ISR and thus the k-factor decreases and tends toward the Z+jets/ γ +jets k-factor ratio, which is $\sim 90\%$ [609]. Based on these calculations, a k-factor ratio of 1 is used for the extrapolation from the $t\bar{t} + \gamma$ CR to the $t\bar{t} + Z$ in the SRs. Uncertainties associated with this choice are described in Sec. 6.2.3.

Reference	$\sigma_{t\bar{t}\gamma}^{\text{LO}}$	$\sigma_{t\bar{t}\gamma}^{\text{NLO}}$	$k_{t\bar{t}\gamma}$	$\sigma_{t\bar{t}Z}^{\text{LO}}$	$\sigma_{t\bar{t}Z}^{\text{NLO}}$	$k_{t\bar{t}Z}$	$k_{t\bar{t}Z}/k_{t\bar{t}\gamma}$
[544]	$1.96^{+0.64}_{-0.45}$	$2.93^{+0.42}_{-0.39}$	1.49				
[505]	$1.203(1)^{+29.6}_{-21.3}$	$1.744(5)^{+9.8}_{-11.0}$	1.45	$0.5273(41)^{+30.5}_{-21.8}$	$0.7598(26)^{+9.7}_{-11.1}$	1.44	0.99
[619]				0.808	1.09	$1.35^{+0.25}_{-0.25}$	
[620]				0.808	1.121(2)	1.39	
[621](1)				0.1035(1)	0.1370(3)	1.32	
[621](2)				0.00379(0)	0.00516(1)	1.36	
[621](3)				0.00325(0)	0.00480(1)	1.48	
[543]				0.1539(1)	0.2057(2)	$1.34^{+0.22}_{-0.27}$	

Table 5.8: NLO QCD corrections to the $t\bar{t} + \gamma$ and $t\bar{t} + Z$ cross-sections. Ref. [544] is for 14 TeV and has a second k-factor given for a second selection that has a harder jet requirement. Since the calculation is fixed-order, this artificially increases the k-factor to the 1.9 value that was used for the 8 TeV analysis. The k-factor in Ref. [543] is for 8 TeV. Both Ref. [619] and Ref. [620] show the differential (in $Z p_T$) k-factor up to 200 GeV, which appears to be relatively flat in that range. The value [621](1) is for the zero-width approximation while [621](2) is for a narrow-width approximation and uses the MSTW08 PDF set for both LO and NLO. The third value [621](3) is for a narrow-width approximation and mixes CTEQ6L1 at LO with CT10 at NLO. Electroweak corrections have also been reported in Ref. [622].

⁸This idea is due to Stefan Hoche.

$p_{T,\text{cut}}^{\text{boson}} [\text{GeV}]$	$\sigma_{t\bar{t}\gamma}^{\text{LO}}$	$\sigma_{t\bar{t}\gamma}^{\text{NLO}}$	$k_{t\bar{t}\gamma}$	$\sigma_{t\bar{t}Z}^{\text{LO}}$	$\sigma_{t\bar{t}Z}^{\text{NLO}}$	$k_{t\bar{t}Z}$	$k_{t\bar{t}Z}/k_{t\bar{t}\gamma}$
100	0.2002(4)	0.329(2)	1.62	0.2330(3)	0.367(1)	1.59	0.98
200	0.0479(1)	0.0784(5)	1.62	0.0812(1)	0.1278(8)	1.58	0.97
300	0.01428(3)	0.0227(2)	1.58	0.02768(5)	0.04244(2)	1.53	0.97
400	0.00489(1)	0.00775(6)	1.59	0.01002(2)	0.01512(8)	1.51	0.95
500	0.001872(6)	0.00291(3)	1.57	0.003917(9)	0.00583(4)	1.49	0.95
600	0.000791(3)	0.00121(2)	1.55	0.001654(5)	0.00240(2)	1.45	0.93

Table 5.9: LO and NLO cross-sections for $t\bar{t}\gamma$ and $t\bar{t}Z$ as a function of the boson p_T threshold computed with SHERPA+OPENLOOPS by Stefan Hoche. All cross-sections are in pb. The numbers in parentheses are the statistical uncertainties. The PDF is CT14. A scale of $H_T = \sum p_T$ over all final state objects is used (the difference between the scalar sum of p_T and m_T was found to be negligible in this range).

$p_{T,\text{cut}}^{\text{boson}} [\text{GeV}]$	$\sigma_{t\bar{t}\gamma}^{\text{LO}}$	$\sigma_{t\bar{t}\gamma}^{\text{NLO}}$	$k_{t\bar{t}\gamma}$	$\sigma_{t\bar{t}Z}^{\text{LO}}$	$\sigma_{t\bar{t}Z}^{\text{NLO}}$	$k_{t\bar{t}Z}$	$k_{t\bar{t}Z}/k_{t\bar{t}\gamma}$
100	0.2634(8)	0.3842(3) ^{+12.9%} _{-13.4%}	1.46	0.3122(10)	0.4209(2)	1.35	0.92
200	0.06305(2)	0.08864(6) ^{+13.2%} _{-13.8%}	1.41	0.1077(3)	0.1433(9)	1.33	0.95
300	0.01842(5)	0.02608(2) ^{+13.6%} _{-14.1%}	1.42	0.03587(1)	0.04760(4)	1.33	0.94
400	0.00615(2)	0.008737(7) ^{+15.5%} _{-15.0%}	1.42	0.01274(4)	0.01673(2)	1.31	0.92
500	0.002305(7)	0.003234(2) ^{+17.0%} _{-15.8%}	1.40	0.00489(2)	0.00643(7)	1.31	0.94
600	0.000947(3)	0.001342(10) ^{+14.4%} _{-15.4%}	1.42	0.002032(8)	0.00258(2)	1.27	0.89

Table 5.10: LO and NLO cross-sections for $t\bar{t}\gamma$ and $t\bar{t}Z$ as a function of the boson p_T threshold computed with MG5_AMC. A custom fortran filter is used to isolate $t\bar{t} + Z$ events with a fixed Z boson threshold. All cross-sections are in pb. The uncertainties on the NLO cross section are from variations of the factorization and renormalization scale. The numbers in parentheses are the statistical uncertainties. The PDF is NNPDF2.3 LO for the LO calculations and NNPDF2.3NLO for the NLO calculations. The for both LO and NLO scale is half the scalar sum of the transverse mass of all out-going partons (default for NLO and scale option 3 for LO [623]). The impact of adding a PS was found to be small ($\lesssim 10\%$).

As stated earlier, given a choice of the k -factor ratio, the actual k -factors themselves do not impact the prediction. Nonetheless, it is useful to make an informed choice for the k -factor in order to directly compare the simulation with the data in the CR. The k -factor used by the $\sqrt{s} = 13$ TeV analysis is 1.33, which is based off of the leading order $t\bar{t} + Z$ cross-section from the ATLAS generation and the NLO cross-section from the MG5_AMC collaboration [505]. Note that the k -factor directly from the MG5_AMC collaboration is about 10% larger because their leading order calculation used a different PDF set (NNPDF2.3LO versus MSTW2008nlo68cl) and top quark mass (172.5 GeV versus 173.2 GeV used by aMC).

In order to justify the use of unity for the cross-section ratio, the leading order simulation for $t\bar{t} + \gamma$ must be as similar as possible to that for $t\bar{t} + Z$. The simulation for both processes are based on MG5_AMC interfaced with PYTHIA 8, but there are some significant differences. In particular, the $t\bar{t} + \gamma$ ($t\bar{t} + Z$) sample uses the CTEQ6L1 (NNPDF2.3) PDF set, a fixed (variable) factorization and renormalization scale of $2 \times m_{\text{top}}$ (transverse mass), and no extra partons (up to two extra partons) are generated in the matrix element. Using the ATLAS simulation framework, small $t\bar{t} + Z$ samples were generated with variations to study the impact of these settings. Changing the PDF from CTEQ6L1 to NNPDF2.3 resulted in a 12% higher cross-section. The cross section is reduced by 2% when no additional partons are considered in the calculation and by 5% when adopting the fixed scale choice of the $t\bar{t}\gamma$ simulation. The combination of the three effects yields a 4% difference in cross section from the choice of generator settings. The $t\bar{t}\gamma$ cross section is increased by 4% to account for these known differences.

5.4.2.5 Event Selection

In order for the $t\bar{t} + \gamma$ process to be as kinematically close as possible to the $t\bar{t} + Z(\rightarrow \nu\bar{\nu})$ process, the γ is added⁹ to the E_T^{miss} , mimicking the lost neutrinos. The sum of $\vec{p}_T^{\text{miss}}$ and $\vec{p}_T^\gamma$ will be denoted $\tilde{\vec{p}}_T^{\text{miss}}$. This new variable is then used to construct $\tilde{E}_T^{\text{miss}}$ and $\tilde{m}_T$ with the standard definitions, replacing $\vec{p}_T^{\text{miss}}$ with $\tilde{\vec{p}}_T^{\text{miss}}$. Table 5.11 shows the event selections used for the $\sqrt{s} = 8$ TeV $t\bar{t} + \gamma$ validation region (VR8) and the $\sqrt{s} = 13$ TeV $t\bar{t} + \gamma$ control region (CR13). The jet p_T requirements are chosen to match the the signal regions. The upper E_T^{miss} requirement for CR13 ensures orthogonality with the $t\bar{t}$ CR. Single lepton and E_T^{miss} triggers are used to collect the data for the VR and a dedicated high p_T photon trigger ($p_T > 120$ GeV) is used for the CR in order to increase the available statistics.

Requirement	VR8	CR13
At least four jets with p_T [GeV] >	80, 60, 40, 25	120, 80, 50, 25
At least one signal photon with $p_T > [\text{GeV}]$	100	125
$\tilde{E}_T^{\text{miss}}$ [GeV] >	120	120
$\tilde{m}_T$ [GeV] >	110	110
$\tilde{H}_{T,\text{sig}}^{\text{miss}} >$	–	5
E_T^{miss} [GeV] <	–	200

Table 5.11: The requirements for the $t\bar{t} + \gamma$ VR ($\sqrt{s} = 8$ TeV) and CR ($\sqrt{s} = 13$ TeV). In both regions, exactly one signal lepton is required with no other baseline leptons. Furthermore the event selections require at least one b -tagged jet. The tilde variables include the photon in the $\tilde{\vec{p}}_T^{\text{miss}}$ as described in the text.

The predicted composition both the VR and CR are summarized in Table 5.12. In addition to the changes in cross-section and integrated luminosity between the two energies, the main difference between the regions is the jet p_T requirements in the CR that are kinematically tighter in order to be close to SR13. In addition, there is a

⁹A more pragmatic adjective would be ‘remove’ instead of ‘add’ since the photon is already part of the E_T^{miss} calculation as a visible object.

photon-electron overlap removal at $\sqrt{s} = 13$ TeV¹⁰ that additionally helps to reduce the $t\bar{t}$ contamination. Both regions have a very high $t\bar{t} + \gamma$ purity, with about 75% in the VR and 92% in the CR.

process	VR8	CR13
$t\bar{t} + \gamma$	75.2 ± 1.6	29.2 ± 1.4
$t\bar{t}$	27.0 ± 1.2	1.6 ± 0.3
Other	1.7 ± 0.5	0.9 ± 0.2
total SM	103.9 ± 2.1	31.6 ± 1.5
data	104	45

Table 5.12: Expected and observed event yields in the $t\bar{t} + \gamma$ validation/control regions. All MC numbers are normalized to 20.3 fb^{-1} for the VR at $\sqrt{s} = 8$ TeV and 3.32 fb^{-1} for the CRs at $\sqrt{s} = 13$ TeV. The $t\bar{t}$ sample at $\sqrt{s} = 8$ TeV is reweighted according to the standard procedure described in Section 5.1. The displayed uncertainties are due to limited statistics.

The remainder of this section shows key kinematic distributions in the VR and CR. Figure 5.36 compares the E_T^{miss} distribution in the VR with the $\tilde{E}_T^{\text{miss}}$ distribution. Nearly all events have $E_T^{\text{miss}} \lesssim 200$ GeV, as the $\tilde{E}_T^{\text{miss}}$ is dominated by the photon momentum, as also expected for $t\bar{t} + Z(\rightarrow \nu\bar{\nu})$ events (see Fig. 5.27). Analogous plots for m_T and $\tilde{m}_T$ are shown in Fig. 5.37. Overall, the simulation agrees well with the data within the large statistical uncertainties, though this is partly coincidental due¹¹ to the large k -factor (see Sec. 5.4.2.4). Supporting plots for the $\sqrt{s} = 13$ TeV CR are in Fig. 5.38 and Fig. 5.39. The photon $p_T > 125$ GeV by construction and has a broad spectrum. Most photons are central, with most photons contained in $|\eta| \lesssim 1$. The $\tilde{E}_T^{\text{miss}}$ and $\tilde{m}_T^{\text{miss}}$ distributions in Fig. 5.39 are similar to the corresponding $\sqrt{s} = 8$ TeV ones. There is no significant evidence for mis-modeling any of the kinematic distributions, though the statistical precision is limited.

¹⁰This idea is due to J. Montejo Berlingen.

¹¹Additionally, the overlap removal described at the end of Sec. 5.4.2.3 is not applied, which would further reduce the total SM by removing approximately 5 $t\bar{t}$ events.




Figure 5.36: The distribution of E_T^{miss} (left) and $\tilde{E}_T^{\text{miss}}$ (right) constructed from $\vec{p}_T^{\text{miss}}$ and $\vec{p}_T^\gamma$. Both distributions use data and simulation in the VR at $\sqrt{s} = 8 \text{ TeV}$. Only statistical uncertainties are included in the error bars and bands.




Figure 5.37: The distribution of m_T^{miss} (left) and $\tilde{m}_T^{\text{miss}}$ (right) constructed from $\vec{p}_T^{\text{miss}}$ and $\vec{p}_T^\gamma$. Both distributions use data and simulation in the VR at $\sqrt{s} = 8 \text{ TeV}$. Only statistical uncertainties are included in the error bars and bands.




Figure 5.38: The distribution of photon p_T and photon η in the $t\bar{t} + \gamma$ CR at $\sqrt{s} + 13$ TeV. Jet energy scale and resolution uncertainties in addition to statistical uncertainties are included in the error bars and bands. A normalization factor of 1.42 is applied. The last bin includes overflow.




Figure 5.39: The distribution of $\tilde{E}_T^{\text{miss}}$ and $\tilde{m}_T$ built from $\tilde{p}_T^{\text{miss}} = \vec{p}_T^{\text{miss}} + \vec{p}_T^\gamma$ in the $t\bar{t} + \gamma$ CR at $\sqrt{s} + 13$ TeV. Uncertainties are the same as Fig. 5.38. A normalization factor of 1.42 is applied. The last bin includes overflow.

5.5 Dibosons

The cross section for the double production of electroweak bosons is suppressed by 3-4 orders of magnitude with respect to the inclusive W +jets cross section. However, unlike generic W +jets, there are some diboson processes which can have a second lepton, possibly with additional neutrinos, that when not identified as such can allow these events to exceed the m_T threshold. Figure 5.40 shows the number of diboson events predicted in simulation after a basic preselection. Due to the large hadronic branching ratio, the semileptonic WW and WZ processes dominate at low m_T . However, with only one leptonically decaying W boson, the semileptonic processes are highly suppressed for $m_T \gtrsim m_W$, after which the dominate processes have multiple leptons/neutrinos. The $WZ \rightarrow l\nu\nu\nu$ process can naturally have large m_T , but its cross-section is slightly lower than the $WW \rightarrow ll\nu\nu$ process. The $ZZ \rightarrow ll\nu\nu$ component of the $VV \rightarrow ll\nu\nu$ is subdominant to the WW part and the split is similar to dilepton $t\bar{t}$: roughly half of the dilepton diboson events have a hadronically decaying τ . Diboson events are a sub-dominant contribution to all signal regions and are estimated using the SHERPA event generator.




Figure 5.40: The m_T distribution of various diboson subprocesses with SHERPA 2.1.

5.6 Overview

This section explores the overall agreement between the data and simulation using all of the background estimation techniques described earlier in this chapter. The final background estimate is performed using the simultaneous fit with systematic uncertainties as nuisance parameters as described in Sec. 7.1. However, it is possible to estimate the per process scale factors by solving the following system of equations:

$$\begin{aligned} N_{\text{WR}}^{\text{data}} &= \mu_W N_{\text{WR}}^{W+\text{jets}} + \mu_{t\bar{t}} N_{\text{WR}}^{t\bar{t}} + \mu_{Wt} N_{\text{WR}}^{Wt} \\ N_{\text{TR}}^{\text{data}} &= \mu_W N_{\text{TR}}^{W+\text{jets}} + \mu_{t\bar{t}} N_{\text{TR}}^{t\bar{t}} + \mu_{Wt} N_{\text{TR}}^{Wt} \\ N_{\text{StR}}^{\text{data}} &= \mu_W N_{\text{StR}}^{W+\text{jets}} + \mu_{t\bar{t}} N_{\text{StR}}^{t\bar{t}} + \mu_{Wt} N_{\text{StR}}^{Wt}, \end{aligned} \quad (5.3)$$

where μ_x is the normalization factor (NF) for process x and N_y^z is the number of simulated or measured events of type z in the y event selection. The data N_y^{data} must be corrected for the simulation-based estimates, $N_y^{\text{data}} = N_y^{\text{data,observed}} - N_y^{\text{VV}} - N_y^{t\bar{t}+V}$. The $t\bar{t} + V$ yield in the $t\bar{t}$, $W+\text{jets}$ and single top control regions is negligible so the data-driven estimate with the photon mostly decouples from the rest of Eq. 5.3. As a system of three equations with three unknowns ($\mu_{t\bar{t}}$, μ_W , and μ_{Wt}), there is a unique solution. Looser versions of the control regions described in Sec. 5.1, 5.2, and 5.3.1 are used in order to study the dependence of the normalization factors on key even kinematic properties. In addition to the preselection, events are required to have the leading four jets with $p_T > 50, 50, 50, 25$ GeV, $E_T^{\text{miss}} > 150$ GeV, and $30 \text{ GeV} < m_T < 90$ GeV. The $t\bar{t}$ enriched region additionally requires $n_{\text{b-jets}} > 0$ and $am_{T2} < 200$ GeV or $n_{\text{b-jets}} = 1$ (orthogonality to the single top region), the $W+\text{jets}$ enriched region requires $n_{\text{b-jets}} = 0$, and the single top enriched region requires $n_{\text{b-jets}} > 1$ and $am_{T2} > 200$ GeV. Figures 5.41, 5.42, 5.43, and 5.44 show the distributions of the leading jet p_T , $H_{T,\text{sig}}^{\text{miss}}$, E_T^{miss} , and the leading large-radius ($R = 1.2$) jet mass in each of the $t\bar{t}$, $W+\text{jets}$, and single top enriched event selections. The $t\bar{t}$ enriched region has about 6000 events, of which about 80% are predicted to be $t\bar{t}$ events; the $W+\text{jets}$ enriched region has about 4000 events, of which which about 75% are predicted to

be W +jets events, and the single top enriched region has about 300 events, of which about 25% are predicted to be single top events.




Figure 5.41: The distribution of the leading jet p_T in the $t\bar{t}$ enriched region (left), the W +jets enriched region (middle), and the single top enriched region (right). See the text for the event selections. The first number in parenthesis after the process in the legend is the estimated event yield and the second number is the number of raw MC events used to make that prediction.




Figure 5.42: The distribution of $H_{T,sig}^{miss}$ in the $t\bar{t}$ enriched region (left), the W +jets enriched region (middle), and the single top enriched region (right). See the text for the event selections. The first number in parenthesis after the process in the legend is the estimated event yield and the second number is the number of raw MC events used to make that prediction.

Each of the kinematic variables in Fig. 5.41, 5.42, 5.43, and 5.44 are scanned to compute the dependence of the normalization factors μ on the variables. Figure 5.45




Figure 5.43: The distribution of E_T^{miss} in the $t\bar{t}$ enriched region (left), the W +jets enriched region (middle), and the single top enriched region (right). See the text for the event selections. The first number in parenthesis after the process in the legend is the estimated event yield and the second number is the number of raw MC events used to make that prediction




Figure 5.44: The distribution of the leading large-radius ($R = 1.2$) jet mass in the $t\bar{t}$ enriched region (left), the W +jets enriched region (middle), and the single top enriched region (right). See the text for the event selections. The first number in parenthesis after the process in the legend is the estimated event yield and the second number is the number of raw MC events used to make that prediction

shows the dependence on E_T^{miss} and $H_{T,\text{sig}}^{\text{miss}}$. The $t\bar{t}$ and W +jets normalization factors are relatively constant as a function of these variables while there is a decreasing trend for the single top normalization factors, albeit with significant statistical uncertainties due to the low yield and purity in the single top enriched region. Similar plots are shown in Fig. 5.46 for the leading jet p_T and the leading large-radius ($R = 1.2$) jet

mass. The $t\bar{t}$ and W +jets normalization factors are relatively constant as a function of the jet mass, but there is a significant decrease in the $t\bar{t}$ normalization factor as a function of the leading jet p_T . This is likely related to the well-known mis-modeling of the top quark p_T (see e.g. Ref. [281]). All of the control regions described in Sec. 5.1, 5.2, and 5.3.1 have the same jet p_T requirements as the signal region in order to remain largely insensitive to this mis-modeling.




Figure 5.45: The dependence of the normalization factors on E_T^{miss} (left) and $H_{T,\text{sig}}^{\text{miss}}$ (right). The error band represents the statistical uncertainty derived from bootstrapping the data in the enriched regions and resolving Eq. 5.3.

Both Fig. 5.45 and 5.46 show that the uncertainty on the single top normalization factor is much larger than the uncertainties for the $t\bar{t}$ and W +jets factors. This is due in part to the low event yield in the single top enriched region but also to the contamination of single top events in the $t\bar{t}$ enriched region and $t\bar{t}$ events in the single top enriched region. The top row of Fig. 5.47 shows the statistical correlations between the various normalization factors when using the enriched samples described earlier without any further requirements. The $t\bar{t}$ and W +jets normalization factors are largely uncorrelated, but the $t\bar{t}$ and single top factors are nearly 100% anti-correlated. Fixing the number of single top events in the single top enriched region, the middle




Figure 5.46: The dependence of the normalization factors on the leading jet p_T (left) and the leading large radius ($R = 1.2$) jet mass (right). The error band represents the statistical uncertainty derived from bootstrapping the data in the enriched regions and resolving Eq. 5.3.

and lower panels of Fig. 5.47 demonstrate the impact of reducing contamination in the $t\bar{t}$ and single top enriched regions. In the middle panel, the $t\bar{t}$ contribution to the single top enriched region is set to zero. This reduces the correlation between the $t\bar{t}$ and single top normalization factors and the overall single top normalization uncertainty by nearly a factor of two. However, the lower panel shows that the correlation between the normalization factors is not the only relevant quantity. When the single top contamination in the $t\bar{t}$ enriched region is set to zero, the single top and $t\bar{t}$ normalization factors are nearly uncorrelated but the uncertainty in the single top factor is only reduced by about 15%. The additional $t\bar{t}$ reduction in the single top control region by the $\Delta R(b, b)$ requirement (Sec. 5.3.1) is therefore a well-motivated technique for reducing the single top normalization factor uncertainty.

The results from the full control region fit used to extract the normalization factors from the control regions described earlier in this chapter are described in Chapter 7. First, Chapter 6 quantifies the accuracy of the transfer factors and MC background estimates with a full assessment of systematic uncertainties.




Figure 5.47: Statistical correlations between the various normalization factors determined from bootstrapping. The data yields are set to the prediction so that the scale factors are centered at $\mu = 1$. In the middle row, the $t\bar{t}$ contribution to the single top enriched region is set to zero and in the bottom row the single top contribution to the $t\bar{t}$ enriched region is set to zero.

Chapter 6

Systematic Uncertainties

The background estimation procedures described in Sec. 5 are only useful if the precision and accuracy are known. The precision is set by the various sources of statistical uncertainty, including the MC statistical uncertainty and the data statistical uncertainty in both the control and signal regions. The MC statistical uncertainty can be reduced by running larger simulations and the uncertainty from the finite control region statistics can be reduced by loosening the selection, at the cost of a larger extrapolation to the signal region. This chapter describes a variety of techniques that are used to estimate potential sources of systematic bias impacting the accuracy of the background estimates. For signal regions with harsh selections, the systematic uncertainties are largely subdominant to the uncertainty from the data statistical uncertainty. This is illustrated quantitatively in Sec. 6.1 with a simple one-bin region. As long as the systematic uncertainty is below the data (Poisson) statistical uncertainty, there is little impact on the sensitivity. However, for the looser signal regions selections (in particular for the shape fits), systematic uncertainties can have a significant impact on the sensitivity to stops.

There are two sources of potential bias: experimental systematic uncertainties and theoretical modeling uncertainties. The experimental uncertainties are related to the reconstructed object efficiencies and resolutions. They are constrained in dedicated auxiliary studies and then the impact on this search is estimated by varying some aspect of the simulation, such as the event weight or per-object kinematic quantities.




Figure 6.1: The normalized significance as a function of the number of background events and the fractional systematic uncertainty for a one-bin fit with log-normal uncertainties and no profiling. The ‘significance’ is computed as $\Phi^{-1}(1 - \text{CL}_s)$ for Φ the cumulative distribution function of the normal distribution. The number of signal events is set to the number of background events (10 events) in the left (right) plot and each column is normalized to be one in the first row. The dashed line represents the data statistical uncertainty: $1/\sqrt{\text{Background Events}}$.

Generally these uncertainties are constructed to cover differences between data and simulation in the auxiliary studies. In contrast, the theoretical modeling uncertainties are largely unconstrained by data since they are used to estimate the background predictions in extreme kinematic regimes. These regions of phase space by construction have little or no data and therefore the modeling uncertainties are constructed to cover all reasonable variations in the simulation.

This chapter is organized as follows. The experimental systematic uncertainties and their impact on the signal region predictions are described in Sec. 6.1. Section 6.2 documents the procedures for the theoretical modeling uncertainties, including several that are unique to this analysis. An overview of all the uncertainties for the various stages of the search is presented in Sec. 6.3.

6.1 Experimental Systematic Uncertainties

Each of the reconstructed objects have an associated uncertainty on their energy scale, energy resolution, and reconstruction efficiency. In addition, there is an uncertainty on the modeling of various global properties such as pileup and the instantaneous luminosity. Even though a diverse set of reconstructed objects is used to construct the signal regions, the total measured energy scale and energy resolution are dominated by jets. Furthermore, due to the complexity and the lack of a conceptually and experimentally clean resonance constraint for jets, the uncertainty on the jet energy scale and resolution are significantly larger than for other reconstructed objects. A comparison of the per-object systematic uncertainties is presented in Table 6.1. The jet energy resolution decreases with p_T because of the Poisson nature of the calorimeter energy resolution while the p_T resolution of track-based objects increases with p_T . The uncertainties on jet properties decrease with p_T due in part to the subdominance of pileup and other effects. In contrast, the uncertainties on electrons, photons, muons, and taus increases with p_T due to the limited availability of resonance decays in data. The resolution and uncertainty on the E_T^{miss} soft term can be significant, but the presence of neutrinos and real missing momentum renders the soft term largely irrelevant even for the more inclusive event selections presented in this analysis. The impact of the per-object and per-event uncertainties on the analysis are detailed in the following sections. Sections 6.1.1 and 6.1.2 cover the jet energy scale and resolution uncertainties. Uncertainties associated with b -tagging and the E_T^{miss} are described in sections 6.1.3 and 6.1.4. All other (minor) uncertainties are summarized in Sec. 6.1.5. An overview of the impact of the leading experimental systematic uncertainties is shown in Fig. 6.2, using SR13 as an example. Normalization factors are extracted using the same technique described in Sec. 5.6, expanded to a fourth equation (the $t\bar{t} + \gamma$ CR) and a fourth unknown ($\mu_{t\bar{t}+Z}$). The $t\bar{t} + Z$ and $t\bar{t} + \gamma$ normalization factors are set equal by construction and the diboson contribution is determined from simulation and subtracted from the data before solving the equations. The resulting normalization factors are the same as the simultaneous fit

described in Sec. 7.1. After multiplying each component by its respective normalization factor, the impact on the total background prediction is about 10%, indicated by the difference between the last and penultimate rows in Fig. 6.2. The statistical uncertainty from the control regions results in a 10% systematic uncertainty on the total background prediction (error band on the *CR stats* row). In general, the uncertainties are reduced using the control region constraint (black versus red error bands) and the dominant systematic uncertainties are from the jet energy scale and jet energy resolution uncertainty.

Object	σ	ϵ	μ uncertainty	σ uncertainty	ϵ uncertainty
Jets	20%/8% [122]	92%/100% [124, 546]	4%/1% [122]	2%/0.5% [122]	2%/0% [124]
b-jets [75]	—	60%/80%	—	—	5%/3%
Electrons	0.025%/0.012% [128]	90%/97% [551]	0.1%/0.3% [128]	5%/17% [128]	3%/0.5% [551]
Photons	0.02%/0.01% [128]	70%/93% [553]	0.2%/0.2% [128]	5%/20% [128]	3.5%/0.5% [553]
Muons [131]	1.5%/2%	99%/99%	0.06%/0.04%	4%/6%	0.1%/0.2%
Taus [550]	22%/6%	75%/80%	4%/4%	—	8%/5%
E_T^{miss} [138]	130%/85%	—	10%/20%	20%/40%	—

Table 6.1: Performance metrics for the various reconstructed objects and their systematic uncertainties at $p_T = 25$ GeV (upper number) and $p_T = 100$ GeV (lower number). The symbol σ denotes the energy resolution (width / mean), μ the energy scale, and ϵ the reconstruction efficiency. Many of the object reconstruction algorithms changed between Run 1 and Run 2; the selection in this table mixes the defaults between the two runs to give an idea of the overall performance. The jet reconstruction efficiency is based on the *loose* quality criteria [546] (and refs. therein) which is over 99.9% efficient and a jet vertex tagger (JVT) threshold of 0.59, which corresponds to a pileup jet efficiency of just over 1% [124]. The b-tagging efficiency is determined using the 70% working point of the MV1 algorithm [75]. Combined (CB) + segment tagged (ST) muons [131] are considered without any further quality criteria and the scale and resolution (uncertainties) are on $m_{\mu\mu}$ from various resonance decays. Electrons are reconstructed with the *loose* criteria. Unconverted *tight* photons [553] are used for photon reconstruction properties. All numbers refer to central objects only (except for E_T^{miss} , which uses objects across η). Medium one-prong tau reconstruction at $\sqrt{s} = 13$ TeV is used for illustration [550]. The last line indicates the properties of the component of the track-based soft term parallel to p_T^{hard} [138] (and refs. therein).




Figure 6.2: An overview of the impact of the experimental uncertainties on the predicted yield in SR13. The Raw MC prediction does not include any data constraints from the control regions. The vertical dashed line indicates the nominal SM prediction after normalizing the simulation to the data in each of the four control regions described in Chapter 5. The black error bands show the size of the uncertainties before the control region constraint and the red bands show the error after using the control region method. Sections 6.1.1, 6.1.2, 6.1.3, and 6.1.4 describe how the various uncertainties are applied. To make them visible on the plot, the **b**-tagging and E_T^{miss} soft term uncertainties are multiplied by 10. The data and CR stats error bands represent the symmetric 68% inter-quantile range centered about the median determined from bootstrapping the data.

6.1.1 Jet Energy Scale

The jet energy scale (JES) and its uncertainty at $\sqrt{s} = 8$ TeV are described in Sec. 1.5.2. A similar procedure is used to calibrate jets and determine the systematic uncertainty at $\sqrt{s} = 13$ TeV [624]. For Run 1, the default jet calibration is LCW+JES while in Run 2, the local cluster weighting is not used by default for $R = 0.4$ jets. At the beginning of Run 1, the EM+JES scheme had significantly larger uncertainties than the LCW+JES scheme at low jet p_T due to the dependence on the quark/gluon composition [121]. However, with the global sequential calibration (GSC) add-on to the calibration procedure, the flavor dependence is significantly reduced (in part because of n_{track} - see Chapter 3) and therefore the uncertainties for the EM+JES scheme in Run 2 are only slightly larger than those with the LCW+JES scheme [122]. The full JES uncertainty has many components that could each be included in the simultaneous background fit as an independent nuisance parameter. The $\mathcal{O}(10)$ parameters introduced in Sec. 1.5.2 is already a reduction from the full $\mathcal{O}(100)$ parameters due to all the auxiliary in-situ measurements [625]. In most of the parameter space probed by the stop search, the fit is not sensitive to the intricate details of the JES uncertainty correlations. Therefore, the Run 1 single-bin regions use a single nuisance parameter to capture the total jet energy scale uncertainty. This uncertainty is largely independent of η and is about 1% for $100 \text{ GeV} < p_T < 1 \text{ TeV}$. Below 100 GeV, the uncertainty grows to about 4% at 25 GeV and above 1 TeV, there is an increase in the uncertainty to about 3% due to a change in the uncertainty method (lack of statistics for an in-situ constraint). The early $\sqrt{s} = 13$ TeV analysis uses three nuisance parameters, which capture most of the relevant correlations [625] while the shape fit region at $\sqrt{s} = 8$ TeV uses the same 17-parameter setup as in in Sec. 1.5.2 because the fit is over-constrained (more bins than normalization parameters) leading to the potential for a reduction in the uncertainty from the fit (*profiling*) and the increased importance of the systematic uncertainties for the more inclusive selection. The impact of the JES uncertainty on the shape fit is revisited at the end of this section.

The JES uncertainty impacts the analysis directly through an uncertainty in the acceptance and indirectly by changing high level variables that depend on jet p_T .

Figure 6.3 shows the jet multiplicity after the preselection at $\sqrt{s} = 13$ TeV for $t\bar{t}$ events. The JES nuisance parameter with the largest variation is shown for illustration. There is a $\pm 10\%$ impact on the total number of events in the four-jet bin, which is the JES-induced uncertainty on the acceptance. The residual impact on the shape of the n_{jets} distribution (middle plot) is less pronounced. A similar trend is observed for the leading jet p_T in Fig. 6.4, where the JES uncertainty is at the percent-level for jets beyond the peak of the distribution at ~ 200 GeV. The JES uncertainty induces a jet mass scale uncertainty for large-radius reclustered jets. Figure 3.55 showed the reclustered jet mass in data after the preselection. The error band is dominated by the JES uncertainty, but seems strangely asymmetric. A similar trend is observed in the jet mass distribution in Fig. 2.20. This is a general trend for resonance peaks and the reason is illustrated by Fig. 6.5. The middle plot of Fig. 6.5 looks as expected for nearly symmetric variations in the JES: the peak position and width are decreased when the JES is reduced and vice versa when the JES is increased. Small bumps in the ratio of the right plot of Fig. 6.5 clearly illustrate the importance of these changes around the resonance peak. The reason for the asymmetric uncertainty band in earlier figures is because of the change in acceptance in addition to the change in the shape. When the JES is shifted down, there are also fewer events that pass the selection. Therefore, the red distribution in the middle plot is scaled down in the left plot of Fig. 6.5 and coincidentally is on top of the nominal histogram, which makes the overall uncertainty seem small. In contrast, when the JES is shifted up, there are more jets that pass the event selection and so the blue histogram in the middle plot is shifted up in the left plot of Fig. 6.5, leading to the large uncertainty only on the right side of the peak.

Figure 6.2 in the previous section showed that a significant fraction of the JES uncertainty cancels from the control region method because the shifts have a similar impact on acceptance in the CR and SR. This cancellation is demonstrated for the shape fit region in Fig. 6.6. This normalization reduces the uncertainties from 20-30% to 1-5% for most bins. However, even after the reduction, the JES uncertainty is comparable or larger than the data statistical uncertainty. Therefore, the fit is expected to be sensitive to the JES uncertainty and in particular can constrain it in




Figure 6.3: The signal jet multiplicity in $t\bar{t}$ events after the preselection before (left) and after (middle) normalizing the distributions to unity. The ratio of the JES $\pm 1\sigma$ to the nominal in the left plot is shown in the right plot.




Figure 6.4: The leading jet p_T in $t\bar{t}$ events after the preselection before (left) and after (middle) normalizing the distributions to unity. The ratio of the JES $\pm 1\sigma$ to the nominal in the left plot is shown in the right plot.

the analysis phase space. More details on this profiling are discussed in Sec. 7.2.




Figure 6.5: The large-radius ($R = 1.2$) jet mass in $t\bar{t}$ events after the preselection before (left) and after (middle) normalizing the distributions to unity. The ratio of the JES $\pm 1\sigma$ to the nominal in the left plot is shown in the right plot.




Figure 6.6: The size of the total JES uncertainty on each bin of the $t\bar{t}$ signal region before (left) and after (right) normalizing the total yields in each E_T^{miss} bin.

6.1.2 Jet Energy Resolution

Similarly to the jet energy scale, the jet energy resolution (JER) uncertainty was introduced already in Sec. 1.5.2. Due to the complexity in reducing the JER (see Sec. 4.1.2.3), the uncertainty is evaluated simply by inflating the resolution and then symmetrizing the effect on the analysis. In addition to changes in acceptance, increasing the JER generally broadens peaks (jet mass) and softens edges (m_T and m_{T2}). Figure 6.7 illustrates the broadening of the top quark mass peak in the jet mass spectrum and Fig. 6.8 shows the softening of the m_T edge when the jet energy resolution is increased. For high m_T , the increase in acceptance from an increase in the JER by 1σ is about 15-20%.




Figure 6.7: Left: The large-radius ($R = 1.2$) jet mass distribution in $t\bar{t}$ events after the preselection with the nominal JER and the JER inflated within its 1σ uncertainty. Right: the ratio of the two histograms in the left plot.

For all one-bin signal regions, a single JER nuisance parameter is used in the fit. To allow for more flexibility in the shape fit, there is one JER nuisance parameter per E_T^{miss} bin. Fig. 6.9 shows the change in each bin after increasing the JER by 1σ before and after normalizing per column. The reduction in the uncertainty is not as large as for the JES uncertainty because of the qualitatively different impact of JER in the CR-like regions and SR-like bins.




Figure 6.8: Left: The m_T distribution in $t\bar{t}$ events after the preselection with the nominal JER and the JER inflated within its 1σ uncertainty. Right: the ratio of the two histograms in the left plot.




Figure 6.9: The size of the total JER uncertainty on each bin of the tN shape signal region before (left) and after (right) normalizing the total yields in each E_T^{miss} bin.

6.1.3 **b**-tagging

Jet flavor tagging is used in two ways: directly in the event selection by requiring at least one **b**-tagged jet and indirectly by specifying which jets are used in higher level variable calculations such as $\mathbf{am}_{T2}$. Uncertainties related to the second use are highly suppressed with respect to the first. In order for a bias in the **b**-tagging to impact variable calculations, the relative ordering of jet **b**-tagging weights needs to be permuted. In contrast, an overall shift in the **b**-tagging weights¹ changes the acceptance but leaves the ordering unchanged. The multi-binned **b**-tagging setup from Sec. 4.2.4 and its uncertainty correctly account for permutations, but since this is a subleading effect for the stop search, only the overall changes in acceptance are considered for **b**-tagging uncertainties.

The uncertainty on the **b**-tagging efficiency is estimated by comparing data and simulation in auxiliary measurements as described in Sec. 4.2.6. As with the JES uncertainty, the **b**-tagging efficiency uncertainty has many components. The high mass stop search is not sensitive to the intricate correlation between the many nuisance parameters and therefore a reduced set is used. Figure 6.10 shows the uncertainty on the number of **b**-tagged jets by varying the **b**-jet² efficiency scale factors within their uncertainties. The JES uncertainty results in a much larger overall difference in acceptance, but has little impact on the normalized **b**-jet multiplicity distribution. A comparison of the various **b**-tagging efficiency scale factor uncertainties at $\sqrt{s} = 13$ TeV is shown in Fig. 6.11. For low **b**-tagged jet multiplicity, the uncertainty on the **b**-jet efficiency is the largest source of uncertainty. Since charm and light jets make a significant contribution to higher **b**-tagged jet multiplicities (only two **b**-quarks expected in $t\bar{t}$ at tree-level), the uncertainties in their **b**-tagging efficiencies are also important at higher multiplicities.

¹Practically, instead of shifting the weights, the uncertainty is estimated by varying the actual **b**-tagging efficiency. This is accomplished by applying event weights.

²This is an unfortunate but standard nomenclature: **b**-jets are jets originating from **b**-quarks, in contrast to **b**-tagged jets (often also called just **b**-jets), which are any jet that is tagged with a **b**-tagging algorithm.




Figure 6.10: The b -jet multiplicity in $t\bar{t}$ events after the preselection before (left) and after (middle) normalizing the distributions to unity. The ratio of the $\pm 1\sigma$ variations to the nominal in the left plot is shown in the right plot. The dashed lines in the right plot correspond to the JES uncertainty while the solid lines are the b -jet efficiency scale factor uncertainties.




Figure 6.11: The relative b -tagging efficiency scale factor uncertainty for various components of the uncertainty as a function of the b -tagged jet multiplicity. The ‘extrapolation from charm’ uncertainty is for τ jets that are b -tagged and the ‘extrapolation’ uncertainty is for extending the in-situ constraints to high p_T where the data statistics are limited.

6.1.4 Missing Transverse Momentum

The missing transverse momentum is a key input to many of the discriminating variables used in the stop search. As it is composed of all objects, the uncertainty in each object is coherently propagated to arrive at an uncertainty on the $\vec{p}_T^{\text{miss}}$ that is correctly correlated with the input object uncertainties. The one component that is not accounted for this way is the momentum not associated with any other reconstructed object (soft term). Uncertainties on the soft term are estimated using auxiliary studies with $Z \rightarrow \mu\mu$ events in simulation ($\sqrt{s} = 13$ TeV [138]) and with additional comparisons in data ($\sqrt{s} = 8$ TeV [136]). These uncertainties are parameterized based on $\vec{p}_T^{\text{hard}}$, which is the sum of all hard objects, including neutrinos (in simulation). An uncertainty is estimated on the scale of the E_T^{miss} soft term parallel to $\vec{p}_T^{\text{hard}}$ and on the resolution parallel and perpendicular to $\vec{p}_T^{\text{hard}}$. Figure 6.12 illustrates the size of the soft term scale uncertainty at $\sqrt{s} = 13$ TeV. For the high E_T^{miss} probed by the stop search, the impact of the soft term scale uncertainty is small compared with the impact of the JES uncertainty on the E_T^{miss} . The uncertainties on the soft term resolution are comparably small.




Figure 6.12: E_T^{miss} in $t\bar{t}$ events after the preselection before (left) and after (middle) normalizing the distributions to unity. The ratio of the $\pm 1\sigma$ variations to the nominal in the left plot is shown in the right plot. The dashed lines in the right plot correspond to the JES uncertainty while the solid lines are the E_T^{miss} soft term scale uncertainties.

6.1.5 Other

In addition to the E_T^{miss} soft term systematic uncertainties, there are a series of subdominant uncertainties related to all reconstructed objects as well as general event properties. The per object uncertainties were summarized in Table 6.1. The electron and muon related energy scale, resolution and reconstruction efficiency uncertainties are all precisely known from resonance decays. Scale factors from these auxiliary measurements are applied to correct the simulation and the uncertainty on the corrections is a source of systematic uncertainty. One purpose of the efficiency scale factors is to correct for any mis-modeling of the isolation criteria. Hadronic tau decays are also constrained from resonance decays (Z boson), but are inherently less clean than for electrons and muons and so the resolutions and uncertainties are generally larger. However, reconstructed taus are only used as a (highly efficient) veto in this analysis. Similarly, there are uncertainties associated with mis-modeling in the trigger, but these are suppressed by working in a regime where the trigger is nearly 100% efficient. The one exception is for the shape fit, where the lowest E_T^{miss} bin is in the trigger turn on region below $E_T^{\text{miss}} = 200$ GeV. This could in principle result in significant systematic uncertainties, but because of the control region method, most of the mis-modeling is absorbed into the normalization of the background at low m_T .

