

Search for Supersymmetry in Final States with 2 or 3 Leptons with Data collected by the ATLAS Experiment

A Thesis submitted for the Degree of Doctor of Philosophy



THE UNIVERSITY
of ADELAIDE

Damir Duvnjak

School of Physical Sciences

University of Adelaide

This dissertation is submitted for the degree of

Doctor of Philosophy

December 2020



Dedicated to the memory of Samuel Huser.

In life, he was my closest friend and a great inspiration for my interest in mathematics
and the sciences.

I can only hope that he would have been proud of all that I have accomplished.

Acknowledgements

I would like to begin by thanking my parents, Mirsad and Branka. Their support throughout the course of my Ph.D was indispensable and without them I doubt it would have been possible. I would especially like to thank my mom for her constant encouragement and interest in my work.

I would also like to express my gratitude to my supervisor: Dr. Paul Jackson. His patience and guidance throughout the years not only helped me develop into a competent physicist, but also into a more well-adjusted adult. I would have been less of a scientist and less of a man had I not known him.

Thirdly, I would like to thank my colleagues that I worked most closely with throughout my Ph.D: Andreas Petridis, Abhishek Sharma and Anum Qureshi. Working with you has been equal parts exhilarating, fascinating, tedious and frantic and it was the best experience that a budding physicist could have.

Beyond this, I would also like to thank the many clever and interesting people I met and worked with in CoEPP and the University of Adelaide: Martin White, Tristan, Albert, Jason and Emily. Additionally, I would like to thank Takashi, Brian, Thor, Carl and Francesca. Living with them while I was working at CERN was a blast and they were some of the best housemates I could've asked for. I would also like to extend a special thanks to Ross Young, Robert, Zeno and Tom, whose help was instrumental when I was (re-)learning theoretical physics in preparation for writing the first part of this thesis.

I would also like to extend my gratitude to the people I worked with at CERN, namely Sarah Williams, Bruce Gallop and Per Johansson. Sarah helped my colleagues and I greatly when she taught us how to use the HistFitter statistical framework, which was instrumental in completing the analysis presented in this thesis. Bruce and Per helped me greatly during my work on the SCT, a project which culminated in my receipt of authorship in the ATLAS collaboration.

I also want to thank my brother, Darian. We may have grown up to be very different people, but I couldn't imagine life without him.

Finally, I would like to mention my oldest companions: Caleb, Luke and Michael. Our fellowship is one of the few places where I've had a sense of community in my life and I'd like to thank all of you for providing me with a place to unwind over the course of my Ph.D.

Declaration

I certify that this work, to the best of my knowledge and belief, contains no material previously published or written by another person, except where due reference has been made in the text. I certify that this work contains no material which has been accepted for the award of any other degree or diploma in my name in any university or other tertiary institution. I certify that no part of this work will, in the future, be used in a submission in my name for any other degree or diploma in any university or other tertiary institution without the prior approval of the University of Adelaide and where applicable, any partner institution responsible for the joint award of this degree. I certify that this dissertation contains less than 80,000 words including footnotes but excluding appendices, the bibliography and table and figure captions. I give permission for the digital version of my thesis to be made available on the web, via the University's digital research repository, the Library Search and also through web search engines, unless permission has been granted by the University to restrict access for a period of time.

Damir Duvnjak
December 2020

Abstract

In this thesis I have presented a search for supersymmetry using data gathered by the ATLAS Detector from the Large Hadron Collider during the data-taking period from 2015 to 2016. In particular, this search looks for chargino and neutralino particles and targets several areas of phase space that haven't been excluded yet. These particles are searched for in two different final states (the 2-lepton, 2-jet + E_T^{miss} final state and the 3-lepton + E_T^{miss} final state). This analysis also uses a novel method of variable construction called the "Recursive Jigsaw Reconstruction" method. This method has the potential to increase the sensitivity of analyses to signatures of SUSY particles without the need to dramatically increase the amount of data analysed.

Table of contents

List of figures	xvii
List of tables	xxvii
Nomenclature	xxxix
Introduction	1
I Theoretical Framework	5
1 The Standard Model	7
1.1 Mathematical Formalism	9
1.1.1 Lagrangian Mechanics	9
1.1.2 Group Theory	12
1.2 Quantum Electrodynamics	13
1.3 Quantum Chromodynamics	16
1.4 Electroweak Unification	20
1.5 Electroweak Symmetry Breaking and the Higgs Mechanism	24
1.6 The Standard Model	29
2 Beyond the Standard Model	33
2.1 Limitations of the Standard Model	34
2.1.1 The Hierarchy Problem	34
2.1.2 Dark Matter	36
2.1.3 Gravity	37
2.2 Supersymmetry	38
2.2.1 The Supersymmetry Transformation	39
2.2.2 SUSY Particle States	39

2.2.3	Constructing a SUSY Lagrangian	41
2.2.4	SUSY Breaking	49
2.3	The MSSM	49
2.3.1	Particle Content and Interactions	50
2.3.2	The MSSM and Symmetry Breaking	54
2.3.3	R-Parity	55
2.4	SUSY and the Limits of the Standard Model	56
2.4.1	The Hierarchy Problem	57
2.4.2	Dark Matter Candidate	57
2.4.3	Gravity	58
II	Experimental Apparatus	59
3	CERN and the LHC	61
3.1	Large Hadron Collider	64
3.1.1	Accelerator Geography	68
3.1.2	Magnet Systems	68
3.1.3	Cryogenic System	72
3.1.4	Vacuum System	73
3.1.5	Powering	74
3.1.6	Accelerator Structure	75
3.1.7	Insertions	78
3.1.8	Accelerator Operations	79
3.1.9	LHC Timelines - History and Future Plans	82
4	The ATLAS Detector	85
4.1	Physics Motivation	86
4.2	Detector Overview	88
4.3	Inner Detector	92
4.3.1	Pixel Detector	95
4.3.2	SemiConductor Tracker	97
4.3.3	Transition Radiation Tracker	108
4.3.4	Insertable B-Layer	111
4.3.5	Inner Detector Magnet System	114
4.4	Calorimeter System	114
4.4.1	ElectroMagnetic Calorimeter	117

4.4.2	Hadronic Calorimeter	117
4.5	Muon System	118
4.5.1	Muon Chambers	118
4.5.2	Muon System Magnet System	119
4.6	Forward Detectors	120
4.7	Structure, Radiation Shielding and Control	121
4.7.1	Beam-Pipe Structure	121
4.7.2	Radiation and Shielding	122
4.7.3	Control	122
5	The ATLAS TDAQ System	125
5.1	Trigger System	128
5.2	Reconstruction	132
5.2.1	Electrons	133
5.2.2	Muons	134
5.2.3	Jets	135
5.2.4	Missing Energy	140
III	Physics Analysis	143
6	Analysis Strategy	145
6.1	The Search for Supersymmetry	145
6.2	Simplified Models	146
6.3	Data and Simulation	148
6.4	Kinematic Variables	149
6.5	Signal, Control and Validation Regions	150
6.6	Statistics and Fits	154
6.7	Unblinding and Results	156
7	Analysis Overview	159
7.1	Data	159
7.2	Monte Carlo	160
7.3	Trigger Strategy	164
7.4	Object Reconstruction and Identification	165
8	Kinematic Variable Construction and the Recursive Jigsaw Method	167
8.1	Standard Variable Construction	167

8.2	Recursive Jigsaw Method	168
8.3	Recursive Jigsaw Strategy	172
8.3.1	Recursive Jigsaw Variables	174
9	Signal Region Optimisation	179
9.1	Preselection Regions	181
9.1.1	2-Lepton Preselection Regions	181
9.1.2	3-Lepton Preselection Regions	186
9.2	Signal Regions: 2-Lepton, 2-Jet and E_T^{miss} Final States	191
9.2.1	Standard Tree Regions	191
9.2.2	Compressed Tree Regions	205
9.2.3	Yields and Significance	211
9.3	Signal Regions: 3-Lepton and E_T^{miss} Final States	214
9.3.1	Standard Tree Regions	214
9.3.2	Compressed Tree Regions	225
9.3.3	Yields and Significance	230
10	Standard Model Background Estimation	233
10.1	Fake Contribution Estimation: the Matrix Method	235
10.1.1	Real Lepton Efficiency	237
10.1.2	Fake Lepton Efficiency	239
10.2	Z+Jets Estimation: Photon Template Method	247
10.2.1	γ +Jets Event Selection	248
10.2.2	γ +Jets Event Contamination	249
10.2.3	γ +Jets Event Reweighting	250
10.2.4	γ +Jets Event Smearing	254
10.2.5	γ +Jets Event $m_{\ell\ell}$ Modelling	257
10.2.6	γ +Jets Photon Splitting	258
10.2.7	Effect of p_T reweighting	258
10.2.8	Limitations of the Photon Template	260
10.3	Control Regions: 2-Lepton, 2-Jet and E_T^{miss} Final States	265
10.3.1	Standard Tree Regions	265
10.3.2	Compressed Tree Regions	270
10.4	Control Regions: 3-Lepton and E_T^{miss} Final States	274
10.4.1	Standard Tree Regions	274
10.4.2	Compressed Tree Regions	276

11 Validation Regions	279
11.1 Validation Regions: 2-Lepton, 2-Jet and E_T^{miss} Final States	280
11.1.1 Standard Tree Regions	280
11.1.2 Compressed Tree Regions	287
11.2 Validation Regions: 3-Lepton and E_T^{miss} Final States	291
11.2.1 Standard Tree Regions	292
11.2.2 Compressed Tree Regions	294
12 Uncertainties	297
12.1 Experimental Systematic Uncertainties	297
12.2 Photon Template Systematics	299
12.3 Fake Estimation Systematics	303
12.4 Systematic Uncertainties: 2-Lepton, 2-Jet and E_T^{miss} Final States	303
12.5 Systematic Uncertainties: 3-Lepton and E_T^{Miss} Final States	304
13 Results	307
13.1 Conclusion	314
References	317
 IV Appendices	 333
Appendix A Gamma Matrices	335
Appendix B Chiral Lagrangian	337
Appendix C History of Particle Physics and CERN	343
C.1 Particle Physics in the 19 th and 20 th Centuries	343
C.2 Foundation of CERN	348
C.3 Timeline of Modern Particle Physics and CERN Milestones	349
Appendix D Jet Reconstruction with Sequential Recombination Algorithms	355
Appendix E Systematic Uncertainty Contributions by Region	361

List of figures

1.1	Standard Model Particles (a) and Interactions (b).	8
1.2	Feynman diagrams of the fundamental interactions of the Standard Model.	32
3.1	The CERN logo. Image retrieved from Ref. [30].	61
3.2	An image showing the locations of major sites, detectors and accelerators overlaid on an aerial photo showing the area around Geneva. Photo taken from the CERN 60 Year Anniversary Celebration [32]. The RED circle represents the Proton Synchrotron, the BLUE circle represents the Super Proton Synchrotron and the YELLOW circle represents the Large Hadron Collider.	62
3.3	An aerial photo taken of the main site near Meyrin. Image retrieved from Ref. [33].	63
3.4	A photo of the LHC in the tunnel.	64
3.5	Schematic diagram showing how the various accelerators at CERN connect to each other. Image retrieved from Ref. [44].	67
3.6	Schematic diagram of a subdivision of the LHC. Image retrieved from Ref. [37].	69
3.7	Distribution of the conductor in the main dipole. This distribution is only of the top-right section, and is mirrored in both the horizontal and vertical axes. The x -axis is in millimetres. Image retrieved from Ref. [37].	70
3.8	Main dipole magnetic flux diagram. Image retrieved from Ref. [37].	71
3.9	Quadrupole magnet fields and layouts	72
3.10	Cross-sections of the cryodipole (a) and cryoquadrupole (b). Images retrieved from Ref. [37].	77
3.11	Schematic of a “FODO Cell” showing the layout of the LHC magnets. Image retrieved from Ref. [37].	78

3.12	Schematic diagram showing the layout of the LHC, numbering the Octants. Beam 1 circulates clockwise and Beam 2 counter-clockwise. Image retrieved from [37]	80
3.13	A histogram showing the production cross-sections of every possible product of the proton-proton collisions that occur in the LHC. Image retrieved from Ref. [51].	82
3.14	Damage to the Beam Pipe from the Magnet Quench.	84
4.1	A cut-away diagram of the ATLAS detector. A pair of humans can be seen standing on the left-hand side for scale. Image retrieved from Ref. [57].	86
4.2	A diagram of a cross-section of the ATLAS detector showing how the different particles interact with different sections of the detector. Image retrieved from Ref. [60].	91
4.3	Schematic diagram showing the ATLAS coordinate system relative to the rest of the LHC. A “top-down” view is shown. Image retrieved from Ref. [61].	92
4.4	Cut-away diagram of the Inner Detector, showing the Pixel detector, SCT and TRT.	93
4.5	Schematic diagrams of the Inner Detector.	94
4.6	Schematic view of a barrel pixel module (top) and a photo of an assembled pixel module (bottom).	96
4.7	A schematic view of the SCT module layout in the barrel (a) and end-cap (b) regions. Both are viewed along the z -axis of the experiment.	98
4.8	A photo of an SCT barrel module. Image retrieved from Ref. [64].	101
4.9	The barrel hybrid, showing the 12 ASIC Chips, the flexible centre section and the off-module connector on the left-hand end. Image retrieved from Ref. [62].	101
4.10	Photo documentation of the SCT subdetector	102
4.11	Components of the ASIC chip. Image retrieved from Ref. [62].	103
4.12	TrimRangeTest Output - Before Masking Channels.	107
4.13	TrimRangeTest Output - After Masking Channels.	109
4.14	A photo of the TRT barrel during integration. The shapes of one inner module, one middle module and one outer module are highlighted in red.	110
4.15	A Cut-away diagram of the ATLAS Calorimeter System	115
4.16	A cut-away diagram of the ATLAS Muon System	119

5.1	A screenshot taken of an oscilloscope screen (left) and a photo of that oscilloscope's controls (right).	128
5.2	Schematic diagram of the ATLAS Trigger and Data Acquisition System used during Run 2.	131
5.3	The transformed likelihood discriminant d'_L for electron candidates with $E_T \in [30, 35]$ GeV and $ \eta < 0.6$. The black histogram is for prompt electrons in a $Z \rightarrow ee$ simulation sample, and the red histogram is for backgrounds in a generic two-to-two process simulation sample. The histograms are both normalised to unit area.	135
5.4	Schema of the calorimeter isolation methodology. The grid represents the second layer of calorimeter cells in $\eta \times \phi$ space. The purple circle represents the cone of isolation with radius $\Delta R = \sqrt{(\Delta\eta)^2 + (\Delta\phi)^2}$. The topological clusters whose centres fall within the cone of isolation are coloured in magenta.	136
5.5	Schematic diagram of jet production and measurement. Figure retrieved from Ref. [76].	137
5.6	A plot showing the results of the anti- k_t jet reconstruction algorithm. The x, y -plane shows the distribution in rapidity-azimuthal angle space of cells in one layer of the calorimeter and the z -axis shows the energy in each cell. The cells are coloured based on which jet they are reconstructed as. Image retrieved from Ref. [80].	139
6.1	Feynman diagrams for the physics scenarios studied in this paper: (a) $\tilde{\chi}_1^\pm \tilde{\chi}_2^0$ decay via leptonically decaying $W^\pm$ and Z^0 bosons, (b) $\tilde{\chi}_1^\pm \tilde{\chi}_2^0$ decay via a hadronically decaying $W^\pm$ and a leptonically decaying Z^0 , (c) $\tilde{\chi}_1^\pm \tilde{\chi}_2^0$ decaying via leptonically decaying $W^\pm$ and Z^0 bosons produced in association with an initial state radiation jet, (d) $\tilde{\chi}_1^\pm \tilde{\chi}_2^0$ decay via a hadronically decaying $W^\pm$ and a leptonically decaying Z^0 produced in association with an initial state radiation jet.	147
6.2	A plot of the masses of the $\tilde{\chi}_1^\pm/\tilde{\chi}_2^0$ and the $\tilde{\chi}_1^0$ with the points that correspond to the chosen masses in the signal samples.	152
6.3	A schematic view of the analysis strategy, showing multiple control, validation and signal regions. The dashed lines represent the different bins in the regions that have them.	155
6.4	Schematic view of the HistFitter analysis flow.	158

7.1	Plots showing the cumulative integrated luminosity delivered by the LHC (green) and recorded by ATLAS (yellow) during 2015 (a) and 2016 (b). Plots retrieved from Ref. [67].	160
7.2	Plot showing the delivered integrated luminosity in terms of the pile-up of the collisions they came from. Plot retrieved from Ref. [67].	160
7.3	Schematic diagram of a proton-proton collision, highlighting how the Monte Carlo simulates the various interactions. Image retrieved from Ref. [92].	161
8.1	A general decay tree for reconstructing Recursive Jigsaw quantities. . .	170
8.2	Recursive Jigsaw decay trees: (a) corresponds to the physics scenario given in Figure 6.1a and (b) corresponds to the physics scenario given in Figure 6.1b.	172
8.3	The ISR decay tree for reconstructing Recursive Jigsaw quantities. . . .	173
9.1	Schematic diagram of the area the various signal regions target in the two-dimensional parameter space: $[m_{\tilde{\chi}_1^\pm/\tilde{\chi}_2^0}, m_{\tilde{\chi}_1^0}]$	180
9.2	Distributions of the 2-lepton standard tree variables in the standard 2-lepton preselection region.	183
9.3	Distributions of the 2-lepton compressed tree variables in the compressed 2-lepton preselection region.	185
9.4	Distributions of the 3-lepton standard tree variables in the standard 3-lepton preselection region.	188
9.5	Distributions of the 3-lepton compressed tree variables in the compressed 3-lepton preselection region.	190
9.6	The SR2 ℓ _Low $m_{\ell\ell}$ distributions.	193
9.7	The SR2 ℓ _Low m_{jj} distributions.	194
9.8	The SR2 ℓ _Low $H_{4,1}^{\text{PP}}$ distributions.	195
9.9	The SR2 ℓ _Low $\frac{p_{T,\text{PP}}^{\text{Lab}}}{p_{T,\text{PP}}^{\text{Lab}} + H_{4,1}^{\text{PP}}}$ distributions.	195
9.10	The SR2 ℓ _Low $\frac{H_{1,1}^{\text{PP}}}{H_{4,1}^{\text{PP}}}$ distributions.	196
9.11	The SR2 ℓ _Low $\min(\Delta\phi_{j_i, E_T^{\text{miss}}})$ Distribution	197
9.12	The SR2 ℓ _Intermediate $m_{\ell\ell}$ distributions.	198
9.13	The SR2 ℓ _Intermediate m_{jj} distributions.	198
9.14	The SR2 ℓ _Intermediate $H_{4,1}^{\text{PP}}$ distributions.	199
9.15	The SR2 ℓ _Intermediate $\frac{p_{T,\text{PP}}^{\text{Lab}}}{p_{T,\text{PP}}^{\text{Lab}} + H_{4,1}^{\text{PP}}}$ distributions.	200
9.16	The SR2 ℓ _Intermediate $\frac{\min(H_{1,1}^{\text{Pa}}, H_{1,1}^{\text{Pb}})}{\min(H_{2,1}^{\text{Pa}}, H_{2,1}^{\text{Pb}})}$ distributions.	200

9.17	The SR2 ℓ _Intermediate $\Delta\phi_V^P$ distribution.	201
9.18	The SR2 ℓ _High $m_{\ell\ell}$ distributions.	202
9.19	The SR2 ℓ _High m_{jj} distributions.	202
9.20	The SR2 ℓ _High $H_{4,1}^{PP}$ distributions.	203
9.21	The SR2 ℓ _High $\frac{p_{T,PP}^{Lab}}{p_{T,PP}^{Lab}+H_{4,1}^{PP}}$ distributions.	203
9.22	The SR2 ℓ _High $\frac{\min(H_{1,1}^{Pa}, H_{1,1}^{Pb})}{\min(H_{2,1}^{Pa}, H_{2,1}^{Pb})}$ distributions.	204
9.23	The SR2 ℓ _High $\Delta\phi_V^P$ distribution.	204
9.24	The SR2 ℓ _Compressed M_Z distributions.	206
9.25	The SR2 ℓ _Compressed m_J distributions.	206
9.26	The SR2 ℓ _Compressed $\Delta\phi_{ISR,I}^{CM}$ distributions.	207
9.27	The SR2 ℓ _Compressed R_{ISR} distributions.	207
9.28	The SR2 ℓ _Compressed $p_{T,ISR}^{CM}$ distributions.	208
9.29	The SR2 ℓ _Compressed $p_{T,I}^{CM}$ distributions.	209
9.30	The SR2 ℓ _Compressed p^{CM}_T distributions.	209
9.31	The SR2 ℓ _Compressed n_{jets}^S distributions.	210
9.32	The SR2 ℓ _Compressed n_{jets}^{ISR} distributions.	210
9.33	Sensitivity Plots for the 2-lepton high- and intermediate-mass signal regions.	212
9.34	Sensitivity Plots for the 2-lepton low-mass region, compressed region and the low-mass + compressed statistical combination.	213
9.35	The SR3 ℓ _Low $m_{\ell\ell}$ distributions.	215
9.36	The SR3 ℓ _Low m_T^W distributions.	215
9.37	The SR3 ℓ _Low $H_{3,1}^{PP}$ distributions.	216
9.38	The SR3 ℓ _Low $\frac{p_{T,PP}^{Lab}}{p_{T,PP}^{Lab}+H_{3,1}^{PP}}$ distributions.	217
9.39	The SR3 ℓ _Low $\frac{H_{T,3,1}^{PP}}{H_{3,1}^{PP}}$ distributions.	217
9.40	The SR3 ℓ _Intermediate $m_{\ell\ell}$ distributions.	218
9.41	The SR3 ℓ _Intermediate m_T^W distributions.	218
9.42	The SR3 ℓ _Intermediate $H_{3,1}^{PP}$ distributions.	219
9.43	The SR3 ℓ _Intermediate $\frac{p_{T,PP}^{Lab}}{p_{T,PP}^{Lab}+H_{3,1}^{PP}}$ distributions.	220
9.44	The SR3 ℓ _Intermediate $\frac{H_{T,3,1}^{PP}}{H_{3,1}^{PP}}$ distributions.	220
9.45	The SR3 ℓ _Intermediate $\frac{H_{1,1}^{Pb}}{H_{2,1}^{Pb}}$ distribution.	221
9.46	The SR3 ℓ _High $m_{\ell\ell}$ distributions.	222
9.47	The SR3 ℓ _High m_T^W distributions.	222
9.48	The SR3 ℓ _High $H_{3,1}^{PP}$ distributions.	223

9.49	The SR3 ℓ _High $\frac{p_{T,PP}^{\text{Lab}}}{p_{T,PP}^{\text{Lab}} + H_{3,1}^{\text{PP}}}$ distributions.	223
9.50	The SR3 ℓ _High $\frac{H_{T,3,1}^{\text{PP}}}{H_{3,1}^{\text{PP}}}$ distributions.	224
9.51	The SR3 ℓ _High $\frac{H_{1,1}^{\text{Pb}}}{H_{2,1}^{\text{Pb}}}$ distribution.	224
9.52	The SR3 ℓ _ISR $m_{\ell\ell}$ distributions.	226
9.53	The SR3 ℓ _ISR m_T^W distributions.	226
9.54	The SR3 ℓ _ISR $\Delta\phi_{\text{ISR,I}}^{\text{CM}}$ distributions.	227
9.55	The SR3 ℓ _ISR R_{ISR} distributions.	227
9.56	The SR3 ℓ _ISR $p_{T,\text{ISR}}^{\text{CM}}$ distributions.	228
9.57	The SR3 ℓ _ISR $p_{T,\text{I}}^{\text{CM}}$ distributions.	229
9.58	The SR3 ℓ _ISR $p_{T,\text{CM}}$ distribution.	229
9.59	Sensitivity Plots for the 3-lepton high- and intermediate-mass signal regions.	231
9.60	Sensitivity Plots for the 3-lepton low-mass region, compressed region and the low-mass + compressed statistical combination.	232
10.1	The real efficiencies for electrons and muons as a function of p_T and η . The black points represent the 2015 data while the red points represent the 2016 data.	237
10.2	2-D Plots of the real-efficiency in terms of p_T and η . The top plots shows the real-efficiency for the electrons while the bottom row shows it for the muons. The left-hand column shows 2015 data while the right-hand column shows the 2016 data.	238
10.3	The relative fake contribution in terms of the category of process that produced the fake lepton and the Monte Carlo background that contained the fake lepton for the dielectron channel.	241
10.4	The relative fake contribution in terms of the category of process that produced the fake lepton and the Monte Carlo background that contained the fake lepton for the electron-muon channel.	242
10.5	The relative fake contribution in terms of the category of process that produced the fake lepton and the Monte Carlo background that contained the fake lepton for the dimuon channel.	243
10.6	The heavy flavour fake rates for electrons and muons used in this analysis.	245
10.7	The photon conversion fake rates for electrons.	246

10.8	The light flavour fake rates for electrons (a) and muons (b) from Monte Carlo (filled coloured circles, open triangle) and data (open blue circles) compared with the heavy flavour fake rate from data in the heavy flavour control region (filled black circles).	246
10.9	Corrected light flavour fake rate for electrons.	246
10.10(a)	The γ +Jets data sample compared with the combined γ +Jets and $V\gamma$ Monte Carlo in the $V\gamma$ control region defined in Table 10.6. (b) The γ +Jets data sample compared with the combined γ +Jets and $V\gamma$ Monte Carlo. It's clear that there is a $V\gamma$ contamination that must be removed.	250
10.11(10.11a)	A Monte Carlo-to-Monte Carlo comparison of the $p_T^{Z^0}$ and p_T^γ distributions before reweighting. The distributions have different shapes and the normalisation is incorrect by a factor of 100, since the γ +Jets cross-section is much larger than the Z+Jets cross-section. (10.11b) A Monte Carlo-to-Monte Carlo comparison of the $p_T^{Z^0}$ and p_T^γ distributions after reweighting. Here the normalisation and kinematics have been properly corrected.	252
10.12	The γ +Jets and Z+Jets distributions of the n_{jets} (left) and H_T (right) variables after the reweighting has been applied. The correction in boson transverse momentum also corrects hadronic kinematics.	253
10.13	An MC-to-MC comparison showing the E_T^{miss} distributions of the γ +Jets and Z+Jets backgrounds, made after the reweighting procedure was done.	254
10.14	An Monte Carlo-to-Monte Carlo comparison showing the $E_{T,\perp}^{\text{Miss}}$ and $E_{T,\parallel}^{\text{Miss}}$ distributions of the γ +Jets and Z+Jets backgrounds in the $\mu\mu$ -channel, made after the reweighting procedure was done.	255
10.15	A data-to-data comparison showing the $E_{T,\perp}^{\text{Miss}}$ and $E_{T,\parallel}^{\text{Miss}}$ distributions of the γ +Jets and Z+Jets backgrounds in the $\mu\mu$ -channel, made after the reweighting procedure was done.	255
10.16	The $E_{T,\parallel}^{\text{Miss}}$ distributions of the Z+Jets and γ +Jets backgrounds in the $\mu\mu$ -channel after both the reweighting and smearing procedures have been done. The two backgrounds have a much better agreement than before the smearing procedure.	256
10.17	The E_T^{Miss} distributions of the Z+Jets and γ +Jets backgrounds in the $\mu\mu$ -channel after both the reweighting and smearing procedures have been done. The two backgrounds have a much better agreement than before the smearing procedure.	257

10.18	The $m_{\ell\ell}$ distributions using Monte Carlo (a) and data (b) comparing the Z+Jets and γ +Jets events.	258
10.19	Plots showing leading lepton transverse momentum distributions using Monte Carlo and data for Z+Jets and γ +Jets events.	259
10.20	Plots showing subleading lepton transverse momentum distributions using Monte Carlo and data for Z+Jets and γ +Jets events.	259
10.21	VR2 ℓ _Low-Zjets Region Plots. Left-hand plots are without reweighting, right-hand plots are with reweighting.	261
10.22	VR2 ℓ _High-Zjets Region Plots. Left-hand plots are without reweighting, right-hand plots are with reweighting.	262
10.23	VR2 ℓ _ISR-Zjets Region Plots. Left-hand plots are without reweighting, right-hand plots are with reweighting.	263
10.24	Schematic diagram of the ABCD Method.	264
10.25	Plots of $p_T^{\ell\ell}$ in terms of m_{jj} and $\frac{p_T^{PP}}{p_T^{PP}+H_{T,4,1}^{PP}}$	264
10.26	Ratio R in term of $\frac{p_T^{PP}}{p_T^{PP}+H_{T,4,1}^{PP}}$	264
10.27	CR2 ℓ _VV - Kinematic Variable Distributions.	267
10.28	CR2 ℓ _Top - Kinematic Variable Distributions.	268
10.29	CR2 ℓ _ISR_VV - Kinematic Variable Distributions.	271
10.30	CR2 ℓ _ISR_Top - Kinematic Variable Distributions.	272
10.31	CR3 ℓ _VV - Kinematic Variable Distributions	275
10.32	CR3 ℓ _ISR_VV	277
11.1	VR2 ℓ _VV - Kinematic Variable Distributions.	282
11.2	VR2 ℓ _Top - Kinematic Variable Distributions.	283
11.3	VR2 ℓ _Z+Jets_Low - Kinematic Variable Distributions.	284
11.4	VR2 ℓ _Z+Jets_High - Kinematic Variable Distributions.	285
11.5	VR2 ℓ _ISR_VV - Kinematic Variable Distributions.	288
11.6	VR2 ℓ _ISR_Top - Kinematic Variable Distributions.	289
11.7	VR2 ℓ _ISR_Z+Jets - Kinematic Variable Distributions.	290
11.8	VR3 ℓ _VV	293
11.9	VR3 ℓ _ISR_VV	295
12.1	“N-1” plots of the m_{jj} distributions in SR2 ℓ _Low and the side-band region ($\frac{p_{T,PP}^{\text{Lab}}}{p_{T,PP}^{\text{Lab}}+H_{4,1}^{\text{PF}}} \in (0.05, 0.15)$).	300
12.2	“N-1” plots of the m_J distributions in SR2 ℓ _ISR and the side-band region ($p_{T,\text{ISR}}^{\text{CM}}$ cut removed and $p_{T,I}^{\text{CM}} \in (40, 80)$ GeV).	301

- 13.1 The observed and expected SM background yields in the CRs, VRs and SRs considered in the 2ℓ channel. The statistical uncertainties in the background prediction are included in the uncertainty band, as well as the experimental and theoretical uncertainties. The bottom panel shows the difference in standard deviations between the observed and expected yields. 308
- 13.2 The observed and expected SM background yields in the CRs, VRs and SRs considered in the 3ℓ channel. The statistical uncertainties in the background prediction are included in the uncertainty band, as well as the experimental and theoretical uncertainties. The bottom panel shows the difference in standard deviations between the observed and expected yields. 309
- 13.3 Distributions of selected kinematic variables in the signal regions of the 2ℓ final state. The histograms show the post-fit background predictions. The right-most bin in each histogram includes the overflow. For SR2 ℓ _Low, the distributions for $H_{4,1}^{\text{PP}}$ (a) and $\frac{\min(H_{1,1}^{\text{Pa}}, H_{1,1}^{\text{Pb}})}{\min(H_{2,1}^{\text{Pa}}, H_{2,1}^{\text{Pb}})}$ (b) are given. For SR2 ℓ _ISR, the distributions for $p_{T,\text{ISR}}^{\text{CM}}$ (c) and R_{ISR} (d) are given. The black, hatched error bars indicate the combined uncertainty: Monte Carlo statistical, theoretical systematic and experimental systematic. The expected distribution for a benchmark signal model, normalised to the NLO+NLL cross-section times integrated luminosity, is also shown for comparison. 311
- 13.4 Distributions of selected kinematic variables in the signal regions of the 3ℓ final state. The histograms show the post-fit background predictions. The right-most bin in each histogram includes the overflow. For SR3 ℓ _Low, the distributions for $p_T^{\ell_1}$ (a) and $H_{3,1}^{\text{PP}}$ (b) are given. For SR3 ℓ _ISR, the distributions for $p_{T,\text{ISR}}^{\text{CM}}$ (c) and R_{ISR} (d) are given. The black, hatched error bars indicate the combined uncertainty: Monte Carlo statistical, theoretical systematic and experimental systematic. The expected distribution for a benchmark signal model, normalised to the NLO+NLL cross-section times integrated luminosity, is also shown for comparison. 312

- 13.5 The “N-1” m_T^W distributions for SR3 ℓ _Low (a) and SR3 ℓ _ISR. The solid red line and arrow shows the cut on m_T^W in each of these regions (the arrow points in the direction of what is *kept* by the cut). The right-most bin in each histogram includes the overflow. The black, hatched error bars indicate the combined uncertainty: Monte Carlo statistical uncertainty, theoretical systematic uncertainties on the VV background and the experimental systematic uncertainty. The expected distribution for a benchmark signal model, normalised to the NLO+NLL cross-section times integrated luminosity, is also shown for comparison. 313
- 13.6 Exclusion limits at the 95% CL on the masses of the $\tilde{\chi}_0^\pm/\tilde{\chi}_2^0$ and $\tilde{\chi}_1^0$ from the analysis of 36.1 fb $^{-1}$ of 13 TeV proton-proton collision data obtained by the ATLAS detector. The contours given in (a) and (b) correspond to the 2 ℓ and 3 ℓ channels, respectively, while the contour given in (c) shows a statistical combination of the 2 ℓ and 3 ℓ channels (assuming a 100% branching ratio of the sparticles decaying into the Standard Model $W^\pm/Z^0$ bosons and the $\tilde{\chi}_1^0$). The dashed line and the shaded band are the expected limit and its $\pm 1\sigma$ uncertainty, respectively. The thick solid line is the observed limit for the central value of the signal cross-section. The dotted lines around the observed limit illustrate the change in the observed limit as the nominal signal cross-section is scaled up and down by the theoretical uncertainty. Finally, (d) compares the exclusion limits from the analysis presented in this thesis and [87]. 315
- C.1 The Crookes tube apparatus. The cathode is on the right end of the tube, with the anode at the base of the tube. When the tube is turned on, the green glow caused by the cathode rays is clear, as is the “shadow” cast by the iron cross. 345
- D.1 A pair of 2-D plots comparing the way that the k_t and anti- k_t algorithms reconstruct the same event. The x, y -plane shows the distribution in rapidity-azimuthal angle space of cells in one layer of the calorimeter and the z -axis shows the energy in each cell. The cells are coloured based on which jet they are reconstructed as. 359

List of tables

2.1	Particle content of the Minimal Supersymmetric Standard Model . . .	53
2.2	Particle content of the MSSM pre- and post-EWSB and the mixing that it induces. Mixing between in the first two generations of sleptons and squarks is assumed to be negligible.	55
4.1	Parameters of the Inner Detector.	95
4.2	Silicon Strip Geometries	99
4.3	The configuration parameters for the modules. The upper table enumerates the values in the chip-level configuration register. The lower table describes other registers and toggles that configure various parameters at the module, chip and channel levels.	105
4.4	Parameters for the TRT barrel and end-cap modules (for one end-cap wheel). Quantities in bold are global quantities, including services and electronics. All other quantities are for individual modules and the active region.	111
4.5	Main parameters of the IBL.	114
4.6	Parameters of the Calorimeter System [59].	116
4.7	Main parameters of the Muon System. The numbers in the brackets for the MDTs and RPCs refers to the final configuration of the detector in 2009. The innermost layer of the MDT modules has an η coverage of $ \eta < 2.0$ and the TGC has special triggering modules that cover $1.05 < \eta < 2.4$	120
7.1	The Standard Model background and SUSY signal Monte Carlo samples used in this analysis.	164

7.2	Summary of the trigger strategy. The dielectron channel requires two electrons, the 2015 run required them to each have $p_T > 12$ GeV while the 2016 required them to each have $p_T > 17$ GeV. The dimuon channel required both muons to have $p_T > 18$ GeV in 2015 and $p_T > 22$ GeV in 2016. Additionally, the triggers for both years have the “mu8noL1” tag, which is a special trigger for events that pass the L1 trigger with only one muon. Finally, the electron-muon channel requires $p_{T,e} > 17$ GeV and $p_{T,\mu} > 14$ GeV during both years of data-taking.	164
9.1	2-lepton preselection region definitions. The f subscripts in the “Lepton Selection” row simply refer to the leptons having the same flavour, whatever that flavour may be.	182
9.2	3-lepton preselection regions definitions. The flavours f_1 and f_2 are the flavours of the leptons that come from the Z^0 and $W^\pm$ decays, respectively.	187
9.3	2-lepton signal region definitions - standard tree	192
9.4	2-lepton signal region definitions - compressed tree. Note that the cuts on n_{jets}^S and $n_{\text{jets}}^{\text{ISR}}$ occur at the end of the cutflow, and are not part of the preselection requirements despite technically being cuts on an amount of jets.	205
9.5	Expected Standard Model yields in the 2-lepton signal regions. Only statistical uncertainties are considered.	211
9.6	3-lepton signal region definitions - standard tree	214
9.7	3-lepton signal region definitions - compressed tree	225
9.8	Expected Standard Model yields in the 3-lepton signal regions. Only statistical uncertainties are considered.	230
10.1	Lepton classification system of the MC truth classifier. The fake lepton source processes are labelled with “Origin” while the category these processes belong to are labelled with “Type”.	240
10.2	Fake rate heavy flavour control region.	244
10.3	Fake rate photon conversion control region.	245
10.4	Z+Jets Event Selection	248
10.5	γ +Jets Event Selection	248
10.6	Definition of the $V\gamma$ control region used by the photon template to estimate the $V\gamma$ contribution in the γ +Jets background.	249
10.7	The definition of the Z+Jets and γ +Jets control regions, from which the reweighting factor for the γ +Jets background is derived.	251

10.8	2-lepton control region definitions - standard tree	266
10.9	Yields in the control regions for the 2-lepton standard tree.	269
10.10	2-lepton control region definitions - compressed tree	270
10.11	Yields in the control regions for the 2-lepton compressed tree.	273
10.12	3-lepton control region definitions - standard tree	274
10.13	3-lepton control region definitions - compressed tree	276
10.14	Yields in the control regions for the 3-lepton standard and compressed trees.	278
11.1	2-lepton validation region definitions - standard tree	281
11.2	Yields in the validation regions for the 2-lepton standard tree.	286
11.3	2-lepton validation region definitions - compressed tree	287
11.4	Yields in the validation regions for the 2-lepton compressed tree.	291
11.5	3-lepton validation region definitions - standard tree	292
11.6	3-lepton validation region definitions - compressed tree	294
11.7	Yields in the validation regions for the 3-lepton standard and compressed trees.	296
12.1	Photon template systematic uncertainties for the SR2 ℓ _Intermediate and SR2 ℓ _High regions. The relative uncertainties (reweighting, smearing, Monte Carlo- and data-closure) are derived from validations regions. The absolute uncertainties ($V\gamma$ subtraction and side-band) are quoted in signal regions.	301
12.2	Photon template systematic uncertainties for the SR2 ℓ _Low region. The relative uncertainties (reweighting, smearing, Monte Carlo- and data-closure) are derived from validations regions. The absolute uncertainties ($V\gamma$ subtraction and side-band) are quoted in signal regions.	302
12.3	Photon template systematic uncertainties for the SR2 ℓ _ISR region. The relative uncertainties (reweighting, smearing, Monte Carlo- and data-closure) are derived from validations regions. The absolute uncertainties ($V\gamma$ subtraction, side-band and photon momentum upper limit) are quoted in signal regions. An additional normalisation factor is derived from the validation region with one cut relaxed ($p_{T,\text{I}}^{\text{CM}} > 90$ GeV for SR2 ℓ _ISR), to correct the total Z+Jets prediction.	302

12.4	Summary of the main systematic uncertainties and their impact (in %) on the total SM background prediction in each of the 2ℓ SRs. The total systematic uncertainty can be different from the sum in quadrature of individual sources due to the correlations between them resulting from the fit to the data.	304
12.5	Summary of the main systematic uncertainties and their impact (in %) on the total SM background prediction in each of the 3ℓ SRs. The total systematic uncertainty can be different from the sum in quadrature of individual sources due to the correlations between them resulting from the fit to the data.	305
13.1	Expected and observed yields from the background-only fit for the 2ℓ SRs. The errors shown are the statistical plus systematic uncertainties. Uncertainties in the predicted background event yields are quoted as symmetric, except where the negative error reaches down to zero predicted events, in which case the negative error is truncated.	308
13.2	Expected and observed yields from the background-only fit for the 3ℓ SRs. The errors shown are the statistical plus systematic uncertainties. Uncertainties in the predicted background event yields are quoted as symmetric, except where the negative error reaches down to zero predicted events, in which case the negative error is truncated.	309
13.3	Model-independent fit results for all signal regions. The first column shows the signal region, the second and third columns show the 95% CL upper limits on the visible cross-section ($\langle\epsilon\sigma\rangle_{\text{obs}}^{95}$) and on the number of signal events (S_{obs}^{95}). The fourth column (S_{exp}^{95}) shows the 95% CL upper limit on the number of signal events, given the expected number (and $\pm 1\sigma$ excursions of the expectation) of background events. The last column indicates the discovery p_0 -value and its associated significance (Z).	310
13.4	Breakdown of the observed and expected (in parentheses) number of events in terms of flavor composition in the SRs with an excess.	310
E.1	A summary of the systematic uncertainty contributions in the SR2 ℓ _Low region.	362
E.2	A summary of the systematic uncertainty contributions in the SR2 ℓ _Int region.	363
E.3	A summary of the systematic uncertainty contributions in the SR2 ℓ _High region.	363

E.4	A summary of the systematic uncertainty contributions in the SR2 ℓ _ISR region.	364
E.5	A summary of the systematic uncertainty contributions in the SR3 ℓ _Low region.	365
E.6	A summary of the systematic uncertainty contributions in the SR3 ℓ _Int region.	366
E.7	A summary of the systematic uncertainty contributions in the SR3 ℓ _High region.	366
E.8	A summary of the systematic uncertainty contributions in the SR3 ℓ _ISR region.	367

Introduction

The European Centre for Nuclear Research laboratory near Geneva, Switzerland operates the worlds most powerful particle accelerator: the Large Hadron Collider. With its activation in 2008 a host of experiments were ready to study the collisions it would perform. In the years since experimental physicists have made ground-breaking discoveries that have greatly expanded the understanding humankind possesses of fundamental phenomena occurring at the smallest length scales ever studied. In the coming years the data-set will monotonically increase in size, providing a myriad of opportunities for experimental physicists to search for physics beyond the Standard Model. My work as a Ph.D student was as an experimentalist on one of these searches for new physics. This thesis will document this search with special focus on the work I performed during my Ph.D.

Part I of this thesis will provide some background information on the theoretical framework used in quantum physics, by theoretical and experimental physicists alike.

Chapter 1 will provide an overview of the Standard Model of Particle Physics, the theoretical framework that describes with singular success the properties of and interactions between the fundamental particles of the universe. This chapter will introduce the mathematical formalism used by the Standard Model before showing how it is used to describe the electromagnetic, weak and strong forces as well as electroweak unification. Immediately following this is a description of electroweak symmetry breaking and the Higgs mechanism. The chapter will end with a summary of the Standard Model in its entirety.

Despite its exceptional success there are still phenomena beyond the explanatory power of the Standard Model, and it will have to be extended in order to account for them. Chapter 2 will explore some of these phenomena as motivation for introducing supersymmetry. SUSY is a symmetry between bosons and fermions and can be used to construct extensions to the Standard Model. One such extension is the “Minimal Supersymmetric Standard Model” (or MSSM) and the search presented in this thesis is for particles that are predicted by the MSSM. Chapter 2 will describe the particle

content of the MSSM before explaining how it solves some of the issues with the Standard Model.

In Part II, I will describe the experimental work I undertook during my Ph.D, exploring the experimental apparatus I both used and worked on as an experimental physicist.

The field of experimental quantum physics has undergone a substantial change over the past century. Gone are the days where a single physicist (or small group of physicists) could construct an experiment by themselves and discover something profound about the universe. As humanity has begun to explore phenomena at smaller and smaller length scales, larger and more powerful experiments have become necessary in order to push the frontier further. Thus, much of 21st Century experimental particle physics (appropriately known as “**H**igh **E**nergy **P**hysics” or “HEP”) is performed by collaborations which have hundreds, thousands or even tens-of-thousands of personnel.

One such collaboration is CERN, of which I am a part. CERN operates the Large Hadron Collider, a particle accelerator that collides protons travelling at speeds within 0.0000001% of the speed of light. These collisions produce literally every particle in the Standard Model (and hopefully some from beyond it) with each process occurring at a rate related to its cross-section, σ . Detectors are set up around the collision points of the Large Hadron Collider to study the products of collisions in attempts to discover new particles or more precisely measure known phenomena. General purpose detectors attempt to capture as many products as possible, providing the most expansive data set that many different groups make use of. Specialised detectors are more discerning, studying collisions that have a good chance of containing a specific (type of) particle they want to draw conclusions about.

The ATLAS detector is one of the general purpose detectors at CERN. It has its own collaboration within CERN, the “ATLAS Collaboration”, of which I am also a part. The general purpose detectors are inundated with collision products, many more than can be reasonably measured and most of which aren’t of particular interest to the analysis teams. As such, ATLAS has a sophisticated system of choosing which collisions to study based on the properties of the products of those collisions. This system needs to be able to decide whether or not to keep an event in the time interval between two collisions, so it is no wonder that it is a complicated mix of both hardware and software.

These topics will be described over the course of three chapters. Chapter 3 will explore CERN and the Large Hadron Collider, Chapter 4 will explore the ATLAS

detector and Chapter 5 will describe how data is read out of the ATLAS detector as well as how the detector chooses which collisions are worth studying.

Part III will document the search for new physics that formed the core of my work during my Ph.D candidature.

Given the growing complexity of both the theoretical models of particle physics and the experiments built to test them, it should come as no surprise that analyses at the frontier of physics research are challenging endeavours that often have teams of a half-dozen physicists working on them over the course of a few years. The main focus of my work as a Ph.D student was on one such analysis, working in a small team of physicists based at the University of Adelaide. Working on such an analysis meant that I often performed a wide variety of physics work (some examples are processing data events, performing statistical tests and general software development), but there were a few topics that I focused on more than others.

The analysis presented in this thesis is a search for new physics. In particular, it was a search for new particles that are predicted by the MSSM. Chapter 6 will describe the methodology that the analysis team (myself included) used to search for these particles. Chapter 7 contains a summary of the data used for this search, as well as the configuration of the detector that recorded it.

Part of the methodology used in this search is the “Recursive Jigsaw Reconstruction” method, described in Chapter 8. This is a novel method that uses boosts to different reference frames when constructing kinematic variables. This method provides a basis of variables that provides a much greater sensitivity to possible new physics processes that would otherwise be unavailable using conventional variables.

The signal, control and validation regions that were introduced in Chapter 6 are documented in Chapters 9, 10 and 11, respectively. Chapter 10 will also cover some of the difficulties with modelling certain background processes and some of the methods used to overcome these difficulties.

Chapter 12 summarises the uncertainties considered by the analysis presented in this thesis. Description of the systematic uncertainties make up the bulk of this discussion, due to the immense complexity of the experimental apparatus.

Finally, Chapter 13 concludes this thesis with a presentation of the results of this analysis.

PART I:

THEORETICAL FRAMEWORK

Mathematics, rightly viewed, possesses not only truth, but supreme beauty - a beauty cold and austere, like that of sculpture, without appeal to any part of our weaker nature, without the gorgeous trappings of painting or music, yet sublimely pure, and capable of a stern perfection such as only the greatest art can show. The true spirit of delight, the exaltation, the sense of being more than Man, which is the touchstone of highest excellence, is to be found in mathematics as surely as in poetry.

Bertrand Russell

Chapter 1

The Standard Model

The Standard Model of particle physics is a quantum field theory that describes all known fundamental constituents and how they interact via the electromagnetic, strong and weak forces[1]. It is the most successful quantum theory that humanity has created, with its theoretical predictions matching experimental measurements with impressive precision in a wide variety of phenomena. Developed over the course of the 20th century, it began as an attempt to quantise the electromagnetic field with the most complete theory being proposed by Dirac in 1927 [2]. Since then the theory has expanded, growing to include the strong and weak forces, seeing the unification of the electromagnetic and weak forces and the development of the Higgs mechanism.


All particles in the Standard Model are treated as being point-like and have an internal angular momentum quantum number called “spin”, $\hat{S}$. Particles of the Standard Model are classified into two categories based on their spin values. Particles with integer-valued spins are called “bosons” and respect Bose-Einstein statistics while particles with half-integer-valued spins are called “fermions” and respect Fermi-Dirac statistics.

Fermions are the constituents of matter and are divided into two groups: quarks and leptons, each of which are divided into three “generations”. The generations are named that because they are ordered: the particles in each generation are heavier than the ones in the previous generations while being identical in all other aspects. There are 12 fermions in total, 6 quarks and 6 leptons. Each quark generation has two quarks: an up-type quark and a down-type quark and each lepton generation has a charged lepton and a neutrino. Quarks interact via all three of the forces described by the Standard Model: strong, weak and electromagnetic. The charged leptons interact via the weak and electromagnetic forces, but not the strong. The neutrinos are neutral,

have a very small, non-zero mass [3] and only interact via the weak force. Each fermion has a corresponding antifermion with the same mass but opposite charge.

Bosons have integer spin values and are divided into two groups: gauge bosons and scalar bosons. The gauge bosons are “force-carriers”. In the Standard Model, when two particles interact via one of the fundamental forces it is understood as those two particles exchanging a virtual force-carrying particle between them. The mediator of the electromagnetic force is the photon, the weak force is mediated by the vector bosons (the W^+ , W^- and Z^0) and the strong force is mediated by the gluons (of which there are eight). The gauge bosons are all spin-1. There is only one scalar boson: the Higgs Boson, which is the particle responsible for giving masses to all of the massive Standard Model particles (including itself). The term “scalar” refers to the spin of the particle, which is zero.

Figure 1.1 shows the particles of the Standard Model and summarises their properties and the interactions they undergo.



(a) A diagram showing all the particles of the Standard Model. Each particle is shown with its mass, spin and electric charge. (b) A schematic diagram summarising the interactions between the particles of the Standard Model.

Fig. 1.1 Standard Model Particles (a) and Interactions (b).

1.1 Mathematical Formalism

As a quantum field theory the Standard Model makes use of many advanced mathematical concepts and it is useful to introduce the formalism for these before discussing the theory proper. These are Lagrangian mechanics and group theory.

1.1.1 Lagrangian Mechanics

Classical mechanics describes the motion of macroscopic objects and one of the ways it does this is through the use of “equations of motion”. These are equations that describe the motion of objects in terms of dynamical variables (the position and velocity/momentum) of particles as well as time. An early attempt to come up with a systematic way of creating equations of motion is Newtons second law of motion:

In an inertial frame of reference, the vector sum of forces acting on an object, $\vec{F}$, is equal to the product of its mass, m and the acceleration it undergoes, $\vec{a}$: $\vec{F} = m\vec{a}$

Applying this law to an object in a physical system generates three equations of motion per object, one for each of the objects cartesian coordinates. This method is somewhat cumbersome since it doesn’t provide a systematic method of dealing with constraints in the ways an object can move (for instance, a simple pendulum only has one degree of freedom).

To solve this issue, Joseph-Louis Lagrange would reformulate classical mechanics as Lagrangian mechanics. This reformulation adds no new physics, but it is much more mathematically sophisticated and systematic. For a system of N particles in 3-dimensional space, the position vector of the i^{th} particle can be written as a cartesian vector (Equation 1.1).

$$\vec{r}_i = (x_i, y_i, z_i), \forall i = 1, \dots, N \quad (1.1)$$

Meaning that there are *at most* $3N$ coordinates that define the physical system. In cases where the particles in the system are constrained in some way, Lagrangian mechanics has a systematic way of determining how many degrees of freedom the system has. A “holonomic constraint” is an equation of the form shown in Equation 1.2.

$$f(\vec{r}_j, t) = 0 \quad (1.2)$$

Note that these equations only involve coordinates of one particle. Each of these constraints removes a coordinate from the system since one coordinate can be written in terms of the other two. The goal is to minimise the amount of coordinates used,

taking full advantage of all known constraints. Therefore, in a system with N particles and C constraints, the number of degrees of freedom n is given by $n = 3N - C$. These n degrees of freedom are thus known as the “generalised coordinates” of the system, denoted $q_i(t)$, and the vector space defined by them is called the “configuration space”. It is in terms of these coordinates that the central quantity of Lagrangian mechanics, the “Lagrangian”, is defined (Equation 1.3).

$$L(\vec{q}(t), \dot{\vec{q}}(t), t) = K(\vec{q}(t), \dot{\vec{q}}(t), t) - V(\vec{q}(t), \dot{\vec{q}}(t), t) \quad (1.3)$$

where:

- $\vec{q}(t) = (q_1(t), \dots, q_n(t))$
- $\dot{\vec{q}}(t) = (\frac{dq_1}{dt}, \dots, \frac{dq_n}{dt})$
- $K(\vec{q}(t), \dot{\vec{q}}(t), t)$ is the summed kinetic energy of all particles in the system in terms of $\vec{q}$, $\dot{\vec{q}}$ and t .
- $V(\vec{q}(t), \dot{\vec{q}}(t), t)$ is the summed potential energy of all particles in the system in terms of $\vec{q}$, $\dot{\vec{q}}$ and t .

The equations of motion of the system are then simply given by the Euler-Lagrange equations (Equation 1.4).

$$\frac{d}{dt} \left(\frac{\partial L}{\partial \dot{q}_i} \right) = \frac{\partial L}{\partial q_i} \quad (1.4)$$

Whether or not the equations of motion can be solved by analytical methods depends on the system, but important things about the system can be learned without needing to solve these equations at all. To this end, the “action” is defined (Equation 1.5).

$$\mathcal{S}[\vec{q}(t)] = \int_{t_1}^{t_2} L(\vec{q}(t), \dot{\vec{q}}(t), t) dt \quad (1.5)$$

This is a “functional”, which takes as input a function and outputs a scalar. While the integral runs from times t_1 to t_2 , $\vec{q}(t_1)$ and $\vec{q}(t_2)$ are positions in configuration space. The integral can thus be understood as a path integral in configuration space from $\vec{q}(t_1)$ to $\vec{q}(t_2)$. The path taken between these two points is simply the one that minimises the action, this is called the “principle of least action”. Additionally, the requirement that the action be minimised is equivalent to the Euler-Lagrange equations holding.

If the system then undergoes some transformation (for instance, a translation through time: $t \rightarrow t' = t + a$) that leaves the action unchanged (or equivalently: the

Euler-Lagrange equations don't change their form) then the action is a “symmetry” of that transformation. Further, if the transformation is continuous (the symmetry is a “continuous symmetry”) it allows the application of Noether's theorem:

Every continuous symmetry of the action of a physical system has a corresponding conservation law.

A conservation law states that there exists some property of a physical system that doesn't change as the system evolves over time. This quantity, C , is known as the “conserved quantity” whose conservation law is rendered mathematically in Equation 1.6.

$$\frac{dC}{dt} = 0 \quad (1.6)$$

In the case of the example above, if a system is invariant under time-translations the corresponding conserved quantity is simply the total energy of the system, E . The *law of conservation of energy* is thus a consequence of a much more fundamental fact: a systems action being invariant under time-translation transformations.

Symmetries are very useful for understanding physical systems, since whenever a symmetry can be found it always means that there is a corresponding conservation law.

Everything discussed in this subsection so far has been for classical particles, but it generalises to quantum fields in a straightforward way. Instead of a particle with coordinates $\vec{r}$, the particle is represented by a generic field $\Phi(x^\mu)$ which permeates through all of spacetime. The Lagrangian is written in terms of Φ , $\partial_\mu \Phi$ and x^μ , the Minkowski spacetime coordinate. Since the theory incorporates special relativity it is more natural to use the “Lagrangian density” $\mathcal{L}$, which must be integrated over all space to find the Lagrangian (Equation 1.7).

$$L = \int \mathcal{L} d^3\mathbf{x} \quad (1.7)$$

In this formalism, the Euler-Lagrange equations take the form shown in Equation 1.8 and the action is defined in Equation 1.9.

$$\partial_\mu \left(\frac{\partial \mathcal{L}}{\partial(\partial_\mu \Phi)} \right) = \frac{\partial \mathcal{L}}{\partial \Phi} \quad (1.8)$$

$$\mathcal{S}[\Phi] = \int \mathcal{L}(\Phi(x^\mu), \partial_\mu \Phi(x^\mu), x^\mu) d^4x \quad (1.9)$$

Noether's theorem can still be applied, but instead of a conserved quantity there is now a *conserved current*, J^μ , with four components. The conservation law reflects this (Equation 1.10).

$$\partial_\mu J^\mu = 0 \quad (1.10)$$

In this formalism, the associated charge, Q , is found by integrating the zeroth component of the conserved current over all space (Equation 1.11). This charge is a conserved quantity (Equation 1.12).

$$Q = \int J^0 d^3\mathbf{x} \quad (1.11)$$

$$\frac{dQ}{dt} = 0 \quad (1.12)$$

The Standard Model is a relativistic quantum field theory that uses the field formalism defined above. Central to the Standard Model is the “gauge transformation”. These are transformations not in the spacetime coordinate x^μ but in the fields themselves, taking the form $\Phi \rightarrow \Phi' = e^{i\alpha}\Phi$. The Lagrangian of a free fermion isn't locally gauge invariant on its own, but adding terms to make it locally gauge invariant (“locally” meaning that the α in the exponent is a function of the spacetime coordinate x^μ) creates new vector fields which correspond to the gauge bosons. This process is called “gauging” and is done for each of the three fundamental forces that the Standard Model describes (see Sections 1.2, 1.3 and 1.4).

1.1.2 Group Theory

A group is a set, G , equipped with a binary operation, $\circ$, that satisfy the four group axioms:

1. **Closure:** If $a, b \in G$ then $a \circ b \in G$
2. **Associativity:** $\forall a, b, c \in G, (a \circ b) \circ c = a \circ (b \circ c)$
3. **Identity:** $\exists I \in G$ such that $a \circ I = I \circ a = a \forall a \in G$. This I is unique and called the “identity element”
4. **Inverse:** For each $a \in G, \exists b \in G$ such that $a \circ b = b \circ a = I$. For each $a \in G, b$ is unique and called the “inverse” of a .

The standard notation to represent this group is $(G, \circ)$.

Every group has a “generating set”, $S \subset G$, defined such that every element of G can be expressed as a combination (under the group operation $\circ$) of a finite number

of the elements of S and their inverses. Additionally, if a groups binary operation $\circ$ satisfies the **commutativity** property (Equation 1.13) then the group is known as an “abelian” group, otherwise it is a “non-abelian” group.

$$\forall a, b \in G, a \circ b = b \circ a \quad (1.13)$$

Quantum field theory makes extensive use of a special type of group called the “Lie group”. A “Lie group” is a group whose underlying set G is a differentiable manifold and whose operation $\circ$ is smooth. Differentiable manifolds are continuous (mathematically speaking, this means that their elements can be expressed in terms of a finite number of real parameters) and so the generating set of a Lie group must also be a differentiable manifold. Thus, when talking about the “generators” of a Lie group it isn’t often helpful to talk about the whole generating set, which is continuous and of uncountably infinite cardinality. It is more useful to refer to the generating set in terms of some vector basis of the set which spans the entire manifold.

Lie groups provide a natural model for continuous symmetries and quantum field theory makes use of two types of Lie groups: $U(n)$ and $SU(n)$. The group $U(n)$ is called the “unitary group”, and is the group of all $n \times n$ unitary matrices under matrix multiplication. The group $SU(n)$ is the “special unitary group” and is the group of all $n \times n$ unitary matrices with determinant 1 under matrix multiplication.

The gauge transformations of the Standard Model have group structure and produce continuous symmetries, with each of the three fundamental forces transforming under a different Lie group. Electromagnetism transforms under the $U(1)$ group, the weak force transforms under the $SU(2)$ group and the strong force transforms under the $SU(3)$. The elements of these groups are the gauge transformations themselves and the generators of the groups will be related to the mediator particles of the interactions.

1.2 Quantum Electrodynamics

Quantum electrodynamics (or “QED”) was the first gauge field theory to be developed and served as the prototype for later, more complex gauge field theories. It began as an attempt to quantise the electromagnetic field [4], with the first reasonably complete theory formulated by Dirac in 1927 [2]. Dirac’s theory modelled both the electron and the electromagnetic field as quantum objects and was significant because it was the first theory to be able to model the emission of a photon by an electron, a process where the number of particles changes. The completion of QED would come with Enrico Fermi’s theory of β^- -decay [5], which incorporated into QED the ability to

create and annihilate fermions. Feynman would later refer to QED as the “jewel of physics” due to its remarkable predictive power.

In QED fermion fields are represented by the four-component Dirac “spinor” (Equation 1.14).

$$\psi(x^\mu) = \begin{bmatrix} \psi_1(x^\mu) \\ \psi_2(x^\mu) \\ \psi_3(x^\mu) \\ \psi_4(x^\mu) \end{bmatrix} \quad (1.14)$$

A fermion can exist in one of two spin states, and its antifermion can also exist in one of two spin states. The four components of the Dirac spinor each represent the particle/antiparticle taking one of the two spin values. Further, the indices labelled 1 through 4 are “Dirac indices” and are a different index from the spacetime index μ .

Equation 1.15 shows the Lagrangian of a free fermionic field.

$$\mathcal{L}_{\text{Free Fermion}} = i\bar{\psi}\gamma^\mu\partial_\mu\psi - m\bar{\psi}\psi \quad (1.15)$$

This Lagrangian is invariant under global gauge transformations (Equation 1.16).

$$\psi(x^\mu) \rightarrow \psi'(x^\mu) = e^{i\alpha}\psi(x^\mu) \quad (1.16)$$

Applying a local gauge transformation (ie, making α a function of the spacetime coordinate x^μ) gives Equation 1.17.

$$\psi(x^\mu) \rightarrow \psi'(x^\mu) = e^{i\alpha(x^\mu)}\psi(x^\mu) \quad (1.17)$$

This gives the transformed Lagrangian an extra term (Equation 1.18).

$$\mathcal{L}_{\text{Free Fermion}} \rightarrow \mathcal{L}'_{\text{Free Fermion}} = i\bar{\psi}\gamma^\mu\partial_\mu\psi - m\bar{\psi}\psi + i\bar{\psi}\gamma^\mu\partial_\mu\alpha\psi \quad (1.18)$$

Gauge invariance can be enforced by adding another term to the Lagrangian (Equation 1.19).

$$\mathcal{L}_{\text{Free Fermion}} \rightarrow \mathcal{L}_{\text{Gauge Inv.}} = \mathcal{L}_{\text{Free Fermion}} - g_e\bar{\psi}\gamma^\mu A_\mu\psi \quad (1.19)$$

The new field A_μ is the gauge field, and its transformation is shown in Equation 1.20.

$$A_\mu \rightarrow A'_\mu = A_\mu - \frac{1}{g_e}\partial_\mu\alpha \quad (1.20)$$

The gauge field needs dynamics of its own, so a non-interacting Lagrangian for the field A_μ is defined (Equation 1.21).

$$\mathcal{L}_\gamma = -\frac{1}{4}F^{\mu\nu}F_{\mu\nu} + \frac{1}{2}m_A A^\mu A_\mu \quad (1.21)$$

Here, $F_{\mu\nu} = \partial_\mu A_\nu - \partial_\nu A_\mu$ is the quantised electromagnetic field strength tensor. Like its classical counterpart it describes the electromagnetic field and is itself gauge invariant. Indeed, the gauge invariance of classical electromagnetism partly inspired the idea that a good quantum field theory should be gauge invariant, and the form of the electromagnetic field tensor would lead to the strong and weak forces having field strength tensors of their own. The second term: $\frac{1}{2}m_A A^\mu A_\mu$ is problematic since it is not locally gauge invariant. To recover local gauge invariance, the gauge boson is taken to be massless ($m_A = 0$). This choice has far reaching consequences, since it implies that the gauge bosons of any gauge invariant theory must be massless.

The final QED Lagrangian is shown in Equation 1.22.

$$\mathcal{L} = i\bar{\psi}\gamma^\mu\partial_\mu\psi - m\bar{\psi}\psi - g_e\bar{\psi}\gamma^\mu A_\mu\psi - \frac{1}{4}F^{\mu\nu}F_{\mu\nu} \quad (1.22)$$

This equation is slightly messy, and can be simplified by introducing the concept of a “covariant derivative”. A covariant derivative is simply the gauge invariant version of the derivative operator. In Equation 1.22, the terms $m\bar{\psi}\psi$ and $-\frac{1}{4}F^{\mu\nu}F_{\mu\nu}$ are gauge invariant while the remaining two terms are gauge invariant when taken together. Therefore, we can use those two terms to define the covariant derivative (Equation 1.23):

$$\mathcal{D}_\mu = \partial_\mu + ig_e A_\mu \quad (1.23)$$

This covariant derivative contains both the dynamics of the fermion field and its interactions with the gauge bosons. Having this covariant derivative allows us to rewrite the QED Lagrangian in a much more concise form (Equation 1.24).

$$\mathcal{L}_{QED} = \bar{\psi}(i\gamma^\mu\mathcal{D}_\mu - m)\psi - \frac{1}{4}F^{\mu\nu}F_{\mu\nu} \quad (1.24)$$

The gauge transformation in this section is based on the abelian $U(1)$ group, the group of 1×1 unitary matrices (ie, the set of all complex numbers of magnitude 1) under matrix multiplication (ie, regular multiplication). Equations 1.16 and 1.17 show the gauge transformations as the field undergoing a complex phase rotation, or alternatively, the field is multiplied by a complex number of magnitude one. That the QED Lagrangian is invariant under these transformations means that the action is also

invariant, so by Noether’s theorem there is a conserved quantity associated with the transformation. This conserved quantity is the electric charge, Q .

QED, like all the gauge field theories of the Standard Model, is re-normalisable. This means that the infinities coming from loop Feynman diagrams that contribute to the self-energies of the fields can be removed. The physical consequence of this is that the QED coupling constant g_e and the fine structure constant $\alpha_e = \frac{g_e^2}{4\pi}$ are both functions of energy, running with the scale of the interaction.

1.3 Quantum Chromodynamics

Quantum chromodynamics is the gauge field theory that describes the strong force. It follows much of the same derivation procedure as QED, with some notable differences such as the number of fermions considered and the structure of the gauge group.

The earliest incarnation of what would eventually become quantum chromodynamics was the “eightfold way” formulated by Murray Gell-Mann in 1961 [6]. The eightfold way was made to organise particles (specifically, the mesons and spin- $\frac{1}{2}$ baryons) with respect to their charge and strangeness. In the following years the idea would be refined into the “quark model”, independently proposed by both Murray Gell-Mann [7] and George Zweig [8, 9] in 1964. The quark model proposed the idea that hadrons were composite particles, and that the various quantum numbers of the hadrons (such as “strangeness” and “charm”) were driven by the fundamental particles that made up the hadrons: the “quarks”.

The quark model would be further extended, with such an extension being simultaneously, independently proposed by Greenberg [10] and Han and Nambu [11] in 1964 and 1965, respectively. While attempting to explain the existence of baryons like the Δ^{++} (which is composed of three up quarks with parallel spins) they proposed that quarks possessed an additional degree of freedom, one that could be modelled by a gauge group. Their ideas here would evolve over time, eventually creating quantum chromodynamics in which the strong force was described as a fully realised gauge field theory.

The degree of freedom mentioned in the previous paragraph is now known as the “colour charge”. The colour charge (or just “colour”) is analogous to the electric charge of the electromagnetic force, being the “source” of the strong force. Unlike the electromagnetic charge, there are *three* possible colours: “red”, “blue” and “green”. The name “colour” was chosen as an analogy between the three “colour” charges and the three primary colours, and doesn’t have anything to do with the lay definition

of colour. Each of these colours behaves analogously to electric charge, that is, they all can be either “positive” or “negative”. Thus, for a “red” colour-charged particle, the corresponding antiparticle would have “antired” as it’s colour charge. Similarly, “blue” particles have “antiblu” antiparticles and “green” particles have “antigreen” antiparticles.

There are six quarks in total (see Figure 1.1a) and each of these can have one of three colours. Therefore, each of the six quark fields are represented by a “triplet” (see Equation 1.25). These “triplets” aren’t really vectors, but rather a way to group together things that are mixed under transformations, as will be seen in a few paragraphs time.

$$\psi = \begin{pmatrix} q_R \\ q_B \\ q_G \end{pmatrix}, \bar{\psi} = (q_R \quad \bar{q}_B \quad \bar{q}_G) \quad (1.25)$$

Here, each of q_R , q_B and q_G is a 4-component Dirac spinor like the one shown in Equation 1.14. This means that the Lagrangian of a free fermion is written in terms of this three-component vector (Equation 1.26).

$$\mathcal{L}_{\text{Free Fermion}} = i\bar{\psi}\gamma^\mu\partial_\mu\psi - m\bar{\psi}\psi \quad (1.26)$$

Here, the derivative is understood to act on *every* element of ψ (Equation 1.27).

$$\partial_\mu\psi = \begin{bmatrix} \partial_\mu q_R \\ \partial_\mu q_B \\ \partial_\mu q_G \end{bmatrix}, \partial_\mu q_i = \begin{bmatrix} \partial_\mu q_{i,1} \\ \partial_\mu q_{i,2} \\ \partial_\mu q_{i,3} \\ \partial_\mu q_{i,4} \end{bmatrix} \quad (1.27)$$

The QCD Lagrangian undergoes a similar gauging procedure to the QED Lagrangian. If the QCD fermion field undergoes the $U(1)$ local gauge transformation shown in Equation 1.17, gauging just produces the electromagnetic force for quarks. Instead, a different gauge group must be used: $SU(3)$. Instead of multiplying the fermion field by a complex phase rotation it is multiplied by a 3×3 matrix (Equation 1.28).

$$\psi(x^\mu) \rightarrow \psi'(x^\mu) = e^{i\alpha_j(x^\mu)\lambda_j}\psi(x^\mu) \quad (1.28)$$

The factor $e^{i\alpha_j(x^\mu)\lambda_j}$ is the 3×3 matrix mentioned in the previous paragraph and introduces two important concepts. First, the index “ j ” is a new index, neither a spacetime index nor a Dirac index. This is the colour index, it runs from 1 to 8 where each value has a corresponding function $\alpha_j(x^\mu)$ and matrix λ_j . The functions

$\alpha_j(x^\mu)$ are analogous to the function $\alpha(x^\mu)$ in Equation 1.17. The matrices λ_j are the Gell-Mann matrices, and are listed in Equation 1.29.

$$\begin{aligned}
\lambda^1 &= \begin{bmatrix} 0 & 1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix} & \lambda^2 &= \begin{bmatrix} 0 & -i & 0 \\ i & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix} & \lambda^3 &= \begin{bmatrix} 1 & 0 & 0 \\ 0 & -1 & 0 \\ 0 & 0 & 0 \end{bmatrix} \\
\lambda^4 &= \begin{bmatrix} 0 & 0 & 1 \\ 0 & 0 & 0 \\ 1 & 0 & 0 \end{bmatrix} & \lambda^5 &= \begin{bmatrix} 0 & 0 & -i \\ 0 & 0 & 0 \\ i & 0 & 0 \end{bmatrix} & \lambda^6 &= \begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & 1 \\ 0 & 1 & 0 \end{bmatrix} \\
\lambda^7 &= \begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & -i \\ 0 & i & 0 \end{bmatrix} & \lambda^8 &= \frac{1}{\sqrt{3}} \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & -2 \end{bmatrix}
\end{aligned} \tag{1.29}$$

The term “ $\alpha_j(x^\mu)\lambda^j$ ” is understood to be a summation over j . Additionally, taking the exponential of this term is understood to be a Taylor expansion (Equation 1.30).

$$e^{i\alpha_j(x^\mu)\lambda^j} = 1 + i \left(\sum_{j=1}^8 \alpha_j(x^\mu)\lambda^j \right) - \frac{1}{2} \left(\sum_{j=1}^8 \alpha_j(x^\mu)\lambda^j \right)^2 - \frac{i}{3} \left(\sum_{j=1}^8 \alpha_j(x^\mu)\lambda^j \right)^3 + \frac{1}{4} \left(\sum_{j=1}^8 \alpha_j(x^\mu)\lambda^j \right)^4 + \dots \tag{1.30}$$

Second, multiplying the triplet ψ by the matrix $e^{i\alpha_j(x^\mu)\lambda^j}$ mixes the components of ψ , ie, the components of ψ' are a linear combination of the components of ψ . This is what an $SU(3)$ transformation actually does.

With the above noted, the gauging procedure follows the same path as for QED. The Lagrangian is transformed under the $SU(3)$ transformation, which leads to extra terms of the form $\bar{\psi}\gamma^\mu\partial_\mu(\alpha_j\lambda^j)\psi$. These new terms are cancelled out with the addition of a gauge field, which receives dynamics of its own. Finally, the derivative is turned into a covariant derivative and the QCD Lagrangian is complete (Equation 1.31).

$$\mathcal{L}_{QCD} = \bar{\psi}(i\gamma^\mu\mathcal{D}_\mu - m)\psi - \frac{1}{4}G^{a,\mu\nu}G_{\mu\nu}^a \tag{1.31}$$

The covariant derivative and field strength tensor are defined in Equations 1.32 and 1.33, respectively.

$$\mathcal{D}_\mu = \partial_\mu + i\frac{g_s}{2}\lambda^a G_\mu^a \tag{1.32}$$

$$G_{\mu\nu}^a = \partial_\mu G_\nu^a - \partial_\nu G_\mu^a + g_s f^{abc} G_\mu^b G_\nu^c \tag{1.33}$$

The Gell-Mann matrices, λ^a are the generators of the $SU(3)$ group. The $SU(3)$ group is non-abelian, so the matrices don't commute with each other, instead obeying the relation in Equation 1.34.

$$[\lambda^a, \lambda^b] = 2if^{abc}\lambda^c \quad (1.34)$$

Each of the Gell-Mann matrices corresponds to one of eight gauge fields G_μ^a . These gauge fields are the gluons. Gluons carry colour charge (compared with the photons of electromagnetism which are electrically neutral). Each gluon has two colour charges which are different from the charges of the other seven as shown in Equation 1.35

$$\begin{array}{ll} \frac{1}{\sqrt{2}}(r\bar{b} + b\bar{r}) & -\frac{i}{\sqrt{2}}(r\bar{b} - b\bar{r}) \\ \frac{1}{\sqrt{2}}(r\bar{g} + g\bar{r}) & -\frac{i}{\sqrt{2}}(r\bar{g} - g\bar{r}) \\ \frac{1}{\sqrt{2}}(b\bar{g} + g\bar{b}) & -\frac{i}{\sqrt{2}}(b\bar{g} - g\bar{b}) \\ \frac{1}{\sqrt{2}}(r\bar{r} + b\bar{b}) & \frac{1}{\sqrt{6}}(r\bar{r} - b\bar{b} - 2g\bar{g}) \end{array} \quad (1.35)$$

The f^{abc} is the structure constant of the $SU(3)$ group and appears in Equation 1.33 as well. The non-abelian nature of $SU(3)$ and the terms involving the structure constant in particular give rise to three- and four-gluon interaction terms. Intuitively this makes sense, the gluon carries colour charge and so would interact with other gluons.