There are also uncertainties on more global properties such as the amount of pileup and the integrated luminosity. All of the techniques designed for the stop search have some corrections or other protection from the effects of pileup and therefore the impact on a mis-modeling of the pileup spectrum is expected to have a small effect. This is quantified by reweighting events so that the average number of interactions per bunch crossing varies by 10% ($\sqrt{s} = 8$ TeV) or 15% ($\sqrt{s} = 13$ TeV). The luminosity is precisely calibrated and measured using techniques described in Ref. [70]. The uncertainty in the luminosity for the analysis of 13 fb^{-1} of $\sqrt{s} = 8$ TeV is 3.6%, of the full 20.3 fb^{-1} of $\sqrt{s} = 8$ TeV is 2.8%, and of the first 3.2 fb^{-1} at $\sqrt{s} = 13$ TeV is 5%. Even though it has an experimental origin, the luminosity uncertainty is mostly relevant for the theory uncertainties as various background components are normalized by the $\sigma \times \int \mathcal{L} dt$. A detailed description of the theory modeling uncertainties is in the next section, Sec. 6.2.

6.2 Theoretical Modeling Uncertainties

In addition to the experimental uncertainties that impact all events, each SM process has an associated uncertainty because the background estimation is performed separately for all the processes. For the background processes that are estimated using the control region method, the theory modeling uncertainty is associated with the extrapolation from the control region to the signal region. All other backgrounds have an additional uncertainty on the inclusive cross-section. Unlike the experimental uncertainties, the theory modeling uncertainties are not usually determined from the difference between data and simulation in auxiliary measurements. By construction, the regions of phase space probed by the search have little or no overlap with previous measurements and so the uncertainties are derived entirely from a complete set of reasonable variations in the simulation. There is no unique way to compute these uncertainties. One way to build a reasonable set of uncertainties for a given process is to decompose the total uncertainty into categories that probe different aspects of the simulation. Such a decomposition might look like the one in Table 6.4. This decomposition provides a quantitative procedure that probes nearly³ independent sources of uncertainty arising from fixed order calculations and phenomenological models describing non-perturbative effects. Individual processes may have additional sources of theoretical modeling uncertainties. Sections 6.2.1, 6.2.2, 6.2.3, 6.2.4, and 6.2.5 describe the application of Table 6.4 as well as any additional uncertainties for the $t\bar{t}$, single top, $t\bar{t} + V$, $W + \text{jets}$, and dibosons processes, respectively.

Even though Table 6.4 describes a clear decomposition for evaluating the theory modeling uncertainties, there is no unique way to ascribe a ‘ 1σ uncertainty’ for each source. The general prescription is to take the difference in the predicted yield between the nominal sample N and a variation V and compute $\sigma = |N - V|$. When a particular process is normalized in a control region, only the difference in the yield after normalizing both samples in the control region is used for the uncertainty (transfer factor). When the procedure calls for a comparison between two samples V_1 and V_2 ,

³For example, varying the factorization and renormalization scales changes both the inclusive and differential cross section. However, the inclusive cross-section is usually known with much higher precision than the differential one.

Source	Procedure
Inclusive cross-section	Uncertainty of the most precise calculation
Parton momentum	PDF uncertainty/compare PDF sets, Vary factorization scale μ_f
Differential cross-section	Vary the renormalization scale μ_r
Merging Scheme (NLO)	Compare POWHEG and MC@NLO
Matching Scheme (LO)	Vary the CKKW or MLM parameters
Fragmentation Model	Compare PYTHIA and HERWIG
‘Extra’ Radiation (ISR/FSR/MPI)	Vary PS tune, vary h_{damp} (POWHEG-Box)

Table 6.2: A decomposition of theory modeling uncertainties into several categories.

neither of which is the nominal, the general strategy is to take $|V_1 - V_2|$ when they are both MC samples with nominal settings and $\frac{1}{2}|V_1 - V_2|$ when V_1 is an ‘up’ variation and V_2 is a ‘down’ variation of some simulation parameter. Furthermore, simulations are computationally expensive and need to be sufficiently large to make MC statistical uncertainties negligible. Therefore, most of the theoretical modeling uncertainties are evaluated at particle-level using particle-level event selections analogous to the detector-level ones. This may still not be sufficient to populate the extreme kinematic tails distributions near the signal regions and therefore an additional strategy is to compare two samples with a looser event selection and then extrapolate the difference to tighter selections.

6.2.1 Top Quark Pair Production

The uncertainty on the $t\bar{t}$ extrapolation from the control region to the signal region closely follows the prescription from Table 6.4. Table 6.3 summarizes the specific procedure for $t\bar{t}$ at both $\sqrt{s} = 8$ TeV and $\sqrt{s} = 13$ TeV. The most important uncertainties are those related to the number and spectra of the ‘extra’ jets that are not from the ME. The extrapolation in m_T changes the $t\bar{t}$ composition from a mostly one-lepton topology with ME-induced jets to a mostly two-lepton topology with extra jets from ISR/FSR. An uncertainty on the fragmentation model is estimated by comparing PYTHIA 6 and HERWIG(++), fixing POWHEG-BOX as the ME generator. This results in a $\sim 10\%$ uncertainty in the extrapolation from the CR to the SR. In addition, the amount of radiation within and around jets is varied using dedicated parton shower tune variations. At $\sqrt{s} = 8$ TeV, ACERMC +PYTHIA 6 with the AUET2B tune [626] was modified based on the measurement of radiation gaps in dilepton $t\bar{t}$ events at $\sqrt{s} = 7$ TeV [295]. PYTHIA 6 parameters related to the value of α_s used in generating ISR and FSR are varied to bracket the measurement [627]. At $\sqrt{s} = 13$ TeV, the PYTHIA 6 Perugia2012 tune variations radHi and radLo, which vary the shower α_s , are used in conjunction with simultaneous variations of the factorization and renormalization scales as well as h_{damp} . The combination (tune, μ_f , μ_r , h_{damp})=(radLo/Hi, $\times 2/0.5$, $\times 2/0.5$, $m_{\text{top}}/2m_{\text{top}}$) is based on the $\sqrt{s} = 7$ TeV gap fraction measurement as well as other $t\bar{t}$ properties measurements [628]. Related to the amount of radiation in the event is the interface between the NLO matrix element and the parton shower. At $\sqrt{s} = 8$ TeV, MC@NLO+HERWIG (NLO) and Alpgen+PYTHIA 6 (LO) with MLM matching were studied, but found to be significantly worse models of the data in inclusive event selections and therefore were not considered for the final uncertainty. For the early Run 2 analysis, the MC@NLO and POWHEG methods are compared, fixing HERWIG++ for fragmentation. Figure 6.13 shows a comparison between the nominal $t\bar{t}$ sample and MG5_aMC+HERWIG++ using a particle-level selection at $\sqrt{s} = 13$ TeV. After updating $h_{\text{damp}} \rightarrow m_{\text{top}}$ and swapping HERWIG with HERWIG++ ($\sqrt{s} = 8 \rightarrow 13$ TeV), both of these simulations are reasonable models of the data inclusively and the fact that they predict similar event yields gives confidence in the extrapolation from the CR to the SR. The

PDF4LHC procedure [504] is used to estimate the PDF uncertainty on the acceptance at $\sqrt{s} = 8$ TeV. As expected, this uncertainty is subdominant to others ($\mathcal{O}(1\%)$) as PDF variations only slightly change the energy and rapidity distributions. These uncertainties were ignored for the early $\sqrt{s} = 13$ TeV analysis. The total uncertainty on the extrapolation from the TCR to the SR is in the range 15%-25%.

Source	Procedure ($\sqrt{8}$ TeV)	Procedure ($\sqrt{13}$ TeV)
Inclusive cross-section	N/A (CR method)	N/A (CR method)
Parton momentum	PDF4LHC	Ignored
Differential cross-section	μ_f, μ_r by $\times \frac{1}{2}$ and $\times 2$	(see last row)
Merging / Matching	MC@NLO*, ALPGEN*	P+HERWIG++ v. M+HERWIG++
Fragmentation Model	P+PYTHIA 6 v. P+HERWIG	P+PYTHIA 6 v. P+HERWIG++
Amount of ‘Extra’ Radiation	AcerMC variations	μ_f, μ_r , PYTHIA 6 tune, h_{damp}

Table 6.3: A summary of the theoretical modeling uncertainties for $t\bar{t}$ at $\sqrt{s} = 8$ TeV and $\sqrt{s} = 13$ TeV. The * indicates that this uncertainty was studied but not applied because of worse agreement between the alternative generators and data in an inclusive event selection. P stands for POWHEG-BOX and M stands for MG5_aMC.




Figure 6.13: The particle-level E_T^{miss} (left) and the number of particle-level jets (right) for POWHEG+PYTHIA 6 and MG5_aMC+HERWIG++.

6.2.2 Single Top

As described in Sec. 5.3, the dominant single top process is the Wt -channel. Systematic uncertainties are estimated for this process and the other single top channels in a similar manner as for $t\bar{t}$, summarized in Table 6.4. The main difference is that without a control region at $\sqrt{s} = 8$, the search is sensitive to the inclusive cross section uncertainty. Furthermore, without any region enriched in single top events, there is no way to conclude that one model is better than another and so the MC@NLO+HERWIG that was not used for $t\bar{t}$ at $\sqrt{s} = 8$ TeV is used to set an uncertainty on the merging scheme for single top. This is a $\sim 10\%$ uncertainty on the yield in the signal region. The dominant uncertainty for single top is due to the modeling of the interference between single top and $t\bar{t}$ (30%), as described below.

Source	Procedure ($\sqrt{8}$ TeV)	Procedure ($\sqrt{13}$ TeV)
Inclusive cross-section	6.8%	N/A (CR method)
Parton momentum	PDF4LHC	Ignored
Differential cross-section	Ignored	(see penultimate row)
Merging / Matching	MC@NLO	Ignored
Fragmentation Model	P+PYTHIA 6 v. P+HERWIG	P+PYTHIA 6 v. P+HERWIG++
Amount of ‘Extra’ Radiation	AcerMC variations	μ_f, μ_r , PYTHIA 6 tune
Interference with $t\bar{t}$	DR/DS*, AcerMC	MG5_aMC

Table 6.4: A summary of the theoretical modeling uncertainties for single top at $\sqrt{s} = 8$ TeV and $\sqrt{s} = 13$ TeV. The * indicates that this uncertainty was studied but not applied (see the text). P stands for POWHEG-BOX.

At NLO accuracy, there is a non-trivial interference between leading order $t\bar{t}$ and NLO Wt with one real emission. For example, the reaction $gg \rightarrow t\bar{t}^* \rightarrow t\bar{t}bW^+$ contributes at LO to $t\bar{t}$ and at NLO to single top. When $m(bW^+) \sim m_t$ this interference is large. The interference is treated by removing a contribution from the Wt simulation using either the Diagram Removal (DR) or Diagram Subtraction (DS) [273] schemes and the difference can be an estimate of the uncertainty. Both schemes result in unphysical results that are either not gauge-invariant (DR) or not intended for exclusive event selections (DS). When the interference between NLO single top and LO $t\bar{t}$ is small, the two interference schemes are comparable [273, 532, 629]. However,

the stop search selects single top events with relatively high purity and such events have kinematic properties that result in a non-trivial interference with LO $t\bar{t}$. Using separated $t\bar{t}$ and Wt processes simulated at NLO may not be meaningful and the difference between the DR and DS interference removal schemes may not give an accurate estimate of the uncertainty on the modeling of the composite process.

Specific examples at particle level with $E_T^{\text{miss}} > 200$ GeV are shown in Fig. 6.15 and Fig. 6.16 for a single lepton and a dilepton selection, respectively. The two-lepton events are a useful complement to the single lepton ones because most Wt events in the SR have two leptonically decaying W bosons. To enhance the interference, events are required to have at least four jets with $p_T > 25$ GeV and at least two such jets must have originated from b -quarks. For both the single and double lepton selections, there is a significant and increasing difference between the single top simulations with the DR and DS schemes for the E_T^{miss} (Fig. 6.15) and the leading jet p_T distributions (Fig. 6.16). In both cases, the DR scheme results in both a harder E_T^{miss} and leading jet p_T spectrum. For this selection the Wt contribution is a small fraction (5–10%) of the total top quark contribution. However, in the E_T^{miss} tail for the dilepton events, the $\gtrsim 50\%$ difference between the single top DR and DS results in 15–20% differences in the combined $t\bar{t} + Wt$ simulation when comparing the two interference schemes. Even higher purities occur for tighter selections where the difference between the two setups can approach 100% and $\gtrsim 50\%$ overall uncertainties.

At $\sqrt{s} = 8$ TeV, the interference in the signal regions was further studied by looking at a LO sample generated by AcerMC [294] with the inclusive $2 \rightarrow 6$ reaction $pp \rightarrow W^+W^-b\bar{b}$ ($WWb\bar{b}$) that includes the double resonant $t\bar{t}$ production, the single resonant Wt production in association with a b -quark and the non-resonant diboson production in association with jets (see Fig. 6.14). The rest of this section presents an analogous study using MG5_aMC interfaced with Pythia 8 at $\sqrt{s} = 13$ TeV. Figures 6.17 and 6.18 include this inclusive $WWb\bar{b}$ sample in comparison with the POWHEG-BOX simulation using the DR scheme. There are significant differences in both the E_T^{miss} and leading jet p_T distributions. This is due in part to the fact that the $t\bar{t}$ component of the $WWb\bar{b}$ sample is LO and there are significant NLO corrections.




Figure 6.14: The inclusive $WWbb$ process includes Feynman diagrams with doubly resonant (left), singly resonant (middle) and non-resonant top quark contributions. Feynman diagrams from MG5_aMC.

Further studies of the full ME comparison will benefit from a multileg $WWbb$ simulation, which requires a non-trivial merging setup, and ultimately a full NLO $WWbb$ simulation interfaced with a parton shower. Both the Sherpa+OpenLoops [630] and MG5_aMC [631] collaborations have calculated fixed order inclusive calculations, but there is currently no general matrix element at NLO matched to a parton shower. It is not even currently possible with the existing frameworks to compute $WWbb$ with extra partons in the matrix element because of the non-trivial overlap in the phase space between the b -quarks at NLO for the single top process and quarks and gluons from the parton shower. The fact that the stop search is so sensitive to the interference between the processes and that a single top control region can be constructed with relatively high purity suggests that the data can be used to directly constrain the existing and future models of higher order interference.

Despite the disadvantages of comparing the leading order simulation with the NLO setup, it may produce a conservative uncertainty and therefore is used as the baseline method⁴. The uncertainty at $\sqrt{s} = 8$ and 13 TeV from comparing the $WWbb$ samples with the NLO $t\bar{t} + Wt$ (DR scheme) simulations using POWHEG-BOX result in $\sim 30\%$ uncertainties on the modeling of the interference.

⁴This is not fully satisfactory and will be an important topic of future study. Fortunately, the MC community has made significant progress very recently - see Ref. [632].




Figure 6.15: The particle-level E_T^{miss} for events passing a one lepton (a) and a two lepton (b) selection. Both selections require at least four jets with $p_T > 25$ GeV, of which two must have originated from b -quarks. All distributions are normalized to have the same integral in the above range. The gray band in the ratio is the statistical uncertainty from the simulation using the DR scheme and the uncertainty on the markers is from the simulation used in the numerator of the ratio. Most of these uncertainties are smaller than the markers.




Figure 6.16: Same as Fig. 6.15 but for the leading jet p_T instead of the E_T^{miss} .




Figure 6.17: Same as Fig. 6.15, but comparing the NLO calculation with interference removal with an inclusive WWbb sample generated with MG5_aMC.




Figure 6.18: Same as Fig. 6.17 but for the leading jet p_T instead of the E_T^{miss} .

6.2.3 Top Quark Pair Production with a Z Boson

Unlike the $t\bar{t}$ and Wt backgrounds, the $t\bar{t} + Z(\rightarrow \nu\bar{\nu})$ background does not need a second lepton in order to pass the event selection. Therefore, the four jets used in all signal regions already exist at tree-level and so the sensitivity to the modeling of additional radiation is significantly reduced. Table 6.5 summarizes the procedure for both the simulation-based approach at $\sqrt{s} = 8$ TeV and the data-driven approach at $\sqrt{s} = 13$ TeV. In the simulation-based case, the uncertainty is directly evaluated on the predicted yield in the signal region (Sec. 6.2.3.1) while the data-driven uncertainties are on the relative yield between the CR and SR and between $t\bar{t} + Z$ and $t\bar{t} + \gamma$ (Sec. 6.2.3.2).

Source	Procedure ($\sqrt{s} = 8$ TeV)	Procedure ($\sqrt{s} = 13$ TeV)
Inclusive cross-section	22%	N/A (CR method)
Parton momentum	PDF4LHC	CT14 and NNPDF3.0
Differential cross-section	μ_f, μ_r	μ_f, μ_r
Merging / Matching	Matching scale, n_{partons}	SHERPA and MG5_aMC (both MEPS)
Fragmentation Model	Ignored	Ignored
Amount of ‘Extra’ Radiation	ISR/FSR variations	Ignored

Table 6.5: A summary of the theoretical modeling uncertainties for $t\bar{t} + Z$ at $\sqrt{s} = 8$ TeV and $\sqrt{s} = 13$ TeV.

6.2.3.1 Simulation-based

Without a CR, there is an overall cross-section uncertainty of 22% [542, 543]. The default $t\bar{t} + V$ simulation at $\sqrt{s} = 8$ TeV was produced by MADGRAPH 5 with up to two extra partons in the matrix element. To assess the impact of the extra partons, an additional sample was generated with only one extra parton in the matrix element. In principle, the nominal two-parton setup should be more accurate, but with no data to constrain the modeling, the difference between the two samples is taken as a crude and likely conservative uncertainty. A similar probe of the extra radiation is from varying the MLM matching scale that connects MADGRAPH 5 and PYTHIA 6. The uncertainty is estimated by changing the scale (`xqcut`) between 15

and 25 GeV (20 GeV is nominal). Additionally, the amount of ISR and FSR can be varied using the renormalization scale of the α_s used for both processes coherently in MADGRAPH 5 and PYTHIA 6. The ISR variation scales `alpsfact` (MADGRAPH 5) and `PARP(64)` (PYTHIA 6) by factors of 2 and 4, respectively. FSR from PYTHIA 6 is varied by changing the value of Λ_{QCD} used in the running α_s (`PARP(72)`) from 0.2635 GeV to 0.7905 GeV (0.527 GeV is nominal). In addition, the infrared cutoff for FSR (`PARJ(82)`) is simultaneously varied between 0.5 GeV and 1.66 GeV (0.83 GeV is nominal). Figure 6.21 illustrates the differences in the distributions of two key kinematic variables when using one or two extra partons in the matrix element and Fig. 6.20 summarizes the slopes from fitting the ratios of all the above variations for $\text{am}_{\text{T}2}$, $E_{\text{T}}^{\text{miss}}$ and m_{T} . The corresponding uncertainties are largely statistically consistent with zero, except for the conservative finite partons slope. A combination of an uncertainty at preselection and an extrapolation into a single-region like selection for all three variables results in a $\sim 20\%$ uncertainty on the $\text{t}\bar{\text{t}} + \text{V}$ yield in addition to the inclusive cross-section uncertainty.




Figure 6.19: The $\text{am}_{\text{T}2}$ (left) and $E_{\text{T}}^{\text{miss}}$ (right) distributions when using one or two extra partons in the ME for $\text{t}\bar{\text{t}} + \text{Z}$ with MADGRAPH 5 + PYTHIA 6.




Figure 6.20: The slope parameter from a linear fit to the ratio of the variations listed on the horizontal axis. For ‘finite partons’, the comparison is between one and two extra partons in the ME; for ‘FSR’, PYTHIA 6 parameters are varied (see the text), for ‘Scales’, the μ_f and μ_r are simultaneously doubled and halved; for α_s , ISR parameters in MADGRAPH 5 and PYTHIA 6 are varied (see the text); for ‘Matching’, the MLM matching scale is varied. The error bars are due to MC statistical uncertainty. The units of the slope parameter are $1/\text{GeV}$.

6.2.3.2 Data-driven

At $\sqrt{s} = 13 \text{ TeV}$, the nominal background estimation method for $t\bar{t} + V$ uses a $t\bar{t} + \gamma$ control region, as described in Sec. 5.4.2. The total systematic uncertainty on the transfer factor from the $t\bar{t} + \gamma$ CR to $t\bar{t} + Z$ in the SR is 12% and consists of⁵:

1. A 10% systematic due to coherent factorization and renormalization scale variations as a function of boson p_T for the LO samples. This value is based on Fig. 6.21, which shows that the 10% in the double ratio $(\mu_{\text{up}}^{t\bar{t}+\gamma}/\mu_{\text{down}}^{t\bar{t}+\gamma})/(\mu_{\text{up}}^{t\bar{t}+Z}/\mu_{\text{down}}^{t\bar{t}+Z})$ is relatively independent of the E_T^{miss} for $E_T^{\text{miss}} \gtrsim 300 \text{ GeV}$.
2. A 5% systematic due to variation of the k-factor ratio resulting from scale variations. There is no uncertainty in the absolute cross-section (and thus k-factor)

⁵All of the calculations in this section are based on fixed order results without a parton shower. In the case of MG5_aMC, it was checked that the addition of PYTHIA 8 does not have a significant impact on the reported cross-section differences.

because of the CR normalization, but there is an uncertainty in the difference in the higher order corrections between the $t\bar{t}+\gamma$ and $t\bar{t}+Z$ processes. Figure 6.23 shows the k-factor ratio as a function of boson p_T for various NLO matrix element generator, PDF, and scale choices⁶. Fixing SHERPA+OPENLOOPS as the matrix element generator and either NNPDF3.0 or CT14 as the PDF set, the variation in the k-factor ratio (up versus down triangles in Fig. 6.23) is about 5% when the scale is varied by a factor of two. The default scale is the sum of the transverse mass $m_T^2 = m^2 + p_T^2$ of all out-going particles. Note that these scale variations are the LO to NLO coherent scale variations so partially related with the first bullet but not fully correlated.

3. A 5% systematic to cover the differences in k-factor ratios between Sherpa/OpenLoops and Madgraph/aMC@NLO as shown in Fig. 6.23. One difference between the setups is the electroweak parameter scheme, i.e. which parameters are taken as input and which are calculated to a fixed order internally.
4. $\sim 1\text{-}2\%$ to cover the differences in k-factor ratios between different PDF sets for a fixed matrix element generator and scale choice as in Fig. 6.23.

⁶Thank you to Stefan Hoche for useful discussions about these uncertainties and for providing the SHERPA+OPENLOOPS numbers.




Figure 6.21: Variations of the leading order renormalization and factorization Mad-Graph scales coherently by factors of two and one-half. The default scale is the transverse mass (see Table 13 in Ref. [623]). The $t\bar{t} + \gamma$ sample has a ME cut of 100 GeV. The $\tilde{E}_T^{\text{miss}}$ is the regular E_T^{miss} in the $t\bar{t} + V$ case and the magnitude of the vector sum of the missing momentum and the photon momentum. The relative normalization of the $t\bar{t} + V$ and $t\bar{t} + \gamma$ samples is arbitrary. The upper ratio compares the scale up with the scale down for the $t\bar{t} + V$ and $t\bar{t} + \gamma$ separately. The lower panel is the ratio of the ratios: $r = \mu_{\text{up}}/\mu_{\text{down}}$.




Figure 6.22: Left: The k-factor for $t\bar{t} + \gamma$ (left) and $t\bar{t} + Z$ (right) as a function of the boson p_T for SHERPA+OPENLOOPS and MG5_aMC. The PDF and scale choice are given in the legend.




Figure 6.23: The k-factor ratio between $t\bar{t} + \gamma$ and $t\bar{t} + Z$ as a function of the boson p_T for SHERPA+ OPENLOOPS and MG5_aMC. The PDF and scale choice are given in the legend.

6.2.4 W +jets

Unlike all previous samples, both the diboson and W +jets backgrounds are estimated using SHERPA as the nominal MC generator. Both of these processes require many extra hard jets to pass the event selection and only MADGRAPH, SHERPA, and ALPGEN have this capability. Both SHERPA and MG5_aMC can model extra jets at NLO and the SHERPA+OPENLOOPS setup is used at $\sqrt{s} = 13$ TeV for two extra partons at NLO and four partons at LO. Table 6.6 summarizes the systematic uncertainties for W +jets at both $\sqrt{s} = 8$ and 13 TeV. At both energies, one of the main systematic uncertainties is from scale variations to probe shape differences that could change the extrapolation from the CR to the SR. The modeling of the extra radiation is probed at $\sqrt{s} = 8$ TeV by varying the number of partons in the matrix element (see Sec. 6.2.3.1). At $\sqrt{s} = 13$ TeV, the extra radiation is varied by changing the SHERPA resummation scale `QSF` as well as comparing the SHERPA sample with a simulation from MG5_aMC+PYTHIA 8 with up to four extra partons in the matrix element. This comparison simultaneously varies the ME setup and the parton shower⁷.

In addition to the standard comparisons listed above, there is an additional source of uncertainty due to the extrapolation from a mostly no b -jet region in the W +jets CR (with a b -tag veto) to a mostly b - and c -jet selection in the SR. Section 5.2.1 showed that there is little flavor dependence on the m_T extrapolation, but nonetheless it is important to estimate the extrapolation in flavor on the yield in the SR. A 25% uncertainty from the ATLAS $W + b\bar{b}$ cross-section measurement [633] is combined with a $\sim 15\%$ uncertainty from ALPGEN 2.14 [634] parameter variations⁸ on the extrapolation from two jets to the four jets as required by all signal regions. This prescription is certainly conservative, as it includes a total cross-section uncertainty that should be canceled by the control region normalization.

⁷Ideally, these variations would be decomposed, but it is not possible to vary only the PS within the SHERPA framework.

⁸ALPGEN is combined with HERWIG and the factorization, renormalization, matching scales are varied. Additionally, the minimum ΔR between and minimum p_T of partons are varied. Strictly speaking these uncertainties only apply for ALPGEN, but because the jets beyond the leading two in SHERPA are also at LO, these uncertainties may be a useful proxy for the SHERPA sample as well.

Source	Procedure ($\sqrt{8}$ TeV)	Procedure ($\sqrt{13}$ TeV)
Inclusive cross-section	N/A (CR method)	N/A (CR method)
Parton momentum	PDF4LHC	Ignored
Differential cross-section	μ_f, μ_r	μ_f, μ_r
Merging / Matching	n_{partons}	Matching Scale; MG5_aMC
Fragmentation Model	Ignored	SHERPA and MG5_aMC+PYTHIA 8
Amount of ‘Extra’ Radiation	Ignored	Resummation scale (QSF)
$W+HF$	28% to $W + bb$	28% to $W + bb$

Table 6.6: A summary of the theoretical modeling uncertainties for W +jets at $\sqrt{s} = 8$ TeV and $\sqrt{s} = 13$ TeV.

6.2.5 Dibosons

The uncertainties for dibosons at $\sqrt{s} = 8$ TeV are similar to the analogous W +jets ones from Sec. 6.2.4. Due to the sub-dominance of dibosons for the $\sqrt{s} = 13$ TeV SR, a crude and likely conservative approach compares the SHERPA sample with up to three extra partons in the ME to a POWHEG-BOX+PYTHIA 8 sample with no extra partons in the matrix element. This is a simultaneous variation of the matrix element calculation and the parton shower. Figure 6.24 shows that the leading jet p_T spectrum is similar for the two generators, but the number of jets and the p_T spectrum of the subleading jets are significantly different. The extra jet activity is likely significantly underestimated by the POWHEG-BOX sample, but serves as a crude and likely conservative approach. Table 6.7 summarizes the diboson uncertainties.

Source	Procedure ($\sqrt{8}$ TeV)	Procedure ($\sqrt{13}$ TeV)
Inclusive cross-section	5-7%	6%
Parton momentum	PDF4LHC	Ignored
Differential cross-section	μ_f, μ_r	Ignored
Merging / Matching	Matching scale, n_{partons}	POWHEG-BOX
Fragmentation Model	Ignored	SHERPA and PYTHIA 8
Amount of ‘Extra’ Radiation	Ignored	Ignored

Table 6.7: Theoretical modeling uncertainties for dibosons.




Figure 6.24: Various comparisons of jet related quantities between SHERPA and POWHEG-BOX+PYTHIA 8 for the process $ll\nu\nu$. There is a $\sim 50\%$ difference between the predicted yields at preselection and an additional $\sim 20\%$ from extrapolating between the preselection and the SR.

6.3 Summary

Table 6.8 presents a summary of the uncertainties for the signal regions with the full $\sqrt{s} = 8$ TeV data and the early $\sqrt{s} = 13$ TeV data. These uncertainties are shown after the control region method is applied, so the any coherent uncertainties between the CR and SR are eliminated. The data statistical uncertainty dominates over the systematic uncertainties for the single bin regions, while the systematic uncertainty is much larger than the statistical uncertainty for the more inclusive shape fit signal region. The JES and JER uncertainties are the largest experimental uncertainties. For the single bin regions that have higher m_T thresholds and are thus more sensitive to the resolution tail beyond $m_T = m_W$, the JER is a bigger uncertainty than the JES. At $\sqrt{s} = 8$ TeV, there are only three components to the b -tagging efficiency uncertainty, while at $\sqrt{s} = 13$ TeV, two additional components describe various extrapolation uncertainties (see Sec. 6.1.3). The b -jet component of the b -tagging efficiency uncertainty is about 2% in all regions and the other components are only relevant for the single bin regions. This is due in part to the presence of c -jets which can allow events to exceed stringent m_{T2} requirements. The luminosity uncertainty is small because most of the backgrounds are normalized using control regions; this is especially true for tN13 (only dibosons are directly from MC) and the shape fit region for which the backgrounds other than $t\bar{t}$ are small. This also explains why the uncertainty from single top and other backgrounds is small for the shape fit region. For example, the interference between Wt and $t\bar{t}$ results in a 30% uncertainty for the $\sqrt{s} = 8$ TeV analyses, but due to the relatively small fraction of single top events in the SR, the total impact of this uncertainty is only a few percent at most. Inclusive cross-section uncertainties are only relevant when the background process is not normalized in a control region. The total systematic uncertainty is between 10% and about 20% in all regions.

Type	Source	tN11	tN12	tN21	tN22	tNmed	tNhigh	tN13
Experimental	JES (leading)	13%	10%	12%	12%	8%	6%	6%
	JES (subleading)	9%	7%	8%	8%	N/A	N/A	3%
	JER	9%	9%	8%	6%	11%	11%	13%
	b-tagging (b-jets)	2%	2%	2%	2%	2%	2%	2%
	b-tagging (c-jets)	–	–	–	–	1%	2%	2%
	b-tagging (light-jets)	–	–	–	–	1%	1%	1%
	b-tagging (other)	N/A	N/A	N/A	N/A	N/A	N/A	2%
	E_T^{miss} soft-scale	5%	3%	5%	2%	1%	2%	1%
	E_T^{miss} soft-resolution	1%	–	–	–	1%	2%	1%
	Other Experimental	2%	2%	2%	1%	1%	1%	1%
	Luminosity	–	–	–	–	1%	1%	–
Theoretical	$t\bar{t}$ Fragmentation	3%	2%	1%	–	1%	3%	2%
	$t\bar{t}$ Extra Radiation	6%	2%	1%	5%	2%	4%	5%
	$t\bar{t}$ Hard-scatter	1%	1%	2%	3%	1%	1%	2%
	Wt Cross-section	–	–	–	–	1%	–	N/A
	Wt Fragmentation	–	–	–	–	–	–	1%
	Wt Extra Radiation	–	–	–	–	–	–	–
	Wt Hard-scatter	–	–	1%	–	1%	1%	–
	Wt/ $t\bar{t}$ Interference	1%	–	2%	1%	2%	3%	3%
	W+jets Modeling	1%	1%	1%	–	1%	2%	3%
	W+HF	1%	–	–	–	2%	2%	5%
	$t\bar{t}$ + V Modeling	–	–	–	1%	3%	3%	6%
	$t\bar{t}$ + V Cross-section	–	–	–	1%	3%	3%	N/A
	VV Total	–	–	–	–	3%	2%	5%
Total Systematic Uncertainty		16%-	13%	13%	12%	17%	19%	21%
Data Statistical Uncertainty		9%	8%	10%	7%	28%	45%	90%

Table 6.8: A summary of the uncertainty in the total signal region yield after the background-only fit from the control region method (see Sec. 7.1). If an uncertainty is less than 1%, it is marked with ‘–’ while uncertainties that are not applicable are labeled N/A. PDF uncertainties are included in the HS uncertainty. tNxy is the $(x+1)^{\text{th}}$ E_T^{miss} bin and $(y+2)^{\text{th}}$ m_T bin of the shape fit. When there is more than one JES nuisance parameter, there are at least three, but only the two biggest ones are shown here. Due to correlations in the uncertainties after the fit, the total systematic uncertainty is not the sum in quadrature of the individual uncertainties.

Chapter 7

Search Results

Unfortunately, despite extensive efforts to search for stops with the $\sqrt{s} = 8$ and early $\sqrt{s} = 13$ TeV datasets, there is no significant evidence for a deviation from the Standard Model. Section 7.1 describes the statistical framework used to quantify the compatibility with the SM and to set limits on models of SUSY. The statistical fit from Sec. 7.1 is exercised in Sec. 7.2 using only the control regions and the fidelity of the predictions are tested in validation regions in Sec. 7.3. Limits on stop models are described in Sec. 7.4 for each of the signal regions, including the evolution of sensitivity with more data and technique improvements. The chapter and Part III ends in Sec. 7.5 with a broad overview of all ATLAS and CMS Run 1 SUSY searches

7.1 Statistical Methods

Formally, the statistical analysis of the search results is a hypothesis test with the null hypothesis $H_1 = \mathbf{SM\ only}$ and the alternative hypothesis $H_0 = \mathbf{SM+stop}$. A given signal model is excluded if the corresponding null hypothesis is rejected¹. By the Neyman-Pearson lemma [582], for a fixed upper bound on the probability of rejecting the null hypothesis when it is true (type 1 error), the likelihood ratio test minimizes the probability of not rejecting the null when the alternative is true (type II error) i.e. maximizes the probability of rejecting the SM-only hypothesis when there is SUSY. The likelihood function is given by

$$\begin{aligned} L(\mathbf{v}, \vec{\theta}, \vec{\mu}) &:= p(\vec{n}, \vec{\theta}_0 | \mathbf{v}, \vec{\mu}, \vec{s}(\vec{\theta}), \vec{b}(\vec{\theta})) = \prod_{i=1}^{n_{\text{bins}}} \text{Pr}(n_i | \mathbf{v}, \vec{\mu}, \vec{s}(\vec{\theta}), \vec{b}(\vec{\theta})) \times p(\vec{\theta} | \vec{\theta}_0) \\ &= \prod_{i=1}^{n_{\text{bins}}} \frac{\left(\mathbf{v} s_i(\vec{\theta}) + \sum_{k=1}^{n_{\text{backs}}} \mu_k b_{ki}(\vec{\theta}) \right)^{n_i}}{n_i!} e^{-\left(\mathbf{v} s_i(\vec{\theta}) + \sum_{k=1}^{n_{\text{backs}}} \mu_k b_{ki}(\vec{\theta}) \right)} \times p(\vec{\theta} | \vec{\theta}_0), \end{aligned} \quad (7.1)$$

where there are n_{bins} total SR and CR bins with MC predictions for s_i signal events and b_{ki} background events of the k^{th} SM background process in bin i . The values μ_k are the normalization factors. For the processes without a data-driven background estimate, μ_k is fixed to unity. The values θ_i are all of the nuisance parameters associated with each systematic uncertainty; the input values of these uncertainties are given by $\vec{\theta}_0$. The last term $p(\vec{\theta} | \vec{\theta}_0)$ is the constraint on the nuisance parameters. The nuisance parameters are constructed so that they are mostly independent and therefore $p(\vec{\theta} | \vec{\theta}_0)$ factorizes for each parameter θ_i ². For all theoretical modeling uncertainties and all systematic uncertainties for the shape fit region, $\theta_i | \theta_{0,i}$ follows a standard normal distribution. The experimental systematic uncertainties in the single-bin regions and the

¹The setup is different when optimizing the sensitivity of the test to *discover* SUSY; in that case the null hypothesis is the SM only case. The focus of Chapter 7 will be on the exclusion of signal models given the lack of a significant excess in any signal region.

²Note that even though the input nuisance parameters are independent, the output θ_i can be correlated given the data.

dedicated signal model uncertainties in all regions are modeled with a standard log-normal distribution, $\theta_i|\theta_{0,i} \sim \exp(\mathcal{N}(0, 1))$. The impact of the nuisance parameters on the yield y (such as b_{ki} or s_i) is given by $y = y_0(1 + \sum_i \theta_i(H(\theta_i)\sigma_i^+ + (1 - H(\theta_i))\sigma_i^-))$, where σ_i is the fractional uncertainty on the yield for systematic uncertainty source i and H is the *Heaviside step function*. The purpose of H is to allow the impact of ‘up’ (σ^+) and ‘down’ (σ^-) shifts of the nuisance parameter to asymmetric effects on the yield. When $\sigma_i^+ = \sigma_i^- = \sigma_i$, the contribution to the yield is simply $\theta_i\sigma_i$ as $H(x) + (1 - H(x)) = 1$. The *parameter of interest* in Eq. 7.1 is ν , which is 1 under H_0 and $\nu = 0$ under H_1 . The *test statistic* used to perform the hypothesis test is the log of the *profile likelihood ratio*:

$$t = -2 \ln \left(\frac{\max_{\vec{\mu}, \vec{\theta}} L(1, \vec{\mu}, \vec{\theta})}{\max_{\nu', \vec{\mu}', \vec{\theta}'} L(\nu', \vec{\mu}', \vec{\theta}')} \right). \quad (7.2)$$

Note that since t is not a monotonic transformation of the likelihood ratio, its type II error is not optimal even in the absence of nuisance parameters³. However, the value of ν that maximizes $L(\nu, \vec{\mu}, \vec{\theta})$ will be close to zero (no evidence for SUSY) and therefore t is close to optimal. Near-optimality is also true in general in the asymptotic limit of large event yields due to a result by A. Wald [635, 636]. Profiling refers to the maximization of the likelihood over the nuisance parameters in Eq. 7.2. The maximized (‘fitted’) values of the θ_i and their post-fit uncertainty can deviate from zero and unity, respectively. When this is significant, the relevant nuisance parameter is said to be ‘profiled’. Profiling is revisited in Sec. 7.2. The calculation of t for the search results presented in Sec. 7.4 are implemented using HistFitter [637] based on RooStats [595], RooFit [638], and ROOT [639] through HistFactory [640].

In addition to the likelihood in Eq. 7.1, a useful related quantity is the CR-only likelihood that is identical to Eq. 7.1, but with $\nu = 0$ and the signal regions removed from the product:

³Interestingly, even though the log profile likelihood ratio is standard for the LHC collaborations, the Tevatron collaborations used the log ratio of the profile likelihoods, which is optimal in the absence of nuisance parameters.

$$L_{\text{CR-only}}(\vec{\theta}, \vec{\mu}) = \prod_{i=1}^{n_{\text{CR bins}}} \text{Pr}(n_i | \vec{\mu}, \vec{b}(\vec{\theta})) \times p(\vec{\theta} | \vec{\theta}_0). \quad (7.3)$$

The CR-only fit referred to in several places in earlier sections is simply $(\vec{\mu}, \vec{\theta}) = \text{argmax}_{\vec{\mu}', \vec{\theta}'} L_{\text{CR-only}}(\vec{\theta}', \vec{\mu}')$. Since $\vec{\mu}$ is not directly constrained by a PDF in the likelihood, when the number of CR bins is equal to the number of normalization factors, the CR fit simply returns values of μ_i that solve the system of equations (or a subset/superset if there are fewer/more data-driven regions) in Eq. 5.3 and $\theta_i = 0$.

The distribution of the test-statistic t can be estimated numerically by sampling from the distributions of the input stochastic variables or with asymptotic formulae [636]. Due to its computational simplicity and accuracy, the asymptotic approximation is used as default and a few signal models are checked with the full numeric approach. The formula is based on the observation that in the asymptotic regime, the log likelihood approaches a (non-central) chi-square distribution [635], for which the p -value can be readily computed:

$$p\text{-value} = 1 - \Phi^{-1} \left(\sqrt{t_{\text{asymptotic}}} - \frac{1 - \nu}{\sigma_{\text{asymptotic}}} \right), \quad (7.4)$$

where Φ is the standard normal cumulative distribution function and $t_{\text{asymptotic}}$ is the value of t when the number of events in each bin is $\nu s_i(\vec{\theta}_0) + \sum_{k=1}^{n_{\text{backs}}} \hat{\mu}_k b_{ki}(\vec{\theta}_0)$ and $\sigma_{\text{asymptotic}}^2 = (1 - \nu)^2 / t_{\text{asymptotic}}$. The parameter $\hat{\mu}$ is the value of μ that maximizes L given ν and $\vec{\theta}_0$; with these values, $\max_{\nu', \vec{\mu}', \vec{\theta}'} L(\nu', \vec{\mu}', \vec{\theta}') = L(\nu, \hat{\mu}, \vec{\theta}_0)$. Amazingly, even though the asymptotic convergence is only $\mathcal{O}(1/\sqrt{N})$, the approximation in Eq. 7.4 well-approximates the full value even when the number of events N is $\gtrsim 10$. Figure 7.1 illustrates the convergence for various values of N using one signal region with one background process and one systematic uncertainty. When there are 3 background events and 2 signal events, the asymptotic formula overestimates the p -value at the measured t -value by 15-20%, while when there are 12 background events and 5 signal events, the p -value under H_1 is only off by 3% and the p -value under H_0

is over-estimated by about 10%.




Figure 7.1: A comparison of the p-value under H_1 (CL_b) and under H_0 (CL_{s+b}) computed using numeric methods (10,000 toys) and the asymptotic formula given in Eq. 7.4. The vertical dashed line indicates the observed value of t and the dashed line is at the p-value corresponding to the observed t -value using the asymptotic formula.