The strong force has two important properties that it doesn't share with the other two forces, they are called "confinement" and "asymptotic freedom".

Confinement states that particles can only exist as colour singlets (in other words, particles with a net colour charge can never be isolated and observed individually). The only objects that can be observed in nature are composite hadrons: baryons and mesons. Baryons are composed of three quarks, with the three quarks either being red, blue and green or antired, antiblue and antigreen. Either way, these three colours combine to give a "colourless" or "white" particle. Mesons are comprised of a quark and antiquark where one quark carries any of the three colours and the antiquark carries the corresponding anticolour, again creating a "colourless" object. The reason only colourless objects are observed in nature has to do with how the strong force scales with distance [1].

If one attempts to separate a colour-singlet state into coloured components (for instance, to separate a meson into a quark and an antiquark) the strong force grows approximately proportionally to the distance between the quarks. Eventually, the distance increases so much that it becomes energetically favourable for a new quark-antiquark pair to appear from the vacuum between the original quark-antiquark pair.

The new antiquark then combines with the old quark and vice-versa, leading to two new mesons [1].

For very high energy hadrons, this process repeats for the newly created hadrons, leading to a collimated spray of particles. A similar thing happens with quarks and gluons. A high energy quark will become a colour singlet as well as emit gluons. Gluons are technically massless, so should propagate forever (barring effects like redshift), but their strong self-interaction means that high-energy single gluons will decay into several gluons or quarks, which in turn decay themselves. Both cases lead to a collimated spray of particles in the direction the original quark/gluon was moving in.

This process, whereby a high-energy strongly-interacting object (be it a hadron, quark or gluon) decays and produces a spray of particles is known as “parton showering”. The spray of particles itself is known as a “parton shower”.

Like QED, QCD is renormalisable and so its coupling constant g_s and fine structure constant α_s both run with the scale of the interaction. However, the QCD β -function (which describes how α_s runs) is negative and decreases logarithmically with increasing energy. In practical terms, this means that the collisions that occur in modern colliders like the LHC are high enough energy that the quarks and gluons in colliding protons can be thought of as interacting directly. This is known as “asymptotic freedom”.

1.4 Electroweak Unification

The weak force (sometimes called “quantum flavourdynamics”) is the third force described by the Standard Model. It is the force that describes radioactive decay, which makes it essential for understanding processes such as nuclear fission.

The first attempt to study a process mediated by the weak force was Fermi’s theory of β -decay, proposed by Enrico Fermi in 1933 [5]. Fermi was attempting to explain the process of β -Decay (in which a nucleon emits a β -particle, shown in Equation 1.36) by modelling it as a four-particle contact interaction (the Hamiltonian approximating it is shown in Equation 1.37).

$$n^0(p^+) \rightarrow p^+(n^0) + e^-(e^+) + \bar{\nu}_e(\nu_e) \quad (1.36)$$

$$H = \frac{G_F}{\sqrt{2}} (\bar{p}\gamma^\mu n)(\bar{e}\gamma^\mu \nu_e) \quad (1.37)$$

Fermis theory would ultimately prove untenable because its calculated cross-section grows as the square of the energy ($\sigma \propto E^2$). Since the cross-section grows without bound the theory isn't valid at high energies.

Unlike QED and QCD, a theory of the weak force would never be designed “on its own”. In principle, such a theory would have the weak force transform under the $SU(2)$ gauge group, and gauging this would give three gauge bosons and a conserved flavour charge. However, such a theory would remain unexplored. Instead, the electromagnetic and weak forces would be unified into one, the “electroweak” force. The electromagnetic and weak forces would then be understood as two components of this single force. This unification would be pioneered by Sheldon Glashow [12], Abdus Salam [13] and Steven Weinberg [14], winning them the 1979 Nobel Prize in Physics [15].

The electroweak force acts on both leptons and quarks, but how it interacts with them is dependent on their chirality. Every fermion field can be written as a sum of its chiral-left and chiral-right components (Equation 1.38) and each of these components can be found using the chiral projection operators defined in Equation 1.39.

$$\psi(x^\mu) = \Psi_L(x^\mu) + \psi_R(x^\mu) \quad (1.38)$$

$$\hat{P}_L = \frac{1-\gamma^5}{2} \quad , \quad \hat{P}_R = \frac{1+\gamma^5}{2} \quad (1.39)$$

The free fermion Lagrangian is then written in terms of the chiral-left and -right fields (see Equation 1.40).

$$\mathcal{L}_{\text{Free Fermion}} = i\bar{\Psi}_L \gamma^\mu \mathcal{D}_\mu \Psi_L + i\bar{\psi}_R \gamma^\mu \mathcal{D}_\mu \psi_R \quad (1.40)$$

Compared to earlier free fermion Lagrangians there is a couple of immediately noticeable differences: the lack of left-right coupled vector terms and the lack of any scalar/mass terms whatsoever. This is because the decomposition of the field ψ into its chiral-left and -right components creates terms in the Lagrangian that vanish, specifically, the left-right vector terms and the left-left and right-right scalar terms. This leaves left-left and right-right vector terms and left-right scalar terms. However, the left-right scalar terms aren't gauge invariant, so to preserve gauge invariance *all* the fermions in the theory must be massless. This leaves only the left-left and right-right vector terms¹.

¹An in-depth derivation of this fact would take up too much space and interrupt the flow of this chapter. Nevertheless, it is interesting and as such is detailed in full in Appendix B.

That the left-right coupling implies that there are no mass terms has profound implications. *Any* theory of the weak force that takes into account the chirality of the particles must either have *all* particles be massless or have some other mechanism to generate mass for the particles. In nature it is clear that some particles have mass, so the Standard Model must have a mechanism to generate the masses of the particles that is gauge invariant. This mechanism is known as the Higgs mechanism and is discussed in detail in Section 1.5.

As the electroweak force is the unification of two forces, the gauging procedure is done through a direct product of two gauge groups: $SU(2) \times U(1)$, where $SU(2)$ is the weak gauge group and $U(1)$ is the familiar electromagnetic gauge group. In general, such a direct product would mean that an “electroweak gauge transformation” would be the product of a weak gauge transformation and an electromagnetic gauge transformation applied to a field ψ (see Equation 1.41).

$$\psi(x^\mu) \rightarrow \psi'(x^\mu) = e^{i\alpha_k(x^\mu)\sigma_k} e^{i\beta(x^\mu)} \psi(x^\mu) \quad (1.41)$$

Each of the exponentials represent a transformation. The first one, $e^{i\alpha_j(x^\mu)\sigma_j}$, is the $SU(2)$ transformation. It is a 2×2 matrix, understood to be a Taylor expansion like Equation 1.30. The term “ $\alpha_k(x^\mu)\sigma_k$ ” is understood to be a summation over *yet another* index, k . This index is neither spacetime, Dirac nor colour, but is a “flavour” index. The index runs from 1 to 3, with each value corresponding to a function $\alpha_k(x^\mu)$ and a matrix σ_k . The σ_k matrices are the generators of $SU(2)$ and are known as the “Pauli matrices”, shown in Equation 1.42. The second term, $e^{i\beta(x^\mu)}$, is the $U(1)$ transformation familiar from Equation 1.17.

$$\sigma^1 = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix} \quad \sigma^2 = \begin{bmatrix} 0 & -i \\ i & 0 \end{bmatrix} \quad \sigma^3 = \begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix} \quad (1.42)$$

In practice, the fermion field transforms differently depending on its chirality. For chiral-left fermions ($\Psi_L = \hat{P}_L(\psi)$) the fields are arranged in $SU(2)$ doublets which contain either a charged lepton and its corresponding neutrino or two quarks of the same generation (see Equation 1.43). For chiral-right fermions ($\psi_R = \hat{P}_R(\psi)$) the fields are just singlets (Equation 1.44).

$$\Psi_L = \begin{pmatrix} \nu_{e,L} \\ e_L \end{pmatrix}, \begin{pmatrix} \nu_{\mu,L} \\ \mu_L \end{pmatrix}, \begin{pmatrix} \nu_{\tau,L} \\ \tau_L \end{pmatrix}, \begin{pmatrix} u_L \\ d_L \end{pmatrix}, \begin{pmatrix} c_L \\ s_L \end{pmatrix}, \begin{pmatrix} t_L \\ b_L \end{pmatrix} \quad (1.43)$$

$$\psi_R = e_R, \mu_R, \tau_R, u_R, d_R, c_R, s_R, t_R, b_R \quad (1.44)$$

The chiral-left fermion fields undergo both the $SU(2)$ and $U(1)$ transformations; the $SU(2)$ mixes the two components of the doublet and the $U(1)$ applies the same complex phase rotation to both components (Equation 1.45). The chiral-right fermion fields undergo only the $U(1)$ transformation (Equation 1.46).

$$\Psi_L(x^\mu) \rightarrow \Psi'_L(x^\mu) = e^{i\alpha_k(x^\mu)\sigma_k} e^{i\beta(x^\mu)} \Psi_L(x^\mu) \quad (1.45)$$

$$\psi_R(x^\mu) \rightarrow \psi'_R(x^\mu) = e^{i\beta(x^\mu)} \psi_R(x^\mu) \quad (1.46)$$

One might have noticed that Equation 1.44 doesn't list any neutrino fields. This is intentional, since the chiral-right fields don't undergo the $SU(2)$ transformation they also don't feel the weak force. This is merely a footnote for the charged leptons (who undergo the $U(1)$ transformation and interact electromagnetically) and the quarks (who undergo both $SU(3)$ and $U(1)$ so interact via both electromagnetism and the strong force). For the chiral-right neutrinos this is something of a dilemma. They don't interact via the weak force, but since they are electrically neutral and don't possess a colour charge they don't interact electromagnetically or through the strong force either. The theory constructed here has created a particle that doesn't interact with any other particles at all. It is impossible to experimentally validate the existence of such a particle, so the Standard Model simply assumes that chiral-right neutrinos don't exist.

Both the $SU(2)$ and $U(1)$ gauge transformations each have their own conserved charges, generators, gauge bosons and field strength tensors. The charges are the “weak hypercharge” Y and the “weak isospin” I , which correspond to the $U(1)$ and $SU(2)$ groups respectively. These charges are defined in such a way that the sum of weak hypercharge and half the third-component of weak isospin I_3 for a specific fermion is that fermion's electric charge (see Equation 1.47).

$$Q = Y + \frac{I_3}{2} \quad (1.47)$$

Similarly, both $SU(2)$ and $U(1)$ have their own generators. For $SU(2)$ the generators are the Pauli matrices mentioned above. Each of these matrices corresponds to one of three $SU(2)$ gauge bosons: $W_\mu^i, i = 1, 2, 3$. The $U(1)$ transformation has one generator like QED, and has only one gauge boson: B_μ . All four of these gauge bosons are spin-1 and massless.

Gauging the electroweak free fermion Lagrangian under the combined $SU(2) \times U(1)$ group gives the covariant derivatives, but since the chiral-left and -right fields undergo

different transformations the covariant derivatives acts on them in different ways as well (Equation 1.48).

$$\mathcal{D}_\mu \Psi_L = (\partial_\mu + i\frac{g_w}{2}\sigma^a W_\mu^a + ig_Y B_\mu)\Psi_L \quad , \quad \mathcal{D}_\mu \psi_R = (\partial_\mu + ig_Y B_\mu)\psi_R \quad (1.48)$$

Here, g_w and g_Y are the $SU(2)$ and $U(1)$ running coupling constants, respectively. Finally, both the weak and electromagnetic interactions have their own field strength tensors. These are given in Equations 1.49 and 1.50, respectively.

$$W_{\mu\nu}^a = \partial_\mu W_\nu^a - \partial_\nu W_\mu^a + g_w f^{abc} W_\mu^b W_\nu^c \quad (1.49)$$

$$B_{\mu\nu} = \partial_\mu B_\nu - \partial_\nu B_\mu \quad (1.50)$$

The complete, gauge invariant electroweak Lagrangian is given in Equation 1.51.

$$\mathcal{L}_{EWK} = i\bar{\Psi}_L \gamma^\mu \mathcal{D}_\mu \Psi_L + i\bar{\psi}_R \gamma^\mu \mathcal{D}_\mu \psi_R - \frac{1}{4} W^{a,\mu\nu} W_{\mu\nu}^a - \frac{1}{4} B^{\mu\nu} B_{\mu\nu} \quad (1.51)$$

1.5 Electroweak Symmetry Breaking and the Higgs Mechanism

The electroweak theory described in the previous section is fully gauge invariant under $SU(2) \times U(1)$ but it has a serious issue. Like all gauge theories discussed so far it can't possess massive gauge bosons without breaking gauge invariance. This is in direct contradiction with measurements of the masses of the $W^\pm$ and Z^0 bosons, which show that they are far from massless (approximately 80 [16] and 91 [17] GeV, respectively). Further, the chiral nature of the theory means that any fermionic mass term would also break gauge invariance. This also doesn't agree with measurements of nature since all fermions are known to be massive [18, 19]. This means a mechanism which gives masses to the vector bosons and all fermions while leaving the photon and gluons massless, all in a way that preserves gauge invariance, must be added to the theory.

Such a mechanism was proposed in 1964, when three papers were published almost simultaneously by three different groups:

1. François Englert and Robert Brout [20]
2. Peter Higgs [21]

3. Gerald Guralnik, Carl Richard Hagen, Tom Kibble [22]

The mechanism proposed by these three groups would become known as the “Higgs mechanism” named after Peter Higgs². In brief, the mechanism introduces a new $SU(2)$ scalar field with its own dynamics and potential. Above a certain, extremely high temperature (such as the temperature of the universe shortly after the big bang) the electroweak symmetry is preserved. Once the temperature falls below that value, the electroweak symmetry is spontaneously broken³ and the scalar field gains a non-zero vacuum expectation value, or “vev”. This in turn gives the particle interacting with the Higgs field their masses. The rest of this section will document how the mechanism works in detail and how the masses of the particles of the Standard Model are calculated.

Consider four real, scalar fields $\phi_1, \dots, \phi_4$ arranged into a complex $SU(2)$ doublet (Equation 1.52).

$$\phi = \begin{pmatrix} \phi^+ \\ \phi^0 \end{pmatrix} = \begin{pmatrix} \frac{1}{\sqrt{2}}(\phi_1 + i\phi_2) \\ \frac{1}{\sqrt{2}}(\phi_3 + i\phi_4) \end{pmatrix} \quad (1.52)$$

Here, ϕ^+ is a charged scalar field (though not necessarily positively charged) while ϕ^0 is a neutral scalar field. The interactions and dynamics of this field is given by the Lagrangian in Equation 1.53.

$$\mathcal{L}_\phi = (\mathcal{D}^\mu \phi)^\dagger (\mathcal{D}_\mu \phi) - V(\phi) \quad (1.53)$$

The covariant derivative $\mathcal{D}$ in Equation 1.53 is the $SU(2) \times U(1)$ covariant derivative that can be seen in Equation 1.48. It acts on the scalar field the same way it acts on a chiral-left fermion (Equation 1.54).

$$\mathcal{D}_\mu \phi = (\partial_\mu + i\frac{g_w}{2}\sigma^a W_\mu^a + ig_Y B_\mu)\phi \quad (1.54)$$

The potential term is unknown, but a physical potential should be well behaved⁴ enough to be written as a Taylor expansion (Equation 1.55).

$$V(\phi) = V_0 + \sum_{i=1}^{\infty} a_i \phi^i + \sum_{j=1}^{\infty} b_j (\phi^\dagger)^j + \sum_{k=1}^{\infty} \sum_{l=1}^{\infty} c_{k,l} (\phi^\dagger)^k \phi^l \quad (1.55)$$

²Alternative names such as the “Brout-Englert-Higgs mechanism”, the “Englert-Brout-Higgs-Guralnik-Hagen-Kibble mechanism” or the “ABEGHHKt’H mechanism” (for Anderson, Brout, Englert, Higgs, Hagen, Kibble and t’Hooft) are also used, albeit much less extensively.

³A process known as “**E**lectro**W**eak **S**ymmetry **B**reaking” or “EWSB”.

⁴That is, the mathematical function representing a physical potential should be smooth (ie, continuous and differentiable to all orders).

Fortunately, this series is far less complicated than it seems. For the potential to be gauge invariant any term of the form ϕ^i and $(\phi^\dagger)^j$ can't be present since they aren't gauge invariant. Further, terms of the form $(\phi^\dagger)^k \phi^l$ where $k \neq l$ can't be present either, for the same reasons. The constant term can be ignored since it doesn't affect the shape of the potential. This only leaves terms of the form $(\phi^\dagger \phi)^n$, but most of these terms also can't be present. If this theory is to be renormalisable it can't contain terms with $n \geq 3$. This leaves the fairly simple potential in Equation 1.56.

$$V(\phi) = \mu^2 \phi^\dagger \phi + \lambda (\phi^\dagger \phi)^2 \quad (1.56)$$

Here $\mu^2 < 0$ and $\lambda > 0$, which gives the potential its famous “sombbrero” or “Mexican hat” shape. This potential has an infinite number of degenerate minimum energy states that satisfy $\phi^\dagger \phi = \frac{\mu^2}{2\lambda} = \frac{v^2}{2}$. Note that this potential is invariant under $SU(2)$ transformations so choosing one ground state breaks the symmetry. The usual choice in literature is given in Equation 1.57.

$$\langle \phi_0 \rangle = \frac{1}{\sqrt{2}} \begin{bmatrix} 0 \\ v \end{bmatrix} \quad (1.57)$$

Here, “ v ” is the actual vacuum value, and is approximately equal to 246 GeV. This choice is chosen because if conservation of electric charge is to remain unbroken one can only allow the neutral scalar to acquire a vev. Expanding around the vacuum gives Equation 1.58.

$$\langle \phi_0 \rangle = \frac{1}{\sqrt{2}} \begin{bmatrix} 0 \\ v + H(x) \end{bmatrix} \quad (1.58)$$

Here, $H(x)$ is the famous Higgs field. Substituting this value of ϕ into Equation 1.53 gives Equations 1.59 and 1.60.

$$(\mathcal{D}_\mu \phi)^\dagger (\mathcal{D}^\mu \phi) = \left| \left(\partial_\mu + i \frac{g_w}{2} \sigma^a W_\mu^a + i g_Y B_\mu \right) \frac{1}{\sqrt{2}} \begin{pmatrix} 0 \\ v + H(x) \end{pmatrix} \right|^2 \quad (1.59)$$

$$V(\phi) = \mu^2 \left| \begin{pmatrix} 0 \\ v + H(x) \end{pmatrix} \right|^2 + \lambda \left| \begin{pmatrix} 0 \\ v + H(x) \end{pmatrix} \right|^4 \quad (1.60)$$

A full expansion of these terms is beyond the scope of this thesis, but from looking at these two equations a few facts can be derived. The covariant derivative gives the Higgs field its interaction terms with the electroweak gauge bosons while also giving

some of them their masses. The potential gives the three- and four-Higgs interaction vertices. The mass of the Higgs boson itself is given in Equation 1.61.

$$m_H = \sqrt{2\lambda}v \quad (1.61)$$

The $SU(2) \times U(1)$ gauge bosons mix to form the vector bosons ($W^\pm$ and Z^0) and the photon. These combinations and the masses of the new bosons are given in Equation 1.62 for the $W^\pm$, Equation 1.63 for the Z^0 and Equation 1.64 for the photon.

$$W_\mu^\pm = \frac{1}{\sqrt{2}}(W_\mu^1 \mp iW_\mu^2) \text{ with mass } m_{W^\pm} = \frac{1}{2}vg_w \quad (1.62)$$

$$Z_\mu^0 = \frac{1}{\sqrt{g_w^2 + g_Y^2}}(g_w W_\mu^3 - g_Y B_\mu) \text{ with mass } m_{Z^0} = \frac{1}{2}v\sqrt{g_w^2 + g_Y^2} \quad (1.63)$$

$$A_\mu = \frac{1}{\sqrt{g_w^2 + g_Y^2}}(g_Y W_\mu^3 + g_w B_\mu) \text{ with mass } m_A = 0 \quad (1.64)$$

To see how the fermions gain their masses a new term must be added to the electroweak Lagrangian. This is called the “Yukawa Lagrangian” (Equation 1.65), named after Hideki Yukawa.

$$\mathcal{L}_{\text{Yukawa}} = g\bar{\psi}\phi\psi \quad (1.65)$$

This Lagrangian is gauge invariant even after the field ψ is decomposed into its chiral-left and -right components, since the term $\bar{\psi}\phi\psi$ is a singlet under $SU(2) \times U(1)$. The field ψ represents all fermion fields, so a fully expanded-out Yukawa Lagrangian would take the form shown in Equation 1.66.

$$\begin{aligned} \mathcal{L}_{\text{Yukawa}} = & \Gamma_{mn}^u q_{m,L}^- \tilde{\phi} u_{n,R} + h.c. + \Gamma_{mn}^d q_{m,L}^- \phi d_{n,R} + h.c. + \\ & \Gamma_{mn}^e l_{m,L}^- \tilde{\phi} e_{n,R} + h.c. \end{aligned} \quad (1.66)$$

Here the Γ_{mn} coefficients are known as the “Yukawa couplings” which describe the strength of the interaction between the single Higgs doublet and the fermions. There are three sets of Yukawa couplings: Γ_{mn}^u for the up-type quarks, Γ_{mn}^d for the down type quarks and Γ_{mn}^e for the charged leptons. There is an implied summation over m and n . The two chiral-left fermion fields $q_{m,L}$ and $l_{m,L}$ are the $SU(2)$ doublets shown in Equation 1.43 and represent the quarks and leptons, respectively. Similarly, the three chiral-right fermion fields $u_{n,R}$, $d_{n,R}$ and $e_{n,R}$ are the singlets shown in

Equation 1.44, representing the up-type quarks, down-type quarks and the charged leptons, respectively. The indices m and n both run from 1 to 3 with each value representing one generation. Equation 1.52 gives the pre-EWSB scalar field ϕ while $\tilde{\phi}$ is given by Equation 1.67. Finally “*h.c.*” stands for “hermitian conjugate” and is used as short-hand for the hermitian conjugate of the previous term. The “*h.c.*” terms thus contain all of the term of the form $\psi_R \phi \Psi_L^\dagger$ (also $SU(2)$ singlets).

$$\tilde{\phi} = \begin{pmatrix} \phi^{0*} \\ -\phi^- \end{pmatrix} = \begin{pmatrix} \frac{1}{\sqrt{2}}(\phi_1 + i\phi_2) \\ \frac{1}{\sqrt{2}}(\phi_3 + i\phi_4) \end{pmatrix} \quad (1.67)$$

It’s important to note that the Yukawa Lagrangian doesn’t have a term for neutrinos. This may seem like a strange thing to draw attention to, since this chapter has already established that the Standard Model assumes that chiral-right neutrinos don’t exist. The reason this is important is that it has consequences for the left-handed neutrinos, which do exist. If the Yukawa Lagrangian had a term for chiral-right neutrinos, it must look like this: $\Gamma_{mn}^\nu l_{m,L}^- \phi \nu_{n,R}$. Such a term is impossible since there are no $\psi_{n,R}$ fields, which means that the Higgs mechanism can’t generate masses for chiral-left neutrinos.

The conclusion here is that there is no mechanism in the Standard Model that can generate mass for neutrinos. This is at odds with what was stated at the beginning of the chapter, which is that according to the latest measurements neutrinos have a very-small, but still non-zero mass [3]. This means that the Standard Model must be extended in some way to give it the ability to generate masses for neutrinos. At the time of writing this is a topic which is still being researched, with a variety of extensions to the Standard Model having been proposed. A full description of these extensions is beyond the scope of this thesis. Suffice it to say that even though the Standard Model on its own doesn’t account for neutrino masses, it can be extended in a straightforward way to include them.

The masses of all the non-neutrino fermions can now be generated from specific terms in the Yukawa Lagrangian using $\tilde{\phi}$ and ϕ . Consider the first generation of charged leptons and quarks as an example (Equation 1.68).

$$\mathcal{L}_{\text{Yukawa, FirstGen}} = y_e \bar{l}_L \phi e_R + h.c. + y_u \bar{q}_L \tilde{\phi} u_R + h.c. + y_d \bar{q}_L \phi d_R + h.c. \quad (1.68)$$

Here, $y_e = \Gamma_{11}^e$, $y_u = \Gamma_{11}^u$ and $y_d = \Gamma_{11}^d$. Choosing the Higgs vev in Equation 1.57 gives the values for ϕ and $\tilde{\phi}$ in Equation 1.69.

$$\phi = \frac{1}{\sqrt{2}} \begin{pmatrix} 0 \\ v \end{pmatrix} \Rightarrow \tilde{\phi} = \frac{1}{\sqrt{2}} \begin{pmatrix} v \\ 0 \end{pmatrix} \quad (1.69)$$

Substituting these values into Equation 1.68 gives Equation 1.70.

$$\mathcal{L}_{\text{Yukawa, FirstGen, EWSB}} = \frac{y_e v}{\sqrt{2}} (\bar{e}_L e_R + \bar{e}_R e_L) + \frac{y_u v}{\sqrt{2}} (\bar{u}_L u_R + \bar{u}_R u_L) + \frac{y_d v}{\sqrt{2}} (\bar{d}_L d_R + \bar{d}_R d_R) \quad (1.70)$$

The terms of from $\bar{\psi}_L \phi_R + \bar{\psi}_R \psi_L$ are all $SU(2)$ singlets and so are gauge invariant. They reduce down to $\bar{\psi}\psi$ and the mass terms can just be read off (Equation 1.71).

$$m_e = \frac{y_e v}{\sqrt{2}} \quad m_u = \frac{y_u v}{\sqrt{2}} \quad m_d = \frac{y_d v}{\sqrt{2}} \quad (1.71)$$

This example generalises to all charged leptons and quarks, with their masses being given in Equation 1.72.

$$m_f = \frac{y_f v}{\sqrt{2}} \quad (1.72)$$

Thus, the Higgs boson is responsible for giving mass to all of the massive fundamental particles of the Standard Model⁵.

1.6 The Standard Model

With all three forces of the Standard Model understood, it is now possible to write a Lagrangian for the full Standard Model. Since the previous two sections have dealt with how all fermions in the Standard Model interact via two of the three fundamental forces described by the Standard Model as well as how their masses are generated using the Higgs mechanism, the work is almost done. All that's left to do is to add the strong force.

⁵It should be noted that the mass of composite particles such as the proton is *not* equal to the sum of the masses of the particles that make it up. Indeed, the sum of the masses of the valence quarks in a proton accounts for around 5% of its mass while the other 95% comes from QCD effects.

The Standard Model begins with a free fermion Lagrangian that is gauged under the complete $SU(3) \times SU(2) \times U(1)$ group. The generic transformation of this group is given in Equation 1.73.

$$\psi(x^\mu) \rightarrow \psi'(x^\mu) = e^{i\alpha_j(x^\mu)\lambda_j} e^{i\alpha_k(x^\mu)\sigma_k} e^{i\beta(x^\mu)} \psi(x^\mu) \quad (1.73)$$

The chiral-left and -right components transform similarly to what is shown in Equation 1.46, with the chiral-left field transforming under $SU(3) \times SU(2) \times U(1)$ and the chiral-right transforming under $SU(3) \times U(1)$. Similarly, only the quarks undergo the $SU(3)$ component of the transformation. The covariant derivative of this transformation is given in Equation 1.74.

$$\mathcal{D}_\mu = \partial_\mu + i\frac{g_s}{2}\lambda^a G_\mu^a + i\frac{g_w}{2}\sigma^b W_\mu^b + ig_Y B_\mu \quad (1.74)$$

Again, the chiral-left and -right components have different covariant derivatives act on them, and the quarks are the only fermions that have the $SU(3)$ term in their covariant derivatives. The indices a and b are the familiar colour and flavour indices, respectively. There are three field strength tensors which can be combined into a total field strength tensor $\mathcal{F}^{\mu\nu}$ (Equation 1.75).

$$\mathcal{F}_{\mu\nu} = \underbrace{G_{\mu\nu}^a}_{SU(3) \text{ Strong}} + \underbrace{W_{\mu\nu}^a}_{SU(2) \text{ Weak}} + \underbrace{B_{\mu\nu}}_{U(1) \text{ Electromagnetic}} \quad (1.75)$$

The scalar field ϕ behaves the same way as described in Section 1.5. It doesn't undergo the $SU(3)$ transformation and so it doesn't interact with quarks and gluons at all.

Using these terms the Standard Model Lagrangian can be written in full. It is shown in Equation 1.76

$$\begin{aligned}
 \mathcal{L} = & - \underbrace{\frac{1}{4} F^{\mu\nu} F_{\mu\nu}}_{\text{Field Strength Tensor}} \\
 & + \underbrace{i \bar{\Psi} \gamma^\mu \mathcal{D}_\mu \Psi}_{\text{Fermion Dynamics and Interactions with Gauge Bosons}} \\
 & + \underbrace{(\mathcal{D}^\mu \Phi)^\dagger (\mathcal{D}_\mu \Phi)}_{\text{Higgs Field Dynamics and Interactions with Gauge Bosons}} \\
 & + \underbrace{V(\Phi)}_{\text{Higgs Potential}} \\
 & + \underbrace{\bar{\Psi}_L \Gamma \Phi \Psi_R}_{\text{Yukawa Couplings}} + h.c.
 \end{aligned} \tag{1.76}$$

This Lagrangian contains terms that describe all the fundamental particle interactions of the Standard Model. Each term corresponds to a Feynman diagram showing a three-or-four particle interaction. The diagrams for the three fundamental forces described by the Standard Model is shown in Figure 1.2. The interactions with the Higgs field generates masses for all the particles after EWSB and these interactions can produce actual Higgs boson particles.

Standard Model Interactions (Forces Mediated by Gauge Bosons)




Fig. 1.2 Feynman diagrams of the fundamental interactions of the Standard Model.

Chapter 2

Beyond the Standard Model

The Standard Model has been remarkably successful at describing fundamental physics phenomena. Experiments have probed energies up to several TeV and haven't found any unambiguous evidence of new physics¹. However, it is clear that the Standard Model is an incomplete description of nature and will have to be extended to describe physics at higher energies. At a bare minimum, an extension will be required once experiments reach the reduced Planck scale, the scale at which quantum gravitational effects become important.

The electroweak scale ($v = 246$ GeV) and the reduced Planck scale ($M_P = 2.4 \times 10^{18}$ GeV) are separated by sixteen orders of magnitude, which means that barring some miraculous new technology humanity won't be able to probe the reduced Planck scale for quite some time. Does this mean humanity is stuck making increasingly precise measurements of Standard Model parameters until we reach that frontier? Not necessarily. After all, nature has surprised even the most brilliant physicist so it isn't unreasonable to assume that new physics will manifest somewhere in those 16 orders of magnitude.

With this in mind, many extensions to the Standard Model have been proposed over the years. Among these is “supersymmetry” (shortened to “SUSY”), an extension that proposes that there exists a symmetry between fermions and bosons. This extension is theoretically well-motivated, and if SUSY proves to be an accurate description of nature then it solves many of the shortcomings of the Standard Model. The core of this thesis is a search for SUSY particles, so it is necessary to spend a chapter defining

¹In the context of particle physics, the term “new physics” refers to any as-yet undiscovered fundamental phenomena. It is used interchangeable with the term “physics beyond the Standard Model”.

what it is, how it solves some of the issues of the Standard Model and what kinds of new particles it predicts.

This chapter mostly uses information drawn from two sources: “A Supersymmetry Primer” by Stephen P. Martin [23] and “Supersymmetry and the MSSM: An Elementary Introduction” by Ian Aitchison [24]. Section 2.1 mostly draws from Ref. [24] and Sections 2.2 and 2.3 mostly draw from Ref. [23].

2.1 Limitations of the Standard Model

The ultimate goal of physics is to construct a so-called “Theory of Everything”; a theoretically self-consistent theory that explains all physical phenomena. While the Standard Model is far from being a theory of everything, its success at describing fundamental phenomena makes it a good candidate to be extended with an eye towards creating one. In this view, whenever a phenomenon isn’t able to be explained by the Standard Model (be it something related to quantum mechanics like the hierarchy problem or something totally beyond the purview of quantum mechanics like dark energy) it can be thought of as a deficiency of the Standard Model².

While there are many examples of beyond the Standard Model (or “BSM”) physics, it is beyond the scope of this thesis to give a detailed description of all of them. Instead, only three will be described here: the hierarchy problem, dark matter and gravity. These three have been chosen because SUSY extends the Standard Model in a way which provides satisfying explanations of all of these phenomena, thereby providing good motivation for SUSY as an extension to the Standard Model.

2.1.1 The Hierarchy Problem

This chapter has already made passing reference to a “hierarchy problem” without explaining what it is. In brief, it has to do with the μ and λ parameters in the Higgs potential (see Section 1.5) and how they’re calculated.

The values of physical parameters (such as masses and coupling constants) that appear in the Standard Model Lagrangian are known as the “fundamental values” of those parameters. These fundamental values aren’t the same as the values of those parameters that are measured by experiments, which are known as the “effective values” of those parameters. The fundamental and effective values of a given parameter

²Alternatively, if one has a more optimistic temperament they can be thought of as areas where the Standard Model could potentially be extended.

typically aren't equal and the effective value is found by taking the fundamental value and adding quantum corrections to it.

A hierarchy problem manifests when a very specific set of quantum corrections is required to produce the effective value of a parameter. To better illustrate this consider the following example. Suppose there is a quantum field theory with a parameter ρ . This parameter has an effective value of 120 GeV and a fundamental value of 100 GeV. The theory gives the following corrections to the fundamental value: 2 GeV, 5 GeV, 10,000,020 GeV and $-10,000,007$ GeV. The theory works perfectly, the sum of the corrections and the fundamental value is the same as the effective value. However, two of the corrections (10,000,020 and $-10,000,007$) are several orders of magnitude larger than ρ and it is quite fortuitous that their difference is only 7 GeV.

This is the essence of a hierarchy problem. The effective value of a parameter is highly dependent on a specific cancellation taking place. This isn't a problem for the theory as it exists since it gives the correct prediction, but one might imagine that the theory will need to be extended at some point. This would lead to possible changes to the correction terms (for instance, the 10,000,020 becoming 10,005,020), which might in turn upset the delicate cancellation and give a wildly different effective value for the parameter.

The actual hierarchy problem³ occurs in the Standard Model. Recall the Higgs potential after expansion around the minimum (Equation 1.60). Taking just the electrically-neutral part gives Equation 2.1.

$$V(H) = -\mu^2|H|^2 + \lambda|H|^4 \quad (2.1)$$

The Higgs discovery in 2012[25, 26] means that, assuming the Standard Model is correct as an effective field theory⁴, the parameters in Equation 2.1 have effective values: $\mu^2 = (92.9 \text{ GeV})^2$ and $\lambda = -0.126$.

³It should be noted that the hierarchy problem described here isn't the only hierarchy problem in physics. In cosmology, there is another hierarchy problem that deals with comparatively large quantum corrections to the very-small but non-zero cosmological constant, Λ . The hierarchy problem described here has a more technical name: the "weak-scale instability problem". While it isn't *technically* correct to refer to the weak-scale instability problem as *the* hierarchy problem, it is so ubiquitous when discussing extensions to the Standard Model that it is accepted convention to refer to the weak-scale instability problem simply as the hierarchy problem. This thesis will follow this convention throughout.

⁴The phrase "the Standard Model is correct as an effective field theory" simply means that, for all parameters in the theory the corrected fundamental value and effective value are the same to within error.

The corrections to the fundamental values of those two terms are many and varied, but to illustrate the hierarchy problem only one needs to be considered (Equation 2.2).

$$\lambda \int_0^\infty d^4k \frac{1}{k^2 - m_H^2} \quad (2.2)$$

This integral is over all energy values and diverges quadratically (the numerator is proportional to k^4 while the denominator is proportional to k^2). To solve this, a “cut-off scale” Λ_{UV} is introduced (Equation 2.3). This scale is used to regulate the loop integral, and should be interpreted as the scale at which new physics manifests and undermines the ability of the Standard Model to give accurate predictions. Above this scale the Standard Model will need to be extended with new fields and interactions which will change the way that calculations are done. In particular, above this scale the integral shown in Equation 2.2 will no longer be accurate.

$$\lambda \int_0^{\Lambda_{UV}} d^4k \frac{1}{k^2 - m_H^2} \quad (2.3)$$

Evaluating this integral gives a correction to the value of μ that is proportional to $\lambda \Lambda_{UV}^2$ ⁵. This is where the hierarchy problem becomes apparent. Suppose that no new physics appears until the reduced Planck scale: $\Lambda_{UV} = M_P = 2.4 \times 10^{18}$ GeV. This means that the term $\lambda \Lambda_{UV}^2 \sim \mathcal{O}(10^{35})$ GeV. This is 35 (!) orders of magnitude greater than the effective value of μ . The only way to obtain such an effective value is to have another term of similar size to $\lambda \Lambda_{UV}^2$ but negative so that they cancel. This “fortuitous” cancellation is the centre of the hierarchy problem, although at these scales it would probably be more accurate to call it “miraculous” rather than simply “fortuitous”.

This problem is extremely disturbing in part because it affects the Higgs potential parameters: μ and λ . Since every massive particle in the Standard Model derives its mass from the Higgs Boson, the mass of every particle is thus dependent on those two parameters.

2.1.2 Dark Matter

Another limitation of the Standard Model can be seen in astrophysical observations of the universe. Galaxies are rotating at such speeds that their observable matter couldn’t possibly hold them together gravitationally. This has lead to physicists postulating

⁵The term in the Lagrangian takes the form: $(\mu^2 + \lambda \Lambda_{UV}^2) \bar{\phi} \phi$

that there is something in these galaxies that isn't directly detectable but nonetheless produces the necessary gravitational force to keep galaxies together.

This “something” is dark matter [27]. Unlike normal or “luminous” matter, dark matter doesn't interact via the electromagnetic force. This means that it is invisible (light passes directly through it without interacting) and intangible (luminous matter passes through it without interacting⁶). It does interact gravitationally and is more abundant than luminous matter (outnumbering luminous matter approximately 6 to 1 [27]), two facts which when combined allows it to hold galaxies together where luminous matter alone wouldn't be able to.

The Standard Model doesn't account for dark matter at all. Dark matter needs to be made up of particles like luminous matter, so a “candidate” particle needs to be put forward to fulfil this role. The only particle that comes close to fulfilling this role in the Standard Model is the neutrino, since it only interacts via the weak force. However, the neutrinos very small mass combined with its measured abundance in the universe prevents it from fulfilling the gravitational role of dark matter, so dark matter probably isn't made up of neutrinos.

With this in mind, the Standard Model needs to be extended so that it can provide a dark matter candidate.

2.1.3 Gravity

The Standard Model explicitly avoids modelling gravity. That said, calling this a “limitation” or a “failing” of the Standard Model understates just how difficult it is to unite quantum field theory (the theory of the electroweak and strong forces) with general relativity (the theory of gravity).

A good example of why the two theories are so difficult to unify is the so-called “vacuum catastrophe” (also known as the “cosmological constant problem”) [28]. In general relativity, the cosmological constant, Λ , present in Einstein's field equations⁷ gives the energy density of empty space (this is also known as the “vacuum energy”). In quantum field theory empty space is defined by the “vacuum state”, the quantum state with the lowest possible energy. This state contains all quantum fields in their ground states, so it generally means that no real particles are present. However, all these fields fluctuate around their ground states, producing pairs of virtual particles

⁶All macroscopic “push and pull” forces are manifestations of the electromagnetic force. Specifically, the electrons in matter repel each other by exchanging photons. This means that dark matter, which doesn't interact via the electromagnetic force, passes right through luminous matter.

⁷ $R_{\mu\nu} - \frac{1}{2}Rg_{\mu\nu} + \Lambda g_{\mu\nu} = \frac{8\pi G}{c^4}T_{\mu\nu}$

that annihilate back into the vacuum quickly enough to still obey the uncertainty principle. In order to use quantum field theory to calculate the vacuum energy, these fluctuations must be taken into account.

The value of the vacuum energy has been measured cosmologically and calculated using quantum field theory. The problem is that these two values differ by *120 order of magnitude*! This is the vacuum catastrophe, and is the single greatest discrepancy between theoretical calculation and experimental measurement in the history of physics [28].

With such a large discrepancy between general relativity and quantum field theory, it should come as no surprise that attempts to unify them are incredibly mathematically involved. Quantum field theory attempts to model gravity as a force in the same vein as the strong and electroweak forces, even theorising the existence of a gauge boson for gravity, the “graviton”. String theory is another theoretical framework which is attempting to explain both quantum and relativistic phenomena.

2.2 Supersymmetry

Supersymmetry is both the name of a symmetry and the name of an extension to the Standard Model. The symmetry is a symmetry between bosons and fermions, and is based on a transformation that turns one into the other. A system preserves supersymmetry (it is said to be “supersymmetric”) if the system is invariant under that transformation. The extension to the Standard Model involves introducing a set of new supersymmetric partner (or “superpartner”) particles to Standard Model particles. These superpartners are the result of applying SUSY transformations to Standard Model particles.

Supersymmetry is a very complex theory with a rich mathematical structure that harbours many unexpected and pleasing results, but this also means that an introduction to the theory with only a mediocre level of detail would be over 100 pages long. Further, there isn’t just one SUSY extension to the Standard Model. Since no SUSY particles have been discovered, there aren’t many restrictions on what a SUSY extension can contain. As such, a host of different extensions to the Standard Model have been codified: the MSSM, MSUGRA, superstring theory, and so on.

This section is written with these two facts in mind. In the interest of being concise it will favour broad overviews over detailed explanations and instead of introducing a specific SUSY model this section will detail how a generic SUSY model is constructed while highlighting characteristics that all SUSY models share.

First, Subsection 2.2.1 will introduce the actual SUSY transformation operator that acts on states and fields. Subsection 2.2.2 will then introduce how quantum fields are organised in SUSY. The largest Subsection, 2.2.3, will detail the construction of a generic SUSY theory with multiple particles and interactions, as well as showing how to write down a SUSY Lagrangian. Finally, Subsection 2.2.4 will briefly describe what it means for SUSY to be a broken symmetry.

2.2.1 The Supersymmetry Transformation

SUSY transformations transform bosons into fermions and vice-versa. The generator of SUSY transformations, labelled $\hat{Q}$, is an anti-commuting spinor. Since spinors are complex objects the hermitian conjugate of $\hat{Q}$, $\hat{Q}^\dagger$ must also be a generator of the transformation. As such, any SUSY transformation can be written as a linear combination of $\hat{Q}$ and $\hat{Q}^\dagger$ (Equation 2.4).

$$(a\hat{Q} + b\hat{Q}^\dagger)|\text{Boson}\rangle = |\text{Fermion}\rangle \quad (\alpha\hat{Q} + \beta\hat{Q}^\dagger)|\text{Fermion}\rangle = |\text{Boson}\rangle \quad (2.4)$$

2.2.2 SUSY Particle States

As a quantum field theory, SUSY uses the same field formalism that the Standard Model does to describe its particles. SUSY arranges its particles into special “doublets”, called “supermultiplets”. Supermultiplets are the single-particle states of a supersymmetric theory, and are the irreducible representations of the SUSY algebra. Each supermultiplet contains two fields: one fermion and one boson, which are each others “superpartners”. By definition, the superpartners are related to each other via a SUSY transformation, so if a supermultiplet contains the fields Ω and Ω' then Ω' is equal to some linear combination of the $\hat{Q}$ and $\hat{Q}^\dagger$ operators acting on Ω (up to a spacetime translation and/or rotation).

The squared-mass operator, $-\hat{P}^2$, commutes with both $\hat{Q}$ and $\hat{Q}^\dagger$ as well as with all of the spacetime translation and rotation operators. This means that particles which inhabit the same supermultiplet must have the same eigenvalues of $-\hat{P}^2$ and thus must have equal masses. Similarly, the SUSY generators $\hat{Q}$ and $\hat{Q}^\dagger$ commute with the generators of all the gauge transformations. This means that particles in the same supermultiplet must also have the same electric charge, weak isospin and colour degrees of freedom. Finally, supermultiplets must have the same number of fermionic

degrees of freedom, n_F , as bosonic degrees of freedom, n_B (Equation 2.5). All of these results are highly consequential and will see use throughout the rest of this chapter.

$$n_F = n_B \quad (2.5)$$

The simplest possibility for a supermultiplet that fulfils Equation 2.5 consists of a single Weyl fermion⁸ with two spin helicity states ($n_F = 2$) and two real scalars ($n_B = 1 + 1$). The two real scalar fields are then assembled into one complex scalar field. The fermions transform under the fundamental representation of the gauge group, so their bosonic superpartners must do the same. This combination of a Weyl fermion and a complex scalar field is called a “chiral supermultiplet”⁹.

The next-simplest option for a supermultiplet would contain a spin-1 vector boson. If SUSY is to be renormalisable, then the gauge boson must be massless (at least prior to any gauge symmetry breaking). A massless spin-1 boson has two degrees of freedom, so $n_B = 2$. Its superpartner must therefore be a massless spin- $\frac{1}{2}$ Weyl fermion with two helicity states so that $n_F = 2$. These superpartners are known as “gauginos”. Gauge bosons transform under the adjoint representation of the gauge group, so the fermionic gauginos must do the same. Further, because the adjoint representation of a gauge group is always its own conjugate, then the chiral-left and -right components of the gauginos must have the same gauge transformation properties. The gauge boson and gaugino are arranged into another supermultiplet, known as a “gauge supermultiplet”¹⁰.

Using only chiral and gauge supermultiplets, one can already construct some of the more prominent SUSY extensions to the Standard Model, such as the MSSM. However, other extensions to Standard Model go even further, introducing more exotic supermultiplets. One might want to extend their theory to include gravity, which means a new gauge boson must be theorised to act as its force carrier. Such a particle has already been proposed: the graviton, a massless spin-2 gauge boson. It has two helicity states, so $n_B = 2$. Its proposed superpartner would be the spin- $\frac{3}{2}$ gravitino, which would also be massless (if SUSY is unbroken) in order to fulfil the requirement that $n_F = 2$.

⁸The term “Weyl fermions” refers to the fermions under the Weyl representation. This is in contrast to the “Dirac representation” used in the previous chapter. The Dirac representation has its fermion field as a four-component spinor (recall Equation 1.14) while the Weyl representation has it as a two-component spinor.

⁹Although this thesis will use chiral supermultiplet throughout, other terms such as “scalar supermultiplet” and “matter supermultiplet” are also used in the literature.

¹⁰This thesis will use the term gauge supermultiplet throughout, but the term “vector supermultiplet” is used throughout the literature as well.

2.2.3 Constructing a SUSY Lagrangian

Constructing a Lagrangian for SUSY is a much more complex endeavour than it is for the Standard Model. This is due to not only the increased particle content, but also because SUSY makes use of mathematical constructs unknown to the Standard Model. As such, this subsection will begin with a few simple example SUSY models to illustrate a few concepts before constructing a complete theory.

The simplest possible SUSY model consists of a single, free chiral supermultiplet: a left-handed, two-component Weyl fermion ψ with a complex scalar field ϕ as its superpartner. Equations 2.6 and 2.7 show the simplest possible action and Lagrangians that can be written down using these two fields. This model is known as the massless, non-interacting Wess-Zumino model.

$$\mathcal{S} = \int d^4x (\mathcal{L}_\psi + \mathcal{L}_\phi) \quad (2.6)$$

$$\mathcal{L}_\psi = i\psi^\dagger \bar{\sigma}^\mu \partial_\mu \psi, \quad \mathcal{L}_\phi = -\partial^\mu \phi^* \partial_\mu \phi \quad (2.7)$$

A candidate for a SUSY transformation should turn a scalar field ϕ into something involving the fermion field ψ . This is given in Equation 2.8.

$$\delta\phi = \epsilon\psi, \quad \delta\phi = \epsilon^\dagger\psi^\dagger \quad (2.8)$$

Here, the factor ϵ is an infinitesimal, non-commuting, two-component Weyl fermion object that parameterises the SUSY transformation. If ϵ is a constant (ie, $\partial_\mu \epsilon = 0$) then SUSY is a global symmetry, if ϵ is a function of the spacetime coordinate x^μ then SUSY is a local symmetry. For the sake of brevity this thesis will only consider global SUSY, but promoting SUSY to a local symmetry has important implications related to gravity that Subsection 2.4.3 will briefly touch on.

Under this transformation, the scalar Lagrangian $\mathcal{L}_\phi$ undergoes the transformation given in Equation 2.9.

$$\delta\mathcal{L}_\phi = -\epsilon\partial^\mu\psi\partial_\mu\phi^* - \epsilon^\dagger\partial^\mu\psi^\dagger\partial_\mu\phi \quad (2.9)$$

In order for this theory to be invariant under SUSY, the action needs to remain unchanged by a SUSY transformation. This means that $\delta\mathcal{L}_\phi$ must cancel $\delta\mathcal{L}_\psi$ (at least,

up to a total derivative). To this end, there is only one choice for the fermion SUSY transformation (Equation 2.10).

$$\delta\psi_\alpha = -i(\sigma^\mu\epsilon^\dagger)_\alpha\partial_\mu\phi, \quad \delta\psi^\dagger_\beta = i(\epsilon\sigma^\mu)_\beta\partial_\mu\phi^* \quad (2.10)$$

Using this transformation (and some mathematical manipulations) one arrives at the fermion Lagrangian transformation given in Equation 2.11.

$$\begin{aligned} \delta\mathcal{L}_\psi &= \epsilon\partial^\mu\psi\partial_\mu\phi^* + \epsilon^\dagger\partial^\mu\psi^\dagger\partial_\mu\phi \\ &\quad - \partial_\mu(\epsilon\sigma^\nu\bar{\sigma}^\mu\psi\partial_\nu\phi^* + \epsilon\psi\partial^\mu\phi^* + \epsilon^\dagger\psi^\dagger\partial^\mu\phi) \end{aligned} \quad (2.11)$$

The first two terms of $\delta\mathcal{L}_\psi$ cancel the two terms in $\delta\mathcal{L}_\phi$, leaving only the total derivative. This means that the theory is indeed invariant under the transformations given in Equations 2.8 and 2.10. Note that I haven't referred to these transformations as "SUSY transformations" yet. This is because SUSY transformations must also obey the fact that the SUSY algebra is closed. In practice, this means that the commutator of two SUSY transformations parameterised by two different spinors ϵ_1 and ϵ_2 is also a symmetry of the theory. Using Equations 2.8 and 2.10 (and some more mathematical manipulations) one gets Equations 2.12 and 2.13.

$$(\delta_{\epsilon_2}\delta_{\epsilon_1} - \delta_{\epsilon_1}\delta_{\epsilon_2})\phi = \delta_{\epsilon_2}(\delta_{\epsilon_1}\phi) - \delta_{\epsilon_1}(\delta_{\epsilon_2}\phi) = i(-\epsilon_1\sigma^\mu\epsilon_2^\dagger + \epsilon_2\sigma^\mu\epsilon_1^\dagger)\partial_\mu\phi \quad (2.12)$$

$$(\delta_{\epsilon_2}\delta_{\epsilon_1} - \delta_{\epsilon_1}\delta_{\epsilon_2})\psi_\alpha = i(-\epsilon_1\sigma^\mu\epsilon_2^\dagger + \epsilon_2\sigma^\mu\epsilon_1^\dagger)\partial_\mu\psi_\alpha + i\epsilon_{1,\alpha}\epsilon_2^\dagger\bar{\sigma}^\mu\partial_\mu\psi - i\epsilon_{2,\alpha}\epsilon_1^\dagger\bar{\sigma}^\mu\partial_\mu\psi \quad (2.13)$$

This is a singular result. Equation 2.12 shows that the commutation of two SUSY operators return a constant times the derivative of the original field. Further, the factor $i\partial_\mu$ is the momentum operator $\hat{P}_\mu$, the generator of spacetime translations. This means that SUSY transformations are intimately related to spacetime translations.

Equation 2.13 is somewhat problematic. The first two terms are the same as the ones in Equation 2.12, but the last two terms only vanish on-shell¹¹. This is something of a problem since this symmetry should hold quantum mechanically¹². Fortunately, there is a method to make the algebra close off-shell. This is through the introduction

¹¹That is, the equations of motion derived from the principle of least action are only enforced.

¹²That is, the symmetry holds both on- and off-shell.

of a new complex field, F . This field is called an *auxiliary* and has a very simple Lagrangian (Equation 2.14).

$$\mathcal{L}_{\text{Auxiliary}} = F^* F \quad (2.14)$$

This auxiliary field has the equations of motion $F = F^* = 0$, which are fundamentally uninteresting¹³. Since it's a scalar field the SUSY transform turns it into a fermion field ψ (Equation 2.15) while its Lagrangian transform vanishes on-shell, but not off-shell (Equation 2.16).

$$\delta F = -i\epsilon^\dagger \bar{\sigma}^\mu \partial_\mu \psi, \quad \delta F^* = i\partial_\mu \psi^\dagger \bar{\sigma}^\mu \epsilon \quad (2.15)$$

$$\delta \mathcal{L}_{\text{Auxiliary}} = -i\epsilon^\dagger \bar{\sigma}^\mu \partial_\mu \psi F^* + i\partial_\mu \psi^\dagger \bar{\sigma}^\mu \epsilon F \quad (2.16)$$

These transformations were chosen so that the fermion transformation could be redefined (Equation 2.17).

$$\delta \psi_\alpha = -i(\sigma^\mu \epsilon^\dagger)_\alpha \partial_\mu \phi + \epsilon_\alpha F, \quad \delta \psi_\beta^\dagger = i(\epsilon \sigma^\mu)_\beta \partial_\mu \phi^* + \epsilon_\beta^\dagger F^* \quad (2.17)$$

This changes $\delta \mathcal{L}_\psi$, which now has an extra term which cancels $\delta \mathcal{L}_{\text{Auxiliary}}$, so this modified theory “ $\mathcal{L}_\psi + \mathcal{L}_\phi + \mathcal{L}_{\text{Auxiliary}}$ ” is still SUSY invariant. Repeating the commutation from earlier, Equation 2.18 is the result.

$$(\delta_{\epsilon_2} \delta_{\epsilon_1} - \delta_{\epsilon_1} \delta_{\epsilon_2})X = i(-\epsilon_1 \sigma^\mu \epsilon_2^\dagger + \epsilon_2 \sigma^\mu \epsilon_1^\dagger)X \quad (2.18)$$

Here, $X = \psi, \psi^\dagger, \phi, \phi^*, F$ and F^* . Thus, with the auxiliary field, the commutation of two transformations acting on any field in the theory is still a symmetry of the theory. In turn, this means that SUSY is a preserved symmetry of this theory both on- and off-shell.

The above example of a single, free chiral supermultiplet will serve as a guide for constructing a more realistic SUSY model. Such a model should contain multiple chiral supermultiplets that interact with one another. Creating a theory with multiple, free chiral supermultiplets is trivial. For each chiral supermultiplet in the theory, simply add another set of ψ, ϕ and F fields (see Equation 2.19).

$$\mathcal{L}_{\text{Free}} = -\partial^\mu ((\phi^i)^*) \partial_\mu \phi_i + i(\psi^i)^\dagger \bar{\sigma}^\mu \partial_\mu \psi_i + (F^i)^* F_i \quad (2.19)$$

¹³Indeed, these auxiliary fields are just bookkeeping devices that allow the SUSY algebra to remain closed off-shell.

The index i runs over all particles in the theory¹⁴. Each fermion ψ_i has its own specific gauge quantum numbers which it shares with its superpartner ϕ_i . The field F_i is the auxiliary field for the i^{th} chiral supermultiplet. Transformations for these fields are given in Equations 2.20 through 2.22.

$$\delta\phi_i = \epsilon\psi_i \qquad \delta(\phi^i)^* = \epsilon^\dagger(\psi^i)^\dagger \qquad (2.20)$$

$$\delta(\psi_i)_\alpha = -i(\sigma^\mu\epsilon^\dagger)_\alpha\partial_\mu\phi_i + \epsilon_\alpha F_i \qquad \delta((\psi^i)^\dagger)_\beta = i(\epsilon\sigma^\mu)_\beta\partial_\mu(\phi^i)^* + \epsilon^\dagger_\beta(F^i)^* \qquad (2.21)$$

$$\delta F_i = -i\epsilon^\dagger\bar{\sigma}^\mu\partial_\mu\psi_i \qquad \delta(F^i)^* = i\partial_\mu(\psi^i)^\dagger\bar{\sigma}^\mu\epsilon \qquad (2.22)$$

The transformations shown in Equations 2.20 through 2.22 are SUSY transformations (both on- and off-shell) and the theory is invariant under them. The next step is to introduce interactions into the theory. This is done through the addition of an interaction Lagrangian $\mathcal{L}_{\text{Int}}$, which contains all of the interaction terms of the theory. Constructing the most general possible interaction Lagrangian might seem daunting, since it can technically contain infinitely many terms. Luckily, requiring the theory to be renormalisable imposes a strong restriction on what terms can appear in it (Equation 2.23).

$$\mathcal{L}_{\text{Int}} = \left(-\frac{1}{2}W^{ij}\psi_i\psi_j + W^iF_i + x^{ij}F_iF_j \right) + c.c. - U(\phi, (\phi^i)^*) \qquad (2.23)$$

Here, W^{ij} , W^i , x^{ij} and U are polynomials in the fields ϕ and $(\phi^i)^*$ of degree 1, 2, 0 and 4, respectively. Their form is unknown, but can be found by imposing another restriction: SUSY invariance. For $\mathcal{L}_{\text{Int}}$ to be SUSY invariant, the function $U(\phi, (\phi^i)^*)$ must vanish since its SUSY transformation can't be cancelled by any other term. The dimensionless coupling x^{ij} must likewise vanish for similar reasons. The result is Equation 2.24.

$$\mathcal{L}_{\text{Int}} = \left(-\frac{1}{2}W^{ij}\psi_i\psi_j + W^iF_i \right) \qquad (2.24)$$

Finding the forms of W^{ij} and W^i is harder (see [23] for details). It turns out that both W^{ij} and W^i are related and are derived from the same object, W , given in Equation 2.25.

$$W = \frac{1}{2}M^{ij}\phi_i\phi_j + \frac{1}{6}y^{ijk}\phi_i\phi_j\phi_k \qquad (2.25)$$

¹⁴This index shouldn't be confused with the spinor indices that were present in previous Lagrangians but have been suppressed here.

This is called the “superpotential”. It isn’t a scalar potential in the regular sense (in fact, it isn’t even real), but a holomorphic function where the scalar fields ϕ_i are treated as complex variables. Here, M^{ij} is a symmetric mass matrix and y^{ijk} are the Yukawa couplings for three field interactions that is totally symmetric under the exchange of i, j and k . As mentioned above, the superpotential defines both W^{ij} and W^i (Equation 2.26).

$$W^i = \frac{\partial W}{\partial \phi_i}, \quad W^{ij} = \frac{\partial^2 W}{\partial \phi_i \partial \phi_j} \quad (2.26)$$

The most significant result that can be derived from the superpotential is that it can be used to remove the need for auxiliary fields. Taking the full theory $\mathcal{L}_{\text{Free}} + \mathcal{L}_{\text{Int}}$ and finding the equations of motion gives Equation 2.27.

$$F_i = -(W_i)^* \quad (F^i)^* = -W^i \quad (2.27)$$

This means that the auxiliary fields can be written in terms of the scalar fields ϕ_i . The total Lagrangian ($\mathcal{L} = \mathcal{L}_{\text{Free}} + \mathcal{L}_{\text{Int}}$) of this theory can now be written (Equation 2.28).

$$\begin{aligned} \mathcal{L} = & -\partial^\mu ((\phi^i)^*) \partial_\mu \phi_i + i(\psi^i)^\dagger \bar{\sigma}^\mu \partial_\mu \psi_i \\ & - \frac{1}{2} (W^{ij} \psi_i \psi_j + (W_{ij})^* (\psi^i)^\dagger (\psi^j)^\dagger) \\ & - W^i (W_i)^* \end{aligned} \quad (2.28)$$

This Lagrangian can be further simplified by introducing the scalar potential $V(\phi, \phi^*) = W^i (W_i)^*$. Using this and substituting the general form of the superpotential into Equation 2.28 gives the full Lagrangian density (Equation 2.29)¹⁵.

$$\begin{aligned} \mathcal{L}_{\text{Chiral}} = & -\partial^\mu ((\phi^i)^*) \partial_\mu \phi_i + i(\psi^i)^\dagger \bar{\sigma}^\mu \partial_\mu \psi_i - V(\phi, \phi^*) \\ & - \frac{1}{2} M^{ij} \psi_i \psi_j - \frac{1}{2} (M_{ij})^* (\psi^i)^\dagger (\psi^j)^\dagger \\ & - \frac{1}{2} y^{ijk} \phi_i \psi_j \psi_k - \frac{1}{2} (y_{ijk})^* (\phi^i)^* (\psi^j)^* (\psi^k)^* \end{aligned} \quad (2.29)$$

¹⁵Using the equations of motion that are derived from this Lagrangian as well as the diagonalised mass matrix M^{ij} one can show that the Weyl fermions and their complex scalar superpartners are indeed mass-degenerate.

This Lagrangian contains every possible chiral-chiral interaction. The next step in constructing a realistic SUSY theory is the inclusion of gauge fields and interactions. This is actually fairly straightforward. The addition of gauge supermultiplets into the theory uses much of the same methodology as chiral supermultiplets, while the gauge interactions are mostly implemented through covariant derivatives like they are in the Standard Model.

Gauge supermultiplets contain a massless gauge boson field A_μ^a and a two-component fermionic gaugino Υ^a . The index a represents the gauge indices familiar from the previous chapter. It runs over $a = 1, \dots, 8$ for $SU(3)$ colour, $a = 1, 2, 3$ for $SU(2)$ weak isospin and $a = 1$ for $U(1)$ weak hypercharge. The gauge supermultiplet fields undergo gauge transformations (Equations 2.30 and 2.31).

$$A_\mu^a \rightarrow A_\mu^{a'} = A_\mu^a + \partial_\mu \alpha^a + g f^{abc} A_\mu^b \alpha^c \quad (2.30)$$

$$\Upsilon^a \rightarrow \Upsilon^{a'} = \Upsilon^a + g f^{abc} \Upsilon^b \alpha^c \quad (2.31)$$

Here α^a is the infinitesimal gauge parameter, g is the gauge coupling and f^{abc} are the totally antisymmetric structure constants of the gauge group¹⁶. The Lagrangian of a theory that contains just these two fields is given in Equation 2.32.

$$\mathcal{L}_{\text{Gauge}} = -\frac{1}{4} F_{\mu\nu}^a F^{\mu\nu,a} + i(\Upsilon^a)^\dagger \bar{\sigma}^\mu \nabla_\mu \Upsilon^a + \frac{1}{2} D^a D^a \quad (2.32)$$

Here, $F_{\mu\nu}^a$ is the familiar gauge field strength (Equation 2.33) and $\mathcal{D}_\mu \Upsilon^a$ is the covariant derivative of the gaugino field (Equation 2.34).

$$F_{\mu\nu}^a = \partial_\mu A_\nu^a - \partial_\nu A_\mu^a + g f^{abc} A_\mu^b A_\nu^c \quad (2.33)$$

$$\mathcal{D}_\mu \Upsilon^a = \partial_\mu \Upsilon^a + g f^{abc} A_\mu^b \Upsilon^c \quad (2.34)$$

The fields D^a are the auxiliary fields, analogous to the F field for chiral supermultiplets. Each gauge field has its own auxiliary field, hence the “ a ” index. In contrast to the fermion auxiliary fields, the fields D^a are all real. The SUSY transformations of all three of these fields is given in Equations 2.35, 2.36 and 2.37.

$$\delta A_\mu^a = -\frac{1}{\sqrt{2}}(\epsilon^\dagger \bar{\sigma}_\mu \Upsilon^a + (\Upsilon^a)^\dagger \bar{\sigma}_\mu \epsilon) \quad (2.35)$$

¹⁶The special case of an abelian group has $f^{abc} = 0, \forall a, b, c$.

$$\delta\Upsilon_\alpha = \frac{i}{2\sqrt{2}}(\sigma^\mu\bar{\sigma}^\nu\epsilon)_\alpha F_{\mu\nu}^a + \frac{1}{\sqrt{2}}\epsilon_\alpha D^a \quad (2.36)$$

$$\delta D^a = \frac{i}{\sqrt{2}}(-\epsilon^\dagger\bar{\sigma}^\mu\nabla_\mu\Upsilon^a + \nabla_\mu(\Upsilon)^\dagger\bar{\sigma}^\mu\epsilon) \quad (2.37)$$

The factors $\frac{1}{\sqrt{2}}$, $\frac{i}{2\sqrt{2}}$ and $\frac{i}{\sqrt{2}}$ are chosen so that the action obtained by integrating $\mathcal{L}_{\text{Gauge}}$ is indeed SUSY invariant. The phase of Υ^a is chosen so that the treatment of the MSSM in the next section is more convenient. Using the above three equations it is straightforward, if tedious, to show that these transformations fulfil the requirement that the SUSY algebra is closed (Equation 2.38).

$$(\delta_{\epsilon_2}\delta_{\epsilon_1} - \delta_{\epsilon_1}\delta_{\epsilon_2})X = i(-\epsilon_1\sigma^\mu(\epsilon_2)^\dagger + \epsilon_2\sigma^\mu(\epsilon_1)^\dagger)\nabla_\mu X \quad (2.38)$$

Here, X can be replaced by any of the gauge-covariant fields $F_{\mu\nu}^a$, Υ^a , $(\Upsilon^a)^\dagger$ and D^a as well as for arbitrary covariant derivatives acting on them.

Now, suppose that the chiral supermultiplets transform under the fundamental representation of the gauge group which has hermitian matrices $(T^a)_i^j$ which satisfies the familiar commutation relation: $[T^a, T^b] = if^{abc}T^c$ ¹⁷. Since the SUSY and gauge transformations commute, the fermion, scalar and auxiliary fields must all transform the same (Equation 2.39).

$$X_i \rightarrow X'_i = X_i + ig\alpha^a(T^a X)_i \quad (2.39)$$

Here, $X = \phi_i, \psi_i, F_i$.

In order for this theory to be gauge invariant, the regular derivatives $\partial_\mu\phi_i$, $\partial_\mu(\phi^i)^*$ and $\partial_\mu\psi_i$ must be replaced everywhere with covariant derivatives (Equations 2.40 through 2.42).

$$\mathcal{D}_\mu\phi_i = \partial_\mu\phi_i - igA_\mu^a(T^a\phi)_i \quad (2.40)$$

$$\mathcal{D}_\mu(\phi^i)^* = \partial_\mu(\phi^i)^* + igA_\mu^a(\phi^*T^a)^i \quad (2.41)$$

$$\mathcal{D}_\mu\psi_i = \partial_\mu\psi_i - igA_\mu^a(T^a\psi)_i \quad (2.42)$$

This actually completes most of the work, since applying the covariant derivatives automatically produces interactions between the fermions/scalars in the chiral

¹⁷For instance, if this were the $SU(3)$ gauge group, the T^a matrices would be the 8 Gell-Mann matrices and f^{abc} would be the structure constants.

supermultiplets and gauge bosons in the gauge supermultiplets. All that's left are the interactions between the chiral supermultiplets and the gaugino and auxiliary fields. However, due to the requirement for SUSY to be renormalisable, there are only three such interactions (Equation 2.43).

$$(\phi^* T^a \psi) \Upsilon^a, \quad (\Upsilon^a)^\dagger (\psi^\dagger T^a \phi) \text{ and } (\phi^* T^a \phi) D^a \quad (2.43)$$

One might think that all that needs to happen now is to add the above terms to the chiral and gauge Lagrangians and demand the result be real and SUSY invariant. While this is the case, there needs to be a few things taken care of first. Namely, the SUSY transformations given in Equations 2.20 through 2.22 must have their regular derivatives replaced with covariant derivatives, and δF_i must have another term strategically added.

$$\delta \phi_i = \epsilon \psi_i \quad (2.44)$$

$$\delta \psi_{i,\alpha} = -i(\sigma^\mu \epsilon^\dagger)_\alpha \mathcal{D}_\mu \phi_i + \epsilon_\alpha F_i \quad (2.45)$$

$$\delta F_i = -i\epsilon^\dagger \bar{\sigma}^\mu \mathcal{D}_\mu \psi_i + \sqrt{2}g(T^a \phi)_i \epsilon^\dagger (\Upsilon^a)^\dagger \quad (2.46)$$

Now a full, completely general SUSY Lagrangian can be written down (Equation 2.47).

$$\begin{aligned} \mathcal{L} = & \mathcal{L}_{\text{Chiral}} + \mathcal{L}_{\text{Gauge}} \\ & - \sqrt{2}g(\phi^* T^a \psi) \Upsilon^a - \sqrt{2}g(\Upsilon^a)^\dagger (\psi^\dagger T^a \phi) + g(\phi^* T^a \phi) D^a \end{aligned} \quad (2.47)$$

Here the factors in front of the three terms shown in Equation 2.43 are fixed by some algebra. The last term can be used to find an equation of motion for the auxiliary field D^a (Equation 2.48).

$$D^a = -g(\phi^* T^a \phi) \quad (2.48)$$

Thus, like the fermion auxiliary field F , D can be expressed purely in terms of the scalar fields. Knowing this, the definition of the scalar potential $V(\phi, \phi^*)$ can be extended to include both auxiliary fields (Equation 2.49).

$$V(\phi, \phi^*) = (F^i)^* F_i + \frac{1}{2} \sum_a D^a D^a = (W^i)^* W_i + \frac{1}{2} \sum_a g_a^2 (\phi^* T^a \phi)^2 \quad (2.49)$$

The two expressions found in this potential are known as the F - and D -terms. Since V is defined as the sum of sums of squares, it is always ≥ 0 . Another interesting feature of SUSY is that the scalar potential is completely determined by other interactions in the theory, the F -terms come from the Yukawa couplings and the D -terms come from the gauge interactions.

2.2.4 SUSY Breaking

This section has tried to be as general as possible in constructing a SUSY model, but it has made extensive use of the fact that such a model should be SUSY invariant. One consequence of this is that Standard Model particles and their superpartners must have the same masses, a fact that this section has explicitly stated a few times. This is problematic since if SUSY partners have the same masses as their Standard Model counterparts, why haven't they been discovered yet?

The obvious answer to this dilemma is that SUSY is a broken symmetry. To put it more precisely, it is expected that SUSY is an exact symmetry that is broken spontaneously. In other words, the SUSY model should have a Lagrangian that is invariant under SUSY, but a vacuum state that isn't. This way, SUSY breaking would behave much the same way that electroweak symmetry breaking does. At high energies (such as in the early universe) SUSY invariance is preserved and superpartners are mass degenerate, but once the universe cools below a specific energy, SUSY is broken and superpartners begin to take on different masses.

That said, there isn't really a consensus on how SUSY breaking should be implemented theoretically. The simplest idea is to take the Lagrangian of whatever extension one is working with and just add extra terms that explicitly break the symmetry. Such terms must be "soft" (ie, their couplings should have positive mass dimensions), as this would create a separation between the masses of Standard Model and SUSY particles in a natural way. For the generic theory described in this section, the most general "Soft SUSY-Breaking" Lagrangian is given in Equation 2.50.

$$\mathcal{L}_{\text{Soft}} = - \left(\frac{1}{2} M_a \Upsilon^a \Upsilon^a + \frac{1}{6} a^{ijk} \phi_i \phi_j \phi_k + \frac{1}{2} b^{ij} \phi_i \phi_j + t^i \phi_i \right) + c.c. - (m^2)_j^i (\phi^j)^* \phi_i \quad (2.50)$$

2.3 The MSSM

The previous section derived the most general SUSY model possible, with the only assumption being that SUSY was an exact symmetry that is broken spontaneously.

While this produced many useful general results, it doesn't provide anything substantive that experimentalists can search for. To do this, generality must be discarded in favour of a distinct SUSY model that predicts the existence of particles/values for parameters that experimentalists can search for/measure.

As mentioned in the previous section there are many different SUSY extensions to the Standard Model, and this thesis will use the “**Minimal Supersymmetric Standard Model**” (or “MSSM”) as its model of choice. The MSSM is the minimal¹⁸ phenomenologically viable SUSY extension to the Standard Model, and serves as a good starting model for SUSY searches.

2.3.1 Particle Content and Interactions

The minimal nature of the MSSM comes from the way it defines SUSY particles. For the most part, each particle in the Standard Model is defined to have a superpartner whose spin differs by $\frac{1}{2}$. This pair of particles is then placed in either a chiral or gauge supermultiplet.

The spin- $\frac{1}{2}$ fermions of the Standard Model are placed into chiral supermultiplets with their spin-0 boson superpartners. Each boson is named by taking the name of their fermionic superpartner and pre-pending the letter “s” (for “scalar”), while they use the same symbol as their superpartners, but with a “ $\sim$ ” above it. Generically, leptons and quarks have “sleptons” and “squarks” as their superpartners, while specific particles such as the electron, e , has the selectron, $\tilde{e}$, as its superpartner.

The Standard Model treats chiral-left and -right components of fermions as different particles which transform differently under the $SU(2) \times U(1)$ gauge group, so it is necessary to have corresponding “left-handed” and “right-handed” superpartners for each of these particles. For instance, the chiral-left and -right electrons e_L and e_R would have as their superpartners the “left-handed” and “right-handed” selectrons $\tilde{e}_L$ and $\tilde{e}_R$, respectively. It should be noted that the bosonic superpartners are spin-0, so by definition they can't have a chirality. Rather, their “handedness” refers to the chirality of the Standard Model fermions they are the superpartners of.

Thus, each of the chiral-left and -right leptons (electrons: e_L and e_R , muons: μ_L and μ_R and taus: τ_L and τ_R) have as their superpartners “left-handed” and “right-handed” sleptons (selectrons: $\tilde{e}_L$ and $\tilde{e}_R$, smuons: $\tilde{\mu}_L$ and $\tilde{\mu}_R$ and staus: $\tilde{\tau}_L$ and $\tilde{\tau}_R$). The Standard Model only has left-handed neutrinos, so the “ L ” subscript is suppressed and they are generically referred to by their lepton flavour: ν_e , ν_μ and ν_τ . These have

¹⁸Minimal in the sense that it is the SUSY model which introduces the smallest number of new particles.

corresponding “sneutrinos”: $\tilde{\nu}_e$, $\tilde{\nu}_\mu$ and $\tilde{\nu}_\tau$. The chiral-left and -right quarks (u_L , u_R , d_L , d_R , s_L , s_R , c_L , c_R , b_L , b_R , t_L and t_R) have their corresponding squarks ($\tilde{u}_L$, $\tilde{u}_R$, $\tilde{d}_L$, $\tilde{d}_R$, $\tilde{s}_L$, $\tilde{s}_R$, $\tilde{c}_L$, $\tilde{c}_R$, $\tilde{b}_L$, $\tilde{b}_R$, $\tilde{t}_L$ and $\tilde{t}_R$). Finally, both the Standard Model fermions and their SUSY superpartners undergo the same gauge interactions. For example, the left-handed sleptons $\tilde{e}_L$, $\tilde{\mu}_L$ and $\tilde{\tau}_L$ couple to the Z boson while the right-handed sleptons $\tilde{e}_R$, $\tilde{\mu}_R$ and $\tilde{\tau}_R$ don't.

The spin-1 gauge bosons of the Standard Model are placed into gauge supermultiplets, which they share with their spin- $\frac{1}{2}$ fermionic superpartners, the gauginos. The gauginos are named by taking the name of their corresponding gauge boson and adding “-ino” to the end. Like the sleptons and squarks, gauginos use the same symbols as their superpartners but with a tilde above them. As discussed in the previous section, each of the four $SU(2) \times U(1)$ and eight $SU(3)$ gauge bosons has their own gaugino.

The gauge bosons of the $SU(3)$ gauge group are called the “gluons” and are denoted g . Each gluon has as its superpartner a “gluino”, denoted $\tilde{g}$. The eight gluinos exist in a colour octet like the gluons do. The (pre-EWSB) gauge bosons of the $SU(2) \times U(1)$ gauge group: W^1 , W^2 , W^3 and B^0 all have corresponding spin- $\frac{1}{2}$ fermionic superpartners: $\tilde{W}^1$, $\tilde{W}^2$, $\tilde{W}^3$ and $\tilde{B}^0$, which are known as the “winos” and the “bino”. After electroweak symmetry breaking the gauginos mix to form new particles in a similar way to their Standard Model superpartners. The $\tilde{W}^3$ and the $\tilde{B}^0$ mix together to form the $\tilde{Z}^0$ and $\tilde{\gamma}$, called the “zino” and “photino”, respectively.

This is where the similarity ends, since there are many more mixings that occur between SUSY particles than between their Standard Model superpartners. Indeed, while the $\tilde{W}^1$ and $\tilde{W}^2$ *do* mix with each other, they also mix with the Higgs fields to form new particles that don't have Standard Model analogues. This is detailed in Subsection 2.3.2.

The use of the term “Higgs fields” isn't an accident. The Higgs boson is spin-0, so it must reside in a chiral supermultiplet which it shares with its spin- $\frac{1}{2}$ fermionic superpartner, the “higgsino”. However, it turns out that having just *one* higgs supermultiplet isn't enough, since it upsets a cancellation that removes a gauge anomaly from the Standard Model.

One of the conditions for the cancellation of gauge anomalies is given in Equation 2.51.

$$Tr[(I_3)^2 Y] = Tr[(Y)^3] = 0 \quad (2.51)$$

Here, I_3 and Y are the third-component of weak isospin and weak hypercharge, respectively, and are familiar from Equation 1.47. These traces runs over *all* left-handed

Weyl fermion degrees of freedom in the theory. In the Standard Model, this condition is (somewhat miraculously) already satisfied by the leptons and quarks. The addition of another Weyl fermion in the form of the higgsino, which by definition has $Y = \frac{1}{2}$ or $Y = -\frac{1}{2}$, means that the condition no longer holds.

This problem can be solved if there are *two* Higgs supermultiplets, where one has $Y = \frac{1}{2}$ and the other has $Y = -\frac{1}{2}$, so the condition is once again met. While this solution might seem slightly contrived, there are other reasons that one might want to have two Higgs supermultiplets. One of these is that due to the structure of supersymmetric theories, only $Y = \frac{1}{2}$ Higgs chiral supermultiplet can give masses to the $+\frac{2}{3}$ charged up-type quarks. Another supermultiplet is needed with $Y = \frac{1}{2}$ to give masses to the $-\frac{1}{3}$ charged down-type quarks.

The end result is that there are now *four* independent Higgs fields: two complex scalar fields with weak hypercharges $Y = \frac{1}{2}$ and $Y = -\frac{1}{2}$, and their two superpartners. The Higgs fields with weak hypercharge $Y = \frac{1}{2}$ and $Y = -\frac{1}{2}$ will be denoted H_u and H_d , respectively, while their superpartners will be denoted $\tilde{H}_u$ and $\tilde{H}_d$, respectively. Each of these four fields is itself an $SU(2)$ doublet, whose components have different values of weak isospin (either $I_3 = \frac{1}{2}$ or $I_3 = -\frac{1}{2}$). These are written in Equation 2.52, with the convention that the first component in the doublet has $I_3 = \frac{1}{2}$ and the second component in the doublet has $I_3 = -\frac{1}{2}$.

$$\begin{aligned} H_u &= \begin{pmatrix} H_u^+ \\ H_u^0 \end{pmatrix} & H_d &= \begin{pmatrix} H_d^0 \\ H_d^- \end{pmatrix} \\ \tilde{H}_u &= \begin{pmatrix} \tilde{H}_u^+ \\ \tilde{H}_u^0 \end{pmatrix} & \tilde{H}_d &= \begin{pmatrix} \tilde{H}_d^0 \\ \tilde{H}_d^- \end{pmatrix} \end{aligned} \quad (2.52)$$

The fields H_u^+ , H_u^0 , H_d^0 and H_d^- are the four spin-0 bosonic Higgs fields, while the $\tilde{H}_u^+$, $\tilde{H}_u^0$, $\tilde{H}_d^0$ and $\tilde{H}_d^-$ are their spin- $\frac{1}{2}$ fermionic superpartners. The Standard Model Higgs field is a linear combination of H_u^0 and H_d^0 since the Standard Model Higgs fields gives mass to all Standard Model particles but has no electric charge.

This completes the particle content of the MSSM, which is summarised in Table 2.1.

The MSSM interactions can be seen in its superpotential (Equation 2.53). Here, the dimensionless Yukawa couplings $\mathbf{y}_u$, $\mathbf{y}_d$ and $\mathbf{y}_e$ are 3×3 matrices in generation space and the term μ is the SUSY equivalent of the Standard Model Higgs mass.

$$\begin{aligned} W_{\text{MSSM}} &= \mathcal{U}_R^\dagger \mathbf{y}_u \mathcal{Q}_L H_u - \mathcal{D}_R^\dagger \mathbf{y}_d \mathcal{Q}_L H_d - \ell_R^\dagger \mathbf{y}_e L_L H_d + \mu H_u H_d \\ &+ \tilde{\mathcal{U}}_R^* \mathbf{y}_u \tilde{\mathcal{Q}}_L \tilde{H}_u - \tilde{\mathcal{D}}_R^* \mathbf{y}_d \tilde{\mathcal{Q}}_L \tilde{H}_d - \tilde{\ell}_R^* \mathbf{y}_e \tilde{L}_L \tilde{H}_d + \mu \tilde{H}_u \tilde{H}_d \end{aligned} \quad (2.53)$$

Table 2.1 Particle content of the Minimal Supersymmetric Standard Model

MSSM Particles							
Chiral Supermultiplets				Gauge Supermultiplets			
Standard Model		SUSY		Gauge Bosons		Gauginos	
Field	Spin	Field	Spin	Field	Spin	Field	Spin
Leptons		Sleptons		$SU(2) \times U(1)$ Gauge Bosons			
$L_L = \begin{pmatrix} \nu_{e,L} \\ e_L \\ \nu_{\mu,L} \\ \mu_L \\ \nu_{\tau,L} \\ \tau_L \end{pmatrix}$	$\frac{1}{2}$	$\tilde{L}_L = \begin{pmatrix} \tilde{\nu}_{e,L} \\ \tilde{e}_L \\ \tilde{\nu}_{\mu,L} \\ \tilde{\mu}_L \\ \tilde{\nu}_{\tau,L} \\ \tilde{\tau}_L \end{pmatrix}$	0	W^1		$\tilde{W}^1$	
				W^2	1	$\tilde{W}^2$	0
				W^3		$\tilde{W}^3$	
$\ell_R = \begin{pmatrix} e_R \\ \mu_R \\ \tau_R \end{pmatrix}$	$\frac{1}{2}$	$\tilde{\ell}_R = \begin{pmatrix} \tilde{e}_R \\ \tilde{\mu}_R \\ \tilde{\tau}_R \end{pmatrix}$	0	B^0	1	$\tilde{B}^0$	$\frac{1}{2}$
Quarks		Squarks		$SU(3)$ Gauge Bosons			
$\mathcal{Q}_L = \begin{pmatrix} u_L \\ d_L \\ c_L \\ s_L \\ t_L \\ b_L \end{pmatrix}$	$\frac{1}{2}$	$\tilde{\mathcal{Q}}_L = \begin{pmatrix} \tilde{u}_L \\ \tilde{d}_L \\ \tilde{c}_L \\ \tilde{s}_L \\ \tilde{t}_L \\ \tilde{b}_L \end{pmatrix}$	0	g	1	$\tilde{g}$	$\frac{1}{2}$
$\mathcal{U}_R = \begin{pmatrix} u_R \\ c_R \\ t_R \end{pmatrix}$	$\frac{1}{2}$	$\tilde{\mathcal{U}}_R = \begin{pmatrix} \tilde{u}_R \\ \tilde{c}_R \\ \tilde{t}_R \end{pmatrix}$	0				
$\mathcal{D}_R = \begin{pmatrix} d_R \\ s_R \\ b_R \end{pmatrix}$	$\frac{1}{2}$	$\tilde{\mathcal{D}}_R = \begin{pmatrix} \tilde{d}_R \\ \tilde{s}_R \\ \tilde{b}_R \end{pmatrix}$	0				
Higgs		Higgsinos					
$H_u = \begin{pmatrix} H_u^+ \\ H_u^0 \end{pmatrix}$	0	$\tilde{H}_u = \begin{pmatrix} \tilde{H}_u^+ \\ \tilde{H}_u^0 \end{pmatrix}$	$\frac{1}{2}$				
$H_d = \begin{pmatrix} H_d^0 \\ H_d^- \end{pmatrix}$		$\tilde{H}_d = \begin{pmatrix} \tilde{H}_d^0 \\ \tilde{H}_d^- \end{pmatrix}$					

2.3.2 The MSSM and Symmetry Breaking

The particles listed in the previous subsection represent the particle content of the MSSM before electroweak symmetry breaking. Electroweak symmetry breaking is more complicated in SUSY since there are two complex Higgs doublets compared to the single doublet in the Standard Model. The Higgs potential of the MSSM is given in Equation 2.54¹⁹.

$$\begin{aligned}
 V = & (|\mu|^2 + m_{H_u}^2)(|H_u^0|^2 + |H_u^+|^2) + (|\mu|^2 + m_{H_d}^2)(|H_d^0|^2 + |H_d^-|^2) \\
 & + [b(H_u^+ H_d^- - H_u^0 H_d^0) + c.c.] \\
 & + \frac{1}{8}(g^2 + (g')^2)(|H_u^0|^2 + |H_u^+|^2 - |H_d^0|^2 - |H_d^-|^2)^2 + \frac{1}{2}g^2 |H_u^+(H_d^0)^* + H_u^0(H_d^-)^*|^2
 \end{aligned} \tag{2.54}$$

This equation can be simplified by performing some gauge transformations. Namely, an $SU(2)$ transformation can be performed that sets $H_u^+ = 0$, and after such a transformation one can find a minimum of the potential satisfying $H_d^- = 0$. This is a welcome result, since these two being zero means that only the neutral scalars acquire vevs and so charge conservation remains unbroken²⁰. Applying this result to Equation 2.54 gives Equation 2.55.

$$\begin{aligned}
 V = & (|\mu|^2 + m_{H_u}^2)|H_u^0|^2 + (|\mu|^2 + m_{H_d}^2)|H_d^0|^2 - [bH_u^0 H_d^0 + c.c.] \\
 & + \frac{1}{8}(g^2 + (g')^2)(|H_u^0|^2 - |H_d^0|^2)^2
 \end{aligned} \tag{2.55}$$

This leaves two Higgs fields, H_u^0 and H_d^0 . Each of these fields are responsible for giving masses to different sections of the particle content of the MSSM. Namely, the H_u^0 field gives masses to the up-type quarks while H_d^0 gives masses to the down-type quarks and the leptons.

Electroweak symmetry breaking also leads to particle mixing, and in SUSY the greater particle content leads to a greater amount of particles being created. One can think of it in terms of degrees of freedom. Pre-EWSB, the two complex Higgs doublets account for 8 degrees of freedom. Post-EWSB, these degrees of freedom move to different particles. Three of them are absorbed by the Standard Model W bosons

¹⁹Actually, this isn't the full MSSM Higgs potential. There are slepton and squark terms that do not get vevs and are thus ignored.

²⁰A similar choice is made in the Standard Model, see Section 1.5.

which become the $W^\pm$ and the Z^0 , while the other five turn into the neutral scalars h^0 , H^0 and A^0 and the charged scalars $H^\pm$. Here, h^0 is the familiar Standard Model Higgs boson with a mass of 125 GeV, while the other four are roughly mass degenerate but much heavier.

SUSY particles mix with each other as well. In addition to the $\tilde{B}^0$ and $\tilde{W}^3$ mixing mentioned in the previous subsection are the higgsino mixings. The neutral higgsinos and the gauginos combine to form four “neutralinos”, while the charged higgsinos and the winos mix to give two “charginos”. Due to having their origins in electroweak symmetry breaking, the neutralinos and charginos are collectively referred to as “electroweakinos”. All of these mixings are summarised in Table 2.2.

Table 2.2 Particle content of the MSSM pre- and post-EWSB and the mixing that it induces. Mixing between in the first two generations of sleptons and squarks is assumed to be negligible.

Names	Spin	Gauge Eigenstates	Mass Eigenstates
Sleptons	0	$\tilde{e}_L, \tilde{e}_R, \tilde{\nu}_e$	(Same)
	0	$\tilde{\mu}_L, \tilde{\mu}_R, \tilde{\nu}_\mu$	(Same)
	0	$\tilde{\tau}_L, \tilde{\tau}_R, \tilde{\nu}_\tau$	$\tilde{\tau}_1, \tilde{\tau}_2, \tilde{\nu}_\tau$
Squarks	0	$\tilde{u}_L, \tilde{u}_R, \tilde{d}_L, \tilde{d}_R$	(Same)
	0	$\tilde{s}_L, \tilde{s}_R, \tilde{c}_L, \tilde{c}_R$	(Same)
	0	$\tilde{t}_L, \tilde{t}_R, \tilde{b}_L, \tilde{b}_R$	$\tilde{t}_1, \tilde{t}_2, \tilde{b}_1, \tilde{b}_2$
Higgs Bosons	0	$H_u^+, H_u^0, H_d^0, H_d^-$	$h^0, H^0, A^0, H^\pm$
Neutralinos	$\frac{1}{2}$	$\tilde{B}^0, \tilde{W}^3, \tilde{H}_u^0, \tilde{H}_d^0$	$\tilde{\chi}_1^0, \tilde{\chi}_2^0, \tilde{\chi}_3^0, \tilde{\chi}_4^0$
Charginos	$\frac{1}{2}$	$\tilde{W}^1, \tilde{W}^2, \tilde{H}_u^+, \tilde{H}_d^-$	$\tilde{\chi}_1^\pm, \tilde{\chi}_2^\pm$
Gluino	$\frac{1}{2}$	$\tilde{g}$	(Same)

2.3.3 R-Parity

The MSSM described in this section is “minimal” in the sense that it contains the smallest number of fields necessary to produce a phenomenologically viable model, but there are other terms that could be added to the theory but aren’t since they violate either Baryon (B) or Lepton (L) number symmetry. Baryon and lepton number symmetry seems to be preserved in nature, with the strongest experimental evidence of this being the extremely long lifetime of the proton. Any good QFT should contain B and L conservation, and indeed the Standard Model has it as a consequence of

renormalisation. SUSY on the other hand doesn't, so either B and L conservation needs to be added to the theory post-hoc, which is a step back from the Standard Model, or another symmetry has to be included which has B and L conservation as a consequence.

Naturally, the second option is chosen for the MSSM through the concept of "R-Parity". Each particle in the theory has an R-Parity value, defined by Equation 2.56.

$$\mathcal{P}_R = (-1)^{3(B-L)+2s} \quad (2.56)$$

Here, s is the spin of the particle. R-Parity has the interesting property that $\mathcal{P}_R = +1$ for all Standard Model particles and Higgs bosons and $\mathcal{P}_R = -1$ for all sleptons, squarks, gauginos and higgsinos. Further, if R-Parity is conserved there are several important consequences:

- Both B and L are conserved
- There is no mixing between $\mathcal{P}_R = +1$ and -1 particles
- Every interaction vertex *must* have an even number of $\mathcal{P}_R = -1$ particles, so collider experiments can only produce them in pairs
- The lightest $\mathcal{P}_R = -1$ particle (also known as the "lightest supersymmetric particle" or "LSP") is absolutely stable and only interacts via the weak interaction
- Each non-LSP SUSY particle must eventually decay into a state that contains an odd number of LSPs

The MSSM is *defined* to conserve R-Parity, so all of the consequences listed above are necessarily true in the MSSM. This is another reason why the MSSM is a popular extension to the Standard Model.

2.4 SUSY and the Limits of the Standard Model

Returning to the starting premise of this chapter: the Standard Model is unable to explain various phenomena and so it must be extended. To be a well-motivated extension the MSSM must be able to explain at least some of the phenomena that the Standard Model can't, otherwise all the work that has been done is merely a mathematical novelty.

Of course, if SUSY and the MSSM *weren't* able to solve these problems they wouldn't have had an entire chapter dedicated to explaining them. This section will consider each of the problems listed in Section 2.1 and explain how the MSSM would solve them.

2.4.1 The Hierarchy Problem

Recall Equation 2.3, the integral that calculates the one-loop contribution to the Higgs potential parameter $-\mu^2$. The hierarchy problem manifests here because the cut-off scale is so large that higher-order contributions must cancel each other in improbable ways to give the comparatively small value of $-\mu^2$. One such contribution comes from the scalar field ϕ decaying into a fermion-antifermion pair which then annihilate back into the ϕ field. This contribution takes the form given in Equation 2.57.

$$\left(-g_f^2 \int^{\Lambda_{UV}} \frac{1}{k\bar{k}} d^4k\right) \hat{\phi}^\dagger \hat{\phi} \sim -g_f^2 \hat{\phi}^\dagger \hat{\phi} \Lambda_{UV} \quad (2.57)$$

The sign of this contribution is crucial, since when it is combined with Equation 2.3 it gives the total contribution shown in Equation 2.58.

$$(\lambda - g_f) \Lambda_{UV} \hat{\phi}^\dagger \hat{\phi} \quad (2.58)$$

There is still a quadratic sensitivity to Λ_{UV} , but now there is also a factor of $\lambda - g_f$ as well. If, for whatever reason it could be shown that $\lambda = g_f$, then the whole factor would disappear and the hierarchy problem would be solved. Unsurprisingly, such a relation between the 4-boson coupling and the the boson-fermion coupling described above is a characteristic of SUSY.

A side-effect of this process is that any new SUSY particles must have a mass that is $\mathcal{O}(10 \text{ TeV})$. In theory, this means that SUSY particles could be discovered using present-day or near-future detector technology [29].

2.4.2 Dark Matter Candidate

While the Standard Model is unable to provide a candidate particle for dark matter, the MSSM actually has one: the LSP. Recall some of the properties of the LSP from Subsection 2.3.3:

- The LSP is *absolutely stable*
- It only interacts via the weak interaction

Add to these two properties the ability to interact gravitationally (which is a reasonable assumption about the particles of the MSSM, even if the MSSM technically doesn't model gravity) makes the LSP an ideal dark matter candidate.

2.4.3 Gravity

While it should be noted that while the MSSM doesn't model gravity, there are other SUSY theories that do.

Recall Equation 2.8, the first definition of a SUSY transformation. The ϵ parameterises the SUSY transformation, and it is assumed to be a constant (ie, SUSY is a global symmetry). This SUSY is an exact symmetry that is spontaneously broken. Much like EWSB, this leads to the creation of a SUSY Goldstone boson, creatively named the “goldstino”. This particle is massless and electrically neutral.

Making ϵ a function of the spacetime coordinate x^μ promotes SUSY to a local symmetry. Doing this gives rise to a new particle, the “gravitino”, which is to the local SUSY symmetry as the gauge bosons are to the local gauge symmetries. The gravitino is a Weyl fermion, so its superpartner, named the “graviton”, is a scalar boson. Before SUSY-breaking, both of these particles are massless, but after SUSY is broken the gravitino absorbs the goldstino and gains a mass. This process is analogous to the Higgs mechanism of gauge theories and has been named the “Super-Higgs mechanism”.

The theory described here is called “supergravity”, and it is noteworthy since the local SUSY-breaking process described above necessarily unifies the regular spacetime symmetries found in general relativity with the local SUSY transformations.

PART II:

EXPERIMENTAL APPARATUS

The world has changed far more in the past 100 years than in any other century in history. The reason is not political or economic but technological - technologies that flowed directly from advances in basic science.

Stephen Hawking

Chapter 3

CERN and the LHC



Fig. 3.1 The CERN logo. Image retrieved from Ref. [30].

The European Centre for Nuclear Research (French: Organisation Européenne pour la Recherche Nucléaire), commonly known as “CERN”, is a European research organisation that operates the largest particle physics laboratory in the world. It was founded in 1954 by 12 European countries: the Federal Republic of Germany, the Kingdom of Belgium, the Kingdom of Denmark, the French Republic, the Hellenic Republic, the Italian Republic, the Kingdom of Norway, the Kingdom of the Netherlands, the United Kingdom of Great Britain and Northern Ireland, the Kingdom of Sweden, the Swiss Confederation and the Socialist Federal Republic of Yugoslavia [31]. Since then it has been continuously active and has expanded, with more European countries

becoming member states and some countries (both European and otherwise) becoming associate members.

The main facility straddles the Franco-Swiss border, near the Swiss town of Meyrin. There are also many other smaller facilities in the area which are often located directly above various points on the below-ground accelerators. Figure 3.2 shows an aerial photo of the area around Geneva with an overlay showing where the accelerators and detectors are, while Figure 3.3 shows an aerial photograph of the main facility near Meyrin, Switzerland.




Fig. 3.2 An image showing the locations of major sites, detectors and accelerators overlaid on an aerial photo showing the area around Geneva. Photo taken from the CERN 60 Year Anniversary Celebration [32]. The **RED** circle represents the Proton Synchrotron, the **BLUE** circle represents the Super Proton Synchrotron and the **YELLOW** circle represents the Large Hadron Collider.

The name “CERN” is used to refer to the laboratory and is an acronym; standing for “**C**onseil **E**uropéenne pour la **R**echerche **N**ucléaire” (English: European Council for Nuclear Research). This was given to the provisional council, established in 1952, that oversaw the creation of the organisation and the construction of the main facility.




Fig. 3.3 An aerial photo taken of the main site near Meyrin. Image retrieved from Ref. [33].

This council was dissolved once the full organisation was established in 1954, but the name was kept [31].

CERNs purpose is to provide the infrastructure necessary for high-energy particle physics experiments: accelerators, data retrieval and storage, computing power, etc. Over the laboratories lifetime it has housed some of the largest particle accelerators made by humankind, the latest of these being the Large Hadron Collider. The analysis presented in this thesis draws its data from collisions that the LHC provides, so this chapter will describe the collider, its operations and how they affect data-taking. Figure 3.4 shows a photo taken of the LHC inside the tunnel.

The history of CERN is intimately related to the history of particle physics spanning the last 2 centuries. While an understanding of these histories isn't necessary to understand the analysis presented, it is nonetheless an interesting subject on its own and can provide some additional context for the design of the LHC. Appendix C summarises the history of the field of particle physics before CERN, details the foundation of CERN and ends by listing major achievements in the field between CERNs foundation and the turn of the millennium.




Fig. 3.4 A photo of the LHC in the tunnel.

The large amount of scientific expertise available to the organisation has also allowed it to advance fields outside of particle physics, such as engineering and computer science. Probably the most important of these advances was the creation of the internet by a CERN scientist named Tim Berners-Lee, which was originally made to satisfy the information sharing needs of the scientists at CERN [34]. Even outside the realm of scientific research, CERN has a presence, such as having observer status at the United Nations [35].

3.1 Large Hadron Collider

The **Large Hadron Collider** (or “LHC”) is the largest and most powerful particle accelerator in the world, as well as the largest machine ever constructed [36]. The sheer size and complexity of the accelerator demands that the rest of this chapter be used to describe it and its operations. The information presented here is largely drawn from volume 1 the original LHC technical design report [37].

The first serious consideration towards the construction of a high-energy hadron collider at CERN came in 1984 at a workshop jointly held at CERN and at Lausanne, Switzerland [38]. There, attendees discussed the possibility of placing a high energy hadron collider in the tunnel that was, at the time, being dug to house LEP. Early on, a popular idea was to have LEP and the LHC occupy the same tunnel, with the LHC

placed on top of LEP. This idea was eventually deemed infeasible and was scrapped. However, this meant that once LEP reached the end of its lifetime in the year 2000 [39] the LHC was a natural choice to replace it. After all, LEP had a 27 km long tunnel dug 100 m deep underground to accommodate it, and it would be wasteful to just abandon it. The construction itself began immediately after the LEP collider was completely removed from the tunnel and would be completed by mid-2008.

The Large Hadron Collider is a synchrotron, albeit the largest in existence, and uses the same method to accelerate charged particles that all other synchrotrons use. Charged particles are accelerated using electric fields while magnetic fields are used to steer the beam so that it maintains a circular path. Magnetic fields are ideal for this due to the nature of the magnetic force (Equation 3.1).

$$\vec{F}_{\text{Magnetic}} = q\vec{v} \times \vec{B} \quad (3.1)$$

By applying a uniform magnetic field in the vertical direction, any particle passing through it will feel a force perpendicular to its direction of motion. This will lead to the particle taking a circular path in the horizontal plane, with the magnetic force acting as the centripetal force. As the particle orbits, it repeatedly passes through the electric field, gaining energy every time it does so. Equating the magnetic and centripetal forces (Equation 3.2) shows that if the mass and charge of the particles in the beam remain constant then in order to keep the radius of orbit constant then the magnetic field strength must increase with the speed of the particle beam. This is exactly what happens in synchrotrons and is what separates them from cyclotrons. In a cyclotron, the magnetic field strength is held constant and so the radius of the orbit of the particle beam increases with its speed until the beam is fired. In synchrotrons, the magnetic field strength is *synchronised* with the speed of the beam, hence the name.

$$|\vec{F}_{\text{Centripetal}}| = \frac{m|\vec{v}|^2}{r} = q|\vec{v}||\vec{B}| = |\vec{F}_{\text{Magnetic}}| \quad (3.2)$$

While the previous paragraph broadly describes how the LHC works, there are a few things that the LHC does that the typical synchrotron doesn't necessarily do. The LHC has two particle beams, which orbit around the LHC in different directions and are steered into each other to create collisions. This is in contrast to many earlier synchrotrons which accelerate a beam and fire it at a stationary target. This is done simply because two high energy beams produce higher energy collisions than a beam

fired at a stationary target. Each of the two beams has a devoted electric field system to accelerate it, which are discussed in detail later in this Subsection 3.1.7.

Finally, the LHC doesn't accelerate particles from a stand-still by itself. The CERN Accelerator Complex contains a total of 11 accelerators that are either currently in use or were used in the past (Figure 3.5 shows how all these accelerators are connected). Of these, there are 5 that make up the series of accelerators that accelerate particles from a stand-still to 99% the speed of light. These are "Linear Accelerator 2", the "Proton Synchrotron Booster", the "Proton Synchrotron", the "Super Proton Synchrotron" and the LHC. The **L**INear **A**Ccelerator **2** (known as "LINAC2") is, unsurprisingly, a linear accelerator that is the starting point for the protons that are used in experiments at CERN. The proton source is a bottle of hydrogen at one end of LINAC2. These hydrogen atoms are subject to an electric field that strips them of their electrons before they are put into the accelerator. LINAC2 then accelerates them to an energy of 50 MeV before sending them into the Proton Synchrotron Booster [40]. The **P**roton **S**ynchrotron **B**ooster (or "PSB") is the first synchrotron in the series and was installed in 1972. It receives protons from LINAC2 and accelerates them to an energy of 1.4 GeV before passes them through to the **P**roton **S**ynchrotron (or "PS") [41]. However, this wasn't always the case. The PS was first used in 1954. LINAC2 fed it 50 MeV protons directly, which it then accelerated to an energy of 25 GeV [42]. This configuration had a problem in that the low injection energy meant that there was a limit to how many protons could be accepted by the PS. With the construction of the PSB, the increased injection energy means that the PS can now accept 100 times more protons than it could without the PSB [41]. For a time, the PS was the most powerful particle accelerator in the world and was the workhorse of CERN's particle physics program, providing the highest energy collisions used in the forefront of research at the time. Today, the PS accelerates the protons before passing them to the penultimate accelerator: the **S**uper **P**roton **S**ynchrotron (or "SPS"). The SPS succeeded the PS as the workhorse of CERN's particle physics program, providing high energy collisions for research from when it was switched on in 1976 until its retirement in 1991. In this time it served as a proton-proton collider, a proton-antiproton collider and an electron-positron accelerator (where it injected into LEP the way it injects into the LHC today). Today, the SPS accelerates protons to an energy of 450 GeV before injecting them into the LHC [43], where they have been accelerated to energies of 5, 7, 8 and most-recently 13 TeV.




Fig. 3.5 Schematic diagram showing how the various accelerators at CERN connect to each other. Image retrieved from Ref. [44].

3.1.1 Accelerator Geography


The LHC is divided into 8 arcs and 8 insertions. The arcs are the long, curved sections of the LHC. Each arc is approximately 3 km in length and spans almost $\frac{\pi}{4}$ radians around the collider. Combined the arcs make up approximately 24 km of the approximately 27 km long circumference of the LHC. The arcs are where the majority of the vacuum, cryogenic and magnet systems of the LHC are located. By contrast, the insertions are shorter, straight sections of the collider. They are much shorter than the arcs, being approximately 528 m in length. An insertion can contain either one of CERNs 4 main experiments (ATLAS, CMS, ALICE or LHCb) or utility equipment related to maintaining the beam (beam insertion points, beam dump points, etc) or sometimes even one of each.

An insertion can be divided further into three or four subsystems. At each end of an insertion there are dispersion suppressors. These are systems of magnets that alter the beam so that it behaves correctly for the section it is entering (either leaving an arc and entering an insertion or vice-versa). Moving closer towards the centre of the insertion, the Dispersion Suppressors give way to the matching sections, a different system of magnets that is used to “aim” the beam so that it enters the correct beam pipe once it leaves the insertion. In insertions that contain experiments, there is a third pair of magnet systems called the inner triplets which focus the beam to maximise the number of collisions for the experiments. Finally, at the centre of an insertion is an interaction point. This is where the beams would collide if the insertion held an experiment. If the insertion in question doesn’t contain an experiment the beams simply pass by each other and exit the insertion.

Taken together, the interaction point, the inner triplets (if they are present) and the matching sections form what is known as a long straight section. The section of the LHC that begins at one interaction point and ends at the next one is known as a sector. Conversely, the section of the LHC that begins in the centre of one arc and ends at the centre of the next arc is called an octant. A schematic diagram of these different types of subdivisions as well as the subsystems mentioned in the previous paragraph is shown in Figure 3.6.

3.1.2 Magnet Systems

There are a total of 9,594 magnet systems in the LHC [37]. These magnet systems perform a variety of tasks towards the control and maintenance of the beam in the LHC and come in a variety of forms (dipoles, quadrupoles, sextupoles, octupoles and



* Inner triplets are only present in insertions with physics experiments

Fig. 3.6 Schematic diagram of a subdivision of the LHC. Image retrieved from Ref. [37].

decapoles). While documenting every single type of magnet in use at the LHC is far beyond the scope of this thesis, it is worth exploring a few of the most important types of magnets used by the LHC.

Probably the most important magnet in use at the LHC is called the main dipole. These are dipole magnets that produce the uniform, vertical magnetic fields that steer the beam in a circle. Each main dipole magnet consists of two sets of coiled superconducting cable. The loops of the coil are rectangular with dimensions 15 m x 15 mm. The 15 m length is parallel to the beam pipe while the 15 mm width points radially outward from the centre of the beam pipe. Figure 3.7 shows the distribution of conductors in the top-right quarter of a main dipole magnet.

In order to steer a beam of particles moving at relativistic speeds, the main dipoles must create enormously strong magnetic fields, far larger than conventional magnets can generate. Thus, superconducting magnets are used. A “superconductor” is a material which exhibits “superconductivity”, the phenomenon where a materials electrical resistance fall to zero when it is cooled to temperatures near absolute zero. Superconductors can then have enormous currents flowing through them and can generate similarly enormous magnetic fields.

The magnets used in the LHC must produce the highest strength field possible while using tried-and-tested technology in the limited space of the LEP tunnel. To this end the LHC uses NbTi magnets, the cutting edge of superconductor technology at the time the LHC was built. These magnets are typically operated by cooling them with supercritical helium to a temperature of 4.2 K having them produce a 5 T magnetic field. In the LHC the magnets are pushed even further, being cooled to 1.9 K (0.8 K




Fig. 3.7 Distribution of the conductor in the main dipole. This distribution is only of the top-right section, and is mirrored in both the horizontal and vertical axes. The x -axis is in millimetres. Image retrieved from Ref. [37].

colder than intergalactic space [45]) with superfluid helium and producing a maximum field of 8.5 T.

The main dipoles are spread throughout the accelerator, with each beam pipe having dipoles of a specific orientation (ie, the dipole produces an upward-pointing or downward-pointing magnetic field). Figure 3.8 shows a magnetic flux diagram for a cross-section of the accelerator structure.

The LHC uses two kinds of particle beams, proton beams and lead ion beams. With these, it performs proton-proton collisions, proton- Pb collisions and Pb - Pb collisions. Both the protons (p^+) and the lead ions ($^{208}_{82}Pb^+$) have positive charges of the same magnitude. While this makes steering the beam easier a problem arises because all of these particles naturally repel each other through the electric force and thus beam dissipates over time. As can be seen in Subsection 3.1.6, the beam pipe is actually fairly small and so dissipation can quickly lead to particles colliding with the beam pipe walls. Additionally, the searches for new physics that the LHC is being used for requires both high energy collisions and a high rate of collisions. The second point in particular means that an intense beam with a very high number of particles in the smallest cross-sectional area possible is necessary.

Both of the above points mean that the beam must be focused as it travels along the beam pipe and as it enters the experiments. This focusing is performed by quadrupole magnets, called main quadrupoles (or less commonly, lattice quadrupoles). The quadrupole magnets consist of four coils located at the 12, 3, 6 and 9 o'clock positions around the beam pipe. Likewise, this means there are four poles, located at the 1:30, 4:30, 7:30 and 10:30 positions around the beam pipe, with the north poles





Fig. 3.8 Main dipole magnetic flux diagram. Image retrieved from Ref. [37].

opposite each other and the south poles opposite each other. This means that there are two possible layouts for the poles: the first is with north poles on the top-left and bottom-right and the south poles on the top-right and bottom-left (for simplicity, this will be referred to as the “right-handed layout”), the second is with the north poles on the top-right and bottom-left and the south poles on the top-left and bottom right (the “left-handed layout”). This design was deliberately chosen in order to exploit the innate physics of the beam. Figure 3.9a shows an example of the magnetic field lines and forces in a right-handed layout. If a particle moves away from the centre in the x -direction, it experiences a restoring force that pushes the particle back into the centre. Thus, the beam is focused in the x -direction. However, the opposite happens in the y -direction. If a particle moves away from the centre in the y -direction it will be pushed even further away. This “defocusing” in the y -direction may seem contradictory to the stated use of the quadrupole magnets, but this is only the case for the right-handed layout. The left-handed layout has the opposite effect, where the beam is focused in the y -direction and defocused in the x -direction. The designers of



(a) A diagram showing the magnetic field (blue lines) and magnetic force (red arrows) acting on a particle moving out of the page. Image retrieved from Ref. [47].
 (b) Two quadrupole magnets laid out in series. The black arrows show the magnetic field lines while the red arrows show the direction of the magnetic force. Image retrieved from Ref. [48].

Fig. 3.9 Quadrupole magnet fields and layouts

the LHC make use of both of these effects, alternating between left- and right-handed layouts so that the beam is focused and defocused in both directions. This is called “strong focusing” and counter-intuitively this process has the net effect of focusing the beam much more strongly than just regularly focusing the beam on its own [46]. This layout of quadrupole magnets can be seen in Figure 3.9b.

Together, the main dipoles and main quadrupoles make up around 15% of the magnet systems in use at the LHC. The remaining 85% come in many different forms (dipoles, quadrupoles, sextupoles, octupoles and decapoles) and are used to solve a range of problems such as linear and non-linear imperfections, chromaticity, collective effects.

3.1.3 Cryogenic System

The cryogenic system of the LHC is responsible for cooling all of the superconducting magnets. These magnets are placed in containers known as “cold-masses”. These containers encapsulate everything that must be cooled by the cryogenic system; the superconducting magnet coils, the iron yoke that provides the structure for the magnets and the helium vessel that contains the superfluid helium.

The refrigeration system was originally built for use by LEP, and had 4 refrigeration plants located around the tunnel. The LHC has a higher demand for cryogenic systems

and so an additional 4 plants had to be constructed, bringing the total number of plants to 8. Each of these eight plants provide cooling for a 3.3 km long section of the accelerator ring.

The choice of superfluid helium as the coolant was motivated by a combination of the useful properties of the superfluid and constraints in how the LHC must be constructed. The LHC magnets must be cooled to temperatures lower than the magnets in any other accelerator. The magnets are immersed in superfluid helium at a temperature of 1.9 K and a pressure of 0.13 MPa (1.3 bar). In these conditions the superfluid helium has several very useful properties:

- **Low Bulk Viscosity:** Superfluids are defined by their very low viscosity. The lower the viscosity of a fluid, the less friction there is between the atoms in the fluid. Superfluids typically have close to zero viscosity, meaning that the superfluid flows without any friction between its atoms or its surroundings. This means that the superfluid is able to permeate to the heart of the windings of the magnets, where fluids with higher viscosity might not.
- **Large Specific Heat:** Superfluid helium has a very high specific heat, meaning that more energy is required to increase its temperature a specific amount than other fluids.
- **Maximum Thermal Conductivity:** The thermal conductivity of helium has a maximum at 1.9 K. High thermal conductivity means that any heat that it absorbs quickly spreads throughout the entirety of the fluid.

All three of these properties mean that superfluid helium is very good at cooling objects down to very low temperatures and keeping them at those temperatures.

The cryogenic system must also operate under several constraints. The first is that the cold-masses must be designed so that the accelerator can still fit in the tunnel. The refrigeration system must be able to fill and cool the entire LHC cold-mass (37×10^6 kg) within 15 days while avoiding thermal gradients greater than 75 K. Additionally, the cryogenic system must be able to warm-up small strings of cold-masses to allow for the repair or replacement of specific cold-masses.

3.1.4 Vacuum System

The LHC has the particularity of having *three* vacuum systems:

- Beam vacuum

- Insulation vacuum for the cryomagnets
- Insulation vacuum for the Helium Distribution Line (QRL)

The beam vacuum handles the evacuation of the beam pipe, while the two insulation vacuums are for evacuating the cryomagnets and the helium distribution line. Naturally, the vacuums have different properties and are measured using different scales. The insulation vacuums are typically measured in either pascals (Pa) or bars (bar). The two insulation vacuums don't need to be particularly intense. At room temperature the vacuum doesn't need to be stronger than 10 Pa (0.1 mbar), while at cryogenic temperatures it only needs to be about 10^{-4} Pa (10^{-6} mbar). The beam vacuum is quantified using “equivalent hydrogen gas density”: $H_2 \text{ m}^{-3}$. A gas with $n H_2 \text{ m}^{-3}$ has the same pressure as n hydrogen gas molecules at the same temperature. The beam vacuum needs to be much more intense than the other two vacuums. In the arcs and dispersion suppressors it needs to be $10^{15} H_2 \text{ m}^{-3}$. This value was chosen in order to give the beam the lifetime it requires (approximately 100 hours) and to minimise the heat that the beam pipe absorbs from relativistic beam particles colliding with ambient gas molecules and knocking them into the beam pipe. Since the area around the beam pipe is cooled to below 2 K any collisions like these that heat up the beam pipe are unacceptable and must be minimised. In the insertions where the experiments are present it has to be even more intense: $10^{13} H_2 \text{ m}^{-3}$, in order to reduce the background for the experiments.

The vacuum systems are both divided into manageable sections. The beam vacuum is divided using “sector valves” and the insulation vacuum are divided using vacuum barriers. The insulation vacuum sectors are different lengths, with the QRLs being 428 m long and the cryomagnet sections being 214 m long. The distance between sector valves is much longer, approximately 2900 m in one area of the accelerator.

3.1.5 Powering

The LHC has almost 10,000 magnet systems. These systems are all electromagnets, either superconductors or the “regular” copper electromagnets, and require power to function. To provide this power, the LHC has a network of 1612 electrical circuits, each of which provide power to a “family” of magnets.

The LHC is divided into eight sectors (these are the “sectors” defined in Subsection 3.1.1) and these sectors are all powered independently of each other. The sectors are further divided into 28 subsectors, of which there are four different types:

- 8 arc-powering subsectors which power all of the magnets in the arc cryostats

- 8 inner-triplet-powering subsectors which power all of the magnets in the arc cryostats
- 12 matching-section-powering subsectors which power the magnets in the smaller cryostats in the matching sections.
- 7 powering subsectors that power the normal conducting magnets around the interaction points

Each of the eight arc-powering subsectors contain one electrical circuit to power the main dipoles (designated “RB” in the technical documentation). In this circuit all of the dipole magnets are connected in series. There are two other circuits (designated “RQF” and “RQD”) that power the focusing and defocusing magnets, respectively. Each of the focusing magnets is connected in series in the RQF circuit, and each of the defocusing magnets is connected in series in the RQD circuit.

The reason that the LHC takes such a modular approach towards its circuitry has to do with keeping the magnet systems safe. The energy present in any of the eight dipole circuits is on the order of 1.22 GJ, which is enough to damage the magnets should the magnets stop being superconducting. The reason that the circuitry is so modular is to avoid depositing enough energy in a given circuit to do significant damage to any of the magnets. For this reason many other circuits, like the “RQF” and “RQD” have much less energy in them than the dipole circuit, around 20 MJ.

3.1.6 Accelerator Structure

The magnet, cryogenic, vacuum and power systems described in the previous four sections are all contained in a single cylindrical structure. This section will briefly describe these structures as well as how they are connected to create the LHC ring.

At the centre of the LHC is the beam tube. The beam tube is a continuous circular tunnel, a few centimetres in diameter, which serves as the passage that the particle beam travels through. There are two beam tubes, each having a single particle beam travelling through it. The two particle beams travel in the opposite direction to the other. The two beam tubes merge into one in every insertion. In the insertions containing experiments, this is to allow the beams to collide. In the other insertions, the beams pass by each other and enter into the opposite beam pipe from which they left. The vacuum in the beam tube is the most intense in the LHC.

The walls of the beam tube are made of an austenitic steel alloy. This steel extends from the beam tube wall a few centimetres in all directions and contains the superconducting magnets within its mass. As such, this mass of steel (known as the


“austenitic steel tubes” or “cold bore tubes” in the technical documentation) serves as both the vacuum vessel and the medium through which the superfluid helium flows to get to the superconducting magnets (while also preventing the same superfluid helium from leaking into the beam vacuum). The alloy is non-magnetic so it is unaffected by even the tremendous magnetic fields produced by the LHCs superconducting magnets. The alloy is also seamless, to ensure that both the material is as leak-tight as possible and the magnetic permeability throughout the material is as uniform as possible. Within this mass of steel is where the magnets are located. The magnets are laid out parallel to the beam tube, with where they’re placed being dependent on what kind of magnet (dipole, quadrupole, etc).

Surrounding the cold bore tubes is the “iron yoke”. This structure is a solid mass of iron that holds the cold bore tubes and magnets in place. Such a robust structure is necessary to hold the magnet system together since the enormous fields produced by the superconducting magnets exert enormous forces on the magnets themselves, and the system would fly apart if the yoke weren’t present.


Surrounding the iron yoke is the helium vessel, within the boundaries of which all of the superfluid helium is contained. The outer layer of this section is thermally insulated, so that superfluid helium stays as cold as possible for as long as possible. Taken together, the beam tube, cold bore tubes, magnet system, iron yoke and helium vessel form the “cold-mass”, the set of systems of the LHC that are cooled using superfluid helium.

Outside the helium vessel is another vacuum vessel. The vacuum here is actually used for insulation (the “insulation vacuum for the cryomagnets” mentioned in Subsection 3.1.4), which vacuums are exceptionally good at provided radiative heat isn’t a big factor. This vacuum vessel is surrounded by yet another layer of insulation.

Finally, outside this layer is the outer layer of the LHC, the blue metal shell familiar from photos. The structure described above is fairly general, but there are small differences in structure based on what magnet is used. Thus, the LHC is divided into cylinders containing the above structures whose lengths are similar to the magnets contained within them. For example, the main dipoles are contained in cylinders called “cryodipoles”. These are roughly 15 m long (same as the main dipole magnets) and are specially designed to hold the main dipoles. By contrast, the main quadrupole magnets are installed in “short straight section cryostats”. The “**Short Straight Sections**” (or “SSS”) are short, straight sections of the accelerator. They are found in the arcs and the cryostats are around 5 m long. Figures 3.10a and 3.10b show cross-sections of the cryodipoles and SSS cryostats, respectively.



(a) Cryodipole cross-section



(b) Cryoquadrupole cross-section

Fig. 3.10 Cross-sections of the cryodipole (a) and cryoquadrupole (b). Images retrieved from Ref. [37].

The accelerator is made up of many cylinders linked together to form a complete circle. The majority of the accelerators circumference is made up of cryodipoles and SSS cryostats, with smaller cryostats containing the rest of the magnet systems. These cylinders are linked end-to-end around the length of the arcs and to the static beam tube in the insertions. The cryostats are installed in a certain order in so-called “FODO-Cells”. The acronym “FODO” stands for “Focusing, O (“Oh”, zero, as in nothing happens), Defocusing, O”. The “focusing” refers to a focusing quadrupole. The “O” means “zero”, that is, nothing happening. This isn’t quite accurate. It isn’t that “nothing” happens, rather there is no focusing/defocusing happening. There are three cryodipoles placed between the focusing and defocusing quadrupoles as shown in Figure 3.11. The “defocusing” refers to a defocusing quadrupole magnet. Figure 3.11 shows a schematic diagram of a FODO-Cell. The FODO-Cells are 106.9 m long and there are 23 of them in each arc. The FODO-Cells can be further divided into 54.45 m long “half-cells” which contain three cryodipoles and an SSS cryostat.




Fig. 3.11 Schematic of a “FODO Cell” showing the layout of the LHC magnets. Image retrieved from Ref. [37].

3.1.7 Insertions

There are 8 insertions around the accelerator, designated as “Point 1” through “Point 8”. The purpose of and equipment within each insertion are detailed below.

Points 1 and 5 are home to the CERNs two multi-purpose experiments: ATLAS (**A** **T**oroidal **L**HC **A**pparatu**S**) and CMS (**C**ompact **M**uon **S**olenoid). Low- β collisions

occur at these insertions and the two experiments attempt to collect as much information about the collisions as they can. Point 1 is also directly below the main CERN facility near Meyrin in Switzerland.

Points 2 and 8 have a dual purpose: Experiment and Beam Injection. ALICE (**A** **L**arge **I**on **C**ollider **E**xperiment) makes use of the lead ion collisions that CERN provides at the end of every run year [49], and is located at Point 2. The LHCb (**LHC** **b**eauty) experiment specialises in studying the “bottom” or “beauty” quark [50], and is located at Point 8. Additionally, Points 2 and 8 are where the proton beam is injected once the SPS accelerates it to 450 GeV.

Point 4 is where the Radio-Frequency (“RF”) Accelerating System resides. This is where the particles in the beam are accelerated by an alternating electric field. There are two RF systems, one for each beam, and each of these systems has four RF cavities. Each RF cavity have a voltage across it of 4 MV, with the entire RF system providing 16 MV. The RF cavities have a frequency of 400 MHz, slightly larger than the 352 MHz of the LEP RF System, but close enough so that the LHC RF System can be made with the same technology.

Point 6 is where the beam dump is located. The function here is fairly straightforward. If the beam is found to be wanting in some way, there are a pair of magnets that deflect the beam vertically into a pair of concrete blocks. These blocks absorb the beam so that it doesn’t harm anything.

Finally, Points 3 and 7 house the beam maintenance apparatus. Figure 3.12 shows a schematic diagram of the LHC, with the octants and insertions labelled.

3.1.8 Accelerator Operations

This section has already briefly discussed how protons are accelerated up to their full energy by the accelerators at CERN. This subsection will explore how the LHC accelerates and collides particle beams more deeply, before detailing how the colliders performance is measured.

When the two proton beams enter the LHC from the SPS, they are said to have an energy of 450 GeV. What this really means is that each of the protons in the two beams have a mass-energy¹ of 450 GeV. As the LHC accelerates the beams, this energy increases. The beams themselves aren’t uniformly distributed streams of protons, but instead are made up of “bunches”. These bunches are groups of protons, with each bunch separated from the bunches before and after it by empty space. Each bunch is

¹As in, $E = \gamma mc^2$.




Fig. 3.12 Schematic diagram showing the layout of the LHC, numbering the Octants. Beam 1 circulates clockwise and Beam 2 counter-clockwise. Image retrieved from [37]

approximately $16 \mu\text{m}$ in width, several centimetres long and contains $\sim 10^{11}$ protons. Each beam contains 2808 bunches.

When accelerated to their full energy, each beam orbits the LHC 40 million times a second. Thus, when colliding the accelerator has a collision rate of 40 MHz. The distance between the bunches is measured in the time interval between them, so at full power the “bunch-spacing” is said to be $t = \frac{1}{f} = \frac{1}{40\text{MHz}} = 25\text{ns}$. Relative to the width of each bunch, protons are effectively point particles. This means that when two bunches “collide” they mostly just pass through each other, with only a handful of proton-proton interactions occurring, approximately 25. When a proton-proton collision occurs it is known as an “event” and the number of events per bunch crossing is called “pile-up”.

One of the features of proton-proton colliders is that whenever two protons collide, they mostly undergo QCD interactions. This means that literally any particle in the Standard Model can be the product of a proton-proton collision. It makes sense to list all of the possible products of a proton-proton collision, ordered by their cross-section, but there are examples of particles that the detector can't feasibly detect. Both gluons and quarks hadronise before they can get too far from the interaction point, producing a spray of more stable particles which enter the detector. It's impossible for a detector to separate out the different particles in such a spray, or to tell what the source particle was, so instead it treats it as a single object called a "jet" (more on this in Chapter 5). The vector bosons $W^\pm$ and Z^0 have a short lifetime so they decay into fermions, which are either picked up by the detector directly (leptons), as a jet (quarks) or not at all (neutrinos). However, in contrast to the jets above it is often possible to tell whether a set of leptons/jets came from a vector boson.

A corollary to this is that it isn't really possible to classify collisions based on the Standard Model interaction vertices present. Instead, they are classified by their intermediate decay products. This is done in Figure 3.13, which shows the production cross-sections of all the processes that occur in the proton-proton collisions in the LHC, where the processes are classified based on a combination of their intermediate and final decay products.

Finally, in order to measure the performance of the collider, it's necessary to have a way of measuring the "output" of the collider. To this end, the "luminosity" is defined (Equation 3.3).

$$L = \frac{(N_b)^2 n_b f_{Rev} \gamma}{4\pi \epsilon_n \beta^*} F \quad (3.3)$$

Here,

- N_b is the number of protons per bunch
- n_b is the number of bunches
- f_{Rev} is the revolution frequency (given as 40 MHz above)
- γ is the familiar Lorentz factor
- ϵ_n is the normalised transverse beam emittance
- β^* is the β -function of the beam at the collision point
- F is a geometric reduction factor for when there is a small angle between the beams (ie, they don't collide *exactly* head on)

The luminosity is calculated from these parameters. It is used in one of two forms: "instantaneous luminosity" and "integrated luminosity". Instantaneous luminosity





Fig. 3.13 A histogram showing the production cross-sections of every possible product of the proton-proton collisions that occur in the LHC. Image retrieved from Ref. [51].

is calculated during a run and is directly proportional to the collision rate. The nominal integrated luminosity of the LHC is $L = 10^{34} \text{ cm}^{-2}\text{s}^{-1}$. Integrated luminosity is simply the instantaneous luminosity integrated over the running time of the LHC. The integrated luminosity is a measure of the total “output” of the collider and is proportional to the total number of collisions.

3.1.9 LHC Timelines - History and Future Plans

- **LEP Shut Down:** The year 2000 is when the LEP Accelerator sees its final shut down and the process of removing it from the tunnel begins. The Large Hadron Collider is to be constructed in its place[39].
- **Inaugural Tests:** At 10:28am on the 10th of September, 2008 the LHC successfully accelerates its first beam, with the beam orbiting in a clockwise direction. This milestone would mark the beginning of the LHCs influential, if colourful, operating history.

- **Magnet Quench Incident:** Once the LHC had begun operations it would not be long before it would be rendered inoperable again. On the 19th of September, 2008, within 2 weeks of the inaugural tests, several of the superconducting magnets quenched during powering tests. As mentioned earlier in this chapter, superconducting magnets operate at very low temperatures with negligible resistances and very large currents. A “Magnet Quench” is when a section of a magnet spontaneously develops a non-zero resistance which leads to the current flowing through that area quickly heating the magnet anywhere between a few dozen and a hundred Kelvin. This occurred at the LHC when the magnets were being slowly powered up to their operating current of 9.3 kA. At 11:18:36 CEST on the 19th of September a resistive zone developed in a magnet between a dipole and a quadrupole, which was first detected as a voltage of approximately 300 mV. After 0.39 s, the resistive voltage had grown to 1 V and the system initiated a Fast Power Abort, aborting the test after 0.46 s. Within the first second of the resistance developing, an electrical arc occurred which punctured the helium vessel, leading to the superfluid helium heating up, returning to a gaseous phase and leaking into the insulation vacuum. This violent expansion of gas knocked the cryostat from its supports and in some locations tore its anchors out of the *concrete floor of the tunnel* [52]. This led to an immediate loss of almost 2 tonnes of helium, but since repair crews couldn’t be sent into the tunnel until it was shown to be safe, an additional 4 tonnes of helium was lost. Large sections of the superconducting magnets were damaged, requiring their transport to the surface for inspection and repair and the beam pipe was contaminated by soot from the initial electrical arc [53]. Prior to this incident, the LHC was set to operate at 10 TeV by the end of that year. The repairs to the accelerator that were required to bring it back up to full operational capacity occupied the rest of 2008 and much of 2009.
- On the 30th of April, 2009, the final magnet (a quadrupole) is repaired and is lowered into the accelerator tunnel [54]. Figure 3.14 shows the damage to the beam pipe
- On the 20th of November, 2009, after more than a year of repairs, the LHC resumes operations. The accident has had lasting effects, and the LHC is only able to operate at a centre-of-mass energy of 7 TeV.
- On the 30th of November, 2009, not even a fortnight after being fully reactivated, the LHC breaks the record for *Worlds Highest Energy Particle Accelerator*. The two proton beams were accelerated to 1.18 TeV per beam ($\sqrt{s} = 2.36$ TeV),



(a) External Damage

(b) Internal Damage

Fig. 3.14 Damage to the Beam Pipe from the Magnet Quench.

beating the previous record holder, the Tevatron at FermiLab, who had 0.98 TeV per beam ($\sqrt{s} = 1.96$ TeV) [55].

- The 2011 run of proton-proton collisions began on the 13th of March, 2011.
- The LHC breaks yet another record, this time for achieving the highest luminosity at a hadron collider, with a Luminosity of $4.67 \times 10^{32} \text{ cm}^{-2}\text{s}^{-1}$. The previous record holder was, again, the Tevatron at FermiLab with a luminosity of $4.024 \times 10^{32} \text{ cm}^{-2}\text{s}^{-1}$ [56].
- The 2011 run of lead ions (and the first lead ion run) occurs in November of 2011.
- 2012 saw another proton-proton run as well as another lead ion run.
- The final proton-proton and lead ion runs occur in 2013. Having reached a luminosity of $7.7 \times 10^{33} \text{ cm}^{-2}\text{s}^{-1}$ and with $\sqrt{s} = 8$ TeV, the accelerator enters what was known as a “long shut-down”. This long shut-down would see the LHC go through maintenance so that it could reach energies even higher than before.
- April 2015 would see the beginning of run 2. The second running period of the LHC would begin in April of 2015 and last until November of 2018. All four runs were at $\sqrt{s} = 13$ TeV.

Chapter 4

The ATLAS Detector

The ATLAS (**A** **T**oroidal **L**HC **A**pparatu**S**) detector is one of the two general-purpose detectors that study the collisions provided by the LHC, the other being the **C**ompact **M**uon **S**olenoid, or “CMS”. These general-purpose detectors need to be able to capture the maximum possible fraction of the products of every collision, so they are designed to cover as big of a solid angle as possible and to be deep enough to capture even the highest energy particles.

The ATLAS detector is roughly cylindrical in shape, with a base diameter of 25 m and a length of 44 m. It has a mass of approximately 7000 tonnes. The detector has a layered design, each layer being designed to detect a specific type of particle. Smaller layers are closer to the centre and are encapsulated by the larger layers. Figure 4.1 shows a cut-away diagram of the detector, showing the layered structure.

Recall that this thesis is presenting a search for new physics, and were that its only end a brief summary of the detector with no special focus would suffice for this chapter. However, a significant part of my work as a Ph.D student has been on the SemiConductor Tracker, a subsection of the ATLAS detector and this chapter has been written to reflect this.

With this in mind, this chapter will be structured as follows. Section 4.1 will discuss the physics motivation behind the design of the detector. Section 4.2 will give a brief overview of the detector in its entirety, listing the different layers within the detector and how they work to detect and identify particles. Of the different parts of the detector described there, only the innermost layers (collectively known as the “Inner Detector”) will be covered in this chapter, done in Section 4.3. The SemiConductor Tracker mentioned above is part of the Inner Detector and as the part of the detector I’ve done extensive work on, this chapter will include a very detailed description of how it works as well as a summary of some of the work I’ve done.




Fig. 4.1 A cut-away diagram of the ATLAS detector. A pair of humans can be seen standing on the left-hand side for scale. Image retrieved from Ref. [57].

Contrast this with any of the other parts of the detector, like the calorimeters or the Muon System. My expertise with these is much more limited than it is with the SemiConductor Tracker and so the description of these parts of the detector will be far less detailed. Section 4.4 will describe the two calorimeter systems that surround the Inner Detector. Section 4.5 will give an overview of the Muon System that surrounds the calorimeters. Section 4.6 will cover the forward detectors and Section 4.7 covers the structural elements of the detector, namely the physical structure, radiation shielding and control.

4.1 Physics Motivation

During the design phase, the design, optimisation and construction of the ATLAS Detector was guided by the ATLAS “physics programme”, a list of experimental goals (usually taking the form of a search for a new particle or a more precise measurement of some Standard Model parameter) that were well motivated by the theoretical physics predictions of the day.

The physics programme was particularly focused on studying the origin of electroweak symmetry breaking, with a major focus being the search for the Higgs boson, which wasn't discovered at the time. As such, the detector was optimised to be able to be sensitive to the Higgs boson over the full range of possible Higgs' masses. In addition to the search for the Higgs, the detector also searched for other phenomena related to EWSB, such as particles predicted by the SUSY and Technicolour theories. Outside of EWSB, the detector was also designed to search for new gauge bosons and evidence for composite quarks and leptons. Investigating CP Violation in b -decays and precision measurements of $W^\pm$ and top quark masses were also important components of the ATLAS physics programme [58].

Each of these goals imposes their own set of requirements on what the ATLAS detector must be able to accomplish:

- The various Higgs boson searches (which study some of the most challenging signatures) were used as the benchmarks for the detector performance. High-resolution electron, muon and photon measurements, excellent secondary vertex detection for taus and b -quarks, high-resolution calorimetry for jets and E_T^{miss} are all essential to be able to explore the full range of possible Higgs Boson masses.
- Searches for SUSY set the benchmarks on the hermeticity¹ and E_T^{miss} capability of the detector, as well as on b -tagging at high-luminosity.
- Searches for new, heavy gauge bosons provided benchmark requirements for high-resolution lepton measurements and charge identification up to p_T values of a few TeV.
- The signatures characteristic of quark compositeness set the requirements for the measurement of very high p_T jets.
- Precision measurements of the $W^\pm$ and top quark masses, gauge boson coupling, CP violation and determination of the Cabbibo-Kobayashi-Maskawa unitarity triangle yield benchmarks that address the need to precisely control the energy scale of jets and leptons, precisely determine secondary vertices, fully reconstruct final states with relatively low p_T particles and trigger on low- p_T leptons.

¹An objects *hermeticity* refers to how *hermetic* (as in, air-tight) it is.

4.2 Detector Overview

This section will give an overview of the ATLAS detector and the coordinate system and nomenclature used to describe it. The information in this section, and in Sections 4.3, 4.4 and 4.5 largely comes from Ref. [59].

The ATLAS detector is made up of almost a dozen “subdetectors”. These are smaller, self-contained detectors that use one kind of detection technology and performs one function. Each subdetector has it’s own dedicated group that runs the detector during data-taking, as well as performing tests and doing maintenance work. Subdetectors can be broadly classified into one of three categories: trackers, calorimeters and forward detectors.

The trackers measure the spatial positions of particles passing through them. Whenever a particle passes through one of the layers of the tracker, its position in that layer is registered as a “hit”. As the particle passes through the multiple layers of the trackers, it leaves multiple hits. These hits are used to recreate the paths that the particles took through the trackers. Trackers are also immersed in a powerful magnetic field, so that charged particles take a curved path through the them. The strength of the magnetic field and the charge-to-mass ratio of the particles are known, so it’s possible to use the curve to calculate the momenta of the particles passing through the tracker. The trackers are the **I**nsertable **B**-**L**ayer (“IBL”), the **P**ixel detector (named for the “pixel” modules that make it up), the **S**emi**C**onductor **T**racker (“SCT”), the **T**ransition **R**adiation **T**racker (“TRT”) and the Muon System.

The calorimeters measure the energies of the particles. Once a particle enters a calorimeter it interacts with the calorimeter, giving a portion of its energy to the calorimeter and creating more particles in the process. The particles created by this process also give some of their energy to the calorimeter and create even more particles. This process repeats until the energy runs out, creating a “shower” inside the calorimeter, where the energy of the initial particle that entered the calorimeter is the same as sum of the energy deposited into it. The calorimeters are the Electromagnetic Calorimeter and the Hadronic Calorimeter.

The forward detectors are a special class of detector, different from the other two. They cover the regions closest to the beam pipe, and are mostly used to measure the luminosity of the collisions. Section 4.6 will cover the forward detectors in full detail, while the rest of this section will focus on the trackers and calorimeters and how they work together to identify what kind of particles enter the detector.

The trackers use a variety of detection technologies but they all have one common limitation: they can only detect charged particles. Of the particles that are produced by the collider only a subset of them are stable enough to enter the detector. These are the electrons, muons, photons and neutrinos as well as the various hadronic objects (protons, neutrons and other hadrons like the pion and kaon). Of these, only the charged particles (electrons, muons, protons, $\pi^\pm$, etc) interact with the trackers, the neutral particles (photons, neutrons, neutrinos, $\pm^0$, etc) pass through undetected.

The two calorimeters use different detection technologies from the trackers and each other and thus also interact with different particles. The Electromagnetic Calorimeter mostly interacts with electrons, photons and pions, which create showers in the calorimeter. The Hadronic Calorimeter interacts with all hadronic objects that enter into it, each of which leaves a shower in the calorimeter. Single hadrons leave compact showers in the Hadronic Calorimeter while the larger groups of hadrons that come from parton showers leave much larger, more complicated areas of activity in the calorimeter.

The calorimeters are encapsulated by the Muon System, the outermost and largest layer of the ATLAS detector. It is a tracker, designed specifically to detect and measure the properties of muons emerging from collisions.

Using all of this information the detector is able to distinguish between different kinds of particles. In practice, this means matching tracks and showers that are likely to have come from the same particle and labelling the particle based on their properties:

- If a track in the trackers is matched to a shower in the Electromagnetic Calorimeter, the particle is likely an electron.
- If a shower in the Electromagnetic Calorimeter isn't able to be matched with any track, the particle is likely a photon.
- If a track in the trackers is matched to a shower in the Hadronic Calorimeter, the particle is likely a proton.
- If a shower in the Hadronic Calorimeter can't be reasonably matched to any track, the particle is probably a neutron.
- Muons are special cases. They mostly produce tracks in the trackers², but they are the only particles that produce tracks in the Muon System.
- Finally, if a particle passes through the detector without interacting with anything, it is treated as missing energy. Neutrinos are the only known particle that does this, although many hypothetical SUSY particles (such as the LSP discussed in Chapter 2) would also behave this way.

²They are understood to also leave some signals in the calorimeters.

This model is complicated by the presence of various other hadronic objects. These objects behave similarly to several of the particles listed above (for instance, the π^0 leaves a similar signal to the photon while the $\pi^\pm$ can mimic electrons and positrons). While the subdetectors themselves employ various methods to differentiate these particles from each other, the process of turning tracks and clusters into fully realised particles is, in general, much more complicated than what has been presented here. This process is called “Reconstruction” and a full explanation of the reconstruction algorithms used by the ATLAS detector is beyond the scope of this thesis, though suffice it to say it also deals with differentiating various types of particles from one-another. However, the particular reconstruction algorithms that are relevant to the analysis presented in this thesis are described in Chapter 5, Section 5.2.

The subdetectors are arranged in a layered structure in the ATLAS detector. At the time of construction the smallest subdetector, the one closest to the beam-pipe, was the Pixel detector. Larger detectors then envelope the smaller ones, with the SCT enveloping the Pixel detector and the TRT enveloping the SCT. These three detectors make up the “Inner Detector”. A superconducting solenoid magnet is wrapped around the Inner Detector, and provides the magnetic field by which charged particles are curved. Around this magnet the calorimeters are placed. The Electromagnetic Calorimeter serves as the inner layer, surrounding the Inner Detector, with the Hadronic Calorimeter surrounding the Electromagnetic Calorimeter. Finally, the Muon System envelopes the Hadronic Calorimeter, and serves as the outer-most layer of detector material. The Muon System itself has several magnet systems embedded within it, which serve the same purpose as the magnets in the Inner Detector. Figure 4.2 shows a cross-sectional diagram of the ATLAS detector, how the various subdetectors are layered and how the various particles interact with it.

The ATLAS detector was constructed in a large open space underground, known as the “ATLAS cavern”, and was completed in 2008. The structure of the detector would remain the same until 2014, when the IBL would be added to the detector. The IBL would become the new inner-most layer of the detector, being placed between the Pixel detector and the beam pipe.

ATLAS uses a special coordinate system to describe both the detector and the positions of particles emerging from collisions. It is summarised here.

The nominal interaction point, the point at the centre of the detector where the primary proton collisions occur, is taken as the origin of the coordinate system. From here the cartesian axes are defined: the positive x -direction points towards the centre of the LHC ring, the positive y -direction points towards the surface (ground-level,




Fig. 4.2 A diagram of a cross-section of the ATLAS detector showing how the different particles interact with different sections of the detector. Image retrieved from Ref. [60].

Geneva) and the positive z -direction points towards Point 8 and the LHCb detector (this side of the detector is known as “Side-A”). Defined like this the coordinate system is right-handed, the z -axis is tangent to the beam pipe at the interaction point and the x, y -plane (known as the “transverse plane”) is perpendicular to the beam-pipe. Figure 4.3 shows this coordinate system relative to the rest of the LHC.

This cartesian coordinate system isn’t very helpful for studying collisions. The structure of the collisions: two coaxial particle beams of almost zero thickness colliding and the detector mostly measuring the solid angle element that the collision products pass through lends itself very well to a polar coordinate system instead. However, the lack of detector coverage near the beam pipe leads to some problems with the θ variable as it approaches 0 or π . Instead of using θ , a new variable, called “pseudorapidity”, is defined: $\eta = -\ln(\tan(\frac{\theta}{2}))$. This variable tends towards $+\infty$ as θ goes to 0 from above, and towards $-\infty$ as θ goes to π from below. Pseudorapidity is only usable




Fig. 4.3 Schematic diagram showing the ATLAS coordinate system relative to the rest of the LHC. A “top-down” view is shown. Image retrieved from Ref. [61].

for massless particles, so for massive particles the equivalent variable is just called “rapidity” and is defined as: $\eta = (\frac{1}{2})\ln(\frac{E+p_z}{E-p_z})$. The coordinates (r, η, ϕ) define the new coordinate system, with the distance in pseudorapidity-azimuthal angle space defined as: $\Delta R = \sqrt{(\Delta\eta)^2 + (\Delta\phi)^2}$.

Finally, there are a few variables that are defined in the transverse plane. For instance, the transverse momentum “ p_T ” of some particle is the projection of the momentum vector of that particles momentum vector onto the transverse plane. The transverse plane refers to the x, y -plane, the plane transverse to the beam pipe.

4.3 Inner Detector

The **Inner Detector** (“ID”) is the combination of the three trackers that are closest to the beam pipe: the Pixel detector, the SCT and the TRT. These trackers are the primary subdetectors that fulfil the precision tracking and resolution requirements that are imposed by the ATLAS physics programme.

These trackers have the layered structure that is shared by the rest of the subdetectors (both trackers and calorimeters). Each subdetector is divided into two sections; the “barrel” region and the “end-cap” region. The barrel is the central region of the subdetector and is made up of multiple layers of detector material arrayed in concentric

cylinders. The end-caps are made up of multiple discs of detector material placed at each end of the barrel. Figure 4.4 shows a cut-away diagram of the ID that shows how the subdetectors are layered. Figure 4.5 then expands on this, showing schematic diagrams of the barrel and end-cap regions which display the precise radial and z -axis positions of the barrel cylinders and end-cap discs, respectively. Finally, Table 4.1 shows some of the parameters of the ID.






Fig. 4.4 Cut-away diagram of the Inner Detector, showing the Pixel detector, SCT and TRT.



(a) Schematic diagram of the ID barrel showing the subdetectors and their radial distances from the beam pipe.



(b) Schematic diagram of the ID barrel and end-cap showing the subdetectors and their radial distances from the beam pipe and their locations along the z -axis.

Fig. 4.5 Schematic diagrams of the Inner Detector.

Table 4.1 Parameters of the Inner Detector.

Item		Radial Extent (mm)	Length (mm)
Beam Pipe		$29 < r < 36$	
Pixel	Overall Envelope	$45.5 < r < 242$	$0 < z < 3092$
3 Cylindrical Layers	Sensitive Barrel	$50.5 < r < 122.5$	$0 < z < 400.5$
2×3 Discs	Sensitive End-Cap	$88.8 < r < 149.6$	$495 < z < 650$
SCT	Overall Envelope	$255 < r < 549$ (Barrel)	$0 < z < 805$
		$251 < r < 610$ (End-Cap)	$810 < z < 2797$
4 Cylindrical Layers	Sensitive Barrel	$299 < r < 514$	$0 < z < 749$
2×9 Discs	Sensitive End-Cap	$275 < r < 560$	$839 < z < 2735$
TRT	Overall Envelope	$554 < r < 1082$ (Barrel)	$0 < z < 780$
		$617 < r < 1106$ (End-Cap)	$827 < z < 2744$
73 Straw Planes	Sensitive Barrel	$563 < r < 1066$	$0 < z < 712$
160 Straw Planes	Sensitive End-Cap	$644 < r < 1004$	$848 < z < 2710$
Overall ID Envelope		$0 < r < 1150$	$0 < z < 3512$

4.3.1 Pixel Detector

The Pixel detector was the innermost subdetector of ATLAS at the time of its construction. The extreme conditions at the LHC impose demanding requirements on the tracking, resolution and vertex measurement capabilities of the ATLAS detector. As the subdetector closest to the interaction point, the Pixel detector has the harshest requirements imposed upon it. To achieve the necessary performance, the Pixel detector was constructed using materials and techniques at the bleeding edge of technology at the time.