A given signal model is excluded if the p-value from the test described above is sufficiently small. One undesirable feature of the p-value under H_0 which is a

general property of two-model hypothesis testing is that the value can be small even if the data are inconsistent with both the SM and SM-only hypotheses. For example, suppose that for a one-bin signal region there are $M \gg 1$ predicted background events and $\epsilon \ll M$ predicted signal events. The p -value under H_0 when there are $N \ll M$ observed events will be small, but the p -value under H_1 will also be small. This is a general feature when comparing two models that do not partition the space of all possible models and in particular when the null and alternative hypotheses are similar. In high energy physics, the standard solution to this problem is to regulate the p -value under the null hypothesis (CL_{s+b}) by dividing by the p -value under the alternative hypothesis (CL_b) to form the $CL_s = CL_{s+b}/CL_b$ value [641, 642]. This new quantity has the property that it will be large in the example described above, i.e. when both the null and alternative hypothesis are inconsistent with the data. The community standard is to treat CL_s as if it were a proper p -value by declaring a model excluded when $CL_s < 0.05$. However, it should be noted that the CL_s is not a p -value and is not unique. Any function $f(x)$ that has the property $\lim_{x \rightarrow 0} f(x) = 0$ will be able to regulate the CL_{s+b} by $CL_{s+b}/f(CL_b)$. One simple function is

$$f_r(x) = \begin{cases} x & x \leq r \\ 1 & x > r \end{cases}, \quad (7.5)$$

where $0 \leq r \leq 1$ is a fixed value. One natural choice is $r = 0.5$. When CL_b is small, this *regulated* CL_s is enlarged just like the usual CL_s ⁴. However, when $CL_b > 0.5$, a regime where presumably there is no need for the correction, the *power* ($= 1 - \text{Pr}(\text{type II error})$) is strictly larger for the regulated CL_s than for the usual CL_s . Figure 7.2 demonstrates the increased power of the regulated CL_s . For a fixed background yield in the left plot of Fig. 7.2 shows that the power of the regulated CL_s interpolates between the baseline CL_s definition and the CL_{s+b} , which is a proper p -value and has maximal power by the Neyman-Pearson lemma (with the caveats

⁴The regulated CL_s is similar to the idea of power-constrained limits in Ref. [643]. Without a proper loss function for Type 1 errors under the background-only hypothesis, there is no unique way to regulate the p -value. Thank you K. Cranmer for pointing out this interesting paper.

discussed above). The right plot of Fig. 7.2 shows the minimum number of signal events that are needed to exclude a model given the number of estimated background events. The regulated CL_s requires about 15% fewer signal events than the baseline CL_s procedure. More sophisticated choices for f are possible to increase the power in the low CL_b regime and still build in protection from the undesirable properties of CL_{s+b} . Despite the promise of the regulated CL_s , the community standard is the baseline CL_s and therefore it is important to use the same definition when comparing results with other analysis. Thus, the baseline CL_s is used for all subsequent results.




Figure 7.2: Left: The power for three exclusion techniques (see the text for details) as a function of the number of signal events, for a fixed number of background events. Right: The minimum number of signal events that are needed to exclude a model given the number of estimated background events. A log-normal constraint is used to model the background uncertainty and 100,000 toys are used for each p-value calculation.

7.2 CR-only Fit

Before describing the compatibility of the predicted yields with the observations in the signal regions, this section documents the outcome of the CR-only fit from maximizing Eq. 7.3. Figure 7.3 shows the normalization factors μ_i for each signal region. Except for the tN1shape regions, each signal region has dedicated $t\bar{t}$ and W +jets control regions that participate in the fit. The $\sqrt{s} = 13$ signal region additionally has the signal top and $t\bar{t} + \gamma$ control regions to constrain $\mu_{\text{single top}}$ and $\mu_{t\bar{t} + Z}$. For each point, the outer error bar is the total uncertainty from the fit, including the impact of systematic uncertainties. The inner error bar represents the control region statistical uncertainty and is determined by bootstrapping the data in the control regions and resolving the system of equations in Eq. 5.3. The single bin regions have the same number of control region bins as normalization parameters and so the central value from these fits are the same. All of the control regions have $\mathcal{O}(100)$ events and so the statistical uncertainty is $\mathcal{O}(10\%)$ and is the dominant uncertainty for most regions. The normalization factor uncertainties for the single top and $t\bar{t} + Z$ processes for SR13 are significantly larger than the corresponding $t\bar{t}$ and W +jets factor uncertainties due to the small yield ($t\bar{t} + Z$) and purity (single top) in the control regions. The shape fit regions have one normalization parameter per E_T^{miss} bin, as motivated in Sec. 4.4. Each E_T^{miss} bin has four m_T bins, so the CR-only fit (which is actually the full likelihood in Eq. 7.1 only with $\mathbf{v} = 0$) is already over-constrained. For this reason, the system of equations method does not apply and so there are no inner error bars in the last three points in Fig. 7.3. Across all bins, the W +jets scale factors are less than unity. This is comparable to the 15-20% over-estimation of the W +jets process by SHERPA in the inclusive phase space probed by the ATLAS cross-section measurement reported in Ref. [559]. The $t\bar{t}$ normalization factors are approximately consistent with unity, though there is a slight trend for $\lesssim 10\%$ upward corrections. The single top normalization factor is much less than one, but its uncertainty is too large to make conclusions. Despite the large uncertainty on the $t\bar{t} + Z$ normalization factor, it is significantly greater than one. The inclusive $t\bar{t} + Z$ cross-section measurement at $\sqrt{s} = 13$ TeV also observes an excess, but there is not enough events yet to determine

if it is significant [599]; the $\sqrt{s} = 8$ TeV measurement does not see the same excess, though the statistical uncertainty is comparably large [600]. If the same (lower) k -factor is used for $\sqrt{s} = 8$ TeV $t\bar{t} + \gamma$ validation region (see Sec. 5.4.2) as for the $t\bar{t} + \gamma$ CR at $\sqrt{s} = 13$ TeV, the data suggests a normalization factor that is also ~ 1.5 .




Figure 7.3: The μ values from the CR-only fit for all single bin signal regions. There is one normalization factor per E_T^{miss} bin for the tN1shape fit. The outer error bars indicate the total post-fit uncertainty while the inner error bars show the statistical uncertainty only (see the text for details).

Another important aspect of the fit to investigate before showing the full results is the level of nuisance parameter profiling. In the signal regions with equal numbers of control regions and normalization factors, the CR-only fit will not profile the nuisance parameters by construction. However, once the signal regions are included and in general for the shape fit signal region, the fit is over-constrained and so the nuisance parameters can change from their initial values. Figure 7.4 shows the impact of a *background-only* fit using the control and signal regions associated with tNmed, tNhigh, and tN1shape for JES and JER nuisance parameters. The background-only fit is identical to the CR-only fit, but including the data and simulation in the signal region (i.e. maximize Eq. 7.1 with $\nu = 0$). In the absence of profiling, the mean is zero and the standard deviation is unity. As expected, since the number of events

in the single bin regions is small compared to the number of events in the control regions, there is essentially no profiling of the jet energy related nuisance parameters. In contrast, there is significant profiling of the JES and JER nuisance parameters for the tN1shape fit. The six reduced nuisance parameters from the in-situ measurements (NP1-6) are ordered from the biggest to smallest impact on the jet energy scale. Therefore, the largest profiling occurs for the first NPs (at the 60% level) whereas there is little sensitivity and thus little profiling for NP5 and NP6. Due to their size, there is also significant profiling for the nuisance parameters associated with the η inter-calibration (increases with $|\eta|$), the jet flavor [644, 645], the technical closure⁵, and pileup⁶ (see Ref. [121] for more details). Interestingly, the largest change in the central value of a nuisance parameter is for the **b**-jet energy scale, which is reduced by about 25%. The last JES nuisance parameter is associated with the high p_T JES derived from the single-hadron response. As this is only relevant for jets with $p_T \gtrsim 1$ TeV, the corresponding nuisance parameter is not profiled. The last three points in Fig 7.4 show the profiling of the JER uncertainty nuisance parameters, one per E_T^{miss} slice. Due to its significant impact on the m_T shape, this parameter is significantly profiled in all three E_T^{miss} regions.




Figure 7.4: The post-fit JES and JER nuisance parameters for the background-only all-bins fit for tNmed, tNhigh, and tN1shape. Red (blue) lines indicate the JER (JES) parameters. Single bin (shape fit) regions are on the left (right).

⁵There was a small change in the simulation from the calibration sample to the applied sample.

⁶This includes an uncertainty on the μ , NPV, and p_T dependence of the pileup corrections and an uncertainty on the modeling of the median pileup density ρ . See Ref. [207] for more detail.

7.3 Validation Regions

Data and predictions in the signal regions using the CR-only fit are shown in Fig. 7.5. In addition to the signal regions, Fig. 7.5 also shows comparisons for a set of *validation regions* that are kinematically between the control regions and signal regions. The validation regions have the same selections as the corresponding control regions, but instead of $60 \text{ GeV} < m_T < 90 \text{ GeV}$, they require $90 \text{ GeV} < m_T < 120 \text{ GeV}$. Both $t\bar{t}$ and W +jets validation regions are associated with tN_{med} , tN_{high} , and tN_{13} . Overall, there is excellent agreement between the predictions and the data; the $\chi^2/\text{NDF} \sim 0.4$ with a p -value of about 98%⁷.




Figure 7.5: A comparison of the data and simulation in the validation and signal regions using the CR-only background fit. The error bar in the ratio panel shows the *pull*, defined as the difference between the data and the prediction, divided by the uncertainty. In this case, the uncertainty is the sum in quadrature of the data statistical uncertainty and the total background uncertainty.

⁷See Sec. 7.5 for a discussion about the possible over-estimation of uncertainties.

7.4 Exclusion Limits

In the absence of a significant excess, limits are set on simplified models with $\tilde{t} \rightarrow t\tilde{\chi}^0$. The following sections will show a series of contours, similar to the schematic one shown in Fig. 7.6. The horizontal axis will be the stop mass, which sets the cross-section (see Fig. 2.1) and the vertical axis will be the neutralino mass or the mass difference between the neutralino mass and the stop mass, quantities which are responsible for determining how much phase space is available for the stop decay products. A black dashed line indicates the *expected exclusion limit*, which is determined by computing the median CL_s assuming that the data follow a Poisson distribution with mean value given by the SM prediction. The 1σ systematic uncertainty is represented by a yellow band around the dashed line. The exclusion limit from the observed data is represented by a solid red line and the theoretical cross-section uncertainty on the signal is represented by red dashed lines around the solid line. Only the total cross-section uncertainty on the signal is included in the red dashed lines; all other uncertainties on the signal model are included in the yellow band. Due to computing and storage limitations, a grid of models with approximately 50 GeV spacing in m_{stop} and 25 GeV in m_{LSP} is used to estimate the full contour. A bilinear interpolation between grid points is performed using the significance, $\sigma = \Phi^{-1}(1 - CL_s)$.




Figure 7.6: A schematic diagram demonstrating how the exclusion limits are presented. See the text for details.

7.4.1 Early $\sqrt{s} = 8$ TeV Results

Figure 7.7 shows the exclusion contour in the $(m_{\text{stop}}, m_{\text{LSP}})$ plane after collecting 13 fb^{-1} of data at $\sqrt{s} = 8$ TeV. The three signal regions SR1-3 are combined using the mapping shown in Fig. 7.8 based on the lowest expected CL_s value. The limits extend significantly beyond the full $\sqrt{s} = 7$ TeV Run sensitivity [501], pushing the limit up to about $m_{\text{stop}} = 625$ GeV for a massless LSP. There are three factors that led to the improved limit. First, the total integrated luminosity at $\sqrt{s} = 7$ TeV was only 4.7 fb^{-1} resulting in a factor of 2.8 more events with the 13 fb^{-1} at $\sqrt{s} = 8$ TeV. Second, the increase in the center-of-mass energy increased the cross-section for stops by about a factor of 2 for $m_{\text{stop}} \sim 500$ GeV. The stop cross section is a factor of 4.6 lower for $m_{\text{stop}} = 625$ GeV ($\sqrt{s} = 8$ TeV limit) than at $m_{\text{stop}} = 500$ GeV ($\sqrt{s} = 7$ TeV limit). Therefore, there would be 20% more $m_{\text{stop}} = 625$ GeV stop events at $\sqrt{s} = 8$ TeV than $m_{\text{stop}} = 500$ GeV events at $\sqrt{s} = 7$ TeV. However, the number of background events also increases with $\sqrt{s}$. The dominant $t\bar{t}$ cross section increases by at least 50% between $\sqrt{s} = 7$ and $\sqrt{s} = 8$ TeV (see Fig. 2.1). To achieve sensitivity to $m_{\text{stop}} = 625$ GeV, a kinematically tighter event selection is required. The third factor that led to the improvement in the early $\sqrt{s} = 8$ TeV analysis is the addition of the m_{T2} variables that allowed for a harsher event selection with a higher background rejection than with the toolkit from the $\sqrt{s} = 7$ TeV analysis. These tools also helped improve the sensitivity for high LSP masses (SR2), where the maximum height of the contour at $\sqrt{s} = 8$ TeV nearly doubled with respect to the $\sqrt{s} = 7$ TeV analysis. Note that the new tools mostly helped to improve the kinematically tight signal regions (SR2, SR3); there was essentially no improvement at low stop mass (SR1).

7.4.2 Full $\sqrt{s} = 8$ TeV Results

New techniques and more integrated luminosity further improved the limits from the partial to the full $\sqrt{s} = 8$ TeV dataset. The updated exclusion limits are presented in Fig. 7.9. For a massless neutralino, the high mass limit extends to about $m_{\text{stop}} = 675$ GeV, a 50 GeV improvement over the expected limit from Sec. 7.4.1 and evidence that the small excess in SR3 from the partial dataset is a statistical fluctuation. For the same integrated luminosity, there would be about 57% fewer stop events with $m_{\text{stop}} = 675$ than for $m_{\text{stop}} = 625$. Accounting for the difference in dataset size, there would be about 10% fewer stops at the limit with the full dataset compared to the number of stops at the (expected) limit for the partial dataset if the acceptance was constant. Figure 7.11 shows the acceptance for the tNmed and tNhigh signal regions. Near $m_{\text{stop}} = 625$ GeV, the acceptance is about 5% for tNhigh and increases by about 20% when m_{stop} is increased by 50 GeV. Therefore, the number of signal events near the exclusion limit with the full dataset would be slightly higher than the number of of signal events near the partial dataset exclusion for a fixed event selection. However, the number of background events increases by about 55%. This means that tNhigh is able to reject about 55% more background events than SR3 while only reducing the signal by about 10%.

For low stop masses, a significant increase in the sensitivity is from the multi-bin signal region (tNshape). Figures 7.12 and 7.13 highlight this challenging region of parameter space. The smallest mass gap $m_{\text{stop}} - m_{\text{top}} - m_{\text{LSP}}$ that is excluded is about 12-14 GeV for a stop mass near 250 GeV. This is a significant improvement of about 12 GeV over previous limits. Over this range, the top quark p_T drops by nearly a factor of two based on Eq. 2.2. For higher stop masses, the limit weakens as the cross-section is too small for the inclusive tNshape event selection to have any sensitivity. The sensitivity also decreases for lower stop masses as the signal is less distinguished from the background. The limit at low stop mass will not improve with more data unless the dominant systematic uncertainties can be reduced, additional variables are identified with a larger variation in s/b , and/or modeling uncertainties are sufficiently small to employ new techniques based on ISR or ME jets mentioned in Sec. 2.1.




Figure 7.9: The exclusion contour and SR mapping of simplified stop models using the full $\sqrt{s} = 8$ TeV dataset based on the lowest expected CL_s value. By construction, tNmed is the most sensitive for intermediate m_{stop} and large m_{LSP} and tNhigh is the most sensitive at m_{stop} .

All of the limits presented thus far assume the stop is mostly the partner of the right-handed top quark. As discussed in Sec. 2, it is expected that the limits are slightly weaker for a mostly left-handed stop due in part to the softer lepton p_T spectrum. The stop mass limit for mostly right-handed stops and $m_{\text{LSP}} = 50$ GeV is about 50 GeV higher than for mostly left-handed stops.

Even though the event selections presented in this section were optimized using simplified stop models with a 100% branching ratio $\tilde{t} \rightarrow t\tilde{\chi}_1^0$, the signal regions are sensitive to many extensions of the SM. The next section will discuss the sensitivity to other models that predict $t\bar{t} + E_T^{\text{miss}}$ topologies. To close this section, consider a slightly less-simplified scenario in which $\mathcal{BR}(\tilde{t} \rightarrow t\tilde{\chi}_1^0) < 100\%$. Figure 7.10 shows the observed exclusion limits for $\mathcal{BR}(\tilde{t} \rightarrow t\tilde{\chi}_1^0) = 50\%$, with the other 50% of the time the




Figure 7.10: The observed limits using the best-expected map from Fig. 7.9 for two different branching ratio assumptions ($\mathcal{BR}(\tilde{t} \rightarrow t\tilde{\chi}_1^0) + \mathcal{BR}(\tilde{t} \rightarrow b\tilde{\chi}_1^\pm) = 1$). The green lines are computed by comparing the observed cross-section limits from Fig. 7.9 to 25%, 50%, or 75% of the predicted cross-section. To make a smooth contour, the significance is set to $2\sigma_{\text{excluded}}/\sigma_{\text{predicted}}$. The wavy line on the left-hand side is due to the lack of signal models above the limit.

stop decays via the flavor-changing decay $\tilde{t} \rightarrow b\tilde{\chi}_1^\pm$ followed by $\tilde{\chi}_1^\pm \rightarrow W^\pm\tilde{\chi}_1^0$. A two-dimensional mass plane is no longer sufficient to uniquely specify such a model. To reduce the parameter space, Fig. 7.10 assumes $m_{\tilde{\chi}_1^\pm} = 2m_{\tilde{\chi}_1^0}$, motivated by gaugino universality (see Chapter 1). For high neutralino mass, the limit is only reduced by about 50 GeV whereas at low neutralino mass, the limit is reduced by over 150 GeV. At low stop mass, there is little impact from the branching ratio reduction. All of the exclusion at high stop mass is due to tNmed and not tNhigh. As one might expect, the tighter event selections reduce the breadth of sensitivity. Interestingly, the exclusion limit for $\mathcal{BR}(\tilde{t} \rightarrow t\tilde{\chi}_1^0) = 50\%$ is significantly better than would be expected if tNmed and tNshape were only sensitive to events where both stops decayed via $t\tilde{\chi}_1^0$ ($\sigma \mapsto 25\%\sigma$). The limit is actually similar to the case where the signal regions are not sensitive to events where both stops decay via $b\tilde{\chi}_1^\pm$ ($\sigma \mapsto 75\%\sigma$).




Figure 7.11: Top: the acceptance for the tNmed (left) and tNhigh (right) event selections as a function of m_{stop} and m_{LSP} using a *particle-level* event selection analogous to the ones described in Sec. 4. The particle-level objects are similar to the ones used throughout Part II and are detailed in the appendix of Ref. [646]. Bottom: the ratio of the acceptance using detector-level objects to the acceptance using particle-level objects. This ratio is a correction for detector-effects in the event selection and is mostly uniform in the sensitive regions of parameter space. The acceptance in the upper plots is defined using particle-level objects in order to facilitate comparisons with other models for which a detector-simulation is not available.




Figure 7.12: Excluded regions in the $(m_{\text{stop}}, m_{\text{stop}} - m_{\text{top}} - m_{\text{LSP}})$ plane. The plots are identical aside from a vertical inversion. The region inside the bounded contour is excluded using the $\text{CL}_s < 0.05$ criteria (N.B. this is **not** the 95% confidence level). In the region below (above) the contour in the left (right) plot, the LSP is a Tachyon ($m < 0$) and in the region above (below) the contour, models are allowed by the data.




Figure 7.13: Left: the excluded region of the $(m_{\text{stop}}, m_{\text{stop}} - m_{\text{top}} - m_{\text{LSP}})$ mass plane (linear-scale version of Fig. 7.12). Right: a one-dimensional projection of the significance $\sigma = \Phi^{-1}(1 - \text{CL}_s)$ as a function of the mass gap $m_{\text{stop}} - m_{\text{top}} - m_{\text{LSP}}$ for $m_{\text{stop}} = 250$ GeV. Crosses indicate grid points and the solid/dashed lines interpolate between points. The horizontal dashed line is the exclusion threshold of $\Phi^{-1}(0.95)$.

7.4.2.1 Recasting Stop Limits

As introduced in Sec. 1.5, there are several extensions of the SM that predict new particles resulting in $t\bar{t} + E_T^{\text{miss}}$ topologies. Table 7.1 categorizes models based on their spin and color charge, which determine the cross-section. As neither the spin nor the color are measured, the cross-section is a sum over states and therefore the cross-section increases with spin and the dimension of the color representation. Scalar leptoquarks have the same cross-section as stops⁸ while vector leptoquarks have a significantly higher cross-section. Figure 7.14 illustrates the cross-section differences as a function of new particle mass⁹. The remainder of this section focuses on GMS, but the methods could be applied to any of the models in Table 7.1.

Name	Spin	Color Charge	Electric Charge	Relative Cross-section
–	0	1	2/3	Tiny
Stop	0	3	2/3	Small
Scalar Leptoquark	0	3	2/3	Small
–	0	8	2/3	Medium
–	1/2	1	2/3	Tiny
T (or T')	1/2	3	2/3	Medium
GMS	1/2	8	2/3	Large
–	1	1	2/3	Tiny
Vector Leptoquark	1	3	2/3	Medium-Large
–	1	8	2/3	Large

Table 7.1: Example models that result in $t\bar{t} + E_T^{\text{miss}}$ categorized by their spin, color charge, electric charge (has to be 2/3 if the missing particles are neutral), and the relative cross-section. The acronym GMS stands for gluino mediated stop. The particle T exists in a variety of models and could decay via $T \rightarrow tA_0$ for a new weakly interacting particle A_0 (see Ref. [500] and the references therein), or it could be a vector-like quark decaying via $T' \rightarrow tZ(\rightarrow \nu\bar{\nu})$, which also has the $t\bar{t} + E_T^{\text{miss}}$ topology. A ‘–’ indicates that there is no standard model with these properties.

⁸With small differences in acceptance that depend on the stop mixing.

⁹Thanks to Marat Freytsis for providing the UFO model for the vector lepto-quark, which was used through MG5_aMC 2.1.1 to compute the cross-sections.




Figure 7.14: The cross-section for various models highlighted in Table 7.1. The parameter k for the vector lepto-quarks corresponds to the κ model parameter in Ref. [482].

As introduced in Sec. 1.5, gluino mediated stops (GMS) with nearly mass degenerate stops and neutralinos have a similar signature to direct stop production. However, the gluino pair production cross section is much larger than the cross section for direct pair produced stops. For example, at $\sqrt{s} = 8$ TeV, the stop pair production cross-section for $m_{\text{stop}} = 800$ GeV is about 0.002 pb whereas the cross-section for stops produced from the decay of 1 TeV pair produced gluinos is about 0.02 pb [503]. Therefore, mass limits in the GMS model will be higher than those for direct stop production. At $\sqrt{s} = 13$ TeV, the effective cross section for the gluino mediated process is twice the direct stop pair production cross section, which is exploited by the early $\sqrt{s} = 13$ TeV search to be sensitive to discover new particles earlier than expected. The results of that search are presented in Sec. 7.4.3. This section describes how the limits on GMS models can be extracted indirectly from the stop limits discussed in Sec. 7.4.2.

General GMS models are well-motivated by naturalness (see Sec. 1.5). Both ATLAS [351, 647–650] and CMS [651–656] have searched extensively for generic GMS models, excluding spectra with large mass splittings up to $m_{\tilde{g}} \lesssim 1.5$ TeV. Compressed

mass spectra are generally more difficult to identify than spectra with large splittings, but are still well-motivated by e.g. dark matter. Gluino pair production with four high energy top or bottom quarks leaves a striking signature in a detector. However, if any of the mass splittings are compressed, the power of traditional techniques may deteriorate. Figure 7.15 shows the possible GMS mass hierarchies, highlighting the presence of direct stop-like and direct sbottom-like signatures. Searches for direct stop/sbotttom pair production can be recast as searches for GMS in order to extend the sensitivity. This section will show that compressed GMS limits at $\sqrt{s} = 8$ TeV can be extended by at least 225 GeV for a 1.1 TeV stop. Before describing the GMS limits, the general notion of *equivalence* for model reinterpretation is described in Sec. 7.4.2.1.1.




Figure 7.15: All possible mass hierarchies for $\tilde{g} \rightarrow t\tilde{t}$, $\tilde{t} \rightarrow b\chi_1^\pm$. The arrows indicate which, if any, high energy particles are produced in the cascade decay. Decay (d) is phenomenologically the same as the compressed GMS and (g) is equivalent to the gluino mediated sbottom. The other possibilities are better covered by dedicated searches than direct stop/sbotttom production. The gluino-mediated sbottom is similar, but with $\tilde{t} \leftrightarrow \tilde{b}$ and $t \leftrightarrow b$.

7.4.2.1.1 Equivalent Models

The two ingredients needed to set experimental limits on a SUSY model are the acceptance and the cross-section. The acceptance is the predicted fraction of SUSY events that pass the experimental event selection and the cross-section is the rate of production for SUSY events¹⁰. Let M be a particular SUSY model and define $M_s = (\epsilon, \sigma)$, where ϵ is the acceptance of the model M under an experimental selection s and σ is the cross section for M . Note that the cross section does not depend on s . Two distinct SUSY models M and M' are defined to be *equivalent* under the experimental event selection s if $M_s = M'_s$. If two models M and M' are equivalent under the experimental selection s , then one is excluded by s if and only if the other is also excluded. In SUSY simplified models, the cross section and acceptance depend on only a few key parameters. For direct stop production ($\mathcal{M}_{\tilde{t}}$), the cross section σ for models in $\mathcal{M}_{\tilde{t}}$ depend only on the stop mass, $m_{\tilde{t}}$ and the acceptance under a given experimental selection depends on both¹¹ the stop mass and the neutralino mass $m_{\tilde{\chi}_1^0}$. For the GMS production ($\mathcal{M}_{\tilde{g}}$), the cross section is set by the gluino mass $m_{\tilde{g}}$ and the acceptance depends on all three masses: $m_{\tilde{g}}$, $m_{\tilde{t}}$, and $m_{\tilde{\chi}_1^0}$. Since models in $\mathcal{M}_{\tilde{t}}$ and $\mathcal{M}_{\tilde{g}}$ both need the stop and neutralino masses as input, for clarity, $m_{\tilde{t}_1}^{\tilde{t}}$ will denote the stop mass in a model $M_{\tilde{t}} \in \mathcal{M}_{\tilde{t}}$ and $m_{\tilde{t}_1}^{\tilde{g}}$ represents the stop mass in a model $M_{\tilde{g}} \in \mathcal{M}_{\tilde{g}}$ (and analogously for the neutralino). The next sections describe a procedure for associating to every model $M_{\tilde{g}} \in \mathcal{M}_{\tilde{g}}$, an equivalent model $M_{\tilde{t}} \in \mathcal{M}_{\tilde{t}}$. Experimental limits on $M_{\tilde{t}}$ can then be used to place limits on $M_{\tilde{g}}$.

¹⁰As with Fig. 7.11, there is often a distinction between the particle-level acceptance (often simply called acceptance) and the detector-level acceptance.

¹¹As noted in Sec. 2.1, the acceptance also depends on the top polarization. Top quarks in the GMS model originate directly from the scalar gluino and thus are unpolarized (the stops are produced on-shell in these models). This will have a small impact on the exclusion which is ignored in the following.

7.4.2.1.2 Limits on Compressed Gluino Mediated Stop Production

The lost sensitivity to compressed $\tilde{g} \rightarrow t\bar{t}$ from direct gluino searches with multi-top quark, multi- b quark, or multi-lepton final states can be recovered by direct stop searches. There are only subtle differences between the models due to the fact that the gluino is a fermionic color octet, instead of a scalar triplet like the stop, so there will be small changes in angular distributions and radiation patterns between jets. However, most analysis techniques are not sensitive to these effects. One non-negligible difference is the electric charge, as stops can have the same charge when from gluinos (as it is a Majorana particle), but must be oppositely charged for direct stop production. For this reason, same-sign lepton searches can retain sensitivity even when the decay chains are compressed. However, the results below indicate that the one- and zero-lepton searches are more powerful, due to the much larger branching ratio.

Given an experimental selection s , for a particular model $M_{\tilde{t}} \in \mathcal{M}_{\tilde{t}}$, the goal is to find an equivalent model $M_{\tilde{g}} \in \mathcal{M}_{\tilde{g}}$. The first step in finding an equivalent model is to match the cross sections $\sigma(M_{\tilde{g}}) = \sigma(M_{\tilde{t}})$. There is a one-to-one correspondence between $m_{\tilde{g}}$ and $\sigma(M_{\tilde{g}})$ and between $m_{\tilde{t}}^{\tilde{t}}$ and $\sigma(M_{\tilde{t}})$. The numerical relationships can be found in Ref. [503]. Therefore, given $m_{\tilde{t}}^{\tilde{t}}$, there is a unique $m_{\tilde{g}}$ such that

$$\sigma(M_{\tilde{t}}) = \sigma(m_{\tilde{t}}^{\tilde{t}}) = \sigma(m_{\tilde{g}}) = \sigma(M_{\tilde{g}}). \quad (7.6)$$

The second step is to find $m_{\tilde{\chi}_0^0}^{\tilde{g}}$ (chosen to be nearly identical to $m_{\tilde{t}}^{\tilde{g}}$) and $m_{\tilde{\chi}_0^0}^{\tilde{t}}$ such that the acceptances under s are the same for $M_{\tilde{t}}$ and $M_{\tilde{g}}$. As described in Sec. 2.1, this can be accomplished by choosing $m_{\tilde{\chi}_0^0}^{\tilde{g}}$ and $m_{\tilde{\chi}_0^0}^{\tilde{t}}$ such that the final state objects have the same top quark momentum spectrum $p(M, m)$ given in Eq. 2.2, where $(M, m) = (m_{\tilde{t}}^{\tilde{t}}, m_{\tilde{\chi}_0^0}^{\tilde{t}})$ for $\mathcal{M}_{\tilde{t}}$ and $(M, m) = (m_{\tilde{g}}, m_{\tilde{\chi}_0^0}^{\tilde{g}})$ for $\mathcal{M}_{\tilde{g}}$. Given $m_{\tilde{t}}^{\tilde{t}}$, $m_{\tilde{\chi}_0^0}^{\tilde{t}}$ and determining $m_{\tilde{g}}$ by the equality of the cross sections between $M_{\tilde{t}}$ and $M_{\tilde{g}}$, $m_{\tilde{t}}^{\tilde{g}}$ is chosen by solving $p(m_{\tilde{t}}^{\tilde{t}}, m_{\tilde{\chi}_0^0}^{\tilde{t}}) = p(m_{\tilde{g}}, m_{\tilde{\chi}_0^0}^{\tilde{g}})$. The solution to this equation is quartic in $m_{\tilde{t}}^{\tilde{g}}$, so in general there can be up to four real solutions. Fortunately, two solutions are negative

(or imaginary) and of the two possible positive solutions, only one is smaller than $m_{\tilde{g}}$ and thus there is at most one physical solution. Various kinematic distributions for one particular set of equivalent models are shown in Fig. 7.16. The model $M_{\tilde{t}}$ is specified by a 700 GeV stop mass and a massless neutralino and the equivalent model $M_{\tilde{g}}$ has a 1.1 TeV gluino and a ~ 650 GeV stop/neutralino. By construction, all of the kinematic distributions are nearly identical between these two models and as a result, any selection s based on kinematic variables should have the same acceptance. The difference between the two models is quantified in Fig. 7.17, which shows that the approximation equating these models is valid only when $\delta \equiv |m_{\tilde{t}}^{\tilde{g}} - m_{\tilde{\chi}_1^0}^{\tilde{g}}|$ is sufficiently small. For $\delta \lesssim 15$ GeV for this representative model, the correction to the efficiency is a few percent and grows to about ten percent when $\delta \sim 10$ GeV. The number of jets with a particular transverse momentum increases when δ increases because the charm quarks can produce measurable jets when there is enough phase space. In contrast, the magnitude of the missing transverse momentum decreases because the charm quarks take energy away from the neutralinos. For this reason, the product of efficiencies for jet variables and missing momentum variables is much less dependent on δ . Note that if the four-body decay of the stop, $\tilde{t} \rightarrow bff'\tilde{\chi}^0$, for fermions f and f' , dominated over the two-body decay $\tilde{t} \rightarrow c\tilde{\chi}^0$, the dependence on δ is also reduced because there are more objects that need to share the sparse phase space.

A set of model equivalences between direct stop and gluino mediated compressed stop models are summarized in Table 7.2. These models are chosen because they are at the edge of the high mass exclusion limit from Sec. 7.4.2.

$m_{\tilde{t}}^{\tilde{t}}$ [GeV]	$m_{\tilde{\chi}_1^0}^{\tilde{t}}$ [GeV]	$\sigma(m_{\tilde{t}}^{\tilde{t}})$ [pb]	$m_{\tilde{g}}^{\tilde{g}}$ [GeV]	$m_{\tilde{t}}^{\tilde{g}}$ [GeV]
675	100	0.011	1090	670
625	220	0.018	1030	690
600	240	0.025	995	680
550	240	0.045	930	660

Table 7.2: A set of direct stop models that are equivalent to GMS models. The third and fourth columns are from Ref. [503].




Figure 7.16: Kinematic distributions for two distinct, but equivalent SUSY models. The model labeled *stop* is direct stop pair production with a 700 GeV stop mass and a massless neutralino. The model labeled *gluino* is a gluino mediated compressed stop model with a 1.1 TeV gluino, a 652 GeV stop and a 650 GeV neutralino. The generation is performed with MadGraph5_aMC@NLO version 5.2.1.1 [505] for the matrix element and Pythia 6.428 [96] for the parton shower and hadronization. The 2 GeV difference between the stop mass and neutralino mass in the gluino mediated stop model is due to a 1.5 GeV charm mass in Pythia (for $\tilde{t} \rightarrow c\tilde{\chi}^0$). A detector simulation is modeled with Delphes v3.1.2 [657]. Jets are clustered with the anti- k_t algorithm [117] with $R = 0.4$ using the fastjet program [203]. The top plot shows the distribution of the leading four jet p_T , the bottom left plot shows the magnitude of the missing transverse momentum and the bottom right plot shows the number of jets with $p_T > 25$ GeV.




Figure 7.17: The efficiency of a 315 GeV threshold on the magnitude of the missing transverse momentum (square markers) and the efficiency of a four jet requirement with transverse momentum thresholds (75, 65, 40, 25) GeV (circles). The markers indicate the efficiency for the compressed gluino model with $m_{\tilde{g}}^g = 1.1 \text{ TeV}$ and $m_{\tilde{t}}^g = 650 \text{ GeV}$. The yellow band is the efficiency for the equivalent stop model with $m_{\tilde{t}}^g = 700 \text{ GeV}$ and a massless neutralino. The band and the error bars represent statistical uncertainties from finite simulated datasets. The generation is performed with MadGraph5_aMC@NLO version 5.2.1.1 [505] for the matrix element and Pythia 6.428 [96] for the parton shower and hadronization. A detector simulation is modeled with Delphes v3.1.2 [657]. Jets are clustered with the anti- k_t algorithm [117] with $R = 0.4$ using the FASTJET program [203].

For every direct stop production model, the procedure above assigns an equivalent gluino model. However, there are gluino models that do not have an equivalent stop pair production model. At a fixed gluino mass $m_{\tilde{g}}$, let $m_{\tilde{t}}^{\tilde{t}}$ be the stop mass such that $\sigma(m_{\tilde{t}}^{\tilde{t}}) = \sigma(m_{\tilde{g}})$. Stop masses $m_{\tilde{t}}^{\tilde{g}}$ with $p(m_{\tilde{g}}, m_{\tilde{t}}^{\tilde{g}}) > p(m_{\tilde{t}}^{\tilde{t}}, 0)$ have no equivalent stop pair production model. For example, at $\sqrt{s} = 8$ TeV, $m_{\tilde{t}}^{\tilde{t}} = 700$ GeV and $m_{\tilde{g}} = 1.1$ TeV have the same cross section, but clearly the gluino model with $m_{\tilde{t}}^{\tilde{g}} = 0$ has no equivalent direct stop pair production model since the available momentum in the gluino model exceeds the direct stop mass. However, this leads to an artificial truncation of gluino models that can be excluded by direct stop searches. Acceptance generally increases with the top/neutralino momentum for a fixed cross section. Therefore, if the point $(m_{\tilde{g}}, m_{\tilde{t}}^{\tilde{g}})$ is excluded by a particular search, then all models specified by $(m_{\tilde{g}}, x)$ with $x < m_{\tilde{t}}^{\tilde{g}}$ will also be excluded.

One can take this argument further to extrapolate to a region of phase space applicable to gluino searches, but forbidden to direct stop searches. Consider a direct stop model with $m_{\tilde{t}}^{\tilde{t}}$ just beyond the exclusion limit. The equivalent gluino model with mass $m_{\tilde{g}}$ will correspondingly not be excluded. However, since the acceptance increases in decreasing $m_{\tilde{t}}^{\tilde{g}}$, there may be a model with gluino mass $m_{\tilde{g}}$ that is excluded, but has no equivalent direct stop model. One way to estimate the excluded region is to fit the acceptance curve from Fig. 2.4 and predict the acceptance of a particular gluino model. For large values of p , the acceptance should be roughly linear in p as the missing momentum in the event is linear in p . Therefore, a linear fit for $p > 200$ GeV is shown in Fig. 2.4 for extrapolating the acceptance to higher values of p . Values of $(m_{\tilde{g}}, m_{\tilde{t}}^{\tilde{g}})$ can be declared excluded if $\mathcal{L}_{\text{int}} \times \sigma(m_{\tilde{g}}) \times \epsilon(p(m_{\tilde{g}}, m_{\tilde{t}}^{\tilde{g}})) \times \kappa > n_{\text{excluded}}$, where κ is the efficiency from Fig. 7.11 (roughly independent of stop/LSP mass) and n_{excluded} is the model-independent limit on the number of BSM events from scanning over the number of BSM events in the SR and then re-running the exclusion fit. For tNmed/tNhigh, the observed (expected) n_{excluded} is 8.5/6.0 (9.2/6.0).

One can do even better than naively recasting limits based on n_{excluded} by tightening thresholds on the key variables (e.g. E_T^{miss} , m_T , and $\mathbf{am}_{T2}$), but this change would require a careful assessment of the change in the background yield which is beyond the scope of this section.

7.4.2.1.3 Derived Limits

Re-casted direct stop limits are shown Fig. 7.18 alongside existing limits from the ATLAS same-sign search [648] and the inclusive one lepton¹² search [650]. The same-sign limits are optimistic because the selection in Ref. [648] requires a third hard jet, which is not part of the leading order description of the final state. Estimates based on calculations with MG5_aMC version 2.1.1 [505] indicate that the fraction of the time an additional jet from initial or final state radiation has enough p_T to pass the jet selection is roughly 40%. This agrees well with the three jet selection efficiency published in auxiliary material Table 64 [658] of the ATLAS search for a model with a large stop mass for which kinematically the soft c -quark jets will not pass the hard jet p_T threshold. As the mass splitting between the stop and the neutralino goes to zero, the reduction in the limit for the highest mass splitting reduces by $\lesssim 100$ GeV (not shown). The inclusive one lepton search is based on generic variables such as E_T^{miss} , m_T , effective mass, etc. and is not optimized for the $t\bar{t} + E_T^{\text{miss}}$ final state (the limits may even degrade as $m_{\tilde{t}} \rightarrow m_{\tilde{\chi}^0}$). The improvement over these existing analyses for the reinterpreted direct search are shown in shaded blue in Fig. 7.18. The darkest blue is from the strict re-interpretation based on the strategy leading up to Table 7.2. The light blue area below the dark blue area is assumed excluded because the signal efficiency increases for the larger mass splitting. The light blue area to the right of the dashed line is from interpolating and extrapolating the efficiency and comparing to the n_{excluded} . For a 1.1 TeV gluino, the inclusive one lepton limit is extended vertically by about 225 GeV.

¹²A similar search exists in the zero lepton final state, with slightly weaker limits [649]




Figure 7.18: A comparison of existing limits and the re-interpretation of the direct stop search limits at $\sqrt{s} = 8$ TeV (see Sec. 7.4.2). The expected limits (based on the CL_s procedure [642]) are used to control for statistical fluctuations in the observations. The blue shaded region is the re-interpretation of the direct stop search. Above the marked diagonal line, every gluino model has an equivalent stop model. Below this line, there is no equivalent stop model and the exclusion limits are estimated by extrapolating the signal region acceptance as described in the text. The hatched region is from the $\sqrt{s} = 8$ TeV ATLAS search for compressed direct stop production via an ISR monojet [513]. The red line is from the ATLAS inclusive one lepton search [650] (Fig. 18a) and the black line is from the ATLAS same-sign lepton search [648].

7.4.2.1.4 Transitioning to $\sqrt{s} = 13$ TeV

The sensitivity of the direct stop search to GMS models is a strong motivation for performing the stop search with the early $\sqrt{s} = 13$ TeV data. Table 7.3 summarizes the relative increase in cross-sections for direct stop models and GMS models from $\sqrt{s} = 8$ to 13 TeV. Larger masses generally have a larger increase in cross-section because they are probing a smaller momentum fraction of the proton. At the edge of the Run 1 sensitivity, the expected increase in the yield of stops from GMS is twice the corresponding yield for directly produced stops.

$m_{\tilde{t}}$	$\sigma^{8 \text{ TeV}}(m_{\tilde{t}})$	$m_{\tilde{g}}$	$\sigma^{13 \text{ TeV}}(m_{\tilde{t}})$	$\sigma^{13 \text{ TeV}}(m_{\tilde{g}})$	$\sigma(m_{\tilde{t}}) \text{ 13 TeV}/8 \text{ TeV}$	$\sigma(m_{\tilde{g}}) \text{ 13 TeV}/8 \text{ TeV}$
600	0.03	1000	0.2	0.3	7.0	13.6
700	0.008	1125	0.07	0.1	8.3	17.1
800	0.003	1250	0.03	0.06	9.8	21.7

Table 7.3: The expected increase in yields for the direct stop search and the re-interpreted gluino search from $\sqrt{s} = 8$ to 13 TeV. The first column is the stop mass in GeV, the second column is the stop cross section at $\sqrt{s} = 8$ TeV in pb from Ref. [503]. The third column is in pb and also uses Ref. [503] to solve $\sigma(m_{\tilde{t}}) = \sigma(m_{\tilde{g}})$. The fourth and fifth columns give the cross sections for stop and gluino production at $\sqrt{s} = 13$ TeV from Ref. [502]. The last two columns give the ratio of the increase in yields for direct stop and GMS production, respectively.

All possibilities for natural SUSY should be targeted, including those with compressed scenarios. If there is a light enough gluino to mediate, more territory for light stops and sbottoms will be accessible to the direct searches with the early data. As discussed in Sec. 4, GMS models were used as benchmarks for optimizing the $\sqrt{s} = 13$ TeV analysis. The results of that search are presented in the next section.

7.4.3 Early $\sqrt{s} = 13$ TeV Results

Figure 7.19 shows the exclusion limits from SR13 using the 3.2 fb^{-1} from the 2015 $\sqrt{s} = 13$ TeV dataset. The signal region was optimized with the GMS benchmark model $(m_{\tilde{g}}, m_{\tilde{t}}, m_{\tilde{\chi}^0}) = (1250, 750, 745)$, which is just on the edge of the exclusion limit in the right plot of Fig. 7.19. For a stop mass of $m_{\tilde{t}} \approx 650$ GeV, the GMS limit is extended over 400 GeV in gluino mass. Part of this gain is in a genuinely new region of parameter space that does not have an equivalent stop model. A $m_{\tilde{g}} \sim 1.45$ TeV has the same cross-section at $\sqrt{s} = 13$ TeV as $m_{\tilde{t}} \sim 850$ GeV; therefore the highest stop mass in the GMS model that corresponds to a physical direct stop model is about 920 GeV. The highest gluino mass for $m_{\tilde{t}} \approx 650$ GeV that has a physical direct stop model equivalent is about 1.2 TeV.

By construction, SR13 is sensitive to $t\bar{t} + E_T^{\text{miss}}$ topologies and therefore it can be used to set limits on direct stop models as well as the target GMS models. The left plot of Fig. 7.19 shows that the observed limit improves by over 75 GeV for a massless LSP, albeit with a large uncertainty. It is likely that with a dedicated optimization, the limits would be even stronger, though the limit statistics of the 2015 dataset is prohibitive for a much stronger direct stop limit. With the full Run 2 dataset, it may be possible to discover or rule out $m_{\tilde{t}} \lesssim 1$ TeV, the benchmark for naturalness.




Figure 7.19: The exclusion limits using SR13 at $\sqrt{s} = 13$ TeV for direct stop pair production (left) and GMS with $m_{\tilde{t}} - m_{\tilde{\chi}^0} = 5$ GeV (right). The blue filled area in the right plot is from the early Run 2 ATLAS mono-jet search [659].

7.5 The LHC Run I SUSY Epilogue

In addition to the lack of evidence for stops, the Run 1 (and early Run 2) data do not support the existence of weak-scale SUSY in general. Both ATLAS and CMS have conducted extensive searches for SUSY in a multitude of final states, with various numbers of jets, leptons, and photons. The kinematic reach of the detectors have been exploited in order to be sensitive to high mass particles, which may be produced with a low cross section. However, with the large number of searches ($\mathcal{O}(100)$ between ATLAS and CMS), some low p -value results are expected due to statistical fluctuations. This section presents¹³ a meta-analysis of the $\sqrt{s} = 8$ TeV ATLAS and CMS SUSY searches, studying the distribution of p -values associated with the SM-only hypothesis.

7.5.1 Constructing the Dataset

Even though the 8 TeV dataset was collected in 2012, both ATLAS and CMS continue(d) to analyze the data. This section presents data from all analyses prior to an arbitrarily cutoff at the SUSY 2014 conference (July 20, 2014). This includes 17 ATLAS papers [351, 352, 414, 513, 646–649, 661–669] and 12 CMS papers [651–653, 670–678]. The difficulty in assembling the dataset is to understand the correlations between measurements. The general strategy is to categorize the various searches by their selections on jets, leptons, and photons. Two analyses which have non-overlapping requirements in the number and properties of these objects are treated as uncorrelated. For the data, this is an excellent assumption and only breaks down in the rare case that the data in one signal region is used for the background estimate of another signal region. If two signal regions are such that one is a subset of the other, then a decorrelation procedure is attempted in order to produce two orthogonal regions. If the yields are $x \pm \sigma_x$ and $y \pm \sigma_y$ with $x < y$, then the decorrelated regions have yields $x \pm \sigma_x$ and $(y - x) \pm \sqrt{\sigma_y^2 - \sigma_x^2}$. In all other cases, it is not possible with the information given to determine the correlations and the signal regions in question are simply not used. In general, if there are two analyses with an unknown correlation,

¹³The analysis presented here is published in Ref. [660] and includes input from T. Rudelius.

the one with more signal regions is preferred unless the one with fewer regions already has orthogonal selections. The regions to be included were selected before looking at any p -values in order to minimize potential biases. Tables 7.4 and 7.5 give some summary information about the dataset construction given the general guidelines from above. In total, there are 124 ATLAS regions and 325 CMS regions.

arXiv reference	Category	Note
1303.2985	Multijets	Regions orthogonal; drop those with $H_T > 800$ GeV due to overlap with 1402.4770
1402.4770	Multijets	Regions orthogonal; drop those with $H_T \in [500, 800]$ GeV due to overlap with 1303.2985
1305.2390	Multijets	Unknown correlation with 1303.2985 and 1402.4770: remove
1311.4937	One Lepton	Regions orthogonal; use the LS method for uncertainties when given a choice
1308.1586	One Lepton	Unknown correlation with 1311.4937. Prefer 1311.4937 as its regions are orthogonal
1212.6194	Same sign leptons	$SR6 \subseteq SR3 \subseteq SR4 \subseteq SR1 \subseteq SR0$. Drop other regions.
1311.6736	Same sign leptons	Drop regions with ≥ 2 b-jets due to overlap with 1212.6194. Arbitrarily pick the low p_T region
1306.6643	Multileptons	Unknown overlap with 1404.5801. Use 1404.5801 as it has more regions.
1404.5801	Multileptons	Regions orthogonal.
1405.3886	Multileptons	Use the two lepton OS regions only.
1405.7570	Multileptons	Use the two lepton OS regions only. Use signal sensitive regions (as described in the text)
1312.3310	Diphoton	Regions orthogonal.

Table 7.4: An overview of the signal regions used in the meta-analysis from 8 TeV CMS searches.

arXiv reference	Category	Note
1308.1841	Multijets	$8j80xb \subseteq 8j50xb$, $x \in \{0, 1, 2\}$. Unknown correlations between M_j^x regions and others, drop
1308.2631	Multijets	$SRA \ m_{CT}(350) \subseteq SRA \ m_{CT}(300) \subseteq \dots \subseteq SRA \ m_{CT}(150)$
1407.0608	Multijets	$M3 \subseteq M2 \subseteq M1$; $C2 \subseteq C1$
1405.7875	Multijets	$2jt \subseteq 2jm \subseteq 2jl$. $2jW \cap 3j$ unknown, drop $2jW$. $6jt \subseteq 6jm \subseteq 5j$ and $6jl$ ($5j \cap 6jl = \emptyset$) once $6jm$ is removed. Drop all other regions due to unknown correlations.
1406.1122	Multijets	$SRA2 \subseteq SRA1$; $SRA4 \subseteq SRA3$. Drop SRB. $SRC3 \subseteq SRC2 \subseteq SRC1$
1407.0600	Multijets	$SR-0l-7j-C \subseteq B \subseteq A$; Drop $4j$ regions due to $4j \cap 7j = ?$
1407.0583	One Lepton	Unknown correlations between shape fit regions, consider only tN_diag (signal sensitive regions). $tN_high \subseteq tN_med \subseteq$ tightest tN_diag region. $bCb_high \subseteq bCb_med1$. Unknown relation between bCa_low and bCa_med , drop low. Unknown correlation between bCd , $tNbC_mix$ and other regions, drop
1407.0603	At Least One τ	$1\tau Tight \subseteq 1\tau Loose$, $2\tau \text{ GMSB} \subseteq 2\tau \text{ nGM} \subseteq 2\tau \text{ Incl}$. Unknown overlap between $2\tau \text{ bRPV}$ and $2\tau \text{ GM}$, drop $bRPV$. $\tau+1 \text{ bGM} \subseteq \tau+1 \text{ mSUGRA}$. Unknown overlap between $\tau+1 \text{ GMSB}$ and $bRPV$, drop $bRPV$
1407.0350	At least two τ s	$C1C1 \cap C1N2 = ?$, drop $C1C1$ $DS-lowMass \cap DS-highMass = ?$, drop $lowMass$
1403.4853	Two OS Leptons	unknown correlation of $L90,120$ with 1403.5294, drop unknown correlation between $L110-100$, drop $L100$ $H160$ orthogonal, drop MVA region
1403.5294	Two OS Leptons	Jet veto regions orthogonal to other searches, drop $Zjets$ $m_{12}^{150}(x) \subseteq m_{12}^{120}(x) \subseteq WWc(x) \subseteq m_{12}^{90}(x)$, $x \in \{SF, DF\}$ Overlap of $WWb(x)$ with $m_{12}^{90}(x)$ unknown, drop
1404.2500	Same Sign Leptons	Regions orthogonal. Drop $SR3Llow/high$ due to unknown overlap with 1402.7029
1403.5222	Multileptons	$SRxb \subseteq SRxa$, $x \in \{2, 3\}$
1402.7029	Three Leptons	Regions orthogonal except $SR2\tau a \cap SR2\tau b = ?$, drop b
1405.5086	≥ 4 Leptons	$SRxnoZb \subseteq SRxnoZa$, $x \in \{0, 1, 2\}$
1310.3675	Disappearing Tracks	Region inclusion by increasing p_T cut
1310.6584	Out-of-time	For the muon veto, inclusion by jet p_T

Table 7.5: An overview of the signal regions used in the meta-analysis from 8 TeV ATLAS searches. Note that OS = opposite sign. The stop search results are part of the ‘One Lepton’ category.

7.5.2 Statistical Analysis

Once the ATLAS and CMS datasets are constructed, the expected and observed distributions of p -values are computed for both a Gaussian and a lognormal distribution of the expected number of counts (the number of counts itself is assumed to be Poisson). A p -value was assigned to each data point according to

$$p\text{-value} = \int_0^\infty \phi(\lambda|\mu, \sigma) P_{\geq n}(\lambda) d\lambda. \quad (7.7)$$

Here, $P_{\geq n}$ is the probability of observing n or more counts given a Poisson distribution with parameter λ ,

$$P_{\geq n}(\lambda) = \sum_{k=n}^{\infty} \frac{e^{-\lambda} \lambda^k}{k!} = 1 - \sum_{k=0}^{n-1} \frac{e^{-\lambda} \lambda^k}{k!}. \quad (7.8)$$

In addition to analyzing the excesses, one can also study the deficits in the SUSY search regions by replacing $P_{\geq n}$ in Eq. 7.7 with $P_{\leq n}$: the probability of observing n or less counts given a Poisson distribution with parameter λ . The function $\phi(\lambda|\mu, \sigma)$ is the probability distribution function of the specified random variable with mean μ and standard deviation σ . These parameters are the expected value for the number of counts (μ) and the uncertainty on that value (σ). For the Gaussian distribution,

$$\phi(\lambda|\mu, \sigma) = \frac{1}{N\sigma\sqrt{2\pi}} e^{-(\lambda-\mu)^2/2\sigma^2}, \quad (7.9)$$

where N is a normalization constant correcting for the fact that λ cannot be negative, and so the negative part of the distribution must be cut off. For the lognormal distribution, whose support is $\mathbb{R}_+$, no such normalization constant is required,

$$\phi(\lambda|\mu, \sigma) = \frac{1}{\lambda\tilde{\sigma}\sqrt{2\pi}} e^{-(\ln \lambda - \tilde{\mu})^2/2\tilde{\sigma}^2}, \quad (7.10)$$

with $\tilde{\mu} := \ln \mu^2 / \sqrt{\mu^2 + \sigma^2}$, $\tilde{\sigma} := \sqrt{\ln 1 + \sigma^2 / \mu^2}$ defined so that the lognormal distribution is precisely the distribution of $Y = e^X$ for a Gaussian random variable X with mean $\tilde{\mu}$ and variance $\tilde{\sigma}^2$.

One might expect the distribution of p -values defined in this way to be uniformly

distributed on the interval $[0, 1]$ under the null hypothesis, in accordance with the usual interpretation of p -values as the probability of observing a more significant result in precisely $p \times 100\%$ of studies. However, this intuitive understanding is only correct when the distribution is continuous [679], not in the case of Poisson distribution considered here. As a result, the first step of the analysis is to compute the *expected distribution* of p -values under the null hypothesis and then compare this with the observed distribution of p -values. The expected distribution of p -values is determined by summing up the probability that each particular trial would fall into one of ten bins, $(\frac{i}{10}, \frac{i+1}{10}]$, $i = 0, \dots, 9$,

$$\Pr\left(\frac{i}{10} < \text{p-value} \leq \frac{i+1}{10}\right) = \int_0^\infty d\lambda f_i(\lambda) \phi(\lambda|\mu, \sigma), \quad (7.11)$$

where

$$f_i(\lambda) = \sum_{m=0}^{\infty} \left[\Pr(X = m) \times \begin{cases} 1 & \text{if } \Pr(X \geq m) \in (\frac{i}{10}, \frac{i+1}{10}] \\ 0 & \text{otherwise} \end{cases} \right]. \quad (7.12)$$

Here, $X \sim \text{Poisson}(\lambda)$ is the random variable measuring the number of counts, and the $\geq$ in Eq. 7.12 is replaced by a $\leq$ when computing deficits below rather than excesses above the expected signal.

Some of the studied signal regions had 0 expected events. There is no lognormal distribution with a mean of 0, so these regions had to be discarded in performing the lognormal analysis. Fortunately, this only applied to seven of the CMS signal regions and none of the ATLAS ones. However, a fairly sizable fraction had an expected mean that was very close to zero. For these trials, it is reasonable to suspect that neither a Gaussian with a cutoff imposed at 0 nor a lognormal will provide a good approximation to the true error distribution. As a check, the analysis was repeated after removing all data points with $\mu - 2\sigma < 0$ ($\approx 10\%$ for ATLAS, 30% for CMS). The results of this second analysis did not differ qualitatively from the first, indicating that the results of the original analysis are not significantly affected by the statistical modeling of these data points.

Note that the both the log-normal and Gaussian distributions are simple approximations to complicated likelihood functions (see e.g. Sec. 7.1); however, they should capture the essential features of the distributions and the difference between the two approaches will give a sense of the robustness of the procedure.

7.5.3 Results and Discussion

The results of the combined ATLAS and CMS analysis are shown in Figures 7.20 and the results of statistical tests are presented in Tables 7.6-7.7. There is a lack of deficits with $p < 0.1$ at a level of 3.23σ and a lack of deficits with $p < 0.3$ at a level of 3.15σ in the Gaussian case and 4.10σ in the lognormal case. This trend is also observed separately in both the ATLAS and CMS datasets. The observed p -value distribution is significantly different from the expected one, but the difference is not concentrated at low p -values. It is interesting to note that the distributions observed here are somewhat different from those observed in the $\sqrt{s} = 7$ TeV version of this study [680]. That analysis also revealed a deficit of p -values in the tails of the distribution, but there were significantly fewer p -values < 0.1 , indicating a possible overestimation of the mean background as well as the uncertainty. Here, there is actually a slight (statistically insignificant) surplus of p -value excesses < 0.1 in the Gaussian case, but a clear lack of p -value deficits < 0.1 in both the Gaussian and lognormal cases.

The results presented here indicate:

1. The uncertainties are not well-modeled by Gaussian or lognormal distributions.
2. SM predictions have an inherent bias not captured by systematic uncertainties.
3. There is a contribution of SUSY or another model of BSM that causes the observed distribution of p -values to deviate from the expected one.

The present analysis cannot distinguish between these three possibilities. At the least, the differences indicate that the true uncertainty distributions are not well described

by Gaussian or lognormal distributions with the reported means and uncertainties¹⁴. It will be interesting to see how this picture changes with the Run 2 dataset; hopefully the new physics is not so subtle that a meta-analysis is required to identify it. The analysis presented in Part III has significantly reduced the parameter space of natural SUSY. In doing so, new discriminating variables and background estimation techniques have been developed that will continue to be useful for probing the high energy nature of the SM and beyond.




Figure 7.20: The distribution of p -values for excesses (left) and deficits (right) for both Gaussian and log-normal uncertainty distributions. For a continuous probability distribution, one expects the distribution of p -values to be uniform on $[0, 1]$.

¹⁴This is hopefully a strong motivation for making additional statistical details about search results public. There is also a sociological aspect of SUSY searches related to (2); in particular, people often ‘worry less’ about deficits than excess so they receive less scrutiny.

Quantity	Dist. under H_0 (T)	Test statistic (t)		Pr(T > t)	
		Gaussian	LN	Gaussian	LN
Trials with $p < 0.1$	N(0,1)	0.23	-0.27	0.82	0.79
Trials with $p < 0.3$	N(0,1)	-0.57	0.31	0.57	0.75
Trials with $p < 0.2$ or $p > 0.8$	N(0,1)	-5.28	-5.77	$\ll 0.001$	$\ll 0.001$
Expected vs. observed dist.	χ^2_9	36.18	45.74	$\ll 0.001$	$\ll 0.001$

Table 7.6: Results for statistical hypothesis tests on combined ATLAS and CMS excesses, under the assumptions of Gaussian and lognormal error distributions.

Quantity	Dist. under H_0 (T)	Test statistic (t)		Pr(T > t)	
		Gaussian	LN	Gaussian	LN
Trials with $p < 0.1$	N(0,1)	-3.23	-3.23	0.001	0.001
Trials with $p < 0.3$	N(0,1)	-3.15	-4.10	0.002	$\ll 0.001$
Trials with $p < 0.2$ or $p > 0.8$	N(0,1)	-3.24	-3.04	0.001	0.002
Expected vs. observed dist.	χ^2_9	29.04	27.11	0.0006	0.001

Table 7.7: Results for statistical hypothesis tests on combined ATLAS and CMS deficits, under the assumptions of Gaussian and lognormal error distributions.

Chapter 8

Conclusions and Future Outlook

Undoubtedly, particle physics will embark on many grand adventures in the near future. Our experimental and theoretical tools allow us to probe the SM to unprecedented precision and we have begun a full expedition of the unexplored TeV landscape. The SM does not predict any undiscovered particles or forces, but they must be there. Something new is expected, but anything new will be a surprise. Supersymmetric models with a light stop still remain some of the most tantalizing theories. There are many extensions of the search presented in Part III that will push the sensitivity to higher mass scales, more complex decay chains, and more compressed mass spectra. These extensions will benefit from and extend the techniques presented here to identify signal-like events and to suppress and estimate background processes. At the same time, all the searches from Sec. 7.5 (and more) have extensively mapped out the tails of kinematic distributions in the $\sqrt{s} = 8$ TeV and the early $\sqrt{s} = 13$ TeV data. Now, we know a lot about where there is nothing, but there is a lot we can learn about where there is something¹. The phenomenology of a multi-TeV jet is mostly governed by a single number²: α_s . Yet, there are qualitatively different physics processes that occur on all scales spanning $\Lambda_{\text{QCD}} \sim 1$ GeV all the way to the

¹Think about how we know the ocean floor with 5 km precision [681] while the surface of Mars has been mapped with 100 m precision [682]. There are likely no Martians amongst high energy quarks and gluons, but there is a lot of rich structure at the bottom of the ocean.

²At LHC energies, we will also be able to probe electroweak radiation in a regime with reduced phase space suppression for W and Z emission during jet formation.

energy of the initiating quark or gluon. With jet substructure techniques, we can probe jet formation by studying the quantum properties of jets. Pushing this frontier beyond what is presented in Part II will require both experimental and theoretical advances³. Track reconstruction inside jets will play an increasingly important role in reconstructing jet substructure and will allow us to push boson and top quark tagging to the multi-TeV regime. The LHC has performed exceedingly well and the ATLAS collaboration, as a team, has shown that we can harness our detector to measure extreme energies with great precision. With more data and new ideas, together we will uncover the next clue in Nature’s captivating mystery.

Benjamin Philip Nachman
Geneva, Switzerland, July 2016

³Significant progress on track reconstruction inside high p_T jets between Run 1 and Run 2 will already boost sensitivity in the future [112].

Appendix A

Radiation Damage

As the closest subdetector to the interaction point, the ATLAS pixel detector will be exposed to an extreme amount of radiation over its lifetime ($\gtrsim 10^{15} \text{ n}_{\text{eq}}/\text{cm}^2$). The modules composing the detector are designed to be radiation tolerant, but their performance will degrade over time. It is therefore critical to model the impact of radiation damage for accurate simulation of tracking in the future. Including a radiation damage model is especially relevant for the high luminosity upgrade of the LHC; the instantaneous and integrated luminosity will significantly exceed current values, but simulations of the upgraded inner detector (ITK) [683] do not include the effects of radiation damage. This section briefly documents a digitization model¹ designed for the ATLAS software system that includes the impact of radiation damage. For a detailed account of the impact of radiation damage to silicon sensors, see Ref. [684] and Sec. 5 in Ref. [81]. The model described here includes two impacts of bulk defects: modifications to the electric field inside the sensors and charge trapping. Energy deposition in the silicon is modeled with GEANT4 and then various effects illustrated in Fig. A.1 are accounted for during *digitization*: the modeling of the detection and read-out of energy deposited by charged particles. A minimum ionizing particle (MIP) is a charged particle that has momentum corresponding to the minimum average energy loss per distance ($\langle dE/dx \rangle$) in a given material. Since the stopping power increases

¹This work is built on previous studies by many people and benefited from direct technical input from M. Benoit, M. Bomben, C. Bertsche, and R. Carney.

only logarithmically for several decades in momentum beyond the minimum, the definition of a MIP is extended to include momenta up to the point where radiative losses become important. For example, muons between about 1 GeV and 1 TeV in silicon are MIPs. When a MIP traverses silicon, it generates electron-hole pairs. The energy required to generate such pairs is about 3.6 eV (depends mildly on temperature). This leads to about 80 electron-hole pairs deposited per micron. The digitization model converts the energy deposited by GEANT4 into discrete charge clumps which are propagated to the electrode. The number of collected charges is then converted into a discrete *time over threshold* (TOT) value, which is exactly the same output of a real pixel module. Due to time constraints, it is not possible for each charge clump to represent a fundamental charge (electron or hole). The implications of this clumping are described at the end of this section.




Figure A.1: A schematic diagram illustrating the components of the digitizer model described in this section. While included in the model, thermal diffusion is a generally small effect and is not discussed in the following section.

After generating clumps of charge, several process are simulated as the clumps propagate to the electrode. The bias voltage applied to the sensor generates a large

electric field that causes the electrons (holes) to drift toward (away from) the collecting electrode. The velocity of this motion is determined by the charge carrier mobility μ via $\mathbf{v} = \mu \mathbf{E}$. The mobility has a small E-field and temperature dependence:

$$\mu^p(\mathbf{E}) = \frac{v_s^p / E_c^p}{\left(1 + \left(\frac{E}{E_c}\right)^{\beta^p}\right)^{1/\beta^p}}, \quad (\text{A.1})$$

where p stands for electron or hole. The values for the saturation velocity v_s , critical E-field E_c and temperature exponent β can be found in Table A.1. In addition to the electric field from the bias voltage, there is a magnetic field generated by the solenoid surrounding the inner detector. One effect from this field is that it modifies the mobility so that the average velocity of charge carriers follows the *Hall mobility*, which is the drift mobility (Eq. A.1) multiplied by the Hall factor r found in Table A.1. Another impact of this field is that the charges do not travel parallel to the electric field: they travel at an angle called the *Lorentz angle*. This angle is only relevant in the direction perpendicular to the beam and is approximately $\tan \theta \sim 0.2$ in the inner detector barrel.

quantity	electrons	holes
v_s ($\mu\text{m}/\text{ns}$)	$116 \times (T/273 \text{ K})^{-0.87}$	$88 \times (T/273 \text{ K})^{-0.52}$
E_c (kV/cm)	$6.0 \times (T/273 \text{ K})^{1.55}$	$15 \times (T/273 \text{ K})^{1.68}$
β	$1.0 \times (T/273 \text{ K})^{0.66}$	$1.1 \times (T/273 \text{ K})^{0.17}$
r	$1.13 + 8 \times 10^{-4} \times (T/K - 273)$	$0.72 - 5 \times 10^{-4} \times (T/K - 273)$

Table A.1: Physical constants describing the mobility of charge carriers in silicon. The first three rows are reformatted from Ref. [685] and the Hall scale factor is from Ref. [686].

The time for a charge chunk to reach the electrode is estimated by integrating the mobility:

$$t_{\text{electrode}} = \int_{z_{\text{initial}}}^{z_{\text{final}}} \frac{dz}{\mu_p(z)E(z)}, \quad (\text{A.2})$$

where z_{final} is the depth of the electrode (200 μm for the ATLAS IBL) for electrons and 0 for holes (the velocity is negative). Figure A.2 shows the average electric field as a function of depth in an IBL sensor with and without radiation damage. There are several models for simulating radiation damage effects on the electric field - for this study, the Chiochia model [687, 688] was simulated using a TCAD model². The main effect of irradiation is that the field strength in the center of the sensor is reduced while the field strength near the edges increases. Using these fields as input, Fig. A.3 shows the projected drift time from Eq. A.2.




Figure A.2: The average (over x and y) electric field as a function of the depth (z) inside an unirradiated planar 200 μm deep planar sensor with a bias voltage of 80 V (1000 V) for an unirradiated ($5 \times 10^{15} \text{ n}_{\text{eq}}/\text{cm}^2$) sensor on the left (right).

As a result of irradiation, defects form in the silicon and are sites for charge trapping. In the simulation, charge chunks are declared trapped if the projected time to reach the electrode from Fig. A.3 exceeds a random trapping time t that is exponentially distributed with mean value $1/(\kappa\Theta)$, where Θ is the fluence. The

²Input E-fields and Ramo potential maps are from M. Bomben.




Figure A.3: The time for an electron or hole to drift to the top (collecting electrode) or bottom of the sensor as a function of the depth (z) using the averaged E fields shown in Fig. A.2 for a planar sensor with a bias voltage of 80 V (1000 V) for an unirradiated ($5 \times 10^{15} \text{ n}_{\text{eq}}/\text{cm}^2$) sensor on the left (right).

constant κ (called β in the literature) has been measured at the 2001 CERN test beam and is approximately $\kappa = 3 \times 10^{-16} \text{ cm}^2/\text{ns}$ [689]. Charge trapping reduces the collected signal and thus degrades track reconstruction efficiency.

However, not all the trapped charge is lost. Charge is induced on the electrode as soon as the electrons or holes start to move. The amount of induced charge can be readily calculated using the *Ramo potential* from the Shockley-Ramo theorem [690, 691]. This theorem states that the amount of induced charge is the particle charge multiplied by the difference in the Ramo potential from its starting and ending (trapped) location. The Ramo potential for a particular electrode is calculated by calculating the electrostatic potential by holding the given electrode at unit voltage and setting all other electrodes to have zero potential. For example, for a an infinite parallel plate capacitor, the field is constant in between the plates, so the Ramo potential is linear (starting at 1 and decreasing to zero). Figure A.7 shows the Ramo potential for parallel plate capacitors that have various widths³. As the width decreases, the area over which the charge is collected becomes increasingly small, i.e. the Ramo

³Example inspired by Ref. [692].

potential is increasingly peaked at zero relative to the rest of the sensor. This trend is also illustrated in two dimensions in Fig. A.5. Note that since the Ramo potential extends beyond the extend of the sensor, charge is also induced in neighboring pixels.




Figure A.4: The Ramo potential at $x = y = 0$ (centered on the collecting electrode ϵ) as a function of the distance z away from ϵ for three sizes of ϵ : $W = L/10$, $W = L/2$, and $W = \infty$, where L is the sensor thickness.




Figure A.5: The Ramo potential at $y = 0$ (centered on the collecting electrode ϵ) as a function of the distance z away from ϵ for three sizes of ϵ : $W = L/10$, $W = L/2$, and $W = \infty$, where L is the sensor thickness.

By construction, without charge trapping, the total induced charge on the primary

electrode must be the total charge. Similarly, the charge induced on the neighboring electrodes must be zero without trapping. However, while the induced charge on the primary electrode increases monotonically with time, the charge induced on the neighboring electrodes increases when the electron or hole is far away and then decreases once it is close enough. To understand this, consider a point unit charge that is a distance z away from an infinite plate that has been cut into strips, where each strip is grounded. The surface charge density is peaked at zero and the peak increases the closer the charge is to the surface. The left plot in Fig. A.6 shows a transverse slice of the surface charge density. A comparison of the various colored lines in this plot shows how the field moves over the neighboring electrode. The middle plot in Fig. A.6 is the integral of the charge density on the electrode neighboring the primary one to the right and the induced charge calculated with the Ramo potential is shown in the right plot.




Figure A.6: A cross-section of the surface charge density from a point charge a distance z away from a series of grounded strips. The ‘primary electrode’ is the strip centered at 0 (extending to $\pm\infty$ in the direction into and out of the page) and the ‘neighboring electrode’ is the one shown between 2 and 3. The middle plot shows the total charge on the neighboring electrode as a function of the distance the charge is from the origin and the plot on the right shows the fraction of the charge that is induced on the neighbor using the Ramo potential.

Figures A.7 and A.8 show the final depth for electrons and holes that start at some depth and are trapped after a time t shown on the vertical axis. The induced charge is computed as the difference in the Ramo potential between the initial and final

depths⁴. The induced charge based on the unirradiated E-field is shown in Fig. A.11. In practice, the time to the trap is always ∞ in this case, but Fig. A.11 provides a technical closure of the setup. As expected, the charge induced on the primary electrode reaches 100% as the trapping time goes to infinity. The asymmetry with respect to the center of the detector is due in part to the difference in mobilities between electrons and holes. Similarly, the charge induced on the neighboring electrodes goes to zero as time goes to infinity. The induced charge is much larger for the electrode that is only 50 μm away (short direction) compared with the one that is 200 μm away (long direction).




Figure A.7: The final location (based only on drift) for electrons starting at a depth z and traveling a time given by the vertical axis. The coordinate z is measured with respect to the back-side (away from the collecting electrode) of the sensor using the averaged E fields shown in Fig. A.2 for a planar sensor with a bias voltage of 80 V (1000 V) for an unirradiated ($5 \times 10^{15} \text{ n}_{\text{eq}}/\text{cm}^2$) sensor on the left (right).

One last effect that can play an important role in modeling the collected charge is related to charge chunking. Representing many fundamental charges as one multi-charged chunk does not change the average charge collected, but does impact the charge collection resolution⁵. Figure A.10 quantifies how the resolution increases as more and more fundamental charges are combined together into one chunk. This

⁴The Ramo potential depends only on geometry and not on the fluence [693].

⁵Thank you to M. Garcia-Sciveres for pointing this out.




Figure A.8: The final location (based only on drift) for holes starting at a depth z and traveling a time given by the vertical axis. The coordinate z is measured with respect to the back-side (away from the collecting electrode) of the sensor using the averaged E fields shown in Fig. A.2 for a planar sensor with a bias voltage of 80 V (1000 V) for an unirradiated ($5 \times 10^{15} \text{ n}_{\text{eq}}/\text{cm}^2$) sensor on the left (right).




Figure A.9: The charge induced on the primary (left) and neighboring electrodes (short direction in the middle and long direction on the right) computed using the Ramo potential.

effect can be corrected by using a method inspired by the forward-folding method from Sec. 4.1.2.3. In particular, if X is a random variable with mean μ and standard deviation σ , then $Y = \mu + \kappa(X - \mu)$ will have mean μ and standard deviation $\kappa\sigma$. The resolution for chunks is larger than for fundamental charges, so $\kappa \leq 1$ (i.e. unsmearing is required). In this case, $\kappa = 1/\mathbf{n}$, where $\mathbf{n}$ is the number of fundamental charges

that one chunk represents. The average value is also known: for a charge Q , the average charge that will be collected is $e^{-t_{\text{electrode}}/t_{\text{trapping}}}Q$. Therefore, the corrected charge is

$$Q \times \delta_{\text{collected}} \mapsto Q \left[e^{-t_{\text{electrode}}/t_{\text{trapping}}} + \kappa \left(\delta_{\text{collected}} - e^{-t_{\text{electrode}}/t_{\text{trapping}}} \right) \right], \quad (\text{A.3})$$

where Q is the charge per chunk and $\delta_{\text{collected}}$ is one if the charge is collected and zero if it is trapped. Note that even if a charge is trapped, it will still contribute to the collected charge.




Figure A.10: An illustration of the impact of representing multiple fundamental charges by a single multi-charged chunk. Displayed is the fraction of collected charge as a function of the exponential trapping time constant. N chunks are propagated to the electrode and some fraction f of them reach the electrode before a random exponential time. If a chunk reaches the electrode, a charge of Q is recorded. The points show $\langle fQN \rangle$ and the error bars are the standard deviation. Red corresponds to $Q = 1$ and blue corresponds to $Q = 100$ (left), $Q = 10$ (middle), and $Q = 1$ (right).

In order to validate the radiation damage model, modules with and without a significant radiation dose⁶ are tested with a dedicated testbeam at the SLAC End Station A⁷. A kicker magnet extracts a 5 Hz electron beam that is incident on a

⁶Irradiated at Ljubljana with neutrons to $5 \times 10^{15} \text{ n}_{\text{eq}}/\text{cm}^2$.

⁷Thank you to M. Benoit for helping acquiring the samples and to Su Dong for an extensive amount of time in the lab/testbeam for preparations and operations. Thank you also to M. McCulloch and R. Carney for help with the setup and operations.

copper target and focused to produce an 11 GeV electron beam. A *telescope* of six planes with Mimosa26 [694] sensors allow for $\mathcal{O}(\mu\text{m})$ precision tracking. Three of these planes are on either side of a gap for the Device Under Test (DUT), which in this case is an irradiated or unirradiated IBL-like planar FEI4 module. Figure A.11 shows the setup inside the SLAC beamline.




Figure A.11: Pictures of the testbeam setup. The left picture shows the six telescope planes, three on each side of the DUT. The unirradiated module does not need to be cooled and sits on a stage that can be automatically raised, lowered, and tilted. The beam goes from left to right, piercing the center of the black squares shown in the middle of each telescope plane. The blue, red, and black cables from the DUT are for the data, low voltage, and high voltage, respectively. The right picture shows a close-up of the irradiated module inside the box used to keep it cool (with dry ice). The orange cables are temperature sensors.

One of the most striking features of the irradiated sensors is the predicted electric field, as shown in Fig. A.2. To expose this dependence, the modules are rotated and tilted so that instead of the electrons traversing the 200 μm depth of the sensor, they pass through the 50 μm edge. Figure A.12 illustrates this configuration: electron-hole pairs from pixels near the beginning or end of the cluster will probe the (large) field closest to and furthest away from the collecting electrode, while those passing through the center of the cluster will see the lower electric field in the middle of the sensor. The angle corresponds to a cluster length of 15 pixels. Representative event displays from sensors with and without irradiation are shown in Fig. A.13. As expected, particles form long streaks in the short pixel direction (Y in these rotated

coordinates). Figure A.14 shows how the TOT distribution depends on the position inside one of the long clusters. For the unirradiated module, the charge distribution is nearly independent of the position inside the cluster, as expected. This also seems to be true for clusters of length 10 inside the irradiated module. Even though the tilt angle is the same for both modules, the irradiated sensor is not fully depleted, so the clusters do not reach the full length. Part of the degradation to the electric field in the middle of the sensor is compensated by the induced charge from the Ramo potential. Further studies using more sophisticated clustering algorithms that can account for gaps in the clusters (from pixels below threshold) may reveal a structure that will be useful for tuning the simulation.




Figure A.12: A illustration of the tilted sensor configuration. A charged particle travels from the bottom left to the top right and traverses about $50\text{ }\mu\text{m}$ of silicon in each sensor. The tilt angle is chosen so that the particle will traverse about 15 pixels, probing different depths along its path.




Figure A.13: Representative event displays from the unirradiated (left) and the irradiated (right) sensors. In both cases, there are three clusters. The sensors are tilted (Fig. A.12) so that one expects about 15 pixel clusters per particle. The unirradiated sensor uses a 2000 electron threshold and 8 TOT is tuned to 11000 collected electrons, while the unirradiated sensor uses a 1500 electron threshold with 8000 electrons corresponding to 5 TOT. The unirradiated sensor is biased with 60 V while the irradiated sensor is biased by 1.1 kV and cooled to about -35°C .




Figure A.14: The TOT distribution as a function of depth inside clusters of length exactly 15 (left) or 10 (right) for the unirradiated and irradiated sensors from Fig. A.13. The tilt angle is the same for both modules, but the pixel size distribution was peaked at lower values (10 versus 15) for the irradiated sensor, indicating that it is not fully depleted (even at 1.1 kV).

Appendix B

Boson Polarizations

B.1 Angular Distributions in W and Z Boson Decays

Without loss of generality, assume that the W or Z boson is moving in the $+z$ direction with momentum $\mathbf{p}_V^\mu = (0, 0, p_z, \sqrt{p_z^2 + m_V^2})$. The massive gauge bosons have three polarization states corresponding to a spin that is anti-aligned (transverse, spin -1), aligned (transverse, spin $+1$), or orthogonal (longitudinal, spin 0) to the boson momentum. The three corresponding polarization vectors are $\epsilon_{-1}^\mu = \frac{1}{\sqrt{2}}(0, 1, -i, 0)$, $\epsilon_{+1}^\mu = -\frac{1}{\sqrt{2}}(0, 1, i, 0)$, and $\epsilon_0^\mu = \frac{1}{m_V}(p_z, 0, 0, \sqrt{p_z^2 + m_V^2})$. The weak charged and neutral currents have the form

$$j^\mu = \bar{u}(f) \frac{1}{2} \gamma^\mu (c_V^f - c_A^f \gamma^5) v(\bar{f}'), \quad (\text{B.1})$$

where the W boson only couples to left-handed fermions and right-handed anti-fermions, whereas the Z boson couples to both left- and right-handed fermions, but with unequal couplings $\mathbf{c}$. The factors $\bar{u}$ and v are the spinors for the out-going fermion f and the out-going anti-fermion $\bar{f}'$. For the W decay, $c_A = c_V = 1$ (i.e. the parenthetical term in Eq. B.1 is a pure projection operator) and for the Z decay, $c_A = \frac{1}{2}$ for up-type quarks (u, c, t) and $-\frac{1}{2}$ for down-type quarks (d, s, b) while $c_V = \frac{1}{2} - 2 \times \frac{2}{3} \times \sin^2 \theta_W \approx 0.19$ for up-type quarks and $\frac{1}{2} + 2 \times \frac{1}{3} \times \sin^2 \theta_W \approx -0.35$ for down-type quarks. In terms of pure left- and right-handed out-going fermions, one can write $j^\mu = c_{L/R} \bar{u}(f) \gamma^\mu v(\bar{f}')$, where the couplings $c_{L/R}$ can be extracted from Eq. B.1 using projection operators and are given by $c_L \approx -0.35$ for up-type quarks, $c_L \approx -0.42$ for down-type quarks, $c_R \approx -0.15$ for up-type quarks, and $c_R \approx 0.08$ for down-type quarks. The matrix element is $M^2 \propto |\epsilon_\mu j^\mu|^2$, where the proportionality constant is a coupling factor for the weak vertices multiplied by the number of colors N_C . Using the setup shown in Fig. B.1 and working in the boson rest frame, the momentum of the fermions are $\mathbf{p}_f^\mu = \frac{m_V}{2}(\sin \theta, 0, \cos \theta, 1)$ and $\mathbf{p}_{f'}^\mu = \frac{m_V}{2}(-\sin \theta, 0, -\cos \theta, 1)$ (ignoring the fermion mass). With these momenta, the leading order matrix element is given by (see e.g. the polarized $e^+e^- \rightarrow \mu^+\mu^-$ calculations from Ref. [20]):

$$|M_{-1}^W|^2 = \frac{3g_2^2 m_W^2 V^2}{4} (1 + \cos \theta)^2 \quad (\text{B.2})$$

$$|M_0^W|^2 = \frac{3g_2^2 m_W^2 V^2}{2} \sin^2 \theta \quad (\text{B.3})$$

$$|M_{+1}^W|^2 = \frac{3g_2^2 m_W^2 V^2}{4} (1 - \cos \theta)^2, \quad (\text{B.4})$$

where V is an element of the CKM matrix. For Z bosons,

$$|M_{-1}^Z|_{\text{up}}^2 = \frac{3g_2^2 m_Z^2}{2 \cos^2 \theta_W} [c_{L,\text{up}}^2 (1 + \cos \theta)^2 + c_{R,\text{up}}^2 (1 - \cos \theta)^2] \quad (\text{B.5})$$

$$|M_0^Z|_{\text{up}}^2 = \frac{3g_2^2 m_Z^2}{\cos^2 \theta_W} [c_{L,\text{up}}^2 + c_{R,\text{up}}^2] \sin^2 \theta \quad (\text{B.6})$$

$$|M_{+1}^Z|^2 = \frac{3g_2^2 m_Z^2}{2 \cos^2 \theta_W} [c_{L,\text{up}}^2 (1 - \cos \theta)^2 + c_{R,\text{up}}^2 (1 + \cos \theta)^2], \quad (\text{B.7})$$

and the equivalent formula for down-type quarks but with $\text{up} \leftrightarrow \text{down}$.




Figure B.1: A diagram illustrating the setup for the calculation described in the text. The boson spin is along the z -axis.

B.2 Polarization of W Bosons

A calculation similar to Appendix B.1 can be used to determine the fractions of transverse and longitudinally polarized W bosons from various production modes. To begin, consider W bosons produced from top quark pair production. Consider a top quark decay $t \rightarrow W^+ b$ from the top quark rest frame with the top spin aligned along the $+z$ axis. For illustration, suppose that the W^+ and b momenta are (anti-)parallel to the $+z$ axis and that the top quark spin is $+\frac{1}{2}$. Neglecting the b -quark mass, there are only two possibilities: (a) the b is moving in the $-z$ direction with spin $+\frac{1}{2}$ (left-handed) and by conservation of angular momentum, the W boson is longitudinally polarized and (b) the b -quark is moving in the $+z$ direction with spin $-\frac{1}{2}$ (left-handed) and by conservation of angular momentum, the W boson has spin $+1$. The matrix element is given by

$$M = g_2 m_W \bar{u}(b) \epsilon_\mu^*(W^+) \gamma^\mu \frac{1}{2} (1 - \gamma^5) u(t) = g_2 m_W u^\dagger(b) \gamma^0 \epsilon_\mu^*(W^+) \gamma^\mu \frac{1}{2} (1 - \gamma^5) u(t). \quad (\text{B.8})$$

In the chiral basis (the one used by Ref. [20]),

$$\gamma^\mu = \left[\begin{pmatrix} 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \\ 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \end{pmatrix}, \begin{pmatrix} 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \\ 0 & -1 & 0 & 0 \\ -1 & 0 & 0 & 0 \end{pmatrix}, \begin{pmatrix} 0 & 0 & 0 & -i \\ 0 & 0 & i & 0 \\ 0 & i & 0 & 0 \\ -i & 0 & 0 & 0 \end{pmatrix}, \begin{pmatrix} 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & -1 \\ -1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \end{pmatrix} \right], \quad (\text{B.9})$$

and

$$u(p) = \frac{1}{2} \begin{pmatrix} (\mathbb{I}_2 - \hat{p} \cdot \sigma) \xi \\ (\mathbb{I}_2 + \hat{p} \cdot \sigma) \xi \end{pmatrix}, \quad (\text{B.10})$$

where $\mathbb{I}_2$ is the 2×2 identity matrix and ξ is a two-component spinor. The top

quark is at rest and has $\xi = (1, 0)^T$ (spin up) so $\mathbf{u}(\mathbf{t}) = \sqrt{2\mathbf{m}_t}(1, 0, 0, 0)^T$. When the $\mathbf{b}$ -quark is spin up, it is moving in the $-z$ direction so $\hat{\mathbf{p}} \cdot \boldsymbol{\sigma} = -\sigma^3$. Therefore, $\mathbf{u}(\mathbf{t}) = \sqrt{2\mathbf{E}_b}(1, 0, 0, 0)^T$. In contrast, when the $\mathbf{b}$ -quark is spin down ($\xi = (0, 1)^T$), it is moving in the $+z$ direction and so $\mathbf{u}(\mathbf{t}) = \sqrt{2\mathbf{E}_b}(0, 1, 0, 0)^T$. The longitudinal W^+ polarization vector is $\frac{1}{m_W}(\mathbf{p}_W, 0, 0, E_W)$ and the transverse (spin -1) W^+ polarization vector is $\frac{1}{\sqrt{2}}(0, 1, -i, 0)$. Putting all of these pieces together with Eq. B.8 produces the following results (dropping constants appearing in both terms):

$$M_0 \propto \frac{1}{m_W} \begin{pmatrix} 0 & 0 & 1 & 0 \end{pmatrix} \begin{pmatrix} 0 & 0 & E_W - p_W & 0 \\ 0 & 0 & 0 & E_W + p_W \\ E_W + p_W & 0 & 0 & 0 \\ 0 & E_W - p_W & 0 & 0 \end{pmatrix} \begin{pmatrix} 1 \\ 0 \\ 0 \\ 0 \end{pmatrix} \quad (\text{B.11})$$

$$= \frac{(E_W + p_W)}{m_W} = \frac{m_t}{m_W}, \quad (\text{B.12})$$

where the last equality holds because $\mathbf{p}_W = \mathbf{p}_b$ in the top quark rest frame by conservation of momentum and by conservation of energy, $m_t = E_b + E_W = \mathbf{p}_b + E_W = \mathbf{p}_W + E_W$ (ignoring the $\mathbf{b}$ -quark mass). Likewise,

$$M_- \propto \frac{1}{\sqrt{2}} \begin{pmatrix} 0 & 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 2 & 0 \\ 0 & 0 & 0 & 0 \\ -2 & 0 & 0 & 0 \end{pmatrix} \begin{pmatrix} 1 \\ 0 \\ 0 \\ 0 \end{pmatrix} \quad (\text{B.13})$$

$$= \frac{2}{\sqrt{2}} \quad (\text{B.14})$$

Therefore, the ratio of the number of longitudinally polarized W^+ bosons to transversely polarized W^+ bosons is

$$\frac{|M_0|^2}{|M_-|^2} = \frac{m_t^2}{2m_W^2} \approx 2.3, \quad (\text{B.15})$$

so there are *more* longitudinally polarized W bosons from top quark decays relative to transversely polarized W bosons. The above observation is true even if the $\mathbf{b}W$ axis is not aligned with the top quark spin axis and has been computed at NNLO in QCD to be 0.21 ± 0.05 [695]. In contrast, W bosons produced via inclusive W +jets processes, are predominately produced with a transverse polarization. This is nicely explained, along with studies of higher order QCD effects, in Ref. [696].