The Pixel detector gets its name from the “pixel” sensors that it uses to detect particles. Each pixel sensor is rectangular and minuscule, measuring only $50 \mu\text{m} \times 400 \mu\text{m}$ ($50 \mu\text{m} \times 600 \mu\text{m}$ in the case of the pixels near the front-end chips on the modules). Their active detector material³ is silicon, a material that was chosen primarily for its high radiation hardness and cost efficiency. These sensors are tessellated together to form rectangles of silicon called “sensor tiles” approximately $63.4 \text{ mm} \times 24.4 \text{ mm}$ in size. These sensor tiles are the main component of the Pixel detector “modules”. Each module is a self-contained piece of electronics that handles detecting particles passing through its pixel sensors and sending information about the detections (called “hits”) upstream to off-detector computer systems. Each pixel module is made up of 5 components layered on top of each other. These are:

1. 16 front-end electronics chips, each with 2880 channels

³The “active detector material” is the material that reacts to particles passing through it

2. The “bump-bonds” (made of either In or PbSn) which connect the chips to the pixel sensors
3. The sensor tile itself
4. A flexible polyimide circuit board (flex-hybrid) with a module-control chip glued to it
5. A polyimide pig-tail with Cu lines and a connector (for barrel modules) or a wire micro-cable (for end-cap modules) bonded to the flex-hybrid

Figure 4.6 shows a schematic view of these layers as well as a photograph of a completed pixel module.




Fig. 4.6 Schematic view of a barrel pixel module (top) and a photo of an assembled pixel module (bottom).

Each sensor tile has 46,080 pixel sensors. These sensors are partitioned into groups of 2,880 and the sensors in each group are connected to one of the 16 front-end chips. Each pixel sensor has a corresponding channel through which hit information is passed from the front-end chip to the off-detector data acquisition system (see Chapter 5). The physical position of each pixel sensor is known and documented alongside the identification number of the channel that connects to the sensor. This way, when a

pixel sensor registers a hit, the physical location of the hit can be found by referencing identification number of the sensor that registered it. There are a total of 1,744 pixel modules in the Pixel detector, meaning that the Pixel detector has around 80 million channels in total.


Both the barrel region and the end-cap region of the Pixel detector make use of the same types of modules. The barrel region consists of Pixel modules tiled in three concentric cylinders. Each end-cap consists of three layers of Pixel modules arranged in an annulus.

4.3.2 SemiConductor Tracker

The SemiConductor Tracker is the second tracker in the ID, enveloping the Pixel detector. As mentioned above, I've worked on the SCT a great deal during my Ph.D, and so this subsection will be significantly longer than the subsections describing the other subdetectors. The information here largely comes from Ref. [59], Ref. [62] and Ref. [63] as well as my own work.

The overall goal of the SCT is to create a tracker which has a high spatial accuracy when measuring hits and a temporal accuracy sufficient to determine which bunch crossing a hit originates from. The requirements of the ATLAS physics programme demand that each SCT sensor has an efficiency of 99% (a low amount of false negatives, ie, situations where a particle passes through a strip but isn't registered) and a noise occupancy of 5×10^{-4} (a low amount of false positives, ie, situations where no particle has passed through a strip but the strip has recognised a hit to have occurred, usually misidentifying noise as a particle hit). In addition to recording each hit in the correct time bin, the modules must be able to discriminate between two pulses separated by 50 ns. Finally, the SCT modules are designed to be able to continue operation in spite of the presence of dead channels. This tolerance is specified to be 1% of all channels in a module, or 15 channels. Additionally, a large number of dead channels in series is also unacceptable, so a limit of 7 consecutive dead channels is also imposed.

The SCT barrel consists of four layers of detector material. The physical structure of each barrel layer is provided by a carbon-fibre sheet rolled into a cylinder. All of the barrel modules are mounted on this sheet using carbon-fibre brackets. The modules are mounted in rows of 12, with each row parallel to the beam line. There are 32 rows of modules mounted around the cylinder. Within each row the modules are mounted in alternating vertical orientations, while each row of modules is rotated around the axis parallel to the beam pipe by $\sim 11^\circ$ with respect to the modules in the previous row. Both of these are done in order to allow the modules to overlap with each other,



(a) Module layout in the barrel region (b) Module layout in the end-cap region

Fig. 4.7 A schematic view of the SCT module layout in the barrel (a) and end-cap (b) regions. Both are viewed along the z -axis of the experiment.

ensuring there are no holes in their coverage. All of the SCT barrel modules are identical.

Each SCT end-cap is made up of 9 discs of detector material, with 18 discs used in total. The physical structure of each end-cap is provided by a 8.7 mm disc consisting of an aramid/phenolic honeycomb core covered by $200\mu\text{m}$ carbon-fibre face skins on either side. The end-cap modules are mounted onto this disc. The larger size of the SCT compared to the Pixel detector means that, unlike the Pixel end-caps, the SCT end-caps can't be made using one type of module without either having large amounts of overlap between modules or large gaps between modules. Thus, the SCT end-caps contain no-less than *five* different types of modules arranged in three concentric rings. These modules are labelled with a “W” (short for “wedge”) followed by a pair of numbers denoting where they are on the disc. The inner ring use one type of module, labelled “W12”, the middle ring uses two types of modules, labelled “W21” and “W22” and the outer ring also uses two types of modules, labelled “W31” and “W32”. Figure 4.7 shows a schematic view of the placement of SCT modules in the barrel and end-cap regions.

The requirements placed upon the SCT aren't as stringent as the requirements placed on the Pixel detector and as a result the SCT doesn't need to be at the bleeding-edge of technology. The SCT sensor tiles are made of a silicon crystal semiconductor, from which the name “SemiConductor Tracker” is derived. The bulk of the sensor

tile consists of a 285 μm thick substrate made of over-doped n^+ silicon covered with a much thicker layer of lightly-doped n-type silicon. On the surface of this bulk are placed the p^+ -doped silicon strips that serve as the sensors of the SCT. On top of these strips are aluminium tracks which connect the silicon strips to the front-end electronics so that information can be read out of them. The six different module types use six differently sized silicon strips, each of which are approximately 6 cm long and vary in strip pitch between 55 and 95 μm . Pairs of these strips are daisy-chained together, creating 12 cm long strips of silicon. In the barrel these pairs are arrayed side-by-side, parallel to one another, into rectangular sensor tiles. In the end-caps, each pair is placed at a small angle with respect to its neighbours in trapezoidal modules so that when they are arrayed in a circle the strips point radially outward. Like the barrel modules, the barrel strips are all identical. Similarly, each of the five different types of end-cap modules uses a uniquely sized silicon strip. Table 4.2 lists the sizes of the sizes of each type of module and the sizes of the strips that make them up.

Table 4.2 Silicon Strip Geometries

Module Type	Barrel	End-Cap Inner	End-Cap Middle		Enc-Cap Outer	
Sensor Type	Barrel	W12	W21	W22	W31	W32
Strip Length (mm)	63.960	61.060	65.085	54.435	65.540	57.515
Strip Pitch (μm)	80.0	56.9 - 69.2	69.9 - 83.0	83.4 - 94.2	70.9 - 81.1	81.5 - 90.4
Outer Width (mm)	63.560	55.488	66.130	74.847	64.635	71.814
Inner Width (mm)	63.560	45.735	55.734	66.152	56.475	64.653
Inter-Strip Angle (μRad)	0	207.0	207.0	207.0	161.5	161.5

While the modules might have differently sized strips and differently shaped sensor tiles, there are two things which they all have in common: the number of strips per module and the number of sensor tiles per module. As was mentioned in the previous paragraph, pairs of the 6 cm long silicon strips are daisy-chained together into 12 cm-long strips. Each of these pairs are connected to a single channel, and hits registered on either 6 cm strip is passed to the DAQ system through that same channel. Like the pixel sensors, each strips location can be found by referencing the channel it is connected to. However, there is a minor limitation in using this approach with the silicon strips, namely their size. When a particle passes through a pixel sensor or silicon strip it is impossible to know *where* the particle passed through it. For the minuscule pixel sensors this isn't a problem, since their tiny size means that the position of any hit is known to a fairly high accuracy. For the silicon strips there is a problem since it is impossible to tell where along the 12 cm length of the strips the particle passed through. The solution to this is to use two sensor tiles per module, with one module rotated 40 mRad with respect to the other. A particle passing through the module passes through both sensor tiles, which means that it passes through two strips at a

small angle to each other. The physical location of any hit can be found by taking the channel numbers of the two strips that registered the hit, finding the locations of those strips using the channel numbers and finding the small area where they intersect. The η and ϕ coordinates of this area are then assigned as the location of the hit.

The six types of SCT module are all different shapes and sizes, but they share a common structure. The base of the module is the “baseboard”. All of the other components are attached to the baseboard, and the baseboard is attached to the physical structure of the detector. The baseboard was designed to use as little material as possible while still having a structure strong enough to support the rest of the module, to be able to conduct heat effectively and to be electrically insulating. To achieve this the baseboards are made of a composite material called “VHCPG” (**V**ery **H**igh **C**onductivity **P**yrolytic **G**raphite) sandwiched between two beryllia plates. The VHCPG has the interesting property that it has a very high thermal conductivity in one plane (about $1700 \text{ Wm}^{-1}\text{K}^{-1}$) while it has a relatively low thermal conductivity in the other (around $6 \text{ Wm}^{-1}\text{K}^{-1}$). This allows it to conduct heat away from the front-end electronics very well. The beryllia plates are very electrically insulating, and the resulting structure of the VHCPG sandwiched between two beryllia plates is both strong and light. As was alluded to in the previous paragraph there are two sensor tiles in each module. These are both attached to the baseboard, with one placed on either side of it. Around this assemblage is wrapped a hybrid, an integrated circuit with the front-end electronics embedded into it. The front-end electronics consist of 12 radiation hardened, ABCD3TA **A**pplication **S**pecific **I**ntegrated **C**ircuit (“ASIC”) chips. There are six of these chips placed on either side of each module, with each group of six connected to the sensor tile on their side of the module. Each sensor tile has 768 strips, each of which is connected to a single channel through a chip. Each chip handles 128 strips/channels. In total, each module has 1,536 strips/channels. Figure 4.8 shows a photo SCT barrel module. Figure 4.9 shows a photo of the SCT hybrid before being attached to an SCT module.

The barrel of the SCT consists of four layers of coaxial cylinders, made up of a total of 2,112 modules. Each end-cap consists of nine layers of end-cap discs, made up of a total of 1,976 modules. In total, the SCT is made up of 4,088 modules, each of which attend to 1,536 channels and so the SCT has a total of 6,156,288 channels.

Figure 4.10 shows photos of the SCT barrel and end-cap regions in storage after their construction but before their placement into the detector.

All of the ASIC chips in the SCT are identical, regardless of where they are located in the detector. All of the chips run synchronously with the 40MHz LHC bunch-crossing




Fig. 4.8 A photo of an SCT barrel module. Image retrieved from Ref. [64].






Fig. 4.9 The barrel hybrid, showing the 12 ASIC Chips, the flexible centre section and the off-module connector on the left-hand end. Image retrieved from Ref. [62].



(a) A photo of the completed SCT barrel during a test construction. Image retrieved from Ref. [65].



(b) A photo of a completed SCT end-cap disc in its testbox in NIKHEF. Image retrieved from Ref. [66].

Fig. 4.10 Photo documentation of the SCT subdetector




Fig. 4.11 Components of the ASIC chip. Image retrieved from Ref. [62].

clock. The structure of the chip is shown in Figure 4.11 and is explained in detail below.

When a charged particle passes through a silicon strip it leaves behind a residual charge which the front-end chip “sees” as a voltage (this voltage will be referred to as the “induced voltage” for the rest of this subsection). This induced voltage is the main quantity that the ASIC chips employ, being used for everything from determining whether a hit has occurred to converting it into digital data to be sent to the off-detector readout system.

First, the induced voltage is amplified by a pre-amplifier so that it can be used effectively by the rest of the circuit. Once it is amplified it is then input into a comparator, where it is compared to another voltage known as the “threshold voltage”. The output of this comparator is a binary voltage whose value depends on which voltage is larger, the threshold or the induced. If $V_{\text{Threshold}} > V_{\text{Induced}}$ then the output is 0 V (binary number: 0). If $V_{\text{Threshold}} < V_{\text{Induced}}$ then the output is 5 V (binary number: 1). Working ideally, this comparator would output “1” if a hit was registered and “0” otherwise. However, in practice there are problems with this method that my work on the SCT sought to address, which will be discussed in detail at the end of this subsection. This bit is then run through a mask register before being stored in the pipeline storage buffer.

The pipeline is used to store hit patterns in a chip while waiting for the trigger system to process the events (more on the trigger system in Chapter 5). A pipeline is made up of 12 12-bit shift registers and each channel has its own pipeline. The clock sends pulses to the 12 registers in sequence. Each pulse shifts one bit in a register one

place and subsequent pulses shift subsequent bits. This process means that a bit takes 132 pulses to travel from one end of a register to the other. When the level one trigger accept (“L1A”) arrives at the chip, the three bits at the end of each channel are moved into the readout buffer. These bits correspond to the hit/miss decision made at the bunch crossing selected by the L1A, as well as the previous and subsequent bunch crossings. This sequence of three bits is then filtered using one of four patterns: “hit” where a hit is present in any of the three positions, “level” which matches a hit in the central hit following a miss and “test” which will pass any pattern.

The data coming off the chip is compressed by reading out only those channels matching the specified pattern. The final bit-stream sent by the chip includes the following: the 4-bit chip number, the number of the first channel in a string of consecutive hits, then the hits from all three time bits for all of the consecutive channels.

The 12 chips on each module are arranged in two chains of six, one chain on each side of the module. Each of these chains is attached to two output streams, with one stream being used to read out information from the module and the other serving as a back-up cable. The first chip in each chain is the “master chip” which passes data from one of the chips in the chain to the output cable. The readout begins when the master chip receives the L1A; first sending a header containing the bunch crossing counter and the L1 counter for the event and then the hit data from the master chip itself. After this, the master chip passes a token to the next chip to tell it to send the bit-stream it has. This chip also sends the token to the next chip in the chain until the end of the chain is reached at which point a “trailer” (a sequence of 0’s that can’t appear any other way) is appended to the bit-stream. This design allows dead chips in the readout chain to be bypassed by sending the token to the next-next chip instead. Alternatively, the token can also be passed to the other side of the module in the case where the master chip is the faulty one.

The ABCD3TA chips have several configurable parameters and changing these allows the detector control team to control the behaviour of the chips. This is primarily done during calibration to allow the detector control team to configure the detector response. There are both parameters that affect the whole chip and some that affect single channels. A summary of all of these is shown in Table 4.3.

At the chip level there are six configuration registers. One is used to set the discriminator threshold. Two are used to control the size and timing of the calibration charge injection. Two others are used to control the currents which modify the performance of the pre-amplifier (usually in response to radiation damage). The final

Table 4.3 The configuration parameters for the modules. The upper table enumerates the values in the chip-level configuration register. The lower table describes other registers and toggles that configure various parameters at the module, chip and channel levels.

Name	Values	Normal Value
Compression Mode	Hit, Level, Edge, Test	Edge
Calibration Index	0, 1, 2, 3	Any
TrimRange	0, 1, 2, 3	Any, but lower is better
Accumulate Mode	On, Off	Off
Edge Detection	On, Off	On
Send Mask	On, Off	Off
Master Toggle	On, Off	On in master
End Toggle	On, Off	On at end
Input Bypass	Normal, Bypass	Normal
Output Bypass	Normal, Bypass	Normal
Clock Feed-through	On, Off	Off

Name	Level	Units	Minimum	Maximum	Step Size
Discriminator Threshold	Chip	mV	0	640	2.5
Calibration Amplitude	Chip	fC	0	16	0.0625
Strobe Delay	Chip	ns	0	24	≈ 1
Pre-Amplifier Bias Current	Chip	μA	0	294.4	9.2
Shaper Current	Chip	μA	0	38.4	1.2
TrimDAC	Channel	mV	0	60-240	4-16
Mask	Channel	Boolean	—	—	—
Redundant Input Select	Module	Boolean	—	—	—

register is a 16-bit register that contains various flags which control the readout and bypass systems. The flags are as follows:

- **Compression Mode:** Chooses which compression pattern to use
- **Calibration Index:** List of channels to send the calibration charge to
- **TrimRange:** Specifies the size of the step when determining the trim voltage
- **Accumulate Mode:** Toggle the accumulation of hits at the end of the pipeline
- **Edge Detection:** Toggle the edge detection circuits
- **Send Mask:** Activate to send the channel mask as data through the pipeline
- **Master Toggle:** Marks the chip as master (this is ignored unless the hardware toggle is set)
- **End Toggle:** Marks the end of the readout chain
- **Input Bypass:** Use the bypass connection for input tokens
- **Output Bypass:** Use the bypass connection for output tokens

- **Clock Feed-through:** The master chip sends the clock-divided-by-2 signal instead of data (set on startup)

There are a very limited number of ways that one can control the ASIC chips. The simplest command is simply the reset signal, used to resynchronise the chip with the clock during a physics run. An L1A command can also be used to force the chip to accept the Level 1 trigger.

Apart from the two commands above, the only way to control the chip is to configure the 6 registers discussed above. My work on the SCT involved the chip-level TrimRange setting, and the two channel-level settings given in Table 4.3: the TrimDAC and the Mask.

The TrimRange and TrimDAC are used to set the threshold voltage described above. TrimDAC takes a 4-bit digital signal and transforms it into the threshold voltage. The magnitude of the threshold voltage is dependent on the TrimRange, which determines how large the step-size between TrimDAC values. For example, the smallest TrimRange value is called “Range_0” and corresponds to a step-size of 4 mV. This way, a TrimDAC value of 1 corresponds to a threshold voltage of 4 mV, a TrimDAC of 2 means a threshold of 8 mV, etc. The maximum TrimDAC value is 15, which corresponds to a threshold of 60 mV. Increasing the TrimRange to “Range_1” increases the step-size to 8 mV, meaning that the TrimDAC values of 1, 2 and 15 corresponds to thresholds of 8, 16 and 120 mV, respectively.

The threshold voltage is thus set for each channel using a chip-wide TrimRange and a channel-specific TrimDAC. Recall that the threshold voltage is used to determine whether or not a strip has detected a particle, based on the induced voltage that the particle leaves behind. Naïvely, this method of setting the threshold doesn’t really make sense since one would think that a voltage is present only when a particle passes through the strip, so the threshold voltage would always just be 0 mV. In practice, all electronic equipment has a non-zero amount of noise present all the time. Using a threshold of 0 mV would lead to every strip picking up hits all the time. Thus, the threshold must be set so that hits can be distinguished from the noise.

This method also allows affords the chip controller the ability to react to changing conditions in the detector. Over time, radiation damage builds up in the strips, which increases their noise levels. The TrimRange and TrimDAC must then be increased so that the larger noise is ignored. However, there is a problem with this method, and it lies with the fact that the TrimRange is set at chip-level while the TrimDAC is set at channel-level.

Suppose that only one of a chips' 128 channels is much noisier than the rest. In order to account for this one troublesome channel, the TrimRange must be chosen at a high value, which means that the rest of the channels will receive a TrimRange far too high for them, and so they will ignore softer hits which they would otherwise be sensitive too. An example of this is given in Figure 4.12. These plots come from the “TrimRangeTest”, a test used to find the TrimDAC and TrimRange values of a given module. In these plots the x -axis is the channel number and the y -axis gives the TrimDAC and VT50⁴ values (represented in red and blue, respectively). The broken vertical lines represent the divisions between different chips.




Fig. 4.12 TrimRangeTest Output - Before Masking Channels.

This plot clearly shows the shortcomings discussed. Most of the chips are using “Range_2” because they have troublesome channels that are artificially increasing the TrimRange. The channels using “Range_1” however, have a much more spread out

⁴The “VT50” value is another measure used by the chip, and in technical terms it represents the discriminator threshold (unrelated to the threshold voltage discussed in this chapter) at which occupancy is at 50%. It is a concept that is somewhat beyond the scope of this thesis, however, so it won't be explained here.

distribution of TrimDAC values since they don't need a large TrimRange to effectively remove noise from all channels.

The core of my work with the SCT was finding a solution to this problem. A physical solution was out of the question, going down into the ATLAS cavern and replacing a noisy strip is simply infeasible given both the amount of modules and the fact that access to the cavern is forbidden during runs for radiation-safety reasons.

The next best solution is to mask the troublesome channels. This is done entirely through the software that runs the SCT, so it can be implemented and deployed while a run is ongoing⁵. In practice, this meant finding the standard deviation of the TrimDAC values for a given chip and masking the channels that were too many standard deviations from the mean. The masking is done through the mask register mentioned above. This register can be controlled using software and is used to remove the results of a given channel. In terms of the detector electronics, if the mask register is active it simply sets all of the bits in the channel to 0.

Implementing this solution and rerunning the TrimRangeTest gives the plots in Figure 4.13.

The results speak for themselves. Removing troublesome channels in a given chip dramatically reduces the TrimRange necessary to remove noise effects from that channel. Before the masking, 10 of the 12 chips used “Range_2”, with the last two using “Range_1”. Now, there are 10 chips using “Range_0”, with only two using “Range_1”. Once again the smaller range means that the TrimDAC can take more spread out values across different channels.

4.3.3 Transition Radiation Tracker

The Transition Radiation Tracker is the third tracker in the ID, serving as its outermost layer. Unlike the Pixel and SCT, the TRT doesn't make use of silicon as its detector material. Instead, it uses tubes containing a special gas mixture which reacts to particles passing through it. Again in contrast to the Pixel and SCT, the TRT is designed to give a large number of hits whenever a particle passes through it (typically around 36) while not being optimised to give perfect information of exactly *where* the hits occurred.

⁵During any given run day, there will be a set of hours dedicated to actually colliding particle beams and there will be another set of hours where collisions aren't occurring but the accelerator is still on. This second time frame is used by the accelerator and detector teams to fix any issues that might have arisen during the collisions.




Fig. 4.13 TrimRangeTest Output - After Masking Channels.

The TRT sensors are polyimide drift tubes (commonly referred to as “straw tubes” or just “straws”) which are approximately 4 mm in diameter and similar in structure to coaxial cables. The outer walls of the tubes are structurally robust and electrically insulating while using the minimum amount of material possible. At the centre of the tubes there is a metallic wire, and between the wire and the walls of the tube there is a gas. The outer wall is 70 μm thick; made of two 35 μm thick multi-layer films bonded back-to-back. These films consist of a roughly 25 μm thick polyimide film, coated on one side by a 0.2 μm thick layer of aluminium with a 5 μm thick graphite-polyimide layer protecting it, and coated on the other side by a 5 μm thick polyurethane layer which is used to heat-seal the two back-to-back films. The central wire is a piece of tungsten 31 μm in diameter plated with a 0.5 - 0.7 μm thick layer of gold. The gas mixture in the tubes is: 70% Xe , 27% CO_2 and 3% O_2 .

The TRT sensors are arrayed in modules, which have a different structure depending on whether they are in the barrel or end-cap regions. In the barrel all of the straws are arrayed parallel to the beam pipe and in layers at specific distances from the beam pipe. The straws are all placed so that their midpoints are located at $\eta = 0$. There are

73 layers of straws in total. Each straw is 144 cm long and 4 mm in diameter, with the mean spacing between straws being 7 mm. Unlike the Pixel and SCT sensors, the TRT sensors have no way of knowing where along the length of the straw a charged particle passed through. As such the TRT sensors are only able to provide information on a hit's r and ϕ coordinates, with an intrinsic $r - \phi$ accuracy of $130 \mu\text{m}$ in every straw. The wires within the straws are divided into two at approximately $\eta = 0$.

The TRT barrel modules are essentially 144 cm long rhombic prisms (Figure 4.14 shows a photo with the rhombic TRT modules highlighted) with an internal array of straw tubes embedded in a matrix of 9 mm-diameter polypropylene fibres. There are three layers of modules and each layer consists of 32 modules. A corollary to this is that the further out a module is, the larger it is and the more straw tubes it contains. The shell of a TRT module is made of $400 \mu\text{m}$ thick carbon fibre, and is the structure to which the straw tubes are attached and the front-end electronics are embedded. The shell serves as a container for the CO_2 gas which circulates around the straw tubes. This gas is used to prevent high-voltage discharges and prevent Xenon build-up in the case of leaks. Heat from the straws is drawn away from them by the CO_2 which transfers it into the carbon-fibre shell, which has a very high thermal conductivity. The heat is then drawn from the shell and away from the TRT by cooling pipes at the acute corners of the TRT modules.




Fig. 4.14 A photo of the TRT barrel during integration. The shapes of one inner module, one middle module and one outer module are highlighted in red.

The TRT end-cap modules are annuli that surround the SCT end-cap discs. Each end-cap is made up of 20 “wheels”, divided into two types. The first type, of which there are 12, are the closest to the collision point. Each of these wheels have 8 layers of straw tubes placed 8 mm apart. The second type, of which there are 8, are the furthest from the collision point. These wheels also have 8 layers of straw tubes but they are placed 15 mm apart. All layers are identical, having 768 straw tubes each. The straw tubes in the end-caps are arrayed pointing radially-outward, azimuthally equally spaced and are all 37 cm long. The spaces between the layers is filled with layers of 15 μm -thick polypropylene radiator foils separated by a polypropylene net. Finally, the end-caps use the same cooling method that the barrel modules do, with the CO_2 passing through each wheel sequentially.

The parameters of the TRT for the barrel and one end-cap are listed in Table 4.4.

Table 4.4 Parameters for the TRT barrel and end-cap modules (for one end-cap wheel). Quantities in **bold** are global quantities, including services and electronics. All other quantities are for individual modules and the active region.

	$ z_{Min} $ (mm)	$ z_{Max} $ (mm)	R_{Min} (mm)	R_{Max} (mm)	Number of Modules	Number of Layers	Straws per Module
Barrel (Both Sides)	0	780	554	1082	96	73	52544
Type-1 Module (Inner)	400	712.1	563	624	32	9	329
Type-1 Module (Outer)	7.5	712.1	625	694		10	
Type-2 Module	7.5	712.1	697	860	32	24	520
Type-3 Module	7.5	712.1	863	1066	32	30	793
End-Cap (One Side)	827	2744	615	1106	20	16	122880
Type-A Wheels	848	1705	644	1004	12	8	6144
Type-B Wheels	1740	2710	644	1004	8	8	6144

4.3.4 Insertable B-Layer

As it was originally designed, the Pixel detectors innermost layer (known as the “B-layer”) was to be the innermost layer of the ATLAS detector and for the first four years of operations (from 2008 to 2012) it was. Having this innermost layer perform well is critical for the ATLAS detector to reach its full physics potential. However, the original Pixel detector designs have a few issues with not only the Pixel detector itself but also the beam pipe it is mechanically attached to. In order to overcome these problems, a new layer of the ATLAS detector was installed (as well as a new beam pipe) called the Insertable B-Layer. This section will document the motivation behind the installation of the detector (and the new beam pipe), the constraints upon its design and what the detector is actually made of.

The motivations behind the IBL project are many and varied:

- **Tracking Robustness:** The innermost Pixel layer is exposed to a harsher radiation environment than the rest of the ATLAS detector. As such, it will suffer the most radiation damage over the course of ATLAS' lifetime, which will lead to the failure of individual pixel sensors in the short term and whole modules in the long term. Failures like this occurring in the B-layer of the Pixel detector will lead to serious deterioration in the impact parameter resolution, which is paramount for good secondary vertex identification and b -tagging. The insertion of a new B-layer overcomes these issues.
- **Luminosity:** The original Pixel B-layer was designed for an instantaneous luminosity of $L = 1 \times 10^{34} \text{ cm}^{-2}\text{s}^{-1}$, but the luminosity provided by the LHC increased over time⁶. The higher luminosity means a higher pile-up which in turn means higher occupancy and reduced readout efficiencies. These inefficiencies will affect the Pixel B-layer disproportionately compared to the other parts of the detector, and will significantly harm b -tagging efficiency. The introduction of a new B-layer, with its small occupancy, will solve these issues and preserve tracking performance.
- **Tracking Precision:** The IBL has a much smaller cylindrical radius than the Pixel B-layer, meaning it will be much closer to the interaction point. As such, it will have an even better impact parameter resolution and will thus give higher quality vertex reconstruction and b -tagging.
- **B-Layer Replacement:** The designers of the ATLAS detector understood that the Pixel B-layer would need to be replaced regularly because of radiation damage. The final design of the detector didn't allow for this, instead a compromise was reached where the Pixel detector and the beam pipe shared the same mechanical structure. This means that replacing the Pixel B-layer also involves removing other parts of the Pixel detector and the beam pipe. This compromise is obviously suboptimal, so in 2008 a task force was assembled to explore alternative designs for the section of the ATLAS detector nearest the interaction point to see if it could be redesigned with easier removal of the Pixel B-layer and beam pipe in mind. It was found that insertion of a new B-layer (the IBL) with a new structure for both the beam pipe and Pixel detector was not just the *best*, but the *only* viable solution.

⁶Indeed, the collisions at the end of Run 2 boasted instantaneous luminosities of $L = 2.1 \times 10^{34} \text{ cm}^{-2}\text{s}^{-1}$ [67].

- **Beam Pipe Replacement:** As mentioned above, the original design of the Pixel detector had it share a structure with the beam pipe, and to remove the beam pipe one would have to remove sections of the Pixel detector as well. This is suboptimal, and the redesign mentioned in the previous point included a change that allows the beam pipe to be removed on its own.
- **Large Radiation Doses:** At the time of installation the radiation doses that the Pixel was predicted to receive in the runs that would take place from 2015 to 2018 weren't thought to be large enough to cause problems for the Pixel detector. Nonetheless, installing a new detector layer will mean that any problems that *do* arise will be more manageable.

However, the IBL project also has a few strong constraints that have a substantial impact on the design of the IBL:

- The smaller radius of the IBL and its closer proximity to the interaction point means that the technology it uses must be even **more** radiation hard than the Pixel
- The smaller radial distance between the new beam pipe and the IBL means that it can't have overlapping, tilted modules like the Pixel and SCT do. Instead, the modules' sensor tiles are designed differently so that there is complete ϕ coverage without overlap.
- Minimising the material used is important for optimisation of tracking and vertexing performance. As such, the radiation length of the IBL is around 60% of the current Pixel B-layer.
- The goal of being able to remove the beam pipe without removing the IBL means that the support structure holding up the detector needs to be redesigned.

The IBL itself is a fairly simple detector. It only consists of a barrel (no end-caps) and only has one barrel layer. The barrel is held up by 14 carbon fibre staves, each of which are 2 cm in diameter, 64 cm long and are arrayed around the beam pipe at a mean radius of 33 mm. Each staves is rotated by approximately 14 degrees in ϕ with respect to the previous stave. Each stave has attached to it pipes for cooling and the optical fibres that carry data to and from the modules. Each stave carries 32 silicon sensors, each of which has a single FE-I4 front-end chip bump-bonded to it. Each of these sensors has 26,880 sensors, each connected to a channel via the FE-I4 chip. The pixel sensors are arrayed in 80 columns and 366 rows. Table 4.5 contains the mains parameters of the IBL.

Table 4.5 Main parameters of the IBL.

Item		Radial Extension [mm]	Length [mm]	Staves	Modules	Pixels ($\times 10^6$)
IBL	Envelope Sensitive	$R \in [31.0, 40.0]$ $\langle R \rangle = 25.7$	$ z < 332$	14	224	6.02

4.3.5 Inner Detector Magnet System

The whole Inner Detector is enveloped by a large solenoid magnet. It is this magnet that generates the axial magnetic field that bends the paths of charged particles in the Inner Detector, so that their momenta can be calculated. The solenoid is approximately 5.8 m long, with an inner radius of 2.46 m and an outer radius of 2.56 m.

The solenoid is made of the same NbTi alloy that all of the LHCs magnets are made out of, and as such is also a superconducting magnet, stored in a cryostat. The magnet produces a 2 T magnetic field within the Inner Detector with a current of approximately 7.7 kA passing through it.

The design of the solenoid needed to be carefully optimised given its position in the detector. As it is enveloped by the calorimeters, it must be optimised to affect them as little as possible. To this end, the solenoid is only 0.66 radiation lengths thick. It is also stored in the same vacuum vessel as the calorimeter system, eliminating two vacuum vessel walls and thus removing even more material between the interaction point and the calorimeters.

4.4 Calorimeter System

In contrast with the Inner Detector, whose purpose it is to track the particles that pass through it and provide momentum measurements, the purpose of the calorimeters is to capture the particles entering them and measure their energies. These energy measurements are used to calculate a variety of quantities used by analysis teams, from simple things like the masses of jets and invariant masses of multi-particle systems to more complicated variables that are used as proxies for the masses of intermediate states (such as the vector bosons or various hypothesised SUSY particles). As such, having accurate measurements of particle energies is essential for the experiment.

The calorimeter system is made up of two subdetectors, the Electromagnetic Calorimeter and the Hadronic Calorimeter. The Electromagnetic Calorimeter captures particles that interact electromagnetically, such as electrons and photons, while the

Hadronic Calorimeter captures particles that interact via the strong force, like protons and neutrons. Together, the calorimeter systems covers a range of $|\eta| < 4.9$, but this range contains many different physics processes spread throughout it and so the calorimeters use different techniques to study different processes in different parts of that η range. The EM calorimeter shares a section of the η range with the ID and has a very fine granularity to allow for the precision measurement of electrons and photons. The Hadronic Calorimeter has a comparatively coarser granularity but one that is still good enough to satisfy the physics requirements needed for jet and missing energy reconstruction [59].

Finally, the calorimeter must fulfil an under-appreciated goal: the ability to capture *all* of the material that enters it. The calorimeter must be able to capture all of the particles that enter it. If it does not, they will “punch through” and travel into the outer most subdetector, the Muon System.

Figure 4.15 shows a cut-away diagram of the ATLAS Calorimeter System, while Table 4.6 shows the parameters of the subdetectors that make up the calorimeter system.




Fig. 4.15 A Cut-away diagram of the ATLAS Calorimeter System

Table 4.6 Parameters of the Calorimeter System [59].

ElectroMagnetic Calorimeter				
	Barrel		End-Cap	
	Number of Layers	$ \eta $ Coverage	Number of Layers	$ \eta $ Coverage
Presampler	1	$ \eta < 1.52$	1	$1.5 < \eta < 1.8$
Calorimeter	3	$ \eta < 1.35$	2	$1.375 < \eta < 1.5$
	2	$1.35 < \eta < 1.475$	3	$1.5 < \eta < 2.5$
			2	$2.5 < \eta < 3.2$
	Granularity ($\Delta\eta \times \Delta\phi$)	$ \eta $ Coverage	Granularity ($\Delta\eta \times \Delta\phi$)	$ \eta $ Coverage
Presampler	0.025×0.1	$ \eta < 1.52$	0.025×0.1	$1.5 < \eta < 1.8$
Calorimeter (1 st Layer)	$0.025/8 \times 0.1$	$ \eta < 1.4$	0.050×0.1	$1.375 < \eta < 1.425$
	0.025×0.025	$1.4 < \eta < 1.475$	0.025×0.1	$1.425 < \eta < 1.5$
			$0.025/8 \times 0.1$	$1.5 < \eta < 1.8$
			$0.025/6 \times 0.1$	$1.8 < \eta < 2.0$
			$0.025/4 \times 0.1$	$2.0 < \eta < 2.4$
			0.025×0.1	$2.4 < \eta < 2.5$
			0.1×0.1	$2.5 < \eta < 3.2$
Calorimeter (2 nd Layer)	0.025×0.025	$ \eta < 1.4$	0.050×0.025	$1.375 < \eta < 1.425$
	0.075×0.025	$1.4 < \eta < 1.475$	0.025×0.025	$1.425 < \eta < 2.5$
			0.1×0.1	$2.5 < \eta < 3.2$
Calorimeter (3 rd Layer)	0.050×0.025	$ \eta < 1.35$	0.050×0.025	$1.5 < \eta < 2.5$
	Number of ReadOut Channels		Number of ReadOut Channels	
Presampler	7808		1536 (Both Sides)	
Calorimeter	101760		62208 (Both Sides)	
Hadronic Calorimeter				
Scintillator Tile Calorimeter				
	Barrel		Extended Barrel	
	Number of Layers	$ \eta $ Coverage	Number of Layers	$ \eta $ Coverage
	3	$ \eta < 1.0$	3	$0.8 < \eta < 1.7$
Outer Layer	Granularity ($\Delta\eta \times \Delta\phi$)		Granularity ($\Delta\eta \times \Delta\phi$)	
	0.1×0.1		0.1×0.1	
	0.2×0.1		0.2×0.1	
	Number of ReadOut Channels		Number of ReadOut Channels	
	5760		4092 (Both Sides)	
LAr End-Cap Calorimeter				
	Barrel		End-Cap	
			Number of Layers	$ \eta $ Coverage
			4	$1.5 < \eta < 3.2$
			Granularity ($\Delta\eta \times \Delta\phi$)	$ \eta $ Coverage
			0.1×0.1	$1.5 < \eta < 2.5$
			0.2×0.2	$2.5 < \eta < 3.2$
			Number of ReadOut Channels	
			5632 (Both Sides)	
LAr Forward Calorimeter				
	Barrel		End-Cap	
			Number of Layers	$ \eta $ Coverage
			3	$3.1 < \eta < 4.9$
			Granularity ($\Delta\eta(cm) \times \Delta\phi(cm)$)	$ \eta $ Coverage
Forward Calorimeter 1	3.0×2.6		$3.15 < \eta < 4.30$	
	Approx. 4x Finer		$3.10 < \eta < 3.15$, $4.30 < \eta < 4.83$	
Forward Calorimeter 2	3.3×4.2		$3.24 < \eta < 4.50$	
	Approx. 4x Finer		$3.20 < \eta < 3.24$, $4.50 < \eta < 4.81$	
Forward Calorimeter 3	5.4×4.7		$3.32 < \eta < 4.60$	
	Approx. 4x Finer		$3.29 < \eta < 3.32$, $4.60 < \eta < 4.75$	
			Number of ReadOut Channels	
			3524 (Both Sides)	

4.4.1 ElectroMagnetic Calorimeter

The Electromagnetic Calorimeter is used to capture and measure the energy of electrons and photons. The calorimeter is a lead/liquid-argon (often abbreviated as “Lead LAr” or just “LAr”) calorimeter and covers $|\eta| < 3.2$. Like the trackers, the EM calorimeter is made up of a barrel part and a pair of end-caps. The barrel is divided into two identical half-barrels, with the separation centred at $z = 0$ and the half-barrels being 4cm apart. Combined the barrel has a coverage of $|\eta| < 1.47$. Each of the end-caps is made of a pair of co-axial discs, with the outer wheel covering $1.375 < |\eta| < 2.5$ and the inner wheel covering $2.5 < |\eta| < 3.2$. The barrel has a finer granularity than the end-caps. Finally, the barrel and each of the end-caps have their own cryostat [59].

4.4.2 Hadronic Calorimeter

The Hadronic Calorimeter is divided into three components: the Tile Calorimeter, the LAr Hadronic End-Caps and the LAr Forward Calorimeter [59].

Tile Calorimeter

The Tile Calorimeter serves as the barrel section of the Hadronic Calorimeter and is made up of three cylinders. The inner most cylinder is unhelpfully known as the “barrel” and is central and symmetric with respect to the z -axis, covering $|\eta| < 1.0$. The other two are collectively known as the “extended barrel” and cover $0.8 < |\eta| < 1.7$, with one covering the positive range and the other covering the negative range. Each of these three cylinders contains three layers of detector modules, which are approximately 1.5, 4.1 and 1.8 interaction lengths thick for the barrel and 1.5, 2.6 and 3.3 interaction lengths for the extended barrel. The calorimeter extends from an inner radius of 2.28 m to an outer radius of 4.25 m. This calorimeter is made of steel (as the absorber) and scintillating tiles (as the active material) [59].

LAr Hadronic End-Caps

The end-caps of the Hadronic Calorimeter are made of the same material as the whole of the Electromagnetic Calorimeter. The hadronic end-cap (HEC) calorimeter has two independent wheels in each end-cap, and each wheel is made up of 32 identical wedge-shaped modules. The HEC covers an η range of $1.5 < |\eta| < 3.2$. The HECs share a cryostat with the Electromagnetic Calorimeter end-caps [59].

LAr Forward Calorimeter

The Forward Calorimeter (FCal) shares a cryostat with the other two end-caps (electromagnetic and LAr hadronic) and serves to extend the η range of the calorimeter system. The FCal is no fewer than 10 interaction lengths deep in the horizontal direction, and single-handedly extends the η range of the calorimeter system all the way to $|\eta| < 4.9$ [59].

4.5 Muon System

While most types particles that enter the detector pass through the trackers and land in the calorimeters, there are two kinds of particles that don't; muons and neutrinos. Neutrinos barely interact with anything at all and so pass through the detector and contribute to the E_T^{miss} , while muons have a special system that measures their properties: the Muon System [59].

The Muon System surrounds the calorimeter system and is used to measure the positions and momenta of muons. It is itself a tracker and it uses the same methods that other trackers use, namely the particles are deflected by powerful magnetic fields and pass through high-precision tracking chambers which measure the positions of the particles. Like the other subdetectors the Muon System can be split into two segments, the barrel and the end-caps. Each of these sections cover different η ranges and have different magnet systems built within them. In total, the Muon System covers the η range: $|\eta| < 2.7$. Figure 4.16 shows a cut-away diagram of the Muon System with the components labelled [59].

4.5.1 Muon Chambers

The Muon System uses a variety of different technologies in its modules, and uses different types of modules for measuring particles and triggering on them.

For measuring, the Muon System uses two different kinds of modules: the **M**onitored **D**rift **T**ubes (MDTs) and the **C**athode **S**trip **C**hambers (CSCs). MDTs consists of a drift tube with a sense wire within it, which allows the MDT modules to give precision measurements of muons while also being robustly constructed. The MDTs are typically used at smaller rapidities ($|\eta| < 2.0$). The CSCs are multiwire proportional chambers with cathode strips segmented into strips, and are much higher granularity than the MDTs. They are used in higher pseudorapidity regions ($2.0 < |\eta| < 2.7$) which have demanding rates and background conditions that the MDTs can't operate in.




Fig. 4.16 A cut-away diagram of the ATLAS Muon System

For triggering, the Muon System has two more kinds of modules: **R**esistive **P**late **C**hambers (RPCs) and **T**hin **G**ap **C**hambers (TGCs). The RPCs are used in the barrel and the TGCs are used in the end-cap, and together the trigger modules cover a η range of $|\eta| < 2.4$. These modules have three purposes: providing well-defined p_T thresholds, providing bunch-crossing identification and to measure the muon coordinate in the direction orthogonal to the direction determined by the precision tracking chambers [59].

Table 4.7 shows the main parameters of the Muon System.

4.5.2 Muon System Magnet System

The magnet system of the Muon System consists of three large air-core toroids, one toroid for each end-cap and one barrel toroid. Each of these three toroids is made up of eight coils of wire arranged radially and symmetrically around the beam pipe, with the end-cap toroids rotated 22.5° with respect to the barrel toroid in order to

Table 4.7 Main parameters of the Muon System. The numbers in the brackets for the MDTs and RPCs refers to the final configuration of the detector in 2009. The innermost layer of the MDT modules has an η coverage of $|\eta| < 2.0$ and the TGC has special triggering modules that cover $1.05 < |\eta| < 2.4$.

	Coverage	Number of Chambers	Number of Channels	Function
MDT	$ \eta < 2.70$	1088 (1150)	339,000 (354,000)	Precision Tracking
CSC	$2.00 < \eta < 2.70$	32	31,000	Precision Tracking
RPC	$ \eta < 1.05$	544 (606)	359,000 (373,000)	Triggering
TGC	$1.05 < \eta < 2.70$	3588	318,000	Triggering

provide radial overlap and optimise the field strength in the interface between the two coil systems. Each of the eight coils of wire in the barrel have their own cryostat with linking elements joining them together and providing mechanical structure. Each of the end-cap toroids have all eight of their coils cold-linked and contained in a single cold-mass, and then that cold-mass is contained in a single cryostat for both cooling and structural support [59].

4.6 Forward Detectors

The subdetectors discussed in this chapter so far provide almost all of the $\approx 4\pi$ solid angle coverage of the ATLAS detector. However, the η range of these subdetectors only extends to about $|\eta| < 5.0$. The region past this ($|\eta| > 5.0$, known as the “forward region”) remains uncovered. The reason for this is that covering this region by placing detector material within the physical extent of the ATLAS detector would require placing detector material *inside* the beam pipe. In order to have detector coverage of the forward region three “forward detectors” are constructed and placed near the beam pipe at various distances from the interaction point [59].

The closest forward detector to the interaction point is the **L**uminosity Measurement using **C**erenkov **I**ntegrating **D**etector, or “LUCID”. LUCID is placed 17 m away from the interaction point, within the end-cap regions. As its name implies LUCID is (one of the two) detectors used to measure the luminosity delivered to the ATLAS detector. In particular, LUCID measures the relative luminosity delivered to ATLAS while the LHC is running [59].

The other luminosity detector is “**A**bsolute **L**uminosity **F**or **A**TLAS” Detector, or “ALFA”. Each end of the detector is placed 240 m away from the interaction point in either direction. ALFA is constructed using scintillating fibre trackers inside of roman

pots which allow parts of the detector to be placed as close as 1 mm to the beam itself [59].

The final forward detector is the **Zero-Degree Calorimeter**, or “ZDC”. The detector is placed at ± 140 m, beyond the point where the short-straight-sections divide into two independent beam pipes. The ZDC is made of alternating quartz rods and tungsten plates and is used to measure the presence of neutral particles at η values as high as 8.2 [59].

4.7 Structure, Radiation Shielding and Control

The extreme experimental conditions at the LHC imposes several challenges on the design and construction of the ATLAS detector. The immense energies and luminosities of the beam means that the amount of radiation the detector is exposed to (both in the form of collision products that the detector is *meant* to detect as well as radiation given off by the beam) is severe. The ATLAS detector is designed with this in mind and has optimised the layout of radiation shielding to simultaneously keep the detector safe from radiation while also minimising the physical material used by the detector. The ambition to place tracking detectors as close to the beam pipe as possible has meant that the beam pipe itself must be carefully designed to accommodate it. Finally, the LHC and ATLAS control rooms need to be in constant communication for all manner of reasons, from beam monitoring to timing to safety. This section will briefly describe these challenges and what is done to overcome them [59].

4.7.1 Beam-Pipe Structure

The section of the beam pipe enclosed by the ATLAS detector is approximately 38 m long and divided into seven sections. These sections are bolted together with flanges and together form a fully in-situ bakeable ultra-high vacuum system. The central chamber is centred around the interaction point, it has a 58 mm inner diameter and 0.8 mm thick beryllium walls. The remaining six chambers are made of stainless steel and are installed symmetrically on both side of the central section. Structurally, the three pairs of chambers (moving outward) are supported by the end-cap LAr cryostats, the end-cap toroids and the forward shielding structures, respectively [59].

4.7.2 Radiation and Shielding

At the LHC, the primary source of radiation at full luminosity comes from collisions at the interaction points. While accurately measuring the products of these collisions is indeed the purpose of the ATLAS Experiment, there is such a large amount of material entering the detector (that isn't necessarily of interest to analysts) that it must be shielded from the very collisions it is attempting to study [59].

The next most intense source of radiation comes from the beam tube. Collisions can occur wherein the protons strike each-other with a glancing blow, resulting in protons that are offset from the beam direction by a very shallow angle and strikes the walls of the beam pipe. These strikes produce particles of their own, which travel into the detector as well [59].

These background radiations have a number of deleterious effects on the detector. The radiation itself causes background levels to increase. Exposure damages the detector components over time, diminishing their ability to perform measurements. Single particles from collisions other than the one at the interaction point can enter into recorded events and upset them. The radiation can even create long-lived radionuclides which present a hazard for the maintenance crews entering the ATLAS cavern once the collider has been shut down [59].

In order to limit the amount of radiation received by the detector and prevent the deleterious effects listed above the ATLAS detector has almost 3000 tonnes radiation shielding installed in it. This shielding is installed in three layers, each of which are made using different materials and perform different functions. The inner-most layer is made of materials like copper and iron which have a large amount of interaction lengths per unit volume. This layer is used to stop high energy hadrons and secondaries. The second layer is made of doped polyethylene enriched with hydrogen, and is used to moderate the neutron radiation escaping from the first layer. The third and outer-most layer is a large section of steel and lead and is the only thing that can stop the photons that escape from the previous two layers [59].

4.7.3 Control

The LHC and the ATLAS detector must continually communicate in order to maintain optimal and safe operations of both machines. The LHC communicates information about the beam; it's position, the intensity of it's bunches, etc. However, the most important thing that the LHC sends to ATLAS (and indeed, all of the experiments) is the 40.08 MHz clock, without which all measurements would be futile. Conversely,

the ATLAS detector sends back information on the luminosity of the beam it gathers with its luminosity detectors, information of the quality of the collisions gained by the detector itself and the ATLAS beam interlock system provides information on whether the beams are safe to be injected or if they need to be dumped [59].

Chapter 5

The ATLAS TDAQ System

Despite its size and complexity, the only direct measurements that the ATLAS detector takes are of the position and energy of particles entering into it.

Trackers record the positions of charged particles passing through them as hits. Calorimeters register showers as a contiguous group of calorimeter cells with energy deposits in them. On their own, these hits and energy deposits aren't usable data. Every hit and deposit only briefly exists in the detector before dissipating, so there needs to be some way to read the data associated with each one off the detector and store it somewhere more permanently. The system that reads data off of the detector and stores it on CERN's computer system is a large aggregation of computers and electronics, and rivals the detector itself in complexity. This system is known as the ATLAS “**T**rigger and **D**ata **A**cquisition” (or TDAQ) system.

This system will be the focus of this chapter, with a small caveat. The ATLAS TDAQ system is so complex that there have been thousands of pages of reports, papers and technical documentation written about its construction, operation, performance and the upgrades that it's undergone over its lifetime. Further, the entire system is incredibly heterogeneous, with each subdetector having its own unique computer system handling the readout of data. All of this is far beyond the scope of this thesis and including it with any significant amount of detail would result in this thesis being well over a thousand pages long. Instead, this thesis will give a brief, qualitative description of the TDAQ system with special focus on the two parts (the trigger and reconstruction) that are specially relevant to the discussion of the physics analysis in the coming chapters.

What is common across the whole detector is the method by which the data is carried from the detector to the various computer systems that process it. The detector inhabits a cavern that is approximately 40 m below the ground, while the

ATLAS computer systems are at sea level. Standard copper wires aren't sufficient for connecting the detector to the computer system, so instead fibre-optic cables are used. To facilitate this, every module has an opto-package that converts the digital signal from hits/deposits to an optical one that is then passed along the optical fibres. Similar technology is present on the other end of the optical fibre, another opto-package converts the optical signal back to digital so that it can be used by a computer.

Returning to the question of how to store the data; the LHC has a collision rate of 40 MHz, meaning that with a pile-up of 25 proton-proton interactions per bunch crossing there are about one billion proton-proton interactions occurring in the ATLAS detector every second. Such a collision rate is necessary to produce statistically significant amounts of the rare processes that the LHC was built to study, but a caveat to this is that the ATLAS detector is constantly being inundated with an absurd amount of particles. Recording the hits and energy deposits of every single particle produced by the collisions would mean recording approximately 70 TB of data per second. Recording this amount of data this quickly is infeasible. The average commercial CD-ROM holds approximately 700 MB [68], if these discs are used to store the outgoing data from the ATLAS detector 100,000 of them would be filled *every second*. It is clear that from a purely logistical point of view, the rate of data-taking must be reduced.

To this end, there are two possible solutions: reducing the amount of data taken per event, or ignoring certain events in favour of others. The first of these solutions isn't useful. In order to bring the data rate down to something manageable, such as $\mathcal{O}(100 \text{ MBs}^{-1})$, each event would have to be reduced in size by a factor of 400,000. This isn't really possible since each event is already $\mathcal{O}(100 \text{ MB})$. Such a reduction would mean that each event would only have one byte of storage space available to it, which isn't nearly enough space to store all of the necessary information to fully describe an event.

The second possible solution: ignoring certain events in favour of others, is much more realistic and the solution chosen for the ATLAS detector. Recording the data for every event that occurs in the ATLAS detector is not just technologically infeasible, it isn't even particularly useful. As stated earlier, the processes that the LHC was built to study are very rare, so most of the processes that occur in collisions aren't of any interest anyway¹. This means that the overwhelming majority of events that occur in the ATLAS detector can probably be ignored. In fact, this is true of so many events

¹This isn't strictly true. The study of rare processes, such as production of the Higgs boson, was at least part of the motivation for building the LHC. However, there are many groups within CERN dedicated to more precise measurements of the properties of already well-understood processes, and these groups are indeed interested in these less-rare processes.

that decreasing the data rate by a factor of 400,000 is not only technically possible, but has been done.

How to implement this solution isn't straightforward. Separating "uninteresting" events from "interesting" events is impossible without detailed knowledge of the events themselves, usually knowing which particles the events contained. This information isn't immediately available from the hits and energy deposits themselves, but can be derived from them with time and computing power. Unfortunately, this is once again at odds with the large collision rate of the LHC. The TDAQ system must be able to decide whether an event is interesting or not in the 25 nanoseconds between two collisions².

This complex problem of balancing fast decision-making against detailed event analysis is at the heart of the design of the ATLAS TDAQ system. The two parts of the TDAQ mentioned above were each made as a solution to one of these problems: the fast decision-making is handled by the ATLAS "trigger" system and the detailed event analysis is done by the ATLAS "reconstruction" algorithms. Of course, in order to be able to solve the complex problem of deciding which events to keep these two systems must be connected and in constant communication with one another.


The trigger system is made up of two parts: a hardware-based trigger and a software-based trigger. The hardware trigger has direct access to the data the detector provides. This allows it to make a (vague) decision on whether to keep an event or not in less than 25 ns. The software trigger is more complex, slower and utilises information from the reconstruction system to make more detailed decision [69]. Section 5.1 will describe the trigger system in detail.

The reconstruction system is a software framework that collects the hits and energy deposits in an event and "reconstructs" back into particles. This process is less straightforward than it sounds, since the reconstruction system only has access to the hits, energy deposits and information from the hardware trigger for any given event and it often isn't easy to turn this information into fully realised particles. The output of the reconstruction system also informs the decision made by the software trigger. Section 5.2 will describe the reconstruction system in detail.

²This isn't strictly true. The front-end electronics on the detector have buffers that allow the data from an event to be temporarily stored at times when any previous event is taking longer than usual to be processed.

5.1 Trigger System

The ATLAS trigger system works similarly to the trigger mechanism in an oscilloscope. In an oscilloscope the signal shown on the screen, such as the example signal given in Figure 5.1a, needs to be read from the same point of the waveform every sweep³ to be stable. If this doesn't happen, every time the signal is updated (in the case of the oscilloscope shown in Figure 5.1, this rate is 60 Hz) it will appear at a new horizontal position and the screen becomes unreadable. The trigger is the mechanism that tells the oscilloscope when to begin reading the waveform so that the signal on the screen is stable, and is controlled by the knobs shown in Figure 5.1b.



(a) A screenshot of an oscilloscope screen showing a sinusoidal alternating current. (b) A photograph of the controls of an oscilloscope.

Fig. 5.1 A screenshot taken of an oscilloscope screen (left) and a photo of that oscilloscope's controls (right).

The ATLAS trigger system has approximately the same function as the trigger mechanism in an oscilloscope. When an event occurs the ATLAS detector records the particles that enter it as hits and energy deposits. Information from these is then used by the trigger to decide if the event is worth recording. If the trigger decides that it is (this is known as the event “passing” the trigger) the data from the hits and deposits is read off the detector and saved to the ATLAS computer system.

As already stated, the ATLAS trigger system is made up of a hardware component and a software component. The hardware component of the trigger is called “level 1” (or “L1”) and the software component is called the “high level trigger” (or “HLT”).

³In the context of oscilloscopes, the term “sweep” used to refer to how the beam in old analogue oscilloscopes would sweep from left-to-right on the screen. In analogue oscilloscopes this was effectively the screen refresh rate. Starting the sweep at a different height on the screen every time would result in an unreadable screen since most oscilloscopes had sweep rates in the dozen of Hertz.

The L1 trigger reduces the event rate from 40 MHz down to 100 kHz and then the HLT takes this 100 kHz event rate and reduces it further down to 1 kHz. In both cases, data from the detector is used to decide whether an event should pass a trigger [69].

The L1 trigger makes use of a simple, fast reconstruction algorithm to turn the hits and energy deposits into more complex objects. Hits in the MS are grouped together into “tracks”, curves in 3-dimensional space that contain the hits associated with them⁴. Energy deposits come from showers in the calorimeters, meaning that most deposits come in contiguous groups in the calorimeters. The reconstruction algorithm turns these contiguous groups into “clusters”. Clusters are limited to exist only in single calorimeters, that is, clusters can’t have part of their deposits in the electromagnetic calorimeter and part in the hadronic calorimeter [70].

These tracks and clusters are then further reconstructed into particle candidates: electrons/photons, hadrons/taus, muons and jets⁵. These candidate particles have approximate values for properties such as momentum calculated for them by the L1 reconstruction algorithm. This approximation is the best it can possibly be given that the trigger that uses it must be able to make decisions in the time between events. The L1 trigger itself contains a list of thresholds for each of the properties of each different type of particle. If an event contains a particle which has a property (for instance, momentum) above the threshold, then the trigger for that particle and property is set to boolean value “true”. At the end of the event analysis the trigger then takes the logical OR of all of the particle triggers. If the result is “true”, then the L1 trigger issues a “Level 1 Accept” (or “L1A”) signal to the HLT, before sending the event data to both the reconstruction software and the HLT. Approximately 1-in-400 events pass the L1 trigger, leading to the 100 kHz outgoing event rate mentioned above [70].

The HLT can afford to be slower than the L1, since it has a much smaller rate of events entering it. This means that it can make use of the objects made by the reconstruction algorithm (discussed in detail in Section 5.2). This means that it has access to higher-quality candidate particles and that it can trigger on more complex variables such as missing energy or statistical likelihoods.

HLT triggers are software objects and so have special, systematic names. These names consist of short strings of characters representing the particle and property

⁴It should be noted that while the trigger only uses tracks from the MS, the hits in the ID can also be made into tracks. These tracks aren’t used by the L1 trigger, but are used in the reconstruction procedure described in Section 5.2.

⁵Since information from the ID isn’t being used, it isn’t possible to tell certain types of particles apart. This is why the electrons and photons (who both leave similar showers in the electromagnetic calorimeter) are reconstructed as the same particle candidate. The same thing happens with hadrons and taus, who both leave similar showers in the hadronic calorimeter.

being triggered on, separated by underscores. An example of such a trigger is: “e60_lhmedium_nod0”. Here, the “e60” means that the event must contain one electron with an energy of at least 60 GeV, the “lhmedium” means that electron must pass the “medium” likelihood requirement (a statistical measure discussed in detail in Section 5.2). The “nod0” is pronounced “no d_0 ”, and means that the event should have $d_0 = 0$. The variable d_0 is the transverse distance from the primary vertex to the reconstructed particles origin vertex [69]. If an event contains an electron with all of these properties, the event passes that trigger.

There are many different triggers used by the HLT, and the properties they can trigger on are even more numerous. There is “j100”, which requires a jet with an energy of 100 GeV, there are “loose” and “tight” likelihoods as well: “lhloose” and “lhtight”. There are multi-particle triggers too: “4j100” means four different jets, each with energy greater than 100 GeV. There are even combined triggers like “e17_lhloose_nod0_mu14”, which is a 17 GeV electron with no d_0 triggered together with a 14 GeV muon.

The HLT uses the same method of determining whether an event passes it as the L1 does. For each of the triggers in the HLT, the event is assigned a boolean value of “true” or “false” depending on whether it passed the trigger or not. A logical OR is then taken of all of these booleans and if the result is “true” then the event passes the HLT.

The combined list of L1 and HLT triggers in use during a run is known as the “trigger menu”. The full trigger menu used during the 2015 and 2016 data-taking periods is beyond the scope of this thesis and won’t be documented here for two reasons. First, the trigger menu isn’t fixed and changes over the course of a run. Including all of the triggers that were used over the course of 2015 and 2016 would result in 5-6 pages of trigger tables. Second, the full trigger menu isn’t really necessary to understand the analysis presented in this thesis. Since the HLT is a logical OR of many triggers analyses tend to impose their own trigger requirements and only consider events that pass them. The analysis presented in this thesis does exactly this and Chapter 7 will discuss the triggers that it uses.

Finally, a schematic diagram of the whole trigger system is given in Figure 5.2.

It should be noted that in the past the trigger system used to have three layers: L1, L2 and the Event Filter, EF. This is the system used in Run 1. In this system, both the L2 and EF were software triggers with different levels of complexity. The old L1 trigger was less advanced than the new one, reducing the event rate from 40 MHz down to 75 kHz instead of 100 kHz. The L2 trigger reduces the rate down to 1 kHz,




Fig. 5.2 Schematic diagram of the ATLAS Trigger and Data Acquisition System used during Run 2.

while the EF reduces the rate further down to 100 Hz. Both these systems have some level of reconstruction capability, with the EF being more advanced (and slower) than the L2 [71].

At the end of Run 1 the trigger was already receiving an event rate much higher than it was designed to handle. This prompted the ATLAS trigger team to perform a comprehensive upgrade to the trigger system. The L1 was made faster so that it could output an event rate of 100 kHz over the old 75 kHz. The L2 and EF were combined into the HLT. The HLT could make decisions on whether to use the L2 or EF reconstruction algorithms, instead of both being used all the time. This flexible decision making made the HLT use up much less computer space and also made it much faster. The final acceptance rate for the HLT was 1 kHz, up from the old EF rate of 400 Hz [72].

5.2 Reconstruction

As mentioned in the introduction to this chapter, the ATLAS detector “sees” particles as hits in the trackers and energy deposits in the calorimeters. There needs to be a software system that takes the hits and energy deposits in a given event and turns them back into the particles that emerged from the collision. This system is appropriately called the “reconstruction” system.

Again, this process isn’t as straightforward as it might seem. Since the only information that the detector has access to is based on position and energy it has no direct way of telling which particles have entered it. The different layers of detector can be used to deduce which particles are which (recall Figure 4.2), but this method has its shortcomings. Since the detector has to reconstruct hits and deposits *into* particles there is a chance that it could reconstruct something incorrectly, and there would be no way of knowing since the detector doesn’t have perfect knowledge of every particle. This problem is exasperated by the sheer amount of particles the collisions produce as there are often multiple hits/deposits in close proximity to one another. The reconstruction process must then be done in such a way that it creates the most likely version of the event that it can, since the actual anatomy of the event is unknowable. This is done through a process of trial-and-error.

Suppose that an event passes the L1 trigger. The event is passed to the reconstruction algorithms with the data from the hits and energy deposits as well as data from the L1 trigger. The algorithm reconstructs the hits in the ID and MS into tracks and the showers in the electromagnetic and hadronic calorimeters into clusters. The reconstruction algorithm then supposes that the event contains only electrons and begins reconstructing the tracks and clusters into electron candidates. This process prioritises the most isolated electron candidates (isolation will be discussed in more detail in Subsection 5.2.1). Once the algorithm has reconstructed as many electrons as possible they are all analysed and each of them is assigned a statistical weight called a “likelihood”. This likelihood will be explained in detail in Subsection 5.2.1, but for the time being it is best to think of it as a measure of how “electron-like” the reconstructed electron candidate is. These electron candidates are stored and this process is repeated for each particle type: muons, photons and jets. Each of these has different reconstruction algorithms and each is given a likelihood.

The event as a whole is then reconstructed by reconstructing tracks and clusters into the particle with the highest possible likelihood until there are no tracks and clusters left. The reconstructed objects then all have “isolation” variables calculated

for them. These variables are a way of measuring how separate a given particle is from the rest of the activity in the detector.

While the reconstruction algorithm outputs finished, reconstructed objects, the data set also contains the RAW data from the detector and L1 trigger. The reason for this is that sometimes an analysis might wish to implement their own reconstruction algorithm different from the one used in the TDAQ system. This is possible because the RAW data is preserved, so any analysis team can use it. Another thing that analysis teams can do is choose events based on the data used by the reconstruction algorithm. This is called “object selection” and the analysis presented in this thesis does this (see Chapter 7 for details).

5.2.1 Electrons

Electrons are reconstructed by matching a track from the ID and a cluster from the Electromagnetic Calorimeter. This process produces a basic electron object, to which is applied small corrections such as reconfiguring the cluster used by the reconstruction algorithm or calibrating it so that its η , ϕ and energy values better matches that of the original object [73] [74].

In order to be reasonably sure that what has been constructed is in fact an electron (as opposed to something else such as a heavy-flavour jet that has been wrongly reconstructed as an electron) the likelihood method described above is used. Each reconstructed electron is assigned a pair of probabilities representing the probability that the electron is a “signal” electron or a “background” electron⁶. These probabilities, called “likelihoods” are calculated by taking a product of a set of probability distribution functions, shown in Equation 5.1.

$$L_{S(B)} = \prod_{i=1}^n P_{S(B),i} \quad (5.1)$$

Each electron is then given a discriminator value based on the likelihoods, given in Equation 5.2.

$$d_L = \frac{L_S}{L_S + L_B} \quad (5.2)$$

⁶The terms “signal” and “background” are used extensively in this thesis and are explained in Chapter 6.

An inverse sigmoid function is used to transform the discriminator into a more useful form (Equation 5.3).

$$d'_L = -\tau^{-1} \ln(d_L^{-1} - 1) \quad (5.3)$$

Here: $\tau = 15$. Figure 5.3 shows distributions of this transformed discriminator for background and signal electrons. From this figure, it is clear that if one were to accept events with a log-transformed likelihood of greater than zero, then most of the electrons entering the analysis would be signal. The ratio of signal electrons to total amount of electrons entering an analysis is called the “signal efficiency”, and it is an option in the HLT trigger menu (recall the “lhloose” string). In fact, the HLT has four triggers on signal efficiency, corresponding to four values of discriminator. These are called “VeryLoose”, “Loose”, “Medium” and “Tight”. The “Tight”, “Medium” and “Loose” correspond to the discriminator levels that generate signal efficiencies of 93%, 88% and 80%, respectively. The “VeryLoose” operating point doesn’t have any requirements and is used for background studies. Finally, there is also an alternate “Loose” operating point, called “LooseAndBLayer”, that uses the same discriminant value as “Loose” but has the same requirement on the innermost PIXEL layer hits that the “Medium” and “Tight” working points do.

In addition to the likelihood method, there is another way of telling signal electrons from background electrons. Signal electrons tend to have a reduced amount of activity in the area of the detector immediately around them. This feature is exploited by the aptly-named electron “isolation” algorithm. This algorithm analyses the area around every electron candidate and derives an “isolation variable”, similar to the likelihood, that measures the amount of activity near the candidate electron. This is done for both tracks and clusters. For clusters, the isolation variable measures the amount of energy in the calorimeter within a cone of radius ΔR centred on the candidate electron. For tracks, the same method is used except with tracks in the ID and with the cone being smaller. Like the signal efficiency, the isolation variable has several operating points [73]. Figure 5.4 shows a schematic diagram of the calorimeter isolation algorithm.

5.2.2 Muons

Muons undergo largely the same reconstruction and isolation procedures as electrons but with a few important differences. Muons don’t interact with the hadronic calorimeter and barely interact with the em calorimeter, while also producing tracks in the




Fig. 5.3 The transformed likelihood discriminant d'_L for electron candidates with $E_T \in [30, 35]$ GeV and $|\eta| < 0.6$. The black histogram is for prompt electrons in a $Z \rightarrow ee$ simulation sample, and the red histogram is for backgrounds in a generic two-to-two process simulation sample. The histograms are both normalised to unit area.

specialised muon system surrounding the calorimeters. The reconstruction for muons matches tracks in the ID with tracks in the MS. Instead of a likelihood, the muon reconstruction has a quality-based identification requirement that makes use of a χ^2 -fit between the ID and MS tracks. Like the electron likelihood, this identification requirement has four operating points: *Loose*, *Medium*, *Tight* and *High- p_T* . The isolation requirement makes use of information from the ID and calos in much the same way electron isolation does [75].

5.2.3 Jets

The leptons discussed in the last two subsections are the particles which have the cleanest decays; that is, they leave highly collimated tracks in the trackers and dense clusters in the calorimeters. However, many of the particles produced in proton-proton collisions don't decay as cleanly. Quarks and gluons are bound by confinement and hadronise after leaving the interaction point. This results in a spray of particles that leaves many tracks in the trackers and many clusters in the calorimeters. This is called a “jet” and a schematic diagram of one is shown in Figure 5.5.




Fig. 5.4 Schema of the calorimeter isolation methodology. The grid represents the second layer of calorimeter cells in $\eta \times \phi$ space. The purple circle represents the cone of isolation with radius $\Delta R = \sqrt{(\Delta\eta)^2 + (\Delta\phi)^2}$. The topological clusters whose centres fall within the cone of isolation are coloured in magenta.

When looking at a jet in an event display a human eye can clearly see a jet, with multiple tracks leading to multiple, discontinuous calorimeter clusters, as originating from a single particle coming from the interaction point. Humans have evolved this pattern recognition ability, but computers don't possess it innately⁷. It is the function of a jet reconstruction algorithm to recognise when tracks and clusters come from a jet, and reconstruct them appropriately.

There are many different jet reconstruction algorithms in use by particle physics experiments⁸. However, it shouldn't be said that there is a single "best" jet reconstruction algorithm. All of them have their own strengths and weaknesses, and which one is used depends on the needs of the experiment. That said, they all must fulfil the most

⁷Although ATLAS has begun using neural networks to do parts of the jet reconstruction process [77][78].

⁸Indeed, there are about a dozen jet reconstruction algorithms, each with their own design philosophies, methods and definitions of what a "jet" is. They are all summarised in [79].




Fig. 5.5 Schematic diagram of jet production and measurement. Figure retrieved from Ref. [76].

basic need of a jet reconstruction algorithm: they must be able to reconstruct a jet from a single, hard, isolated particle.

The jet reconstruction algorithm used by this analysis is called the “anti- k_t algorithm”, and is part of a class of algorithms called “sequential recombination algorithms” [80]. In brief, these algorithms work by defining a generalised distance measure between two particles⁹ in the detector and then using this distance measure to group particles into jets.

While this chapter will focus on the anti- k_t algorithm and its properties there is some worth in exploring the properties of sequential recombination algorithms more generally, if only to provide more context for why the anti- k_t algorithm has the properties it does. To this end, Appendix D will briefly describe some other sequential recombination algorithms and their properties.

The anti- k_t algorithm uses the generalised distance measures defined in Equations 5.4 and 5.5.

$$d_{i,j} = \min\left(\frac{1}{k_{T,i}^2}, \frac{1}{k_{T,j}^2}\right) \frac{\Delta R_{i,j}^2}{R^2} \quad (5.4)$$

$$d_{i,B} = \frac{1}{k_{T,i}^2} \quad (5.5)$$

Here:

⁹The word “particles” here is somewhat misleading. It actually can stand for any detector-level object (hits and tracks in the trackers, energy deposits and clusters in the calorimeters) or any reconstructed object.

- y_i , ϕ_i and $k_{T,i}$ are the rapidity coordinate, azimuthal angle coordinate and transverse energy of the i^{th} particle.
- $\Delta R_{i,j}^2 = (y_i - y_j)^2 + (\phi_i - \phi_j)^2$, the distance in rapidity-azimuth space between particles i and j .
- $d_{i,j}$ is the distance between “particles” i and j in the detector.
- $d_{i,B}$ is the distance between the i^{th} “particle” and the beam, represented by B
- R is the radius of the cone defined by the jet. It’s value is chosen by the analysis team and usually takes the value of 0.4.

The anti- k_t algorithm then works as follows:

1. Calculate $d_{i,j}$ and $d_{i,B}$ for all possible particles
2. Find the minimum of all the calculated
3. If the minimum is a $d_{i,j}$, then combine the i^{th} and j^{th} particles into a jet and repeat.
4. If the minimum is a $d_{i,B}$, then declare i to be a final state jet and remove it from the list of particles before continuing.
5. Repeat until no particles remain.

The results of applying the anti- k_t algorithm to a set of test data is shown in Figure 5.6. The algorithm favours reconstructing jets as cones with stable circular boundaries. There are a few jets with non-circular shapes (the magenta and yellow jets for instance), but the shapes of these can be explained by other jets being reconstructed first and their energy deposits being removed from the algorithms consideration (in the case of the magenta jet the green jet would’ve been reconstructed first, while in the case of the yellow jet the blue jet would’ve been reconstructed first).

b -Tagging

While jets can be created by any particle that hadronises, jets created by the hadronisation of b -quarks are studied in detail due to the unusual properties of the b -quark decay. The enormous mass of the top quark means that it has an extremely short lifetime, having a predicted lifetime of 5×10^{-25} s [81]. This short lifetime means that the top quark doesn’t have enough time to hadronise before it decays, leading to all of the top quark decays taking the form shown in Equation 5.6.

$$t \rightarrow W^\pm q \quad (q = d, c, b) \quad (5.6)$$




Fig. 5.6 A plot showing the results of the anti- k_t jet reconstruction algorithm. The x, y -plane shows the distribution in rapidity-azimuthal angle space of cells in one layer of the calorimeter and the z -axis shows the energy in each cell. The cells are coloured based on which jet they are reconstructed as. Image retrieved from Ref. [80].

Further, the majority of top quark decays produce b -quarks. Consider Equation 5.7.

$$\frac{\Gamma(t \rightarrow W^\pm b)}{\Gamma(t \rightarrow W^\pm q)} \quad (5.7)$$

This is the ratio of the branching ratio of the top quark decaying into a b -quark to the branching ratio of the top quark decaying into any quark, and it measures the percentage of $t \rightarrow W^\pm q$ decays that produce a b -quark. The best measured value of this ratio is 0.957 ± 0.034 [19], so whenever a top quark decays into another quark, most of the time that other quark will be a b -quark.

In contrast to top quarks, b -quarks don't have very short lifetimes. In fact, the opposite is true. The b -quarks create b -hadrons¹⁰ which have unusually long lifetimes compared to other particles. Further, b -quarks are intimately tied with several important physics processes being studied, such as measurement of the elements of the CKM-matrix. Exploiting the long lifetimes of b -hadrons to study interesting physics processes is an intriguing prospect and something that the ATLAS collaboration actively pursues.

When a b -hadron finally hadronises and produces a jet, the origin point of the jet is displaced from the primary vertex. This property is exploited by the ATLAS

¹⁰A “ b -hadron” is simply a hadron containing a b -quark.

reconstruction software, which has algorithms that analyse jets and tags them as coming from a b -hadron if the origin point of the reconstructed jet is displaced from the primary vertex. Jets that are identified as coming from the decay of b -hadrons are called “ b -jets”, and the process of identifying them is known as “ b -tagging”.