Appendix C

Additional Statistical Considerations

C.1 Uncertainty Ellipses

In one dimension, for a random variable $X \sim \mathcal{N}(\mu, \sigma^2)$, $(X - \mu)/\sigma \sim \mathcal{N}(0, 1)$ and so an interval centered at the mean that contains p -percent of the probability distribution of X is given by $\mu \pm Z\sigma$, where $p = \frac{1}{\sqrt{2\pi}} \int_{-Z}^Z dx \exp(-x^2/2)$. An equivalent way to arrive at the same interval that generalizes to higher dimensions is to note that $(X - \mu)^2/\sigma^2 \sim \chi_1^2$, a chi-squared distribution with one-degree of freedom. Then, the same interval can be constructed as $\mu \pm \sqrt{C}\sigma$, where $p = \frac{1}{\sqrt{2\pi}} \int_0^C dx x^{-1/2} \exp(-x/2)$. Now, suppose that $\vec{X} \sim \mathcal{N}(\vec{\mu}, \Sigma)$ where μ is an n -dimensional vector and Σ is the $n \times n$ covariance matrix (symmetric, positive semi-definite). Then, $(\vec{X} - \vec{\mu})^T \Sigma^{-1} (\vec{X} - \vec{\mu}) \sim \chi_n^2$, a chi-squared distribution with n degrees of freedom. When $n = 1$, this reduced to the one-dimensional case above. An ellipsoid centered about the mean which contains p -percent of the probability distribution is then given implicitly by $(\vec{x} - \vec{\mu})^T \Sigma^{-1} (\vec{x} - \vec{\mu}) \leq C$, where $p = \frac{1}{2^{n/2} \Gamma(n/2)} \int_0^C dx x^{n/2-1} \exp(-x/2)$. In two dimensions $(X, Y) \sim \mathcal{N}((\mu_x, \mu_y), \Sigma)$, this is an ellipse. One can write

$$\Sigma = \begin{pmatrix} \sigma_x^2 & \rho \sigma_x \sigma_y \\ \rho \sigma_x \sigma_y & \sigma_y^2 \end{pmatrix}, \quad (\text{C.1})$$

where ρ is the correlation between X and Y . The matrix Σ is diagonalizable such that after a suitable rotation of X and Y ,

$$\Sigma = \begin{pmatrix} \lambda_+ & 0 \\ 0 & \lambda_- \end{pmatrix}, \quad (\text{C.2})$$

where $\lambda_{\pm}$ are the eigenvalues of Σ and are found by solving $\text{Det}(\Sigma - \lambda I) = 0$:

$$\lambda_{\pm} = \frac{1}{2}(\sigma_x^2 + \sigma_y^2) \pm \frac{1}{2}\sqrt{(\sigma_x^2 + \sigma_y^2)^2 - 4(1 - \rho^2)\sigma_x^2\sigma_y^2}. \quad (\text{C.3})$$

Let $\mathbf{X}'$ and $\mathbf{Y}'$ be the centered and rotated versions of $\mathbf{X}$ and $\mathbf{Y}$. In these transformed coordinates, the uncertainty ellipse is given by

$$\frac{(\mathbf{x}')^2}{C\lambda_+} + \frac{(\mathbf{y}')^2}{C\lambda_-} \leq 1, \quad (\text{C.4})$$

which is the standard form of an ellipse with radii $\sqrt{C\lambda_{\pm}}$. The tilt of the ellipse with respect to the original coordinates can be computed from the orientation of the eigenvectors $\mathbf{v}_{\pm}$ of Σ . For example, $\Sigma\mathbf{v}_+ = \lambda_+\mathbf{v}_+$ gives the condition

$$(\mathbf{v}_+)_x = \left(\frac{\rho\sigma_x\sigma_y}{\lambda_+ - \sigma_x^2} \right) (\mathbf{v}_+)_y. \quad (\text{C.5})$$

Therefore, the rotation angle counter close-wise from the x -axis is $\theta = \tan^{-1}[(\lambda_+ - \sigma_x^2)/(\rho\sigma_x\sigma_y)]$. As expected, as $\rho \rightarrow 0$, $\lambda_+ \rightarrow \sigma_x^2$, $\lambda_- \rightarrow \sigma_y^2$ and $\theta \rightarrow 0$. Figure C.1 shows the general form of the uncertainty ellipse.




Figure C.1: A schematic diagram of an uncertainty ellipse for a bivariate normal distribution. See the text for details.

C.2 Wrapped Gaussian

The EM algorithm (Sec. 4.3.1.2) depends on the event topology. For instance, if a Gaussian density is used to model ϕ , then, in the E step, a particle with ϕ_i near 2π will be deemed far from a cluster with location ϕ_j near 0. To avoid this undesirable behavior and enforce the equivalence of the angles 0 and 2π , ϕ is associated with a *wrapped Gaussian density* and $\mathbf{y}$ with a standard Gaussian density:

$$\Phi(\mathbf{y}, \phi | \mu_\phi, \mu_y, \sigma^2) = \Phi_y(\mathbf{y} | \mu_y, \sigma^2) \frac{1}{\sqrt{2\pi\sigma^2}} \sum_{I=-\infty}^{\infty} \exp \left[\frac{-(\phi - \mu_\phi(I))^2}{2\sigma^2} \right], \quad (\text{C.6})$$

where Φ_y is a normal distribution and $\mu_\phi(I) = \mu_\phi + 2\pi I$. In order to approximate the sum in Eq. (C.6), only the leading contribution is retained by choosing $\mu_\phi(I^*)$ for $I^* = \text{argmin}_I |\phi - \mu_\phi + 2\pi I|$. Other contributions are exponentially suppressed and this part recovers continuity near 0 and 2π , as illustrated in Fig. C.2.




Figure C.2: A three-particle event display illustrating the results of fuzzy jet clustering using a Gaussian density for ϕ (left) and a wrapped Gaussian density approximation for ϕ (right). Figure from C. Stansbury.

C.3 The EM algorithm

This appendix contains two derivations: the modified EM algorithm updates in Eq. (4.41) and the proof that the modified EM algorithm generically improves the original modified log likelihood Eq. (4.39) with every iteration. Recall the expected modified complete log likelihood (mmCLL) from Eq. (4.40):

$$\sum_{i=1}^n \sum_{j=1}^k p_{\tilde{\pi}}^{\alpha} (q_{ij} \log \Phi(\vec{\rho}_i; \vec{\mu}_j, \Sigma_j) + q_{ij} \log \pi_j).$$

Viewing the mCLL as a function of $\vec{\mu}, \Sigma$ and π for fixed λ and $\vec{\rho}$ we can maximize. For π , we optimize

$$\sum_{i=1}^n \sum_{j=1}^k p_{\tilde{\pi}}^{\alpha} (q_{ij} \log \pi_j) + \lambda \left(\sum_{j=1}^k \pi_j - 1 \right),$$

where the last term is needed so that the optimal π^* is a probability. The derivative of this expression with respect to π_j is

$$\pi_j = -\frac{1}{\lambda} \sum_{i=1}^n p_{\tilde{\pi}}^{\alpha} q_{ij},$$

and then summing the equation over j and using $\sum_{j=1}^k q_{ij} = 1$ and the constraint equation $\sum_{j=1}^k \pi_j = 1$, we find that

$$\pi_j^* = \frac{1}{\sum_{i=1}^n p_{\tilde{\pi}}^{\alpha}} \sum_{i=1}^n p_{\tilde{\pi}}^{\alpha} q_{ij}$$

The updates for $\vec{\mu}$ and Σ follow from the standard derivation (by similarly taking

derivatives of the mCLL with respect to components of these multi-dimensional objects) by noting that the only difference is that $\mathbf{q}_{ij} \mapsto \mathbf{q}_{ij} \mathbf{p}_{\Pi}^{\alpha}$ and there are no Lagrange multipliers needed unlike for π_j^* .

Finally, we prove the claim that the modified EM algorithm described in the body of the text monotonically improves the modified log likelihood in Eq. (4.39). First, we note that we can rewrite the (log) likelihood as

$$\begin{aligned} \mathbf{p}_{\Gamma}^{\alpha} \log \mathbf{p}(\rho|\theta) &= \mathbf{p}_{\Gamma}^{\alpha} \log \left(\sum_{\lambda \in \{1,2,\dots,k\}} \mathbf{p}(\rho, \lambda; \theta) \right) \\ &= \mathbf{p}_{\Gamma}^{\alpha} \log \left(\sum_{\lambda \in \{1,2,\dots,k\}} \frac{\mathbf{q}(\lambda) \mathbf{p}(\rho, \lambda; \theta)}{\mathbf{q}(\lambda)} \right) \\ &= \mathbf{p}_{\Gamma}^{\alpha} \log \mathbb{E}_{\mathbf{q}} \left[\frac{\mathbf{p}(\rho, \lambda; \theta)}{\mathbf{q}(\lambda)} \right] \\ &\geq \mathbb{E}_{\mathbf{q}} \left[\mathbf{p}_{\Gamma}^{\alpha} \log \left(\frac{\mathbf{p}(\rho, \lambda; \theta)}{\mathbf{q}(\lambda)} \right) \right] \equiv \mathcal{L}(\mathbf{q}, \theta), \end{aligned}$$

where the inequality in the last line follows from Jensen's inequality. Now, we are ready to prove the claim that $\mathbf{p}_{\Gamma}^{\alpha} \mathbf{p}(\rho|\theta^{(t)})$ improves monotonically with $\mathbf{t}$, the index for the iteration of the EM algorithm. First, note that

$$\begin{aligned} \mathcal{L}(\mathbf{q}, \theta) &= \mathbb{E}_{\mathbf{q}} \left[\mathbf{p}_{\Gamma}^{\alpha} \log \left(\frac{\mathbf{p}(\rho, \lambda; \theta)}{\mathbf{q}(\lambda)} \right) \right] \\ &= \mathbb{E}_{\mathbf{q}} [\mathbf{p}_{\Gamma}^{\alpha} \log (\mathbf{p}(\rho, \lambda; \theta))] - \mathbb{E}_{\mathbf{q}} [\mathbf{p}_{\Gamma}^{\alpha} \log (\mathbf{q}(\lambda))], \end{aligned}$$

where the first term is the mCLL and the second term has no θ dependence and so maximize $\mathcal{L}(\mathbf{q}, \theta)$ over θ is equivalent to maximize the mCLL over θ . Therefore, $\mathcal{L}(\mathbf{q}^{(t+1)}, \theta^{(t)}) \leq \mathcal{L}(\mathbf{q}^{(t+1)}, \theta^{(t+1)})$. By the inequality above, $\mathcal{L}(\mathbf{q}^{(t+1)}, \theta^{(t+1)}) \leq \mathbf{p}_{\Gamma}^{\alpha} \mathbf{p}(\rho|\theta^{(t+1)})$. The E step can be recast as choosing

$$\mathbf{q}^{(t+1)}(\lambda_i = j) = \mathbf{q}_{ij}(\boldsymbol{\theta}^{(t)}) = \mathbb{E}_{\boldsymbol{\theta}^{(t)}}[\mathbf{q}_{ij}] = \mathbf{p}(\lambda|\boldsymbol{\rho}, \boldsymbol{\theta}^{(t)}).$$

This enforces:

$$\begin{aligned} \mathcal{L}(\mathbf{p}(\lambda|\boldsymbol{\rho}, \boldsymbol{\theta}^{(t)}), \boldsymbol{\theta}^{(t)}) &= \mathbb{E}_{\mathbf{p}(\lambda|\boldsymbol{\rho}, \boldsymbol{\theta}^{(t)})} \left[\mathbf{p}_T^\alpha \log \left(\frac{\mathbf{p}(\boldsymbol{\rho}, \lambda; \boldsymbol{\theta}^{(t)})}{\mathbf{p}(\lambda|\boldsymbol{\rho}, \boldsymbol{\theta}^{(t)})} \right) \right] \\ &= \mathbb{E}_{\mathbf{p}(\lambda|\boldsymbol{\rho}, \boldsymbol{\theta}^{(t)})} [\mathbf{p}_T^\alpha \log (\mathbf{p}(\boldsymbol{\rho}; \boldsymbol{\theta}^{(t)}))] \\ &= \mathbf{p}_T^\alpha \log (\mathbf{p}(\boldsymbol{\rho}; \boldsymbol{\theta}^{(t)})) \end{aligned}$$

Putting this together with the bounds from the M step, we arrive at the desired result: $\mathbf{p}_T^\alpha \mathbf{p}(\boldsymbol{\rho}|\boldsymbol{\theta}^{(t)}) \leq \mathbf{p}_T^\alpha \mathbf{p}(\boldsymbol{\rho}|\boldsymbol{\theta}^{(t+1)})$, i.e., every step of the modified EM algorithm improves or leaves the same the original likelihood.

C.4 A Leading Order Description of Fuzzy Jet σ

Section 4.3.1.3.4 demonstrated that the fuzzy jet σ is correlated with $\rho = m/p_T$. One can build some intuition for this relationship by considering a leading order QCD calculation of σ . Consider an isolated quark jet with energy E which radiates a gluon with angle $\theta \ll 1$ from the jet axis and with energy fraction $z \ll 1$. Without loss of generality, suppose the quark is moving in the $\phi = 0$ direction and the splitting happens in the $\phi = \pi/2$ direction so that the four vector of the quark is $q^\mu = E(1-z)(1, 0, 0, 1)$, and the gluon four-vector is $g^\mu = Ez(1, \theta, 0, 1)$, to leading order. To this order, the jet mass is simply $m = Ez\theta^2$. What is σ ? Consider $k = 1$ and something like the event-jet applied so that we can treat this jet in isolation from other hadronic activity in the event. Since $k = 1$, the soft memberships are all one, i.e., $q_{i1} = 1$ and there is only one step of the EM algorithm. The anti- k_t jet has (y, ϕ) coordinates $(0, \theta)$, which could be used for the seed, but since $k = 1$, the seed is not used. The quark has coordinates $(0, 0)$, and the gluon has coordinates $(0, \theta)$. One can compute the fuzzy jet coordinates in the (single) M step:

$$\mu_y = 0 \tag{C.7}$$

$$\mu_\phi = \frac{0 \times E(1-z) + \theta \times Ez}{E(1-z) + Ez} = z\theta \tag{C.8}$$

$$\sigma^2 = \frac{(0 - z\theta)^2 \times E(1-z) + (\theta - z\theta)^2 \times Ez}{2(E(1-z) + Ez)} \tag{C.9}$$

$$= z\theta^2 + \mathcal{O}(\theta^2 z^2). \tag{C.10}$$

Therefore, to leading order and $k = 1$, the learned σ is the jet mass. For $k = 2$, there are enough degrees of freedom to resolve the substructure of the hard splitting and so the relationship between the jet mass and σ breaks down.

C.5 Controlling Jet Multiplicity with p_T

In contrast to most uses of hierarchical-agglomerative clustering algorithms, the number of fuzzy jets is fixed before clustering begins. Whereas a single traditional jet can reasonably be considered to correspond to a parton in appropriate cases, mGMM jets should not be, as several mGMM jets can together express structure of what would be one or several jets according to another algorithm. The choice of the number of jets used in mGMM jet clustering therefore controls the expressive power of the algorithm to look at the event structure. In practice, choosing too many jets does not greatly affect the value of the leading learned σ variable, because the additional jets learn finer features of the event structure. On the other hand, choosing too few jets is often problematic as can be seen in Figure C.3 - the fuzzy jets need to grow in order to cover the full energy distribution in the event. Using anti- k_t jets as seeds for fuzzy jets has the feature that the number of fuzzy jets change dynamically with the complexity of the event. The algorithm is not very sensitive to the exact locations of the anti- k_t jets - studies which randomly perturbed the initial jet locations inside a disc of radius 1.0 found that σ was robust to such fluctuations, even on an event by event basis. However, the p_T threshold for the seed anti- k_t jets can have a significant impact on the fuzzy jets as this alters the number of seeds. The p_T threshold for the anti- k_t seeds is typically lower than the p_T threshold one would use to consider anti- k_t jets alone because the fuzzy jets algorithm needs enough seeds to populate the low energy regions of the detector. One way of mitigating the impact of the p_T cut on the fuzzy jet clustering is to introduce an *event jet*, described in Section 4.3.1.4.




Figure C.3: Changing the choice of the p_T cut used to select seeds can make a vast difference in the values of the constructed variables, like σ . In this event, clustered on the left with a cut of 5 GeV resulting in five jets, and on the right with a cut of 50 GeV resulting four jets. Fewer degrees of freedom in the four jet case means a much larger learned value for the σ variable. Figures from C. Stansbury.

C.6 Computation of Significance Variables

This section briefly describes how to numerically compute the significance variables introduced in Sec. 3.2.2. If Gaussian approximations to the input object resolution functions are valid and the observable is sufficiently close to a linear combination of the input object kinematic quantities, then an analytic approximation using linear error propagation should be sufficient. However, to capture non-Gaussian attributes or important non-linear behavior of the kinematic variable, numeric propagation may be necessary. In particular, if $\mathbf{m}$ is a mass-like variable with a restriction $\mathbf{m} > 0$, the resolution function will necessarily be non-Gaussian near $\mathbf{m} = 0$. In such cases, one can estimate how many random draws are necessary to accurately compute $\sigma_{\mathbf{m}}$. If s^2 is the sample variance, then the variance of the sample variance is given by Eq. C.11, where κ is the excess kurtosis [697].

$$\text{Var}[s^2] = \sigma^4 \left(\frac{2}{n-1} + \frac{\kappa}{n} \right). \quad (\text{C.11})$$

For an absolute uncertainty on the standard deviation f and an $\mathcal{O}(1)$ standard deviation, one needs

$$n = \frac{2 + \kappa + f^2 + \sqrt{4 + 4\kappa + 4f^2 + \kappa^2 - 2f^2\kappa + f^4}}{2f^2}. \quad (\text{C.12})$$

For $f \ll 1$ and an order 1 or smaller κ (this is zero for a Gaussian),

$$n \approx \frac{2 + \kappa + \sqrt{4 + 4\kappa + \kappa^2}}{2f^2} \sim \frac{3}{f^2}. \quad (\text{C.13})$$

For example, one needs $n \approx 300$ for an accuracy of 0.1 GeV.

C.7 The Non-closure of Numerical Inversion

The jet calibration procedures of ATLAS [289] and CMS [349] involve several steps to correct for pileup, the non-linear detector response, the η -dependence of the jet response, flavor-dependence of the jet response, and residual data/simulation differences in the jet response. The simulation-based corrections to correct for the calorimeter non-linearities in $\mathbf{p}_T$ and η are accounted for using *numerical inversion*. Let X will be a random variable representing the particle-jet $\mathbf{p}_T$ and Y will be a random variable representing the reconstructed jet $\mathbf{p}_T$. Define¹

$$f(x) = \mathbb{E}[Y|X = x] \quad (\text{C.14})$$

$$R(x) = \mathbb{E}\left[\frac{Y}{x} \middle| X = x\right] = \frac{f(x)}{x}. \quad (\text{C.15})$$

Often, the normal approximation is valid: $Y|X = x \sim \mathcal{N}(f(x), \sigma(x))$, where this notation means ‘ Y given $X = x$ is normally distributed with mean $f(x)$ and standard deviation $\sigma(x)$.’ The function $R(x)$ is called the *response function*. Formally, numerical inversion is the following procedure:

1. Compute $f(x)$, $R(x)$.
2. Let $\tilde{R}(Y) = R(f^{-1}(Y))$.
3. Apply a jet-by-jet correction: $Y \mapsto Y/\tilde{R}(Y)$.

The intuition for the second step is that $f^{-1}(Y)$ is an estimate for x and then $R(f^{-1}(Y))$ is an estimate for the response at the estimate of x that gives rise to Y . Note that $\mathbb{E}[X|Y]$ is not useful instead of $f^{-1}(Y)$ because the former depends on $\mathbf{p}(x)$, the underlying $\mathbf{p}_T$ spectrum, whereas f (and thus f^{-1}) to do not depend on $\mathbf{p}(x)$, by construction.

In principle, a biased jet calibration is usable, even beneficial if the resolution can be made small. However, for a variety of reasons, it is desirable for the calibration procedure to *close*:

¹Capital letters represent random variables and lower case letters represent realizations of those random variables, i.e. $X = x$ means the random variable X takes on the (non-random) value x .

$$\mathbb{E} \left[\frac{Y}{\tilde{R}(Y)|_X} \middle| X = x \right] = 1. \quad (\text{C.16})$$

The random variable $Y/\tilde{R}(Y) = f^{-1}(Y)$. To see this, let $\tilde{x} = f^{-1}(Y)$, Then,

$$\frac{Y}{\tilde{R}(Y)} = \frac{f(\tilde{x})}{R(\tilde{x})} = \tilde{x} = f^{-1}(Y). \quad (\text{C.17})$$

Now, suppose that $Y|X = x \sim \mathcal{N}(f(x), \sigma(x))$. One can calculate the non-closure for a given function f . First, a lemma:

Lemma. Suppose that $X \sim \mathcal{N}(\mu, \sigma)$. Then, $f(X) \sim \mathcal{N}(\mu', \sigma')$ if and only if $f(x)$ is linear in x .

The proof is in Sec. C.8. Now a corollary for numerical inversion:

Corollary. Suppose that $Y|X = x \sim \mathcal{N}(f(x), \sigma(x))$. Then, the calibrated jet p_T response $Y/\tilde{R}(Y)|X = x$ is normally distributed if and only if f is linear in x .

This corollary is surprising because the ATLAS response function is non-linear and therefore numerical inversion *spoils* normality. However, for Run 1 conditions with moderate pileup, this is a small effect. Figure C.4 shows the *theoretical* non-closure as a function of x . The most relevant curve is the red one, which closely models the ATLAS response function. The method clearly does not close, but the amount of non-closure is already less than 0.5% at 20 GeV (and decreases with x).




Figure C.4: The calculated non-closure for numerical inversion assuming $Y|X = x \sim \mathcal{N}(f(x), \sigma(x))$, where $\sigma(x)/x = 1/\sqrt{x}$ estimated from Ref. [289] and $f(x)$ has several possibilities. The blue line is $f(x) = ax + b$ (which closes identically), red is $f(x) = a + b \log(x)$ (derived by estimating the curve from the lowest $|\eta|$ bin in Ref. [289] - $a = 0.5$ and $b = 0.09$), purple is $f(x) = \sqrt{x}$, and green is $f(x) = ax^2 + bx + c$ with $a = .73, b = 0.002, c = 0.0$.

C.8 Gaussian Invariance Lemma

Lemma. Suppose that $X \sim \mathcal{N}(\mu, \sigma)$. Then, $f(X) \sim \mathcal{N}(\mu', \sigma')$ if and only if $f(x)$ is linear in x .

Proof. The converse is trivial. For the other direction, suppose that $f(X) \sim \mathcal{N}(\mu', \sigma')$. Let $Y = (X - \mu)/\sigma$ and define

$$g(y) = \frac{f(\sigma y + \mu) - \mu'}{\sigma'}. \quad (\text{C.18})$$

so that Y and $g(Y)$ both have a standard normal distribution. First, note that for $Z = g(Y)$, the following relation holds amongst the probability distributions for Z and Y :

$$f_Z(z) = f_Y(g^{-1}(z)) \frac{\partial g^{-1}(z)}{\partial z}. \quad (\text{C.19})$$

In particular, since the normal probability distribution is never non-positive, g has to be monotonic (the derivative term can never be zero). Then, we can write for any c :

$$\begin{aligned} \Phi(c) &= \Pr(Z < c) = \Pr(g(Y) < c) \\ &= \Pr(Y < g^{-1}(c)) = \Phi(g^{-1}(c)), \end{aligned} \quad (\text{C.20})$$

where the second line holds because g preserves ordering. Since the normal distribution cumulative distribution function is invertible, we then have that $g(c) = c$. Inserting the definition of g then gives us the final result:

$$f(x) = \frac{\sigma'}{\sigma}(x - \mu) + \mu' \quad (\text{C.21})$$

□

Bibliography

- [1] B. Efron, *Bootstrap Methods: Another Look at the Jackknife*, *Ann. Statist.* **7** (1979) 1–26.
- [2] J. L. Horowitz, *Chapter 52 - The Bootstrap*, in *Handbook of Econometrics* (J. J. Heckman and E. Leamer, eds.), vol. 5, pp. 3159 – 3228. Elsevier, 2001.
- [3] T. Gleisberg *et. al.*, *Event generation with SHERPA 1.1*, *JHEP* **02** (2009) 007 [[arXiv:0811.4622](#)].
- [4] A. Pais, *Inward Bound: Of Matter and Forces in the Physical World*. Oxford University Press, USA, 1988.
- [5] S. S. Schweber, *QED and the men who made it: Dyson, Feynman, Schwinger, and Tomonaga*. 1994.
- [6] R. Assmann, M. Lamont and S. Myers, *A brief history of the LEP collider*, *Nucl. Phys. Proc. Suppl.* **109B** (2002) 17–31.
- [7] C. Campagnari and M. Franklin, *The Discovery of the top quark*, *Rev. Mod. Phys.* **69** (1997) 137–212 [[arXiv:hep-ex/9608003](#)].
- [8] J. Ellis, M. K. Gaillard and D. V. Nanopoulos, *An Updated Historical Profile of the Higgs Boson*, [arXiv:1504.07217](#).
- [9] G. Zweig, *An SU_3 model for strong interaction symmetry and its breaking; Version 2*, *CERN-TH-412* (1973) [[cds.cern.ch:570209](#)].

- [10] M. Gell-Mann, *A Schematic Model of Baryons and Mesons*, *Phys. Lett.* **8** (1964) 214–215.
- [11] H. Fritzsch, M. Gell-Mann and H. Leutwyler, *Advantages of the Color Octet Gluon Picture*, *Phys. Lett.* **B47** (1973) 365–368.
- [12] F. J. Dyson, *The Radiation theories of Tomonaga, Schwinger, and Feynman*, *Phys. Rev.* **75** (1949) 486–502.
- [13] S. L. Glashow, *Partial Symmetries of Weak Interactions*, *Nucl. Phys.* **22** (1961) 579–588.
- [14] S. Weinberg, *A Model of Leptons*, *Phys. Rev. Lett.* **19** (1967) 1264–1266.
- [15] A. Salam, *Weak and Electromagnetic Interactions*, *Conf. Proc.* **C680519** (1968) 367–377.
- [16] F. Englert and R. Brout, *Broken Symmetry and the Mass of Gauge Vector Mesons*, *Phys. Rev. Lett.* **13** (1964) 321–323.
- [17] P. W. Higgs, *Broken Symmetries and the Masses of Gauge Bosons*, *Phys. Rev. Lett.* **13** (1964) 508–509.
- [18] P. W. Higgs, *Broken symmetries, massless particles and gauge fields*, *Phys. Lett.* **12** (1964) 132–133.
- [19] G. S. Guralnik, C. R. Hagen and T. W. B. Kibble, *Global Conservation Laws and Massless Particles*, *Phys. Rev. Lett.* **13** (1964) 585–587.
- [20] M. E. Peskin and D. V. Schroeder, *An Introduction to quantum field theory*. Addison-Wesley, Reading USA, 1995.
- [21] R. Ticciati, *Quantum field theory for mathematicians*. Cambridge University Press, Cambridge UK, 1999.
- [22] R. K. Ellis, W. J. Stirling and B. R. Webber, *QCD and collider physics*. Cambridge University Press, Cambridge UK, 1996.

- [23] **ATLAS** Collaboration, *Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC*, *Phys. Lett. B* **716** (2012) 1–29 [[arXiv:1207.7214](#)].
- [24] **CMS** Collaboration, *Observation of a new boson at a mass of 125 GeV with the CMS experiment at the LHC*, *Phys. Lett. B* **716** (2012) 30–61 [[arXiv:1207.7235](#)].
- [25] N. Cabibbo, *Unitary symmetry and leptonic decays*, *Phys. Rev. Lett.* **10** (1963) 531–532.
- [26] M. Kobayashi and T. Maskawa, *CP Violation in the Renormalizable Theory of Weak Interaction*, *Prog. Theor. Phys.* **49** (1973) 652–657.
- [27] **Particle Data Group**, *The Review of Particle Physics*, *Phys. Rev. D* **86** (2012) 010001.
- [28] D. J. Gross and F. Wilczek, *Ultraviolet Behavior of Nonabelian Gauge Theories*, *Phys. Rev. Lett.* **30** (1973) 1343–1346.
- [29] H. D. Politzer, *Reliable Perturbative Results for Strong Interactions?*, *Phys. Rev. Lett.* **30** (1973) 1346–1349.
- [30] C. G. Callan, Jr., *Broken scale invariance in scalar field theory*, *Phys. Rev. D* **2** (1970) 1541–1547.
- [31] K. Symanzik, *Small distance behavior in field theory and power counting*, *Commun. Math. Phys.* **18** (1970) 227–246.
- [32] R. D. Ball *et. al.*, *Parton distributions with LHC data*, *Nucl. Phys. B* **867** (2013) 244–289 [[arXiv:1207.1303](#)].
- [33] <http://hepdata.cedar.ac.uk/>.
- [34] Y. L. Dokshitzer, *Calculation of the Structure Functions for Deep Inelastic Scattering and e^+e^- Annihilation by Perturbation Theory in Quantum Chromodynamics.*, *Sov. Phys. JETP* **46** (1977) 641–653.

- [35] V. N. Gribov and L. N. Lipatov, *Deep inelastic ep scattering in perturbation theory*, *Sov. J. Nucl. Phys.* **15** (1972) 438–450.
- [36] G. Altarelli and G. Parisi, *Asymptotic Freedom in Parton Language*, *Nucl. Phys.* **B126** (1977) 298–318.
- [37] A. Hoecker, *The Hadronic Contribution to the Muon Anomalous Magnetic Moment and to the Running Electromagnetic Fine Structure Constant at M_Z - Overview and Latest Results*, *Nucl. Phys. Proc. Suppl.* **218** (2011) 189–200 [[arXiv:1012.0055](https://arxiv.org/abs/1012.0055)].
- [38] M. Baak *et. al.*, *The Electroweak Fit of the Standard Model after the Discovery of a New Boson at the LHC*, *Eur. Phys. J.* **C72** (2012) 2205 [[arXiv:1209.2716](https://arxiv.org/abs/1209.2716)].
- [39] **ATLAS** Collaboration, “Summary plots from the ATLAS Standard Model physics group.” <https://atlas.web.cern.ch/Atlas/GROUPS/PHYSICS/CombinedSummaryPlots/SM/>.
- [40] J. H. Schwarz and N. Seiberg, *String theory, supersymmetry, unification, and all that*, *Rev. Mod. Phys.* **71** (1999) S112–S120 [[arXiv:hep-th/9803179](https://arxiv.org/abs/hep-th/9803179)].
- [41] T. S. van Albada *et. al.*, *Distribution of dark matter in the spiral galaxy NGC 3198*, *Astrophys. J.* **295** (1985) 305–313.
- [42] D. Clowe, A. Gonzalez and M. Markevitch, *Weak lensing mass reconstruction of the interacting cluster 1E0657-558: Direct evidence for the existence of dark matter*, *Astrophys. J.* **604** (2004) 596–603 [[arXiv:astro-ph/0312273](https://arxiv.org/abs/astro-ph/0312273)].
- [43] M. Markevitch *et. al.*, *Direct constraints on the dark matter self-interaction cross-section from the merging galaxy cluster 1E0657-56*, *Astrophys. J.* **606** (2004) 819–824 [[arXiv:astro-ph/0309303](https://arxiv.org/abs/astro-ph/0309303)].
- [44] **Particle Data Group**, *Dark Matter and Cosmological Parameters*, *Chin. Phys.* **C38** (2014) 090001.

- [45] **Super-Kamiokande** Collaboration, Y. Fukuda *et. al.*, *Evidence for oscillation of atmospheric neutrinos*, *Phys. Rev. Lett.* **81** (1998) 1562–1567 [[arXiv:hep-ex/9807003](#)].
- [46] R. A. Battye and A. Moss, *Evidence for massive neutrinos from cosmic microwave background and lensing observations*, *Phys. Rev. Lett.* **112** (2014) 051303.
- [47] S. F. King, *Neutrino mass models*, *Rept. Prog. Phys.* **67** (2004) 107–158 [[arXiv:hep-ph/0310204](#)].
- [48] M. Dine, *TASI lectures on the strong CP problem*, *Flavor physics for the millennium* (2000) 349–369 [[arXiv:hep-ph/0011376](#)].
- [49] C. A. Baker *et. al.*, *An Improved experimental limit on the electric dipole moment of the neutron*, *Phys. Rev. Lett.* **97** (2006) 131801 [[arXiv:hep-ex/0602020](#)].
- [50] V. Baluni, *CP Violating Effects in QCD*, *Phys. Rev.* **D19** (1979) 2227–2230.
- [51] R. Crewther *et. al.*, *Chiral estimate of the electric dipole moment of the neutron in quantum chromodynamics*, *Phys. Lett.* **B88** (1979) 123–127.
- [52] R. D. Peccei and H. R. Quinn, *CP Conservation in the Presence of Instantons*, *Phys. Rev. Lett.* **38** (1977) 1440–1443.
- [53] R. D. Peccei and H. R. Quinn, *Constraints Imposed by CP Conservation in the Presence of Instantons*, *Phys. Rev.* **D16** (1977) 1791–1797.
- [54] S. Weinberg, *A New Light Boson?*, *Phys. Rev. Lett.* **40** (1978) 223–226.
- [55] F. Wilczek, *Problem of Strong p and t Invariance in the Presence of Instantons*, *Phys. Rev. Lett.* **40** (1978) 279–282.
- [56] R. Gupta, *Introduction to lattice QCD: Course, Probing the standard model of particle interactions* (1997) 83–219 [[arXiv:hep-lat/9807028](#)].

- [57] G. Plass, *The CERN Proton Synchrotron: 50 years of reliable operation and continued development*, *Eur. Phys. J.* **H36** (2012) 439–454.
- [58] **Gargamelle Neutrino** Collaboration, *Observation of Neutrino Like Interactions Without Muon Or Electron in the Gargamelle Neutrino Experiment*, *Phys. Lett.* **B46** (1973) 138–140.
- [59] **UA1** Collaboration, *Experimental Observation of Isolated Large Transverse Energy Electrons with Associated Missing Energy at $\sqrt{s} = 540$ GeV*, *Phys. Lett.* **B122** (1983) 103–116.
- [60] **UA1** Collaboration, *Experimental Observation of Lepton Pairs of Invariant Mass Around 95 GeV/c² at the CERN SPS Collider*, *Phys. Lett.* **B126** (1983) 398–410.
- [61] **UA2** Collaboration, *Observation of Single Isolated Electrons of High Transverse Momentum in Events with Missing Transverse Energy at the CERN $\bar{p}p$ Collider*, *Phys. Lett.* **B122** (1983) 476–485.
- [62] **UA2** Collaboration, *Evidence for $Z^0 \rightarrow e^+e^-$ at the CERN $\bar{p}p$ Collider*, *Phys. Lett.* **B129** (1983) 130–140.
- [63] M. A. Heald and J. B. Marion, *Classical Electromagnetic Radiation*. Brooks Cole, 1994.
- [64] J. D. Jackson, *Classical Electrodynamics Third Edition*. Wiley, 1998.
- [65] E. Keil, *Synchrotron radiation dominated hadron colliders*, *LHC-Project-Report-102* (1997) [[cds.cern.ch:327302](https://cds.cern.ch/327302)].
- [66] P. Lebrun, *Cryogenics for the Large Hadron Collider; 1999 ed.*, *LHC-Project-Report-338* (1999) [[cds.cern.ch:411139](https://cds.cern.ch/411139)].
- [67] H. Wiedemann, *Particle Accelerator Physics*. Springer, 2007.
- [68] M. Lamont, *Status of the LHC*, *Journal of Physics: Conference Series* **455** (2013) 012001.

- [69] S. van der Meer, *Calibration of the effective beam height in the ISR*, *ISR-PO-68-31* (1968) [[cds.cern.ch:296752](https://cds.cern.ch/296752)].
- [70] **ATLAS** Collaboration, *Improved luminosity determination in pp collisions at $\sqrt{s} = 7$ TeV using the ATLAS detector at the LHC*, *Eur. Phys. J.* **C73** (2013) 2518 [[arXiv:1302.4393](https://arxiv.org/abs/1302.4393)].
- [71] C. Grupen and I. Buvat, eds., *Handbook of particle detection and imaging, vol. 1 and vol. 2*. Springer, Berlin Germany, 2012.
- [72] S. Tavernier, *Experimental Techniques in Nuclear and Particle Physics*. Springer, Berlin Germany, 2010.
- [73] R. L. Gluckstern, *Uncertainties in track momentum and direction, due to multiple scattering and measurement errors*, *Nucl. Instrum. Meth.* **24** (1963) 381–389.
- [74] R. Wigmans, *Calorimetry: Energy measurement in particle physics*, *Int. Ser. Monogr. Phys.* **107** (2000) 1–726.
- [75] **ATLAS** Collaboration, *Performance of b -Jet Identification in the ATLAS Experiment*, *JINST* **11** (2016) P04008 [[arXiv:1512.01094](https://arxiv.org/abs/1512.01094)].
- [76] **ATLAS** Collaboration, *dE/dx measurement in the ATLAS Pixel Detector and its use for particle identification*, *ATLAS-CONF-2011-016* (2011) [[cds.cern.ch:1336519](https://cds.cern.ch/1336519)].
- [77] **ATLAS** Collaboration, *The ATLAS Experiment at the CERN Large Hadron Collider*, *JINST* **3** (2008) S08003.
- [78] **ATLAS** Collaboration, *Studies of the performance of the ATLAS detector using cosmic-ray muons*, *Eur. Phys. J.* **C71** (2011) 1593 [[arXiv:1011.6665](https://arxiv.org/abs/1011.6665)].
- [79] S. Parker, C. Kenney and J. Segal, *3D - A proposed new architecture for solid-state radiation detectors*, *Nucl. Instr. and Meth.* **A395** (1997) 328 – 343.

- [80] **ATLAS** Collaboration, *Track Reconstruction Performance of the ATLAS Inner Detector at $\sqrt{s} = 13$ TeV*, *ATL-PHYS-PUB-2015-018* (2015) [[cds.cern.ch:2037683](https://cds.cern.ch/record/2037683)].
- [81] **ATLAS** Collaboration, *ATLAS pixel detector electronics and sensors*, *JINST* **3** (2008) P07007.
- [82] **ATLAS** Collaboration, *Topological cell clustering in the ATLAS calorimeters and its performance in LHC Run 1*, [arXiv:1603.02934](https://arxiv.org/abs/1603.02934).
- [83] **ATLAS** Collaboration, *Performance of the ATLAS Trigger System in 2010*, *Eur. Phys. J.* **C72** (2012) 1849 [[arXiv:1110.1530](https://arxiv.org/abs/1110.1530)].
- [84] **ATLAS** Collaboration, *2015 start-up trigger menu and initial performance assessment of the ATLAS trigger using Run-2 data*, *ATL-DAQ-PUB-2016-001* (2016) [[cds.cern.ch:2136007](https://cds.cern.ch/record/2136007)].
- [85] **ATLAS** Collaboration, *Measurement of the inclusive $W^\pm$ and Z/γ cross sections in the electron and muon decay channels in pp collisions at $\sqrt{s} = 7$ TeV with the ATLAS detector*, *Phys. Rev.* **D85** (2012) 072004 [[arXiv:1109.5141](https://arxiv.org/abs/1109.5141)].
- [86] **ATLAS** Collaboration, *Measurement of the total cross section from elastic scattering in pp collisions at $\sqrt{s} = 7$ TeV with the ATLAS detector*, *Nucl. Phys.* **B889** (2014) 486–548 [[arXiv:1408.5778](https://arxiv.org/abs/1408.5778)].
- [87] J. Alwall *et. al.*, *Comparative study of various algorithms for the merging of parton showers and matrix elements in hadronic collisions*, *Eur. Phys. J.* **C53** (2008) 473–500 [[arXiv:0706.2569](https://arxiv.org/abs/0706.2569)].
- [88] S. Frixione and B. R. Webber, *Matching NLO QCD computations and parton shower simulations*, *JHEP* **06** (2002) 029 [[arXiv:0204244](https://arxiv.org/abs/0204244)].
- [89] P. Nason, *A New method for combining NLO QCD with shower Monte Carlo algorithms*, *JHEP* **0411** (2004) 040 [[arXiv:0409146](https://arxiv.org/abs/0409146)].

- [90] P. Nason and B. Webber, *Next-to-Leading-Order Event Generators*, *Ann. Rev. Nucl. Part. Sci.* **62** (2012) 187–213 [[arXiv:1202.1251](#)].
- [91] L. Lönnblad and S. Prestel, *Merging Multi-leg NLO Matrix Elements with Parton Showers*, *JHEP* **03** (2013) 166 [[arXiv:1211.7278](#)].
- [92] R. Frederix and S. Frixione, *Merging meets matching in MC@NLO*, *JHEP* **12** (2012) 061 [[arXiv:1209.6215](#)].
- [93] T. Gehrmann, S. Hoeche, F. Krauss, M. Schonherr and F. Siegert, *NLO QCD matrix elements + parton showers in $e^+e^- \rightarrow \text{hadrons}$* , *JHEP* **01** (2013) 144 [[arXiv:1207.5031](#)].
- [94] S. Hoeche, F. Krauss, M. Schonherr and F. Siegert, *QCD matrix elements + parton showers: The NLO case*, *JHEP* **04** (2013) 027 [[arXiv:1207.5030](#)].
- [95] T. Gleisberg, S. Hoeche, F. Krauss, A. Schaliche, S. Schumann and J.-C. Winter, *SHERPA 1. alpha: A Proof of concept version*, *JHEP* **02** (2004) 056 [[arXiv:0311263](#)].
- [96] T. Sjöstrand, S. Mrenna and P. Z. Skands, *PYTHIA 6.4 Physics and Manual*, *JHEP* **0605** (2006) 026 [[arXiv:hep-ph/0603175](#)].
- [97] G. Corcella *et. al.*, *HERWIG 6.5 release note*, [arXiv:hep-ph/0210213](#).
- [98] T. Sjöstrand and P. Z. Skands, *Transverse-momentum-ordered showers and interleaved multiple interactions*, *Eur. Phys. J.* **C39** (2005) 129–154 [[arXiv:hep-ph/0408302](#)].
- [99] S. Gieseke, P. Stephens and B. Webber, *New formalism for QCD parton showers*, *JHEP* **0312** (2003) 045 [[arXiv:hep-ph/0310083](#)].
- [100] P. Golonka and Z. Was, *PHOTOS Monte Carlo: A Precision tool for QED corrections in Z and W decays*, *Eur. Phys. J.* **C45** (2006) 97–107 [[arXiv:hep-ph/0506026](#)].