There are many different b -tagging algorithms used by the ATLAS collaboration, and the one being used in the analysis presented by this thesis is known by the opaque name of “MV2c10” [82]. This algorithm is part of a class of algorithms known as “multi-variate classifiers”, hence the “MV2” part. These algorithms use multiple variables in a boosted decision tree to decide whether a jet is a b -jet, a c -jet or a light-flavour (ie, coming from u -, d - or s -hadrons or gluons) jet. The BDT is trained on a background sample containing c -jets and light-flavour jets and a signal sample of b -jets. Once the BDT has “learned” to tell the difference between them, it can be used to identify b -jets in data events.

Like the selections in the electron and muon reconstruction algorithms, the MV2 algorithm uses different operating points to achieve different b -tagging efficiencies. The specific value of the multi-variate classifier used to get certain efficiencies vary from analysis to analysis but the efficiencies of the operating points remain the same; they are 60%, 70%, 77% and 85%.

Finally, the MV2 algorithm allows analysis teams to choose the fraction of the background that is made up of c -jets. There are three variants of MV2 algorithm depending on the choice: MV2c00, MV2c10 and MV2c20. These correspond to background c -jet percentages of 0%, 10% and 20%, respectively.

5.2.4 Missing Energy

While most of the particles produced in collisions at the LHC interact with the ATLAS Detector in some way, there are those that don’t. Neutrinos pass through the detector without interacting with it, as would the LSP mentioned in Chapter 2. For any pair of colliding protons the total amount of momentum in the z -direction is unknowable, but the amount of momentum in the transverse plane is negligible. By conservation of momentum, the sum of the momenta of all of the particles coming from a single collision must be zero in the transverse plane. This means that if an event is reconstructed, but the momenta of the particles in that event don’t all sum to zero, then that event must

have had at least one invisible particle present. Equation 5.8 shows a naïve way of calculating this missing transverse energy¹¹.

$$E_T^{\text{miss}} = - \sum_{i \in \{\text{Measured Objects}\}} p_{T,i} \quad (5.8)$$

In practice, every event contains a combination of fully reconstructed objects and leftover tracks and calorimeter clusters. These two kinds of contributions are known as “hard terms” and “soft terms” respectively, and both must be taken into account when calculating the missing transverse momentum. This is given in Equation 5.9.

$$E_{x(y)}^{\text{miss}} = - \sum_{i \in \{\text{Hard Objects}\}} p_{x(y),i} - \sum_{j \in \{\text{Soft Signals}\}} p_{x(y),j} \quad (5.9)$$

The hard term is the sum of momentum contributions from all of the fully reconstructed objects. Equation 5.10 gives the most general form of this equation.

$$E_{x(y)}^{\text{miss, Hard Term}} = - \sum_{\text{Selected Electrons}} p_{x(y)}^e - \sum_{\text{Selected Photons}} p_{x(y)}^\gamma - \sum_{\text{Selected Taus}} p_{x(y)}^\tau - \sum_{\text{Selected Muons}} p_{x(y)}^\mu - \sum_{\text{Selected Jets}} p_{x(y)}^{\text{jets}} \quad (5.10)$$

After the reconstruction procedure is complete, there are still likely to be tracks and clusters in the event that don’t belong to any reconstructed particle. These terms must be incorporated into the E_T^{miss} soft term so that they are accounted for when calculating the E_T^{miss} value¹².

With all of these factors fully understood, an equation can be written down which fully encompasses all of the objects that are considered when calculating E_T^{miss} . This is given in Equation 5.11.

$$E_{x(y)}^{\text{miss}} = - \underbrace{\sum_{\text{Selected Electrons}} p_{x(y)}^e - \sum_{\text{Selected Photons}} p_{x(y)}^\gamma - \sum_{\text{Selected Taus}} p_{x(y)}^\tau - \sum_{\text{Selected Muons}} p_{x(y)}^\mu - \sum_{\text{Selected Jets}} p_{x(y)}^{\text{jets}}}_{E_{x(y)}^{\text{miss, Hard Term}}} - \sum_{\text{Unused Objects}} p_{x(y)}^{\text{track}} \quad (5.11)$$

¹¹The terms “missing transverse energy” and “missing transverse momentum” are interchangeable and used as such in this thesis.

¹²It should be noted that the soft-term doesn’t use contributions from *all* of the objects in a given event. There are always stray tracks and clusters that don’t originate from the primary vertex and so are usually ignored by the E_T^{miss} reconstruction algorithm.

PART III:

PHYSICS ANALYSIS

The scientific man does not aim at an immediate result. He does not expect that his advanced ideas will be readily taken up. His work is like that of the planter - for the future. His duty is to lay the foundation for those who are to come, and point the way. He lives and labours and hopes.

Nikola Tesla

Chapter 6

Analysis Strategy

The remainder of this thesis will now present the search for supersymmetry that served as the core of the work I performed during my Ph.D. The search is for pair-produced electroweakinos (specifically, the $\tilde{\chi}_1^\pm$ and the $\tilde{\chi}_2^0$) which decay into a vector boson pair (a $W^\pm$ and Z^0 , respectively) which themselves decay into either two or three leptons, accompanied by jets and missing transverse energy.

6.1 The Search for Supersymmetry

As was discussed in Chapter 2, there are many extensions to the Standard Model that make use of supersymmetry and one needs to be chosen in order to have a concrete model that can be tested. So while the search presented in this thesis is technically a “search for supersymmetry”, it is more accurate to specify the SUSY model being studied and say that this search is for “pair-produced electroweakinos that are predicted by the MSSM”.

While the MSSM is a model that I, as an experimentalist, can *technically* test, there is another issue. Taking the MSSM as a realistic theory, with R -parity conservation and soft SUSY-breaking, means that it has 105 free parameters¹ [23]. Trying to draw conclusions on the full 105-dimensional parameter space of the MSSM isn’t possible due to the inability to constrain all of the degrees of freedom.

¹A “free parameter” in the context of quantum field theory, is a parameter whose value can’t be derived from the theory itself and must be measured by experiment.

6.2 Simplified Models

This problem is solved through the use of “simplified models” [83]. Simplified models are specially designed SUSY models that have most of the particle content of the MSSM removed. The particle content of a typical simplified model consists of all of the Standard Model particles and a few MSSM particles. These models are designed in a way that allows them to have a reduced particle content (which makes them easier to test) while still providing a rough description of SUSY physics² [83].

This analysis makes use of one of these simplified models. The model in question contains three SUSY particles: two neutralinos ($\tilde{\chi}_2^0, \tilde{\chi}_1^0$) and one chargino ($\tilde{\chi}_1^\pm$). These particles are said to undergo wino-like pair-production and decay into the Standard Model vector bosons and a bino-like LSP. The LSP then escapes undetected while the vector bosons decay, with the Z^0 decaying into leptons ($Z^0 \rightarrow \ell\ell$) and the $W^\pm$ decaying into either leptons ($W^\pm \rightarrow \ell\ell$) or quarks ($W^\pm \rightarrow qq$). The detector records these leptons, quarks and LSPs as leptons, jets and missing energy. These detector-reconstructed objects form the “final state” of this process. The final state of an event is the set of objects that the detector reconstructs after all the particles in an event have decayed into stable states.

Two physics scenarios are considered in this search, with each one based on whether the $W^\pm$ decays hadronically or leptonically. The first scenario involves the Z^0 decaying into a pair of leptons ($Z^0 \rightarrow \ell^+\ell^-$) and the $W^\pm$ decaying into a pair of quarks ($W^\pm \rightarrow qq$). This scenario is given in Equation 6.1 and gives a final state with 2 leptons, 2 jets and missing energy. The second scenario involves the Z^0 decaying into a pair of leptons and the $W^\pm$ decaying into a lepton and its corresponding neutrino ($W^\pm \rightarrow \ell^\pm\nu_\ell$). This scenario is given in Equation 6.2 and gives a final state with 3 leptons and missing energy.

$$\tilde{\chi}_1^\pm \tilde{\chi}_2^0 \rightarrow W^\pm Z^0 \tilde{\chi}_1^0 \tilde{\chi}_1^0 \rightarrow qq\ell^\pm\ell^\mp\tilde{\chi}_1^0\tilde{\chi}_1^0 \quad (6.1)$$

$$\tilde{\chi}_1^\pm \tilde{\chi}_2^0 \rightarrow W^\pm Z^0 \tilde{\chi}_1^0 \tilde{\chi}_1^0 \rightarrow \ell^\pm\nu_\ell\ell^\pm\ell^\mp\tilde{\chi}_1^0\tilde{\chi}_1^0 \quad (6.2)$$

Two additional physics scenarios are considered in this search. These additional scenarios are the same as the two described above, but with initial state radiation

²In practice, a single simplified model won’t cover all of the interesting SUSY phenomenology. Instead, many models are defined, each with different particle content. Each of these models is used in an analysis like this one and the expectation is that doing this with enough models will cover all of the interesting SUSY phenomenology.

produced in association with the decay. This **Initial State Radiation** (“ISR”) takes the form of a quark or gluon that is directly produced by the collision of two protons and recoils away from the chargino-neutralino system. This quark/gluon then hadronises and is reconstructed as a jet in the detector. Figure 6.1 shows the Feynman diagrams of all of four of the processes this analysis studies.




Fig. 6.1 Feynman diagrams for the physics scenarios studied in this paper: (a) $\tilde{\chi}_1^\pm \tilde{\chi}_2^0$ decay via leptonically decaying $W^\pm$ and Z^0 bosons, (b) $\tilde{\chi}_1^\pm \tilde{\chi}_2^0$ decay via a hadronically decaying $W^\pm$ and a leptonically decaying Z^0 , (c) $\tilde{\chi}_1^\pm \tilde{\chi}_2^0$ decaying via leptonically decaying $W^\pm$ and Z^0 bosons produced in association with an initial state radiation jet, (d) $\tilde{\chi}_1^\pm \tilde{\chi}_2^0$ decay via a hadronically decaying $W^\pm$ and a leptonically decaying Z^0 produced in association with an initial state radiation jet.

Finally, the simplified model that provides these four scenarios sets the branching fractions of both the $\tilde{\chi}_1^\pm \rightarrow W^\pm \tilde{\chi}_1^0$ and $\tilde{\chi}_2^0 \rightarrow Z^0 \tilde{\chi}_1^0$ decays to 100%.

6.3 Data and Simulation

Chapters 4 and 5 have both discussed the fact the detector only “sees” particles as hits in the trackers and energy deposits in the calorimeters. This fact becomes even more egregious when trying to discover new particles, since the detector won’t “see” the $\tilde{\chi}_1^\pm$, $\tilde{\chi}_2^0$ or the $\tilde{\chi}_1^0$. Instead, it will see the final states, which are comprised of leptons, jets and missing energy. This is problematic because there are many other processes (called “background processes”) that have the same final state as the physics scenarios given above (called “signal processes”) and each data event can’t know if their reconstructed final state came from one or the other.

The way to overcome this issue is to simulate the background processes (which are, by definition, well understood). This way the data can be compared to a simulated background in a statistical counting experiment: if there are a statistically significantly larger number of data events than Monte Carlo events then it can be said that some physics beyond that simulated by the Monte Carlo exists. This is the way that all searches for new physics are done, and this search is no exception. The rest of this chapter will discuss this method and its nuances in detail.

Firstly, while there are millions of data and background events, the signal processes are much rarer, occurring only a few hundred times over the course of the entire data-taking period. Applying these numbers to the counting experiment defined above, the data would only have a few hundred more events than the background. Such an excess doesn’t even approach statistical significance.

The naïve solution to this is to discard background events in the same way that the trigger discards uninteresting events. However, there are two issues with this idea. First, there is no obvious way to decide which requirement ought to be imposed. Imposing a requirement with the goal of removing more background events than signal would require that there is some value of some quantity where there is an excess of data events over background, but given that there are only a few hundred signal events *in total*, even if *all* of these events had the same value for that quantity, such an excess would still be indistinguishable from a statistical fluctuation. Therefore, using *just* the data and Monte Carlo background to select signal events is impossible. Secondly, not only is it impossible, but it also isn’t good practice. Using the data to select signal events biases the analyser, which causes any statistical significances on discovery to be overstated (more on this in Sections 6.5, 6.6 and 6.7).

The solution to this is to once again use simulation. The simplified models provide everything that is required (more on this in Section 6.5) in order to simulate what

the signal process would look like should it occur. This means that the signal can be compared with the background so that requirements can be defined which remove background events but keep signal events. The signal sample being used is a simulation, so these requirements can be defined without looking at the data (this is known as the analysis being “blinded”).

The most common method of simulating both signal and background processes is to use a Monte Carlo algorithm and is the method by which the signal and background are generated for this search. Chapter 7 gives an overview of the data used in this search and the settings of the detector (in particular the trigger and the reconstruction) during the data-taking period as well a summary of the different backgrounds and the Monte Carlo generators used to construct them.

6.4 Kinematic Variables

Even at the highest level, the trigger only has access to a few variables (the multiplicities of different reconstructed particles in an event, lepton charge and flavour, particle transverse momentum and energy) in its decision making process. These variables on their own aren’t sufficient to separate signal and background processes to the degree required to make a discovery.

These basic variables are thus used to construct more complex kinematic variables. These variables are much more useful and flexible than the variables that the reconstruction algorithm defines. For example, while a particles position in $\eta \times \phi$ space isn’t of particular interest, the difference in the ϕ values (ie, the azimuthal angle between their momentum vectors) between two particles is very useful. Likewise, the transverse momenta of a given particle might be useful in passing some quality requirement, but the invariant mass of a two-particle system can be used to make an educated guess of the particle decay they came from.

The basis of kinematic variables is expanded even further through the use of the “Recursive Jigsaw Reconstruction” technique. This is a novel method that analyses each event and constructs special variables by using boosts to different reference frames. Chapter 8 will introduce this method of constructing variables, as well as list the kinematic variables (both “standard” and Recursive Jigsaw) that are used.

6.5 Signal, Control and Validation Regions

Each event recorded by the ATLAS detector has unique values for all of the variables described in Chapter 8. This list of variables themselves can be thought of as a giant, multi-dimensional phase space, where each axis of the space corresponds to the real values that a single variable can take. In this view, an event is simply a point in this phase space whose coordinates correspond to the values that the variables take in that event.

Events are chosen or discarded based on the values that these variables take. For instance, suppose that the requirement “ $p_T^{\ell_1} > 25 \text{ GeV}$ ” was imposed upon some dataset. This means that any event whose leading lepton had a transverse momentum greater than 25 GeV would be kept, while the rest of the events would be discarded. In terms of the phase space, this requirement defines a region of the phase space within which all events pass the requirement. The technical term for this requirement is “cut”, since it “cuts out” events that are to be discarded, and imposing the cut on some variable is known as “cutting on” that variable.

There is no reason why more than one cut can’t be applied to the same dataset. Suppose the following cuts were applied:

1. The event must have two leptons, which have opposite electrical charges but the same flavour.
2. $p_T^{\ell_1} > 25 \text{ GeV}$
3. $p_T^{\ell_2} > 25 \text{ GeV}$
4. $m_{\ell\ell} \in [80, 100] \text{ GeV}$

This choice is fairly devious. At face-value this set of cuts creates a region of phase space within which all events have two leptons who are opposite in the sign of their electric charge but with the same flavour, who both have a transverse momentum of greater than 25 GeV and when combined have an invariant mass of between 80 and 100 GeV. However, these are also the most probable properties present in the pair of leptons that emerge from the $Z^0 \rightarrow \ell\ell$ decay. By making the cuts listed above, a set of events has been created that are much more likely to contain a $Z^0 \rightarrow \ell\ell$ decay than not.

This simplified example is the genesis of the concept of a “signal region” (or “SR”). By applying a series of cuts that are consistent with the properties of some signal process, a region of phase space can be found which is enriched with events likely to contain a signal event.

In practice, the process of deciding which cuts are used is an optimisation problem. By studying the distributions of all of the variables for both the signal and the background, a signal region can be defined which contains a statistically significantly larger number of signal events than background events. The “statistically significantly” qualifier is quantified using a function [84] provided by ROOT, a data analysis framework for high energy physics [85] [86]. This function outputs a variable called “ Z_n ”, which in technical statistical terminology is a significance that was derived by applying an inverse cumulative distribution function (also known as a “quantile function”) to a p -value. Equation 6.3 shows this.

$$Z_n = D^{-1}(p) = \sqrt{2} \operatorname{erf}^{-1}(2p - 1) \quad (6.3)$$

Here:

- $D^{-1}(p)$ is the quantile function
- p is the p -value from which Z_n is derived
- $\operatorname{erf}^{-1}(x)$ is the inverse of the error function

The p -value is derived from the numbers of signal and background events and the error on these values (Equation 6.4).

$$p = \frac{\int_0^{1+\tau} t^{N_s+N_b-1} (1-t)^{\frac{1}{\varepsilon_R^2}} dt}{\int_0^1 t^{N_s+N_b-1} (1-t)^{\frac{1}{\varepsilon_R^2}} dt} \quad (6.4)$$

Here:

- N_S is the number of signal events
- N_B is the number of background events
- $\varepsilon_R = \frac{\varepsilon_T}{N_B}$ is the relative error
- $\varepsilon_T = \sqrt{\varepsilon_{stat}^2 + \varepsilon_{syst}^2}$ is the total error
- ε_{stat} is the statistical uncertainty
- ε_{syst} is the systematic uncertainty³
- $\tau = \frac{1}{N_B \varepsilon_R^2}$

The significance “ Z_n ” serves to measure the sensitivity of the signal region to the signal model. If Z_n is greater than some predetermined value (in the case of this

³For the purposes of determining significances for signal models, this is taken to be 30% of the statistical uncertainty. The more statistically rigorous analysis described later in this chapter uses the actual systematic uncertainties derived from theory and the experimental apparatus.

analysis: 1.6) then the signal region is considered sensitive enough to be able to make a discovery.

The signal Monte Carlo is defined by the parameters in the simplified model, but there is some ambiguity. While the simplified model lists the SUSY particles, decays and branching ratios, the masses of the SUSY particles aren't known and could take a wide range of values. To account for this, a set of signal samples are generated, where each sample has a different mass for the $\tilde{\chi}_1^\pm$, $\tilde{\chi}_2^0$ and $\tilde{\chi}_1^0$, although to simplify the $\tilde{\chi}_1^\pm$ and $\tilde{\chi}_2^0$ are taken to be mass degenerate. The values of these masses are taken to be multiples of 50 GeV. The result of this process is the creation of the so-called “signal-grid”, where the masses of the SUSY particles can be thought of as occupying points on a plane defined by $m_{\tilde{\chi}_1^\pm, \tilde{\chi}_2^0}$ and $m_{\tilde{\chi}_1^0}$ (Figure 6.2).




Fig. 6.2 A plot of the masses of the $\tilde{\chi}_1^\pm/\tilde{\chi}_2^0$ and the $\tilde{\chi}_1^0$ with the points that correspond to the chosen masses in the signal samples.

Once a suitable signal region is created, its sensitivity to other signal models can be found. Doing this for every signal model associates a Z_n value to every point on the signal-grid. An interpolation can be performed on these Z_n values, giving a continuous distribution of Z_n values that is a function of $m_{\tilde{\chi}_1^\pm/\tilde{\chi}_2^0}$ and $m_{\tilde{\chi}_1^0}$. There is now a subset of the mass-plane, a neighbourhood of points around the signal model that the signal region was optimised on, within which the Z_n values are all above 1.6. This means that the signal region is sensitive to all signal models within that neighbourhood. This is usually shown by plotting the mass-plane and using a colour scale to represent the

Z_n value. These are called “sensitivity plots” and the ones corresponding to the signal regions used in this analysis are given in Subsections 9.2.3 and 9.3.3.

There are two possibilities for each of the models in this neighbourhood: discovery or exclusion. Their fates are ultimately decided when the analysis is unblinded and the data and background are compared in the signal region. A discovery is claimed if there is a statistically significant excess of data events in the signal region. If there is no such excess then the signal models in the neighbourhood are rejected and limits are placed on the potential masses of the sparticles in the simplified model.

This method of placing limits on the masses of the sparticles in the simplified models is the norm when a search for new particles doesn’t discover anything. These limits take the form of a curve, known as an “exclusion contour”, in the mass-plane that bounds the set of all excluded models. With each analysis that is performed new limits are found and the exclusion contour in the mass-plane expands. The most recent analysis that targets the same simplified model as the analysis presented in this thesis is given in [87]. Section 6.7 will describe the statistically rigorous ways in which the data and background is compared in the signal regions and how any excesses (or lack thereof) are translated into discoveries or exclusion limits.

The sensitivity plots described earlier in this section serve as prototype exclusion contours: estimating the areas that would be excluded if no new particles are found. These sensitivity plots don’t have the same level of statistical rigour that the exclusion contours do, instead using rough estimates for parameters like the total systematic uncertainty.

In order for the significances calculated for each signal region to be accurate, all of the Monte Carlo backgrounds must be modelled well so that their distributions have the correct shapes and their yields are accurate. To check this another set of regions are defined, called “control regions” (or “CRs”), that are enriched with background events.

Each control region is made to study a single background and is defined in a way that maximises the contribution from that one background while minimising the contribution from other backgrounds and the signal. Control regions are also designed to be statistically independent of the signal regions while being “close” to them in phase space. This means that there is no overlap in events between the two regions: no event occupies both a signal region and a control region. It should also be noted that not every background contribution (listed in Table 7.1) needs to have a control region defined for it. Some backgrounds (for example, the Higgs background) have a negligible contribution in the signal regions and so don’t need to be modelled.

Once the control regions are defined a data-to-Monte Carlo comparison is done in each region. If the data and Monte Carlo distributions match then the Monte Carlo backgrounds are validated and the analysis proceeds. If this isn't the case, this usually means that the Monte Carlo has some limitation that doesn't allow it to model a given background accurately and some other method is needed. The background yields within the control regions are now used to estimate the yields of the same backgrounds in the signal regions, but before this can be done, the idea that the background estimate from a control region *can* be used to estimate the same background outside of that control region must itself be validated.

For this process a third type of region is defined: the “validation region” (or “VR”). These regions are placed “in-between” the signal and control regions, but are orthogonal⁴ to both. The purpose of the validation regions is to check whether the Monte Carlo that predicts the background in the control region still works in regions with more than one background contribution and signal contamination. This is done in the same way as the control region check: the Monte Carlo background and the data are compared. If their distributions and yields match then the analysis proceeds as normal. If they don't match, then the fit used to predict the background contribution in the validation region needs to be studied further.

A schematic view of the propagation of background estimates through the various region types is given in Figure 6.3.

Chapters 10 and 11 document the control and validation regions used in this search, respectively. Additionally, Chapter 10 describes how several backgrounds are modelled poorly by Monte Carlo and document the alternative methods that are used to model them.

6.6 Statistics and Fits

Section 6.5 gave a qualitative description of the process by which searches for new physics are performed. This section will now describe the quantitative, statistically rigorous method that is used in practice.

All statistical analysis is done using a package called HistFitter [88], which provides particle physicists with a set of probability distribution functions and fitting procedures from which the results are derived.

⁴If two (signal, control or validation) regions are “orthogonal” it means that there is no overlap between, which means that they are statistically independent.




Fig. 6.3 A schematic view of the analysis strategy, showing multiple control, validation and signal regions. The dashed lines represent the different bins in the regions that have them.

The signal and control regions are statistically independent, so they can be modelled using separate PDFs and then combined into a simultaneous fit. HistFitter shares all PDF parameters across all regions, which allows it to use every signal and background component, as well as the systematics, consistently in all regions.

Results are obtained by using three different types of likelihood fits: background-only fits, model-dependent fits and model-independent fits.

The background-only fits are done first. These are the fits that are used to estimate the background contributions in all of the regions. These fits are only performed in the control regions and only use background samples. The control regions are defined to have minimal signal contamination so the background-only fits assume the control regions to be free of it. Similarly, each control region has one dominant background source provide most of the events, and that is the background that is normalised by the fitting procedure in that region.

The background-only fit outputs a scale factor, that is applied to the Monte Carlo simulated background to normalise it to the observed data event count in the region. Since these fits are performed in regions with minimal signal contamination and deliberately neglect any signal contribution, the background estimations derived from this method are independent from any signal observations in the signal regions.

These background estimations are then extrapolated into the signal regions to predict the background contributions in them. This is done through the use of a “transfer factor”, T_{FT} . Equation 6.5 defines the extrapolation.

$$N_B(\text{Est}, \text{SR}) = N_B(\text{Obs}, \text{CR}) \times \left[\frac{\text{MC}_B(\text{Obs}, \text{SR})}{\text{MC}_B(\text{Obs}, \text{CR})} \right] \quad (6.5)$$

Here,

- $\text{MC}_B(\text{Obs}, \text{CR})$ is the observed amount of Monte Carlo simulated events from background B in the control region
- $\text{MC}_B(\text{Obs}, \text{SR})$ is the observed amount of Monte Carlo simulated events from background B in the signal region
- $N_B(\text{Obs}, \text{CR})$ is the observed amount of data events in the control region
- $N_B(\text{Est}, \text{SR})$ is the estimated amount of events coming from background B in the signal region.

The ratio appearing in square brackets in Equation 6.5 is the transfer factor between the control and signal regions.

Before the background estimations can be done in the signal regions, they must first be validated in the validation regions. This is done using the transfer factor method above, but with an equivalent equation that relates the control and validation regions. If a satisfactory agreement can be found between the normalised background predictions and the number of observed events in the validation regions, then the extrapolation is further extended to the signal regions.

6.7 Unblinding and Results

Up until this point the entire search has been performed without looking at the data yields and distributions in the signal regions, the analysis is still blinded. Performing the model-dependent and -independent fits requires looking at the data in the signal regions, also known as “unblinding” the analysis. This procedure of physicists blinding themselves is useful because being blinded until the last possible moment means that the analysis isn’t biased by any premature signal-based predictions.

The model-independent fits provide an array of methods for quantitatively comparing the signal and background yields in the signal region that aren’t dependent on any given signal model. The model-independent fits can:

1. Quantitatively measure how well the background prediction and observed yields agree in the signal region
2. Quantify the number of possible new physics events in the signal region
3. Quantify the statistical significance of the excess

This fit proceeds in a similar way to the background-only fit, with the difference being that the number of observed events in the signal region and a “dummy” signal model parameter (constrained to be non-negative) that represents the expected signal strength for new physics processes are both added to the fit. The observed/expected 95% confidence level (“CL”) upper limit on the number of new physics events for each signal region (denoted as S_{Obs}^{95} and S_{Exp}^{95} , respectively) are derived using the CL_s procedure [89]. It should be noted that this method neglects signal contamination in the control region.

Normalising these limits to the integrated luminosity of the data sample gives the 95% confidence level on the visible cross-section of the new physics process ($\langle\epsilon\sigma\rangle_{\text{Obs}}^{95}$). The visible cross-section is defined as the product of the processes cross-section, the detector acceptance and the detector efficiency.

Finally, the model-independent fit is used to perform the background-only hypothesis test, which is used to calculate the significance, σ , of the excess. The threshold for discovery of a new particle is set at 5σ significance⁵. If the fit returns a significance between 2σ and 5σ it doesn’t qualify as a discovery, but one may be able to be made with further analysis. If the model-independent fit gives a lower significance value, below 2σ then the excess is indistinguishable from a statistical fluctuation and the signal model is excluded using a model-dependent fit.

The model-dependent fits are used to study a single signal model. This fit behaves similarly to the model-independent fit, but with the yields in both the signal and control regions taken into account. Signal yield systematic uncertainties are also included in the fit, with correlations between signal and background systematic uncertainties being taken into account if applicable. This fit produces limits on the cross-sections of the new physics processes, which are then transformed into limits on the masses of the sparticles $\tilde{\chi}_1^\pm/\tilde{\chi}_2^0$ and $\tilde{\chi}_1^0$. These limits are shown as Chapter 13 as exclusion contours.

Figure 6.4 gives a schematic diagram of the fitting process described in Sections 6.6 and 6.7.

⁵The 2012 Higgs discovery had a 5.9σ expected excess [25, 26].




Fig. 6.4 Schematic view of the HistFitter analysis flow.

Chapter 7

Analysis Overview

This chapter will summarise the data and Monte Carlo simulation used in this analysis. It will also summarise the trigger and reconstruction used during the data-taking period.

7.1 Data

This analysis uses data that were collected by the ATLAS detector during its 2015 and 2016 data-taking runs. During this time the collider was running at an energy of $\sqrt{s} = 13$ TeV, with a collision rate of 40 MHz with a bunch spacing of 25 ns.

The 2015 data-taking period ran from May to November of that year. It had a peak instantaneous luminosity of $L = 5.2 \times 10^{33} \text{ cm}^{-2}\text{s}^{-1}$ and a mean pile-up of $\langle\mu\rangle = 14$ interactions per bunch crossing. For that year, the integrated luminosity provided by the LHC was 4.2 fb^{-1} and integrated luminosity recorded by ATLAS was 3.9 fb^{-1} [67].

The 2016 data-taking period ran from April to October of that year. It had a peak instantaneous luminosity of $L = 1.37 \times 10^{34} \text{ cm}^{-2}\text{s}^{-1}$ and a mean pile-up of $\langle\mu\rangle = 24$ interactions per bunch crossing. For that year, the integrated luminosity provided by the LHC was 38.5 fb^{-1} and integrated luminosity recorded by ATLAS was 35.6 fb^{-1} [67].

Figure 7.1 shows the cumulative luminosity delivered by the LHC and recorded by ATLAS over 2015 (Figure 7.1a) and 2016 (Figure 7.1b).

Figure 7.2 shows the delivered integrated luminosity in terms of pile-up for the 2015 run, the 2016 run and for both runs combined.




Fig. 7.1 Plots showing the cumulative integrated luminosity delivered by the LHC (green) and recorded by ATLAS (yellow) during 2015 (a) and 2016 (b). Plots retrieved from Ref. [67].




Fig. 7.2 Plot showing the delivered integrated luminosity in terms of the pile-up of the collisions they came from. Plot retrieved from Ref. [67].

7.2 Monte Carlo

At ATLAS a Monte Carlo (“MC”) simulation is used to simulate both the signal and background processes. This simulation process encompasses everything that occurs in a collision: the proton-proton collision, the interactions of the quarks and gluons therein, the products of a collision, their interactions with the detector and any pile-up collisions occurring as well. Each of these tasks is performed by a separate part of the simulation software. The protons in the beam are simulated using parton distribution functions. These are probability distributions for the momenta of the partons inside of the protons in the beam, a quantity that is essential for calculating cross-sections. After two protons collide the partons that make them up interact. This requires a

matrix element (“ME”) to be calculated to describe the products of the collision. Any free partons hadronise, a process which is simulated using special software that mimics parton showers. Finally, all of the particles that emerge from an event have their interactions with the detector simulated [90] with a package based on GEANT4 [91]. A schematic diagram of a proton-proton collision with the simulation methods of various processes labelled is given in Figure 7.3.




Fig. 7.3 Schematic diagram of a proton-proton collision, highlighting how the Monte Carlo simulates the various interactions. Image retrieved from Ref. [92].

A variety of different Monte Carlo generators are used within the ATLAS collaboration. Each of these generators has their own strengths and weaknesses, with some generators being better suited to simulating certain backgrounds than others. This analysis uses several Monte Carlo generators such as MADGRAPH, SHERPA and POWHEG to simulate different backgrounds, and the rest of this section will document which generators are used for which backgrounds and why.

The simulation of the production of Z^0 bosons in association with jets [93] was simulated with the SHERPA v2.2.1 generator [94]. The NNPDF3.0NNLO [95] PDF was used with the in-built SHERPA parton shower tuning. The ME was calculated for up to two partons at next-to-leading order (“NLO”) and with up to two additional partons at leading-order (“LO”) using the COMIX [96] and OPEN LOOPS [97] ME generators and then merged with the SHERPA parton shower (“PS”) algorithms [98] using the

ME+PS@NLO prescription [99]. The Z+Jets background was also simulated using a data-driven method (described in Chapter 10) that was based on γ +Jets events that were generated using the SHERPA v2.1.1 generator with the CT10 [100] PDF set.

The POWHEG-BOX v2 generator [101] was used to generate the $t\bar{t}$ events as well as the single-top events in the Wt and s -channels. The single-top t -channel [102] production was done with the POWHEG-BOX v1 generator [103], with MadSpin [104] being used to simulate the top-quark decays in a way that preserved spin correlations. For all processes the same CT10 [100] PDF set as above was used to calculate the ME while the parton shower, fragmentation and underlying event was generated using the PYTHIA 6.428 [105] with the CTEQ6L1 [106] PDF set and a set of tuned parameters called the “PERUGIA2012” tune [107]. The $t\bar{t}$ and Wt -channel single-top events were normalised to cross-sections that were calculated at next-to-next-to-leading-order plus next-to-next-to-leading-logarithm (“NNLO+NNLL”) [108, 109, 110, 111] accuracy, while the s - and t -channel single-top events were normalised to NLO cross-sections [112, 113]. The production of Zt events was done with the MG5_aMC@NLO 2.2.1 [114] generator at LO with the CTEQ6L1 PDF set.

The final top-quark based background this analysis considers is the $t\bar{t}$ +EW ($t\bar{t} + W$, $t\bar{t} + Z$ or $t\bar{t} + WZ$ [115]) with up to two ($t\bar{t} + W$, $t\bar{t} + Z(\rightarrow \nu\nu/qq)$), one ($t\bar{t} + Z(\rightarrow \ell\ell)$) or zero ($t\bar{t} + WW$) extra partons included in the ME. It is generated using the MG5_aMC@NLO 2.2.2 generator (2.2.3 for $t\bar{t} + Z/\gamma$) at LO, interfaced with the PYTHIA 8.186 [116] and all processes are normalised to their respective NLO cross-sections [117, 118]. All samples that use a top-quark have its mass set at 172.5 GeV.

The diboson processes (ZZ , WZ , WW) [119] were simulated using SHERPA 2.2.1 generator and contain off-shell contributions. The processes with 4 charged leptons (4ℓ), three charged leptons and a neutrino ($3\ell + \nu$) and 2 charged leptons and two neutrinos ($2\ell + 2\nu$) had their matrix elements contain all diagrams with four electroweak couplings and were calculated with up to one (for the 4ℓ and $2\ell + 2\nu$) or no ($3\ell + \nu$) extra partons at NLO. All the diboson samples were also simulated with up to three additional partons at LO using the COMIX and OPEN LOOPS ME generators, which were merged with the SHERPA parton shower using the ME+PS@NLO prescription. The diboson events were normalised to their NLO cross-sections [120, 121].

Samples of a leptonically decaying vector boson and a photon ($V\gamma$) were generated at LO using SHERPA 2.1.1 [94]. Matrix elements including all diagrams with three electroweak couplings were calculated with up to three partons at LO and merged with the SHERPA PS [122], again, according to the ME+PS@NLO prescription [123]. The CT10 PDF set is used along with the SHERPA in-built PS tuning.

Triboson processes (WWW , WWZ , WZZ and ZZZ) were simulated using SHERPA 2.2.1 with the matrix elements being calculated at LO with up a single additional parton. Triboson events were normalised to their LO cross-sections [124].

Since the discovery of the Higgs boson in 2012 [25, 26] it has been demoted to a lowly background source in searches for new physics. The Higgs boson production processes (such as vector boson fusion, or “VBF”, associated vector-boson production, VH , or gluon-gluon fusion) were all generated using POWHEG-BOX v2 [102] interfaced with PYTHIA 8.186 and normalised to cross-sections calculated at NNLO with soft gluon emission effects added at NNLL accuracy. The $t\bar{t}H$ samples are produced using MG5_aMC@NLO 2.2.2 interfaced with HERWIG 2.7.1 [125] and normalised to the NLO cross-section [126]. All Higgs boson samples assume a Higgs mass of 125 GeV.

Finally, the signal samples were generated from LO matrix elements with up to two additional partons using MADGRAPH v2.2.3 [127] interfaced with PYTHIA 8.186. Pythia used the “A14” parameter tune [128] for the modelling of the SUSY decay chain, parton showering, hadronisation and modelling of the underlying event. Parton luminosities were provided by the NNPDF2.3LO PDF set [100]. The jet-parton matching follows the CKKW-L prescription [129] with the matching scale set to be equal to one quarter of the $\tilde{\chi}_1^\pm/\tilde{\chi}_2^0$ mass. Signal cross-sections were calculated at NLO with soft gluon emission effects added at next-to-leading-logarithm (NLL) accuracy [130, 131, 132, 133, 134, 135]. The nominal cross-section and the uncertainty were taken from an envelope of cross-section predictions using different PDF sets, factorisation and renormalisation scales, described in [136].

Table 7.1 summarises the signal and backgrounds simulated by Monte Carlo, including which generator was used, the cross-section normalisation, the PDFs used and the tuning of the generator.

The EvtGen v1.2.0 program [137] was used to model the decays of the b and c -hadrons in the Standard Model background samples which weren’t simulated using SHERPA. All simulated events were overlaid with multiple proton-proton collisions simulated by the soft QCD process of PYTHIA 8.186 using the A2 tune [138] and the MSTW2008LO PDF [139]. The Monte Carlo samples were generated with a variable of additional proton-proton interactions in the same and neighbouring bunch crossings and were re-weighted to match the distribution of the mean number of interactions that was observed in the data.

Table 7.1 The Standard Model background and SUSY signal Monte Carlo samples used in this analysis.

Physics Process	Generator	Cross-Section Normalisation	PDF Set	Parton Shower	Tune
SUSY Processes	MADGRAPH v2.2.3	NLO+NLL	NNPDF2.3LO	PYTHIA 8.186	A14
$Z/\gamma^*(\rightarrow \ell\bar{\ell}) + \text{Jets}$	SHERPA 2.2.1	NNLO	NNPDF3.0NNLO	SHERPA	SHERPA default
$\gamma + \text{Jets}$	SHERPA 2.1.1	LO	CT10	SHERPA	SHERPA default
$H(\rightarrow \tau\tau), H(\rightarrow WW)$	POWHEG-Box v2	NLO	CTEQ6L1	PYTHIA 8.186	A14
HW, HZ	MG5_aMC@NLO 2.2.2	NLO	NNPDF2.3LO	PYTHIA 8.186	A14
$t\bar{t} + H$	MG5_aMC@NLO 2.2.2	NLO	CTEQ6L1	HERWIG 2.7.1	A14
$t\bar{t}$	POWHEG-Box v2	NNLO+NNLL	CT10	PYTHIA 6.428	PERUGIA2012
Single-top (Wt -channel)	POWHEG-Box v2	NNLO+NNLL	CT10	PYTHIA 6.428	PERUGIA2012
Single-top (s -channel)	POWHEG-Box v2	NLO	CT10	PYTHIA 6.428	PERUGIA2012
Single-top (t -channel)	POWHEG-Box v1	NLO	CT10f4	PYTHIA 6.428	PERUGIA2012
Single-top (Zt -channel)	MG5_aMC@NLO 2.2.1	LO	CTEQ6L1	PYTHIA 6.428	PERUGIA2012
$t\bar{t} + W/WW$	MG5_aMC@NLO 2.2.2	NLO	NNPDF2.3LO	PYTHIA 8.186	A14
$t\bar{t} + Z$	MG5_aMC@NLO 2.2.3	NLO	NNPDF2.3LO	PYTHIA 8.186	A14
WW, WZ, ZZ	SHERPA 2.2.1	NLO	NNPDF30NNLO	SHERPA	SHERPA default
$V\gamma$	SHERPA 2.1.1	LO	CT10	SHERPA	SHERPA default
Triboson	SHERPA 2.2.1	NLO	NNPDF30NNLO	SHERPA	SHERPA default

7.3 Trigger Strategy

The trigger strategy used by this analysis was designed to take advantage of the presence of a lepton pair in the final states of all four of the signal processes shown in Figure 6.1. This analysis uses three dilepton triggers: a dielectron trigger, a dimuon trigger and an electron-muon trigger. The dielectron and dimuon triggers each correspond to one of the leptonic decays of the Z^0 boson: $Z^0 \rightarrow ee$ and $Z^0 \rightarrow \mu\mu$, respectively. While the lepton flavour violating decay $Z^0 \rightarrow e\mu$ doesn't occur¹ events that contain different flavour lepton pairs are still useful for this analysis (see Chapter 11). The triggers used by this analysis are summarised in Table 7.2.

Table 7.2 Summary of the trigger strategy. The dielectron channel requires two electrons, the 2015 run required them to each have $p_T > 12$ GeV while the 2016 required them to each have $p_T > 17$ GeV. The dimuon channel required both muons to have $p_T > 18$ GeV in 2015 and $p_T > 22$ GeV in 2016. Additionally, the triggers for both years have the “mu8noL1” tag, which is a special trigger for events that pass the L1 trigger with only one muon. Finally, the electron-muon channel requires $p_{T,e} > 17$ GeV and $p_{T,\mu} > 14$ GeV during both years of data-taking.

Channel	Trigger	
	Data: 2015 Run	Data: 2016 Run
Dielectron	HLT_2e12_lhloose_L12EM10VH	HLT_2e17_lhvloose_nod0
Dimuon	HLT_mu18_mu8noL1	HLT_mu22_mu8noL1
Electron-Muon	HLT_e17_lhloose_mu14	HLT_e17_lhloose_nod0_mu14

¹The current upper-limit of the branching ratio of $Z^0 \rightarrow e\mu$ is 7.5×10^{-7} [140].

7.4 Object Reconstruction and Identification

This section will describe the object reconstruction/identification process used during the 2015 and 2016 data-taking periods.

Each event has a “primary vertex” that is reconstructed as follows. It must be consistent with the luminous region and must have at least two associated tracks with $p_T > 400$ GeV. If more than one such vertex is found, the primary vertex is defined as the vertex whose associated tracks have the greatest summed transverse momentum².

This analysis divides lepton candidates (electrons and muons) into two different classes: “baseline” and “high-purity”.

Baseline muons must have $p_T > 10$ GeV, $|\eta| < 2.7$ and must pass the *Medium* identification requirement. High-purity muons must pass stricter requirements. They must have $|\eta| < 2.4$. The significance of the transverse impact parameter relative to the primary vertex must fulfil $\frac{|d_0^{\text{PV}}|}{\sigma_{d_0^{\text{PV}}}} < 3$ and the longitudinal impact parameter relative to the primary vertex must fulfil $|z_0^{\text{PV}} \sin\theta| < 0.5$ mm. Finally, high-purity muons must also fulfil the *GradientLoose* isolation requirement [75].

Baseline electrons must have $p_T > 10$ GeV, $|\eta| < 2.47$ and must pass the *Loose-AndBLayer* signal efficiency. High-purity electrons must satisfy the *Medium* likelihood requirement, must have $\frac{|d_0^{\text{PV}}|}{\sigma_{d_0^{\text{PV}}}} < 5$ and $|z_0^{\text{PV}} \sin\theta| < 0.5$ mm and must satisfy the same *GradientLoose* isolation requirement that high-purity muons must pass.

Jets are reconstructed using the anti- k_t algorithm with a jet radius parameter $R = 0.4$. Once reconstructed, the jets undergo several corrections. The energy of the jet must be corrected for energy coming from pile-up interactions [141]. Another set of energy corrections, called “jet energy scale corrections”, are derived from a comparison between data and Monte Carlo and are used to calibrate the energy of jets to the scale of their constituent particles [142]. Jets originating from pile-up are removed by the jet vertex tagger [143]. After these corrections are applied, jets with $p_T > 20$ GeV and $|\eta| < 4.5$ are designated as baseline jets. High-purity jets must pass the stricter requirement: $|\eta| < 2.4$. High-purity jets also interact with the jet vertex tagger differently, where jets with $p_T < 60$ GeV must pass a JVT cut of $\text{JVT} > 0.59$.

The MV2c10 algorithm is used to identify jets that come from b -hadrons, using the 77% signal efficiency operating point. Candidate b -tagged jets must have $p_T > 20$ GeV and $|\eta| < 2.4$.

After the above selections are made, ambiguities can arise when more than one particle is reconstructed using a given track or cluster. They are resolved as follows:

²The transverse momentum is summed as a scalar, ie, $\sum_{\text{Tracks}} |p_T|$.

1. Any electron sharing an ID track with a muon is discarded.
2. If any b -tagged jet (which for the purpose of removing ambiguities is defined using the 85% operating point) is within $\Delta R < 0.2$ of an electron candidate, the electron is discarded. The reason for this is that electrons this close to b -jets are likely to originate from the semileptonic b -hadron decay.
3. If a non- b -tagged jet is within $\Delta R < 0.2$ of an electron it is discarded since it is probably an electron-induced shower.
4. If electrons are within $\Delta R < 0.4$ of the remaining jet candidates are discarded, since they likely originate from the semileptonic decays of c - and b -hadrons.
5. Jets with less than three associated tracks that are close to a muon which has a significant fraction of the transverse momentum of the jet (the exact inequality is $p_{T,\mu} > 0.7 \sum_{\text{Jet Tracks}} p_{T,\text{Jet Track}}$) are discarded if they are within $\Delta R < 0.2$ of the muon or if the muon contains a track associated with the jet.
6. Muons within $\Delta R < 0.4$ of any remaining jet candidates are discarded to suppress muons coming from the semileptonic decays of c - and b -hadrons.

Finally, the missing transverse momentum vector $\vec{p}_T^{\text{miss}}$ (and its magnitude E_T^{miss}) is based on the transverse momenta of all electron, muon, photon and jet candidates as well as the transverse momenta of any tracks originating from the primary vertex that aren't associated with any of these objects. The missing transverse momentum is defined as the negative of the vector sum of the transverse momenta of all of these objects.

Chapter 8

Kinematic Variable Construction and the Recursive Jigsaw Method

Chapter 6 described how kinematic variables are used to construct signal, control and validation regions and this chapter will list the variables used to construct these regions. This chapter will be divided into four sections. Section 8.1 will list the “standard” variables, these are variables that don’t require any sort of special framework to construct and can be found in most particle physics analyses. Section 8.2 will introduce the Recursive Jigsaw Reconstruction technique, a method for constructing kinematic variables that have greater discriminatory power than the “standard” variables described in Section 8.1. Section 8.3 will then describe the Recursive Jigsaw techniques used study events for the analysis presented in this thesis, as well as list the kinematic variables constructed using the Recursive Jigsaw method that were used to construct the signal, control and validation regions.

8.1 Standard Variable Construction

This section will list the “standard” variables used by to construct the signal, control and validation regions.

- *Lepton Charge and Flavour*: the simplest variables available in reconstructed lepton objects, the electric charge and flavour (whether the lepton is an electron, a muon or a tau). These variables are used to determine whether a given lepton pair is from a Z^0 decay; if two leptons have opposite electric charge and the same flavour it is likely they came from a Z^0 decay, which must conserve charge and flavour [140].

- p_T^ℓ : the transverse momentum of the lepton ℓ . Since there can be multiple leptons in an event, they are usually labelled by a subscript: “ ℓ_i ”. The highest p_T lepton is called the “leading” lepton and is labelled ℓ_1 , the second highest is called the subleading lepton and is labelled ℓ_2 and the third highest is called the third-leading lepton and is labelled ℓ_3 .
- n_{Jets} : the number of jets in an event.
- p_T^j : the transverse momentum of the jet j . The jets are labelled the same way as leptons are. The leading jet is labelled j_1 , the subleading j_2 and so-on.
- $m_{\ell\ell}$: the invariant mass of a pair of leptons. This is also used to determine whether a lepton pair originated from a leptonically decaying Z^0 boson. If it did, the invariant mass of the pair would be close to the mass of the Z^0 boson: $m_{Z^0} = 91 \text{ GeV}$ [17].
- m_{jj} : the invariant mass of a pair of jets. This is used in a similar vein to $m_{\ell\ell}$, but for the hadronically decaying $W^\pm$ boson. As above, if the $W^\pm$ decayed into a pair of quarks, which then produced jets, those jets should have an invariant mass close to the mass of the $W^\pm$ boson: $m_{W^\pm} = 80 \text{ GeV}$ [16].
- m_T^W : In events with three leptons, two of the leptons are assigned to have come from the Z^0 decaying leptonically, the other comes from the $W^\pm$. The variable m_T^W is the transverse mass of this remaining lepton.
- $\min(\Delta\phi)$: The azimuthal angle between the E_T^{Miss} four-vector and the four vector of the closest jet.

8.2 Recursive Jigsaw Method

The **Recursive Jigsaw Reconstruction** [144, 145, 146] (“RJR”, sometimes just “**Recursive Jigsaw**” or “RJ”) technique is a novel method for analysing events that uses boosts between different reference frames to construct kinematic variables. The method was invented by Christopher Rogan and Paul Jackson, my supervisor during my Ph.D.

The reconstruction system described in Chapter 5 and the variables listed at the beginning of this chapter ultimately only make use of objects that are present in the final state of each event. Even the selection given at the beginning of Section 6.5, despite being constructed to select events that are likely to contain Z^0 bosons decaying to a lepton pair, only uses information from the reconstructed leptons in the final

state. This limitation is at odds with the goals of this search (as well many searches generally) which is trying to infer the presence of a particle as one of the intermediate decay states.

If the detector were able to provide physicists with perfect knowledge of every event, then those data would be able to be combined to estimate the properties of intermediate states fairly well. Unfortunately, no experiment can give perfect measurements and a variety of complications arise when attempting to estimate the properties of intermediate states using incomplete informations.

The first of these complications is “kinematic ambiguity”. If an event contains invisible particles (such as neutrinos or the hypothesised LSP) they will escape undetected and carry information, not just about themselves but also about whatever intermediate decays they came from, with them. This is a problem when trying to estimate the properties of the intermediate states because the loss of information means that any estimated quantities can take a variety of values and still be consistent with the rest of the event.

The second complication is known as “combinatoric ambiguity” and occurs in events that have more than one of a given type of particle. Since the detector is unable to tell which processes produced the particles it detects, events with multiple of a given type of particle can reconstruct the decay in several different ways, all of which would be consistent. As an example, the physics scenario given in Figure 6.1b has a final state with three leptons. Two of these come from the Z^0 decaying into an opposite-sign lepton pair. The final state can, without loss of generality, be assumed to contain two positive leptons and one negative lepton. The ambiguity is the choice of which positive lepton to assign as coming from the Z^0 decay, either choice is consistent and there is no way of knowing which one should be chosen.

The RJR technique is a methodology for analysing reconstructed events which incorporates a group of algorithms to resolve any kinematic or combinatoric ambiguities that might arise.

The RJR technique begins by imposing a “decay tree” onto the event. The decay tree is a schematic view of the reference frames that are used in the RJR. Figure 8.1 shows a generic RJ decay tree. Each of the circles corresponds to a reference frame: the grey circle corresponds to the laboratory frame wherein all measurements are made. The red circles are frames corresponding to the intermediate particle states. The frame labelled “PP” is the rest frame of the two-particle system that emerges from the proton-proton collision (taking the physics scenarios in Figures 6.1a and 6.1b as examples, the PP frame would be the rest frame of the $\tilde{\chi}_1^\pm\text{-}\tilde{\chi}_2^0$ system). The frames

labelled “ P_a ” and “ P_b ” are each the rest frames of one of the particles in the PP-frame. The blue circles labelled “ V_a ” and “ V_b ” represent rest frames of the visible decay products of the particles P_a and P_b , respectively. Similarly, the green circles labelled “ I_a ” and “ I_b ” represent the rest frames of the invisible decay products of the particles P_a and P_b , respectively.




Fig. 8.1 A general decay tree for reconstructing Recursive Jigsaw quantities.

This formulation also allows an alternate way of parameterising an event. Events can be fully described using the set of momentum four-vectors of all of the particles in the final state, but with the decay tree the set of all particle masses and decay angles also describes the event fully. These masses and decay angles are a more natural basis of variables for the RJ method since it uses masses and the angles between boosts rather than momenta to analyse decays.

Similarly, it is useful to think of the ambiguities described above in terms of the degrees of freedom associated with the decay tree. The total number of degrees of freedom in an event needn’t be the same as the number of kinematic measurements made in the reconstruction process. If there are invisible particles present in the decay tree, some of the variables that describe them kinematically may be under-constrained. Similarly, if two of the same particle appear in a decay tree, then there is a combinatoric ambiguity coming from where each particle should be assigned in the decay tree. The remaining masses and decay angles are then functions of these unknowns and analysing the event amounts to choosing values for them.

Resolving these ambiguities isn’t a new problem and the particle physics literature already has solutions to all of these. For combinatoric ambiguities, the common solution

is to find some variable to extremise with respect to the combinatoric assignments. Popular variables to extremise include invariant masses (such as $m_{\ell\ell}$) and distance metrics (such as the ΔR used in jet reconstruction). The RJR technique uses this method for resolving its combinatoric ambiguities.

Kinematic ambiguities are more complicated, with a wide variety of strategies for resolving them appearing in the literature [147, 148, 149, 150, 151]. These methods treat any quantities that are of interest as functions that are dependent on the masses/momenta of both the visible and invisible particles. To resolve kinematic ambiguities, they attempt to remove any dependence that the quantities of interest have on the properties of the invisible particles, which are unmeasured. In practice, this usually involves imposing a constraint on the masses of the particles in an event and then solving the resulting system of equations. While this may seem a straightforward solution, such systems of equations typically involve higher-order polynomials, which have multiple solutions, none of which are guaranteed to be real [152].

As an alternative to the above, under-constrained degrees of freedom that are known to be small can be ignored. This constrains multiple unknowns simultaneously but can cause issues with how accurate the rest of the observables are when this constraint is applied.

The RJR approach makes use of methods similar to the above to resolve kinematic ambiguities through the use of its “jigsaw rules”. The RJ method is set apart from the methods described above through its observation that there needn’t be only one quantity of interest in each event. Instead, it makes use of many observables in concert, with each observable depending only on a subset of the under-constrained degrees of freedom. This list of observables form the basis of RJ variables, but care must be taken during their construction. The optimal scenario would be to have all observables be maximally uncorrelated. This is achieved through careful consideration of how the under-constrained degrees of freedom affect the determination of the velocities that define the boosts between the decay frames.

The “jigsaw rules” (“JRs” or just “jigsaws”) are algorithms that resolve the unknowns in a single decay frame. After a decay tree is imposed on an event, a jigsaw is used to analyse the laboratory frame and resolve any ambiguities. The RJR then boosts to the next frame (in the case of Figure 8.1, this would be the PP frame) and applies another jigsaw. This is repeated for each frame, with each frame having a different jigsaw applied, until the ends of the tree are reached. This factorised approach means that the JRs used are interchangeable; different JRs can be applied to resolve the same unknowns, although applying different jigsaws means that the resulting

observables will behave differently. Finally, the recursive application of the JRs means that the observables that correspond to the various reference frames are all maximally uncorrelated.

The RJR framework is simply a library of JRs which, like the pieces of a jigsaw puzzle, can be assembled to analyse a decay tree. The RJR algorithm used is simply defined as the set of JRs chosen.

8.3 Recursive Jigsaw Strategy

Two decay trees are constructed, one corresponding to the 2-lepton physics scenario given in Figure 6.1a and the other corresponding to the 3-lepton physics scenario given in Figure 6.1b.




Fig. 8.2 Recursive Jigsaw decay trees: (a) corresponds to the physics scenario given in Figure 6.1a and (b) corresponds to the physics scenario given in Figure 6.1b.

In the 2-lepton physics scenario, each event is evaluated as if two SUSY particles (the $\tilde{\chi}_1^\pm$ and the $\tilde{\chi}_2^0$) are produced in the collision. Each of these particles are then assigned to their own frame and hemisphere. The $\tilde{\chi}_1^\pm$ then decays into an invisible $\tilde{\chi}_1^0$ and two leptons through an intermediate $W^\pm$. Each of the leptons is assigned to a visible frame while the $\tilde{\chi}_1^0$ is assigned to an invisible frame. Similarly, the $\tilde{\chi}_2^0$ decays into an invisible $\tilde{\chi}_1^0$ and two quarks through an intermediate Z^0 . The quarks are recorded as jets, who are each assigned to a visible frame, while the $\tilde{\chi}_1^0$ is once again assigned to an invisible frame.

In the 3-lepton physics scenario, each event is again evaluated as if two SUSY particles are produced in the collision, with each particle being assigned its own frame and hemisphere. The $\tilde{\chi}_2^0$ decays identically to its 2-lepton counterpart and has the same decay tree structure. The $\tilde{\chi}_1^\pm$ decays into a lepton and a neutrino through an intermediate $W^\pm$, as well as a $\tilde{\chi}_1^0$. The lepton is assigned its own visible frame, but since both the neutrino and the $\tilde{\chi}_1^0$ are unmeasured (and therefore their momenta are unknown) they must both be assigned to the same invisible frame.

The two ISR scenarios given in Figures 6.1c and 6.1d have a specialised decay tree constructed for them. It is given in Figure 8.3.




Fig. 8.3 The ISR decay tree for reconstructing Recursive Jigsaw quantities.

The two ISR scenarios contain an ISR jet that recoils away from the $\tilde{\chi}_1^\pm\tilde{\chi}_0^2$ system. This is modelled in the decay tree with the combined $\tilde{\chi}_1^\pm/\tilde{\chi}_0^2$ -ISR jet system being assigned the centre-of-mass (or “CM”) frame. From there, two more frames are introduced: a visible frame for the ISR jet, and an intermediate frame “S” representing all of the intermediate states of the sparticle decay tree. This frame is then divided into visible and invisible frames, containing the visible and invisible final state particles from the sparticle decay, respectively.

This decay tree serves a different purpose from the trees given in Figure 8.2. It doesn’t attempt to resolve the entirety of every event, but instead it’s used to accurately model the $\tilde{\chi}_1^\pm\tilde{\chi}_0^2$ system recoiling away from the ISR jet. This decay topology is studied for the unique insight it provides into the “low-mass/compressed” area of the signal grid (recall Figure 6.2). This area is where the “mass-splitting” between the $\tilde{\chi}_1^\pm/\tilde{\chi}_0^2$ and the $\tilde{\chi}_1^0$ is small, where the mass-splitting is defined as the mass difference between

the $\tilde{\chi}_1^\pm/\tilde{\chi}_0^2$ and the $\tilde{\chi}_1^0$: $\Delta m = m_{\tilde{\chi}_1^\pm/\tilde{\chi}_0^2} - m_{\tilde{\chi}_1^0}$. This mass-splitting is related to the amount of energy available to the vector boson-LSP system. Accurately measuring the momentum of the objects in the visible frame allows accurate estimation of the momentum of the vector boson, which in turn allows accurate estimation of the LSP momentum.

Accurately estimating the momentum of the LSP is crucial in the low-mass/compressed area. In this region the signal and background processes are difficult to distinguish from one another because the LSP carries so little momentum that it's addition doesn't change the missing energy of the event when compared to background processes. Having the sparticle system recoil off an ISR jet means that the LSP carries more momentum and is easier to tell apart from the background.

The ISR decay tree has its own set of variables that are defined with it. These variables are used to define specialised ISR signal regions optimised for the low-mass/compressed region. The definitions of these regions and a description of how they're used is given in Chapter 9.

8.3.1 Recursive Jigsaw Variables

After the imposition of the decay tree and the application of the jigsaw rules to resolve any unknowns the relevant properties of all particles are defined in the relevant reference frames. These can now be used to construct specialised variables that are only accessible through the RJ method.

The simplest variables available from the Recursive Jigsaw method are simply the momenta of particles or systems of particles. These variables take the form given in Equation 8.1

$$\vec{p}_A^{\text{F}} \quad (8.1)$$

This quantity is the momentum vector of the object ‘‘A’’ evaluated in the reference frame ‘‘F’’. The object in question can be a single object, such as a lepton ($\vec{p}_\ell$), or it can be a collection of objects, such as all invisible objects ($\vec{p}_I$) or all the ISR system ($\vec{p}_{\text{ISR}}$). The frame can be any one of the frames defined in the decay trees given in Figures 8.2 and 8.3, such as the laboratory frame ($\vec{p}^{\text{Lab}}$).

The primary energy-scale variables produced by the RJ method are denoted using the letter H and are defined in Equation 8.2.

$$H_{n,m}^{\text{F}} = \sum_{i=1}^n |\vec{p}_{V,i}^{\text{F}}| + \sum_{j=1}^m |\vec{p}_{I,j}^{\text{F}}| \quad (8.2)$$

Once again the superscript “**F**” refers to the frame that H is evaluated in. The indices n and m represent the number of visible and invisible momentum objects considered, respectively. If an event has $k < n$ visible objects, the sum only runs up to k . If an event has more than n objects, the sum first runs over all of the leptons (of which there are n_ℓ) before considering only the highest-energy jets (of which there are $n - n_\ell$).

There exists a “transverse” version of the variable given in Equation 8.2. This version of the variable takes only the transverse component of the momentum vectors (that is, the projection of the momentum vector onto the transverse plane) of visible and invisible objects. The “transverse plane” in the frame **F** is defined as follows: the boost that transforms the laboratory frame into the frame **F** is decomposed into a longitudinal (ie, parallel to the beam axis) boost followed by a transverse boost. The transverse plane is then defined as the plane perpendicular to the longitudinal boost. In practice this is the plane orthogonal to the beam-axis. Equation 8.3 defines this variable.

$$H_{T,n,m}^{\mathbf{F}} = \sum_{i=1}^n |\vec{p}_{T,V,i}^{\mathbf{F}}| + \sum_{j=1}^m |\vec{p}_{T,I,j}^{\mathbf{F}}| \quad (8.3)$$

The variables used in the definitions of the signal, control and validation regions are given below. The index n present in some of the variables takes different values depending on which physics scenario is being considered. For the scenarios where the $W^\pm$ decays leptonically, $n = 3$ and for the scenarios where the $W^\pm$ hadronically, $n = 4$.

- $H_{n,1}^{\mathbf{PP}}$: the scale variable described above.
- $\frac{H_{1,1}^{\mathbf{PP}}}{H_{4,1}^{\mathbf{PP}}}$: provides a good way of testing how “balanced” an event is. An event is said to be “unbalanced” if it contains particle that is much higher energy than the other particles in the event, causing this variable to take a value closer to unity. Signal events are typically balanced, but Z+Jets background events often have a single jet that dominates the event and so are typically unbalanced. This variable is useful for discriminating between the two.
- $\frac{p_{T,PP}^{\mathbf{Lab}}}{p_{T,PP}^{\mathbf{Lab}} + H_{n,1}^{\mathbf{PP}}}$: $p_{T,PP}^{\mathbf{Lab}}$ is the magnitude of the vector sum of the transverse momenta of all of the objects associated with the PP system evaluated in the **Lab** frame. This variable is used to test for significant boosts in the transverse direction: the distribution of this variable for signal events peaks sharply near zero while the background has a more spread out distribution.
- $\frac{H_{T,n,1}^{\mathbf{PP}}}{H_{n,1}^{\mathbf{PP}}}$: measures the fraction of momentum that lies in the transverse plane.

- $\frac{\min(H_{1,1}^{\mathbf{P}_a}, H_{1,1}^{\mathbf{P}_b})}{\min(H_{2,1}^{\mathbf{P}_a}, H_{2,1}^{\mathbf{P}_b})}$: compares the scale of one visible object and E_T^{miss} ($H_{1,1}^{\mathbf{P}_a}$ and $H_{1,1}^{\mathbf{P}_b}$ in their respective production frames) as opposed to two visible objects ($H_{2,1}^{\mathbf{P}_a}$ and $H_{2,1}^{\mathbf{P}_b}$). The numerator and denominator use the minima of the two values, respectively. In the case of the 3-lepton scenarios, this minimum always corresponds to the hemisphere that contains the Z^0 boson, since it's the only one with two visible objects. In this case, the variable reduces to $\frac{H_{1,1}^{\mathbf{P}_b}}{H_{2,1}^{\mathbf{P}_b}}$. This variable tests whether a single object is carrying a large percentage of the total momentum, and like $\frac{H_{1,1}^{\mathbf{P}\mathbf{P}}}{H_{4,1}^{\mathbf{P}\mathbf{P}}}$ is good at discriminating against Z+Jets events where this is likely to be the case
- $\Delta\phi_V^{\mathbf{P}}$: the azimuthal angle between the direction of the visible system V in frame $\mathbf{P}$ and the direction of the boost from frame $\mathbf{P}\mathbf{P}$ to frame $\mathbf{P}$. The topology of Standard Model processes is different from the one imposed by the decay tree and their distributions tend to peak near zero and π , whereas the signal is more broadly spread out.

In addition to these, the ISR decay tree produces its own set of ISR variables (also called “compressed” variables). As E_T^{miss} is used when deciding whether a jet belongs to the ISR system, transverse view of the reconstructed event is used where the longitudinal momenta of the leptons and jets are ignored [145].

- $p_{T,\text{CM}}$: the magnitude of the vector-sum of all of the transverse momenta of all objects associated with the centre-of-mass system, evaluated in the $\mathbf{Lab}$ frame.
- $p_{T,\text{ISR}}^{\mathbf{CM}}$: the magnitude of the vector-summed transverse momenta of all of the jets assigned to the ISR system, evaluated in the centre-of-mass frame.
- $p_{T,\text{I}}^{\mathbf{CM}}$: the magnitude of the vector-summed transverse momenta of all of objects assigned to the invisible system, evaluated in the centre-of-mass frame.
- $\Delta\phi_{\text{ISR,I}}^{\mathbf{CM}}$: the azimuthal opening angle between the ISR system and the invisible system evaluated in the centre-of-mass frame.
- $R_{\text{ISR}} \equiv \frac{\vec{p}_{\text{I}}^{\mathbf{CM}} \cdot \vec{p}_{T,S}^{\mathbf{CM}}}{|\vec{p}_{T,S}^{\mathbf{CM}}|}$: serves as an estimate of $\frac{m_{\tilde{\chi}_1^0}}{m_{\tilde{\chi}_1^\pm/\tilde{\chi}_2^0}}$. This corresponds to the fraction of the momentum in the S system ($\vec{p}_{T,S}^{\mathbf{CM}}$ being the vector sum of the transverse momenta of the jets and missing energy associated with the S system evaluated in the centre-of-mass frame) that is carried by its invisible system I (with momentum $\vec{p}_{\text{I}}^{\mathbf{CM}}$ in the centre-of-mass frame). As $\vec{p}_{T,S}^{\mathbf{CM}}$ increases, the backgrounds tend to smaller values of R_{ISR} while the signal tends to larger values [145].
- $n_{\text{jets}}^{\text{S}}$: the number of jets assigned to the “signal” or “sparticle” (denoted “S”) system.

- $n_{\text{jets}}^{\text{ISR}}$: the number of jets assigned to the ISR system.
- m_Z : the invariant mass of the dilepton pair assigned to the signal system. In the 3-lepton scenario, the pair of leptons with invariant mass closest to the mass of the Z^0 boson are used to calculate this value. This variable is the ISR equivalent of “ $m_{\ell\ell}$ ” defined in Section 8.1. The two variables aren’t identical since the ISR tree demands that m_Z be defined differently from $m_{\ell\ell}$ (although in the limit of $\eta \rightarrow 0$, $m_{\ell\ell} \rightarrow m_Z$), but they are used for the same purpose; finding events where the lepton pair likely comes from a leptonically decaying Z^0 -boson.
- m_J : the invariant mass of the jet system assigned to the signal system. This variable is the ISR equivalent of m_{jj} , fulfils the same purpose and has the same relationship with m_{jj} that m_Z has with $m_{\ell\ell}$.

Chapter 9

Signal Region Optimisation

The analysis described in this thesis has 8 signal regions in total, split into four different types:

1. High-Mass: Regions that are optimised to the $[m_{\tilde{\chi}_1^\pm/\tilde{\chi}_2^0}, m_{\tilde{\chi}_1^0}] = [600 \text{ GeV}, 0 \text{ GeV}]$ point on the mass-plane. These regions are optimised to search for charginos/neutralinos which are very heavy and produce very boosted final states.
2. Intermediate-Mass: Regions that are optimised to the $[500 \text{ GeV}, 200 \text{ GeV}]$ point. These regions are optimised to search for heavier charginos/neutralinos, but also heavier LSPs. This means a boosted final state (though, not as boosted as the high-mass region) with wider angles between certain decay products.
3. Low-Mass: Regions that are optimised to the $[200 \text{ GeV}, 100 \text{ GeV}]$ point. This area of the mass-plane is hard to study since all of the decay products are very low energy, and any signal region optimised for the $[200 \text{ GeV}, 100 \text{ GeV}]$ point are only able to produce high significances for a very small neighbourhood around the $[200 \text{ GeV}, 100 \text{ GeV}]$ point.
4. Compressed: Regions that are also optimised to the $[200 \text{ GeV}, 100 \text{ GeV}]$ point. These regions require that the sparticle systems recoils from ISR as described in Chapter 8.

There are two of each of these types of signal regions: one targeting the 2-lepton, 2-jet and E_T^{miss} final state and the other targeting the 3-lepton and E_T^{miss} final state. The high-, intermediate- and low-Mass signal regions study the physics scenarios given in Figures 6.1a and 6.1b using the decay trees given in Figure 8.2. The compressed




Fig. 9.1 Schematic diagram of the area the various signal regions target in the two-dimensional parameter space: $[m_{\tilde{\chi}_1^\pm/\tilde{\chi}_2^0}, m_{\tilde{\chi}_1^0}]$.

regions study the ISR scenarios given in Figures 6.1c and 6.1d using the ISR decay tree given in Figure 8.3.

Each of these four types of signal region study a different part of the mass-plane; the neighbourhood around the point it's optimised for. This is shown schematically in Figure 9.1. In it, the high- and intermediate-mass signal regions are fairly separated, coming from the fact that they were optimised on different points on the plane. Conversely, the low-mass and compressed regions overlap with each other because they were optimised on the same point. This arrangement was chosen because on their own the low-mass and compressed regions aren't able to achieve the necessary statistical significance, but when the two are statistically combined they can. In order to perform this statistical combination, the low-mass and compressed regions must be statistically independent in the way that signal and control regions must be, as described in Chapter 6.

9.1 Preselection Regions

Before covering the construction of signal regions in Sections 9.2 and 9.3, it is important to make sure that the Monte Carlo simulations of the various backgrounds model the data correctly. This is done rigorously when constructing control regions, a topic discussed in detail in Chapter 10. However, the process of constructing control regions is time-consuming and analysis teams often require a way to quickly determine whether the simulation is approximately correct before going forward and constructing signal regions. This purpose is fulfilled by the “preselection regions”¹. These regions aren’t optimised on any particular signal point, and contain a minimal amount of cuts which are typically done on the simplest variables (such as the transverse momentum of various particles or the constructed invariant masses). Each of these regions is instead constructed to accept only events that have the same final states as one of the new physics processes being studied. Further, there are only four preselection regions, one for each final state: 2-lepton, 2-lepton+ISR, 3-lepton and 3-lepton+ISR.

Subsection 9.1.1 will document the 2-lepton preselection regions and Subsection 9.1.2 will document the 3-lepton preselection regions.

9.1.1 2-Lepton Preselection Regions

The preselection regions constructed to check the modelling for the 2-lepton signal regions make use of a pair of properties unique to the two 2-lepton signal processes:

- It contains a pair of leptons that come from a Z^0 decay
- It contains a pair of quarks that come from a $W^\pm$ decay

Each of these properties provides several requirements for the preselection regions.

The $Z^0 \rightarrow \ell^+ \ell^-$ decay can make use of the properties listed in Section 6.5. The two leptons must have oppositely signed electric charges and the same flavour. Both of the leptons must be above a floor of energy (given as 25 GeV). Finally, both leptons must have an invariant mass between 80 and 100 GeV. An event that passes these requirements can reasonably be said to contain a Z^0 that decays to a pair of leptons.

The $W^\pm \rightarrow qq$ decay has a lot of the same properties, but with jets instead of leptons. Since jets can’t have a coherent flavour or charge, this requirement is ignored.

¹The term “preselection region” isn’t even technical jargon. It is a term of my own invention to describe the regions defined in this section. The regions described in this section aren’t typically published in papers, although they are usually documented in the various supporting documentation for particle physics analyses.

Instead, the event is required to have either two or three jets. For two jet events, the jets have their invariant mass calculated and if it falls between 60 and 100 GeV then it is likely that the jets came from the $W^\pm$ decaying into two quarks². The three jet events are used to study the 2-lepton, ISR process given in Figure 6.1c. In this case, there is a combinatoric ambiguity in how to assign which jets have come from the $W^\pm$ decay. The method for resolving this is a variation of the extremisation method mentioned in Section 8.2: the invariant mass of every possible pair of jets is calculated, and the pair that produces the invariant mass closest to the mass of the $W^\pm$ are assigned as having come from the $W^\pm$ decay.

Table 9.1 2-lepton preselection region definitions. The f subscripts in the “Lepton Selection” row simply refer to the leptons having the same flavour, whatever that flavour may be.

Preselection Regions: 2-Lepton		
Cut	Standard Tree Selection	Compressed Tree Selection
Lepton Selection	$\ell_f^+ \ell_f^-$	$\ell_f^+ \ell_f^-$
n_{jets}	≥ 2	≥ 3
$p_T^{\ell_1}$ [GeV]	> 25	> 25
$p_T^{\ell_2}$ [GeV]	> 25	> 25
$p_T^{j_1}$ [GeV]	> 30	> 30
$p_T^{j_2}$ [GeV]	> 30	> 30
$m_{\ell\ell}$ [GeV]	$\in (80, 100)$	$\in (80, 100)$
m_{jj} [GeV]	$\in (60, 100)$	$\in (60, 100)$

The distributions of variables used in the 2-lepton signal regions are shown in Figures 9.2 and 9.3. It is clear from the plots that the data is modelled very well by the Monte Carlo³.

²The wider window here is because the energy of jets is harder to reconstruct correctly than the energy of leptons, so there is some room for error in the selection.

³In addition to Monte Carlo simulation, these plots make use of other methods to model certain backgrounds: a photon template is used to model the “Z+Jets” background while the background labelled “Reducible” is a combination of backgrounds that is modelled using the matrix method. These methods, and backgrounds more generally, are discussed in detail in Chapter 10.













Fig. 9.2 Distributions of the 2-lepton standard tree variables in the standard 2-lepton preselection region.

(a) $p_T^{\ell_1}$ (b) $p_T^{\ell_2}$ (c) $p_T^{j_1}$ (d) $p_T^{j_2}$ (e) $p_T^{j_3}$ (f) n_{jets} (g) n_{jets}^S (h) n_{jets}^{ISR} (i) m_Z




Fig. 9.3 Distributions of the 2-lepton compressed tree variables in the compressed 2-lepton preselection region.

9.1.2 3-Lepton Preselection Regions

The two 3-lepton processes have their own distinct properties that can be exploited in the construction of preselection regions.

- It contains a pair of leptons that come from a Z^0 decay.
- It contains a third lepton that comes from the $W^\pm$ decaying leptonically.

Similar to the 2-lepton case, each of these properties provide several requirements for the preselection regions.

The most obvious of these is that any event in the preselection regions should contain exactly three leptons. Each of these leptons must achieve a minimum transverse momentum value to maintain reconstructed lepton quality (this requirement should be familiar from the previous section) with the second leptons requirement being softer than the first, and the third leptons requirement being softer than the second.

Two of the leptons come from the Z^0 decay, so the analysis must choose which of the three leptons in the event came from the Z^0 to be able to construct the $m_{\ell\ell}$ variable. This is yet another combinatoric ambiguity and is handled in much the same way as the jet ambiguity described in Subsection 9.1.1. The invariant mass of all possible lepton pairs are considered and the pair with invariant mass closest to the mass of the Z^0 is assigned as coming from the Z^0 decay. These two leptons are used to calculate the $m_{\ell\ell}$, which has a similar requirement to the one in the previous section placed upon it. The remaining lepton is used to calculate a new variable called m_T^W . This variable uses the remaining lepton to estimate the transverse mass of the $W^\pm$ and fulfils a similar role to m_{jj} in the previous section. There is a caveat in this case since the $W^\pm$ decays into both a lepton and neutrino, so an indeterminable amount of energy is lost when the neutrino passes through the detector. In turn, this means that the energy of the lepton that comes from this decay doesn't correlate very strongly with the energy of the $W^\pm$. Indeed, the energy of the lepton can only be less than the energy of the $W^\pm$, so the only cut on m_T^W is enforcing a lower limit on its possible value.

The requirements on jets in these regions are far more limited in scope than in the 2-lepton regions. The only requirement is that the 3-lepton ISR preselection region contain at least *one* jet. This is so that there is at least one jet that can function as a candidate for the ISR jet. The standard 3-lepton preselection region doesn't have any restriction on the number of jets whatsoever since none of the variables used by the

3-lepton standard signal/control/validation regions make use of any variable that is dependent on jets.

Table 9.2 gives the definitions of the two preselection regions.

Table 9.2 3-lepton preselection regions definitions. The flavours f_1 and f_2 are the flavours of the leptons that come from the Z^0 and $W^\pm$ decays, respectively.

Preselection Regions: 3-Lepton		
Cut	Standard Tree Selection	Compressed Tree Selection
Lepton Selection	$\ell_{f_1}^+ \ell_{f_1}^- (Z^0 \text{ pair}) \ell_{f_2}^\pm$	$\ell_{f_1}^+ \ell_{f_1}^- (Z^0 \text{ pair}) \ell_{f_2}^\pm$
n_{jets}	≥ 0	≥ 1
$p_T^{\ell_1}$ [GeV]	> 25	> 25
$p_T^{\ell_2}$ [GeV]	> 25	> 25
$p_T^{\ell_3}$ [GeV]	20	> 20
$m_{\ell\ell}$ [GeV]	$\in (75, 105)$	$\in (75, 105)$
m_T^W [GeV]	> 50	> 50

The distributions of the variable used in the 3-lepton signal regions are shown in Figures 9.4 and 9.5. Again, the data is modelled by the Monte Carlo very well.





Fig. 9.4 Distributions of the 3-lepton standard tree variables in the standard 3-lepton preselection region.






Fig. 9.5 Distributions of the 3-lepton compressed tree variables in the compressed 3-lepton preselection region.

9.2 Signal Regions: 2-Lepton, 2-Jet and E_T^{miss} Final States

The division of signal regions into 2-lepton and 3-lepton regions has already been discussed. Signal regions can further be divided by which process they are constructed to study. This section will define the 2-lepton signal regions whereas Section 9.3 will define the 3-lepton signal regions.

Three signal regions (the high-, intermediate- and low-mass regions) are constructed to study the standard process (Figure 6.1a) while a fourth (the compressed region) is constructed to study the corresponding ISR process (Figure 6.1c). The former three regions will be referred to as “standard tree regions” while the latter one will be referred to as the “compressed tree region”.

9.2.1 Standard Tree Regions

The three 2-lepton standard tree regions are defined in Table 9.3. The left-most column lists the variables that are cut on, while each column gives the requirement on that variable for each signal region⁴. For example, for an event to be in the “SR2 ℓ _High” region, it must have an $m_{\ell\ell}$ value between 80 GeV and 100 GeV, and a $H_{4,1}^{\text{PP}}$ value greater than 400 GeV.

The first several requirements listed in the table are mostly the same for all three regions. The requirements include a lepton selection⁵, a requirement on the number of jets and requirements on the transverse momentum of the leptons and jets. All these requirements are the same for all three regions, except for the requirement on the number of jets: SR2 ℓ _Intermediate and SR2 ℓ _High require *at least* two jets while SR2 ℓ _Low requires *exactly* two jets. The requirement that SR2 ℓ _Low have *exactly* two jets instead of *at least* two jets is imposed to stop it from overlapping with the 2-lepton compressed tree region defined in Subsection 9.2.2, as the low-mass and compressed regions need to be statistically independent so they can be statistically combined.

⁴NB: Not every region will have a cut placed on every variable listed in the left-hand column. For example, “SR2 ℓ _Low” doesn’t have a requirement on the variable “ $\Delta\phi_V^{\text{P}}$ ”. In these cases, the table will simply have a “–” in the column.

⁵A “lepton selection” is a requirement on the charge, flavour and other properties that leptons in an event must fulfil. In this case, the selection “ $\ell_f^+ \ell_f^-$ ” means that two leptons are selected, they have opposite charges, but the same flavour. A variety of selections are used in the regions described in this thesis.

These requirements, taken together, are colloquially known as “preselection requirements”⁶ and form the basis for all signal regions. These cuts aren’t used to isolate signal events and remove background events so much as they are used to ensure the events are of suitable quality. For instance, the lower limit on the transverse momenta of the leptons and jets is primarily so that particles that are too soft don’t enter the regions (this is familiar from Section 9.1), while the lepton requirement (opposite-sign, same-flavour) is to ensure that the lepton pair is likely to come from a Z^0 decay (as was discussed in Section 6.5).

Table 9.3 2-lepton signal region definitions - standard tree

Signal Regions: 2-Lepton Standard Tree			
Cut	SR2 ℓ _Low	SR2 ℓ _Intermediate	SR2 ℓ _High
Lepton Selection	$\ell_f^+ \ell_f^-$	$\ell_f^+ \ell_f^-$	$\ell_f^+ \ell_f^-$
n_{jets}	$= 2$	≥ 2	≥ 2
$n_{b\text{-tagged jets}}$	$= 0$	$= 0$	$= 0$
$p_T^{\ell_1}$ [GeV]	> 25	> 25	> 25
$p_T^{\ell_2}$ [GeV]	> 25	> 25	> 25
$p_T^{j_1}$ [GeV]	> 30	> 30	> 30
$p_T^{j_2}$ [GeV]	> 30	> 30	> 30
$m_{\ell\ell}$ [GeV]	$\in (80, 100)$	$\in (80, 100)$	$\in (80, 100)$
m_{jj} [GeV]	$\in (70, 90)$	$\in (60, 100)$	$\in (60, 100)$
$\min(\Delta\phi_{j_i, E_T^{\text{miss}}})$	> 2.4	—	—
$H_{4,1}^{\text{PP}}$ [GeV]	> 400	> 600	> 800
$\frac{p_{T,\text{PP}}^{\text{Lab}}}{p_{T,\text{PP}}^{\text{Lab}} + H_{4,1}^{\text{PP}}}$	< 0.05	< 0.05	< 0.05
$\frac{H_{1,1}^{\text{PP}}}{H_{4,1}^{\text{PP}}}$	$\in (0.35, 0.6)$	—	—
$\frac{H_{1,1}^{\text{P}_b}}{H_{2,1}^{\text{P}_b}}$	—	> 0.8	> 0.8
$\Delta\phi_V^{\text{P}}$	—	$\in (0.6, 2.6)$	$\in (0.3, 2.8)$


A final note on optimisation before moving onto the discussion of particular signal regions. The definitions of the signal regions listed in Table 9.3 (as well as tables 9.4, 9.6 and 9.7) are the final cuts that are determined by the optimisation process, and the plots shown will have those cuts drawn on them. This is done to show why those cuts work, and isn’t meant to imply that the cut is chosen based on how a plot looks.

⁶This is where the word “preselection” originally comes from, and where it derives its meaning as a set of selections that are used to confirm that the data is behaving correctly. The “preselection regions” defined in Section 9.1 are named for this concept.

SR2 ℓ _Low

The SR2 ℓ _Low region (or “low-mass region”) is optimised on the [200, 100] point on the mass-plane and has cuts that reflect this.


The first cut that’s applied after the preselection requirements is the cut on the invariant mass of the dilepton system, $m_{\ell\ell}$. This cut is the same across the three standard 2-lepton regions, and is used (along with the same-flavour, opposite-sign requirement) to ensure that the lepton pair is likely to have come from a Z^0 decay. The $m_{\ell\ell}$ distribution after the preselection requirements are applied but before the cut on $m_{\ell\ell}$ is shown in Figure 9.6a. A second distribution, the $m_{\ell\ell}$ distribution with all cuts applied *except* the cut on $m_{\ell\ell}$ is shown in Figure 9.6b. This second distribution is called an “N-1” Plot.



(a) The $m_{\ell\ell}$ distribution after the preselection requirements. (b) The “N-1” Plot of the $m_{\ell\ell}$ distribution.

Fig. 9.6 The SR2 ℓ _Low $m_{\ell\ell}$ distributions.

The next cut is on the invariant mass of the dijet system, m_{jj} . This cut is similar to the previous one, in that it is made to increase the likelihood that the jets come from the $W^\pm$ decay. Figure 9.7a shows the distribution of this variable after the application of the preselection requirements *and* the $m_{\ell\ell}$ cut. It can be clearly seen that the signal distribution has a peak at around $m_{jj} = 80$ GeV, and while the cut removes several dozen signal events, it removes more background events by a margin of five-to-six orders of magnitude.



(a) The m_{jj} distribution after the preselection requirements and $m_{\ell\ell}$ cut. (b) The “N-1” Plot of the m_{jj} distribution.


Fig. 9.7 The SR2ℓ_Low m_{jj} distributions.

The next cut is performed on the $H_{4,1}^{PP}$ variable. Figure 9.8a shows the $H_{4,1}^{PP}$ distribution with the preselection, $m_{\ell\ell}$ and m_{jj} requirements applied. In this distribution it can be seen that the signal and background have similar distributions, albeit with the signal being scaled down by a few orders of magnitude. The $H_{4,1}^{PP}$ variable is sensitive to the difference in mass between the $\tilde{\chi}_1^\pm/\tilde{\chi}_2^0$ and the $\tilde{\chi}_1^0$. The greater the mass difference, the greater the values that $H_{4,1}^{PP}$ will take. Since the difference is smaller in the SR2ℓ_Low region then the SR2ℓ_Intermediate or SR2ℓ_High regions, the distributions of the background and signal are in line and so this cut isn't as efficient as in the other two regions. In spite of this the cut is still done at 400 GeV, as it removes a much larger amount of background while keeping 2 dozen or so signal events.

The left-hand-side plots will continue to show the distributions of variables after the cuts on all of the previous variables are applied. These plots will take the name of “Progressive Plots” since they show the distributions of the variables as *progress* is made down the cut-flow.


The next cut is applied to the variable: $\frac{p_{T,PP}^{\text{Lab}}}{p_{T,PP}^{\text{Lab}} + H_{4,1}^{PP}}$. The signal events favour smaller values of this variable, so a very tight cut is placed on it, keeping only events where it takes values less than 0.05. This removes several signal events while at the same time removes around 10,000 background events. Figure 9.9 shows the progressive and “N-1” plots for this cut.

The penultimate cut is done on another ratio variable: $\frac{H_{1,1}^{PP}}{H_{4,1}^{PP}}$. This time, the signal distribution is fairly flat, with it forming a plateau between 0.0 and 0.6. The background



(a) The $H_{4,1}^{\text{PP}}$ distribution after the preselection, $m_{\ell\ell}$ and m_{jj} cuts. (b) The “N-1” Plot of the $H_{4,1}^{\text{PP}}$ distribution.

Fig. 9.8 The SR2 ℓ _Low $H_{4,1}^{\text{PP}}$ distributions.



(a) $\frac{p_{T,PP}^{\text{Lab}}}{p_{T,PP}^{\text{Lab}} + H_{4,1}^{\text{PP}}}$ Distribution: Progressive Plot. (b) $\frac{p_{T,PP}^{\text{Lab}}}{p_{T,PP}^{\text{Lab}} + H_{4,1}^{\text{PP}}}$ Distribution: “N-1” Plot.

Fig. 9.9 The SR2 ℓ _Low $\frac{p_{T,PP}^{\text{Lab}}}{p_{T,PP}^{\text{Lab}} + H_{4,1}^{\text{PP}}}$ distributions.




Fig. 9.10 The $\text{SR2}\ell_{\text{Low}}$ $\frac{H_{1,1}^{PP}}{H_{4,1}^{PP}}$ distributions.

on the other hand as a large excess of Z+Jets events between 0.0 and 0.35. This means that the optimal cut is to keep all events in the interval $[0.35, 0.6]$. Removing all events with the ratio between 0.0 and 0.35 will remove only a few signal events compared to 1000’s of background, while removing everything between 0.6 and 1.0 will only remove background. This is shown in Figure 9.10

The final cut is placed on a variable that is only used in this region: $\min(\Delta\phi_{j_i, E_T^{\text{miss}}})$. Here, the majority of the signal is near either 0 or π , with none in the centre. On the left hand side, the signal and background distributions match, so keeping the events there isn’t helpful. Hence, the cut is placed at 2.4, and all events above that are kept. It’s worth noting that this cut removed the majority of Z+Jets background events. Figure 9.11 shows that since it has at most a few dozen Z+Jets events compared to a few thousand in the previous plot.

It’s also worth noting that at the end of all of these cuts, there is still more background in the signal region than signal itself. This usually means that the signal region should be redesigned to try and achieve a higher signal-to-background ratio. This isn’t done here; the significance gained in this region is taken in statistical combination with the significance from the $\text{SR2}\ell_{\text{ISR}}$ region, which also doesn’t have the necessary significance on its own (see Section 9.2.2).

Finally, there is only one plot in the figure showing the $\min(\Delta\phi_{j_i, E_T^{\text{miss}}})$ distribution. This is because the “progressive” and “N-1” plots are the same plot, since the cut on $\min(\Delta\phi_{j_i, E_T^{\text{miss}}})$ is the last one to be done in the signal region.








Fig. 9.11 The SR2 ℓ _Low $\min(\Delta\phi_{ji}, E_T^{\text{miss}})$ Distribution


SR2 ℓ _Intermediate

The SR2 ℓ _Intermediate and SR2 ℓ _High regions (or the “intermediate-mass” and “high-mass” regions, respectively) are optimised on points in a very different part of the mass-plane than the SR2 ℓ _Low region. The increase in the masses of the $\tilde{\chi}_1^\pm/\tilde{\chi}_2^0$ allows for very different physics than the small masses of the SR2 ℓ _Low region. The SR2 ℓ _Intermediate region is optimised on the [500, 200] point, which means that the difference between the masses of the $\tilde{\chi}_1^\pm/\tilde{\chi}_2^0$ and the $\tilde{\chi}_1^0$ is 300 GeV. Hence, there is potentially a greater pool of energy that the Z^0 or the $W^\pm$ can draw from. This in turn leads to a number of new properties such as harder, more collimated leptons, better defined Recursive Jigsaw hemispheres and greater angular separation between them. The first cut is again on the invariant mass of the dilepton system. This cut is already familiar from the previous section so it doesn’t need to be discussed.

The second cut is again on the invariant mass of the dijet system. This cut is similar to the m_{jj} cut in the SR2 ℓ _Low region, but it covers a much wider interval (Figure 9.13).

(a) $m_{\ell\ell}$ Distribution: Progressive Plot(b) $m_{\ell\ell}$ Distribution: "N-1" PlotFig. 9.12 The SR2 ℓ _Intermediate $m_{\ell\ell}$ distributions.(a) m_{jj} Distribution: Progressive Plot(b) m_{jj} Distribution: "N-1" PlotFig. 9.13 The SR2 ℓ _Intermediate m_{jj} distributions.

The third cut is on $H_{4,1}^{\text{PP}}$, and Figure 9.14a shows the progressive plot for this variable. Here, it can be clearly seen that the backgrounds have a peak at around 200 GeV, whereas the signal has a peak at around 800 GeV. This happens because the more energetic vector bosons that are produced by the heavier $\tilde{\chi}_1^\pm/\tilde{\chi}_2^0$ particles in turn lead to harder leptons and jets. Since $H_{4,1}^{\text{PP}}$ is similar to a sum of the lepton transverse momenta, jet transverse momenta and E_T^{miss} , it takes larger values as a result. This means that the cut on $H_{4,1}^{\text{PP}}$ is placed further along the x-axis, at a higher value of 600 GeV.

Fig. 9.14 The SR2ℓ_Intermediate $H_{4,1}^{\text{PP}}$ distributions.

The cut on $\frac{p_{T,PP}^{\text{Lab}}}{p_{T,PP}^{\text{Lab}} + H_{4,1}^{\text{PP}}}$ is the same as in the SR2ℓ_Low region. Again, this is because the signal resides larger in the interval $[0.0, 0.5]$ while the background has many events beyond it (Figure 9.15). After this is a cut on another ratio variable, this one not used in the SR2ℓ_Low region: $\frac{\min(H_{1,1}^{\text{Pa}}, H_{1,1}^{\text{Pb}})}{\min(H_{2,1}^{\text{Pa}}, H_{2,1}^{\text{Pb}})}$. The numerator and denominator are both minimising function. The numerator takes the minimum value of $H_{1,1}^{\text{Pa}}$ and $H_{1,1}^{\text{Pb}}$ (which are the $H_{1,1}$ values calculated for each hemisphere that the PP -Frame is divided into). Similarly, the denominator takes the minimum value of the $H_{2,1}^{\text{Pa}}$ and $H_{2,1}^{\text{Pb}}$ (which are the $H_{2,1}$ values calculated for each hemisphere). A ratio of these minima is then taken. Also, since $H_{1,1}^{\text{Pa,b}} < H_{2,1}^{\text{Pa,b}}$ for all possible combinations of hemispheres, then the ratio is always bounded below one. Further, $\frac{\min(H_{1,1}^{\text{Pa}}, H_{1,1}^{\text{Pb}})}{\min(H_{2,1}^{\text{Pa}}, H_{2,1}^{\text{Pb}})} = 1$ when $\min(H_{1,1}^{\text{Pa}}, H_{1,1}^{\text{Pb}}) = \min(H_{2,1}^{\text{Pa}}, H_{2,1}^{\text{Pb}})$. The signal has a slight preference for higher values of this variable, while large sections of the background (particularly the Z+Jets and $V\gamma$ backgrounds) have low values of this variable. As such, the cut is placed at 0.8, and everything above it is kept (Figure 9.16).

The final cut is on an angular variable: $\Delta\phi_{VP}$. The signal peaks at the centre of the distribution while the background peaks at the periphery, thus the cut is to keep events with $\Delta\phi_V^{\text{P}}$ (Figure 9.17).




Fig. 9.15 The $\text{SR}2\ell_{\text{Intermediate}}$ $\frac{p_{T,PP}^{\text{Lab}}}{p_{T,PP}^{\text{Lab}} + H_{4,1}^{\text{PP}}}$ distributions.









Fig. 9.16 The $\text{SR}2\ell_{\text{Intermediate}}$ $\frac{\min(H_{1,1}^{P_a}, H_{1,1}^{P_b})}{\min(H_{2,1}^{P_a}, H_{2,1}^{P_b})}$ distributions.



Fig. 9.17 The SR2 ℓ _Intermediate $\Delta\phi_V^P$ distribution.

SR2 ℓ _High

The SR2 ℓ _High region is optimised on the [600, 0] point. This point has an even greater mass difference than the SR2 ℓ _Intermediate region (300 GeV for that region compared to 600 GeV for this one). This means that the effects caused by the mass difference is even greater for this region than for SR2 ℓ _Intermediate. The preselection requirements and the $m_{\ell\ell}$ and m_{jj} cuts are the same for this region as they are in the SR2 ℓ _Intermediate region. Figures 9.18 and 9.19 show the progressive and “N-1” distributions for $m_{\ell\ell}$ and m_{jj} , respectively.

The $H_{4,1}^{\text{PP}}$ distribution is even more extreme than it was in SR2 ℓ _Intermediate. Now the distribution has moved even further right, with the peak around 1100 GeV. The cut on $H_{4,1}^{\text{PP}}$ has been increased to 800 GeV to exploit this (Figure 9.20). The cut on $\frac{p_{T,\text{PP}}^{\text{Lab}}}{p_{T,\text{PP}}^{\text{Lab}} + H_{4,1}^{\text{PP}}}$ is the same as in the last two regions, for the same reasons as well. This can be clearly seen in Figure 9.21. The cut after this is on $\frac{\min(H_{1,1}^{\text{Pa}}, H_{1,1}^{\text{Pb}})}{\min(H_{2,1}^{\text{Pa}}, H_{2,1}^{\text{Pb}})}$, and is also the same as in SR2 ℓ _Intermediate. The distributions of the signal and background are similar to what they are in SR2 ℓ _Intermediate, so the cut remains the same (Figure 9.22). The final cut is on $\Delta\phi_V^P$. This time it’s slightly looser since more of the background is present at the extreme left/right of the plot (Figure 9.23).

(a) $m_{\ell\ell}$ Distribution: Progressive Plot(b) $m_{\ell\ell}$ Distribution: "N-1" PlotFig. 9.18 The SR2 ℓ _High $m_{\ell\ell}$ distributions.(a) m_{jj} Distribution: Progressive Plot(b) m_{jj} Distribution: "N-1" PlotFig. 9.19 The SR2 ℓ _High m_{jj} distributions.

Fig. 9.20 The SR2ℓ_High $H_{4,1}^{\text{PP}}$ distributions.Fig. 9.21 The SR2ℓ_High $\frac{p_{T,PP}^{\text{Lab}}}{p_{T,PP}^{\text{Lab}} + H_{4,1}^{\text{PP}}}$ distributions.




Fig. 9.22 The SR2 ℓ _High $\frac{\min(H_{1,1}^{P_a}, H_{1,1}^{P_b})}{\min(H_{2,1}^{P_a}, H_{2,1}^{P_b})}$ distributions.




Fig. 9.23 The SR2 ℓ _High $\Delta\phi_V^P$ distribution.

9.2.2 Compressed Tree Regions





The compressed region makes use of the ISR decay tree (Figure 8.3) and the ISR variables that are derived from it. Like SR2 ℓ _Low, the SR2 ℓ _ISR region is optimised on the [200, 100] point. Table 9.4 gives the cuts that define the region.

Table 9.4 2-lepton signal region definitions - compressed tree. Note that the cuts on $n_{\text{jets}}^{\text{S}}$ and $n_{\text{jets}}^{\text{ISR}}$ occur at the end of the cutflow, and are not part of the preselection requirements despite technically being cuts on an amount of jets.


Signal Regions: 2-Lepton Compressed Tree	
Cut	SR2 ℓ _ISR
Lepton Selection	$\ell_f^+ \ell_f^-$
n_{jets}	$\in [3, 4]$
$n_{b\text{-tagged jets}}$	$= 0$
$p_T^{\ell_1}$ [GeV]	> 25
$p_T^{\ell_2}$ [GeV]	> 25
$p_T^{j_1}$ [GeV]	> 30
$p_T^{j_2}$ [GeV]	> 30
m_Z [GeV]	$\in (80, 100)$
m_J [GeV]	$\in (50, 110)$
$\Delta\phi_{\text{ISR},\text{I}}^{\text{CM}}$	> 2.8
R_{ISR}	$\in (0.4, 0.75)$
$p_{T,\text{ISR}}^{\text{CM}}$ [GeV]	> 180
$p_{T,\text{I}}^{\text{CM}}$ [GeV]	> 100
$p_{T,\text{CM}}$ [GeV]	< 20
$n_{\text{jets}}^{\text{S}}$	$= 2$
$n_{\text{jets}}^{\text{ISR}}$	≥ 1

The compressed region doesn't cut on the familiar $m_{\ell\ell}$ and m_{jj} variables. Rather, there are a pair of specially constructed variables called m_Z and m_J . The m_Z variable replaces $m_{\ell\ell}$, and serves as a measure of the mass of the Z^0 boson. The variable m_J replaces m_{jj} , and is constructed differently due to the different jet requirements of this region. Since SR2 ℓ _ISR is required to have either 3 or 4 jets, that means that the leading and subleading jets don't necessarily both come from the $W^\pm$. Again, a choice must be made in assigning which jets come from the $W^\pm$. The method used is identical to the one used in Subsection 9.1.1: the invariant masses of all possible jet pairs is calculated and the pair whose invariant mass is closest to the mass of the $W^\pm$ is taken as coming from the $W^\pm$ decay.

Figures 9.24 and 9.25 show the distributions of the m_Z and m_J variables, respectively.


(a) m_Z Distribution: Progressive Plot(b) m_Z Distribution: "N-1" PlotFig. 9.24 The SR2 ℓ _Compressed M_Z distributions.(a) m_J Distribution: Progressive Plot(b) m_J Distribution: "N-1" PlotFig. 9.25 The SR2 ℓ _Compressed m_J distributions.

The next cut is the first one done on an ISR variable: $\Delta\phi_{\text{ISR},I}^{\text{CM}}$. This azimuthal opening angle between the ISR system and the invisible system, which is larger for the signal than the background (Figure 9.26). Next is a cut on R_{ISR} , which serves as an estimate for the fraction of $\tilde{\chi}_1^\pm/\tilde{\chi}_2^0$ that the LSP mass is. The signal is biased toward lower values of this variable and the cut reflects this (Figure 9.28).



(a) $\Delta\phi_{\text{ISR},I}^{\text{CM}}$ Distribution: Progressive Plot (b) $\Delta\phi_{\text{ISR},I}^{\text{CM}}$ Distribution: “N-1” Plot

Fig. 9.26 The SR2 ℓ _Compressed $\Delta\phi_{\text{ISR},I}^{\text{CM}}$ distributions.



(a) R_{ISR} Distribution: Progressive Plot (b) R_{ISR} Distribution: “N-1” Plot

Fig. 9.27 The SR2 ℓ _Compressed R_{ISR} distributions.

The variable $p_{T,ISR}^{CM}$ is cut on after this. The signal distribution doesn't really favour any particular $p_{T,ISR}^{CM}$ value, while there are many background events with low $p_{T,ISR}^{CM}$. As such, the requirement to keep events with $p_{T,ISR}^{CM} > 180$ GeV removes only a few signal events while removing several thousand background events (Figure 9.28). After $p_{T,ISR}^{CM}$ is the cut on $p_{T,I}^{CM}$, the transverse momentum of the invisible system measured in the centre-of-mass frame. The cut is set at 100 GeV, keeping events with $p_{T,I}^{CM}$ greater than this (Figure 9.29). The reason for this choice is that the invisible system will necessarily have *at least* 100 GeV of energy in it because the $\tilde{\chi}_1^0$ weighs 100 GeV. The final cut on ISR variables is on $p_{T,CM}$, the total transverse momentum of the system in the centre-of-mass frame. In an ideal world this variable should always be zero, but this isn't the case since the detector can mismeasure the various properties of the particles it finds. As such, a cut of $p_{T,CM} < 25$ GeV is used to remove all of the erroneous backgrounds (Figure 9.30).




Fig. 9.28 The SR2 ℓ _Compressed $p_{T,ISR}^{CM}$ distributions.

The last two cuts are on specialised jet counting variables. These are n_{jets}^S , the number of jets in the sparticle hemisphere, and n_{jets}^{ISR} , the number of jets in the ISR hemisphere. The number of sparticle jets is restricted to be equal to two. Since the sparticle system is just the familiar $\tilde{\chi}_1^\pm \tilde{\chi}_2^0$ decay chain, it should only have two jets; the ones that come from the $W^\pm \rightarrow qq$ decay. Hence the cut on n_{jets}^S is to restrict it to two. The number of jets in the ISR hemisphere could theoretically be any non-zero integer, although the energy of the event limits the total amount that can realistically appear. For this analysis, the SR2 ℓ _ISR region requires *at least* one ISR jet in order for the

region to make sense. As such, the cut is simply $n_{\text{jets}}^{\text{ISR}} \geq 1$. Figures 9.31 and 9.32 show the progressive and “N-1” plots for the $n_{\text{jets}}^{\text{S}}$ and $n_{\text{jets}}^{\text{ISR}}$ variables respectively.




Fig. 9.29 The SR2 ℓ _Compressed $p_{T,I}^{\text{CM}}$ distributions.




Fig. 9.30 The SR2 ℓ _Compressed p_T^{CM} distributions.




Fig. 9.31 The SR2 ℓ _Compressed n_{jets}^S distributions.




Fig. 9.32 The SR2 ℓ _Compressed $n_{\text{jets}}^{\text{ISR}}$ distributions.


9.2.3 Yields and Significance

Table 9.5 shows the yields of the major backgrounds in each region.


Table 9.5 Expected Standard Model yields in the 2-lepton signal regions. Only statistical uncertainties are considered.

	SR2 ℓ _High	SR2 ℓ _Intermediate	SR2 ℓ _Low	SR2 ℓ _Compressed
Higgs	0.0 $\pm$ 0.0	0.0 $\pm$ 0.0	0.011 $\pm$ 0.008	0.0 $\pm$ 0.0
Reducible (MM)	0.0 $\pm$ 0.0	0.0 $\pm$ 0.0	0.0 $\pm$ 0.0	0.0 $\pm$ 0.157
$V\gamma$	0.0 $\pm$ 0.0	0.0 $\pm$ 0.0	0.0 $\pm$ 0.0	0.01 $\pm$ 0.007
$Wt + t\bar{t} + t\bar{t}V$	0.011 $\pm$ 0.006	0.04 $\pm$ 0.022	0.643 $\pm$ 0.291	0.331 $\pm$ 0.166
VVV	0.008 $\pm$ 0.004	0.009 $\pm$ 0.004	0.0 $\pm$ 0.0	0.0 $\pm$ 0.0
VV	1.93 $\pm$ 0.26	2.60 $\pm$ 0.30	1.61 $\pm$ 0.26	2.41 $\pm$ 0.21
Z +Jets (MC)	0.0 $\pm$ 0.0	0.0 $\pm$ 0.0	10.18 $\pm$ 15.92	0.93 $\pm$ 0.98
Total Bkg	1.95 $\pm$ 0.26	2.65 $\pm$ 0.30	12.44 $\pm$ 15.93	3.68 $\pm$ 1.03

Figures 9.33 and 9.34 show the sensitivity plots of all of the 2-lepton signal regions. Figure 9.33 shows the sensitivity plots for the SR2 ℓ _High and SR2 ℓ _Intermediate regions. Both of these regions have large exclusion contours, with the SR2 ℓ _High region excluding the [550, 0] point at its right-most limit and the SR2 ℓ _Intermediate region excluding the [500, 200] point at its most extreme value. Figure 9.34 gives the sensitivity plots for the SR2 ℓ _Low region, the SR2 ℓ _Compressed region and their statistical combination. Finally, these significance plots use the data-driven Z +Jets background estimate instead of the Monte Carlo estimate.



(a) SR2ℓ_High Exclusion Contour



(b) SR2ℓ_Intermediate Exclusion Contour

Fig. 9.33 Sensitivity Plots for the 2-lepton high- and intermediate-mass signal regions.



(a) SR2 ℓ _Low Exclusion Contour(b) SR2 ℓ _Compressed Exclusion Contour(c) Combined SR2 ℓ _Low + SR2 ℓ _Compressed Exclusion Contour

Fig. 9.34 Sensitivity Plots for the 2-lepton low-mass region, compressed region and the low-mass + compressed statistical combination.

9.3 Signal Regions: 3-Lepton and E_T^{miss} Final States

The 3-lepton regions are divided into two groups the same way the 2-lepton regions are: three regions studying the standard process (Figure 6.1b) and a fourth for the corresponding ISR process (Figure 6.1d).

9.3.1 Standard Tree Regions



The three 3-lepton standard tree regions are defined in Table 9.6. All 3-lepton regions have the same combinatoric ambiguity that the 3-lepton preselection region described in Subsection 9.1.2 does: the inclusion of a third lepton means that a pair of leptons must be assigned as coming from the Z^0 decay. This is resolved the same way as it is in Subsection 9.1.2 as well: the pair of leptons with invariant mass closest to m_{Z^0} is assigned as coming from the Z^0 decay. The lepton selection is then applied; the pair of leptons that come from the Z^0 decay required to be same-flavour, opposite-sign and are used to calculate $m_{\ell\ell}$ while the third lepton doesn't have any charge or flavour requirement and is used to calculate m_T^W . Apart from this the rest of the preselection requirements are fairly straight-forward: a requirement on the number of jets/ b -tagged jets and the transverse momenta of the leptons.

Table 9.6 3-lepton signal region definitions - standard tree

Signal Regions: 3-Lepton Standard Tree			
Cut	SR3 ℓ _Low	SR3 ℓ _Intermediate	SR3 ℓ _High
Lepton Selection	$\ell_{f_1}^+ \ell_{f_1}^- (Z^0 \text{ pair}) \ell_{f_2}^\pm$	$\ell_{f_1}^+ \ell_{f_1}^- (Z^0 \text{ pair}) \ell_{f_2}^\pm$	$\ell_{f_1}^+ \ell_{f_1}^- (Z^0 \text{ pair}) \ell_{f_2}^\pm$
n_{jets}	= 0	< 3	< 3
$n_{b\text{-taggedjets}}$	= 0	= 0	= 0
$p_T^{\ell_1}$ [GeV]	> 60	> 60	> 60
$p_T^{\ell_2}$ [GeV]	> 40	> 50	> 60
$p_T^{\ell_3}$ [GeV]	> 30	> 30	> 40
$m_{\ell\ell}$ [GeV]	$\in (75, 105)$	$\in (75, 105)$	$\in (75, 105)$
m_T^W [GeV]	> 100	> 130	> 150
$H_{3,1}^{\text{PP}}$ [GeV]	> 250	> 450	> 550
$\frac{p_{T,\text{PP}}^{\text{Lab}}}{p_{T,\text{PP}}^{\text{Lab}} + H_{3,1}^{\text{PP}}}$	< 0.05	< 0.15	< 0.2
$\frac{H_{T,3,1}^{\text{PP}}}{H_{3,1}^{\text{PP}}}$	> 0.9	> 0.8	> 0.75
$\frac{H_{1,1}^{\text{Pb}}}{H_{2,1}^{\text{Pb}}}$	—	> 0.75	> 0.8

SR3 ℓ _Low

The SR3 ℓ _Low region is optimised to the [200, 100] point, like the SR2 ℓ _Low region. The cut on $m_{\ell\ell}$ is similar to the $m_{\ell\ell}$ cuts in the 2-lepton standard regions, being [75, 105] rather than [80, 100]. The m_T^W cut is one-sided, unlike the cuts on m_{jj} in the 2-lepton standard regions. This is because some of the $W^\pm$ bosons energy is carried away by the neutrino that is also part of the decay. Figures 9.35 and 9.36 show the progressive and “N-1” plots for the $m_{\ell\ell}$ and m_{jj} variables, respectively.

Fig. 9.35 The SR3 ℓ _Low $m_{\ell\ell}$ distributions.Fig. 9.36 The SR3 ℓ _Low m_T^W distributions.

After the cuts on the Z^0 and $W^\pm$ mass proxy variables, the 2-lepton standard regions would have a cut on $H_{4,1}^{\text{PP}}$. The equivalent variable for the 3-lepton regions is $H_{3,1}^{\text{PP}}$ where the “3” denotes the three visible leptons and the “1” once again denotes the E_T^{miss} . The function of this variable for the 3-lepton standard regions is otherwise identical to the use of $H_{4,1}^{\text{PP}}$ for the 2-lepton regions; it compares the “scale” of the signal and background and a cut is made that removes more background than signal. The signal and background distributions, like in SR2 ℓ _Low, occupy the same space, with the signal being a few orders of magnitude smaller than the background. As such, the cut on $H_{3,1}^{\text{PP}}$ occurs at a low value of 250 GeV (Figure 9.37). Another variable that has an analogue in the 2-lepton regions, $\frac{p_{T,\text{PP}}^{\text{Lab}}}{p_{T,\text{PP}}^{\text{Lab}} + H_{3,1}^{\text{PP}}}$, is cut on next. The cut is same as the corresponding cut on $\frac{p_{T,\text{PP}}^{\text{Lab}}}{p_{T,\text{PP}}^{\text{Lab}} + H_{4,1}^{\text{PP}}}$ in the 2-lepton regions, as can be seen in Figure 9.38.






Fig. 9.37 The SR3 ℓ _Low $H_{3,1}^{\text{PP}}$ distributions.

The final cut is on a new ratio variable: $\frac{H_{T,3,1}^{\text{PP}}}{H_{3,1}^{\text{PP}}}$. Here the signal favours values nearer to unity, whereas the background distribution is more evenly distributed across all values of $\frac{H_{T,3,1}^{\text{PP}}}{H_{3,1}^{\text{PP}}}$, hence the cut keeping all events with $\frac{H_{T,3,1}^{\text{PP}}}{H_{3,1}^{\text{PP}}} > 0.9$. Figure 9.39 shows the relevant plot.

Fig. 9.38 The SR3 ℓ _Low $\frac{p_{T,PP}^{\text{Lab}}}{p_{T,PP}^{\text{Lab}} + H_{3,1}^{\text{PP}}}$ distributions.Fig. 9.39 The SR3 ℓ _Low $\frac{H_{T,3,1}^{\text{PP}}}{H_{3,1}^{\text{PP}}}$ distributions.

SR3 ℓ _Intermediate

The SR3 ℓ _Intermediate region is optimised on the $[500, 100]$ point, and uses the same variables that the SR3 ℓ _Low uses, with the addition of $\frac{H_{1,1}^{P_b}}{H_{2,1}^{P_b}}$. The first cut is, again, on $m_{\ell\ell}$, which is identical to the one in SR3 ℓ _Low (Figure 9.35). The second cut is on m_T^W , this time done on a higher value, 130 GeV (Figure 9.41).




Fig. 9.40 The SR3 ℓ _Intermediate $m_{\ell\ell}$ distributions.




Fig. 9.41 The SR3 ℓ _Intermediate m_T^W distributions.

The next cut is on $H_{3,1}^{\text{PP}}$, and in a similar trend to the cuts on $H_{4,1}^{\text{PP}}$ seen in the 2-lepton regions, the SR3 ℓ _Intermediate cut on $H_{3,1}^{\text{PP}}$ is harder than it was in SR3 ℓ _Low. This is for the same reason as it was for the 2-lepton regions, the increase in mass of the $\tilde{\chi}_1^\pm/\tilde{\chi}_2^0$ and the $\tilde{\chi}_1^0$ gives harder leptons that give greater values of $H_{3,1}^{\text{PP}}$ than background. As such, the $H_{3,1}^{\text{PP}}$ distribution moves to the right and allows to the cut to be placed at a higher value, in this case, 450 GeV. This can be seen in Figure 9.42.




Fig. 9.42 The SR3 ℓ _Intermediate $H_{3,1}^{\text{PP}}$ distributions.

The next cut is on $\frac{p_{T,\text{PP}}^{\text{Lab}}}{p_{T,\text{PP}}^{\text{Lab}} + H_{3,1}^{\text{PP}}}$. This time the signal is more spread out, so a looser cut is required, this time taking the value 0.15 (Figure 9.43). Another familiar cut, on $\frac{H_{T,3,1}^{\text{PP}}}{H_{3,1}^{\text{PP}}}$, is changed compared to the SR3 ℓ _Intermediate region. This time the signal peaks less violently at zero, so a looser cut is necessary (Figure 9.44). The final cut of this region on a new variable called $\frac{H_{1,1}^{\text{Pb}}}{H_{2,1}^{\text{Pb}}}$, a 3-lepton analogue of the 2-lepton $\frac{\min(H_{1,1}^{\text{Pa}}, H_{1,1}^{\text{Pb}})}{\min(H_{2,1}^{\text{Pa}}, H_{2,1}^{\text{Pb}})}$. Indeed, while these two variables are different for the 2-lepton regions, they are the same for the 3-lepton regions. The signal peaks near unity while the background is fairly evenly distributed so a cut of 0.75 is used, keeping events above that threshold (Figure 9.45).




Fig. 9.43 The SR3 ℓ _Intermediate $\frac{p_{T,PP}^{\text{Lab}}}{p_{T,PP}^{\text{Lab}} + H_{3,1}^{\text{PP}}}$ distributions.




Fig. 9.44 The SR3 ℓ _Intermediate $\frac{H_{T,3,1}^{\text{PP}}}{H_{3,1}^{\text{PP}}}$ distributions.








Fig. 9.45 The SR3 ℓ _Intermediate $\frac{H_{1,1}^{P_b}}{H_{2,1}^{P_b}}$ distribution.



SR3 ℓ _High

The SR3 ℓ _High region is optimised on the $[600, 0]$ point, and cuts on the same variables as the SR3 ℓ _Intermediate region. The cut on $m_{\ell\ell}$ is the same as the last two regions, and doesn't merit further discussion (Figure 9.46). The m_T^W cut is once again made harder, to 150 GeV now (Figure 9.47).

Once again, the heavier $\tilde{\chi}_1^\pm/\tilde{\chi}_2^0$ and $\tilde{\chi}_1^0$ particles give even harder leptons and thus a rightward shifted $H_{3,1}^{\text{PP}}$ distribution. Following this, the cut on $H_{3,1}^{\text{PP}}$ is made even harder, sitting at 550 GeV. It can clearly be seen in Figure 9.48a that this cut removes a large amount of background while removing one-tenth of one event of signal.

The next cut is on $\frac{p_{T,\text{PP}}^{\text{Lab}}}{p_{T,\text{PP}}^{\text{Lab}} + H_{3,1}^{\text{PP}}}$. In keeping with the trend, the signal is more spread out and the cut is made even looser than before, this time keeping any event with $\frac{p_{T,\text{PP}}^{\text{Lab}}}{p_{T,\text{PP}}^{\text{Lab}} + H_{3,1}^{\text{PP}}} < 0.2$ (Figure 9.49). The penultimate cut is on $\frac{H_{T,3,1}^{\text{PP}}}{H_{3,1}^{\text{PP}}}$, which is again loosened and the cut occurs at 0.75 (Figure 9.50). The final cut is on the new ratio: $\frac{H_{1,1}^{\text{Pb}}}{H_{2,1}^{\text{Pb}}}$. This time the variable peaks a little more extremely near unity, so a tighter cut is allowed.

(a) $m_{\ell\ell}$ Distribution: Progressive Plot(b) $m_{\ell\ell}$ Distribution: "N-1" PlotFig. 9.46 The SR3 ℓ _High $m_{\ell\ell}$ distributions.(a) m_T^W Distribution: Progressive Plot(b) m_T^W Distribution: "N-1" PlotFig. 9.47 The SR3 ℓ _High m_T^W distributions.

Fig. 9.48 The SR3ℓ_High $H_{3,1}^{\text{PP}}$ distributions.Fig. 9.49 The SR3ℓ_High $\frac{p_{T,PP}^{\text{Lab}}}{p_{T,PP}^{\text{Lab}} + H_{3,1}^{\text{PP}}}$ distributions.




Fig. 9.50 The SR3 ℓ _High $\frac{H_{T,3,1}^{PP}}{H_{3,1}^{PP}}$ distributions.




Fig. 9.51 The SR3 ℓ _High $\frac{H_{1,1}^{P_b}}{H_{2,1}^{P_b}}$ distribution.

9.3.2 Compressed Tree Regions





The final signal region to discuss is the 3-lepton compressed region, SR3 ℓ _ISR. Like SR2 ℓ _ISR, this region uses the compressed tree and the variables associated with it. Actually, it's even simpler than SR2 ℓ _ISR, using many of the same variables but not the specialised jet variables. Table 9.7 defines the region.

SR3 ℓ _ISR is optimised on the same point as SR3 ℓ _Low and has a similar pre-selection. First, the combinatoric ambiguity from the presence of a third lepton is resolved exactly how it was for the three 3-lepton regions described in Subsection 9.3.1. Following this, there is a requirement on jet multiplicity, needing either 1, 2 or 3 jets. The lower limit of 1 is so that there exists a jet that can be assigned as the ISR jet, while the upper limit of 3 is to allow events where one ISR particle produces 2 or 3 jets, but remove events with more than a few ISR particles. This is done because the purpose of adding the ISR is to have something that the rest of the system can recoil off of and become boosted. With too many ISR jets the ISR system doesn't have a strong directionality and so the sparticle system doesn't have a coherent direction to recoil in.



Table 9.7 3-lepton signal region definitions - compressed tree

Signal Regions: 3-Lepton Compressed Tree	
Cut	SR3 ℓ _ISR
Lepton Selection	$\ell_{f_1}^+ \ell_{f_1}^- (Z^0 \text{ pair}) \ell_{f_2}^\pm$
n_{jets}	$\in [1, 3]$
$n_{b\text{-tagged jets}}$	$= 0$
$p_T^{\ell_1}$ [GeV]	> 25
$p_T^{\ell_2}$ [GeV]	> 25
$p_T^{\ell_3}$ [GeV]	> 20
$m_{\ell\ell}$ [GeV]	$\in (75, 105)$
m_T^W [GeV]	> 100
$\Delta\phi_{\text{ISR,I}}^{\text{CM}}$	> 2.0
R_{ISR}	$\in (0.55, 1.0)$
$p_{T,\text{ISR}}^{\text{CM}}$ [GeV]	> 100
$p_{T,\text{I}}^{\text{CM}}$ [GeV]	> 80
$p_{T,\text{CM}}$ [GeV]	< 25

The first cuts are on $m_{\ell\ell}$ and m_T^W . These are the same as they are for SR3 ℓ _Low and are done for the same reasons. Figures 9.52 and 9.53 show the relevant plots.

(a) $m_{\ell\ell}$ Distribution: Progressive Plot(b) $m_{\ell\ell}$ Distribution: "N-1" PlotFig. 9.52 The SR3 ℓ _ISR $m_{\ell\ell}$ distributions.(a) m_T^W Distribution: Progressive Plot(b) m_T^W Distribution: "N-1" PlotFig. 9.53 The SR3 ℓ _ISR m_T^W distributions.

The next cut is on $\Delta\phi_{\text{ISR},I}^{\text{CM}}$. Unfortunately, the signal and background have vaguely the same shape in this variable, and to make things worse both have a roughly even distribution of events across all values of $\Delta\phi_{\text{ISR},I}^{\text{CM}}$, as can be seen in the progressive plot, Figure 9.54a. Thus, making a choice for a cut is difficult but the signal seems to increase as $\Delta\phi_{\text{ISR},I}^{\text{CM}}$ increases. Therefore, a cut which keeps events with $\Delta\phi_{\text{ISR},I}^{\text{CM}} > 2.0$ is the optimal choice.

(a) $\Delta\phi_{\text{ISR},I}^{\text{CM}}$ Distribution: Progressive Plot (b) $\Delta\phi_{\text{ISR},I}^{\text{CM}}$ Distribution: “N-1” PlotFig. 9.54 The SR3ℓ_ISR $\Delta\phi_{\text{ISR},I}^{\text{CM}}$ distributions.(a) R_{ISR} Distribution: Progressive Plot (b) R_{ISR} Distribution: “N-1” PlotFig. 9.55 The SR3ℓ_ISR R_{ISR} distributions.

After this, there is a cut on R_{ISR} . In this variable as well the signal and background seem to have the same shape, although there are some exploitable features such as the peak at around 0.8 or the low amounts of signal compared to background near values of zero. Figure 9.55 shows the distributions. The next cut is on $p_{T,\text{ISR}}^{\text{CM}}$. The cut here for SR3ℓ_ISR is slightly softer than the equivalent cut in the SR2ℓ_ISR region, landing at 100 GeV rather than 180 (Figure 9.56). The penultimate cut is on $p_{T,I}^{\text{CM}}$. Here, the cut is at 80 GeV, again slightly softer than its SR2ℓ_ISR counterpart. Figure 9.57






Fig. 9.56 The SR3ℓ_ISR $p_{T,ISR}^{CM}$ distributions.

shows the distribution. Finally, is a cut on $p_{T,CM}$, the total transverse momentum of the system measured in the centre-of-mass frame. This cut is made tighter in comparison to SR2ℓ_ISR, at only 20 GeV. Figure 9.58 shows this.

In contrast to SR2ℓ_ISR, this region doesn't have more background in it than signal by the end of the cutflow. That said, this region will be statistically combined with SR3ℓ_Low to produce a significant result, as will be discussed in Chapter 13.

Fig. 9.57 The SR3 ℓ _ISR $p_{T,I}^{\text{CM}}$ distributions.Fig. 9.58 The SR3 ℓ _ISR $p_{T,\text{CM}}$ distribution.

9.3.3 Yields and Significance

Table 9.8 shows the yields of the major backgrounds in each region.

Table 9.8 Expected Standard Model yields in the 3-lepton signal regions. Only statistical uncertainties are considered.

	SR3L_High	SR3L_Intermediate	SR3L_Low	SR3L-Compressed
Higgs	0.0 ± 0.0	0.0 ± 0.0	0.0 ± 0.0	0.0 ± 0.0
Reducible (MM)	0.0 ± 0.074	0.0 ± 0.001	0.0 ± 0.156	0.0 ± 0.192
$t\bar{t}V$	0.03 ± 0.01	0.037 ± 0.009	0.018 ± 0.009	0.06 ± 0.013
VVV	0.193 ± 0.068	0.318 ± 0.073	0.254 ± 0.047	0.081 ± 0.027
VV	0.758 ± 0.358	1.757 ± 0.35	9.182 ± 0.781	3.367 ± 0.366
Total Bkg	0.98 ± 0.37	2.11 ± 0.36	9.45 ± 0.80	3.51 ± 0.41

Figures 9.59 and 9.60 show the sensitivity plots for the 3-lepton regions. The sensitivity plots for the SR3 ℓ _High and SR2 ℓ _Intermediate regions are given in Figure 9.59. Once again, these regions have large exclusion contours, albeit smaller than their 2-lepton counterparts. The SR2 ℓ _High signal region is only barely able to exclude the $[500, 0]$ point while the SR2 ℓ _Intermediate region excludes the $[450, 150]$ point at its most extreme. The 3-lepton signal regions are much stronger in the SR3 ℓ _Low and SR3 ℓ _ISR regions (Figure 9.60). Both of them exclude a non-trivial section of the low-mass/ISR region of the mass-plane, with their statistical combination (Figure 9.60c) excluding a very large portion of that part of the mass-plane.



(a) SR3 ℓ _High Exclusion Contour(b) SR3 ℓ _Intermediate Exclusion Contour

Fig. 9.59 Sensitivity Plots for the 3-lepton high- and intermediate-mass signal regions.



(a) SR3 ℓ _Low Exclusion Contour(b) SR3 ℓ _Compressed Exclusion Contour(c) Combined SR3 ℓ _Low + SR3 ℓ _Compressed Exclusion Contour

Fig. 9.60 Sensitivity Plots for the 3-lepton low-mass region, compressed region and the low-mass + compressed statistical combination.

Chapter 10

Standard Model Background Estimation

A variety of processes contributed to the backgrounds in the signal regions defined in Chapter 9. This chapter will describe the background processes present, the methods used to model them and the control regions used to validate the modelling.

Chapter 7 listed the background processes and the generators that were used to generate them in Table 7.1. In the plots in Chapter 9 these backgrounds are grouped together based on shared characteristics and designated using a specific colour.

- **Z+Jets:** This is the $Z(\rightarrow \ell\ell) + \text{Jets}$ background listed in Table 7.1. This is when the collision produces a Z^0 boson (which then decays into a lepton pair) and some number of jets (that don't come from a vector boson decay).
- **Top:** The background labelled “top” is a composite background made up of the “Single-top”, $t\bar{t}$ and $t\bar{t} + V$ ($t\bar{t} + Z$, $t\bar{t} + W/WW$) backgrounds. As such, it is alternatively labelled “ $Wt + t\bar{t} + ttV$ ” on some of the plots in Chapters 9, 10 and 11.
- **VV:** This is the “diboson” background and comes from collisions that produce a pair of vector bosons: WW , WZ or ZZ .
- **VVV:** Similarly, this is the “triboson” background, which comes from collisions that produce three vector bosons: WWW , WWZ , WZZ and ZZZ .
- **$V\gamma$:** A leptonically decaying vector boson produced in association with a photon.
- **Higgs:** Another composite background, made up of all backgrounds that contain a Higgs boson ($H \rightarrow \tau\tau$, $H \rightarrow WW$, HV , $t\bar{t} + H$).
- **γ +Jets:** A photon produced in association with any amount of jets. This isn't a background seen in the signal regions, but it is used by the photon template method (see Section 10.2).

- Reducible: This is another composite background, a background is considered “reducible” if it produces fake leptons.

The Z+Jets, Top, diboson and reducible backgrounds are the major backgrounds in the signal regions shown in Chapter 9 and are the ones that must be checked to make sure their contributions are modelled properly. Conversely, the VVV , $V\gamma$ and Higgs backgrounds are negligible in the signal regions so it isn’t necessary to check them. The diboson background is a major background in both 2-lepton and 3-lepton signal regions and in both standard and compressed signal regions. As such, this background has a total of four control regions defined for it: a pair of control regions for the 2-lepton final state and another pair for the 3-lepton final state. Each of these pairs is made up of a control region that uses standard variables and another that uses ISR variables. The top background only has a negligible contribution in the 3-lepton signal regions, so there are only two control regions designed for it: one using standard variables and one using ISR variables. The 2-lepton and 3-lepton control regions are defined in Sections 10.3 and 10.4, respectively.

The Z+Jets and reducible backgrounds are more complicated. The Z+Jets background isn’t well modelled by the Monte Carlo, some other method is required to model it accurately. The method used for this analysis is known as the “photon template method”. This method is a data-driven approach that uses photon events in data to emulate the behaviour of Z^0 bosons decaying into a pair of leptons. This method is fully described in Section 10.2.

The term “reducible” doesn’t really refer to a background or even a conglomeration of backgrounds, but is one-half of a binary classification system for backgrounds, the other being “irreducible”. These classifications depend on the properties of the leptons in those backgrounds. A lepton is called “prompt” if it originates directly from the collision or from a vector boson decay. A “real” lepton is a prompt lepton that is suitably isolated from the rest of the activity in the detector (recall the electron isolation from Chapter 5). A lepton that isn’t prompt is called a “non-prompt” lepton and leptons that aren’t real are referred to as “fake” leptons.

A background is defined as an “irreducible” background if all of its leptons and its E_T^{miss} is real. Conversely, a background is considered “reducible” if it contains one or more fake or non-prompt (“FNP”) leptons or where experimental effects (such as mismeasurement of the energies of leptons/jets or from imperfect overlap removal) lead to significant “fake” E_T^{miss} . FNP leptons can come from the semileptonic decay of b - and c - hadrons, decays in flight of light hadrons, misidentification of light flavour jets and photon conversions. For the 2-lepton final state, such backgrounds originate

from the W +Jets, Single-Top and $t\bar{t}$ processes. For the 3-lepton final state, there are additional contributions from the Z +Jets and WW backgrounds as well as any other process that doesn't produce three real leptons. The reducible background is modelled using the “matrix method” described in Section 10.1.

10.1 Fake Contribution Estimation: the Matrix Method

An FNP lepton is the result of a particle that isn't a lepton being misreconstructed as a lepton. This typically occurs when a lepton in a b - or c -jet is reconstructed on its own rather than as a part of the jet, when a light flavour jet is small and misreconstructed as a lepton or from a photon conversion. The fake contribution to the background is calculated using the matrix method [153].

This method makes use of two identification criteria, called “*loose*” and “*tight*”, which are identical to the “baseline” and “high-purity” criteria defined in Chapter 7, respectively. Let T denote the leptons that pass the *tight* criteria and L denote leptons that pass the *loose* criteria. Since a lepton that passes the *tight* criteria also necessarily passes the *loose* criteria, the *loose* criteria is referred to as the “*inclusive-loose*” criteria. The leptons that pass the *loose* criteria but *not* the *tight* criteria is denoted by E , the *exclusive-loose* criteria.

For the 2-lepton final states all events that have both leptons passing the *loose* criteria have their leptons ordered by p_T and are used by the method. For the 3-lepton final states the hardest lepton is assumed to be real while the remaining pair are checked to see if they both the *loose* criteria. If they do, they are ordered by p_T and used in the method.

All events are classified into one of four categories based on which of the two selection criteria each lepton satisfies: TT (both pass *tight*), TE (leading passes *tight*, subleading passes *loose*), ET (leading *tight*, subleading *loose*) and EE (both pass *loose*). The number of events in each category is N_{TT} , N_{TE} , N_{ET} and N_{EE} , respectively.

After this two probabilities¹ are defined: r and f . The first is the probability that a real lepton that has passed the *loose* criteria will also pass the *tight* criteria. The second is the probability that a fake lepton that has passed the *loose* criteria will also pass the *tight*. Since the events being considered contain two leptons each, four probabilities are required: r_1 , r_2 , f_1 and f_2 . The subscripts “1” and “2” refer to the leading and subleading leptons as ordered by p_T , respectively. Further, these probabilities are generally functions of p_T and η .

¹These are also sometimes called *efficiencies*, the terms are used somewhat interchangeably.

Using these probabilities, the number of leptons in each category can be written in terms of the number of events with two real (N_{LL}^{RR}), one real and one fake (N_{LL}^{RF} and N_{LL}^{FR}) and two fake (N_{LL}^{FF}) leptons passing the *loose* criteria. This is done in Equation 10.1.

$$\begin{bmatrix} N_{TT} \\ N_{TL} \\ N_{LT} \\ N_{LL} \end{bmatrix} = \begin{bmatrix} r_1 r_2 & r_1 f_2 & f_1 r_2 & f_1 f_2 \\ r_1(1-r_2) & r_1(1-f_2) & f_1(1-r_2) & f_1(1-f_2) \\ (1-r_1)r_2 & (1-r_1)f_2 & (1-f_1)r_2 & (1-f_1)f_2 \\ (1-r_1)(1-r_2) & (1-r_1)(1-f_2) & (1-f_1)(1-r_2) & (1-f_1)(1-f_2) \end{bmatrix} \begin{bmatrix} N_{LL}^{RR} \\ N_{LL}^{RF} \\ N_{LL}^{FR} \\ N_{LL}^{FF} \end{bmatrix} \quad (10.1)$$

The quantities N_{TT} , N_{TL} , N_{LT} and N_{LL} are known, taken from data. The quantities N_{LL}^{RR} , N_{LL}^{RF} , N_{LL}^{FR} and N_{LL}^{FF} are the quantities that are unknown but are desired, so inverting the matrix in Equation 10.1 gives the unknown quantities in terms of the known quantities (Equations 10.2 through 10.5).

$$N_{LL}^{RR} = (1-f_1)(1-f_2)N_{T,T} - [(1-f_1)f_2]N_{T,E} - [f_1(1-f_2)]N_{E,T} + f_1f_2N_{E,E} \quad (10.2)$$

$$N_{LL}^{RF} = -(1-f_1)(1-r_2)N_{T,T} + [(1-f_1)r_2]N_{T,E} + [f_1(1-r_2)]N_{E,T} + f_1r_2N_{E,E} \quad (10.3)$$

$$N_{LL}^{FR} = -(1-r_1)(1-f_2)N_{T,T} + [(1-r_1)f_2]N_{T,E} + [r_1(1-f_2)]N_{E,T} + r_1f_2N_{E,E} \quad (10.4)$$

$$N_{LL}^{FF} = (1-r_1)(1-r_2)N_{T,T} - [(1-r_1)r_2]N_{T,E} - [r_1(1-r_2)]N_{E,T} + r_1r_2N_{E,E} \quad (10.5)$$

In order to make use of leptons that pass the *tight* criteria, the four quantities above need to be converted into the corresponding yields for an event that contains two *tight* leptons. This is done in Equations 10.6 through 10.9.

$$N_{TT}^{RR} = r_1 r_2 N_{LL}^{RR} \quad (10.6)$$

$$N_{TT}^{RF} = r_1 f_2 N_{LL}^{RF} \quad (10.7)$$

$$N_{TT}^{FR} = f_1 r_2 N_{LL}^{FR} \quad (10.8)$$

$$N_{TT}^{FF} = f_1 f_2 N_{LL}^{FF} \quad (10.9)$$

To calculate the event yields the probabilities need to be quantified. These probabilities aren't known by default and are calculated using information from the data. Subsection 10.1.1 gives how the real efficiencies r_1 and r_2 are calculated while Subsection 10.1.2 gives how the fake efficiencies f_1 and f_2 are calculated.

10.1.1 Real Lepton Efficiency

The real lepton efficiencies are extracted directly from the data samples. This is done through the construction of a special control region, referred to as the “real lepton control region”, whose only requirements are that all events entering it contain an opposite-sign, same-flavour lepton pair whose invariant mass is within 10 GeV of the Z^0 mass. The extracted real efficiency in terms of p_T and η is shown in Figure 10.1. These plots make it clear that the real efficiency (1) is a function of both p_T and η , (2) takes different values for electrons and muons and (3) is different during different runs. With this in mind, Figure 10.2 shows 2-dimensional plots of the real efficiency in terms of p_T and η . The left hand column in this figure gives the real efficiencies for the 2015 data and the right hand column gives them for the 2016 data. Similarly, the top row gives the real efficiencies for electrons while the bottom row gives the real efficiencies for the muons.




Fig. 10.1 The real efficiencies for electrons and muons as a function of p_T and η . The black points represent the 2015 data while the red points represent the 2016 data.




Fig. 10.2 2-D Plots of the real-efficiency in terms of p_T and η . The top plots show the real-efficiency for the electrons while the bottom row shows it for the muons. The left-hand column shows 2015 data while the right-hand column shows the 2016 data.

10.1.2 Fake Lepton Efficiency

Fake leptons are produced by several different sources and each of these sources have their own fake efficiencies, which are also generally dependent on p_T and η . This means that the fake efficiencies need to be combined into one global, weighted-average fake efficiency (Equation 10.10) that can be used as the fake efficiencies in the matrix given in Equation 10.1.

$$f_{\text{Total}}(p_T, \eta) = \sum_i f_i(p_T, \eta) w_i(p_T, \eta) s_i(p_T, \eta) \quad (10.10)$$

Here:

- The efficiency: $f_i(p_T, \eta)$ is the fake efficiency of the i^{th} fake source
- The weight: $w_i(p_T, \eta)$ reflects the relative amount each fake source contributes to the global fake efficiency
- The scale factor: $s_i(p_T, \eta)$ is applied to take care of the discrepancy between the fake efficiency extracted from data and the fake efficiency extracted from Monte Carlo.

Fake Composition

The relative amount of fake leptons that each fake source contributes isn't knowable from the data alone, so Monte Carlo simulation is used. The Monte Carlo simulates both the exact processes that emerge from collisions (known as the Monte Carlo “truth” particles) and the detector response to these processes²(see Chapter 7, Section 7.2). Thus, the MC can be used to estimate the rate of fake leptons and the relative contributions from each fake source.

The relative amount of fake leptons that each fake source contributes changes from region-to-region and so the weight $w_i(p_T, \eta)$ that is applied to the final fake rate calculation in Equation 10.10 needs to be calculated separately for each region. In practice, this method of calculating the fake efficiency isn't used. Instead, the weights for each fake contribution are calculated in a new control region, dubbed the “inclusive dilepton control region”. This region isn't very discriminatory; only requiring the presence of exactly two *baseline* leptons after the application of the trigger requirements given in Chapter 7.

Fake leptons are produced by one of three categories of processes: misreconstruction of light flavour jets, the semileptonic decays of heavy quarks (b - and c -quarks) and

²That is, how the detector would reconstruct these simulated processes.

photon conversions. Each of these categories contains a number of specific processes that are mismeasured as a lepton. The processes within each category are given in Table 10.1 alongside their Monte Carlo truth classification number.

Table 10.1 Lepton classification system of the MC truth classifier. The fake lepton source processes are labelled with “**Origin**” while the category these processes belong to are labelled with “**Type**”.

Category	Origin		Type	
	Name	Number	Name	Number
Conversion	PhotonConv	5	BkgElectron	4
	ISRPhot	39		
	FSRPhot	40		
	UndrPhot	38		
Heavy Flavour	CharmedMeson	25	Any	
	BottomMeson	26		
	CCbarMeson	27		
	JPsi	28		
	BBarMeson	29		
	CharmedBaryon	32		
	BottomBaryon	33		
Light Flavour	NonDefined	0	Any	
	PionDecay	34		
	KaonDecay	35		
	LightBaryon	30		
	StrangeBaryon	31		
	LightMeson	23		
	StrangeMeson	24		
	DalitzDecay	6		
	PiZero	42		
	Mu	8	NonIsoElectron	3

Each fake lepton is classified in two ways: by the Monte Carlo simulated background that produced it and by the category its source process belongs to. The composition of the fake background in terms of both of these is given in the plots shown in Figures 10.3, 10.4 and 10.5. All of these figures show the composition in terms of lepton p_T , E_T^{miss} and $n_{\text{light jets w/ } p_T > 20 \text{ GeV}}$. Further, the three plots use different lepton selections: Figure 10.3 shows the distributions for dielectron events, Figure 10.4 shows the distributions for electron-muon events and Figure 10.5 shows the distributions for dimuon events. An important result shown by these plots is that the dimuon channel has no fake contribution from the photon conversions, which is as expected.

The fake contribution from each of the categories (and thus the weights $w_i(p_T, \eta)$) used in the final dilepton analyses are extracted directly from the inclusive dilepton control region. The relative contribution of each of the fake categories is computed event-by-event using the plots given in Figures 10.3 through 10.5. The average fraction obtained from each of the three variables is used as the final weight value.






Fig. 10.3 The relative fake contribution in terms of the category of process that produced the fake lepton and the Monte Carlo background that contained the fake lepton for the dielectron channel.



(a) Relative contribution in terms of p_T . (b) Relative contribution in terms of E_T^{miss} .



(c) Relative contribution in terms of $n_{\text{light jets w/ } p_T > 20 \text{ GeV}}$.

Fig. 10.4 The relative fake contribution in terms of the category of process that produced the fake lepton and the Monte Carlo background that contained the fake lepton for the electron-muon channel.




Fig. 10.5 The relative fake contribution in terms of the category of process that produced the fake lepton and the Monte Carlo background that contained the fake lepton for the dimuon channel.

Fake Rates

In addition to the relative contributions of the three fake sources, the rates of each fake source must be found so that Equation 10.10 can be used to calculate the global fake efficiency. Each of the three fake contributions has their own rate, which is extracted directly from a control region defined specifically for that fake source. The heavy flavour and photon conversion have their rates extracted directly from data while the light flavour rates are more complex and use a combination of data and Monte Carlo.

The heavy flavour control regions are defined in Table 10.2. There are two because the heavy flavour component can be produced by both electrons and muons and the triggers need to reflect this. The fake rates are extracted from the data in these control regions, with Figure 10.6 showing the fake rates for electrons and muons over the 2015 and 2016 data taking periods in terms of p_T . The photon conversion control region is defined in Table 10.3. Again, the fake rates are extracted from the data in this control region, with the relevant plot in Figure 10.7. It has been shown that the photon conversions don't contribute fake muons, so a dedicated muon control region isn't defined for the photon conversion like it is for the heavy flavour. Finally, the expected contribution from real leptons (once again defined using the MC truth classifier) is subtracted from the fake rate.

Table 10.2 Fake rate heavy flavour control region.

Cut	Electron	Muon
Lepton Selection	<i>Exactly one baseline muon before overlap removal with $\Delta R_{\text{lep,jet}} < 0.4$.</i>	
$n_{b\text{-tagged jets}}$	= 1	
E_T^{miss} [GeV]	> 50	
$m_T(e, E_T^{\text{miss}})$ [GeV]	< 50	
Trigger	$e\mu$ -Trigger	Di- μ Trigger
Lepton Selection	<i>Exactly one baseline</i> Electron with $\Delta R_{\text{lep,jet}} < 0.3$	<i>Exactly one baseline</i> Muon with $\Delta R_{\text{lep,jet}} < 0.3$
Z^0 -Veto	false	true

The light flavour fake rate makes use of both data and Monte Carlo. Another control region is defined, this time only using the requirement that an event must have a pair of leptons with the same sign, but opposite flavour to be included. The fakes rates in this region are shown in Figure 10.8. Here:

- The filled, coloured circles show the fake rates of the various Monte Carlo backgrounds, using MC truth to identify the light flavour component




Fig. 10.6 The heavy flavour fake rates for electrons and muons used in this analysis.

Table 10.3 Fake rate photon conversion control region.

Cut	Electron
Lepton Selection	<i>Exactly</i> two opposite-sign <i>baseline</i> Muons before overlap removal with $p_T^{\mu 1} > 18$ GeV. <i>Exactly</i> one Baseline Electron .
Trigger	Di- μ Trigger
$n_{b\text{-tagged jets}}$	= 0
E_T^{miss} [GeV]	< 50
$ m_{e\mu} - m_{Z^0} $ [GeV]	< 10

- The open pink rectangles show the total Monte Carlo fake rate
- The open blue circles show the fake rate extracted from data
- The filled black circles show the fake rate from data in the heavy flavour control region

For muons the light flavour fake rates agree with the heavy flavour fake rates from data. For electrons the agreement is much worse, with a much smaller fake rate for leptons from light flavour jets, especially at low transverse momentum. To compensate for this an extra light flavour component is added to the total fake rate for electrons. The light flavour fake rate for elections after this correction is added is shown in Figure 10.9.




Fig. 10.7 The photon conversion fake rates for electrons.




Fig. 10.8 The light flavour fake rates for electrons (a) and muons (b) from Monte Carlo (filled coloured circles, open triangle) and data (open blue circles) compared with the heavy flavour fake rate from data in the heavy flavour control region (filled black circles).




Fig. 10.9 Corrected light flavour fake rate for electrons.

10.2 Z+Jets Estimation: Photon Template Method

The Z+Jets background is one of the primary backgrounds in the 2-lepton signal regions. The Z+Jets events that enter these regions often have fake E_T^{miss} from the mismeasurement of the jets that are also present in them. The mismeasurement is hard to model using Monte Carlo because of the complexity of the detector effects, and the effects of this poor modelling can be seen in Monte Carlo generators in regions that are dominated by the Z+Jets background.

In order to properly model the Z+Jets background³ a data-driven method is employed, called the “photon template” method. This method uses data events with a single photon produced in association with jets and missing energy (these events are labelled “ γ +Jets” events) in order to extract accurate variable distributions.

Z+Jets events and γ +Jets events have a similar topology: a well-measured boson that recoils against jets. To use the γ +Jets events to model the Z+Jets background, the photon momentum must be scaled so that the photon momentum distribution matches the Z^0 momentum distribution. This is done by applying a reweighting factor to each event depending on the value of the photon momentum in that event. This process allows an E_T^{miss} distribution (with the fake contribution already incorporated) to be extracted from the γ +Jets events and used as the Z+Jets E_T^{miss} distribution. Subsection 10.2.3 explains the reweighting process in detail.

In addition to the jet mismeasurement, lepton mismeasurement also contributes to the E_T^{miss} , a contribution that becomes more pronounced as the lepton transverse momentum becomes larger. This effect comes from differences in resolution between the Z^0 and γ ; the γ is *too* well resolved when compared to the Z^0 and a smearing function is used to bring its resolution closer to the Z^0 resolution. This is documented in Subsection 10.2.4.

After the reweighting and smearing procedures the E_T^{miss} distribution, as well as the distributions of the angles between the E_T^{miss} vector and the jets, are well modelled. The next variable that must be modelled is the invariant mass $m_{\ell\ell}$, which is done by using Monte Carlo. This is documented in Subsection 10.2.5.

Knowing the $m_{\ell\ell}$ distributions, the final step is to model the behaviour of the single leptons that emerge from the Z^0 decay. A variety of variables (standard, Recursive Jigsaw and ISR) all require information from single leptons, but the photons in the γ +Jets events don’t decay into lepton pairs. Instead, the photon is split into two

³In this context, the term “properly model” means to create a model of the Z+Jets background which produces distributions that by default contain the contributions from events with fake E_T^{miss} .

“imaginary” leptons whose properties are used to calculate more complex variables. This procedure is described in Subsection 10.2.6.

The photon template method makes use of SHERPA 2.1 Z+Jets and γ +Jets samples. The γ +Jets data has a Monte Carlo simulated $V\gamma$ background subtracted from it (see Subsection 10.2.2).

10.2.1 γ +Jets Event Selection

Events that are used by the photon template must pass some basic selection criteria to ensure that they contain the relevant particles and decay topologies. There are two sets of criteria, one for Z+Jets events and one for γ +Jets events, which are given in Tables 10.4 and 10.5, respectively.

Table 10.4 Z+Jets Event Selection

Trigger	Dilepton Triggers
Lepton Selection	At least two leptons
n_{jets}	≥ 2
$m_{\ell\ell}$ [GeV]	> 12

Table 10.5 γ +Jets Event Selection

Trigger	2015	If $p_T \in (37, 45)$ GeV: HLT_g35_loose_L1EM15 If $p_T \in (45, 50)$ GeV: HLT_g40_loose_L1EM15 If $p_T \in (50, 55)$ GeV: HLT_g45_loose_L1EM15 If $p_T \in (55, 125)$ GeV: HLT_g50_loose_L1EM15 If $p_T > 125$ GeV: HLT_g120_loose
	2016	If $p_T \in (37, 65)$ GeV: HLT_g20_loose_L1EM12 If $p_T \in (65, 75)$ GeV: HLT_g60_loose If $p_T \in (75, 85)$ GeV: HLT_g70_loose If $p_T \in (85, 105)$ GeV: HLT_g80_loose If $p_T \in (105, 145)$ GeV: HLT_g100_loose If $p_T > 145$ GeV: HLT_g140_loose
Lepton Selection	No <i>baseline</i> leptons	
Photon Selection	Leading photon must pass electron-photon ambiguity solver	
n_{Photons}	≥ 1	
n_{jets}	≥ 2	

The triggers in the γ +Jets criteria with $p_T < 125$ GeV (for the 2015 data) and $p_T < 145$ GeV (for the 2016 data) are pre-scaled. This means that only a fraction of events that pass the trigger are recorded and those that are are weighted according to

special “prescale factors” that are provided by the ATLAS Collaboration. The lepton selection given in Table 10.5, “No *baseline* leptons” means that an event is rejected if it contains any leptons that pass the *baseline* selection defined in Chapter 7. This is equivalent to saying that to pass the selection criteria in Table 10.5 the event can’t have any leptons.