- [101] B. Andersson *et. al.*, *Parton Fragmentation and String Dynamics*, *Phys. Rep.* **97** (1983) 31.
- [102] B. Webber, *A QCD Model for Jet Fragmentation Including Soft Gluon Interference*, *Nucl. Phys.* **B238** (1984) 492.
- [103] T. Sjöstrand and P. Z. Skands, *Multiple interactions and the structure of beam remnants*, *JHEP* **0403** (2004) 053 [[arXiv:hep-ph/0402078](#)].
- [104] M. Bähr, S. Gieseke and M. H. Seymour, *Simulation of multiple partonic interactions in Herwig++*, *JHEP* **0807** (2008) 076 [[arXiv:0803.3633](#)].
- [105] D. J. Lange, *The EvtGen particle decay simulation package*, *Nucl. Instrum. Meth.* **A462** (2001) 152.
- [106] S. Jadach, Z. Was, R. Decker and J. H. Kuhn, *The tau decay library TAUOLA: Version 2.4*, *Comput. Phys. Commun.* **76** (1993) 361–380.
- [107] Z. Marshall, *Simulation of Pile-up in the ATLAS Experiment*, *J. Phys. Conf. Ser.* **513** (2014) 022024.
- [108] **GEANT4** Collaboration, *GEANT4 Collaboration, GEANT4: A Simulation toolkit*, *Nucl. Instrum. Meth.* **A506** (2003) 250–303.
- [109] **ATLAS** Collaboration, *The simulation principle and performance of the ATLAS fast calorimeter simulation FastCaloSim*, *ATL-PHYS-PUB-2010-013* (2010) [[cds.cern.ch:1300517](#)].
- [110] **ATLAS** Collaboration, *The ATLAS Simulation Infrastructure*, *Eur. Phys. J.* **C70** (2010) 823–874 [[arXiv:1005.4568](#)].
- [111] T. Cornelissen *et. al.*, *Concepts, Design and Implementation of the ATLAS New Tracking (NEWT)*, *ATL-SOFT-PUB-2007-007* (2007) [[cds.cern.ch:1020106](#)].
- [112] **ATLAS** Collaboration, *A neural network clustering algorithm for the ATLAS silicon pixel detector*, *JINST* **9** (2014) P09009 [[arXiv:1406.7690](#)].

- [113] **ATLAS** Collaboration, *ATLAS event at 900 GeV - 6 May 2015 - Run 264034 Event 11475271, ATLAS-PHO-Event-2015-007* (2015) [[cds.cern.ch:2014666](https://cds.cern.ch/record/2014666)].
- [114] S. D. Ellis and D. E. Soper, *Successive combination jet algorithm for hadron collisions*, *Phys. Rev.* **D48** (1993) 3160–3166 [[arXiv:hep-ph/9305266](https://arxiv.org/abs/hep-ph/9305266)].
- [115] Y. L. Dokshitzer, G. Leder, S. Moretti and B. Webber, *Better jet clustering algorithms*, *JHEP* **9708** (1997) 001 [[arXiv:hep-ph/9707323](https://arxiv.org/abs/hep-ph/9707323)].
- [116] M. Wobisch and T. Wengler, *Hadronization corrections to jet cross-sections in deep inelastic scattering*, [arXiv:hep-ph/9907280](https://arxiv.org/abs/hep-ph/9907280).
- [117] M. Cacciari, G. P. Salam and G. Soyez, *The Anti- $k(t)$ jet clustering algorithm*, *JHEP* **0804** (2008) 063 [[arXiv:0802.1189](https://arxiv.org/abs/0802.1189)].
- [118] M. Dasgupta, L. Magnea and G. P. Salam, *Non-perturbative QCD effects in jets at hadron colliders*, *JHEP* **02** (2008) 055 [[arXiv:0712.3014](https://arxiv.org/abs/0712.3014)].
- [119] G. P. Salam, *Towards Jetography*, *Eur. Phys. J.* **C67** (2010) 637–686 [[arXiv:0906.1833](https://arxiv.org/abs/0906.1833)].
- [120] M. Cacciari, G. P. Salam and G. Soyez, *The Catchment Area of Jets*, *JHEP* **04** (2008) 005 [[arXiv:0802.1188](https://arxiv.org/abs/0802.1188)].
- [121] **ATLAS** Collaboration, *Jet energy measurement and its systematic uncertainty in proton-proton collisions at $\sqrt{s} = 7$ TeV with the ATLAS detector*, *Eur. Phys. J.* **C75** (2015) 17 [[arXiv:1406.0076](https://arxiv.org/abs/1406.0076)].
- [122] **ATLAS** Collaboration, *Monte Carlo Calibration and Combination of In-situ Measurements of Jet Energy Scale, Jet Energy Resolution and Jet Mass in ATLAS*, *ATLAS-CONF-2015-037* (2015) [[cds.cern.ch:2044941](https://cds.cern.ch/record/2044941)].
- [123] M. Cacciari and G. P. Salam, *Pileup subtraction using jet areas*, *Phys. Lett.* **B659** (2008) 119–126 [[arXiv:0707.1378](https://arxiv.org/abs/0707.1378)].

- [124] **ATLAS** Collaboration, *Performance of pile-up mitigation techniques for jets in pp collisions at $\sqrt{s} = 8$ TeV using the ATLAS detector*, [arXiv:1510.03823](#).
- [125] **ATLAS** Collaboration, *Jet global sequential corrections with the ATLAS detector in proton-proton collisions at $\sqrt{s} = 8$ TeV*, ATLAS-CONF-2015-002 (2015) [[cds.cern.ch:2001682](#)].
- [126] **ATLAS** Collaboration, *Data-driven determination of the energy scale and resolution of jets reconstructed in the ATLAS calorimeters using dijet and multijet events at $\sqrt{s} = 8$ TeV*, ATLAS-CONF-2015-017 (2015) [[cds.cern.ch:2008678](#)].
- [127] **ATLAS** Collaboration, *Determination of the jet energy scale and resolution at ATLAS using Z/ γ -jet events in data at $\sqrt{s} = 8$ TeV*, ATLAS-CONF-2015-057 (2015) [[cds.cern.ch:2059846](#)].
- [128] **ATLAS** Collaboration, *Electron and photon energy calibration with the ATLAS detector using LHC Run 1 data*, *Eur. Phys. J.* **C74** (2014) 3071 [[arXiv:1407.5063](#)].
- [129] **ATLAS** Collaboration, *Electron reconstruction and identification efficiency measurements with the ATLAS detector using the 2011 LHC proton-proton collision data*, *Eur. Phys. J.* **C74** (2014) 2941 [[arXiv:1404.2240](#)].
- [130] **ATLAS** Collaboration, *Electron performance measurements with the ATLAS detector using the 2010 LHC proton-proton collision data*, *Eur. Phys. J.* **C72** (2012) 1909 [[arXiv:1110.3174](#)].
- [131] **ATLAS** Collaboration, *Measurement of the muon reconstruction performance of the ATLAS detector using 2011 and 2012 LHC proton-proton collision data*, *Eur. Phys. J.* **C74** (2014) 3130 [[arXiv:1407.3935](#)].
- [132] **ATLAS** Collaboration, *Muon reconstruction efficiency and momentum resolution of the ATLAS experiment in proton-proton collisions at $\sqrt{s} = 7$ TeV in 2010*, *Eur. Phys. J.* **C74** (2014) 3034 [[arXiv:1404.4562](#)].

- [133] **ATLAS** Collaboration, *Muon reconstruction performance of the ATLAS detector in proton–proton collision data at $\sqrt{s} = 13$ TeV*, *Eur. Phys. J.* **C76** (2016) 292 [[arXiv:1603.05598](#)].
- [134] **ATLAS** Collaboration, *Identification and energy calibration of hadronically decaying tau leptons with the ATLAS experiment in pp collisions at $\sqrt{s} = 8$ TeV*, *Eur. Phys. J.* **C75** (2015) 303 [[arXiv:1412.7086](#)].
- [135] **ATLAS** Collaboration, *Reconstruction of hadronic decay products of tau leptons with the ATLAS experiment*, *Eur. Phys. J.* **C76** (2015) 295 [[arXiv:1512.05955](#)].
- [136] **ATLAS** Collaboration, *Performance of Missing Transverse Momentum Reconstruction in ATLAS studied in Proton-Proton Collisions recorded in 2012 at 8 TeV*, *ATLAS-CONF-2013-082* (2013) [[cds.cern.ch:1570993](#)].
- [137] **ATLAS** Collaboration, *Performance of Missing Transverse Momentum Reconstruction in Proton-Proton Collisions at 7 TeV with ATLAS*, *Eur. Phys. J.* **C72** (2012) 1844 [[arXiv:1108.5602](#)].
- [138] **ATLAS** Collaboration, *Expected performance of missing transverse momentum reconstruction for the ATLAS detector at $\sqrt{s} = 13$ TeV*, *ATL-PHYS-PUB-2015-023* (2015) [[cds.cern.ch:2037700](#)].
- [139] **ATLAS** Collaboration, *Performance of missing transverse momentum reconstruction for the ATLAS detector in the first proton-proton collisions at $\sqrt{s} = 13$ TeV*, *ATL-PHYS-PUB-2015-027* (2015) [[cds.cern.ch:2037904](#)].
- [140] J. Gallicchio and M. D. Schwartz, *Seeing in Color: Jet Superstructure*, *Phys. Rev. Lett.* **105** (2010) 022001 [[arXiv:1001.5027](#)].
- [141] R. D. Field and R. P. Feynman, *A Parameterization of the properties of Quark Jets*, *Nucl. Phys.* **B136** (1978) 1.
- [142] J. Berge *et. al.*, *Net charge in deep inelastic antineutrino-nucleon scattering*, *Phys. Lett.* **B91** (1980) 311.

- [143] J. Berge *et. al.*, *Quark jets from antineutrino interactions: Net charge and factorization in the quark jets*, *Nucl. Phys.* **B184** (1981) 13.
- [144] P. Allen *et. al.*, *Multiplicity distributions in neutrino-hydrogen interactions*, *Nucl. Phys.* **B181** (1981) 385.
- [145] P. Allen *et. al.*, *Charge properties of the hadronic system in νp and $\bar{\nu} p$ interactions*, *Phys. Lett.* **B112** (1982) 88.
- [146] J. Albanese *et. al.*, *Quark charge retention in final state hadrons from deep inelastic muon scattering*, *Phys. Lett.* **B144** (1984) 302.
- [147] S. Barlag *et. al.*, *Charged hadron multiplicities in high energy $\bar{\nu}_\mu n$ and $\bar{\nu}_\mu p$ interactions*, *Z. Phys.* **C11** (1982) 283.
- [148] R. Erickson *et. al.*, *Charge retention in deep-inelastic electroproduction*, *Phys. Rev. Lett.* **42** (1979) 822.
- [149] **SLD** Collaboration, *Measurement of the Parity-Violation Parameter A_b from the Left-Right Forward-Backward Asymmetry of b Quark Production in Z^0 Decays Using a Momentum-Weighted Track-Charge Technique*, *Phys. Rev. Lett.* **74** (1995) 2890.
- [150] **TASSO** Collaboration, *Measurement of the asymmetry of b quark production in e^+e^- annihilation of $\sqrt{s} = 35$ GeV*, *Z. Phys.* **C48** (1990) 433.
- [151] **DELPHI** Collaboration, *A Measurement of $\sin^2\theta(W)$ from the charge asymmetry of hadronic events at the Z^0 peak*, *Phys. Lett.* **B277** (1992) 371.
- [152] **ALEPH** Collaboration, *Measurement of charge asymmetry in hadronic Z decays*, *Phys. Lett.* **B259** (1991) 377.
- [153] **OPAL** Collaboration, *A measurement of the forward-backward charge asymmetry in hadronic decays of the Z^0* , *Phys. Lett.* **B294** (1992) 436.
- [154] **OPAL** Collaboration, *Measurement of the time dependence of $B_d^0 \leftrightarrow \bar{B}_d^0$ mixing using a jet charge technique*, *Phys. Lett. B* **327** (1994) 411.

- [155] **DELPHI** Collaboration, *Measurement of the B_d^0 oscillation frequency using kaons, leptons and jet charge*, *Z. Phys.* **C72** (1996) 17.
- [156] **CDF** Collaboration, *Measurement of $B^0 - \bar{B}^0$ flavor oscillations using jet-charge and lepton flavor tagging in $p\bar{p}$ collisions at $\sqrt{s} = 1.8$ TeV*, *Phys. Rev.* **D60** (1999) 072003.
- [157] **D0** Collaboration, *Experimental discrimination between charge $2e/3$ top quark and charge $4e/3$ exotic quark production scenarios*, *Phys. Rev. Lett.* **98** (2007) 041801 [[arXiv:hep-ex/0608044](#)].
- [158] **CDF** Collaboration, *The cdf measurement of the top quark charge using the top decay products in lepton+jet channel*, *CDF note 10460* (2011).
- [159] **ATLAS** Collaboration, *Measurement of the top quark charge in $p\bar{p}$ collisions at $\sqrt{s} = 7$ TeV with the ATLAS detector*, *JHEP* **1311** (2013) 031 [[arXiv:1307.4568](#)].
- [160] **ALEPH** Collaboration, *Measurement of triple gauge-boson couplings at 172 GeV*, *Phys. Lett.* **B422** (1998) 369.
- [161] **DELPHI** Collaboration, *Measurement of trilinear gauge boson couplings WWV , ($V = Z, \gamma$) in $e^+ e^-$ collisions at 189 GeV*, *Phys. Lett.* **B502** (2001) 9–23 [[arXiv:hep-ex/0102041](#)].
- [162] **L3** Collaboration, *Measurement of triple gauge boson couplings of the W boson at LEP*, *Phys. Lett.* **B467** (1999) 171–184 [[arXiv:hep-ex/9910008](#)].
- [163] **OPAL** Collaboration, *Measurement of triple gauge boson couplings from W^+W^- production at LEP energies up to 189 GeV*, *Eur. Phys. J.* **C19** (2001) 1–14 [[arXiv:hep-ex/0009022](#)].
- [164] **ATLAS** Collaboration, *Jet Charge Studies with the ATLAS Detector Using $\sqrt{s} = 8$ TeV Proton-Proton Collision Data*, *ATLAS-CONF-2013-086* (2013) [[cds.cern.ch:1572980](#)].

- [165] **CMS** Collaboration, *Identification techniques for highly boosted W bosons that decay into hadrons*, *JHEP* **1412** (2014) 017 [[arXiv:1410.4227](#)].
- [166] **CMS** Collaboration, *Identification techniques for highly boosted W bosons that decay into hadrons*, *JHEP* **12** (2014) 017 [[arXiv:1410.4227](#)].
- [167] **ATLAS** Collaboration, *Light-quark and gluon jet discrimination in pp collisions at $\sqrt{s} = 7$ TeV with the ATLAS detector*, *Eur. Phys. J.* **C74** (2014) 3023 [[arXiv:1405.6583](#)].
- [168] **OPAL** Collaboration, *A Study of the electric charge distributions of quark and gluon jets in hadronic Z^0 decays*, *Phys. Lett.* **B302** (1993) 523–532.
- [169] **UA1** Collaboration, *Analysis of the Fragmentation Properties of Quark and Gluon Jets at the CERN SPS $p\bar{p}$ Collider*, *Nucl. Phys.* **B276** (1986) 253–271.
- [170] **ATLAS** Collaboration, *Measurement of jet charge in dijet events from $\sqrt{s} = 8$ TeV pp collisions with the ATLAS detector*, *Phys. Rev.* **D93** (2016) 052003 [[arXiv:1509.05190](#)].
- [171] D. Krohn, M. D. Schwartz, T. Lin and W. J. Waalewijn, *Jet Charge at the LHC*, *Phys. Rev. Lett.* **110** (2013) 212001 [[arXiv:1209.2421](#)].
- [172] W. J. Waalewijn, *Calculating the charge of a jet*, *Phys. Rev.* **D86** (2012) 094030 [[arXiv :1209.3019](#)].
- [173] C. W. Bauer, D. Pirjol and I. W. Stewart, *Soft collinear factorization in effective field theory*, *Phys. Rev.* **D65** (2002) 054022 [[arXiv:hep-ph/0109045](#)].
- [174] C. W. Bauer and I. W. Stewart, *Invariant operators in collinear effective theory*, *Phys. Lett.* **B516** (2001) 134–142 [[arXiv:hep-ph/0107001](#)].
- [175] C. W. Bauer, S. Fleming, D. Pirjol and I. W. Stewart, *An Effective field theory for collinear and soft gluons: Heavy to light decays*, *Phys. Rev.* **D63** (2001) 114020 [[arXiv:0011336](#)].

- [176] C. W. Bauer, S. Fleming and M. E. Luke, *Summing Sudakov logarithms in $B \rightarrow X_s \gamma$ in effective field theory*, *Phys. Rev.* **D63** (2000) 014006 [[arXiv:hep-ph/0005275](#)].
- [177] M. Procura and W. J. Waalewijn, *Fragmentation in Jets: Cone and Threshold Effects*, *Phys. Rev.* **D85** (2012) 114041 [[arXiv:1111.6605](#)].
- [178] A. Jain, M. Procura and W. J. Waalewijn, *Parton Fragmentation within an Identified Jet at NNLL*, *JHEP* **05** (2011) 035 [[arXiv:1101.4953](#)].
- [179] M. Procura and I. W. Stewart, *Quark Fragmentation within an Identified Jet*, *Phys. Rev.* **D81** (2010) 074009 [[arXiv:0911.4980](#)].
- [180] S. D. Ellis, C. K. Vermilion, J. R. Walsh, A. Hornig and C. Lee, *Jet Shapes and Jet Algorithms in SCET*, *JHEP* **11** (2010) 101 [[arXiv:1001.0014](#)].
- [181] **ATLAS** Collaboration, *A new tagger for the charge identification of b-jets*, *ATL-PHYS-PUB-2015-040* (2015) [[cds.cern.ch:2048132](#)].
- [182] G. Watt and R. Thorne, *Study of Monte Carlo approach to experimental uncertainty propagation with MSTW 2008 PDFs*, *JHEP* **1208** (2012) 052 [[arXiv:1205.4024](#)].
- [183] L. Harland-Lang, A. Martin, P. Motylinski and R. Thorne, *Parton distributions in the LHC era: MMHT 2014 PDFs*, *Eur. Phys. J.* **C75** (2015) 204 [[arXiv:1412.3989](#)].
- [184] **H1**, **ZEUS** Collaborations, *Combined Measurement and QCD Analysis of the Inclusive $e^\pm p$ Scattering Cross Sections at HERA*, *JHEP* **1001** (2010) 109 [[arXiv:0911.0884](#)].
- [185] (**H1** and **ZEUS** Collaborations), *Combination of measurements of inclusive deep inelastic $e^\pm p$ scattering cross sections and QCD analysis of HERA data*, *Eur. Phys. J.* **C75** (2015), no. 12 580 [[arXiv:1506.06042](#)].

- [186] **NNPDF** Collaboration, R. D. Ball *et. al.*, *Parton distributions for the LHC Run II*, *JHEP* **1504** (2015) 040 [[arXiv:1410.8849](#)].
- [187] **ATLAS** Collaboration, *Data-Quality Requirements and Event Cleaning for Jets and Missing Transverse Energy Reconstruction with the ATLAS Detector in Proton-Proton Collisions at a Center-of-Mass Energy of $\sqrt{s} = 7$ TeV*, *ATLAS-CONF-2010-038* (2010) [[cds.cern.ch:1277678](#)].
- [188] **ATLAS** Collaboration, *Performance of primary vertex reconstruction in proton-proton collisions at $\sqrt{s}=7$ TeV in the ATLAS experiment*, *ATLAS-CONF-2010-069* (2010) [[cds.cern.ch:1281344](#)].
- [189] T. Sjostrand, S. Mrenna and P. Z. Skands, *A Brief Introduction to PYTHIA 8.1*, *Comput. Phys. Commun.* **178** (2008) 852 [[arXiv:0710.3820](#)].
- [190] **ATLAS** Collaboration, *Summary of ATLAS Pythia 8 tunes*, *ATL-PHYS-PUB-2012-003* (2012) [[cds.cern.ch:1474107](#)].
- [191] J. M. Campbell, J. Huston and W. Stirling, *Hard Interactions of Quarks and Gluons: A Primer for LHC Physics*, *Rept. Prog. Phys.* **70** (2007) 89 [[arXiv:hep-ph/0611148](#)].
- [192] A. Sherstnev and R. Thorne, *Parton Distributions for LO Generators*, *Eur. Phys. J.* **C55** (2008) 553–575 [[arXiv:0711.2473](#)].
- [193] H.-L. Lai *et. al.*, *New parton distributions for collider physics*, *Phys. Rev.* **D82** (2010) 074024 [[arXiv:1007.2241](#)].
- [194] J. Gao *et. al.*, *CT10 next-to-next-to-leading order global analysis of QCD*, *Phys. Rev.* **D89** (2014) 033009 [[arXiv:1302.6246](#)].
- [195] M. Bahr *et. al.*, *Herwig++ Physics and Manual*, *Eur. Phys. J.* **C58** (2008) 639–707 [[arXiv:0803.0883](#)].
- [196] K. Arnold *et. al.*, *Herwig++ 2.6 Release Note*, [arXiv:1205.4902](#).

- [197] S. Gieseke, C. Röhr and A. Siodmok, *Colour reconnections in Herwig++*, *Eur. Phys. J.* **C72** (2012) 2225 [[arXiv:1206.0041](#)].
- [198] J. Pumplin *et. al.*, *New generation of parton distributions with uncertainties from global QCD analysis*, *JHEP* **0207** (2002) 012 [[arXiv:0201195](#)].
- [199] R. D. Ball, L. Del Debbio, S. Forte, A. Guffanti, J. I. Latorre *et. al.*, *A first unbiased global NLO determination of parton distributions and their uncertainties*, *Nucl. Phys.* **B838** (2010) 136–206 [[arXiv:1002.4407](#)].
- [200] S. Frixione, P. Nason and C. Oleari, *Matching NLO QCD computations with Parton Shower simulations: the POWHEG method*, *JHEP* **0711** (2007) 070 [[arXiv:0709.2092](#)].
- [201] S. Alioli *et. al.*, *A general framework for implementing NLO calculations in shower Monte Carlo programs: the POWHEG BOX*, *JHEP* **06** (2010) 043 [[arXiv:1002.2581](#)].
- [202] S. Frixione, P. Nason and G. Ridolfi, *A Positive-weight next-to-leading-order Monte Carlo for heavy flavour hadroproduction*, *JHEP* **0709** (2007) 126 [[arXiv:0707.3088](#)].
- [203] M. Cacciari, G. P. Salam and G. Soyez, *FastJet User Manual*, *Eur. Phys. J.* **C72** (2012) 1896 [[arXiv:1111.6097](#)].
- [204] W. Lampl *et. al.*, *Calorimeter clustering algorithms: description and performance*, *ATL-LARG-PUB-2008-002* (2008) [[cds.cern.ch:1099735](#)].
- [205] C. Cojocaru *et. al.*, *Hadronic calibration of the ATLAS liquid argon end-cap calorimeter in the pseudorapidity region $1.6 < |\eta| < 1.8$ in beam tests*, *Nucl. Instrum. Meth.* **A531** (2004) 481.
- [206] **ATLAS** Collaboration, *Local hadronic calibration*, *ATL-LARG-PUB-2009-001* (2009) [[cds.cern.ch:1112035](#)].

- [207] **ATLAS** Collaboration, *Pile-up subtraction and suppression for jets in ATLAS*, *ATLAS-CONF-2013-083* (2013) [[cds.cern.ch:1570994](https://cds.cern.ch/record/1570994)].
- [208] A. Banfi, G. P. Salam and G. Zanderighi, *Infrared safe definition of jet flavor*, *Eur. Phys. J.* **C47** (2006) 113–124 [[arXiv:hep-ph/0601139](https://arxiv.org/abs/hep-ph/0601139)].
- [209] **ATLAS** Collaboration, *Light-quark and Gluon Jets: Calorimeter Response, Jet Energy Scale Systematics and Properties*, *ATLAS-CONF-2012-138* (2012) [[cds.cern.ch:1480629](https://cds.cern.ch/record/1480629)].
- [210] G. Corcella *et. al.*, *HERWIG 6: An Event generator for hadron emission reactions with interfering gluons (including supersymmetric processes)*, *JHEP* **0101** (2001) 010.
- [211] J. M. Butterworth, J. R. Forshaw and M. H. Seymour, *Multiparton interactions in photoproduction at hera*, *Z. Phys.* **C72** (1996) 637.
- [212] P. Z. Skands, *Tuning Monte Carlo Generators: The Perugia Tunes*, *Phys. Rev.* **D82** (2010) 074018 [[hep-ph/1005.3457](https://arxiv.org/abs/hep-ph/1005.3457)].
- [213] M. Cacciari *et. al.*, *Top-pair production at hadron colliders with next-to-next-to-leading logarithmic soft-gluon resummation*, *Phys. Lett.* **B710** (2012) 612–622 [[arXiv:1111.5869](https://arxiv.org/abs/1111.5869)].
- [214] P. Bärnreuther, M. Czakon and A. Mitov, *Percent Level Precision Physics at the Tevatron: First Genuine NNLO QCD Corrections to $q\bar{q} \rightarrow t\bar{t} + X$* , *Phys. Rev. Lett.* **109** (2012) 132001 [[arXiv:1204.5201](https://arxiv.org/abs/1204.5201)].
- [215] M. Czakon and A. Mitov, *NNLO corrections to top-pair production at hadron colliders: the all-fermionic scattering channels*, *JHEP* **1212** (2012) 054 [[arXiv:1207.0236](https://arxiv.org/abs/1207.0236)].
- [216] M. Czakon and A. Mitov, *NNLO corrections to top pair production at hadron colliders: the quark-gluon reaction*, *JHEP* **1301** (2013) 080 [[arXiv:1210.6832](https://arxiv.org/abs/1210.6832)].

- [217] M. Czakon, P. Fiedler and A. Mitov, *Total Top-Quark Pair-Production Cross Section at Hadron Colliders Through $\mathcal{O}(\alpha_s^4)$* , *Phys. Rev. Lett.* **110** (2013) 252004 [[arXiv:1303.6254](#)].
- [218] M. Czakon and A. Mitov, *Top++: A Program for the Calculation of the Top-Pair Cross-Section at Hadron Colliders*, *Comput. Phys. Commun.* **185** (2014) 2930 [[arXiv:1112.5675](#)].
- [219] M. Botje *et. al.*, *The PDF4LHC Working Group Interim Recommendations*, [arXiv:1101.0538](#).
- [220] A. Martin, W. Stirling, R. Thorne and G. Watt, *Parton distributions for the LHC*, *Eur. Phys. J.* **C63** (2009) 189.
- [221] A. Martin, W. Stirling, R. Thorne and G. Watt, *Uncertainties on α_s in global PDF analyses and implications for predicted hadronic cross sections*, *Eur. Phys. J.* **C64** (2009) 653.
- [222] M. L. Mangano *et. al.*, *ALPGEN, a generator for hard multiparton processes in hadronic collisions*, *JHEP* **0307** (2003) 001.
- [223] B. P. Kersevan and E. Richter-Was, *The Monte Carlo event generator AcerMC version 2.0 with interfaces to PYTHIA 6.2 and HERWIG 6.5*, *Comput. Phys. Commun.* **184** (2013) 919–985 [[arXiv:hep-ph/0405247](#)].
- [224] **ATLAS** Collaboration, *Measurement of the b-tag efficiency in a sample of jets containing muons with 5 fb⁻¹ of data from the ATLAS detector*, *ATLAS-CONF-2012-043* (2012) [[cds.cern.ch:1435197](#)].
- [225] **ATLAS** Collaboration, *Jet energy scale and its systematic uncertainty in proton-proton collisions at $\sqrt{s} = 7$ TeV with ATLAS 2011*, *ATLAS-CONF-2013-004* (2011) [[cds.cern.ch:1423602](#)].
- [226] **ATLAS** Collaboration, *Charged-particle multiplicities in pp interactions measured with the ATLAS detector at the LHC*, *New J. Phys.* **13** (2011) 053033 [[arXiv:1012.5104](#)].

- [227] D. Krohn, J. Thaler and L.-T. Wang, *Jet Trimming*, *JHEP* **1002** (2010) 084 [[arXiv:0912.1342](#)].
- [228] A. Hocker and V. Kartvelishvili, *SVD approach to data unfolding*, *Nucl.Instrum.Meth.* **A372** (1996) 469–481 [[arXiv:hep-ph/9509307](#)].
- [229] G. D’Agostini, *A Multidimensional unfolding method based on Bayes’ theorem*, *Nucl. Instrum. Meth.* **A362** (1995) 487–498.
- [230] T. Adye, *Unfolding algorithms and tests using RooUnfold*, *Proceedings of the PHYSTAT 2011 Workshop* (2011) 313–318 [[arXiv:1105.1160](#)].
- [231] B. Malaescu, *An Iterative, dynamically stabilized method of data unfolding*, [arXiv:0907.3791](#).
- [232] **ATLAS** Collaboration, *Single hadron response measurement and calorimeter jet energy scale uncertainty with the ATLAS detector at the LHC*, *Eur. Phys. J.* **C73** (2013) 2305 [[arXiv:1203.1302](#)].
- [233] **ATLAS** Collaboration, *Jet energy resolution in proton-proton collisions at $\sqrt{s} = 7$ TeV recorded in 2010 with the ATLAS detector*, *Eur. Phys. J.* **C73** (2013) 2306 [[arXiv:1210.6210](#)].
- [234] **ATLAS** Collaboration, *A study of the material in the ATLAS inner detector using secondary hadronic interactions*, *JINST* **7** (2012) P01013 [[arXiv:1110.6191](#)].
- [235] **ATLAS** Collaboration, *Measurement of the production of a W boson in association with a charm quark in pp collisions at $\sqrt{s} = 7$ TeV with the ATLAS detector*, *JHEP* **05** (2014) 068 [[arXiv:1402.6263](#)].
- [236] **ATLAS** Collaboration, *Measurement of the jet fragmentation function and transverse profile in proton-proton collisions at a center-of-mass energy of 7 TeV with the ATLAS detector*, *Eur. Phys. J.* **C71** (2011) 1795 [[arXiv:1109.5816](#)].

- [237] **ATLAS** Collaboration, *Measurement of track reconstruction inefficiencies in the core of jets via pixel dE/dx with the ATLAS experiment using $\sqrt{s} = 13$ TeV pp collision data*, *ATL-PHYS-PUB-2016-007* (2016) [[cds.cern.ch:2140460](https://cds.cern.ch/record/2140460)].
- [238] **ATLAS** Collaboration, *Search for anomalous production of prompt same-sign lepton pairs and pair-produced doubly charged Higgs bosons with $\sqrt{s} = 8$ TeV pp collisions using the ATLAS detector*, *JHEP* **03** (2015) 041 [[arXiv:1412.0237](https://arxiv.org/abs/1412.0237)].
- [239] J. H. Friedman, *Data analysis techniques for high energy particle physics*, *Proceedings of the CERN School of Computing* (1974) [SLAC-0176].
- [240] D. Krohn, T. Lin, M. D. Schwartz, and W. J. Waalewijn, *Jet Charge at the LHC*, *Phys. Rev. Lett.* **110** (2013) 212001 [[arXiv:1209.2421](https://arxiv.org/abs/1209.2421)].
- [241] **TASSO** Collaboration, *Evidence for Planar Events in e^+e^- Annihilation at High-Energies*, *Phys. Lett.* **B86** (1979) 243–249.
- [242] **MARK-J** Collaboration, *Discovery of Three Jet Events and a Test of Quantum Chromodynamics at PETRA Energies*, *Phys. Rev. Lett.* **43** (1979) 830.
- [243] **PLUTO** Collaboration, *Evidence for Gluon Bremsstrahlung in e^+e^- Annihilations at High-Energies*, *Phys. Lett.* **B86** (1979) 418.
- [244] **JADE** Collaboration, *Observation of planar three jet events in e^+e^- annihilation and evidence for gluon bremsstrahlung*, *Phys. Lett.* **B91** (1980) 142.
- [245] **JADE** Collaboration, *Particle Distribution in Three Jet Events Produced by e^+e^- Annihilation*, *Z. Phys.* **C21** (1983) 37.
- [246] M. Althoff *et. al.*, *A study of 3-jet events in e^+e^- annihilation into hadrons at 34.6 GeV c.m. energy*, *Z. Phys.* **C29** (1985) 29–40.

- [247] H. Aihara *et. al.*, *Tests of models for quark and gluon fragmentation in e^+e^- annihilation at $\sqrt{s} = 29$ GeV*, *Z. Phys.* **C28** (1985) 31–44.
- [248] H. Aihara and others, *Tests of Models for Parton Fragmentation by Means of Three-Jet Events in e^+e^- Annihilation at $\sqrt{s} = 29$ GeV*, *Phys. Rev. Lett.* **54** (1985) 270–273.
- [249] (TPC/Two-Gamma) Collaboration, *Comparison of the Particle Flow in $q\bar{q}g$ and $q\bar{q}\gamma$ Events in e^+e^- Annihilation*, *Phys. Rev. Lett.* **57** (Aug, 1986) 945–948.
- [250] P. D. Sheldon *et. al.*, *Comparison of the particle flow in Three-Jet and radiative Two-Jet Events from e^+e^- Annihilation at E c.m.=29 GeV*, *Phys. Rev. Lett.* **57** (1986) 1398–1401.
- [251] CDF Collaboration, *Evidence for color coherence in $p\bar{p}$ collisions at $\sqrt{s} = 1.8$ TeV*, *Phys. Rev.* **D50** (1994) 5562–5579.
- [252] D0 Collaboration, *Color coherent radiation in multijet events from $p\bar{p}$ collisions at $\sqrt{s} = 1.8$ TeV*, *Phys. Lett.* **B414** (1997) 419–427 [[arXiv:hep-ex/9706012](#)].
- [253] CMS Collaboration, *Probing color coherence effects in pp collisions at $\sqrt{s} = 7$ TeV*, *Eur. Phys. J.* **C74** (2014) 2901 [[arXiv:1311.5815](#)].
- [254] L3 Collaboration, *Search for color reconnection effects in $e^+e^- \rightarrow W^+W^- \rightarrow$ hadrons through particle flow studies at LEP*, *Phys. Lett.* **B561** (2003) 202–212 [[arXiv:hep-ex/0303042](#)].
- [255] DELPHI Collaboration, *Investigation of colour reconnection in WW events with the DELPHI detector at LEP-2*, *Eur. Phys. J.* **C51** (2007) 249–269 [[arXiv:0704.0597](#)].
- [256] L3 Collaboration, *Evidence for gluon interference in hadronic Z decays*, *Phys. Lett.* **B353** (1995) 145–154.

- [257] **ALEPH** Collaboration, *Test of Colour Reconnection Models using Three-Jet Events in Hadronic Z Decays*, *Eur. Phys. J.* **C48** (2006) 685–698 [[arXiv:hep-ex/0604042](#)].
- [258] **D0** Collaboration, *Evidence of color coherence effects in $W + \text{jets}$ events from $p\bar{p}$ collisions at $\sqrt{s} = 1.8$ TeV*, *Phys. Lett.* **B464** (1999) 145–155 [[arXiv:hep-ex/9908017](#)].
- [259] J. Gallicchio *et. al.*, *Multivariate discrimination and the Higgs + W/Z search*, *JHEP* **1104** (2011) 069 [[arXiv:1010.3698](#)].
- [260] **D0** Collaboration, *Search for the standard-model Higgs boson in the $ZH \rightarrow b\bar{b}\nu\bar{\nu}$ channel in 6.4 fb^{-1} of $p\bar{p}$ collisions at $\sqrt{s} = 1.96$ TeV*, *D0 Note 6087-CONF* (2011).
- [261] **CMS** Collaboration, *Search for invisible decays of Higgs bosons in the vector boson fusion and associated ZH production modes*, *Eur. Phys. J.* **C74** (2014) 2980 [[arXiv:1404.1344](#)].
- [262] **CMS** Collaboration, *Search for the standard model Higgs boson produced in association with a W or a Z boson and decaying to bottom quarks*, *Phys. Rev.* **D89** (2014) 012003 [[arXiv:1310.3687](#)].
- [263] **D0** Collaboration, *Measurement of color flow in $t\bar{t}$ events from $p\bar{p}$ collisions at $\sqrt{s} = 1.96$ TeV*, *Phys. Rev.* **D83** (2011) 092002 [[arXiv:1101.0648](#)].
- [264] **ATLAS** Collaboration, *Reconstruction and Modelling of Jet Pull with the ATLAS Detector*, *ATLAS-CONF-2014-048* (2014) [[cds.cern.ch:1741708](#)].
- [265] **ATLAS** Collaboration, *Measurement of colour flow with the jet pull angle in $t\bar{t}$ events using the ATLAS detector at $\sqrt{s} = 8$ TeV*, *Phys. Lett.* **B750** (2015) 475–493 [[arXiv:1506.05629](#)].
- [266] Y. L. Dokshitzer *et. al.*, *Basics of perturbative QCD*. Editions Frontieres, Gif-sur-Yvette, France, 1991.

- [267] S. Catani *et. al.*, *Vector boson production at hadron colliders: a fully exclusive QCD calculation at NNLO*, *Phys. Rev. Lett.* **103** (2009) 082001 [[arXiv:0903.2120](#)].
- [268] N. Kidonakis, *NNLL resummation for s-channel single top quark production*, *Phys. Rev.* **D81** (2010) 054028 [[arXiv:1001.5034](#)].
- [269] N. Kidonakis, *Two-loop soft anomalous dimensions for single top quark associated production with a W or H*, *Phys. Rev.* **D82** (2010) 054018 [[arXiv:1005.4451](#)].
- [270] N. Kidonakis, *Next-to-next-to-leading-order collinear and soft gluon corrections for t-channel single top quark production*, *Phys. Rev.* **D83** (2011) 091503 [[arXiv:1103.2792](#)].
- [271] J. M. Campbell and R. K. Ellis, *An Update on vector boson pair production at hadron colliders*, *Phys. Rev.* **D60** (1999) 113006 [[arXiv:hep-ph/9905386](#)].
- [272] J. M. Campbell, R. K. Ellis and C. Williams, *Vector boson pair production at the LHC*, *JHEP* **07** (2011) 018 [[arXiv:1105.0020](#)].
- [273] S. Frixione, E. Laenen, P. Motylinski, B. R. Webber and C. D. White, *Single-top hadroproduction in association with a W boson*, *JHEP* **0807** (2008) 029 [[arXiv:0805.3067](#)].
- [274] **ATLAS** Collaboration, *First tuning of HERWIG+JIMMY to ATLAS data*, *ATL-PHYS-PUB-2010-014* (2010) [[cds.cern.ch:1303025](#)].
- [275] S. Frixione, P. Nason and B. R. Webber, *Matching NLO QCD and parton showers in heavy flavor production*, *JHEP* **08** (2003) 007 [[arXiv:0305252](#)].
- [276] **ATLAS** Collaboration, *Measurement of the top quark-pair production cross section with ATLAS in pp collisions at $\sqrt{s} = 7$ TeV*, *Eur. Phys. J.* **C71** (2011) 1577 [[arXiv:1012.1792](#)].

- [277] **ATLAS** Collaboration, *Estimation of non-prompt and fake lepton backgrounds in final states with top quarks produced in proton-proton collisions at $\sqrt{s} = 8$ TeV with the ATLAS detector*, ATLAS-CONF-2014-058 (2014) [[cds.cern.ch:1951336](https://cds.cern.ch/1951336)].
- [278] T. P. S. Gillam and C. G. Lester, *Improving estimates of the number of ‘fake’ leptons and other mis-reconstructed objects in hadron collider events: BoB’s your UNCLE*, *JHEP* **11** (2014) 031 [[arXiv:1407.5624](https://arxiv.org/abs/1407.5624)].
- [279] **ATLAS** Collaboration, *Observation and measurement of Higgs boson decays to WW^* with the ATLAS detector*, *Phys. Rev.* **D92** (2015) 012006 [[arXiv:1412.2641](https://arxiv.org/abs/1412.2641)].
- [280] C.-H. Kom and W. J. Stirling, *Charge asymmetry in $W +$ jets production at the LHC*, *Eur. Phys. J.* **C69** (2010) 67–73 [[arXiv:1004.3404](https://arxiv.org/abs/1004.3404)].
- [281] **ATLAS** Collaboration, *Measurements of normalized differential cross sections for $t\bar{t}$ production in pp collisions at $\sqrt{s} = 7$ TeV using the ATLAS detector*, *Phys. Rev.* **D90** (2014) 072004 [[arXiv:1407.0371](https://arxiv.org/abs/1407.0371)].
- [282] J. Alwall *et. al.*, *A Standard format for Les Houches event files*, *Comput. Phys. Commun.* **176** (2007) 300 [[arXiv:hep-ph/0609017](https://arxiv.org/abs/hep-ph/0609017)].
- [283] **ATLAS** Collaboration, *Differential top-antitop cross-section measurements as a function of observables constructed from final-state particles using pp collisions at $\sqrt{s} = 7$ TeV in the ATLAS detector*, *JHEP* **06** (2015) 100 [[arXiv:1502.05923](https://arxiv.org/abs/1502.05923)].
- [284] **ATLAS** Collaboration, *Measurement of Dijet Azimuthal Decorrelations in pp Collisions at $\sqrt{s} = 7$ TeV*, *Phys. Rev. Lett.* **106** (2011) 172002 [[arXiv:1102.2696](https://arxiv.org/abs/1102.2696)].
- [285] **ATLAS** Collaboration, *ATLAS measurements of the properties of jets for boosted particle searches*, *Phys. Rev.* **D86** (2012) 072006 [[arXiv:1206.5369](https://arxiv.org/abs/1206.5369)].

- [286] **ATLAS** Collaboration, *Study of Jet Shapes in Inclusive Jet Production in pp Collisions at $\sqrt{s} = 7$ TeV using the ATLAS Detector*, *Phys. Rev.* **D83** (2011) 052003 [[arXiv:1101.0070](#)].
- [287] **ATLAS** Collaboration, *A measurement of single hadron response using data at $\sqrt{s} = 8$ TeV with the ATLAS detector*, *ATL-PHYS-PUB-2014-002* (2014) [[cds.cern.ch:1668961](#)].
- [288] **ATLAS** Collaboration, *Measurements of jet vetoes and azimuthal decorrelations in dijet events produced in pp collisions at $\sqrt{s} = 7$ TeV using the ATLAS detector*, *Eur. Phys. J.* **C74** (2014) 3117 [[arXiv:1407.5756](#)].
- [289] **ATLAS** Collaboration, *Jet energy measurement with the ATLAS detector in proton-proton collisions at $\sqrt{s} = 7$ TeV*, *Eur. Phys. J.* **C73** (2013) 2304 [[arXiv:1112.6426](#)].
- [290] **ATLAS** Collaboration, *Measurement of the Mistag Rate with 5 fb^{-1} of Data Collected by the ATLAS Detector*, *ATLAS-CONF-2012-040* (2012) [[cds.cern.ch:1435194](#)].
- [291] **ATLAS** Collaboration, *Calibration of b-tagging using dileptonic top pair events in a combinatorial likelihood approach with the ATLAS experiment*, *ATLAS-CONF-2014-004* (2014) [[cds.cern.ch:1664335](#)].
- [292] **ATLAS** Collaboration, *b-jet tagging calibration on c-jets containing D^{*+} mesons*, *ATLAS-CONF-2012-039* (2012) [[cds.cern.ch:1435193](#)].
- [293] A. Altheimer *et. al.*, *Boosted objects and jet substructure at the LHC. Report of BOOST2012, held at IFIC Valencia, 23rd-27th of July 2012*, *Eur. Phys. J.* **C74** (2014) 2792 [[arXiv:1311.2708](#)].
- [294] B. P. Kersevan and E. Richter-Was, *The Monte Carlo event generator AcerMC versions 2.0 to 3.8 with interfaces to PYTHIA 6.4, HERWIG 6.5 and ARIADNE 4.1*, *Comput. Phys. Commun.* **184** (2013) 919–985 [[arXiv:hep-ph/0405247](#)].

- [295] **ATLAS** Collaboration, *Measurement of $t\bar{t}$ production with a veto on additional central jet activity in pp collisions at $\sqrt{s} = 7$ TeV using the ATLAS detector*, *Eur. Phys. J.* **C72** (2012) 2043 [[arXiv:1203.5015](#)].
- [296] **ATLAS** Collaboration, *Comparison of Monte Carlo generator predictions from Powheg and Sherpa to ATLAS measurements of top pair production at 7 TeV*, *ATL-PHYS-PUB-2015-011* (2015) [[cds.cern.ch:2020602](#)].
- [297] A. Buckley *et. al.*, *Rivet user manual*, *Comput. Phys. Commun.* **184** (2013) 2803–2819 [[arXiv:1003.0694](#)].
- [298] **ATLAS** Collaboration, *Search for high-mass diboson resonances with boson-tagged jets in proton-proton collisions at $\sqrt{s} = 8$ TeV with the ATLAS detector*, *JHEP* **12** (2015) 055 [[arXiv:1506.00962](#)].
- [299] **CMS** Collaboration, *Search for the standard model Higgs boson produced through vector boson fusion and decaying to $b\bar{b}$* , *Phys. Rev.* **D92** (2015) 032008 [[arXiv:1506.01010](#)].
- [300] **CMS** Collaboration, *Measurement of electroweak production of two jets in association with a Z boson in proton-proton collisions at $\sqrt{s} = 8$ TeV*, *Eur. Phys. J.* **C75** (2015) 66 [[arXiv:1410.3153](#)].
- [301] **ATLAS** Collaboration, *Measurement of the charged-particle multiplicity inside jets from $\sqrt{s} = 8$ TeV pp collisions with the ATLAS detector*, *Eur. Phys. J.* **C76** (2016) [[arXiv:1602.00988](#)].
- [302] **UA1** Collaboration, *Analysis of the fragmentation properties of quark and gluon jets at the CERN SPS $p\bar{p}$ collider*, *Nucl. Phys.* **B276** (1986) 253.
- [303] **UA2** Collaboration, *Measurement of jet fragmentation properties at the cern $p\bar{p}$ collider*, *Phys. Lett.* **B144** (1984) 291.
- [304] **UA2** Collaboration, *Measurement of production and properties of jets at the CERN $p\bar{p}$ collider*, *Z. Phys.* **C20** (1983) 117.

- [305] **JADE** Collaboration, *Experimental Evidence for Differences in p_T Between Quark Jets and Gluon Jets*, *Phys. Lett.* **B123** (1983) 460.
- [306] **TASSO** Collaboration, *Charged multiplicity distributions and correlations in e^+e^- annihilation at PETRA energies*, *Z. Phys.* **C45** (1989) 193–208.
- [307] **HRS** Collaboration, *Comparison of charged particle multiplicities in quark and gluon jets produced in e^+e^- annihilation at 29 GeV*, *Phys. Lett.* **B165** (1985) 449.
- [308] **MARK II** Collaboration, *Inclusive Charged-Particle Distribution in Nearly Threefold-Symmetric Three-Jet Events at $E_{\text{c.m.}} = 29$ GeV*, *Phys. Rev. Lett.* **55** (1985) 1954.
- [309] **SLD** Collaboration, *Measurement of the charged multiplicity of $Z^0 \rightarrow b\bar{b}$ events*, *Phys. Rev. Lett.* **72** (1994) 3145 [[arXiv:](#)].
- [310] **SLD** Collaboration, *Measurement of the charged multiplicities in b , c and light quark events from Z^0 decays*, *Phys. Lett.* **B386** (1996) 475 [[arXiv:hep-ex/9608008](#)].
- [311] **AMY** Collaboration, *Comparison of quark and gluon jets produced in high-energy e^+e^- annihilations*, *Phys. Rev. Lett.* **63** (1989) 1772.
- [312] **CLEO** Collaboration, *Study of gluon versus quark fragmentation in $\Upsilon \rightarrow g g \gamma$ and $e^+e^- \rightarrow q\bar{q}\gamma$ events at $\sqrt{s}=10$ gev*, *Phys. Rev.* **D56** (1997) 17 [[arXiv:hep-ex/9701006](#)].
- [313] **OPAL** Collaboration, *A Direct observation of quark - gluon jet differences at LEP*, *Phys. Lett.* **B265** (1991) 462–474.
- [314] **OPAL** Collaboration, *A Study of differences between quark and gluon jets using vertex tagging of quark jets*, *Z. Phys.* **C58** (1993) 387.
- [315] **OPAL** Collaboration, *A model independent measurement of quark and gluon jet properties and differences*, *Z. Phys.* **C68** (1995) 179.

- [316] **ALEPH** Collaboration, *Study of the subjet structure of quark and gluon jets*, *Phys. Lett.* **B346** (1995) 389.
- [317] **OPAL** Collaboration, *Test of qcd analytic predictions for the multiplicity ratio between gluon and quark jets*, *Phys. Lett.* **B388** (1996) 659.
- [318] **ALEPH** Collaboration, *Quark and gluon jet properties in symmetric three-jet events*, *Phys. Lett.* **B384** (1996) 353.
- [319] **DELPHI** Collaboration, *Energy dependence of the differences between the quark and gluon jet fragmentation*, *Z. Phys.* **C70** (1996) 179.
- [320] **OPAL** Collaboration, *Multiplicity distributions of gluon and quark jets and tests of QCD analytic predictions*, *Eur. Phys. J.* **C1** (1998) 479 [[arXiv:hep-ex/9708029](#)].
- [321] **DELPHI** Collaboration, *The scale dependence of the hadron multiplicity in quark and gluon jets and a precise determination of C_A/C_F* , *Phys. Lett.* **B449** (1999) 383 [[arXiv:hep-ex/9903073](#)].
- [322] **OPAL** Collaboration, *Experimental properties of gluon and quark jets from a point source*, *Eur. Phys. J.* **C11** (1999) 217 [[arXiv:hep-ex/9903027](#)].
- [323] **OPAL** Collaboration, *A study of coherence of soft gluons in hadron jets*, *Phys. Lett.* **B247** (1990) 617.
- [324] **OPAL** Collaboration, *Scaling violations of quark and gluon jet fragmentation functions in e^+e^- annihilations at $\sqrt{s} = 91.2$ and $183 - 209$ GeV*, *Eur. Phys. J.* **C37** (2004) 25 [[arXiv:0404026](#)].
- [325] **CDF** Collaboration, *Charged particle multiplicity in jets in $p\bar{p}$ collisions at $\sqrt{s} = 1.8$ TeV*, *Phys. Rev. Lett.* **87** (2001) 211804.
- [326] **ATLAS** Collaboration, *Properties of jets measured from tracks in proton-proton collisions at center-of-mass energy $\sqrt{s} = 7$ TeV with the ATLAS detector*, *Phys. Rev.* **D84** (2011) 054001 [[arXiv:1107.3311](#)].

- [327] CMS Collaboration, *Shape, Transverse Size, and Charged Hadron Multiplicity of Jets in pp Collisions at 7 TeV*, *JHEP* **06** (2012) 160 [[arXiv:1204.3170](#)].
- [328] A. J. Larkoski, S. Marzani and J. Thaler, *Sudakov Safety in Perturbative QCD*, *Phys. Rev.* **D91** (2015) 111501 [[arXiv:1502.01719](#)].
- [329] A. J. Larkoski and J. Thaler, *Unsafe but Calculable: Ratios of Angularities in Perturbative QCD*, *JHEP* **09** (2013) 137 [[arXiv:1307.1699](#)].
- [330] A. Capella *et. al.*, *Evolution of average multiplicities of quark and gluon jets*, *Phys. Rev.* **D61** (2000) 074009 [[arXiv:hep-ph/9910226](#)].
- [331] I. M. Dremin and J. W. Gary, *Energy dependence of mean multiplicities in gluon and quark jets at the next-to-next-to-next-to leading order*, *Phys. Lett.* **B459** (1999) 341 [[arXiv:hep-ph/9905477](#)].
- [332] P. Skands, S. Carrazza and J. Rojo, *Tuning PYTHIA 8.1: the Monash 2013 Tune*, *Eur. Phys. J.* **C74** (2014) 3024 [[arXiv:1404.5630](#)].
- [333] ATLAS Collaboration, *ATLAS Run 1 Pythia 8 tunes*, *ATL-PHYS-PUB-2014-021* (2014) [[cds.cern.ch:1966419](#)].
- [334] D. Stump *et. al.*, *Inclusive jet production, parton distributions, and the search for new physics*, *JHEP* **10** (2003) 046 [[arXiv:hep-ph/0303013](#)].
- [335] J. Bellm *et. al.*, *Herwig++ 2.7 Release Note*, [arXiv:1310.6877](#).
- [336] M. H. Seymour and A. Siodmok, *Constraining MPI models using σ_{eff} and recent Tevatron and LHC Underlying Event data*, *JHEP* **10** (2013) 113 [[arXiv:1307.5015](#)].
- [337] A. Buckley *et. al.*, *LHAPDF6: parton density access in the LHC precision era*, *Eur. Phys. J.* **C75** (2015) 132 [[arXiv:1412.7420](#)].
- [338] Z. L. Liu *et. al.*, *Resummation prediction on the jet mass spectrum in one-jet inclusive production at the LHC*, *JHEP* **04** (2015) 005 [[arXiv:1412.1337](#)].

- [339] Y.-T. Chien and M. D. Schwartz, *Resummation of heavy jet mass and comparison to LEP data*, *JHEP* **08** (2010) 058 [[arXiv:1005.1644](#)].
- [340] M. Dasgupta, S. Marzani and G. P. Salam, *QCD calculations for jet substructure*, *Nuovo Cim.* **C37** (2014) 131 [[arXiv:1311.6514](#)].
- [341] A. J. Larkoski, S. Marzani, G. Soyez and J. Thaler, *Soft Drop*, *JHEP* **05** (2014) 146 [[arXiv:1402.2657](#)].
- [342] **ATLAS** Collaboration, *Identification of Boosted, Hadronically Decaying W Bosons and Comparisons with ATLAS Data Taken at $\sqrt{s} = 8$ TeV*, *Eur. Phys. J.* **C76** (2015) [[arXiv:1510.05821](#)].
- [343] **ATLAS** Collaboration, *Performance of jet substructure techniques for large- R jets in proton-proton collisions at $\sqrt{s} = 7$ TeV using the ATLAS detector*, *JHEP* **1309** (2013) 076 [[arXiv:1306.4945](#)].
- [344] **ATLAS** Collaboration, *Study of energy response and resolution of the ATLAS barrel calorimeter to hadrons of energies from 20 GeV to 350 GeV*, *Nucl. Instrum. Meth.* **A621** (2010) 134–150.
- [345] **ATLAS** Collaboration, *Measurement of the jet mass scale and resolution uncertainty for large radius jets at $\sqrt{s} = 8$ TeV using the ATLAS detector*, *ATLAS-CONF-2016-008* (2016) [[cds.cern.ch:2139642](#)].
- [346] **ATLAS** Collaboration, *A search for $t\bar{t}$ resonances using lepton-plus-jets events in proton-proton collisions at $\sqrt{s} = 8$ TeV with the ATLAS detector*, *JHEP* **08** (2015) 148 [[arXiv:1505.07018](#)].
- [347] **ATLAS** Collaboration, *Jet mass reconstruction with the ATLAS Detector in early Run 2 data*, *ATLAS-CONF-2016-035* (2016) [[cds.cern.ch:2200211](#)].
- [348] B. Nachman *et. al.*, *Jets from Jets: Re-clustering as a tool for large radius jet reconstruction and grooming at the LHC.*, *JHEP* **02** (2015) 075 [[arXiv:1407.2922](#)].

- [349] **CMS** Collaboration, *Determination of Jet Energy Calibration and Transverse Momentum Resolution in CMS*, *JINST* **6** (2011) P11002 [[arXiv:1107.4277](#)].
- [350] D. Krohn, J. Thaler and L.-T. Wang, *Jets with Variable R* , *JHEP* **06** (2009) 059 [[arXiv:0903.0392](#)].
- [351] **ATLAS** Collaboration, *Search for new phenomena in final states with large jet multiplicities and missing transverse momentum at $\sqrt{s} = 8$ TeV proton-proton collisions using the ATLAS experiment*, *JHEP* **10** (2013) 130 [[arXiv:1308.1841](#)].
- [352] **ATLAS**. Collaboration, *Search for direct pair production of the top squark in all-hadronic final states in proton-proton collisions at $\sqrt{s} = 8$ TeV with the ATLAS detector*, *JHEP* **1409** (2014) 015 [[arXiv:1406.1122](#)].
- [353] **CMS** Collaboration, *Search for supersymmetry in hadronic final states using MT_2 in pp collisions at $\sqrt{s} = 7$ TeV*, *JHEP* **10** (2012) 018 [[arXiv:1207.1798](#)].
- [354] **CMS** Collaboration, *Inclusive search for squarks and gluinos in pp collisions at $\sqrt{s} = 7$ TeV*, *Phys. Rev.* **D85** (2012) 012004 [[arXiv:1107.1279](#)].
- [355] **CMS**. Collaboration, *Inclusive search for supersymmetry using the razor variables in pp collisions at $\sqrt{s} = 7$ TeV*, *Phys. Rev. Lett.* **111** (2013) 081802 [[arXiv:1212.6961](#)].
- [356] **CMS** Collaboration, *Search for supersymmetry with razor variables in pp collisions at $\sqrt{s} = 7$ TeV*, *Phys. Rev.* **D90** (2014) 112001 [[arXiv:1405.3961](#)].
- [357] **ATLAS** Collaboration, *Boosted hadronic top identification at ATLAS for early 13 TeV data*, *ATL-PHYS-PUB-2015-053* (2015) [[cds.cern.ch:2116351](#)].
- [358] M. Gouzevitch *et. al.*, *Scale-invariant resonance tagging in multijet events and new physics in Higgs pair production*, *JHEP* **07** (2013) 148 [[arXiv:1303.6636](#)].

- [359] CMS Collaboration, *Pileup Jet Identification*, CMS-PAS-JME-13-005 (2013) [[cds.cern.ch:1581583](https://cds.cern.ch/1581583)].
- [360] D. Bertolini *et. al.*, *Pileup Per Particle Identification*, *JHEP* **1410** (2014) 59 [[arXiv:1407.6013](https://arxiv.org/abs/1407.6013)].
- [361] P. Berta *et. al.*, *Particle-level pileup subtraction for jets and jet shapes*, *JHEP* **1406** (2014) 092 [[arXiv:1403.3108](https://arxiv.org/abs/1403.3108)].
- [362] M. Cacciari, G. P. Salam and G. Soyez, *SoftKiller, a particle-level pileup removal method*, *Eur. Phys. J.* **C75** (2015) 59 [[arXiv:1407.0408](https://arxiv.org/abs/1407.0408)].
- [363] M. Cacciari and G. P. Salam, *Dispelling the N^3 myth for the k_t jet-finder*, *Phys. Lett.* **B641** (2006) 57–61 [[arXiv:hep-ph/0512210](https://arxiv.org/abs/hep-ph/0512210)].
- [364] S. D. Ellis, C. K. Vermilion and J. R. Walsh, *Recombination Algorithms and Jet Substructure: Pruning as a Tool for Heavy Particle Searches*, *Phys. Rev.* **D81** (2010) 094023 [[arXiv:0912.0033](https://arxiv.org/abs/0912.0033)].
- [365] S. D. Ellis, C. K. Vermilion and J. R. Walsh, *Techniques for improved heavy particle searches with jet substructure*, *Phys. Rev.* **D80** (2009) 051501 [[arXiv:0903.5081](https://arxiv.org/abs/0903.5081)].
- [366] CMS Collaboration, *Particle-Flow Event Reconstruction in CMS and Performance for Jets, Taus, and MET*, CMS-PAS-PFT-09-001 (2009).
- [367] M. Son, C. Spethmann and B. Tweedie, *Diboson-Jets and the Search for Resonant Zh Production*, *JHEP* **08** (2012) 160 [[arXiv:1204.0525](https://arxiv.org/abs/1204.0525)].
- [368] A. Katz, M. Son and B. Tweedie, *Jet Substructure and the Search for Neutral Spin-One Resonances in Electroweak Boson Channels*, *JHEP* **03** (2011) 011 [[arXiv:1010.5253](https://arxiv.org/abs/1010.5253)].
- [369] S. Schaetzel and M. Spannowsky, *Tagging highly boosted top quarks*, *Phys. Rev.* **D89** (2014) 014007 [[arXiv:1308.0540](https://arxiv.org/abs/1308.0540)].

- [370] T. Plehn *et. al.*, *Stop Reconstruction with Tagged Tops*, *JHEP* **10** (2010) 078 [[arXiv:1006.2833](#)].
- [371] T. Plehn, G. P. Salam and M. Spannowsky, *Fat Jets for a Light Higgs*, *Phys. Rev. Lett.* **104** (2010) 111801 [[arXiv:0910.5472](#)].
- [372] A. J. Larkoski, F. Maltoni and M. Selvaggi, *Tracking down hyper-boosted top quarks*, *JHEP* **06** (2015) 032 [[arXiv:1503.03347](#)].
- [373] S. Bressler *et. al.*, *Hadronic Calorimeter Shower Size: Challenges and Opportunities for Jet Substructure in the Superboosted Regime*, *Phys. Lett. B* **756** (2015) 137–141 [[arXiv:1506.02656](#)].
- [374] ATLAS Collaboration, *A new method to distinguish hadronically decaying boosted Z bosons from W bosons using the ATLAS detector*, [arXiv:1509.04939](#).
- [375] ATLAS Collaboration, *Performance of Boosted W Boson Identification with the ATLAS Detector*, *ATL-PHYS-PUB-2014-004* (2014) [[cds.cern.ch:1690048](#)].
- [376] S. L. Glashow, J. Iliopoulos and L. Maiani, *Weak interactions with lepton-hadron symmetry*, *Phys. Rev.* **D2** (Oct, 1970) 1285–1292.
- [377] ATLAS Collaboration, *A search for flavour changing neutral currents in top-quark decays in pp collision data collected with the ATLAS detector at $\sqrt{s} = 7$ TeV*, *JHEP* **1209** (2012) 139 [[arXiv:1206.0257](#)].
- [378] CMS. Collaboration, *Search for Flavor-Changing Neutral Currents in Top-Quark Decays $t \rightarrow Zq$ in pp Collisions at $\sqrt{s} = 8$ TeV*, *Phys. Rev. Lett.* **112** (2014) 171802 [[arXiv:1312.4194](#)].
- [379] J. Thaler and K. Van Tilburg, *Identifying Boosted Objects with N-subjettiness*, *JHEP* **1103** (2011) 015 [[arXiv:1011.2268](#)].
- [380] J. S. Marshall, A. Münnich and M. A. Thomson, *Performance of Particle Flow Calorimetry at CLIC*, *Nucl. Instrum. Meth.* **A700** (2013) 153–162 [[arXiv:1209.4039](#)].

- [381] N. Cabibbo, *Unitary symmetry and leptonic decays*, *Phys. Rev. Lett.* **10** (1963) 531–533.
- [382] **ATLAS** Collaboration, *Jet mass and substructure of inclusive jets in $\sqrt{s} = 7$ TeV pp collisions with the ATLAS experiment*, *JHEP* **05** (2012) 128 [[arXiv:1203.4606](#)].
- [383] **ATLAS** Collaboration, *Studies of b-tagging performance and jet substructure in a high p_T $g \rightarrow b\bar{b}$ rich sample of large-R jets from pp collisions at $\sqrt{s} = 8$ TeV with the ATLAS detector*, *ATLAS-CONF-2016-002* (2016) [[cds.cern.ch:2135187](#)].
- [384] **CMS** Collaboration, *Measurement of masses in the $t\bar{t}$ system by kinematic endpoints in pp collisions at $\sqrt{s} = 7$ TeV*, *Eur. Phys. J.* **C73** (2013) 2494 [[arXiv:1304.5783](#)].
- [385] **ATLAS** Collaboration, *Measurement of the $WW + WZ$ cross section and limits on anomalous triple gauge couplings using final states with one lepton, missing transverse momentum, and two jets with the ATLAS detector at $\sqrt{s} = 7$ TeV*, *JHEP* **01** (2015) 049 [[arXiv:1410.7238](#)].
- [386] **ATLAS** Collaboration, *Flavor Tagging with Track Jets in Boosted Topologies with the ATLAS Detector*, *ATL-PHYS-PUB-2014-013* (2014).
- [387] L. Mackey *et. al.*, *Fuzzy Jets*, *JHEP* **06** (2016) 010 [[arXiv:1509.02216](#)].
- [388] G. J. McLachlan and D. Peel, *Finite mixture models*. Wiley series in probability and statistics. J. Wiley & Sons, New York, 2000.
- [389] G. Milligan and M. Cooper, *An examination of procedures for determining the number of clusters in a data set*, *Psychometrika* **50** (1985) 159–179.
- [390] R. Tibshirani, G. Walther and T. Hastie, *Estimating the number of clusters in a data set via the gap statistic*, *J. R. Stat. Soc. B* **63** (2001) 411–423.

- [391] S. Catani *et. al.*, *Longitudinally invariant K_t clustering algorithms for hadron hadron collisions*, *Nucl. Phys.* **B406** (1993) 187–224.
- [392] I. W. Stewart *et. al.*, *XCone: N-jettiness as an Exclusive Cone Jet Algorithm*, *JHEP* **11** (2015) 072 [[arXiv:1508.01516](#)].
- [393] J. Thaler and T. F. Wilkason, *Resolving Boosted Jets with XCone*, *JHEP* **12** (2015) 051 [[arXiv:1508.01518](#)].
- [394] D. Y. Grigoriev, E. Jankowski and F. Tkachov, *Optimal jet finder*, *Comput. Phys. Commun.* **155** (2003) 42–64 [[arXiv:hep-ph/0301226](#)].
- [395] S. D. Ellis *et. al.*, *Qjets: A Non-Deterministic Approach to Tree-Based Jet Substructure*, *Phys. Rev. Lett.* **108** (2012) 182003 [[arXiv:1201.1914](#)].
- [396] M. Cacciari and G. P. Salam, *Pileup subtraction using jet areas*, *Phys. Lett.* **B659** (2008) 119–126 [[arXiv:0707.1378](#)].
- [397] J. E. Huth, *et al.*, *Jet energy flow project*, *Proceedings of Snowmass 2001* (1990) 134 [[arXiv:hep-ph/0202207](#)].
- [398] S. D. Ellis *et. al.*, *Jets in hadron-hadron collisions*, *Prog. Part. Nucl. Phys.* **60** (2008) 484–551 [[arXiv:0712.2447](#)].
- [399] H. Hartley, *Maximum Likelihood Estimation from Incomplete Data*, *Biometrics* **14** (1958) 174–194.
- [400] A. Dempster, N. Laird and D. Rubin, *Maximum Likelihood from Incomplete Data via the EM Algorithm*, *Journal of the Royal Society Series B (Methodological)* **39** (1977) 1–38.
- [401] M. G. and T. Krishnan, *The EM Algorithm and Extensions*. Wiley, New York, 1997.
- [402] J. MacQueen, *Some methods for classification and analysis of multivariate observations*, *Proceedings of the Fifth Berkeley Symposium on Mathematical Statistics and Probability, Volume 1: Statistics* (1967) 281–297.

- [403] J. Cogan *et. al.*, *Jet-Images: Computer Vision Inspired Techniques for Jet Tagging*, *JHEP* **02** (2015) 118 [[arXiv:1407.5675](#)].
- [404] L. G. Almeida *et. al.*, *Playing Tag with ANN: Boosted Top Identification with Pattern Recognition*, *JHEP* **07** (2015) 086 [[arXiv:1501.05968](#)].
- [405] L. de Oliveira *et. al.*, *Jet-Images – Deep Learning Edition*, *JHEP* **07** (2016) 1–31 [[arXiv:1511.05190](#)].
- [406] K. Simonyan and A. Zisserman, *Very deep convolutional networks for large-scale image recognition*, [arXiv:1409.1556](#).
- [407] I. J. Goodfellow *et. al.*, *Maxout Networks*, *JMLR WCP 28* (2013) 1319 [[arXiv:1302.4389](#)].
- [408] G. E. Hinton *et. al.*, *Improving neural networks by preventing co-adaptation of feature detectors*, [arXiv:1207.0580](#).
- [409] A. J. Larkoski, D. Neill and J. Thaler, *Jet Shapes with the Broadening Axis*, *JHEP* **04** (2014) 017 [[arXiv:1401.2158](#)].
- [410] I. G. Y. Bengio and A. Courville, “Deep learning.” Book in preparation for MIT Press, 2016.
- [411] J. Schmidhuber, *Deep learning in neural networks: An overview*, *Neural Networks* **61** (2015) 85–117.
- [412] **ATLAS** Collaboration, *Search for direct top squark pair production in final states with one isolated lepton, jets, and missing transverse momentum in $\sqrt{s} = 8$ TeV pp collisions using 13.0 fb⁻¹ of ATLAS data*, *ATLAS-CONF-2012-166* (2012) [[cds.cern.ch:1497732](#)].
- [413] **ATLAS** Collaboration, *Search for direct top squark pair production in final states with one isolated lepton, jets, and missing transverse momentum in $\sqrt{s} = 8$ TeV pp collisions using 21 fb⁻¹ of ATLAS data*, *ATLAS-CONF-2013-037* (2013) [[cds.cern.ch:1532431](#)].

- [414] **ATLAS** Collaboration, *Search for top squark pair production in final states with one isolated lepton, jets, and missing transverse momentum in $\sqrt{s}=8$ TeV pp collisions with the ATLAS detector*, *JHEP* **1411** (2014) 118 [[arXiv:1407.0583](#)].
- [415] **ATLAS** Collaboration, *Search for top squarks in final states with one isolated lepton, jets, and missing transverse momentum in $\sqrt{s}=13$ TeV pp collisions with the ATLAS detector*, *CERN-EP-2016-113* (2016) [[arXiv:1606.03903](#)].
- [416] S. Coleman and J. Mandula, *All Possible Symmetries of the S Matrix*, *Phys. Rev.* **159** (1967) 1251–1256.
- [417] R. Haag, J. T. Lopuszanski and M. Sohnius, *All Possible Generators of Supersymmetries of the S Matrix*, *Nucl. Phys.* **B88** (1975) 257.
- [418] B. Nachman, *1+1 Super Quantum Electrodynamics (QED) from the Symmetries of Superspace, Department of Applied Mathematics and Theoretical Physics Mathematical Tripos Part III Essay* (2013).
- [419] M. Green, J. Schwarz, and E. Witten, *Superstring Theory Volume 1*. Cambridge University Press, Cambridge UK, 1987.
- [420] S. Ferrara, *Supersymmetric gauge theories in two dimensions*, *Nuovo Cimento, Lett.* **13** (1975) 629–35.
- [421] P. Vecchia and S. Ferrara, *Classical solutions in two-dimensional supersymmetric field theories*, *Nuclear Physics* **B130** (1977) 93–104.
- [422] A. Bengtsson and I. Bengtsson, *Some Properties of Supersymmetric QED in 1+1 Dimensions*, *Nuclear Physics* **B231** (1984) 157–171.
- [423] I. Todorov, *Clifford Algebras and Spinors*, *Bulg. J. Phys.* **38** (2011) 3–28 [[arXiv:1106.3197](#)].
- [424] R. Barbieri and G. F. Giudice, *Upper Bounds on Supersymmetric Particle Masses*, *Nucl. Phys.* **B306** (1988) 63.

- [425] B. de Carlos and J. A. Casas, *One loop analysis of the electroweak breaking in supersymmetric models and the fine tuning problem*, *Phys. Lett.* **B309** (1993) 320–328 [[arXiv:hep-ph/9303291](#)].
- [426] A. G. Cohen, D. B. Kaplan and A. E. Nelson, *The More minimal supersymmetric standard model*, *Phys. Lett.* **B388** (1996) 588–598 [[arXiv:hep-ph/9607394](#)].
- [427] S. Dimopoulos and G. F. Giudice, *Naturalness constraints in supersymmetric theories with nonuniversal soft terms*, *Phys. Lett.* **B357** (1995) 573–578 [[arXiv:hep-ph/9507282](#)].
- [428] H. Baer, V. Barger, P. Huang and X. Tata, *Natural Supersymmetry: LHC, dark matter and ILC searches*, *JHEP* **05** (2012) 109 [[arXiv:1203.5539](#)].
- [429] C. Wymant, *Optimising Stop Naturalness*, *Phys. Rev.* **D86** (2012) 115023 [[arXiv:1208.1737](#)].
- [430] M. Papucci, J. T. Ruderman and A. Weiler, *Natural SUSY Endures*, *JHEP* **09** (2012) 035 [[arXiv:1110.6926](#)].
- [431] R. Kitano and Y. Nomura, *Supersymmetry, naturalness, and signatures at the LHC*, *Phys. Rev.* **D73** (2006) 095004 [[arXiv:hep-ph/0602096](#)].
- [432] R. Kitano and Y. Nomura, *A Solution to the supersymmetric fine-tuning problem within the MSSM*, *Phys. Lett.* **B631** (2005) 58–67 [[arXiv:hep-ph/0509039](#)].
- [433] C. Brust, A. Katz, S. Lawrence and R. Sundrum, *SUSY, the Third Generation and the LHC*, *JHEP* **03** (2012) 103 [[arXiv:1110.6670](#)].
- [434] **Planck** Collaboration, *Planck 2013 results. I. Overview of products and scientific results*, *Astron. Astrophys.* **571** (2014) A1 [[arXiv:1303.5062](#)].
- [435] G. Jungman, M. Kamionkowski and K. Griest, *Supersymmetric dark matter*, *Phys. Rept.* **267** (1996) 195–373 [[arXiv:hep-ph/9506380](#)].

- [436] D. Baumann, *Cosmology, Part III Mathematical Tripos* (2014).
- [437] M. Srednicki, *Quantum Field Theory*. Cambridge University Press, Cambridge UK, Cambridge, 2007.
- [438] I. J. R. Aitchison, *Supersymmetry in Particle Physics. An Elementary Introduction*. Cambridge University Press, Cambridge UK, 2007.
- [439] H. Georgi and S. L. Glashow, *Unity of all elementary-particle forces*, *Phys. Rev. Lett.* **32** (1974) 438–441.
- [440] **Particle Data Group**, *Review of Particle Physics*, *Chin. Phys.* **C38** (2014) 090001.
- [441] L. E. Ibanez and G. G. Ross, *Low-Energy Predictions in Supersymmetric Grand Unified Theories*, *Phys. Lett.* **B105** (1981) 439.
- [442] L. E. Ibanez and G. G. Ross, *$SU(2)$ - $L \times U(1)$ Symmetry Breaking as a Radiative Effect of Supersymmetry Breaking in Guts*, *Phys. Lett.* **B110** (1982) 215–220.
- [443] S. P. Martin, *A Supersymmetry primer*, [arXiv:hep-ph/9709356](https://arxiv.org/abs/hep-ph/9709356).
- [444] G. F. Giudice and A. Masiero, *A Natural Solution to the μ Problem in Supergravity Theories*, *Phys. Lett.* **B206** (1988) 480–484.
- [445] J. E. Kim and H. P. Nilles, *The μ Problem and the Strong CP Problem*, *Phys. Lett.* **B138** (1984) 150.
- [446] J. F. Gunion and H. E. Haber, *Higgs Bosons in Supersymmetric Models*, *Nucl. Phys.* **B272** (1986) 1.
- [447] **Super-Kamiokande Collaboration**, *Search for proton decay via $p \rightarrow \nu K^+$ using 260 kiloton-year data of Super-Kamiokande*, *Phys. Rev.* **D90** (2014) 072005 [[arXiv:1408.1195](https://arxiv.org/abs/1408.1195)].

- [448] G. R. Farrar and P. Fayet, *Phenomenology of the Production, Decay, and Detection of New Hadronic States Associated with Supersymmetry*, *Phys. Lett.* **B76** (1978) 575–579.
- [449] S. Dimopoulos and D. W. Sutter, *The Supersymmetric flavor problem*, *Nucl. Phys.* **B452** (1995) 496–512 [[arXiv:9504415](#)].
- [450] P. Fayet, *Supersymmetry and Weak, Electromagnetic and Strong Interactions*, *Phys. Lett.* **B64** (1976) 159.
- [451] P. Fayet, *Spontaneously Broken Supersymmetric Theories of Weak, Electromagnetic and Strong Interactions*, *Phys. Lett.* **B69** (1977) 489.
- [452] P. Fayet, *Relations Between the Masses of the Superpartners of Leptons and Quarks, the Goldstino Couplings and the Neutral Currents*, *Phys. Lett.* **B84** (1979) 416.
- [453] S. Dimopoulos and H. Georgi, *Softly Broken Supersymmetry and $SU(5)$* , *Nucl. Phys.* **B193** (1981) 150.
- [454] J. A. Casas, J. R. Espinosa and H. E. Haber, *The Higgs mass in the MSSM infrared fixed point scenario*, *Nucl. Phys.* **B526** (1998) 3–20 [[arXiv:hep-ph/9801365](#)].
- [455] B. Allanach, *SOFTSUSY: a program for calculating supersymmetric spectra*, *Comput.Phys.Comm.* **143** (2002) 305–331 [[arXiv:hep-ph/0104145](#)].
- [456] H. Baer *et. al.*, *Simulating Supersymmetry with ISAJET 7.0 / ISASUSY 1.0*, [arXiv:hep-ph/9305342](#).
- [457] W. Porod, *SPheno, a program for calculating supersymmetric spectra, SUSY particle decays and SUSY particle production at e^+e^- colliders*, *Comput.Phys.Comm.* **153** (2003) 275–315 [[arXiv:hep-ph/0301101](#)].
- [458] A. Djouadi, J.-L. Kneur and G. Moultaka, *SuSpect: A Fortran code for the supersymmetric and Higgs particle spectrum in the MSSM*, *Comput. Phys. Commun.* **176** (2007) 426–455 [[arXiv:hep-ph/0211331](#)].

- [459] B. C. Allanach, D. P. George and B. Gripaios, *The dark side of the μ : on multiple solutions to renormalisation group equations, and why the CMSSM is not necessarily being ruled out*, *JHEP* **07** (2013) 098 [[arXiv:1304.5462](#)].
- [460] B. C. Allanach, D. P. George and B. Nachman, *Investigating Multiple Solutions in the Constrained Minimal Supersymmetric Standard Model*, *JHEP* **02** (2014) 031 [[arXiv:1311.3960](#)].
- [461] **MSSM Working Group**, *The Minimal supersymmetric standard model: Group summary report*, [arXiv:hep-ph/9901246](#).
- [462] C. F. Berger *et. al.*, *Supersymmetry Without Prejudice*, *JHEP* **02** (2009) 023 [[arXiv:0812.0980](#)].
- [463] **ATLAS Collaboration**, *Summary of the ATLAS experiment's sensitivity to supersymmetry after LHC Run 1 – interpreted in the phenomenological MSSM*, *JHEP* **10** (2015) 134 [[arXiv:1508.06608](#)].
- [464] M. W. Cahill-Rowley *et. al.*, *The New Look pMSSM with Neutralino and Gravitino LSPs*, *Eur. Phys. J.* **C72** (2012) 2156 [[arXiv:1206.4321](#)].
- [465] M. W. Cahill-Rowley *et. al.*, *More energy, more searches, but the phenomenological MSSM lives on*, *Phys. Rev.* **D88** (2013) 035002 [[arXiv:1211.1981](#)].
- [466] M. Cahill-Rowley *et. al.*, *Lessons and prospects from the pMSSM after LHC Run I*, *Phys. Rev.* **D91** (2015) 055002 [[arXiv:1407.4130](#)].
- [467] A. Buckley, *PySLHA: a Pythonic interface to SUSY Les Houches Accord data*, *Eur. Phys. J.* **C75** (2015) 467 [[arXiv:1305.4194](#)].
- [468] J. Alwall *et. al.*, *Searching for Directly Decaying Gluinos at the Tevatron*, *Phys. Lett.* **B666** (2008) 34–37 [[arXiv:0803.0019](#)].
- [469] J. Alwall, P. Schuster and N. Toro, *Simplified Models for a First Characterization of New Physics at the LHC*, *Phys. Rev.* **D79** (2009) 075020 [[arXiv:0810.3921](#)].

- [470] **LHC New Physics Working Group**, *Simplified Models for LHC New Physics Searches*, *J. Phys.* **G39** (2012) 105005 [[arXiv:1105.2838](#)].
- [471] J. R. Ellis *et. al.*, *Supersymmetric Relics from the Big Bang*, *Nucl. Phys.* **B238** (1984) 453–476.
- [472] P. Draper and H. Rzehak, *A Review of Higgs Mass Calculations in Supersymmetric Models*, *Phys. Rept.* **619** (2016) 1–24 [[arXiv:1601.01890](#)].
- [473] A. Bartl *et. al.*, *Production of stop, sbottom, and stau at LEP-2*, *Z. Phys.* **C73** (1997) 469–476 [[arXiv:hep-ph/9603410](#)].
- [474] G. Belanger *et. al.*, *Top Polarization in Stop Production at the LHC*, *JHEP* **05** (2013) 167 [[arXiv:1212.3526](#)].
- [475] M. Perelstein and A. Weiler, *Polarized Tops from Stop Decays at the LHC*, *JHEP* **03** (2009) 141 [[arXiv:0811.1024](#)].
- [476] I. Low, *Polarized charginos (and top quarks) in scalar top quark decays*, *Phys. Rev.* **D88** (2013) 095018 [[arXiv:1304.0491](#)].
- [477] S. Profumo and C. E. Yaguna, *A Statistical analysis of supersymmetric dark matter in the MSSM after WMAP*, *Phys. Rev.* **D70** (2004) 095004 [[arXiv:hep-ph/0407036](#)].
- [478] J. Edsjo *et. al.*, *Accurate relic densities with neutralino, chargino and sfermion coannihilations in mSUGRA*, *JCAP* **0304** (2003) 001 [[arXiv:hep-ph/0301106](#)].
- [479] C. Boehm, A. Djouadi and M. Drees, *Light scalar top quarks and supersymmetric dark matter*, *Phys. Rev.* **D62** (2000) 035012 [[arXiv:hep-ph/9911496](#)].
- [480] J. R. Ellis, K. A. Olive and Y. Santoso, *Calculations of neutralino stop coannihilation in the CMSSM*, *Astropart. Phys.* **18** (2003) 395–432 [[arXiv:0112113](#)].

- [481] W. Buchmuller, R. Ruckl and D. Wyler, *Leptoquarks in Lepton - Quark Collisions*, *Phys. Lett.* **B191** (1987) 442–448.
- [482] M. Freytsis, Z. Ligeti and J. T. Ruderman, *Flavor models for $\bar{B} \rightarrow D^{(*)} \tau \bar{\nu}$* , *Phys. Rev.* **D92** (2015) 054018 [[arXiv:1506.08896](#)].
- [483] F. del Aguila and M. J. Bowick, *The Possibility of New Fermions With $\Delta I = 0$ Mass*, *Nucl. Phys.* **B224** (1983) 107.
- [484] M. Perelstein, *Little Higgs models and their phenomenology*, *Prog. Part. Nucl. Phys.* **58** (2007) 247–291 [[arXiv:0512128](#)].
- [485] M. Perelstein, M. E. Peskin and A. Pierce, *Top quarks and electroweak symmetry breaking in little Higgs models*, *Phys. Rev.* **D69** (2004) 075002 [[arXiv:0310039](#)].
- [486] N. Arkani-Hamed *et. al.*, *The Lightest Higgs*, *JHEP* **07** (2002) 034 [[arXiv:hep-ph/0206021](#)].
- [487] C. T. Hill, *Topcolor assisted technicolor*, *Phys. Lett.* **B345** (1995) 483–489 [[arXiv:9411426](#)].
- [488] D. B. Kaplan and H. Georgi, *$SU(2) \times U(1)$ Breaking by Vacuum Misalignment*, *Phys. Lett.* **B136** (1984) 183.
- [489] D. B. Kaplan, H. Georgi and S. Dimopoulos, *Composite Higgs Scalars*, *Phys. Lett.* **B136** (1984) 187.
- [490] T. Banks, *CONSTRAINTS ON $SU(2) \times U(1)$ BREAKING BY VACUUM MISALIGNMENT*, *Nucl. Phys.* **B243** (1984) 125.
- [491] H. Georgi and D. B. Kaplan, *Composite Higgs and Custodial $SU(2)$* , *Phys. Lett.* **B145** (1984) 216.
- [492] M. J. Dugan, H. Georgi and D. B. Kaplan, *Anatomy of a Composite Higgs Model*, *Nucl. Phys.* **B254** (1985) 299.

- [493] H. Georgi, *A Tool Kit for Builders of Composite Models*, *Nucl. Phys.* **B266** (1986) 274.
- [494] B. Bellazzini, C. Csáki and J. Serra, *Composite Higgses*, *Eur. Phys. J.* **C74** (2014) 2766 [[arXiv:1401.2457](#)].
- [495] H. Georgi, D. B. Kaplan and P. Galison, *Calculation of the Composite Higgs Mass*, *Phys. Lett.* **B143** (1984) 152.
- [496] B. Bhattacharjee *et. al.*, *Model Independent Analysis of Interactions between Dark Matter and Various Quarks*, *JHEP* **04** (2013) 031 [[arXiv:1212.5013](#)].
- [497] T. Lin, E. W. Kolb and L.-T. Wang, *Probing dark matter couplings to top and bottom quarks at the LHC*, *Phys. Rev.* **D88** (2013) 063510 [[arXiv:1303.6638](#)].
- [498] **CDF** Collaboration, *Search for Production of Heavy Particles Decaying to Top Quarks and Invisible Particles in $p\bar{p}$ collisions at $\sqrt{s} = 1.96$ TeV*, *Phys. Rev. Lett.* **106** (2011) 191801 [[arXiv:1103.2482](#)].
- [499] **CDF** Collaboration, *Search for New T' Particles in Final States with Large Jet Multiplicities and Missing Transverse Energy in $p\bar{t}$ Collisions at $\sqrt{s} = 1.96$ TeV*, *Phys. Rev. Lett.* **107** (2011) 191803 [[arXiv:1107.3574](#)].
- [500] **ATLAS** Collaboration, *Search for New Phenomena in $t\bar{t}$ Events With Large Missing Transverse Momentum in Proton-Proton Collisions at $\sqrt{s} = 7$ TeV with the ATLAS Detector*, *Phys. Rev. Lett.* **108** (2012) 041805 [[arXiv:1109.4725](#)].
- [501] **ATLAS** Collaboration, *Search for direct top squark pair production in final states with one isolated lepton, jets, and missing transverse momentum in $\sqrt{s} = 7$ TeV pp collisions using 4.7 fb^{-1} of ATLAS data*, *Phys. Rev. Lett.* **109** (2012) 211803 [[arXiv:1208.2590](#)].
- [502] C. Borschensky *et. al.*, *Squark and gluino production cross sections in pp collisions at $\sqrt{s} = 13, 14, 33$ and 100 TeV*, *Eur. Phys. J.* **C74** (2014) 3174 [[arXiv:1407.5066](#)].