10.2.2 γ +Jets Event Contamination

The criteria defined in Table 10.5 is used to select γ +Jets events and it’s important for the resulting data sample be as pure as possible. However, $V\gamma$ events are also able to pass this criteria and contaminate the γ +Jets data sample with events containing photons and real E_T^{miss} whenever the vector bosons decay contains an invisible object (such as $W^\pm \rightarrow \ell^\pm \nu$ or $Z^0 \rightarrow \nu\nu$). This contribution must be subtracted from the γ +Jets data sample in order for the E_T^{miss} distribution extracted from it to be accurate.

The $V\gamma$ contribution is simulated using Monte Carlo. First, a control region is defined (Table 10.6) to extract a normalisation factor for the $V\gamma$ Monte Carlo. The comparison between the γ +Jets data and the $V\gamma$ Monte Carlo in this control region is given in Figure 10.10a. It’s clear from this plot that the $V\gamma$ Monte Carlo has a higher event yield than the data, so the normalisation factor will be between 0 and 1. The exact value it takes is 0.79.

After this the γ +Jets data sample is compared with the combined Monte Carlo γ +Jets and $V\gamma$ samples. This is shown in Figure 10.10b. It is clear that the γ +Jets data sample contains both γ +Jets events as well as $V\gamma$ events. The $V\gamma$ Monte Carlo contribution is then subtracted from the γ +Jets data sample before it is used in the rest of this method.

Table 10.6 Definition of the $V\gamma$ control region used by the photon template to estimate the $V\gamma$ contribution in the γ +Jets background.

Photon Template $V\gamma$ Control Region	
n_{leptons}	≥ 1
n_{photons}	≥ 1
p_T^γ [GeV]	> 140 GeV (photon is unprescaled)
E_T^{miss} [GeV]	> 150 GeV




Fig. 10.10 (a) The γ +Jets data sample compared with the combined γ +Jets and $V\gamma$ Monte Carlo in the $V\gamma$ control region defined in Table 10.6. (b) The γ +Jets data sample compared with the combined γ +Jets and $V\gamma$ Monte Carlo. It's clear that there is a $V\gamma$ contamination that must be removed.

10.2.3 γ +Jets Event Reweighting

The difference in the masses of the photon and Z^0 boson means that the kinematic variable distributions in the γ +Jets and Z +Jets events are different in both their shape and normalisation. In order for the γ +Jets data sample to accurately model the Z +Jets background, the γ +Jets events must be reweighted so that their photon momentum distribution matches the Z^0 momentum distribution of the Z +Jets events. To derive the reweighting factor, first a pair of control regions is defined: CR_RW_ Z +Jets and CR_RW_ γ +Jets. These regions are defined in Table 10.7.

The Z +Jets and γ +Jets yields in these regions are used to determine the reweighting factors. Since the photon and Z^0 transverse momentum distributions have different

Table 10.7 The definition of the Z+Jets and γ +Jets control regions, from which the reweighting factor for the γ +Jets background is derived.

Photon Template - Event Reweighting Control Regions		
	CR_RW_Z+Jets	CR_RW_ γ +Jets
Lepton Selection	OS-SF	—
$p_T^{\ell_1}$ [GeV]	> 25	—
$p_T^{\ell_2}$ [GeV]	> 25	—
Photon Selection	—	At least one photon
p_T^γ [GeV]	—	> 37
$n_{b\text{-tagged jets}}$	$= 0$	$= 0$
n_{jets}	≥ 2	≥ 2
$p_T^{j_1}$ [GeV]	> 30	> 30
$p_T^{j_2}$ [GeV]	> 30	> 30
m_{Eff} [GeV] (only for Standard Tree Regions)	> 400	> 400

shapes, the reweighting factors are dependent on p_T . This reweighting is done for both data and Monte Carlo, using Equations 10.11 and 10.12, respectively.

$$f(x) = \frac{n_{\text{Z+Jets}}(x)}{n_{\gamma+\text{Jets}}(x)} \quad (10.11)$$

$$f(x) = \frac{n_{\text{Data}}(x) - n_{t\bar{t}}(x) - n_{VV}(x)}{n_{\gamma+\text{Jets Data}}(x)} \quad (10.12)$$

Here, x is simply the kinematic variable of interest (in this case, the transverse momentum of the boson). In the Monte Carlo reweighting factor, $n_{\text{Z+Jets}}(x)$ and $n_{\gamma+\text{Jets}}(x)$ are the Z+Jets and γ +Jets Monte Carlo yields in their respective control regions. In the data reweighting factor, $n_{\text{Data}}(x)$ is the total data yield, $n_{t\bar{t}}(x)$ is the $t\bar{t}$ data event yield and $n_{VV}(x)$ is the diboson data event yield, all in CR_RW_Z+Jets. The term $n_{\text{Data}}(x) - n_{t\bar{t}}(x) - n_{VV}(x)$ is then the Z+Jets data event yield in CR_RW_Z+Jets. Finally, $n_{\gamma+\text{Jets Data}}(x)$ is the γ +Jets data event yield in CR_RW_ γ +Jets.

Figure 10.11a shows a Monte Carlo-to-Monte Carlo comparison of the boson transverse momentum distributions before reweighting takes place, while Figure 10.11b shows the same distributions after reweighting.

This reweighting affects other variables as well. Figure 10.12 shows the n_{jets} and H_T distributions after the reweighting has been applied and both distributions show good agreement.




Fig. 10.11 (10.11a) A Monte Carlo-to-Monte Carlo comparison of the $p_T^{Z^0}$ and p_T^γ distributions before reweighting. The distributions have different shapes and the normalisation is incorrect by a factor of 100, since the $\gamma+Jets$ cross-section is much larger than the $Z+Jets$ cross-section. (10.11b) A Monte Carlo-to-Monte Carlo comparison of the $p_T^{Z^0}$ and p_T^γ distributions after reweighting. Here the normalisation and kinematics have been properly corrected.




Fig. 10.12 The $\gamma+\text{Jets}$ and $Z+\text{Jets}$ distributions of the n_{jets} (left) and H_T (right) variables after the reweighting has been applied. The correction in boson transverse momentum also corrects hadronic kinematics.

10.2.4 γ +Jets Event Smearing

After the reweighting procedure good agreement is expected between the Z+Jets and γ +Jets E_T^{miss} distributions. To check this, Figure 10.13 shows the E_T^{miss} distributions for both the Z+Jets and γ +Jets distributions in the ee - and $\mu\mu$ -channels. In this figure, it can clearly be seen that the distributions in the ee -channel have good agreement while the $\mu\mu$ -channel has some disagreement in the $[200, 250]$ GeV E_T^{miss} region.




Fig. 10.13 An MC-to-MC comparison showing the E_T^{miss} distributions of the γ +Jets and Z+Jets backgrounds, made after the reweighting procedure was done.

This discrepancy comes from the difference in resolution between the photon and the muons. To show this, E_T^{miss} is decomposed into components parallel ($E_{T,\parallel}^{\text{miss}}$) and perpendicular ($E_{T,\perp}^{\text{miss}}$) to the direction the boson is travelling in. $E_{T,\perp}^{\text{miss}}$ comes solely from the mismeasurement of jet energies in the perpendicular direction, while $E_{T,\parallel}^{\text{miss}}$ is a convolution of boson and jet mismeasurements. Figures 10.14 and 10.15 show Monte Carlo-to-Monte Carlo and data-to-data comparisons, respectively, of $E_{T,\parallel}^{\text{miss}}$ and $E_{T,\perp}^{\text{miss}}$ distributions. From these distributions it is clear that the perpendicular distributions match fairly well while the parallel distributions have large disagreements, proving that the original discrepancy comes from the resolution difference between the photon and muons and that it affects the E_T^{miss} distributions.




Fig. 10.14 An Monte Carlo-to-Monte Carlo comparison showing the $E_{T,\perp}^{Miss}$ and $E_{T,\parallel}^{Miss}$ distributions of the γ +Jets and Z+Jets backgrounds in the $\mu\mu$ -channel, made after the reweighting procedure was done.





Fig. 10.15 A data-to-data comparison showing the $E_{T,\perp}^{Miss}$ and $E_{T,\parallel}^{Miss}$ distributions of the γ +Jets and Z+Jets backgrounds in the $\mu\mu$ -channel, made after the reweighting procedure was done.


The difference between the resolutions of the bosons is corrected using a special “smearing function” that is extracted from the $E_{T,||}^{\text{miss}}$ distribution. This smearing function is derived by deconvolving the Z+Jets and γ +Jets $E_{T,||}^{\text{miss}}$ distributions, which removes the effect of jet mismeasurement from the $E_{T,||}^{\text{miss}}$ distribution. The output smearing function is then simply the difference between the γ +Jets and Z+Jets response functions.

The smearing function is used by sampling a random number from it (derived from Monte Carlo) which is then added to the photon p_T . The E_T^{miss} of the same event has the same random number added to it, but with the opposite sign, since the photon recoils away from the jet system. The result of this smearing is shown in Figures 10.16 and 10.17, wherein it is clearly seen that the smearing procedure has led to better agreement in both the Monte Carlo-to-Monte Carlo and data-to-data comparisons in both the E_T^{miss} and $E_{T,||}^{\text{miss}}$ distributions.



(a) The Monte Carlo-to-Monte Carlo comparison. (b) The data-to-data comparison.

Fig. 10.16 The $E_{T,||}^{\text{miss}}$ distributions of the Z+Jets and γ +Jets backgrounds in the $\mu\mu$ -channel after both the reweighting and smearing procedures have been done. The two backgrounds have a much better agreement than before the smearing procedure.




(a) The Monte Carlo-to-Monte Carlo comparison. (b) The data-to-data comparison.

Fig. 10.17 The E_T^{Miss} distributions of the Z+Jets and γ +Jets backgrounds in the $\mu\mu$ -channel after both the reweighting and smearing procedures have been done. The two backgrounds have a much better agreement than before the smearing procedure.

10.2.5 γ +Jets Event $m_{\ell\ell}$ Modelling

Photons in γ +Jets events don't decay into lepton pairs but instead enter the electromagnetic calorimeter where they produce a shower (see Chapter 4). This means that the γ +Jets events need to have all lepton-derived information added to them "by hand". The first step in this process is modelling the dilepton invariant mass distribution. The $m_{\ell\ell}$ distribution is used to construct individual leptons (see Subsection 10.2.6) so it is crucial for it to be modelled correctly.

In a real event, $m_{\ell\ell}$ is constructed from two leptons, which means that it's value is connected to the boson p_T mismeasurement. To quantify this connection, a 2-D map is made of $m_{\ell\ell}$ in terms of $E_{T,||}^{miss}$ and $p_T^{Z^0}$ using Z+Jets Monte Carlo. This 2-D map is used to assign $m_{\ell\ell}$ values to each γ +Jets event by mapping the $E_{T,||}^{miss}$ value from a γ +Jets event to the $E_{T,||}^{miss}$ value of a Z+Jets event and by mapping the photon transverse momentum to the Z^0 transverse momentum. The final $m_{\ell\ell}$ modelling is shown in Figure 10.18.



(a) The Monte Carlo-to-Monte Carlo comparison. (b) The data-to-data comparison.

Fig. 10.18 The $m_{\ell\ell}$ distributions using Monte Carlo (a) and data (b) comparing the Z+Jets and γ +Jets events.

10.2.6 γ +Jets Photon Splitting

Kinematic variables such as m_T^W and $H_{n,1}^{\text{PP}}$ require informations from individual leptons to be calculated, so this information must be added to γ +Jets events manually. This is done by “splitting” the photon into two “imaginary” leptons.


In each event the system is boosted to the rest frame of an imaginary boson with the same p_T , η and ϕ as the photon in that event and with a mass consistent with the $m_{\ell\ell}$ value assigned to that event⁴. The photon is then split into two imaginary leptons, each of which has $\frac{m_{\ell\ell}}{2}$ assigned as its momentum and random values assigned as its η and ϕ values in that rest frame. The system is then boosted back into the laboratory frame and the imaginary leptons are used to calculate Lab frame 4-vectors for the leptons.

Figures 10.19 and 10.20 show Monte Carlo and data distributions of the transverse momenta of the leading and subleading leptons, respectively. It’s clear that the agreement is good.

10.2.7 Effect of p_T reweighting


Performing the photon splitting completes the photon template method: a set of events has been constructed using γ +Jets data events that mimic the behaviour of Z+Jets

⁴This is done as an approximation of a boost to the rest frame of the photon, since photons by definition don’t have rest frames.



(a) Monte Carlo-to-Monte Carlo Comparison.

(b) Data-to-data Comparison.

Fig. 10.19 Plots showing leading lepton transverse momentum distributions using Monte Carlo and data for Z+Jets and γ +Jets events.

(a) Monte Carlo-to-Monte Carlo Comparison.

(b) Data-to-Data Comparison.

Fig. 10.20 Plots showing subleading lepton transverse momentum distributions using Monte Carlo and data for Z+Jets and γ +Jets events.

events closely enough that they can be used as a substitute for the Z+Jets background in signal, control and validation regions. The final step is to validate this method and check if any additional normalisation is necessary by using the photon template events as the Z+Jets background in several regions.

To do this, a data-to-Monte Carlo comparison is done in three Z+Jets validation regions: “VR2 ℓ _Low-Zjets”, “VR2 ℓ _High-Zjets” and “VR2 ℓ _ISR-Zjets”. These regions are defined in detail in Chapter 11. Figures 10.21, 10.22 and 10.23 show plots of various variables in these respective regions. A reweighting factor of 0.015 is applied to the photon template Z+Jets background.

10.2.8 Limitations of the Photon Template

The photon template method has some limitations when dealing with events with lower energy photons. The $p_T^\gamma < 55$ GeV trigger has very large pre-scale and almost no events pass the $p_T^\gamma < 37$ GeV trigger. This means that the photon templates estimate in the low-mass and compressed signal regions is not precise and has a large statistical uncertainty. To account for this, the “ABCD” method is used to estimate the yields in this region.

The ABCD method works by defining four sections of a plane in phase space based on a pair of variables. In this case, the variables used are $p_T^{\ell\ell}$ and $H_{4,1}^{\text{PP}}$. The four regions, labelled “A”, “B”, “C” and “D” respectively, correspond to $p_T^{\ell\ell}$ taking values greater or less than 100 GeV and $H_{4,1}^{\text{PP}}$ taking values greater or less than 400 GeV. A schematic diagram showing these regions is given in Figure 10.24.

The number of events in each region are labelled n_A , n_B , n_C and n_D . In terms of these variables, this method is trying to find the total Z+Jets estimate: $n_A + n_B$, and the low $p_T^{\ell\ell}$ Z+Jets estimate: n_A . These are estimated by calculating the ratio $R = \frac{n_C}{n_D}$ where the estimate is known to be accurate and taking it to be the same as $\frac{n_A}{n_B}$. The ratio is calculated with some relaxed requirements on dijet invariant mass ($m_{jj} \in [60, 100]$) and $\frac{p_{T,\text{PP}}^{\text{Lab}}}{p_{T,\text{PP}}^{\text{Lab}} + H_{4,1}^{\text{PP}}} < 0.1$. The dependence of $p_T^{\ell\ell}$ on m_{jj} and $\frac{p_{T,\text{PP}}^{\text{Lab}}}{p_{T,\text{PP}}^{\text{Lab}} + H_{4,1}^{\text{PP}}}$ is given in Figure 10.25. The dependence of the ratio on $H_{4,1}^{\text{PP}}$ is given in Figure 10.26.

The ratio value for $H_{4,1}^{\text{PP}} > 400$ GeV is 3.9 ± 2.1 and the direct γ +Jets estimate for high $p_T^{\ell\ell}$ is 1.29 ± 0.5 . The total Z+Jets estimate is then given by $n_A + n_B = Rn_B + n_B = (R + 1)n_B = 6.32 \pm 3.65$ (statistical uncertainty only).




Fig. 10.21 VR2ℓ_Low-Zjets Region Plots. Left-hand plots are without reweighting, right-hand plots are with reweighting.




Fig. 10.22 VR2 ℓ _High-Zjets Region Plots. Left-hand plots are without reweighting, right-hand plots are with reweighting.




Fig. 10.23 VR2 ℓ _ISR-Zjets Region Plots. Left-hand plots are without reweighting, right-hand plots are with reweighting.







Fig. 10.24 Schematic diagram of the ABCD Method.

(a) $p_T^{\ell\ell}$ in terms of m_{jj} .(b) $p_T^{\ell\ell}$ in terms of $\frac{p_T^{PP}}{p_T^{PP} + H_{4,1}^{PP}}$.Fig. 10.25 Plots of $p_T^{\ell\ell}$ in terms of m_{jj} and $\frac{p_T^{PP}}{p_T^{PP} + H_{4,1}^{PP}}$.Fig. 10.26 Ratio R in term of $\frac{p_T^{PP}}{p_T^{PP} + H_{4,1}^{PP}}$.

10.3 Control Regions: 2-Lepton, 2-Jet and E_T^{miss} Final States

As mentioned at the beginning of this chapter, there are four backgrounds with significant contributions to the signal regions defined in Chapter 9: Z+Jets, VV, Top and the combined reducible background. The Z+Jets and reducible backgrounds have already been reasonably validated by the methods that generate them, so only the diboson and top backgrounds need control regions.

The diboson and top backgrounds have two control regions each: one using the standard tree variables and one using the ISR variables. This is done so that each of these backgrounds can be validated in both the standard tree regions and the compressed tree regions. The diboson control regions are labelled “CR2 ℓ _VV” and “CR2 ℓ _ISR_VV” for the standard and compressed regions, respectively. The top control regions are likewise labelled “CR2 ℓ _Top” and “CR2 ℓ _ISR_Top”.

These control regions (and all control regions generally) are designed with two properties in mind: they must be orthogonal to all signal regions and they must be “close” to the signal regions in phase space. Orthogonality is required for the benefit of the fitting algorithms used (see Chapter 6). For two regions to be “close” to one-another means that those regions occupy similar places in phase space. This is once again done for the benefit of the fit and this closeness is often achieved simply by inverting one of the cuts in a signal region.

10.3.1 Standard Tree Regions

This subsection will define the two 2-lepton control regions that use standard tree variables: CR2 ℓ _VV and CR2 ℓ _Top. Their full definitions are given in Table 10.8.

CR2 ℓ _VV cuts on that same variables that define the SR2 ℓ _Intermediate and SR2 ℓ _High regions and has an additional cut on m_T^W . The lepton selection “ $n_{\text{leptons}} \in [3, 4]$ ” is used to select diboson events (such as $ZZ \rightarrow \ell\ell\ell\ell$ and $WZ \rightarrow \ell\nu\ell\ell$) while also making the region orthogonal to all the 2-lepton signal regions. The reconstruction assigns lepton pairs with opposite-sign, same-flavour and invariant mass closest to m_{Z^0} as coming from the Z^0 boson decay while third leptons are assigned as coming from a $W^\pm$ decay. Any remaining leptons have their p_T added to the E_T^{miss} . For events with 3-leptons, a special requirement is placed on m_T^W to ensure orthogonality with the 3-lepton control regions. Figure 10.27 shows the kinematic variable distributions for this control region.

Table 10.8 2-lepton control region definitions - standard tree

Control Regions: 2-Lepton Standard Tree		
Cut	CR2 ℓ _VV	CR2 ℓ _Top
Lepton Selection	$n_{\text{leptons}} \in [3, 4]$	$\ell_f^+ \ell_f^-$
n_{jets}	≥ 2	≥ 2
$n_{b\text{-tagged jets}}$	$= 0$	$= 1$
$p_T^{\ell_1}$ [GeV]	> 25	> 25
$p_T^{\ell_2}$ [GeV]	> 25	> 25
$p_T^{j_1}$ [GeV]	> 30	> 30
$p_T^{j_2}$ [GeV]	> 30	> 30
$m_{\ell\ell}$ [GeV]	$\in (80, 100)$	$\in (80, 100)$
m_{jj} [GeV]	> 20	$\in (40, 250)$
m_T^W [GeV]	$\in (70, 100)$, in events with $n_{\text{leptons}} = 3$	—
$H_{4,1}^{\text{PP}}$ [GeV]	> 200	> 400
$\frac{p_{T,\text{PP}}^{\text{Lab}}}{p_{T,\text{PP}}^{\text{Lab}} + H_{4,1}^{\text{PP}}}$	< 0.05	< 0.05
$\frac{\min(H_{1,1}^{\text{Pa}}, H_{1,1}^{\text{Pb}})}{\min(H_{2,1}^{\text{Pa}}, H_{2,1}^{\text{Pb}})}$	> 0.2	> 0.5
$\Delta\phi_V^{\text{P}}$	$\in (0.3, 2.8)$	$\in (0.3, 2.8)$

CR2 ℓ _Top cuts on the exact same variables as SR2 ℓ _Intermediate and SR2 ℓ _High do. The requirement that there be exactly one b -tagged jet is to select events which contain top-quarks (recall the discussion on b -tagging in Subsection 5.2.3) and to ensure orthogonality between CR2 ℓ _Top and the 2-lepton signal regions. Figure 10.28 shows the kinematic variable distributions for this control region.

Table 10.9 shows the output of the background-only fit for these two regions.




Fig. 10.27 CR2ℓ_VV - Kinematic Variable Distributions.




Fig. 10.28 CR2ℓ_Top - Kinematic Variable Distributions.

Table 10.9 Yields in the control regions for the 2-lepton standard tree.

table.results.yields channel	CR2 ℓ _VV	CR2 ℓ _Top
Observed events	60	97
Fitted bkg events	59.98 ± 7.74	96.97 ± 16.82
Fitted MM events	0.89 ± 0.27	0.00 ± 0.00
Fitted HIGGS events	0.41 ± 0.33	0.01 ± 0.01
Fitted VGAMMA events	0.90 ± 0.04	0.00 ± 0.00
Fitted TOP events	0.00 ± 0.00	60.19 ± 10.82
Fitted TTV events	1.05 ± 0.19	1.40 ± 0.24
Fitted VVV events	0.24 ± 0.05	0.01 ± 0.01
Fitted VV events	56.49 ± 7.75	3.97 ± 0.97
Fitted Z events	0.00 ± 0.00	31.39 ± 15.07
MC exp. SM events	65.34	103.24
MC exp. MM events	0.89	0.00
MC exp. HIGGS events	0.41	0.01
MC exp. VGAMMA events	0.90	0.00
MC exp. TOP events	0.00	66.07
MC exp. TTV events	1.05	1.41
MC exp. VVV events	0.24	0.01
MC exp. VV events	61.85	4.35
MC exp. Z events	0.00	31.39


10.3.2 Compressed Tree Regions

Each of the 2-lepton standard tree control regions has a corresponding compressed tree control region. These are CR2 ℓ _ISR_VV and CR2 ℓ _ISR_Top and are defined in Table 10.10.

Table 10.10 2-lepton control region definitions - compressed tree

Control Regions: 2-Lepton Compressed Tree		
Cut	CR2 ℓ _ISR_VV	CR2 ℓ _ISR_Top
Lepton Selection	$n_{\text{leptons}} \in [3, 4]$	$\ell_f^+ \ell_f^-$
n_{jets}	> 2	$\in [3, 4]$
$n_{b\text{-tagged jets}}$	$= 0$	$= 1$
n_{jets}^S	≥ 2	$= 2$
$n_{\text{jets}}^{\text{ISR}}$	≥ 1	≥ 1
$p_T^{\ell_1}$ [GeV]	> 25	> 25
$p_T^{\ell_2}$ [GeV]	> 25	> 25
$p_T^{j_1}$ [GeV]	> 30	> 30
$p_T^{j_2}$ [GeV]	> 30	> 30
m_Z [GeV]	$\in (80, 100)$	$\in (50, 200)$
m_J [GeV]	> 20	$\in (50, 200)$
$\Delta\phi_{\text{ISR},I}^{\text{CM}}$	> 2.0	> 2.8
R_{ISR}	$\in (0.0, 0.5)$	$\in (0.4, 0.75)$
$p_{T,\text{ISR}}^{\text{CM}}$ [GeV]	> 50	> 180
$p_{T,I}^{\text{CM}}$ [GeV]	> 50	> 100
$p_{T,\text{CM}}$ [GeV]	< 30	< 20

The 2-lepton compressed control regions cut on the same variables as SR2 ℓ _ISR. CR2 ℓ _ISR_VV uses the same strategy as CR2 ℓ _VV: requiring 3 or 4 leptons to enrich the region in VV events and to have it be orthogonal with SR2 ℓ _ISR, which requires exactly 2 leptons. CR2 ℓ _ISR_Top uses the same “single b -tagged jet” requirement that CR2 ℓ _Top uses, again to select events containing top-quarks and to render CR2 ℓ _ISR_Top orthogonal to SR2 ℓ _ISR. In addition to this, CR2 ℓ _ISR_Top is defined with broader m_Z and m_J intervals compared to CR2 ℓ _Top. This is because if CR2 ℓ _ISR_Top used the smaller intervals that CR2 ℓ _Top uses, the region would have too few events to populate it, which leads to statistical issues. Figures 10.29 and 10.30 show the kinematic variable distributions for the CR2 ℓ _ISR_VV and CR2 ℓ _ISR_Top control regions, respectively. Table 10.11 shows the output yields of the background-only fits for these regions.

Fig. 10.29 CR2 ℓ _ISR_VV - Kinematic Variable Distributions.


Fig. 10.30 CR2 ℓ _ISR_Top - Kinematic Variable Distributions.

Table 10.11 Yields in the control regions for the 2-lepton compressed tree.

table.results.yields channel	CR2 ℓ _ISR_VV	CR2 ℓ _ISR_Top
Observed events	28	93
Fitted bkg events	28.07 ± 5.29	93.09 ± 9.64
Fitted MM events	0.00 ± 0.00	0.00 ± 0.00
Fitted HIGGS events	0.21 ± 0.17	0.03 ± 0.01
Fitted VGAMMA events	0.04 ± 0.00	0.00 ± 0.00
Fitted TOP events	—	89.50 ± 9.65
Fitted TTV events	0.45 ± 0.08	0.46 ± 0.14
Fitted VVV events	0.02 ± 0.01	0.01 ± 0.01
Fitted VV events	27.35 ± 5.30	0.99 ± 0.31
Fitted Z events	0.00 ± 0.00	2.10 ± 1.03
MC exp. SM events	29.65	94.19
MC exp. MM events	0.00	0.00
MC exp. HIGGS events	0.00	0.03
MC exp. VGAMMA events	0.04	0.00
MC exp. TOP events	—	90.54
MC exp. TTV events	0.45	0.46
MC exp. VVV events	0.02	0.01
MC exp. VV events	29.14	1.06
MC exp. Z events	0.00	2.09

10.4 Control Regions: 3-Lepton and E_T^{miss} Final States

In the 3-leptons signal regions defined in Chapter 9, every background except for the diboson background has a negligible contribution to the total background yield. This means that it's only necessary to create two 3-lepton control regions: one for the standard tree variables and one for the compressed tree variables. These are called “CR3 ℓ _VV” and “CR3 ℓ _ISR_VV”, respectively.

10.4.1 Standard Tree Regions

The construction of the 3-lepton standard tree control region follows a similar strategy to the construction of 3-lepton standard tree signal regions. CR3 ℓ _VV cuts on the same variables as these three regions, except for $\frac{H_{1,1}^{\text{Pb}}}{H_{2,1}^{\text{Pb}}}$ which SR3 ℓ _Low doesn't cut on either. This region requires three leptons, with the lepton pair whose invariant mass is closest to m_{Z^0} assigned as coming from the Z^0 decay and the remaining lepton assigned as coming from the $W^\pm$. The remaining lepton is used to calculate m_T^W , which is required to be less than 70 GeV to ensure orthogonality between CR3 ℓ _VV and the 3-lepton signal regions. Table 10.12 contains the full definition of CR3 ℓ _VV. Figure 10.31 shows the kinematic variable distributions for CR3 ℓ _VV.

Table 10.12 3-lepton control region definitions - standard tree

Control Regions: 3-Lepton Standard Tree	
Cut	CR3 ℓ _VV
Lepton Selection	$(Z^0 \rightarrow) \ell_{f_1}^+ \ell_{f_1}^- \ell_{f_2}^\pm$
n_{jets}	< 3
$n_{b\text{-tagged jets}}$	$= 0$
$p_T^{\ell_1}$ [GeV]	> 60
$p_T^{\ell_2}$ [GeV]	> 40
$p_T^{\ell_3}$ [GeV]	> 30
$m_{\ell\ell}$ [GeV]	$\in (75, 105)$
m_T^W [GeV]	< 70
$H_{3,1}^{\text{PP}}$ [GeV]	> 250
$\frac{p_{T,\text{PP}}^{\text{Lab}}}{p_{T,\text{PP}}^{\text{Lab}} + H_{3,1}^{\text{PP}}}$	< 0.2
$\frac{H_{T,3,1}^{\text{PP}}}{H_{3,1}^{\text{PP}}}$	> 0.75




Fig. 10.31 CR3ℓ_VV - Kinematic Variable Distributions

10.4.2 Compressed Tree Regions

CR3 ℓ _ISR_VV cuts on all of the same variables that SR3 ℓ _ISR cuts on and most of the cuts have the same values. The biggest difference between the two is that the cut on m_T^W is inverted to ensure orthogonality between CR3 ℓ _ISR_VV and SR3 ℓ _ISR. Table 10.13 gives the full definition of CR3 ℓ _ISR_VV and Figure 10.32 shows its kinematic variable plots.

Table 10.13 3-lepton control region definitions - compressed tree

Control Regions: 3-Lepton Compressed Tree	
Cut	CR3 ℓ _ISR_VV
Lepton Selection	$(Z^0 \rightarrow) \ell_{f_1}^+ \ell_{f_1}^- \ell_{f_2}^\pm$
n_{jets}	≥ 1
$n_{b\text{-tagged jets}}$	$= 0$
$p_T^{\ell_1}$ [GeV]	> 25
$p_T^{\ell_2}$ [GeV]	> 25
$p_T^{\ell_3}$ [GeV]	> 20
$m_{\ell\ell}$ [GeV]	$\in (75, 105)$
m_T^W [GeV]	< 100
$\Delta\phi_{\text{ISR,I}}^{\text{CM}}$	> 2.0
R_{ISR}	$\in [0.55, 1.0]$
$p_{T,\text{ISR}}^{\text{CM}}$ [GeV]	> 80
$p_{T,\text{I}}^{\text{CM}}$ [GeV]	> 60
$p_{T,\text{CM}}$ [GeV]	< 25




Fig. 10.32 CR3ℓ_ISR_VV

Finally, Table 10.14 shows the yields given by the background-only fits for CR3 ℓ _VV and CR3 ℓ _ISR_VV.

Table 10.14 Yields in the control regions for the 3-lepton standard and compressed trees.

table.results.yields channel	CR3 ℓ _VV	CR3 ℓ _ISR_VV
Observed events	331	98
Fitted bkg events	330.99 ± 18.19	98.05 ± 9.90
Fitted HIGGS events	13.79 ± 1.18	$0.01^{+0.16}_{-0.01}$
Fitted MM events	36.83 ± 11.14	3.49 ± 1.05
Fitted TTV events	1.15 ± 0.17	0.90 ± 0.16
Fitted VVV events	1.13 ± 0.06	0.22 ± 0.14
Fitted VV events	278.09 ± 18.32	93.43 ± 9.90
MC exp. SM events	307.35	87.28
MC exp. HIGGS events	13.79	0.01
MC exp. MM events	36.83	3.49
MC exp. TTV events	1.15	0.90
MC exp. VVV events	1.13	0.22
MC exp. VV events	254.45	82.66

Chapter 11

Validation Regions

Chapter 10 described how control regions are used to derive statistically rigorous background yields for the major backgrounds present in the signal regions. These yields are needed to provide an accurate estimate of the background yields in the signal regions, which are derived using transfer factors (recall Equation 6.5). Before this can be done, this method of using transfer factors to estimate the background yields in other parts of the phase space must itself be validated.

This validation is performed using a new set of regions, called “validation regions” (or “VRs”). These regions are defined to be both close to and statistically independent of both the control and signal regions. These properties allow these regions to be used to validate the background yields obtained in the control regions while being reasonably confident that the background yields will be correct once they are transferred into the signal regions.

In practice, this validation is done by redefining the transfer factor equation (Equation 6.5) to use validation regions instead of signal regions (Equation 11.1).

$$N_B(\text{Est}, \text{VR}) = N_B(\text{Obs}, \text{CR}) \times \left[\frac{\text{MC}_B(\text{Obs}, \text{VR})}{\text{MC}_B(\text{Obs}, \text{CR})} \right] \quad (11.1)$$

This equation is used to estimate the yields of a given background in a given validation region. Once the yields have been calculated a data-to-Monte Carlo comparison is done in the validation region using a background-only fit. If the data and Monte Carlo yields and distributions match then the backgrounds are considered validated and the analysis moves on to the unblinding stage.

Like the signal and control regions, the validation regions are divided according to which final state they study and whether they use standard tree variables or compressed

tree variables. Section 11.1 describes the validation regions studying the 2-lepton final state while Section 11.2 describes to validation regions studying the 3-lepton final state.

11.1 Validation Regions: 2-Lepton, 2-Jet and E_T^{miss} Final States

The typical analysis strategy is to assign to every control region a corresponding validation region. This is implemented for this search, but some additional nuance is introduced by the photon template method.

Chapter 10 defined four 2-lepton control regions: CR2ℓ_VV, CR2ℓ_Top, CR2ℓ_ISR_VV and CR2ℓ_ISR_Top. Each of these control regions have a validation region defined for them. Respectively, they are labelled: VR2ℓ_VV, VR2ℓ_Top, VR2ℓ_ISR_VV and VR2ℓ_ISR_Top. Each of these regions are typically defined by inverting a single cut to ensure orthogonality between the validation and control regions.

In addition to these, three special 2-lepton regions are defined: VR2ℓ_Z+Jets_Low, VR2ℓ_Z+Jets_High and VR2ℓ_ISR_Z+Jets. These regions are made to validate the Z+Jets estimate that the photon template method provides. The VR2ℓ_Z+Jets_Low and VR2ℓ_Z+Jets_High regions both use standard tree variables but target different parts of phase space. VR2ℓ_Z+Jets_High is made to validate the direct photon template estimate that is considered accurate for high $p_T^{\ell\ell}$ events, while VR2ℓ_Z+Jets_Low is used to validate the low $p_T^{\ell\ell}$ estimate that comes from the ABCD method (recall Subsection 10.2.8). Finally, VR2ℓ_ISR_Z+Jets is used to validate the Z+Jets estimate in the 2-lepton compressed region.

Subsection 11.1.1 summarises the four standard tree validation regions and Subsection 11.1.2 summarises the three compressed tree validation regions.

11.1.1 Standard Tree Regions

The four standard tree regions are: VR2ℓ_VV, VR2ℓ_Top, VR2ℓ_Z+Jets_Low and VR2ℓ_Z+Jets_High and are defined in Table 11.1.

VR2ℓ_VV cuts on most of the same variables as CR2ℓ_VV and the 2-lepton standard tree regions. It's lepton selection, events are to contain exactly two lepton, with opposite-sign and same-flavour, is chosen so that it is orthogonal to its corresponding control region. Orthogonality to the signal regions is done using cuts on m_{jj} and $\frac{\min(H_{1,1}^{Pa}, H_{1,1}^{Pb})}{\min(H_{2,1}^{Pa}, H_{2,1}^{Pb})}$. The requirement that $m_{jj} \in (40, 70) \cup (90, 500)$ is almost an inversion of the m_{jj} cut in SR2ℓ_Low, and renders SR2ℓ_Low and VR2ℓ_VV orthogonal. The cut

$\frac{\min(H_{1,1}^{\text{Pa}}, H_{1,1}^{\text{Pb}})}{\min(H_{2,1}^{\text{Pa}}, H_{2,1}^{\text{Pb}})} \in (0.4, 0.8)$ likewise renders VR2 ℓ _VV orthogonal to SR2 ℓ _Intermediate and SR2 ℓ _High. Finally, VR2 ℓ _VV has a unique cut on $H_{1,1}^{\text{PP}}$ which is used to suppress the Z+Jets background. Figure 11.1 shows the kinematic distributions of the variables in this region.

Like CR2 ℓ _Top, VR2 ℓ _Top is defined by its requirement of exactly one b -tagged jet to ensure the presence of top-quarks and renders it orthogonal to the 2-lepton signal regions. Additionally, it has a quasi-inverted requirement on the dilepton invariant mass: $m_{\ell\ell} \in (20, 80) \cup (100, \infty)$. This requirement is chosen to make VR2 ℓ _Top orthogonal to CR2 ℓ _Top. Figure 11.2 shows the kinematic distributions of the variables in this region.

The VR2 ℓ _Z+Jets_Low and VR2 ℓ _Z+Jets_High regions take all but one of their cuts from the 2-lepton signal regions. The only cut that distinguishes them is the requirement on the dijet invariant mass: $m_{jj} \in (60, 100) \cup (100, 180)$. This requirement is chosen to remove any contamination from W+Jets events and also renders these two regions orthogonal to the signal regions. Figures 11.3 and 11.4 show the kinematic variable distributions in VR2 ℓ _Z+Jets_Low and VR2 ℓ _Z+Jets_High, respectively.

Finally, Table 11.2 gives the yields of the background-only fits in these four regions.

Table 11.1 2-lepton validation region definitions - standard tree

Validation Regions: 2-Lepton Standard Tree				
Cut	VR2 ℓ _VV	VR2 ℓ _Top	VR2 ℓ _Z+Jets_Low	VR2 ℓ _Z+Jets_High
Lepton Selection	$\ell_f^+ \ell_f^-$	$\ell_f^+ \ell_f^-$	$\ell_f^+ \ell_f^-$	$\ell_f^+ \ell_f^-$
n_{jets}	≥ 2	≥ 2	$= 2$	≥ 2
$n_{b\text{-tagged jets}}$	$= 0$	$= 1$	$= 0$	$= 0$
$p_T^{\ell_1}$ [GeV]	> 25	> 25	> 25	> 25
$p_T^{\ell_2}$ [GeV]	> 25	> 25	> 25	> 25
$p_T^{j_1}$ [GeV]	> 30	> 30	> 30	> 30
$p_T^{j_2}$ [GeV]	> 30	> 30	> 30	> 30
$m_{\ell\ell}$ [GeV]	$\in (80, 100)$	$\in (20, 80) \cup (100, \infty)$	$\in (80, 100)$	$\in (80, 100)$
m_{jj} [GeV]	$\in (40, 70) \cup (90, 500)$	$\in (40, 250)$	$\in (0, 60) \cup (100, 180)$	$\in (0, 60) \cup (100, 180)$
$H_{1,1}^{\text{PP}}$ [GeV]	> 250	—	—	—
$H_{4,1}^{\text{PP}}$ [GeV]	> 400	> 400	> 400	> 600
$\frac{p_{T,\text{PP}}^{\text{Lab}}}{p_{T,\text{PP}} + H_{4,1}^{\text{PP}}}$	< 0.05	< 0.05	< 0.05	< 0.05
$\frac{H_{1,1}^{\text{PP}}}{H_{4,1}^{\text{PP}}}$	—	—	$\in (0.35, 0.6)$	—
$\frac{\min(H_{1,1}^{\text{Pa}}, H_{1,1}^{\text{Pb}})}{\min(H_{2,1}^{\text{Pa}}, H_{2,1}^{\text{Pb}})}$	$\in (0.4, 0.8)$	> 0.5	—	> 0.4
$\Delta\phi_V^{\text{P}}$	$\in (0.3, 2.8)$	$\in (0.3, 2.8)$	—	$\in (0.3, 2.8)$




Fig. 11.1 VR2ℓ_VV - Kinematic Variable Distributions.





Fig. 11.2 VR2l_Top - Kinematic Variable Distributions.

Fig. 11.3 VR2 ℓ _Z+Jets_Low - Kinematic Variable Distributions.


Fig. 11.4 VR2 ℓ _Z+Jets_High - Kinematic Variable Distributions.

Table 11.2 Yields in the validation regions for the 2-lepton standard tree.

table.results.yields channel	VR2 ℓ _Z+Jets_Low	VR2 ℓ _Z+Jets_High	VR2 ℓ _VV	VR2 ℓ _Top
Observed events	263	77	72	491
Fitted bkg events	260.67 $\pm$ 130.32	69.49 $\pm$ 25.50	60.68 $\pm$ 13.45	423.63 $\pm$ 105.22
Fitted MM events	0.00 $\pm$ 0.00	0.00 $\pm$ 0.00	0.00 $\pm$ 0.00	0.61 $\pm$ 0.18
Fitted HIGGS events	0.43 ^{+1.52} _{-0.43}	0.00 $\pm$ 0.00	0.23 $\pm$ 0.02	0.12 $\pm$ 0.05
Fitted VGAMMA events	2.85 $\pm$ 0.40	0.02 $\pm$ 0.00	0.01 $\pm$ 0.00	0.51 $\pm$ 0.02
Fitted TOP events	14.96 $\pm$ 4.81	1.72 $\pm$ 0.69	11.95 $\pm$ 3.70	415.21 $\pm$ 105.22
Fitted TTV events	0.17 $\pm$ 0.07	0.17 $\pm$ 0.04	0.41 $\pm$ 0.07	1.08 $\pm$ 0.09
Fitted VVV events	0.07 $\pm$ 0.02	0.06 $\pm$ 0.01	0.15 $\pm$ 0.01	0.01 ^{+0.02} _{-0.01}
Fitted VV events	30.40 $\pm$ 7.26	16.26 $\pm$ 3.08	39.49 $\pm$ 12.74	3.72 $\pm$ 0.87
Fitted Z events	211.79 $\pm$ 128.43	51.26 $\pm$ 24.83	8.44 $\pm$ 0.61	2.37 $\pm$ 0.11
MC exp. SM events	265.02	71.19	65.59	463.63
MC exp. MM events	0.00	0.00	0.00	0.61
MC exp. HIGGS events	0.43	0.00	0.23	0.12
MC exp. VGAMMA events	2.85	0.02	0.01	0.51
MC exp. TOP events	16.39	1.88	13.09	454.85
MC exp. TTV events	0.17	0.17	0.41	1.08
MC exp. VVV events	0.07	0.06	0.15	0.01
MC exp. VV events	33.32	17.81	43.25	4.08
MC exp. Z events	211.79	51.26	8.44	2.37

11.1.2 Compressed Tree Regions

The three compressed tree regions are: VR2 ℓ _ISR_VV, VR2 ℓ _ISR_Top and VR2 ℓ _ISR_Z+Jets. VR2 ℓ _ISR_VV and VR2 ℓ _ISR_Top are defined as the VR counterpart to the CR2 ℓ _ISR_VV and CR2 ℓ _ISR_Top control regions while VR2 ℓ _ISR_Z+Jets serves as the region in which the ISR variables output by the photon template are validated. Table 11.3 shows the definition of these three regions.

Table 11.3 2-lepton validation region definitions - compressed tree

Validation Regions: 2-Lepton Compressed Tree			
Cut	VR2 ℓ _ISR_VV	VR2 ℓ _ISR_Top	VR2 ℓ _ISR_Z+Jets
Lepton Selection	$\in [3, 4]$	$\ell_f^+ \ell_f^-$	$\ell_f^+ \ell_f^-$
n_{jets}	≥ 3	$\in [3, 4]$	$\in [3, 5]$
$n_{b\text{-tagged jets}}$	$= 0$	$= 1$	$= 0$
n_{jets}^S	≥ 2	$= 2$	≥ 1
$n_{\text{jets}}^{\text{ISR}}$	≥ 1	≥ 1	≥ 1
$p_T^{\ell_1}$ [GeV]	> 25	> 25	> 25
$p_T^{\ell_2}$ [GeV]	> 25	> 25	> 25
$p_T^{j_1}$ [GeV]	> 20	> 30	> 30
$p_T^{j_2}$ [GeV]	> 20	> 30	> 30
m_Z [GeV]	$\in (20, 80) \cup (100, \infty)$	$\in (50, 200)$	$\in (80, 100)$
m_J [GeV]	> 20	$\in (50, 200)$	$\in (0, 50) \cup (110, \infty)$
$\Delta\phi_{\text{ISR},I}^{\text{CM}}$	> 2.0	> 2.8	—
R_{ISR}	$\in (0.0, 1.0)$	$\in (0.4, 0.75)$	—
$p_{T,\text{ISR}}^{\text{CM}}$ [GeV]	> 70	> 180	> 180
$p_{T,I}^{\text{CM}}$ [GeV]	> 70	> 100	> 100
$p_{T,\text{CM}}$ [GeV]	< 30	> 20	< 20

VR2 ℓ _ISR_VV cuts on the same variables as CR2 ℓ _ISR_VV. It shares the CR2 ℓ _ISR_VV requirement that events have three or four leptons, both to enrich the region with VV events and to make it orthogonal to the 2-lepton compressed signal region. Additionally, VR2 ℓ _ISR_VV is made orthogonal to CR2 ℓ _ISR_VV by quasi-inverting the m_Z cut: $m_Z \in (20, 80) \cup (100, \infty)$. Figure 11.5 shows the kinematic variable distributions for this region.

VR2 ℓ _ISR_Top is almost identical to the CR2 ℓ _ISR_Top region, with all but one cut having the same values. This cut is on $p_{T,\text{CM}}$, which is inverted with respect to the $p_{T,\text{CM}}$ cut in CR2 ℓ _ISR_Top. This renders VR2 ℓ _ISR_Top orthogonal to its control region counterpart. Again, this region requires events to possess a b -tagged jet

to ensure the presence of top-quarks and to render it orthogonal to the compressed signal region. Figure 11.6 shows the relevant distributions.

The validation region for ISR variables produced by the photon template, $\text{VR}2\ell_ \text{ISR_Z+Jets}$ is defined with exactly two leptons and between three and five jets. Both the sparticle system and the ISR system has to have at least one of these jets assigned to it, while none of the jets are allowed to have a b -tag. An inverted cut on the mass of the jet system: $m_J \in (0, 50) \cup (110, \infty)$ is done to make this region orthogonal to $\text{SR}2\ell_ \text{ISR}$. Figure 11.7 shows the relevant kinematic variable distributions.

Finally, Table 11.4 gives the results of the background-only fits in these regions.





Fig. 11.5 $\text{VR}2\ell_ \text{ISR_VV}$ - Kinematic Variable Distributions.

Fig. 11.6 VR2 ℓ _ISR_Top - Kinematic Variable Distributions.


Fig. 11.7 VR2 ℓ _ISR_Z+Jets - Kinematic Variable Distributions.

Table 11.4 Yields in the validation regions for the 2-lepton compressed tree.

table.results.yields channel	VR2 ℓ _ISR_VV	VR2 ℓ _ISR_Top	VR2 ℓ _ISR_Z+Jets
Observed events	13	113	248
Fitted bkg events	12.09 ± 3.62	110.45 ± 17.57	308.43 ± 101.41
Fitted MM events	3.22 ± 1.02	0.00 ± 0.00	0.00 ± 0.00
Fitted HIGGS events	0.12 ± 0.03	$0.01^{+0.18}_{-0.01}$	0.42 ± 0.18
Fitted VGAMMA events	0.14 ± 0.01	0.01 ± 0.00	1.30 ± 0.10
Fitted TOP events	0.00 ± 0.00	107.16 ± 17.52	39.96 ± 7.60
Fitted TTV events	0.57 ± 0.11	0.66 ± 0.06	1.11 ± 0.24
Fitted VVV events	0.15 ± 0.04	0.00 ± 0.00	0.15 ± 0.03
Fitted VV events	7.89 ± 3.60	0.97 ± 0.25	66.73 ± 15.22
Fitted Z events	0.00 ± 0.00	1.64 ± 0.80	198.76 ± 98.70
MC exp. SM events	12.62	111.82	313.25
MC exp. MM events	3.22	0.00	0.00
MC exp. HIGGS events	0.12	0.01	0.42
MC exp. VGAMMA events	0.14	0.01	1.30
MC exp. TOP events	0.00	108.48	40.45
MC exp. TTV events	0.57	0.66	1.11
MC exp. VVV events	0.15	0.00	0.15
MC exp. VV events	8.42	1.03	71.06
MC exp. Z events	0.00	1.64	198.76

11.2 Validation Regions: 3-Lepton and E_T^{miss} Final States


As discussed in Chapter 10, there are only two 3-lepton control regions: the standard tree region CR3 ℓ _VV and the compressed tree region CR3 ℓ _ISR_VV. Following this trend, each of these control regions has a corresponding validation region defined for it. These are “VR3 ℓ _VV” and “VR3 ℓ _ISR_VV”. Since the only major background in the 3-lepton signal regions is the diboson background, there is no need for any special VRs to handle the Z+Jets contribution like in Section 11.1.

11.2.1 Standard Tree Regions

VR3 ℓ _VV is defined in a similar manner to its control region counterpart. The region demands events have three leptons: the pair of leptons whose invariant mass is closest to the Z^0 mass are taken as coming from the Z^0 decay while the remaining lepton is assigned as coming from the $W^\pm$ decay and is used to calculate m_T^W . The cut on m_T^W is used to make VR3 ℓ _VV orthogonal to both CR3 ℓ _VV and the signal regions. The cut is defined as $m_T^W \in (70, 100)$, which sits between the “ $m_T^W < 70$ GeV” cut in CR3 ℓ _VV and the “ $m_T^W > 100/130/150$ GeV” cuts in the 3-lepton standard tree regions. Table 11.5 gives the definition of VR3 ℓ _VV and Figure 11.8 shows its kinematic variable distributions.

Table 11.5 3-lepton validation region definitions - standard tree

Validation Regions: 3-Lepton Standard Tree	
Cut	VR3 ℓ _VV
Lepton Selection	$(Z^0 \rightarrow) \ell_{f_1}^+ \ell_{f_1}^- \ell_{f_2}^\pm$
n_{jets}	< 3
$n_{b\text{-tagged jets}}$	$= 0$
$p_T^{\ell_1}$ [GeV]	> 60
$p_T^{\ell_2}$ [GeV]	> 40
$p_T^{\ell_3}$ [GeV]	> 30
$m_{\ell\ell}$ [GeV]	$\in (75, 105)$
m_T^W [GeV]	$\in (70, 100)$
$H_{3,1}^{\text{PP}}$ [GeV]	> 250
$\frac{p_{T,\text{PP}}^{\text{Lab}}}{p_{T,\text{PP}}^{\text{Lab}} + H_{3,1}^{\text{PP}}}$	< 0.2
$\frac{H_{T,3,1}^{\text{PP}}}{H_{3,1}^{\text{PP}}}$	> 0.75

Fig. 11.8 VR3 ℓ _VV

11.2.2 Compressed Tree Regions

VR3 ℓ _ISR_VV is defined in a similar way to CR3 ℓ _ISR_VV and SR3 ℓ _ISR, cutting on all of the same variables and using most of the same values for the cuts. In fact, apart from two cuts VR3 ℓ _ISR_VV is identical to CR3 ℓ _ISR_VV. Where CR3 ℓ _ISR_VV has an inverted M_T^W cut to make it orthogonal to SR ℓ _ISR, VR3 ℓ _ISR_VV doesn't invert the cut and has a relaxed requirement for it. This alone would mean that VR3 ℓ _ISR_VV wouldn't be orthogonal to either SR3 ℓ _ISR or CR3 ℓ _ISR_VV. To achieve orthogonality the $p_{T,CM}$ cut is inverted instead, making SR3 ℓ _ISR, CR3 ℓ _ISR_VV and VR3 ℓ _ISR_VV all mutually orthogonal. Table 11.6 contains the full definition of VR3 ℓ _ISR_VV while Figure 11.9 gives the kinematic variable distributions in that region.

Table 11.6 3-lepton validation region definitions - compressed tree

Validation Regions: 3-Lepton Compressed Tree	
Cut	VR3 ℓ _ISR_VV
Lepton Selection	$(Z^0 \rightarrow) \ell_{f_1}^+ \ell_{f_1}^- \ell_{f_2}^\pm$
n_{jets}	≥ 1
$n_{b\text{-tagged jets}}$	$= 0$
$p_T^{\ell_1}$ [GeV]	> 25
$p_T^{\ell_2}$ [GeV]	> 25
$p_T^{\ell_3}$ [GeV]	> 20
m_{ll} [GeV]	$\in (75, 105)$
m_T^W [Gev]	> 60
$\Delta\phi_{\text{ISR},I}^{\text{CM}}$	> 2.0
R_{ISR}	$\in (0.55, 1.0)$
$p_{T,\text{ISR}}^{\text{CM}}$ [GeV]	> 80
$p_{T,I}^{\text{CM}}$ [GeV]	> 60
$p_{T,CM}$ [GeV]	> 25




Fig. 11.9 VR3ℓ_ISR_VV

Finally, Table 11.7 shows the yields from the background-only fits for the VR3 ℓ _VV and VR3 ℓ _ISR_VV regions.

Table 11.7 Yields in the validation regions for the 3-lepton standard and compressed trees.

table.results.yields channel	VR3ℓ_VV	VR3ℓ_ISR_VV
Observed events	160	83
Fitted bkg events	159.20 ± 37.49	109.01 ± 23.90
Fitted HIGGS events	1.08 ± 0.12	0.27 ± 0.15
Fitted MM events	4.05 ± 1.22	4.20 ± 1.26
Fitted TTV events	0.44 ± 0.10	2.61 ± 0.44
Fitted VVV events	0.44 ± 0.03	0.42 ± 0.04
Fitted VV events	153.19 ± 37.48	101.51 ± 23.89
MC exp. SM events	146.23	97.31
MC exp. HIGGS events	1.08	0.27
MC exp. MM events	4.06	4.21
MC exp. TTV events	0.44	2.61
MC exp. VVV events	0.44	0.42
MC exp. VV events	140.21	89.80

Chapter 12

Uncertainties

All physics experiments have to contend with statistical and systematic uncertainties in their measurements, and this analysis is no different. Of these two categories, the systematic uncertainty contributes the greater share to the total uncertainty¹. The statistical uncertainty is just an uncertainty on the number of data events and is stored in the same data structure that the properties of an event are stored in. In contrast, there are many different sources of systematic uncertainty that must be summed together to obtain the total systematic uncertainty.

This chapter will document the sources of systematic uncertainty that were considered for this search, divided into several categories. Section 12.1 will list the experimental systematics uncertainties, which are associated with the ATLAS detector and its readout system. Each of the two special methods used to estimate the reducible and Z+Jets backgrounds described in Chapter 10 have their own associated systematic uncertainties. The systematics associated with the photon template method and matrix method are described in Sections 12.2 and 12.3, respectively. Finally, Sections 12.4 and 12.5 will summarise the systematic uncertainties in the 2-lepton and 3-lepton final states, respectively.

12.1 Experimental Systematic Uncertainties

The experimental systematic uncertainties considered by this search are as follows:

- Muon Momentum Resolution Uncertainty ($\text{MUONS_ID_}\{\text{UP}, \text{DN}\}$, $\text{MUON_MS_}\{\text{UP}, \text{DN}\}$): $\pm 1\sigma$ variations in the smearing of the ID/MS track

¹This is as expected, given that this search can be thought of as a simple counting experiment but the apparatus performing the experiment is incredibly complicated.

- Muon Momentum Scale Uncertainty ($\text{MUONS_SCALE}_{\{\text{UP},\text{DN}\}}$): $\pm 1\sigma$ variations in the scale of the muon momentum
- Muon Reconstruction Efficiency Uncertainty ($\text{MUON_EFF}_{\{\text{STAT},\text{STAT_LOWPT},\text{SYS},\text{SYS_LOWPT}\}}$): $\pm 1\sigma$ variations in the statistical/systematic error on the reconstruction scale factors
- Muon Isolation Efficiency Uncertainty ($\text{MUON_ISO}_{\{\text{STAT},\text{SYS}\}}$): $\pm 1\sigma$ variations in the statistical/systematic error on the muon isolation WP scale factors
- Electron Energy Scale Uncertainty ($\text{EG_SCALE_ALL}_{\{\text{UP},\text{DN}\}}$): $\pm 1\sigma$ variations in the scale of the electron momentum
- Electron Identification Efficiency Uncertainty (EL_EFF_ID): $\pm 1\sigma$ variations in the error on the electron identification WP scale factors
- Electron Reconstruction Efficiency Uncertainty (EL_EFF_Reco): $\pm 1\sigma$ variations in the error on the reconstruction scale factors
- Electron Isolation Efficiency Uncertainty (EL_EFF_Iso): $\pm 1\sigma$ variations in the error on the electron isolation WP scale factors
- Jet Energy Resolution Uncertainty (JER): $\pm 1\sigma$ variations in the smearing of jets (smear Monte Carlo jets to correct the Monte Carlo resolution)
- Jet Energy Scale Uncertainty ($\text{JET_GroupedNP_1}_{\{\text{UP},\text{DN}\}}$): $\pm 1\sigma$ variations in the scale of the jet momentum.
- Flavor-Tagging Efficiency Uncertainties ($\text{FT_EFF}_{\{\text{B},\text{C},\text{Light}\}}$, $\text{FT_EFF_extrapolation}\{\text{charm}\}$): $\pm 1\sigma$ variations in the error on the scale factor that corrects the tagging rate in simulation to match the one in data
- JVT Uncertainty (JET_JVTEff): $\pm 1\sigma$ variations which account for the residual contamination from pile-up jets after pile-up suppression and the MC generator choice
- E_T^{miss} Soft-Term Resolution Uncertainty ($\text{MET_SoftTrk_Reso}\{\text{Para},\text{Perp}\}$): $\pm 1\sigma$ variations in the E_T^{miss} resolution derived from comparisons of 2015 data to Monte Carlo, with generator (PP8 vs MG) and simulation (Fullsim vs AF2) comparisons

- Pile-up uncertainty (PILEUP): The uncertainties on the pile-up scaling are evaluated, changing the nominal re-scaling ($1/1.16$) to 1.00 and $1/1.23$. The interval covers the full difference between applying and not-applying the correction as well as uncertainty on the luminosity measurement which is expected to be the dominant uncertainty

12.2 Photon Template Systematics

The Z+Jets background estimate that comes from the photon template method has its own set of systematic uncertainties. These are summarised below:

- Reweighting Uncertainty: Nominal uses Boson p_T . Variation uses H_T .
- Smearing Uncertainty: Nominally derived from Monte Carlo. Variation derived from a special data control region.
- Monte Carlo-Closure Uncertainty: The difference between the Z+Jets and γ +Jets Monte Carlo in a special validation region.
- Data-Closure Uncertainty: The difference between the Z+Jets and γ +Jets Data in a special validation region.
- $V\gamma$ Contamination Uncertainty: The $V\gamma$ contribution to the γ +Jets events has a 100% uncertainty applied to it.
- Photon Event Statistics: The uncertainty is the raw number of photon events in the signal region.
- Side-Band Uncertainty: An uncertainty based on extrapolating the m_{jj} distribution into the signal regions to estimate the yields.
- Low- p_T Photon Event Number Upper Limit (SR2 ℓ _ISR Only): An uncertainty that comes from the method for estimating the contribution from events with low photon p_T (see Subsection 10.2.8).

Most of these sources of systematic uncertainty are inherent to the photon template method and are considered by all analyses that use it[154]. The last two contributions, the side-band and low photon transverse momentum, are unique to this search and the rest of this section will describe them in detail.

The photon template method has limited statistics when modelling the m_{jj} distributions in the signal regions. An uncertainty for this is calculated by taking the difference between the nominal event yield and the event yield derived from the side-band method. These yields take different values in each signal region, so the uncertainty must be calculated separately for each signal region. In practice, this is only done for

the SR2 ℓ _Low and SR2 ℓ _ISR regions. The side-band method works by defining a “side-band region” for each signal region. The event yields in these side-band regions are then used to calculate the event yields in the signal regions.

For SR2 ℓ _Low, the side-band region is defined using the same cuts as SR2 ℓ _Low (see Table 9.3) except for the cut on $\frac{p_{T,PP}^{\text{Lab}}}{p_{T,PP}^{\text{Lab}} + H_{4,1}^{\text{PP}}} \in (0.05, 0.15)$. The ratio $R = \frac{n_{\text{Window}}}{n_{\text{Spectrum}}}$ is calculated; n_{Window} is the number of events in the side-band region and n_{Spectrum} is the number of events in the side-band region with the cut on m_{jj} removed. The side-band estimate in SR2 ℓ _Low is thus the ratio R multiplied by the number of events in SR2 ℓ _Low with the m_{jj} cut removed. Figure 12.1 gives the “N-1” plots showing the m_{jj} variable in SR2 ℓ _Low and the side-band region, both with the m_{jj} cut removed.

For SR2 ℓ _ISR, the side-band region is defined by removing the $p_{T,ISR}^{\text{CM}}$ cut and changing the $p_{T,I}^{\text{CM}}$ cut to $p_{T,I}^{\text{CM}} \in (40, 80)$ GeV. The ratio $R = \frac{n_{\text{Window}}}{n_{\text{Spectrum}}}$ is calculated again; n_{Window} is the number of events in this side-band region and n_{Spectrum} is the number of events in the side-band region with the m_J cut removed. Figure 12.2 gives the “N-1” plots showing the m_J variable in SR2 ℓ _ISR and the side-band region, both with the m_J cut removed.




Fig. 12.1 “N-1” plots of the m_{jj} distributions in SR2 ℓ _Low and the side-band region ($\frac{p_{T,PP}^{\text{Lab}}}{p_{T,PP}^{\text{Lab}} + H_{4,1}^{\text{PP}}} \in (0.05, 0.15)$).

The photon template method has a limitation when dealing with γ +Jets events that have low photon transverse momentum (discussed in Subsection 10.2.8). This is accounted for by introducing an upper limit to the Z+Jets estimate as an additional systematic uncertainty. This uncertainty is calculated in a modified version of SR2 ℓ _ISR, which has the n_{jets} , m_J and $\Delta\phi_{\text{ISR},I}^{\text{CM}}$ cuts changed to $n_{\text{jets}} > 2$, $m_J < 300$ GeV and > 2.4 , respectively. The ratio $R_\gamma = \frac{n_{p_T^\gamma < 100 \text{ GeV}}}{n_{p_T^\gamma > 100 \text{ GeV}}}$ is calculated, where $n_{p_T^\gamma < 100 \text{ GeV}}$ is the




Fig. 12.2 “N-1” plots of the m_J distributions in SR2 ℓ _ISR and the side-band region ($p_{T,ISR}^{CM}$ cut removed and $p_{T,I}^{CM} \in (40, 80)$ GeV).

number of events in the region with $p_T^\gamma < 100$ GeV and $n_{p_T^\gamma > 100 \text{ GeV}}$ is the number of events with $p_T^\gamma > 100$ GeV. R_γ is calculated using the γ +Jets events that come from the photon template. The final uncertainty is then calculated by multiplying R_γ with the sum of the nominal Z+Jets estimate from the photon template and the statistical uncertainty (corresponding to a 1σ variation) associated with it.

Tables 12.1, 12.2 and 12.3 give the uncertainties from the sources described above in the SR2 ℓ _High/SR2 ℓ _Intermediate, SR2 ℓ _Low and SR2 ℓ _ISR regions, respectively.

Table 12.1 Photon template systematic uncertainties for the SR2 ℓ _Intermediate and SR2 ℓ _High regions. The relative uncertainties (reweighting, smearing, Monte Carlo and data-closure) are derived from validations regions. The absolute uncertainties ($V\gamma$ subtraction and side-band) are quoted in signal regions.

SR2 ℓ _Intermediate/ SR2 ℓ _High	Nominal Prediction	Variation Method	Relative Uncertainty (%)
Reweighting uncertainty	p_T -reweighting 51.61	H_T -reweighting 39.4	24
Smearing uncertainty	MC smearing function 51.6	Data smearing function 59.7	16
MC-closure uncertainty	γ +Jets MC prediction 35.3	Z+Jets Sherpa MC 52.2	48
Data-closure uncertainty	γ +Jets Data prediction 51.6	Dilepton Data/ $t\bar{t}$ / VV 56.4	9
$V\gamma$ uncertainty	$V\gamma$ subtracted 0.00 (SR2 ℓ _Intermediate) 0.07 (SR2 ℓ _High)	No subtraction 0.72 0.81	
Z+Jets prediction	Nominal 0.00 (SR2 ℓ _Intermediate) 0.07 (SR2 ℓ _High)	Statistical Uncertainty 0.18 0.20	Total Uncertainty 0.74 0.78

Table 12.2 Photon template systematic uncertainties for the SR2 ℓ _Low region. The relative uncertainties (reweighting, smearing, Monte Carlo- and data-closure) are derived from validations regions. The absolute uncertainties ($V\gamma$ subtraction and side-band) are quoted in signal regions.

SR2 ℓ _Low	Nominal Prediction	Variation Method	Uncertainty
Reweighting uncertainty	p_T -reweighting 210.4	H_T -reweighting 145.0	1.96
Smearing uncertainty	MC smearing function 210.4	Data smearing function 190.5	0.57
MC-closure uncertainty	γ +Jets MC prediction 438.9	Z+Jets Sherpa MC 696.5	3.73
Data-closure uncertainty	γ +Jets Data prediction 210.4	Dilepton Data/ $t\bar{t}/VV$ 213.6	0.06
$V\gamma$ uncertainty	$V\gamma$ subtracted 6.32	No subtraction 5.00	1.32
Side-band uncertainty (γ +Jets method, Sherpa 2.2.1)	Nominal 6.32	Side-band extrapolation 5.85	0.47
Z+Jets prediction	Nominal 6.32 (NF= 1.0)	Statistical Uncertainty 3.65	Total Uncertainty 5.77

Table 12.3 Photon template systematic uncertainties for the SR2 ℓ _ISR region. The relative uncertainties (reweighting, smearing, Monte Carlo- and data-closure) are derived from validations regions. The absolute uncertainties ($V\gamma$ subtraction, side-band and photon momentum upper limit) are quoted in signal regions. An additional normalisation factor is derived from the validation region with one cut relaxed ($p_{T,I}^{\text{CM}} > 90$ GeV for SR2 ℓ _ISR), to correct the total Z+Jets prediction.

SR2 ℓ _ISR	Nominal Prediction	Variation Method	Uncertainty
Reweighting uncertainty	p_T -reweighting 198.5	H_T -reweighting 206.6	0.004
Smearing uncertainty	MC smearing function 198.5	Data smearing function 193.2	0.003
MC-closure uncertainty	γ +Jets MC prediction 130.7	Z+Jets Sherpa MC 195.1	0.049
Data-closure uncertainty	γ +Jets data prediction 198.5	Dilepton Data/ $t\bar{t}/VV$ 132.5	0.033
$V\gamma$ uncertainty	$V\gamma$ subtracted 0.10	No subtraction 0.69	0.59
Side-band uncertainty (γ +Jets method, Sherpa 2.2.1)	Nominal 0.10	Side-band extrapolation 0.17	0.07
Upper limit	—	—	2.46
Z+Jets prediction	Nominal 0.10 (NF= 1.06)	Statistical uncertainty 0.50	Total uncertainty 2.58

12.3 Fake Estimation Systematics

Like the photon template method, the matrix method used to estimate the fake contribution has its own set of systematic uncertainties. The sources of these systematics are described in this section.

The limited statistics in the control regions where the fake rate and real efficiencies are calculated is the first source of systematic uncertainty. The fake rate and real efficiency are varied by $\pm 1\sigma$ around the nominal value, from which two estimates are generated: an “upper limit” from both the fake rate and real efficiency having 1σ added to them and a “lower limit” from the fake rate and real efficiency having 1σ subtracted from them. The difference between these limits and the nominal value then define the asymmetric upper and lower bounds of the uncertainty.

The weight of each fake component also has an uncertainty associated with it. The effect of this uncertainty on the fake estimate is found using the method described in the previous paragraph.

The final set of uncertainties are on N_{TT} , N_{TL} , N_{LT} and N_{LL} , the numbers of events used in the matrix inversion described in Section 10.1. The total uncertainty here is calculated by taking the sum of the squares of the weights in each region.

The above three sources of systematic uncertainty are then summed in quadrature and given as the total systematic uncertainty on the fake estimate.

12.4 Systematic Uncertainties: 2-Lepton, 2-Jet and E_T^{miss} Final States

The various systematic uncertainties have different effects on the different backgrounds and the data. This means that the different final states also have different sources of systematic uncertainty contributing to their total uncertainty. This section will summarise the systematics in the 2-lepton final state and the next will summarise the systematics in the 3-lepton final state.

The jet energy scale (“JES”) and jet energy resolution (“JER”) uncertainties are derived as functions of jet p_T and η , as well as of the pile-up conditions and the jet flavour composition of each jet sample. These uncertainties are determined, using a combination of simulated events and data samples, through measurements of the jet response balance in multi-jet, Z+Jets and γ +Jets events [155].

The systematic uncertainties related to the Monte Carlo modelling of E_T^{miss} are estimated by propagating the uncertainties in the energy and momentum scale of the physics objects through the E_T^{miss} calculation. Additional uncertainties related to the E_T^{miss} modelling come directly from the soft-term scale and resolution [156].

The detector related systematics associated with the 2-lepton final state such as: the lepton reconstruction efficiency, b -tagging efficiency [157, 158], lepton energy scale, energy resolution and the modelling of the trigger [75, 159] are all included but are found to be negligible.

The systematic uncertainties associated with the VV Monte Carlo are estimated by varying the renormalisation, factorisation, merging scales and PDFs used to generate the samples.

Finally, both the photon template and matrix method systematic uncertainties are considered in the 2-lepton final state. Table 12.4 lists the largest sources of systematic uncertainties affecting the 2-lepton final state as percentages.

Table 12.4 Summary of the main systematic uncertainties and their impact (in %) on the total SM background prediction in each of the 2ℓ SRs. The total systematic uncertainty can be different from the sum in quadrature of individual sources due to the correlations between them resulting from the fit to the data.

Signal Region	SR2 ℓ _High	SR2 ℓ _Int	SR2 ℓ _Low	SR2 ℓ _ISR
Total uncertainty [%]	42	38	70	103
Z+jets data-driven estimate	42	31	69	96
VV theoretical uncertainties	28	27	6	34
MC statistical uncertainties	16	12	5	9
VV fitted normalization	13	14	2	16
FNP leptons	-	5	13	12
Jet energy resolution	5	10	4	3
Jet energy scale	1	2	< 1	3
E_T^{miss} modelling	3	4	< 1	< 1
$t\bar{t}$ fitted normalization	< 1	< 1	2	2
Lepton reconstruction / identification	< 1	< 1	< 1	< 1

12.5 Systematic Uncertainties: 3-Lepton and E_T^{Miss} Final States

The 3-lepton final state is affected by roughly the same systematic uncertainties as the 2-lepton final state. The main difference is that, since the Z+Jets background contribution

Table 12.5 Summary of the main systematic uncertainties and their impact (in %) on the total SM background prediction in each of the 3ℓ SRs. The total systematic uncertainty can be different from the sum in quadrature of individual sources due to the correlations between them resulting from the fit to the data.

Signal Region	SR3 ℓ _High	SR3 ℓ _Int	SR3 ℓ _Low	SR3 ℓ _ISR
Total uncertainty [%]	44	22	19	26
VV theoretical uncertainties	18	9	12	19
MC statistical uncertainties	37	17	8	10
VV fitted normalisation	8	7	9	11
FNP leptons	7	< 1	3	5
Jet energy resolution	4	< 1	7	3
Jet energy scale	7	< 1	2	3
E_T^{miss} modelling	2	< 1	1	4
Lepton reconstruction / identification	3	4	2	2

is negligible in the 3-lepton signal regions, the systematic uncertainty associated with the photon template method needn't be considered. Table 12.5 summarises the largest sources of systematic uncertainties affecting the 3-lepton final state as percentages.

Finally, a more detailed summary of the systematic uncertainties is given in Appendix E.

Chapter 13

Results

Chapter 9 defined the signal regions used by this search for new physics. Chapter 10 defined the control regions and gave the background yields in each of these regions that came from the background-only fits. Chapter 11 defined the validation regions used to validate the background yields from the background-only fits. Finally, an account of all the relevant uncertainties is given in Chapter 12.

This chapter will document the results of the analysis presented in this thesis: the process of unblinding and the performance of the model-independent fits and the model-dependent fits.

The model-independent fits are performed first, with the observed number of data events being compared to the expected background contribution in each signal region. The 2-lepton channel comparisons are given in Table 13.1 and Figure 13.1. Table 13.1 gives the observed data yields, the estimated background yields (both the total background estimate and the estimate from each contribution) and the Single-top and VV fit parameters. Figure 13.1 observed data yields and background estimates in all of the signal, control and validation regions. The 3-lepton channel comparisons are given in Table 13.2 and Figure 13.2.

For both the high- and intermediate-mass regions in both the 2-lepton and 3-lepton channels, no significant excesses above the expected Standard Model yields are observed. An excess of events *is* observed in the four low-mass and ISR regions. The model-independent fit is used to quantify the level of agreement between the observed data and the background expectation, with Table 13.3 giving the relevant parameters from the results of that fit. Further, kinematic distributions of selected variables in the low-mass and ISR regions are given in Figures 13.3 and 13.4.

Table 13.1 Expected and observed yields from the background-only fit for the 2ℓ SRs. The errors shown are the statistical plus systematic uncertainties. Uncertainties in the predicted background event yields are quoted as symmetric, except where the negative error reaches down to zero predicted events, in which case the negative error is truncated.