- [503] M. Kramer *et. al.*, *Supersymmetry production cross sections in pp collisions at $\sqrt{s} = 7$ TeV*, [arXiv:1206.2892](#).
- [504] M. Botje *et. al.*, *The PDF4LHC Working Group Interim Recommendations*, [arXiv:1101.0538](#).
- [505] J. Alwall *et. al.*, *The automated computation of tree-level and next-to-leading order differential cross sections, and their matching to parton shower simulations*, *JHEP* **1407** (2014) 079 [[arXiv:1405.0301](#)].
- [506] **OPAL** Collaboration, *Search for scalar top and scalar bottom quarks at $\sqrt{s} = 189$ GeV at LEP*, *Phys. Lett.* **B456** (1999) 95–106 [[arXiv:hep-ex/9903070](#)].
- [507] **L3** Collaboration, *Search for scalar leptons and scalar quarks at LEP*, *Phys. Lett.* **B580** (2004) 37–49 [[arXiv:hep-ex/0310007](#)].
- [508] **DELPHI** Collaboration, *Searches for supersymmetric particles in e^+e^- collisions up to 208 GeV and interpretation of the results within the MSSM*, *Eur. Phys. J.* **C31** (2003) 421–479 [[arXiv:hep-ex/0311019](#)].
- [509] **ALEPH** Collaboration, *Search for scalar quarks in e^+e^- collisions at $\sqrt{s}$ up to 209 GeV*, *Phys. Lett.* **B537** (2002) 5–20 [[arXiv:hep-ex/0204036](#)].
- [510] **D0** Collaboration, *Search for scalar top quarks in the acoplanar charm jets and missing transverse energy final state in $p\bar{p}$ collisions at $\sqrt{s} = 1.96$ TeV*, *Phys. Lett.* **B665** (2008) 1–8 [[arXiv:0803.2263](#)].
- [511] **CDF** Collaboration, *Search for Scalar Top Quark Production in $p\bar{p}$ Collisions at $\sqrt{s} = 1.96$ TeV*, *JHEP* **10** (2012) 158 [[arXiv:1203.4171](#)].
- [512] **ATLAS** Collaboration, *Performance and Calibration of the JetFitterCharm Algorithm for c-Jet Identification*, *ATL-PHYS-PUB-2015-001* (2015) [[cds.cern.ch:1980463](#)].

- [513] **ATLAS** Collaboration, *Search for pair-produced third-generation squarks decaying via charm quarks or in compressed supersymmetric scenarios in pp collisions at $\sqrt{s} = 8$ TeV with the ATLAS detector*, *Phys. Rev.* **D90** (2014) 052008 [[arXiv:1407.0608](#)].
- [514] R. Grober *et. al.*, *Light stop decays into $Wb\tilde{\chi}_1^0$ near the kinematic threshold*, *Phys. Lett.* **B747** (2015) 144–151 [[arXiv:1502.05935](#)].
- [515] K. Hagiwara and T. Yamada, *Equal-velocity scenario for hiding dark matter at the LHC*, *Phys. Rev.* **D91** (2015) 094007 [[arXiv:1307.1553](#)].
- [516] H. An and L.-T. Wang, *Opening up the compressed region of top squark searches at 13 TeV LHC*, *Phys. Rev. Lett.* **115** (2015) 181602 [[arXiv:1506.00653](#)].
- [517] S. Macaluso *et. al.*, *Revealing Compressed Stops Using High-Momentum Recoils*, *JHEP* **03** (2016) 151 [[arXiv:1506.07885](#)].
- [518] H.-C. Cheng *et. al.*, *Stop Search in the Compressed Region via Semileptonic Decays*, *JHEP* **05** (2016) 036 [[arXiv:1604.00007](#)].
- [519] B. Dutta *et. al.*, *Probing compressed top squark scenarios at the LHC at 14 TeV*, *Phys. Rev.* **D90** (2014) 095022 [[arXiv:1312.1348](#)].
- [520] **ATLAS** Collaboration, *Measurement of the $t\bar{t}$ production cross-section using $e\mu$ events with b-tagged jets in pp collisions at $\sqrt{s} = 7$ and 8 TeV with the ATLAS detector*, *Eur. Phys. J.* **C74** (2014), no. 10 3109 [[arXiv:1406.5375](#)].
- [521] M. Czakon *et. al.*, *Closing the stop gap*, *Phys. Rev. Lett.* **113** (2014) 201803 [[arXiv:1407.1043](#)].
- [522] **ATLAS** Collaboration, *Measurement of Spin Correlation in Top-Antitop Quark Events and Search for Top Squark Pair Production in pp Collisions at $\sqrt{s} = 8$ TeV Using the ATLAS Detector*, *Phys. Rev. Lett.* **114** (2015) 142001 [[arXiv:1412.4742](#)].

- [523] **ATLAS** Collaboration, *ATLAS Run 1 searches for direct pair production of third-generation squarks at the Large Hadron Collider*, *Eur. Phys. J.* **C75** (2015) 510 [[arXiv:1506.08616](#)].
- [524] **CMS** Collaboration, *Measurement of the $t\bar{t}$ production cross section in the $e\mu$ channel in proton-proton collisions at $\sqrt{s} = 7$ and 8 TeV*, [arXiv:1603.02303](#).
- [525] T. Eifert and B. Nachman, *Sneaky light stop*, *Phys. Lett. B* **B743** (2015) [[arXiv:1410.7025](#)].
- [526] **CMS** Collaboration, *Measurement of the top-quark mass in $t\bar{t}$ events with lepton+jets final states in pp collisions at $\sqrt{s} = 8$ TeV*, *CMS-PAS-TOP-14-001* (2014) [[cds.cern.ch:1690093](#)].
- [527] S. Moch, *Precision determination of the top-quark mass*, *PoS LL2014* (2014) 054 [[arXiv:1408.6080](#)].
- [528] **ATLAS** Collaboration, *Measurement of the top quark mass in the $t\bar{t} \rightarrow \text{lepton} + \text{jets}$ and $t\bar{t} \rightarrow \text{dilepton}$ channels using $\sqrt{s} = 7$ TeV ATLAS data*, *Eur. Phys. J.* **C75** (2015) 330 [[arXiv:1503.05427](#)].
- [529] J. Alwall *et. al.*, *The automated computation of tree-level and next-to-leading order differential cross sections, and their matching to parton shower simulations*, *JHEP* **2014** (2014) 158 [[arXiv:1405.0301](#)].
- [530] S. Catani *et. al.*, *QCD matrix elements + parton showers*, *JHEP* **11** (2001) 063 [[arXiv:hep-ph/0109231](#)].
- [531] L. Lonnblad, *Correcting the color dipole cascade model with fixed order matrix elements*, *JHEP* **05** (2002) 046 [[arXiv:hep-ph/0112284](#)].
- [532] E. Re, *Single-top Wt -channel production matched with parton showers using the POWHEG method*, *Eur. Phys. J.* **C71** (2011) 1547 [[arXiv:1009.2450](#)].
- [533] R. Frederix, E. Re and P. Torrielli, *Single-top t -channel hadroproduction in the four-flavour scheme with POWHEG and aMC@NLO*, *JHEP* **1209** (2012) 130 [[arXiv:1207.5391](#)].

- [534] S. Alioli *et. al.*, *NLO single-top production matched with shower in POWHEG: s- and t-channel contributions*, *JHEP* **09** (2009) 111 [[arXiv:0907.4076](#)].
[Erratum: JHEP02,011(2010)].
- [535] T. Gleisberg and S. Höche, *Comix, a new matrix element generator*, *JHEP* **0812** (2008) 039 [[arXiv:0808.3674](#)].
- [536] F. Cascioli, P. Maierhofer and S. Pozzorini, *Scattering Amplitudes with Open Loops*, *Phys. Rev. Lett.* **108** (2012) 111601 [[arXiv:1111.5206](#)].
- [537] S. Schumann and F. Krauss, *A Parton shower algorithm based on Catani-Seymour dipole factorisation*, *JHEP* **0803** (2008) 038 [[arXiv:0709.1027](#)].
- [538] **ATLAS** Collaboration, *Simulation of top quark production for the ATLAS experiment at $\sqrt{s} = 13$ TeV*, *ATL-PHYS-PUB-2016-004* (2016) [[cds.cern.ch:2120417](#)].
- [539] **ATLAS** Collaboration, *Monte Carlo Generators for the Production of a W or Z/ γ^* Boson in Association with Jets at ATLAS in Run 2*, *ATL-PHYS-PUB-2016-003* (2016) [[cds.cern.ch:2120133](#)].
- [540] **ATLAS** Collaboration, *Modelling of the $t\bar{t}H$ and $t\bar{t}V$ ($V = W, Z$) processes for $\sqrt{s} = 13$ TeV ATLAS analyses*, *ATL-PHYS-PUB-2016-005* (2016) [[cds.cern.ch:2120826](#)].
- [541] **ATLAS** Collaboration, *Multi-Boson Simulation for 13 TeV ATLAS Analyses*, *ATL-PHYS-PUB-2016-002* (2016) [[cds.cern.ch:2119986](#)].
- [542] J. M. Campbell and R. K. Ellis, *$t\bar{t}W^\pm$ production and decay at NLO*, *JHEP* **07** (2012) 052 [[arXiv:1204.5678](#)].
- [543] M. V. Garzelli *et. al.*, *$t\bar{t} W^\pm$ and $t\bar{t} Z$ Hadroproduction at NLO accuracy in QCD with Parton Shower and Hadronization effects*, *JHEP* **11** (2012) 056 [[arXiv:1208.2665](#)].

- [544] K. Melnikov, M. Schulze and A. Scharf, *QCD corrections to top quark pair production in association with a photon at hadron colliders*, *Phys. Rev.* **D83** (2011) 074013 [[arXiv:1102.1967](#)].
- [545] **ATLAS** Collaboration, *Tagging and suppression of pileup jets with the ATLAS detector*, *ATLAS-CONF-2014-018* (2014) [[cds.cern.ch:1700870](#)].
- [546] **ATLAS** Collaboration, *Selection of jets produced in 13 TeV proton-proton collisions with the ATLAS detector*, *ATLAS-CONF-2015-029* (2015) [[cds.cern.ch:2037702](#)].
- [547] **ATLAS** Collaboration, *Expected performance of the ATLAS b-tagging algorithms in Run-2*, *ATL-PHYS-PUB-2015-022* (2015) [[cds.cern.ch:2037697](#)].
- [548] **ATLAS** Collaboration, *Commissioning of the ATLAS b-tagging algorithms using $t\bar{t}$ events in early Run-2 data*, *ATL-PHYS-PUB-2015-039* (2015) [[cds.cern.ch:2047871](#)].
- [549] **ATLAS** Collaboration, *Commissioning of the reconstruction of hadronic tau lepton decays in ATLAS using pp collisions at $\sqrt{s} = 13$ TeV*, *ATL-PHYS-PUB-2015-025* (2015) [[cds.cern.ch:2037716](#)].
- [550] **ATLAS** Collaboration, *Reconstruction, Energy Calibration, and Identification of Hadronically Decaying Tau Leptons in the ATLAS Experiment for Run-2 of the LHC*, *ATL-PHYS-PUB-2015-045* (2015) [[cds.cern.ch:2064383](#)].
- [551] **ATLAS** Collaboration, *Electron efficiency measurements with the ATLAS detector using the 2012 LHC proton-proton collision data*, *ATLAS-CONF-2014-032* (2014) [[cds.cern.ch:1706245](#)].
- [552] **ATLAS** Collaboration, *Electron identification measurements in ATLAS using $\sqrt{s} = 13$ TeV data with 50 ns bunch spacing*, *ATL-PHYS-PUB-2015-041* (2015) [[cds.cern.ch:2048202](#)].

- [553] **ATLAS** Collaboration, *Measurements of the photon identification efficiency with the ATLAS detector using 4.9 fb^{-1} of pp collision data collected in 2011*, *ATLAS-CONF-2012-123* (2012) [[cds.cern.ch:1473426](https://cds.cern.ch/record/1473426)].
- [554] **ATLAS** Collaboration, *Measurement of the inclusive isolated prompt photon cross section in pp collisions at $\sqrt{s} = 7 \text{ TeV}$ with the ATLAS detector*, *Phys. Rev.* **D83** (2011) 052005 [[arXiv:1012.4389](https://arxiv.org/abs/1012.4389)].
- [555] **ATLAS** Collaboration, *Performance of the Reconstruction and Identification of Hadronic Tau Decays with ATLAS*, *ATLAS-CONF-2011-152* (2011) [[cds.cern.ch:1398195](https://cds.cern.ch/record/1398195)].
- [556] **ATLAS** Collaboration, *Performance of the Reconstruction and Identification of Hadronic Tau Decays in ATLAS with 2011 Data*, *ATLAS-CONF-2012-142* (2012) [[cds.cern.ch:1485531](https://cds.cern.ch/record/1485531)].
- [557] **CDF** Collaboration, *Precise measurement of the W-boson mass with the CDF II detector*, *Phys. Rev. Lett.* **108** (2012) 151803 [[arXiv:1203.0275](https://arxiv.org/abs/1203.0275)].
- [558] **D0** Collaboration, *Measurement of the W Boson Mass with the D0 Detector*, *Phys. Rev. Lett.* **108** (2012) 151804 [[arXiv:1203.0293](https://arxiv.org/abs/1203.0293)].
- [559] **ATLAS** Collaboration, *Measurements of the W production cross sections in association with jets with the ATLAS detector*, *Eur. Phys. J.* **C75** (2015) 82 [[arXiv:1409.8639](https://arxiv.org/abs/1409.8639)].
- [560] C. G. Lester and D. J. Summers, *Measuring masses of semi-invisibly decaying particles pair produced at hadron colliders*, *Phys. Lett.* **B463** (1999) 99–103 [[arXiv:hep-ph/9906349](https://arxiv.org/abs/hep-ph/9906349)].
- [561] A. Barr, C. Lester and P. Stephens, *$m(T2)$: The Truth behind the glamour*, *J. Phys.* **G29** (2003) 2343–2363 [[arXiv:hep-ph/0304226](https://arxiv.org/abs/hep-ph/0304226)].
- [562] A. J. Barr *et. al.*, *Guide to transverse projections and mass-constraining variables*, *Phys. Rev.* **D84** (2011) 095031 [[arXiv:1105.2977](https://arxiv.org/abs/1105.2977)].

- [563] A. J. Barr and C. G. Lester, *A Review of the Mass Measurement Techniques proposed for the Large Hadron Collider*, *J. Phys.* **G37** (2010) 123001 [[arXiv:1004.2732](#)].
- [564] P. Konar *et. al.*, *Superpartner Mass Measurement Technique using 1D Orthogonal Decompositions of the Cambridge Transverse Mass Variable M_{T2}* , *Phys. Rev. Lett.* **105** (2010) 051802 [[arXiv:0910.3679](#)].
- [565] M. Burns *et. al.*, *Using Subsystem MT_2 for Complete Mass Determinations in Decay Chains with Missing Energy at Hadron Colliders*, *JHEP* **03** (2009) 143 [[arXiv:0810.5576](#)].
- [566] C. Lester and A. Barr, *MTGEN: Mass scale measurements in pair-production at colliders*, *JHEP* **12** (2007) 102 [[arXiv:0708.1028](#)].
- [567] W. S. Cho *et. al.*, *Measuring superparticle masses at hadron collider using the transverse mass kink*, *JHEP* **02** (2008) 035 [[arXiv:0711.4526](#)].
- [568] P. Konar *et. al.*, *Dark Matter Particle Spectroscopy at the LHC: Generalizing $M(T_2)$ to Asymmetric Event Topologies*, *JHEP* **04** (2010) 086 [[arXiv:hep-ph/0911.4126](#)].
- [569] H.-C. Cheng and Z. Han, *Minimal Kinematic Constraints and $m(T_2)$* , *JHEP* **12** (2008) 063 [[arXiv:0810.5178](#)].
- [570] C. G. Lester and B. Nachman, *Bisection-based asymmetric M_{T2} computation: a higher precision calculator than existing symmetric methods*, *JHEP* **1503** (2015) 100 [[arXiv:1411.4312](#)].
- [571] F. Etayo, L. Gonzalez-Vega and N. del Rio, *A new approach to characterizing the relative position of two ellipses depending on one parameter*, *Computer Aided Geometric Design* **23** (2006) 324 – 350.
- [572] A. J. Barr, B. Gripaios and C. G. Lester, *Transverse masses and kinematic constraints: from the boundary to the crease*, *JHEP* **0911** (2009) 096 [[arXiv:0908.3779](#)].

- [573] Y. Bai *et. al.*, *Stop the Top Background of the Stop Search*, *JHEP* **07** (2012) 110 [[arXiv:hep-ph/1203.4813](#)].
- [574] D. R. Tovey, *On measuring the masses of pair-produced semi-invisibly decaying particles at hadron colliders*, *JHEP* **04** (2008) 034 [[arXiv:0802.2879](#)].
- [575] K. T. Matchev and M. Park, *A General method for determining the masses of semi-invisibly decaying particles at hadron colliders*, *Phys. Rev. Lett.* **107** (2011) 061801 [[arXiv:0910.1584](#)].
- [576] G. Polesello and D. R. Tovey, *Supersymmetric particle mass measurement with the boost-corrected contranverse mass*, *JHEP* **03** (2010) 030 [[arXiv:0910.0174](#)].
- [577] **BABAR** Collaboration, *The BABAR Physics Book*, *SLAC-R-0504* (1998).
- [578] A. Hocker *et. al.*, *TMVA - Toolkit for Multivariate Data Analysis*, *PoS ACAT* (2007) 040 [[physics/0703039](#)].
- [579] B. Nachman and C. G. Lester, *Significance Variables*, *Phys. Rev.* **D88** (2013) 075013 [[arXiv:1303.7009](#)].
- [580] Bruce Knuteson and others, *$p(\text{met})$: The missing transverse energy resolution of an event*, *D0-3629* (1999) [[inspirehep:1199577](#)].
- [581] **CMS** Collaboration, *Missing transverse energy performance of the CMS detector*, *JINST* **6** (2011) P09001 [[arXiv:1106.5048](#)].
- [582] J. Neyman and E. S. Pearson, *On the problem of the most efficient tests of statistical hypotheses*, *Philosophical Transactions of the Royal Society of London A: Mathematical, Physical and Engineering Sciences* **231** (1933) 289–337.
- [583] E. J. Gumbel, *Les valeurs extrêmes des distributions statistiques*, *Annales de l'Institut Henri Poincaré* **5** (1935) 115–158.

- [584] **ATLAS** Collaboration, *ATLAS search for a heavy gauge boson decaying to a charged lepton and a neutrino in pp collisions at $\sqrt{s} = 7$ TeV*, *Eur. Phys. J.* **C72** (2012) 2241 [[arXiv:1209.4446](#)].
- [585] **CMS** Collaboration, *Search for new physics in final states with a lepton and missing transverse energy in pp collisions at the LHC*, *Phys. Rev.* **D87** (2013) 072005 [[arXiv:1302.2812](#)].
- [586] **ATLAS** Collaboration, *Search for the Standard Model Higgs boson in the H to $\tau^+\tau^-$ decay mode in $\sqrt{s} = 7$ TeV pp collisions with ATLAS*, *JHEP* **09** (2012) 070 [[arXiv:1206.5971](#)].
- [587] **CMS** Collaboration, *Search for neutral Higgs bosons decaying to tau pairs in pp collisions at $\sqrt{s} = 7$ TeV*, *Phys. Lett.* **B713** (2012) 68–90 [[arXiv:1202.4083](#)].
- [588] C. G. Lester, *The stransverse mass, MT_2 , in special cases*, *JHEP* **05** (2011) 076 [[arXiv:1103.5682](#)].
- [589] J. A. Nelder and R. Mead, *A simplex method for function minimization*, *The Computer Journal* **7** (1965) 308–313.
- [590] F. James and M. Roos, *Minuit: A System for Function Minimization and Analysis of the Parameter Errors and Correlations*, *Comput. Phys. Commun.* **10** (1975) 343–367.
- [591] W. C. Davidon, *Variable metric method for minimization*, *SIAM Journal on Optimization* **1** (1991) 1–17.
- [592] R. Fletcher and M. J. D. Powell, *A rapidly convergent descent method for minimization*, *The Computer Journal* **6** (1963) 163–168.
- [593] R. Fletcher, *A new approach to variable metric algorithms*, *The Computer Journal* **13** (1970) 317–322.

- [594] M. L. Graesser and J. Shelton, *Hunting Mixed Top Squark Decays*, *Phys. Rev. Lett.* **111** (2013) 121802 [[arXiv:1212.4495](#)].
- [595] L. Moneta *et. al.*, *The RooStats Project*, *PoS ACAT2010* (2010) 057 [[arXiv:1009.1003](#)].
- [596] **ATLAS** Collaboration, *Measurement of jets produced in top quark events using the di-lepton final state with 2 b-tagged jets in pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector*, *ATLAS-CONF-2015-065* (2015) [[cds.cern.ch:2114832](#)].
- [597] **ATLAS** Collaboration, *Measurement of the production cross-section of a single top quark in association with a W boson at 8 TeV with the ATLAS experiment*, *JHEP* **01** (2016) 064 [[arXiv:1510.03752](#)].
- [598] **CMS** Collaboration, *Observation of the associated production of a single top quark and a W boson in pp collisions at $\sqrt{s} = 8$ TeV*, *Phys. Rev. Lett.* **112** (2014), no. 23 231802 [[arXiv:1401.2942](#)].
- [599] **ATLAS** Collaboration, *Measurement of the $t\bar{t}Z$ and $t\bar{t}W$ production cross sections in multilepton final states using 3.2 fb^{-1} of pp collisions at 13 TeV at the LHC*, *ATLAS-CONF-2016-003* (2016) [[cds.cern.ch:2138947](#)].
- [600] **ATLAS** Collaboration, *Measurement of the $t\bar{t}W$ and $t\bar{t}Z$ production cross sections in pp collisions at $\sqrt{s} = 8$ TeV with the ATLAS detector*, *JHEP* **11** (2015) 172 [[arXiv:1509.05276](#)].
- [601] **CMS** Collaboration, *Measurement of the cross section of top quark pair production in association with a Z boson in pp collisions at 13 TeV*, *CMS-PAS-TOP-16-009* (2016) [[cds.cern.ch:2139263](#)].
- [602] **CMS** Collaboration, *Observation of top quark pairs produced in association with a vector boson in pp collisions at $\sqrt{s} = 8$ TeV*, *JHEP* **01** (2016) 096 [[arXiv:1510.01131](#)].

- [603] **CDF** Collaboration, *First Observation of Vector Boson Pairs in a Hadronic Final State at the Tevatron Collider*, *Phys. Rev. Lett.* **103** (2009) 091803 [[arXiv:0905.4714](#)].
- [604] **ATLAS** Collaboration, *Search for squarks and gluinos using final states with jets and missing transverse momentum with the ATLAS detector in $\sqrt{s} = 7$ TeV proton-proton collisions*, *Phys. Lett.* **B701** (2011) 186–203 [[arXiv:1102.5290](#)].
- [605] **CMS** Collaboration, *Search for Supersymmetry in pp Collisions at 7 TeV in Events with Jets and Missing Transverse Energy*, *Phys. Lett.* **B698** (2011) 196–218 [[arXiv:1101.1628](#)].
- [606] **CMS** Collaboration, *Search for New Physics with Jets and Missing Transverse Momentum in pp collisions at $\sqrt{s} = 7$ TeV*, *JHEP* **08** (2011) 155 [[arXiv:1106.4503](#)].
- [607] **CMS** Collaboration, *Data-Driven Estimation of the Invisible Z Background to the SUSY MET Plus Jets Search*, *CMS-PAS-SUS-08-002* (2009) [[cds.cern.ch:1194471](#)].
- [608] Z. Bern *et. al.*, *Driving Missing Data at Next-to-Leading Order*, *Phys. Rev.* **D84** (2011) 114002 [[arXiv:1106.1423](#)].
- [609] Z. Bern *et. al.*, *Missing Energy and Jets for Supersymmetry Searches*, *Phys. Rev.* **D87** (2013) 034026 [[arXiv:1206.6064](#)].
- [610] S. Ask *et. al.*, *Using γ +jets Production to Calibrate the Standard Model $Z(\rightarrow \nu\bar{\nu})$ +jets Background to New Physics Processes at the LHC*, *JHEP* **10** (2011) 058 [[arXiv:1107.2803](#)].
- [611] **CMS** Collaboration, *Comparison of the $Z/\gamma + jets$ to $\gamma + jets$ cross sections in pp collisions at $\sqrt{s} = 8$ TeV*, *JHEP* **10** (2015) 128 [[arXiv:1505.06520](#)].
- [612] **CDF** Collaboration, *Evidence for $t\bar{t}\gamma$ Production and Measurement of $\sigma_{t\bar{t}\gamma}/\sigma_{t\bar{t}}$* , *Phys. Rev.* **D84** (2011) 031104 [[arXiv:1106.3970](#)].

- [613] **ATLAS** Collaboration, *Observation of top-quark pair production in association with a photon and measurement of the $t\bar{t}\gamma$ production cross section in pp collisions at $\sqrt{s} = 7$ TeV using the ATLAS detector*, *Phys. Rev.* **D91** (2015) 072007 [[arXiv:1502.00586](#)].
- [614] **CMS** Collaboration, *Measurement of the inclusive top-quark pair + photon production cross section in the muon + jets channel in pp collisions at 8 TeV*, *CMS-PAS-TOP-13-011* (2014) [[cds.cern.ch:1644573](#)].
- [615] A. Hook and A. Katz, *Unbroken $SU(2)$ at a 100 TeV collider*, *JHEP* **09** (2014) 175 [[arXiv:1407.2607](#)].
- [616] M. Dobbs and J. B. Hansen, *The HepMC C++ Monte Carlo event record for High Energy Physics*, *Comput. Phys. Commun.* **134** (2001) 41–46.
- [617] **ATLAS** Collaboration, *Search for $t\bar{t}H$ production in the $H \rightarrow \gamma\gamma$ channel at $\sqrt{s} = 8$ TeV with the ATLAS detector*, *ATLAS-CONF-2013-080* (2013) [[cds.cern.ch:1564319](#)].
- [618] S. Frixione, *Isolated photons in perturbative QCD*, *Phys. Lett.* **B429** (1998) 369–374 [[arXiv:9801442](#)].
- [619] A. Lazopoulos, T. McElmurry, K. Melnikov and F. Petriello, *Next-to-leading order QCD corrections to $t\bar{t}Z$ production at the LHC*, *Phys. Lett.* **B666** (2008) 62–65 [[arXiv:0804.2220](#)].
- [620] A. Kardos, Z. Trocsanyi and C. Papadopoulos, *Top quark pair production in association with a Z-boson at NLO accuracy*, *Phys. Rev.* **D85** (2012) 054015 [[arXiv:1111.0610](#)].
- [621] R. Röntsch and M. Schulze, *Constraining couplings of top quarks to the Z boson in $t\bar{t} + Z$ production at the LHC*, *JHEP* **07** (2014) 091 [[arXiv:1404.1005](#)].

- [622] S. Frixione *et. al.*, *Electroweak and QCD corrections to top-pair hadroproduction in association with heavy bosons*, *JHEP* **06** (2015) 184 [[arXiv:1504.03446](#)].
- [623] V. Hirschi and O. Mattelaer, *Automated event generation for loop-induced processes*, *JHEP* **10** (2015) 146 [[arXiv:1507.00020](#)].
- [624] **ATLAS** Collaboration, *Jet Calibration and Systematic Uncertainties for Jets Reconstructed in the ATLAS Detector at $\sqrt{s} = 13$ TeV*, *ATL-PHYS-PUB-2015-015* (2015) [[cds.cern.ch:2037613](#)].
- [625] **ATLAS** Collaboration, *A method for the construction of strongly reduced representations of ATLAS experimental uncertainties and the application thereof to the jet energy scale*, *ATL-PHYS-PUB-2015-014* (2015) [[cds.cern.ch:2037436](#)].
- [626] **ATLAS** Collaboration, *ATLAS tunes of PYTHIA 6 and Pythia 8 for MC11*, *ATL-PHYS-PUB-2011-009* (2011) [[cds.cern.ch:1363300](#)].
- [627] **ATLAS** Collaboration, *Comparison of Monte Carlo generator predictions for gap fraction and jet multiplicity observables in top-antitop events*, *ATL-PHYS-PUB-2014-005* (2014) [[cds.cern.ch:1703034](#)].
- [628] **ATLAS** Collaboration, *Comparison of Monte Carlo generator predictions to ATLAS measurements of top pair production at 7 TeV*, *ATL-PHYS-PUB-2015-002* (2015) [[cds.cern.ch:1981319](#)].
- [629] C. D. White *et. al.*, *Isolating Wt production at the LHC*, *JHEP* **0911** (2009) 074 [[arXiv:0908.0631](#)].
- [630] F. Cascioli *et. al.*, *A unified NLO description of top-pair and associated Wt production*, *Eur. Phys. J.* **C74** (2014) 2783 [[arXiv:1312.0546](#)].
- [631] R. Frederix, *Top Quark Induced Backgrounds to Higgs Production in the $WW^{(*)} \rightarrow l\bar{\nu}\nu$ Decay Channel at Next-to-Leading-Order in QCD*, *Phys. Rev. Lett.* **112** (2014) 082002 [[arXiv:1311.4893](#)].

- [632] T. Ježo *et. al.*, *An NLO+PS generator for $t\bar{t}$ and Wt production and decay including non-resonant and interference effects*, [arXiv:1607.04538](#).
- [633] **ATLAS** Collaboration, *Measurement of the cross-section for W boson production in association with b -jets in pp collisions at $\sqrt{s} = 7$ TeV with the ATLAS detector*, *JHEP* **06** (2013) 084 [[arXiv:1302.2929](#)].
- [634] M. L. Mangano, M. Moretti and R. Pittau, *Multijet matrix elements and shower evolution in hadronic collisions: $Wb\bar{b} + n$ jets as a case study*, *Nucl. Phys.* **B632** (2002) 343–362 [[arXiv:hep-ph/0108069](#)].
- [635] A. Wald, *Tests of statistical hypotheses concerning several parameters when the number of observations is large*, *Transactions of the American Mathematical Society* **54** (1943) 426–482.
- [636] G. Cowan *et. al.*, *Asymptotic formulae for likelihood-based tests of new physics*, *Eur. Phys. J.* **C71** (2011) 1554 [[arXiv:1007.1727](#)].
- [637] M. Baak *et. al.*, *HistFitter software framework for statistical data analysis*, *Eur. Phys. J.* **C75** (2015) 153 [[arXiv:1410.1280](#)].
- [638] W. Verkerke and D. P. Kirkby, *The RooFit toolkit for data modeling*, *CHEP03 Proceedings* (2003) [[arXiv:physics/0306116](#)].
- [639] R. Brun and F. Rademakers, *ROOT: An object oriented data analysis framework*, *Nucl. Instrum. Meth.* **A389** (1997) 81–86.
- [640] **ROOT** Collaboration, *HistFactory: A tool for creating statistical models for use with RooFit and RooStats*, *CERN-OPEN-2012-016* (2012) [[cds.cern.ch:1456844](#)].
- [641] T. Junk, *Confidence level computation for combining searches with small statistics*, *Nucl. Instrum. Meth.* **A434** (1999) 435–443 [[arXiv:9902006](#)].
- [642] A. L. Read, *Presentation of search results: The $CL(s)$ technique*, *J. Phys.* **G28** (2002) 2693–2704. [,11(2002)].

- [643] G. Cowan *et. al.*, *Power-Constrained Limits*, [arXiv:1105.3166](#).
- [644] **ATLAS** Collaboration, *Light-quark and Gluon Jets in ATLAS: Calorimeter Response, Jet Energy Scale Systematics, and Sample Characterization*, *ATLAS-CONF-2011-053* (2011) [[cds.cern.ch:1342550](#)].
- [645] **ATLAS** Collaboration, *Jet energy measurement and systematic uncertainties using tracks for jets and for b-quark jets produced in proton-proton collisions at $\sqrt{s} = 7$ TeV in the ATLAS detector*, *ATLAS-CONF-2013-002* (2013) [[cds.cern.ch:1504739](#)].
- [646] **ATLAS** Collaboration, *Search for direct top-squark pair production in final states with two leptons in pp collisions at $\sqrt{s} = 8$ TeV with the ATLAS detector*, *JHEP* **06** (2014) 124 [[arXiv:1403.4853](#)].
- [647] **ATLAS** Collaboration, *Search for strong production of supersymmetric particles in final states with missing transverse momentum and at least three b-jets at $\sqrt{s} = 8$ TeV proton-proton collisions with the ATLAS detector*, [arXiv:1407.0600](#).
- [648] **ATLAS** Collaboration, *Search for supersymmetry at $\sqrt{s} = 8$ TeV in final states with jets and two same-sign leptons or three leptons with the ATLAS detector*, *JHEP* **1406** (2014) 035 [[arXiv:1404.2500](#)].
- [649] **ATLAS** Collaboration, *Search for squarks and gluinos with the ATLAS detector in final states with jets and missing transverse momentum using $\sqrt{s} = 8$ TeV proton-proton collision data*, *JHEP* **09** (2014) 176 [[arXiv:1405.7875](#)].
- [650] **ATLAS** Collaboration, *Search for squarks and gluinos in events with isolated leptons, jets and missing transverse momentum at $\sqrt{s} = 8$ TeV with the ATLAS detector*, *JHEP* **1504** (2015) 116 [[arXiv:1501.03555](#)].
- [651] **CMS** Collaboration, *Search for new physics in the multijet and missing transverse momentum final state in proton-proton collisions at $\sqrt{s} = 8$ TeV*, *JHEP* **1406** (2014) 055 [[arXiv:1402.4770](#)].

- [652] **CMS** Collaboration, *Search for supersymmetry in pp collisions at $\sqrt{s} = 8$ TeV in events with a single lepton, large jet multiplicity, and multiple b jets*, *Phys. Lett.* **B733** (2014) 328–353 [[arXiv:1311.4937](#)].
- [653] **CMS** Collaboration, *Search for new physics in events with same-sign dileptons and jets in pp collisions at $\sqrt{s} = 8$ TeV*, *JHEP* **1401** (2014) 163 [[arXiv:1311.6736](#)].
- [654] **CMS** Collaboration, *Exclusion limits on gluino and top-squark pair production in natural SUSY scenarios with inclusive razor and exclusive single-lepton searches at 8 TeV.*, *CMS-PAS-SUS-14-011* (2014) [[cds.cern.ch:1745586](#)].
- [655] **CMS** Collaboration, *Search for supersymmetry in pp collisions at $\sqrt{s} = 8$ TeV in events with two opposite sign leptons, large number of jets, b -tagged jets, and large missing transverse energy.*, *CMS-PAS-SUS-13-016* (2013) [[cds.cern.ch:1623818](#)].
- [656] **CMS** Collaboration, *Search for supersymmetry in pp collisions at $\sqrt{s} = 8$ TeV in events with three leptons and at least one b -tagged jet*, *CMS-PAS-SUS-13-008* (2013) [[cds.cern.ch:1547560](#)].
- [657] **DELPHES 3** Collaboration, *DELPHES 3, A modular framework for fast simulation of a generic collider experiment*, *JHEP* **1402** (2014) 057 [[arXiv:1307.6346](#)].
- [658] **ATLAS** Collaboration, *Susy-2013-09 auxillary material*, 2014. [URL Link](#).
- [659] **ATLAS** Collaboration, *Search for new phenomena in final states with an energetic jet and large missing transverse momentum in pp collisions at $\sqrt{s} = 13$ TeV using the ATLAS detector*, [arXiv:1604.07773](#).
- [660] B. Nachman and T. Rudelius, *A Meta-analysis of the 8 TeV ATLAS and CMS SUSY Searches*, *JHEP* **02** (2015) 004 [[arXiv:1410.2270](#)].
- [661] **ATLAS** Collaboration, *Search for supersymmetry in events with large missing transverse momentum, jets, and at least one tau lepton in 20 fb^{-1} of $\sqrt{s} = 8$ TeV*, *JHEP* **1502** (2015) 018 [[arXiv:1408.0291](#)].

- TeV proton-proton collision data with the ATLAS detector*, *JHEP* **09** (2014) 103 [[arXiv:1407.0603](#)].
- [662] **ATLAS** Collaboration, *Search for the direct production of charginos, neutralinos and staus in final states with at least two hadronically decaying taus and missing transverse momentum in pp collisions at $\sqrt{s} = 8$ TeV with the ATLAS detector*, *JHEP* **10** (2014) 096 [[arXiv:1407.0350](#)].
- [663] **ATLAS** Collaboration, *Search for supersymmetry in events with four or more leptons in $\sqrt{s} = 8$ TeV pp collisions with the ATLAS detector*, *Phys. Rev. D* **90** (2014) 052001 [[arXiv:1405.5086](#)].
- [664] **ATLAS** Collaboration, *Search for direct production of charginos, neutralinos and sleptons in final states with two leptons and missing transverse momentum in pp collisions at $\sqrt{s} = 8$ TeV with the ATLAS detector*, *JHEP* **1405** (2014) 071 [[arXiv:1403.5294](#)].
- [665] **ATLAS** Collaboration, *Search for direct top squark pair production in events with a Z boson, b-jets and missing transverse momentum in $\sqrt{s} = 8$ TeV pp collisions with the ATLAS detector*, *Eur.Phys.J.* **C74** (2014) 2883 [[arXiv:1403.5222](#)].
- [666] **ATLAS** Collaboration, *Search for direct production of charginos and neutralinos in events with three leptons and missing transverse momentum in $\sqrt{s} = 8$ TeV pp collisions with the ATLAS detector*, *JHEP* **1404** (2014) 169 [[arXiv:1402.7029](#)].
- [667] **ATLAS** Collaboration, *Search for long-lived stopped R-hadrons decaying out-of-time with pp collisions using the ATLAS detector*, *Phys. Rev. D* **88** (2013) 112003 [[arXiv:1310.6584](#)].
- [668] **ATLAS** Collaboration, *Search for charginos nearly mass degenerate with the lightest neutralino based on a disappearing-track signature in pp collisions at $\sqrt{s} = 8$ TeV with the ATLAS detector*, *Phys. Rev. D* **88** (2013) 112006 [[arXiv:1310.3675](#)].

- [669] **ATLAS** Collaboration, *Search for direct third-generation squark pair production in final states with missing transverse momentum and two b -jets in $\sqrt{s} = 8$ TeV pp collisions with the ATLAS detector*, *JHEP* **1310** (2013) 189 [[arXiv:1308.2631](#)].
- [670] **CMS** Collaboration, *Searches for electroweak production of charginos, neutralinos, and sleptons decaying to leptons and W , Z , and Higgs bosons in pp collisions at 8 TeV*, *Eur. Phys. J.* **C74** (2014) 3036 [[arXiv:1405.7570](#)].
- [671] **CMS** Collaboration, *Search for top-squark pairs decaying into Higgs or Z bosons in pp collisions at $\sqrt{s} = 8$ TeV*, *Phys. Lett.* **B736** (2014) 371–397 [[arXiv:1405.3886](#)].
- [672] **CMS** Collaboration, *Search for anomalous production of events with three or more leptons in pp collisions at $\sqrt{s} = 8$ TeV*, *Phys. Rev.* **D90** (2014) 032006 [[arXiv:1404.5801](#)].
- [673] **CMS** Collaboration, *Search for top squark and higgsino production using diphoton Higgs boson decays*, *Phys. Rev. Lett.* **112** (2014) 161802 [[arXiv:1312.3310](#)].
- [674] **CMS** Collaboration, *Search for top-squark pair production in the single-lepton final state in pp collisions at $\sqrt{s} = 8$ TeV*, *Eur. Phys. J.* **C73** (2013) 2677 [[arXiv:1308.1586](#)].
- [675] **CMS** Collaboration, *Search for top squarks in R -parity-violating supersymmetry using three or more leptons and b -tagged jets*, *Phys. Rev. Lett.* **111** (2013) 221801 [[arXiv:1306.6643](#)].
- [676] **CMS** Collaboration, *Search for gluino mediated bottom- and top-squark production in multijet final states in pp collisions at 8 TeV*, *Phys. Lett.* **B725** (2013) 243–270 [[arXiv:1305.2390](#)].
- [677] **CMS** Collaboration, *Search for supersymmetry in hadronic final states with missing transverse energy using the variables α_T and b -quark multiplicity in pp collisions at $\sqrt{s} = 8$ TeV*, *Eur. Phys. J.* **C73** (2013) 2568 [[arXiv:1303.2985](#)].

- [678] CMS Collaboration, *Search for new physics in events with same-sign dileptons and b jets in pp collisions at $\sqrt{s} = 8$ TeV*, *JHEP* **1303** (2013) 037 [[arXiv:1212.6194](#)].
- [679] J. Hartung, G. Knapp and B. Sinha, *Statistical Meta-Analysis with Applications*. Wiley, 2008.
- [680] B. Nachman and T. Rudelius, *Evidence for conservatism in LHC SUSY searches*, *Eur. Phys. J. Plus* **127** (2012) 157 [[arXiv:1209.3522](#)].
- [681] D. T. Sandwell *et. al.*, *New global marine gravity model from CryoSat-2 and Jason-1 reveals buried tectonic structure*, *Science* **346** (2014) 65.
- [682] NASA, *THEMIS 100 meter global daytime mosaic*, *NASA Press Release 10-176* (2010) [[mars.asu.edu/data](#)].
- [683] ATLAS Collaboration, *Letter of Intent for the Phase-II Upgrade of the ATLAS Experiment*, *CERN-LHCC-2012-022* (2012) [[cds.cern.ch:1502664](#)].
- [684] M. Moll, *Radiation damage in silicon particle detectors: Microscopic defects and macroscopic properties*, *DESY-THESIS-1999-040* (1999) [[library.desy.de:thesis99-040](#)].
- [685] C. Jacoboni *et. al.*, *A review of some charge transport properties of silicon*, *Solid State Electronics* **20** (1977) 77–89.
- [686] I. G. Kirnas *et. al.*, *Concentration dependence of the hall factor in n-type silicon*, *Phys. Stat. Sol.* **A23** (1974) K123–K127.
- [687] V. Chiochia *et. al.*, *Simulation of Heavily Irradiated Silicon Pixel Sensors and Comparison With Test Beam Measurements*, *IEEE Trans on Nucl. Sc.* **52** (2005) 1067–1075.
- [688] V. Chiochia *et. al.*, *A double junction model of irradiated silicon pixel sensors for LHC*, *Nucl. Instr. and Meth.* (2006).

- [689] G. Alimonti *et. al.*, *A study of charge trapping in irradiated silicon with test beam data*, *ATL-INDET-2003-014* (2003) [[cds.cern.ch:685542](https://cds.cern.ch/record/685542)].
- [690] W. Shockley, *Currents to Conductors Induced by a Moving Point Charge*, *Journal of Applied Physics* **9** (1938) 635.
- [691] S. Ramo, *Currents Induced by Electron Motion*, *Proceedings of the IRE* **27** (1939) 584–585.
- [692] L. Rossi *et. al.*, *Pixel Detectors: From Fundamentals to Applications*. Springer-Verlag, 2006.
- [693] Z. Li, *Modeling and simulation of neutron induced changes and temperature annealing of N_{eff} and changes in resistivity in high resistivity silicon detectors*, *Nucl. Instr. and Meth.* **A342** (1994) 105–118.
- [694] Ch. Hu-Guo, *et. al.*, *CMOS pixel sensor development: a fast read-out architecture with integrated zero suppression*, *JINST* **4** (2009) P04012.
- [695] A. Czarnecki, J. G. Korner and J. H. Piclum, *Helicity fractions of W bosons from top quark decays at NNLO in QCD*, *Phys. Rev.* **D81** (2010) 111503 [[arXiv:1005.2625](https://arxiv.org/abs/1005.2625)].
- [696] Z. Bern *et. al.*, *Left-Handed W Bosons at the LHC*, *Phys. Rev.* **D84** (2011) 034008 [[arXiv:1103.5445](https://arxiv.org/abs/1103.5445)].
- [697] E. Cho and M. J. Cho, *Variance of Sample Variance, Section on Survey Research Methods - JSM* (2008).