Signal Region	SR2 ℓ _High	SR2 ℓ _Int	SR2 ℓ _Low	SR2 ℓ _ISR
Total Observed Events	0	1	19	11
Total Background Events	1.9 ± 0.8	2.4 ± 0.9	8.4 ± 5.8	$2.7^{+2.8}_{-2.7}$
Z+Jets	$0.07^{+0.78}_{-0.07}$	$0.00^{+0.74}_{-0.00}$	6.3 ± 5.8	$0.10^{+2.58}_{-0.10}$
Fit Output: VV	1.8 ± 0.7	2.4 ± 0.8	1.5 ± 0.9	2.3 ± 1.1
Fit Output: $Wt + t\bar{t}$	0.00 ± 0.00	0.00 ± 0.00	0.57 ± 0.20	$0.28^{+0.34}_{-0.28}$
Other	0.02 ± 0.01	$0.05^{+0.12}_{-0.05}$	$0.02^{+1.07}_{-0.02}$	$0.06^{+0.33}_{-0.06}$
Fit Input: VV	1.9	2.6	1.6	2.4
Fit Input: $Wt + t\bar{t}$	0.00	0.00	0.63	0.28




Fig. 13.1 The observed and expected SM background yields in the CRs, VRs and SRs considered in the 2ℓ channel. The statistical uncertainties in the background prediction are included in the uncertainty band, as well as the experimental and theoretical uncertainties. The bottom panel shows the difference in standard deviations between the observed and expected yields.

Table 13.2 Expected and observed yields from the background-only fit for the 3ℓ SRs. The errors shown are the statistical plus systematic uncertainties. Uncertainties in the predicted background event yields are quoted as symmetric, except where the negative error reaches down to zero predicted events, in which case the negative error is truncated.

Signal Region	SR3 ℓ _High	SR3 ℓ _Int	SR3 ℓ _Low	SR3 ℓ _ISR
Total Observed Events	2	1	20	12
Total Background Events	1.1 ± 0.5	2.3 ± 0.5	10.0 ± 2.0	3.9 ± 1.0
Fit Output: VV	0.83 ± 0.39	1.9 ± 0.5	10.0 ± 2.0	3.8 ± 1.0
Triboson	0.19 ± 0.07	0.32 ± 0.06	0.25 ± 0.03	0.08 ± 0.04
Other	$0.03^{+0.07}_{-0.03}$	0.04 ± 0.02	$0.02^{+0.34}_{-0.02}$	$0.06^{+0.19}_{-0.06}$
Fit Input: VV	0.76	1.8	9.2	3.4




Fig. 13.2 The observed and expected SM background yields in the CRs, VRs and SRs considered in the 3ℓ channel. The statistical uncertainties in the background prediction are included in the uncertainty band, as well as the experimental and theoretical uncertainties. The bottom panel shows the difference in standard deviations between the observed and expected yields.

Table 13.3 Model-independent fit results for all signal regions. The first column shows the signal region, the second and third columns show the 95% CL upper limits on the visible cross-section ($\langle\epsilon\sigma\rangle_{\text{obs}}^{95}$) and on the number of signal events (S_{obs}^{95}). The fourth column (S_{exp}^{95}) shows the 95% CL upper limit on the number of signal events, given the expected number (and $\pm 1\sigma$ excursions of the expectation) of background events. The last column indicates the discovery p_0 -value and its associated significance (Z).

Signal Region	$\langle\epsilon\sigma\rangle_{\text{obs}}^{95}[\text{fb}]$	S_{obs}^{95}	S_{exp}^{95}	p_0 (Z)
SR3 ℓ _ISR	0.42	15.3	$6.9^{+3.1}_{-2.2}$	0.001 (3.02)
SR2 ℓ _ISR	0.43	15.4	$9.7^{+3.6}_{-2.5}$	0.02 (1.99)
SR3 ℓ _Low	0.53	19.1	$9.5^{+4.2}_{-1.8}$	0.016 (2.13)
SR2 ℓ _Low	0.66	23.7	$16.1^{+6.3}_{-4.3}$	0.08 (1.39)
SR3 ℓ _Int	0.09	3.3	$4.4^{+2.5}_{-1.5}$	0.50 (0.00)
SR2 ℓ _Int	0.09	3.3	$4.6^{+2.6}_{-1.5}$	0.50 (0.00)
SR3 ℓ _High	0.14	5.0	$3.9^{+2.2}_{-1.3}$	0.23 (0.73)
SR2 ℓ _High	0.09	3.2	$4.0^{+2.3}_{-1.2}$	0.50 (0.00)

Given that the SR3 ℓ _Low and SR3 ℓ _ISR are statistically independent, it is possible to find whether a given event falls into one region or the other. To this end, Figure 13.5 shows the “N-1” plots of m_T^W in both SR3 ℓ _Low and SR3 ℓ _ISR. This variable is chosen because these two regions make use of two different trees and there aren’t many variables common to both. Figure 13.5a shows the “N-1” plot of m_T^W in SR3 ℓ _Low and Figure 13.5b shows the “N-1” plot of m_T^W in SR3 ℓ _ISR.

Table 13.4 shows the lepton flavour composition in the SR2 ℓ _Low, SR2 ℓ _ISR, SR3 ℓ _Low and SR3 ℓ _ISR regions. The composition of the observed data and the background expectation are both given in Table 13.4, with the background expectation included in parentheses with its uncertainty.

Table 13.4 Breakdown of the observed and expected (in parentheses) number of events in terms of flavor composition in the SRs with an excess.

Signal Region	SR2 ℓ _Low	SR2 ℓ _ISR
ee	9 (4.5 $\pm$ 3.9)	3 (1.2 $\pm$ 1.2)
$\mu\mu$	10 (3.9 $\pm$ 2.6)	8 (1.5 $\pm$ 1.5)
Signal Region	SR3 ℓ _Low	SR3 ℓ _ISR
eee	6 (3.5 $\pm$ 0.7)	3 (1.1 $\pm$ 0.3)
$ee\mu$	6 (2.0 $\pm$ 0.4)	3 (0.9 $\pm$ 0.3)
$\mu\mu\mu$	7 (2.7 $\pm$ 0.6)	4 (1.5 $\pm$ 0.4)
$\mu\mu e$	1 (1.9 $\pm$ 0.4)	2 (0.4 $\pm$ 0.1)




Fig. 13.3 Distributions of selected kinematic variables in the signal regions of the 2ℓ final state. The histograms show the post-fit background predictions. The right-most bin in each histogram includes the overflow. For SR2 ℓ _Low, the distributions for $H_{4,1}^{PP}$ (a) and $\frac{\min(H_{1,1}^{Pa}, H_{1,1}^{Pb})}{\min(H_{2,1}^{Pa}, H_{2,1}^{Pb})}$ (b) are given. For SR2 ℓ _ISR, the distributions for $p_{T,ISR}^{CM}$ (c) and R_{ISR} (d) are given. The black, hatched error bars indicate the combined uncertainty: Monte Carlo statistical, theoretical systematic and experimental systematic. The expected distribution for a benchmark signal model, normalised to the NLO+NNL cross-section times integrated luminosity, is also shown for comparison.




Fig. 13.4 Distributions of selected kinematic variables in the signal regions of the 3ℓ final state. The histograms show the post-fit background predictions. The right-most bin in each histogram includes the overflow. For SR3 ℓ _Low, the distributions for $p_T^{\ell_1}$ (a) and $H_{3,1}^{PP}$ (b) are given. For SR3 ℓ _ISR, the distributions for $p_{T,ISR}^{CM}$ (c) and R_{ISR} (d) are given. The black, hatched error bars indicate the combined uncertainty: Monte Carlo statistical, theoretical systematic and experimental systematic. The expected distribution for a benchmark signal model, normalised to the NLO+NLL cross-section times integrated luminosity, is also shown for comparison.




Fig. 13.5 The “N-1” m_T^W distributions for SR3 ℓ _Low (a) and SR3 ℓ _ISR. The solid red line and arrow shows the cut on m_T^W in each of these regions (the arrow points in the direction of what is *kept* by the cut). The right-most bin in each histogram includes the overflow. The black, hatched error bars indicate the combined uncertainty: Monte Carlo statistical uncertainty, theoretical systematic uncertainties on the VV background and the experimental systematic uncertainty. The expected distribution for a benchmark signal model, normalised to the NLO+NLL cross-section times integrated luminosity, is also shown for comparison.

The results of the model-independent fits are summarised in Table 13.3 and are used to inform the model-dependent fits. The model-dependent fits are used to generate the exclusion contours discussed in Chapter 6 and are shown in Figure 13.6. Figures 13.6a and 13.6b show the exclusion limits obtained for the 2 ℓ and 3 ℓ channels, respectively, where each signal model is tested in the signal region with the highest sensitivity to it. In both cases, the low-mass and ISR regions are statistically combined. Figure 13.6c shows the combined exclusion contour for the statistical combination of the 2 ℓ and 3 ℓ channels. This combination is done by statistically combining the signal regions in each channel that target the same region of the mass-plane (ie, SR2 ℓ _High is combined with SR3 ℓ _High, etc). After the combination, the combined 2 ℓ /3 ℓ signal region that gives the best expected CL_s value for each signal model assumption is chosen. The low-mass and ISR regions are once again statistically combined. Finally, Figure 13.6d compares the expected and observed exclusion limits obtained from this analysis with the results of a recent, similar analysis that makes use of the same data set but uses more conventional variables [87].

These results extend the exclusion limits in the high- and intermediate-mass regions of the mass-plane compared to Ref. [87]. The results also show that the low-mass region (where the mass-splitting $m_{\tilde{\chi}_1^\pm/\tilde{\chi}_2^0} - m_{\tilde{\chi}_1^0}$ is only around 100 GeV) can't be excluded due to the observed excess of events. This result is of special interest because it is in disagreement with the results of Ref. [87]. The observed data excesses in SR2 ℓ _Low, SR2 ℓ _ISR, SR3 ℓ _Low and SR3 ℓ _ISR have associated significances of 1.4, 2.0, 2.1 and 3.0 standard deviations, respectively. These deviations are the reason why this analysis isn't able to exclude the areas of the mass-plane where exclusion was expected to occur, since for exclusion to occur necessitates that no excess is found. The analysis given in Ref. [87] excludes this region while this analysis can't achieve this with the data it uses.

13.1 Conclusion

This thesis has presented a search for the electroweak production of charginos and neutralinos decaying into final states containing either two or three leptons, jets and missing energy. This search was performed using proton-proton collision data at $\sqrt{s} = 13$ TeV corresponding to an integrated luminosity of 36.1 fb⁻¹ recorded by the ATLAS detector at the LHC. The search targets the 2 ℓ and 3 ℓ channels using slightly different signal models for each (the 2 ℓ channel has the $W^\pm$ boson decay into a quark-antiquark pair while the 3 ℓ channel has it decay into a lepton and neutrino) and using the Recursive Jigsaw Reconstruction method to select signal events.

The statistical interpretation of the two search channels attempts to place new exclusion limits on associated $\tilde{\chi}_1^\pm\tilde{\chi}_2^0$ production with gauge boson mediated decays. For a $\tilde{\chi}_1^0$ with a mass of less than 100 GeV, $\tilde{\chi}_1^\pm/\tilde{\chi}_2^0$ masses of up to 600 GeV have been excluded. This result extends the exclusion limit in the supersymmetric parameter space from previous LHC searches in the high- and intermediate-mass regions. In the low-mass and compressed regions an excess of events above the Standard Model expectation is found and the region of parameter space below $m_{\tilde{\chi}_1^\pm/\tilde{\chi}_2^0} = 220$ GeV cannot be excluded. The excesses observed in the 2 ℓ channel regions SR2 ℓ _Low and SR2 ℓ _ISR correspond to local significances of 1.4σ and 2.0σ , respectively. The excesses observed in the 3 ℓ channel regions SR3 ℓ _Low and SR3 ℓ _ISR correspond to local significances of 2.1σ and 3.0σ , respectively.




Fig. 13.6 Exclusion limits at the 95% CL on the masses of the $\tilde{\chi}_0^\pm/\tilde{\chi}_2^0$ and $\tilde{\chi}_1^0$ from the analysis of 36.1 fb⁻¹ of 13 TeV proton-proton collision data obtained by the ATLAS detector. The contours given in (a) and (b) correspond to the 2 ℓ and 3 ℓ channels, respectively, while the contour given in (c) shows a statistical combination of the 2 ℓ and 3 ℓ channels (assuming a 100% branching ratio of the sparticles decaying into the Standard Model $W^\pm/Z^0$ bosons and the $\tilde{\chi}_1^0$). The dashed line and the shaded band are the expected limit and its $\pm 1\sigma$ uncertainty, respectively. The thick solid line is the observed limit for the central value of the signal cross-section. The dotted lines around the observed limit illustrate the change in the observed limit as the nominal signal cross-section is scaled up and down by the theoretical uncertainty. Finally, (d) compares the exclusion limits from the analysis presented in this thesis and [87].

References

- [1] Michael E. Peskin and Daniel V. Schroeder. Reading, USA: Addison-Wesley, 1995.
- [2] Paul Adrien Maurice Dirac. “The quantum theory of the emission and absorption of radiation”. In: *Proceedings of the Royal Society of London* 114 (1927), pp. 243–265.
- [3] Susanne Mertens. “Direct Neutrino Mass Experiments”. In: *Journal of Physics: Conference Series* 718 (May 2016).
- [4] Louis de Broglie. “Recherches sur la théorie des Quanta”. In: *Annales de Physique* 10 (3 1925), pp. 22–128.
- [5] Enrico Fermi. “Versuch einer theorie der β -strahlen”. In: *Zeitschrift für Physik* 88 (1934), pp. 161–171.
- [6] Murray Gell-Mann. “The Eightfold Way: A Theory of Strong Interaction Symmetry”. 1961.
- [7] Murray Gell-Mann. “A schematic model of baryons and mesons”. In: *Physics Letters* 8 (1964), pp. 214–215.
- [8] George Zweig. *An $SU(3)$ Model for Strong Interaction Symmetry and its Breaking, Version 1*. 1964.
- [9] George Zweig. *An $SU(3)$ Model for Strong Interaction Symmetry and its Breaking, Version 2*. Ed. by D. B. Lichtenberg and Simon Peter Rosen. Vol. 1. 1964.
- [10] Oscar Wallace Greenberg. “Spin and Unitary Spin Independence in a Paraquark Model of Baryons and Mesons”. In: *Physical Review Letters* 13 (1964), pp. 598–602.
- [11] Moo-Young Han and Yoichiro Nambu. “Three Triplet Model with Double $SU(3)$ Symmetry”. In: *Physical Review* 139 (1965).
- [12] Sheldon Lee Glashow. “The Renormalizability of Vector Meson Interactions”. In: *Nuclear Physics* 10 (1959), pp. 107–117.
- [13] Abdus Salam and John Clive Ward. “Weak and Electromagnetic Interactions”. In: *Il Nuovo Cimento (1955-1965)* 11.4 (Feb. 1, 1959), pp. 568–577.
- [14] Steven Weinberg. “A Model of Leptons”. In: *Physical Review Letters* 19.21 (Nov. 20, 1967), pp. 1264–1266.
- [15] NobelPrize.org. *The Nobel Prize in Physics 1979*. URL: <https://www.nobelprize.org/prizes/physics/1979/summary/> (visited on 11/26/2019).
- [16] Masaharu Tanabashi et al. “W Mass”. In: *Physical Review D* 98 (2018).

- [17] Masaharu Tanabashi et al. “Z Mass”. In: *Physical Review D* 98 (2018).
- [18] Masaharu Tanabashi and other. “Leptons”. In: *Physical Review D* 98 (2018).
- [19] Masaharu Tanabashi and other. “Quarks”. In: *Physical Review D* 98 (2018).
- [20] François Englert and Robert Brout. “Broken Symmetry and the Mass of Gauge Vector Mesons”. In: *Physical Review Letters* 13 (1964), pp. 321–323.
- [21] Peter Higgs. “Broken Symmetries and the Masses of Gauge Bosons”. In: *Physical Review Letters* 13 (1964), pp. 508–509.
- [22] Gerald Guralnik, Carl Richard Hagen, and Tom Kibble. “Global Conservation Laws and Massless Particles”. In: *Physical Review Letters* 13 (1964), pp. 585–587.
- [23] Stephen P. Martin. *A Supersymmetry Primer*. Jan. 2016.
- [24] Ian J. R. Aitchison. *Supersymmetry and the MSSM: An Elementary Introduction*. 2005.
- [25] ATLAS Collaboration. “Observation of a new particle in the search for the Standard Model Higgs Boson with the ATLAS Detector at the LHC”. In: *Physics Letters B* 716 (Sept. 2012), pp. 1–29.
- [26] CMS Collaboration. “Observation of a new boson at a mass of 125 GeV with the CMS Experiment at the LHC”. In: *Physics Letters B* 716 (Sept. 2012), pp. 30–61.
- [27] CERN Organisation. *Dark Matter*. URL: <https://home.cern/science/physics/dark-matter> (visited on 01/03/2020).
- [28] R. J. Adler, B. Casey, and O. C. Jacob. “Vacuum catastrophe: An Elementary Exposition of the Cosmological Constant Problem”. In: *American Journal of Physics* 63 (1995), pp. 620–625.
- [29] CERN Organisation. *The 2018 data-taking run at the LHC has begun*. URL: <https://home.cern/news/news/accelerators/2018-data-taking-run-lhc-has-begun-0> (visited on 01/04/2020).
- [30] CERN Organisation. *CERN Design Guidelines - Downloads*. URL: <https://design-guidelines.web.cern.ch/downloads> (visited on 02/02/2020).
- [31] CERN Organisation. *The European Organization for Nuclear Research is born*. URL: <https://timeline.web.cern.ch/european-organization-nuclear-research-born> (visited on 02/02/2020).
- [32] CERN Organisation. *Aerial Views*. URL: <https://cern60.web.cern.ch/en/exhibitions/aerial-views> (visited on 09/22/2019).
- [33] Patrice Loiez and Maximilien Brice. *Aerial view of CERN’s Meyrin Site*. URL: <https://cds.cern.ch/record/872960> (visited on 03/22/2019).
- [34] CERN Organisation. *The Birth of the Web*. URL: <https://home.cern/science/computing/birth-web> (visited on 02/02/2020).
- [35] United Nations. *Intergovernmental and Other Organizations*. URL: <https://www.un.org/en/sections/member-states/intergovernmental-and-other-organizations/index.html> (visited on 03/22/2019).

- [36] CERN Organisation. *The Large Hadron Collider*. URL: <https://home.cern/science/accelerators/large-hadron-collider> (visited on 05/09/2019).
- [37] *LHC Design Report*. Vol. 1. June 2004.
- [38] CERN Organisation. *A Large Hadron Collider in the LEP Tunnel?* URL: <https://timeline.web.cern.ch/large-hadron-collider-lep-tunnel> (visited on 05/10/2019).
- [39] CERN Organisation. *LEP's Final Shutdown*. URL: <https://timeline.web.cern.ch/leps-final-shutdown> (visited on 04/30/2019).
- [40] CERN Organisation. *Linear Accelerator 2*. URL: <https://home.cern/science/accelerators/linear-accelerator-2> (visited on 05/31/2019).
- [41] CERN Organisation. *The Proton Synchrotron Booster*. URL: <https://home.cern/science/accelerators/proton-synchrotron-booster> (visited on 05/31/2019).
- [42] CERN Organisation. *The Proton Synchrotron*. URL: <https://home.cern/science/accelerators/proton-synchrotron> (visited on 05/31/2019).
- [43] CERN Organisation. *The Super Proton Synchrotron*. URL: <https://home.cern/science/accelerators/super-proton-synchrotron> (visited on 05/31/2019).
- [44] CERN Organisation. *TE-EPC-LPC in LHC*. URL: <https://te-epc-lpc.web.cern.ch/te-epc-lpc/machines/lhc/general.stm> (visited on 10/27/2020).
- [45] David T. Chuss. *Cosmic Background Explorer*. National Aeronautics and Space Administration. URL: <https://lambda.gsfc.nasa.gov/product/cobe> (visited on 05/17/2019).
- [46] Ernest D. Courant, M. Stanley Livingston, and Hartland S. Snyder. "The Strong-Focusing Synchrotron - A New High Energy Accelerator". In: *Physical Review* 88 (Dec. 1952), pp. 1190–1196.
- [47] Verena Kain. "Beam Dynamics and Beam Losses - Circular Machines". In: (Aug. 2016). DOI: 10.5170/CERN-2016-002.21.
- [48] Xabier Cid Vidal and Ramon Cid Manzano. *Magnetic Multipoles*. URL: https://www.lhc-closer.es/taking_a_closer_look_at_lhc/0.magnetic_multipoles (visited on 10/27/2020).
- [49] CERN Organisation. *ALICE*. URL: <https://home.cern/science/experiments/alice> (visited on 06/16/2019).
- [50] CERN Organisation. *LHCb*. URL: <https://home.cern/science/experiments/lhcb> (visited on 06/16/2019).
- [51] ATLAS Collaboration. *Standard Model Summary Plots Spring 2020*. URL: <https://atlas.web.cern.ch/Atlas/GROUPS/PHYSICS/PUBNOTES/ATL-PHYS-PUB-2020-010/> (visited on 10/31/2020).
- [52] CERN Organisation. *Interim Summary Report on the Analysis of the 19 September 2008 Incident at the LHC*. Oct. 2008.
- [53] CERN Organisation. *Follow up of the Incident of 19 September 2008 at the LHC*. Dec. 2008.
- [54] CERN Organisation. *Final Magnet goes Underground after LHC Repair*. URL: <https://timeline.web.cern.ch/final-magnet-goes-underground-after-lhc-repair> (visited on 06/16/2019).

- [55] CERN Organisation. *LHC Sets New World Record*. URL: <https://home.cern/news/press-release/cern/lhc-sets-new-world-record> (visited on 06/16/2019).
- [56] CERN Organisation. *LHC Sets World Record Beam Intensity*. URL: <https://home.cern/news/press-release/cern/lhc-sets-world-record-beam-intensity> (visited on 06/16/2019).
- [57] ATLAS Collaboration. *Detector and Technology | ATLAS Experiment at CERN*. URL: <https://atlas.cern/discover/detector> (visited on 09/26/2019).
- [58] Ovsat Abidinov et al. “ATLAS Detector and Physics Performance - Technical Design Report - Volume 2”. In: (2008).
- [59] Georges Aad et al. “The ATLAS Experiment at the CERN Large Hadron Collider”. In: *Journal of Instrumentation* 3 (Aug. 2008).
- [60] ATLAS Collaboration. *Event Cross Section in a computer generated image of the ATLAS detector*. URL: <https://cds.cern.ch/record/1096081> (visited on 03/22/2019).
- [61] Giordon Holtsberg Stark. “The Search for Supersymmetry in Hadronic Final States using Boosted Object Reconstruction”. 2018.
- [62] Bruce Joseph Gallop. “ATLAS SCT Readout and Barrel Macro-Assembly Testing and Design of MAPS Test Structures”. 2005.
- [63] D. Robinson et al. “Silicon Micro-Strip Detectors for the ATLAS SCT”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 485 (June 2002), pp. 84–88.
- [64] Alan J. Barr. *Calibrating the ATLAS Semiconductor Tracker Front End Electronics*. Tech. rep. ATL-INDET-CONF-2006-001. ATL-COM-INDET-2006-001. CERN-ATL-COM-INDET-2006-001. Geneva: CERN, Oct. 2004. URL: <http://cds.cern.ch/record/926377> (visited on 11/12/2020).
- [65] CERN Organisation. *Déplacement du sub-detector Silicon Tracker de l’expérience ATLAS dans l’axe du Barrel Silicon Tracker*. Sept. 2005. URL: <http://cdsweb.cern.ch/search?recid=883303> (visited on 11/12/2020).
- [66] CERN Organisation. *A disk for the ATLAS SCT end-cap in its testbox at NIKHEF*. Dec. 2005. URL: <https://cds.cern.ch/record/968738> (visited on 11/12/2020).
- [67] ATLAS Collaboration. *LuminosityPublicResultsRun2*. URL: <https://twiki.cern.ch/twiki/bin/view/AtlasPublic/LuminosityPublicResultsRun2> (visited on 02/05/2020).
- [68] Wikipedia Foundation. *CD-ROM*. URL: <https://en.wikipedia.org/wiki/CD-ROM> (visited on 02/12/2020).
- [69] ATLAS Collaboration. “Performance of the ATLAS Trigger System in 2015”. In: *The European Physical Journal C* 77 (May 2017).
- [70] ATLAS Collaboration. “The ATLAS Central level-1 Trigger Logic and TTC System”. In: *Journal of Instrumentation* (Aug. 2008).
- [71] ATLAS Collaboration. “Performance of the ATLAS Trigger System in 2010”. In: *The European Physical Journal C* 72 (Jan. 2012).

- [72] ATLAS Collaboration. “The Run 2 ATLAS Trigger System”. In: (Feb. 2016).
- [73] Morad Aaboud et al. “Electron reconstruction and identification in the ATLAS Experiment using 2015 and 2016 LHC proton-proton collision data at $\sqrt{s} = 13$ TeV”. In: *European Physical Journal C* (2019). eprint: 1902.04655.
- [74] Morad Aaboud et al. “Electron and photon energy calibration with the ATLAS detector using LHC Run 1 data”. In: *European Physics Journal C* 74 (2014).
- [75] Georges Aad et al. “Muon reconstruction performance of the ATLAS Detector in proton-proton collision data at $\sqrt{s} = 13$ TeV”. In: *European Physics Journal C* 76 (2016), p. 292. eprint: 1603.05598.
- [76] Henning Kirschenmann. *Jets at CMS and the determination of their energy scale*. July 9, 2012. URL: <http://cms.web.cern.ch/news/jets-cms-and-determination-their-energy-scale> (visited on 10/17/2018).
- [77] Andrew J. Larkoski, Ian Moutl, and Benjamin Nachman. “Jet substructure at the Large Hadron Collider: A review of recent advances in theory and machine learning”. In: *Physics Reports* 841 (Jan. 2020), pp. 1–63. ISSN: 0370-1573. DOI: 10.1016/j.physrep.2019.11.001. URL: <http://dx.doi.org/10.1016/j.physrep.2019.11.001> (visited on 11/17/2020).
- [78] *Generalized Numerical Inversion: A Neural Network Approach to Jet Calibration*. Tech. rep. ATL-PHYS-PUB-2018-013. Geneva: CERN, July 2018. URL: <https://cds.cern.ch/record/2630972> (visited on 11/17/2020).
- [79] Gavin P. Salam. “Towards Jetography”. In: *European Physics Journal C* 67 (2010), pp. 637–686. eprint: 0906.1833.
- [80] Matteo Cacciari, Gavin P. Salam, and Gregory Soyez. “The anti- k_t jet clustering algorithm”. In: *Journal of High Energy Physics* 04 (2008), p. 063. eprint: 0802.1189.
- [81] Arnulf Quadt. *Top quark physics at Hadron Colliders*. Vol. 48. 2006.
- [82] “Expected performance of the ATLAS b -tagging algorithms in Run 2”. In: (July 2015). URL: <https://cds.cern.ch/record/2037697> (visited on 12/14/2018).
- [83] Johan Alwall, Philip C. Schuster, and Natalia Toro. “Simplified Models for a First Characterization of New Physics at the LHC”. In: *Physical Review D* 79 (Apr. 2009).
- [84] CERN Organisation. *NumberCountingUtils.cxx*. URL: https://root.cern.ch/doc/v606/NumberCountingUtils_8cxx_source.html#l00126 (visited on 02/22/2020).
- [85] CERN Organisation. *ROOT, a Data Analysis Framework*. URL: <https://root.cern.ch/> (visited on 02/22/2020).
- [86] I. Antcheva et al. “ROOT - A C++ Framework for petabyte Data Storage, Statistical Analysis and Visualization”. In: *Computer Physics Communication* 180.12 (2009), pp. 2499–2512. ISSN: 0010-4655. DOI: <https://doi.org/10.1016/j.cpc.2009.08.005>. URL: <http://www.sciencedirect.com/science/article/pii/S0010465509002550> (visited on 02/22/2020).

- [87] ATLAS Collaboration. “Search for electroweak production of supersymmetric particles in final states with two or three leptons at $\sqrt{s} = 13$ TeV with the ATLAS detector”. In: *The European Physical Journal C* 78.12 (Dec. 2018). ISSN: 1434-6052. DOI: 10.1140/epjc/s10052-018-6423-7. URL: <http://dx.doi.org/10.1140/epjc/s10052-018-6423-7> (visited on 02/22/2020).
- [88] M. Baak et al. “HistFitter software framework for statistical data analysis”. In: *The European Physical Journal C* 75.4 (Apr. 2015). ISSN: 1434-6052. DOI: 10.1140/epjc/s10052-015-3327-7. URL: <http://dx.doi.org/10.1140/epjc/s10052-015-3327-7> (visited on 06/16/2019).
- [89] A. L. Read. “Presentation of search results: the CL_s technique”. In: *Journal of Physics G: Nuclear and Particle Physics* 28.10 (Sept. 2002), pp. 2693–2704. DOI: 10.1088/0954-3899/28/10/313. URL: <https://doi.org/10.1088/0954-3899/28/10/313> (visited on 02/28/2020).
- [90] G. Aad et al. “The ATLAS Simulation Infrastructure”. In: *The European Physical Journal C* 70.3 (Sept. 2010), pp. 823–874.
- [91] S. Agostinelli et al. “GEANT4 - A Simulation Toolkit”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 506 (2003), pp. 250–303. URL: <http://www.sciencedirect.com/science/article/pii/S0168900203013688> (visited on 02/21/2020).
- [92] Elisabeth Schopf. “Search for the Higgs Boson Decay to Bottom and Charm Quarks”. Apr. 2018. URL: <https://cds.cern.ch/record/2637125> (visited on 11/16/2020).
- [93] ATLAS Collaboration. “Monte Carlo Generators for the production of a W or Z/γ^* Boson in association with Jets at ATLAS in Run 2”. In: ATL-PHYS-PUB-2016-003 (Jan. 2016).
- [94] T Gleisberg et al. “Event generation with SHERPA 1.1”. In: *Journal of High Energy Physics* 2009.02 (Feb. 2009), p. 007.
- [95] Richard D. Ball et al. “Parton Distributions for the LHC Run II”. In: *Journal of High Energy Physics* 2015.4 (Apr. 2015).
- [96] Tanju Gleisberg and Stefan Höche. “Comix, a new matrix element generator”. In: *Journal of High Energy Physics* 2008.12 (Dec. 2008), pp. 039–039.
- [97] F. Cascioli, P. Maierhöfer, and S. Pozzorini. “Scattering Amplitudes with Open Loops”. In: *Physical Review Letters* 108.11 (Mar. 2012).
- [98] S Schumann and F Krauss. “A parton shower algorithm based on Catani-Seymour dipole factorisation”. In: *Journal of High Energy Physics* 2008.03 (Mar. 2008), pp. 038–038. ISSN: 1029-8479. DOI: 10.1088/1126-6708/2008/03/038. URL: <http://dx.doi.org/10.1088/1126-6708/2008/03/038>.
- [99] Stefan Höche et al. “QCD matrix elements + parton showers. The NLO case”. In: *Journal of High Energy Physics* 2013.4 (Apr. 2013). ISSN: 1029-8479. DOI: 10.1007/jhep04(2013)027. URL: [http://dx.doi.org/10.1007/JHEP04\(2013\)027](http://dx.doi.org/10.1007/JHEP04(2013)027).
- [100] Hung-Liang Lai et al. “New parton distributions for collider physics”. In: *Physical Review D* 82.7 (Oct. 2010). ISSN: 1550-2368. DOI: 10.1103/physrevd.82.074024. URL: <http://dx.doi.org/10.1103/PhysRevD.82.074024>.

- [101] Simone Alioli et al. “A general framework for implementing NLO calculations in shower Monte Carlo programs: the POWHEG BOX”. In: *Journal of High Energy Physics* 2010.6 (June 2010). issn: 1029-8479. DOI: 10.1007/jhep06(2010)043. URL: [http://dx.doi.org/10.1007/JHEP06\(2010\)043](http://dx.doi.org/10.1007/JHEP06(2010)043).
- [102] ATLAS Collaboration. “Simulation of the top quark production for the ATLAS experiment at $\sqrt{s} = 13$ TeV”. In: (Jan. 2016). URL: <https://cds.cern.ch/record/2120417> (visited on 20/22/2020).
- [103] Stefano Frixione, Paolo Nason, and Carlo Oleari. “Matching NLO QCD computations with parton shower simulations: the POWHEG method”. In: *Journal of High Energy Physics* 2007.11 (Nov. 2007), p. 070. issn: 1029-8479. DOI: 10.1088/1126-6708/2007/11/070. URL: <http://dx.doi.org/10.1088/1126-6708/2007/11/070>.
- [104] Pierre Artoisenet et al. “Automatic spin-entangled decays of heavy resonances in Monte Carlo simulations”. In: *Journal of High Energy Physics* 2013.3 (Mar. 2013). issn: 1029-8479. DOI: 10.1007/jhep03(2013)015. URL: [http://dx.doi.org/10.1007/JHEP03\(2013\)015](http://dx.doi.org/10.1007/JHEP03(2013)015).
- [105] Torbjörn Sjöstrand, Stephen Mrenna, and Peter Skands. “PYTHIA 6.4 physics and manual”. In: *Journal of High Energy Physics* 2006.05 (May 2006), p. 026. issn: 1029-8479. DOI: 10.1088/1126-6708/2006/05/026. URL: <http://dx.doi.org/10.1088/1126-6708/2006/05/026>.
- [106] Jonathan Pumplin et al. “New Generation of Parton Distributions with Uncertainties from Global QCD Analysis”. In: *Journal of High Energy Physics* 2002.07 (July 2002), p. 012. issn: 1029-8479. DOI: 10.1088/1126-6708/2002/07/012. URL: <http://dx.doi.org/10.1088/1126-6708/2002/07/012>.
- [107] Peter Z. Skands. “Tuning Monte Carlo generators: The Perugia tunes”. In: *Physical Review D* 82.7 (Oct. 2010). issn: 1550-2368. DOI: 10.1103/physrevd.82.074018. URL: <http://dx.doi.org/10.1103/PhysRevD.82.074018>.
- [108] Michał Czakon, Paul Fiedler, and Alexander Mitov. “Total Top-Quark Pair-Production Cross Section at Hadron Colliders Through $\mathcal{O}(\alpha_s^4)$ ”. In: *Physical Review Letters* 110.25 (June 2013). issn: 1079-7114. DOI: 10.1103/physrevlett.110.252004. URL: <http://dx.doi.org/10.1103/PhysRevLett.110.252004>.
- [109] Michał Czakon and Alexander Mitov. “Top++: A program for the calculation of the top-pair cross-section at hadron colliders”. In: *Computer Physics Communications* 185.11 (Nov. 2014), pp. 2930–2938. issn: 0010-4655. DOI: 10.1016/j.cpc.2014.06.021. URL: <http://dx.doi.org/10.1016/j.cpc.2014.06.021>.
- [110] Nikolaos Kidonakis. “Two-loop soft anomalous dimensions for single top quark associated production with a W^- or H^- ”. In: *Physical Review D* 82.5 (Sept. 2010). issn: 1550-2368. DOI: 10.1103/physrevd.82.054018. URL: <http://dx.doi.org/10.1103/PhysRevD.82.054018>.
- [111] Nikolaos Kidonakis. “Next-to-next-to-leading-order collinear and soft gluon corrections for channel single top quark production”. In: *Physical Review D* 83.9 (May 2011). issn: 1550-2368. DOI: 10.1103/physrevd.83.091503. URL: <http://dx.doi.org/10.1103/PhysRevD.83.091503>.

- [112] M. Aliev et al. “HATHOR – HAdronic Top and Heavy quarks crOss section calculatoR”. In: *Computer Physics Communications* 182.4 (Apr. 2011), pp. 1034–1046. ISSN: 0010-4655. DOI: 10.1016/j.cpc.2010.12.040. URL: <http://dx.doi.org/10.1016/j.cpc.2010.12.040>.
- [113] P. Kant et al. “HatHor for single top-quark production: Updated predictions and uncertainty estimates for single top-quark production in hadronic collisions”. In: *Computer Physics Communications* 191 (June 2015), pp. 74–89. ISSN: 0010-4655. DOI: 10.1016/j.cpc.2015.02.001. URL: <http://dx.doi.org/10.1016/j.cpc.2015.02.001>.
- [114] 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”. In: *Journal of High Energy Physics* 2014.7 (July 2014). ISSN: 1029-8479. DOI: 10.1007/jhep07(2014)079. URL: [http://dx.doi.org/10.1007/JHEP07\(2014\)079](http://dx.doi.org/10.1007/JHEP07(2014)079).
- [115] 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.” In: (Jan. 2016). URL: <https://cds.cern.ch/record/2120826> (visited on 02/22/2020).
- [116] Torbjörn Sjöstrand et al. “An introduction to PYTHIA 8.2”. In: *Computer Physics Communications* 191 (June 2015), pp. 159–177. ISSN: 0010-4655. DOI: 10.1016/j.cpc.2015.01.024. URL: <http://dx.doi.org/10.1016/j.cpc.2015.01.024>.
- [117] Achilleas Lazopoulos et al. “Next-to-leading order QCD corrections to $t\bar{t}Z$ production at the LHC”. In: *Physics Letters B* 666.1 (Aug. 2008), pp. 62–65. ISSN: 0370-2693. DOI: 10.1016/j.physletb.2008.06.073. URL: <http://dx.doi.org/10.1016/j.physletb.2008.06.073>.
- [118] John M. Campbell and R. Keith Ellis. “ $t\bar{t}W$ Production and Decay at NLO”. In: *Journal of High Energy Physics* 2012.7 (July 2012). ISSN: 1029-8479. DOI: 10.1007/jhep07(2012)052. URL: [http://dx.doi.org/10.1007/JHEP07\(2012\)052](http://dx.doi.org/10.1007/JHEP07(2012)052) (visited on 02/22/2020).
- [119] ATLAS Collaboration. “Multi-Boson Simulation for 13 TeV ATLAS Analyses”. In: ATL-PHYS-PUB-2016-002 (Jan. 2016). URL: <https://cds.cern.ch/record/2119986> (visited on 02/22/2020).
- [120] J. M. Campbell and R. K. Ellis. “Update on vector boson pair production at hadron colliders”. In: *Physical Review D* 60.11 (Nov. 1999). ISSN: 1089-4918. DOI: 10.1103/physrevd.60.113006. URL: <http://dx.doi.org/10.1103/PhysRevD.60.113006>.
- [121] John M. Campbell, R. Keith Ellis, and Ciaran Williams. “Vector boson pair production at the LHC”. In: *Journal of High Energy Physics* 2011.7 (July 2011). ISSN: 1029-8479. DOI: 10.1007/jhep07(2011)018. URL: [http://dx.doi.org/10.1007/JHEP07\(2011\)018](http://dx.doi.org/10.1007/JHEP07(2011)018).
- [122] S Schumann and F Krauss. “A parton shower algorithm based on Catani-Seymour dipole factorisation”. In: *Journal of High Energy Physics* 2008.03 (Mar. 2008), pp. 038–038. ISSN: 1029-8479. DOI: 10.1088/1126-6708/2008/03/038. URL: <http://dx.doi.org/10.1088/1126-6708/2008/03/038>.

- [123] Stefan Höche et al. “QCD matrix elements and truncated showers”. In: *Journal of High Energy Physics* 2009.05 (May 2009), pp. 053–053. ISSN: 1029-8479. DOI: 10.1088/1126-6708/2009/05/053. URL: <http://dx.doi.org/10.1088/1126-6708/2009/05/053>.
- [124] ATLAS Collaboration. “Multi-Boson Simulation for 13 TeV ATLAS Analyses”. In: ATL-PHYS-PUB-2017-005 (May 2017). URL: <https://cds.cern.ch/record/2261933> (visited on 02/22/2020).
- [125] Gennaro Corcella et al. “HERWIG 6: an event generator for hadron emission reactions with interfering gluons (including supersymmetric processes)”. In: *Journal of High Energy Physics* 2001.01 (Jan. 2001), pp. 010–010. ISSN: 1029-8479. DOI: 10.1088/1126-6708/2001/01/010. URL: <http://dx.doi.org/10.1088/1126-6708/2001/01/010>.
- [126] S Dittmaier and et all. “Handbook of LHC Higgs Cross-Sections: 2. Differential Distributions”. In: (2012).
- [127] Fabio Maltoni and Tim Stelzer. “MadEvent: automatic event generation with MadGraph”. In: *Journal of High Energy Physics* 2003.02 (Feb. 2003), pp. 027–027. ISSN: 1029-8479. DOI: 10.1088/1126-6708/2003/02/027. URL: <http://dx.doi.org/10.1088/1126-6708/2003/02/027>.
- [128] ATLAS Collaboration. “ATLAS Run 1 Pythia8 Tunes”. In: (Nov. 2014). URL: <https://cds.cern.ch/record/1966419> (visited on 02/22/2020).
- [129] Leif Lönnblad and Stefan Prestel. “Matching tree-level matrix elements with interleaved showers”. In: *Journal of High Energy Physics* 2012.3 (Mar. 2012). ISSN: 1029-8479. DOI: 10.1007/jhep03(2012)019. URL: [http://dx.doi.org/10.1007/JHEP03\(2012\)019](http://dx.doi.org/10.1007/JHEP03(2012)019).
- [130] W. Beenakker et al. “Squark and gluino production at hadron colliders”. In: *Nuclear Physics B* 492.1-2 (May 1997), pp. 51–103. ISSN: 0550-3213. DOI: 10.1016/S0550-3213(97)80027-2. URL: [http://dx.doi.org/10.1016/S0550-3213\(97\)80027-2](http://dx.doi.org/10.1016/S0550-3213(97)80027-2).
- [131] A. Kulesza and L. Motyka. “Threshold Resummation for Squark-Antisquark and Gluino-Pair Production at the LHC”. In: *Physical Review Letters* 102.11 (Mar. 2009). ISSN: 1079-7114. DOI: 10.1103/physrevlett.102.111802. URL: <http://dx.doi.org/10.1103/PhysRevLett.102.111802>.
- [132] A. Kulesza and L. Motyka. “Threshold Resummation for Squark-Antisquark and Gluino-Pair Production at the LHC”. In: *Physical Review Letters* 102.11 (Mar. 2009). ISSN: 1079-7114. DOI: 10.1103/physrevlett.102.111802. URL: <http://dx.doi.org/10.1103/PhysRevLett.102.111802>.
- [133] A. Kulesza and L. Motyka. “Soft gluon resummation for the production of gluino-gluino and squark-antisquark pairs at the LHC”. In: *Physical Review D* 80.9 (Nov. 2009). ISSN: 1550-2368. DOI: 10.1103/physrevd.80.095004. URL: <http://dx.doi.org/10.1103/PhysRevD.80.095004>.
- [134] Wim Beenakker et al. “Soft-gluon resummation for squark and gluino hadroproduction”. In: *Journal of High Energy Physics* 2009.12 (Dec. 2009), pp. 041–041. ISSN: 1029-8479. DOI: 10.1088/1126-6708/2009/12/041. URL: <http://dx.doi.org/10.1088/1126-6708/2009/12/041>.

- [135] WIM BEENAKKER et al. “SQUARK AND GLUINO HADROPRODUCTION”. In: *International Journal of Modern Physics A* 26.16 (June 2011), pp. 2637–2664. ISSN: 1793-656X. DOI: 10.1142/s0217751x11053560. URL: <http://dx.doi.org/10.1142/S0217751X11053560>.
- [136] Christoph Borschensky et al. “Squark and gluino production cross sections in pp collisions at $\sqrt{s} = 13, 14$ and 100 TeV”. In: *The European Physical Journal C* 74.12 (Dec. 2014). ISSN: 1434-6052. DOI: 10.1140/epjc/s10052-014-3174-y. URL: <http://dx.doi.org/10.1140/epjc/s10052-014-3174-y>.
- [137] David J. Lange. “The EvtGen Particle Decay Simulation Package”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 462 (2001), pp. 152–155. ISSN: 0168-9002. DOI: [https://doi.org/10.1016/S0168-9002\(01\)00089-4](https://doi.org/10.1016/S0168-9002(01)00089-4). URL: <http://www.sciencedirect.com/science/article/pii/S0168900201000894>.
- [138] ATLAS Collaboration. “Summary of ATLAS Pythia 8 Tunes”. In: ATL-PHYS-PUB-2012-003 (Aug. 2012). URL: <https://cds.cern.ch/record/1474107> (visited on 02/22/2020).
- [139] A. D. Martin et al. “Parton distributions for the LHC”. In: *The European Physical Journal C* 63.2 (July 2009), pp. 189–285. ISSN: 1434-6052. DOI: 10.1140/epjc/s10052-009-1072-5. URL: <http://dx.doi.org/10.1140/epjc/s10052-009-1072-5>.
- [140] ATLAS Collaboration. “Search for the lepton flavor violating decay $Z \rightarrow e\mu$ in pp collisions at $\sqrt{s}=8$ TeV with the ATLAS detector”. In: *Physical Review D* ().
- [141] Matteo Cacciari and Gavin P. Salam. “Pileup subtraction using jet areas”. In: *Physics Letters B* 659 (Jan. 2008), pp. 119–126.
- [142] Morad Aaboud et al. “Jet energy scale measurements and their systematic uncertainties in proton-proton collisions at $\sqrt{s} = 13$ TeV with the ATLAS Detector”. In: *Physical Review D* 96 (7 Oct. 2017). eprint: 1703.09665.
- [143] “Tagging and suppression of pileup jets with the ATLAS detector”. In: (May 2014). URL: <https://cds.cern.ch/record/1700870> (visited on 02/22/2020).
- [144] Paul Jackson and Christopher Rogan. “Recursive Jigsaw Reconstruction: HEP Event Analysis in the presence of Kinematic and Combinatoric Ambiguities”. In: *Physical Review D* 96.11 (Dec. 2017). ISSN: 2470-0029. DOI: 10.1103/physrevd.96.112007. URL: <http://dx.doi.org/10.1103/PhysRevD.96.112007> (visited on 02/28/2020).
- [145] Paul Jackson, Christopher Rogan, and Marco Santoni. “Sparticles in Motion: Analyzing Compressed SUSY Scenarios with a New Method of Event Reconstruction”. In: *Physical Review D* 95.3 (Feb. 2017). ISSN: 2470-0029. DOI: 10.1103/physrevd.95.035031. URL: <http://dx.doi.org/10.1103/PhysRevD.95.035031> (visited on 02/28/2020).
- [146] Marco Santoni. “Probing Supersymmetry with Recursive Jigsaw Reconstruction”. Ph.D Thesis. University of Adelaide, 2017.

- [147] Christopher G Lester and Alan J Barr. “mTGen: mass scale measurements in pair-production at colliders”. In: *Journal of High Energy Physics* 2007.12 (Dec. 2007), pp. 102–102. ISSN: 1029-8479. DOI: 10.1088/1126-6708/2007/12/102. URL: <http://dx.doi.org/10.1088/1126-6708/2007/12/102>.
- [148] Ben Gripaios. “Transverse observables and mass determination at hadron colliders”. In: *Journal of High Energy Physics* 2008.02 (Feb. 2008), pp. 053–053. ISSN: 1029-8479. DOI: 10.1088/1126-6708/2008/02/053. URL: <http://dx.doi.org/10.1088/1126-6708/2008/02/053>.
- [149] Daniel R Tovey. “On measuring the masses of pair-produced semi-invisibly decaying particles at hadron colliders”. In: *Journal of High Energy Physics* 2008.04 (Apr. 2008), pp. 034–034. ISSN: 1029-8479. DOI: 10.1088/1126-6708/2008/04/034. URL: <http://dx.doi.org/10.1088/1126-6708/2008/04/034>.
- [150] Partha Konar, Kyoungchul Kong, and Konstantin T Matchev. “ $\sqrt{\hat{s}_{\min}}$: a global inclusive variable for determining the mass scale of new physics in events with missing energy at hadron colliders”. In: *Journal of High Energy Physics* 2009.03 (Mar. 2009), pp. 085–085. ISSN: 1029-8479. DOI: 10.1088/1126-6708/2009/03/085. URL: <http://dx.doi.org/10.1088/1126-6708/2009/03/085>.
- [151] Tao Han, Ian-Woo Kim, and Jeonghyeon Song. “Kinematic Cusps: Determining the Missing Particle Mass at Colliders”. In: *Physics Letters B* 693.5 (2010), pp. 575–579. ISSN: 0370-2693. DOI: <https://doi.org/10.1016/j.physletb.2010.09.010>. URL: <http://www.sciencedirect.com/science/article/pii/S0370269310010610> (visited on 03/01/2020).
- [152] Ben Gripaios, Kazuki Sakurai, and Bryan Webber. “Polynomials, Riemann surfaces, and reconstructing missing-energy events”. In: *Journal of High Energy Physics* 2011.9 (Sept. 2011). ISSN: 1029-8479. DOI: 10.1007/jhep09(2011)140. URL: [http://dx.doi.org/10.1007/JHEP09\(2011\)140](http://dx.doi.org/10.1007/JHEP09(2011)140).
- [153] Georges Aad et al. “Measurement of the top quark-pair production cross section with ATLAS in pp collisions at $\sqrt{s} = 7$ TeV”. In: *The European Physical Journal C* 71 (3 2011), pp. 1–36. eprint: 1012.1792.
- [154] M. Aaboud and et al. “Search for new phenomena in events containing a same-flavour opposite-sign dilepton pair, jets and large missing transverse momentum in $\sqrt{s} = 13$ TeV pp collisions with the ATLAS detector”. In: *The European Physical Journal C* 77.3 (Mar. 2017). ISSN: 1434-6052. DOI: 10.1140/epjc/s10052-017-4700-5. URL: <http://dx.doi.org/10.1140/epjc/s10052-017-4700-5> (visited on 04/09/2020).
- [155] M. Aaboud and et al. “Jet energy scale measurements and their systematic uncertainties in proton-proton collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector”. In: *Physical Review D* 96.7 (Oct. 2017). ISSN: 2470-0029. DOI: 10.1103/physrevd.96.072002. URL: <http://dx.doi.org/10.1103/PhysRevD.96.072002>.
- [156] ATLAS Collaboration. “Expected performance of missing transverse momentum reconstruction for the ATLAS detector at $\sqrt{s} = 13$ TeV”. In: ATL-PHYS-PUB-2015-023 (July 2015). URL: <https://cds.cern.ch/record/2037700> (visited on 04/15/2020).

- [157] ATLAS Collaboration. “Calibration of b -tagging using dileptonic top pair events in a combinatorial likelihood approach with the ATLAS experiment”. In: ATLAS-CONF-2014-004 (Feb. 2014). URL: <https://cds.cern.ch/record/1664335> (visited on 04/15/2020).
- [158] ATLAS Collaboration. “Calibration of the performance of b -tagging for c and light-flavour jets in the 2012 ATLAS data”. In: ATLAS-CONF-2014-046 (July 2014). URL: <https://cds.cern.ch/record/1741020> (visited on 04/15/2020).
- [159] ATLAS Collaboration. “Electron efficiency measurements with the ATLAS detector using the 2015 LHC proton-proton collision data”. In: ATLAS-CONF-2016-024 (June 2016). URL: <http://cds.cern.ch/record/2157687> (visited on 04/15/2020).
- [160] CERN Organisation. *CERN is celebrating 60 years of science for peace in 2014*. URL: <https://cern60.web.cern.ch/en> (visited on 03/22/2019).
- [161] William Prout. “On the Relation between the Specific Gravities of Bodies in their Gaseous State and the Weights of their Atoms”. In: *Annals of Philosophy* 6 (1815), pp. 321–330. URL: <http://web.lemoyne.edu/~giunta/PROUT.HTML> (visited on 04/19/2019).
- [162] Ernest Rutherford. “LIV. Collision of α particles with light atoms. IV. An anomolous effect in nitrogen”. In: *The London, Edinburgh and Dublin Philosophical Magazine and Journal of Science* 37 (222 1919), pp. 581–587.
- [163] Michael Faraday. “VIII. Experimental researches in electricity. - Thirteenth series.” In: *Philosophical Transactions of the Royal Society of London* 128 (1838), pp. 125–168.
- [164] Joseph John Thomson. *The Discharge of Electricity through Gases*. Charles Scribners sons, 1903.
- [165] Michael A. Grayson. *Measuring Mass: From Positive Rays to Proteins*. Chemical Heritage Press, 2002.
- [166] Joseph John Thomson. “Cathode Rays”. In: *Philosophical Magazine* 44 (1897), pp. 293–316.
- [167] Antoine Henri Becquerel. “Déviation du rayonnement du radium dans un champ électrique”. In: *Comptes rendus de l’Académie des sciences* 130 (1900), pp. 809–815.
- [168] Ernest Rutherford and Thomas Royds. “Spectrum of the Radium emanation”. In: *Philosophical Magazine* 16 (1908), pp. 313–317.
- [169] Albert Einstein. “Über einen die Erzeugung und Verwandlung des Lichtes betreffenden heuristischen Gesichtspunkt”. In: *Annalen der Physik* 322 (1905), pp. 132–148.
- [170] Hans Geiger. “On the scattering of α -Particles by matter”. In: *Proceedings of the Royal Society of London A* 81 (1908), pp. 174–177.
- [171] Hans Geiger and Ernest Marsden. “On a Diffuse Relfection of the α -Particles”. In: *Procceedings of the Royal Society of London A* 82 (1909), pp. 495–500.
- [172] Hans Geiger. “The Scattering of the α -Particles by Matter”. In: *Proceedings of the Royal Society of London A* 83 (1910), pp. 492–504.

- [173] Hans Geiger and Ernest Marsden. “The Laws of Deflexion of α -Particles through Large Angles”. In: *Philosophical Magazine* 25 (1913), pp. 604–623.
- [174] Herbert Becker and Walter Bothe. “Künstliche Erregung von Kern- γ -Strahlen”. In: *Zeitschrift für Physik* 66 (1930), pp. 289–306.
- [175] Herbert Becker and Walter Bothe. “Die in Bor und Beryllium erregten γ -Strahlen”. In: *Zeitschrift für Physik* 76 (1932), pp. 421–438.
- [176] James Chadwick. “Possible Existence of a Neutron”. In: *Nature* 129 (1932).
- [177] James Chadwick. “Bakerian Lecture. The Neutron”. In: *Proceedings of the Royal Society A* 142 (1933).
- [178] NobelPrize.org. *The Nobel Prize in Physics 1935*. URL: <https://www.nobelprize.org/prizes/physics/1935/summary/> (visited on 04/19/2019).
- [179] Paul Adrien Maurice Dirac. “The Quantum Theory of the Electron”. In: *Proceedings of the Royal Society A* 117 (1928), pp. 610–624.
- [180] Paul Adrien Maurice Dirac. “Quantised Singularities in the Quantum Field”. In: *Proceedings of the Royal Society A* 133 (1931), pp. 60–72.
- [181] Carl D. Anderson. “The Positive Electron”. In: *Physical Review* 43 (1933), pp. 491–494.
- [182] NobelPrize.org. *The Nobel Prize in Physics 1936*. URL: <https://www.nobelprize.org/prizes/physics/1936/summary/> (visited on 04/19/2019).
- [183] Carl D. Anderson and Seth H. Neddermeyer. “Cloud Chamber Observations of Cosmic Rays at 4300 Meters Elevation and Near Sea-Level”. In: *Physical Review* 50 (1936), pp. 263–271.
- [184] J. C. Street and E. C. Stevenson. “New Evidence for the Existence of a Particle of Mass Intermediate Between the Proton and Electron”. In: *Physical Review* 52 (1937).
- [185] Hideki Yukawa. “On the Interaction of Elementary Particles”. In: *Proceedings of the Physico-Mathematical Society of Japan* 17 (1935), pp. 139–148.
- [186] César Lattes, Giuseppe Occhialini, and Cecil Powell. “Observations on the Tracks of Slow Mesons in Photographic Emulsions”. In: *Nature* 160 (1947), pp. 486–492.
- [187] George Dixon Rochester and Clifford Charles Butler. “Evidence for the Existence of New Unstable Elementary Particles”. In: *Nature* 160 (1947), pp. 855–857.
- [188] CERN Organisation. *Origins*. URL: <https://timeline.web.cern.ch/origins> (visited on 04/25/2019).
- [189] CERN Organisation. *Early Beginnings*. URL: <https://timeline.web.cern.ch/early-beginnings> (visited on 04/25/2019).
- [190] CERN Organisation. *Where to Build*. URL: <https://timeline.web.cern.ch/where-build> (visited on 04/25/2019).
- [191] CERN Organisation. *Convention Complete*. URL: <https://timeline.web.cern.ch/convention-complete> (visited on 04/25/2019).
- [192] Owen Chamberlain et al. “Observation of Antiprotons”. In: *Physical Review* 100 (1955).

- [193] C. L. Cowan et al. “Detection of the Free Neutrino: A Confirmation”. In: *Science* 124 (1956), pp. 103–104.
- [194] CERN Organisation. *CERN’s First Accelerator - the Synchrocyclotron - Starts Up*. URL: <https://timeline.web.cern.ch/cerns-first-accelerator-synchrocyclotron-starts> (visited on 04/29/2019).
- [195] CERN Organisation. *Proton Synchrotron Starts Up*. URL: <https://timeline.web.cern.ch/proton-synchrotron-starts> (visited on 04/29/2019).
- [196] G. Danby et al. “Observation of high-energy neutrino reactions and the existence of two kinds of neutrinos”. In: *Physical Review Letters* 9 (1962).
- [197] CERN Organisation. *First Observations of Antinuclei*. URL: <https://timeline.web.cern.ch/first-observations-antinuclei>.
- [198] CERN Organisation. *Georges Charpak Revolutionises Detection*. URL: <https://timeline.web.cern.ch/georges-charpak-revolutionizes-detection> (visited on 04/30/2019).
- [199] CERN Organisation. *First Proton Collisions: The Intersecting Storage Rings*. URL: <https://timeline.web.cern.ch/first-proton-collisions-intersecting-storage-rings> (visited on 04/30/2019).
- [200] CERN Organisation. *Council Commissions the Super Proton Synchrotron*. URL: <https://timeline.web.cern.ch/council-commissions-super-proton-synchrotron> (visited on 04/30/2019).
- [201] CERN Organisation. *Super Proton Synchrotron Tunnel Completed*. URL: <https://timeline.web.cern.ch/super-proton-synchrotron-tunnel-completed> (visited on 04/30/2019).
- [202] J. J. Aubert et al. “Experimental Observation of a Heavy Particle J ”. In: *Physical Review Letters* 33 (Dec. 1974), pp. 1404–1406.
- [203] CERN Organisation. *The Super Proton Synchrotron Starts Up*. URL: <https://timeline.web.cern.ch/super-proton-synchrotron-starts> (visited on 04/30/2019).
- [204] S. W. Herb et al. “Observation of a Dimuon Resonance at 9.5 GeV in 400 GeV Proton-Nucleus Collisions”. In: *Physical Review Letters* 39 (1977), pp. 252–255.
- [205] M. L. Perl et al. “Evidence for Anomalous Lepton Production in e^+e^- Annihilation”. In: *Physical Review Letters* 35 (1975), pp. 1489–1492.
- [206] Martin Lewis Perl. “Evidence for, and Properties of, the New Charged Heavy Lepton”. In: *Proceedings: 12th Rencontres de Moriond, Leptons and Multi-Leptons. Flaine-Haute Savoie, France Mar 6-Mar 18 1977. 1*. 1977, pp. 75–97. URL: <http://www-public.slac.stanford.edu/sciDoc/docMeta.aspx?slacPubNumber=SLAC-PUB-1923> (visited on 04/30/2019).
- [207] Ch. Berger et al. “First observation of hadron production in e^+e^- collisions at 13 and 17 GeV CMS energy with the pluto detector at petra: Pluto Collaboration”. In: *Physics Letters B* 81 (1979), pp. 410–415.
- [208] CERN Organisation. *First Proton-Antiproton Collisions*. URL: <https://timeline.web.cern.ch/first-proton-antiproton-collisions> (visited on 04/30/2019).
- [209] CERN Organisation. *W and Z Particles Discovered*. URL: <https://timeline.web.cern.ch/w-and-z-particles-discovered> (visited on 04/30/2019).

- [210] CERN Organisation. *Heavy-Ion Collisions Begin*. URL: <https://timeline.web.cern.ch/heavy-ion-collisions-begin> (visited on 04/30/2019).
- [211] CERN Organisation. *LEP Tunnel Completed*. URL: <https://timeline.web.cern.ch/lep-tunnel-completed> (visited on 04/30/2019).
- [212] CERN Organisation. *Large Electron-Positron Collider: First Injection*. URL: <https://timeline.web.cern.ch/large-electron-positron-collider-first-injection> (visited on 04/30/2019).
- [213] CERN Organisation. *The World's First Browser/Editor, Website and Server go Live at CERN*. URL: <https://timeline.web.cern.ch/worlds-first-browsereditor-website-and-server-go-live-cern> (visited on 04/30/2019).
- [214] F. Abe et al. "Observation of Top Quark Production in $p\bar{p}$ Collisions with the Collider Detector at Fermilab". In: *Physical Review Letters* 74 (1995), pp. 2626–2631.
- [215] S. Abachi et al. "Search for High Mass Top Quark Production in $p\bar{p}$ Collisions at $\sqrt{s} = 1.8$ TeV". In: *Physical Review Letters* 74 (1995), pp. 2422–2426.
- [216] CERN Organisation. *First Antiatoms Produced: Antihydrogen, at CERN*. URL: <https://timeline.web.cern.ch/first-antiatoms-produced-antihydrogen-cern> (visited on 04/30/2019).
- [217] CERN Organisation. *CMS and ATLAS Experiments Approved*. URL: <https://timeline.web.cern.ch/cms-and-atlas-experiments-approved> (visited on 04/30/2019).
- [218] CERN Organisation. *ALICE Experiment Approved*. URL: <https://timeline.web.cern.ch/alice-experiment-approved> (visited on 04/30/2019).
- [219] CERN Organisation. *Antiproton Decelerator Approved*. URL: <https://timeline.web.cern.ch/antiproton-decelerator-approved> (visited on 04/30/2019).
- [220] CERN Organisation. *LHCb Experiment Approved*. URL: <https://timeline.web.cern.ch/lhcb-experiment-approved> (visited on 04/30/2019).
- [221] K. Kodama et al. "Observation of Tau Neutrino Interactions". In: *Physics Letters B* 504 (2001), pp. 218–224.
- [222] W. Bartel et al. "Experimental Studies on Multi-Jet Production in e^+e^- Annihilation at PETRA Energies". In: *Zeitschrift für Physik* 33 (1986), p. 23.
- [223] Yu L. Dokshitzner et al. "Better Jet Clustering Algorithms". In: *Journal of High Energy Physics* 08 (1997), p. 001.

PART IV:

APPENDICES

It is not in the nature of things for any one man to make a sudden violent discovery; science goes step by step, and every man depends on the work of his predecessors. When you hear of a sudden unexpected discovery - a bolt from the blue, as it were - you can always be sure that it has grown up by the influence of one man on another, and it is this mutual influence which makes the enormous possibility of scientific advance. Scientists are not dependent on the ideas of a single man, but on the combined wisdom of thousands of men, all thinking of the same problem, and each doing his little bit to add to the great structure of knowledge which is gradually being erected.

Ernest Rutherford

Appendix A

Gamma Matrices

The **Gamma Matrices** (also known as the **Dirac Matrices**) are a set of matrices that follow a very specific set of commutation relations which ensures they generate a matrix representation of the Clifford algebra $CL_{1,3}(R)$. These matrices are used extensively in Chapter 1 and Appendix B so some of the properties of these matrices are listed here.

There are four contravariant gamma matrices: γ^0 , γ^1 , γ^2 and γ^3 . Their defining property is the anticommutation relation:

$$\{\gamma^\mu, \gamma^\nu\} = \gamma^\mu \gamma^\nu + \gamma^\nu \gamma^\mu = 2g^{\mu\nu} \mathcal{I}_4 \quad (\text{A.1})$$

where $\{\cdot, \cdot\}$ is the anticommutator, $g^{\mu\nu}$ is the Minkowski metric with signature $(+, -, -, -)$ and $\mathcal{I}$ is the 4×4 identity matrix. The corresponding covariant matrices are given by:

$$\gamma_\mu = g_{\mu\nu} \gamma^\nu \quad (\text{A.2})$$

Note that the sign convention of the Minkowski metric matters. If the chosen metric signature is $(-, +, +, +)$ the anticommutation relation becomes:

$$\{\gamma^\mu, \gamma^\nu\} = \gamma^\mu \gamma^\nu + \gamma^\nu \gamma^\mu = -2g^{\mu\nu} \mathcal{I}_4 \quad (\text{A.3})$$

There is also a fifth gamma matrix, γ^5 , which is proportional to the product of the first four gamma matrices:

$$\gamma^5 = i\gamma^0\gamma^1\gamma^2\gamma^3 \quad (\text{A.4})$$

The anticommutation property is more fundamental than the numerical values of any given representation of the matrices, and results proved using it are stronger than results proved using a one representation. Below are some identities that are used in this thesis that are true for all representations (ie, they can be proven true using the anticommutation relation).

$$\gamma^5 = (\gamma^5)^\dagger \quad (\text{A.5})$$

$$\gamma^5 \gamma^5 = \mathcal{I}_4 \quad (\text{A.6})$$

$$\{\gamma^5, \gamma^\mu\} = \gamma^5 \gamma^\mu + \gamma^\mu \gamma^5 = 0 \quad (\text{A.7})$$

Finally, one of the most commonly used representations of the gamma matrices is the Dirac representation, the numerical values of which are given below.

$$\begin{aligned} \gamma^0 &= \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & -1 & 0 \\ 0 & 0 & 0 & -1 \end{bmatrix}, & \gamma^1 &= \begin{bmatrix} 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \\ 0 & -1 & 0 & 0 \\ -1 & 0 & 0 & 0 \end{bmatrix}, \\ \gamma^2 &= \begin{bmatrix} 0 & 0 & 0 & -i \\ 0 & 0 & i & 0 \\ 0 & i & 0 & 0 \\ -i & 0 & 0 & 0 \end{bmatrix}, & \gamma^3 &= \begin{bmatrix} 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & -1 \\ -1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \end{bmatrix}, \end{aligned} \quad (\text{A.8})$$

$$\gamma^5 = \begin{bmatrix} 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \\ 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \end{bmatrix} \quad (\text{A.9})$$

Appendix B

Chiral Lagrangian

Chapter 1, Section 1.4 presents a free fermion Lagrangian whose only terms are left-left and right-right coupled dynamics terms. There is a brief description of why this is in the chapter but no mathematical proof. This appendix will supply that proof in full detail.

First, we restate the definitions of the chiral-left and -right fermion fields and projection operators.

$$\psi(x^\mu) = \Psi_L(x^\mu) + \psi_R(x^\mu) \quad (\text{B.1})$$

$$\hat{P}_L = \frac{1-\gamma^5}{2} \quad , \quad \hat{P}_R = \frac{1+\gamma^5}{2} \quad (\text{B.2})$$

The base free fermion Lagrangian is given by:

$$\mathcal{L} = i\bar{\psi}\gamma^\mu\mathcal{D}_\mu\psi - m\bar{\psi}\psi \quad (\text{B.3})$$

Expanding the field ψ in terms of its chiral-left and -right components gives:

$$\begin{aligned} \mathcal{L} &= i(\bar{\Psi}_L + \bar{\psi}_R)\gamma^\mu\mathcal{D}_\mu(\Psi_L + \psi_R) - m(\bar{\Psi}_L + \bar{\psi}_R)(\Psi_L + \psi_R) \\ &= i(\bar{\Psi}_L\gamma^\mu\mathcal{D}_\mu\Psi_L + \bar{\Psi}_L\gamma^\mu\mathcal{D}_\mu\psi_R + \bar{\psi}_R\gamma^\mu\mathcal{D}_\mu\Psi_L + \bar{\psi}_R\gamma^\mu\mathcal{D}_\mu\psi_R) \\ &\quad - m(\bar{\Psi}_L\Psi_L + \bar{\Psi}_L\psi_R + \bar{\psi}_R\Psi_L + \bar{\psi}_R\psi_R) \end{aligned} \quad (\text{B.4})$$

To see which terms can be removed from this equation consider the following properties of Ψ_L and ψ_R :

$$\gamma^5\Psi_L = -\Psi_L \quad , \quad \gamma^5\psi_R = \psi_R \quad (\text{B.5})$$

Proof:

$$\begin{aligned}
 \gamma^5 \Psi_L &= \gamma^5 \hat{P}_L \psi \\
 &= \gamma^5 \frac{1}{2} (1 - \gamma^5) \psi \\
 &= \gamma^5 \times \left(\frac{1}{2} \psi - \frac{1}{2} \gamma^5 \psi \right) \\
 &= \frac{1}{2} \gamma^5 \psi - \frac{1}{2} \underbrace{\gamma^5 \gamma^5}_{\gamma^5 \text{ is its own inverse}} \psi \\
 &= \frac{1}{2} \gamma^5 \psi - \frac{1}{2} \psi \\
 &= -\frac{1}{2} (1 - \gamma^5) \psi \\
 &= -\Psi_L
 \end{aligned} \tag{B.6}$$

$$\begin{aligned}
 \gamma^5 \psi_R &= \gamma^5 \hat{P}_R \psi \\
 &= \gamma^5 \frac{1}{2} (1 + \gamma^5) \psi \\
 &= \gamma^5 \times \left(\frac{1}{2} \psi + \frac{1}{2} \gamma^5 \psi \right) \\
 &= \frac{1}{2} \gamma^5 \psi + \frac{1}{2} \gamma^5 \gamma^5 \psi \\
 &= \frac{1}{2} \gamma^5 \psi + \frac{1}{2} \psi \\
 &= \frac{1}{2} (1 + \gamma^5) \psi \\
 &= \psi_R
 \end{aligned} \tag{B.7}$$

Also of use is how the fields $\bar{\Psi}_L$ and $\bar{\psi}_R$ react when being post-multiplied by γ^5 :

$$\begin{aligned}
\bar{\Psi}_L \gamma^5 &= (\Psi_L)^\dagger \gamma^0 \gamma^5 \\
&= \left(\frac{1}{2} \psi - \frac{1}{2} \gamma^5 \psi \right)^\dagger \gamma^0 \gamma^5 \\
&= - \left(\frac{1}{2} \psi^\dagger - \frac{1}{2} \psi^\dagger (\gamma^5)^\dagger \right) \underbrace{\gamma^5 \gamma^0}_{\gamma^\mu \gamma^5 = -\gamma^5 \gamma^\mu} \\
&= - \underbrace{\left(\frac{1}{2} \psi^\dagger (\gamma^5)^\dagger - \frac{1}{2} \gamma^5 \gamma^5 \right)}_{(\gamma^5)^\dagger = \gamma^5} \gamma^0 \\
&= \left(\frac{1}{2} \psi^\dagger - \frac{1}{2} \psi^\dagger (\gamma^5)^\dagger \right) \gamma^0 \\
&= \left(\frac{1}{2} (1 - \gamma^5) \psi \right)^\dagger \gamma^0 \\
&= (\Psi_L)^\dagger \gamma^0 \\
&= \bar{\Psi}_L
\end{aligned} \tag{B.8}$$

$$\begin{aligned}
\bar{\psi}_R \gamma^5 &= (\psi_R)^\dagger \gamma^0 \gamma^5 \\
&= \left(\frac{1}{2} \psi + \frac{1}{2} \gamma^5 \psi \right)^\dagger \gamma^0 \gamma^5 \\
&= - \left(\frac{1}{2} \psi^\dagger + \frac{1}{2} \psi^\dagger (\gamma^5)^\dagger \right) \gamma^5 \gamma^0 \\
&= - \left(\frac{1}{2} \psi^\dagger (\gamma^5)^\dagger + \frac{1}{2} \psi^\dagger \gamma^5 \gamma^5 \right) \gamma^0 \\
&= - \left(\frac{1}{2} \psi^\dagger + \frac{1}{2} \psi^\dagger (\gamma^5)^\dagger \right) \gamma^0 \\
&= - \left(\frac{1}{2} (1 + \gamma^5) \psi \right)^\dagger \gamma^0 \\
&= -(\psi_R)^\dagger \gamma^0 \\
&= -\bar{\psi}_R
\end{aligned} \tag{B.9}$$

Knowing the above, several of the terms in the Lagrangian have an interesting property.

$$\begin{aligned}
\bar{\Psi}_L \gamma^\mu \mathcal{D}_\mu \psi_R &= \bar{\Psi}_L \gamma^\mu \mathcal{D}_\mu \gamma^5 \gamma^5 \psi_R \\
&= \bar{\Psi}_L \gamma^\mu \gamma^5 \mathcal{D}_\mu \gamma^5 \psi_R \\
&= \bar{\Psi}_L \times (-\gamma^5 \gamma^\mu) \times \mathcal{D}_\mu \psi_R \\
&= -\bar{\Psi}_L \gamma^5 \gamma^\mu \mathcal{D}_\mu \psi_R \\
&= -\bar{\Psi}_L \gamma^\mu \mathcal{D}_\mu \psi_R
\end{aligned} \tag{B.10}$$

$$\begin{aligned}
\bar{\psi}_R \gamma^\mu \mathcal{D}_\mu \Psi_L &= \bar{\psi}_R \gamma^\mu \mathcal{D}_\mu \gamma^5 \gamma^5 \Psi_L \\
&= \bar{\psi}_R \gamma^\mu \gamma^5 \mathcal{D}_\mu \gamma^5 \Psi_L \\
&= \bar{\psi}_R \times (-\gamma^5 \gamma^\mu) \times \mathcal{D}_\mu \times (-\Psi_L) \\
&= \bar{\psi}_R \gamma^5 \gamma^\mu \mathcal{D}_\mu \Psi_L \\
&= -\bar{\psi}_R \gamma^\mu \mathcal{D}_\mu \Psi_L
\end{aligned} \tag{B.11}$$

$$\begin{aligned}
\bar{\Psi}_L \Psi_L &= \bar{\Psi}_L \gamma^5 \gamma^5 \Psi_L \\
&= \bar{\Psi}_L \times (-\Psi_L) \\
&= -\bar{\Psi}_L \Psi_L
\end{aligned} \tag{B.12}$$

$$\begin{aligned}
\bar{\psi}_R \psi_R &= \bar{\psi}_R \gamma^5 \gamma^5 \psi_R \\
&= (-\bar{\psi}_R) \times \psi_R \\
&= -\bar{\psi}_R \psi_R
\end{aligned} \tag{B.13}$$

All four of the above terms are equal to their negatives $\forall x^\mu$. The only way this can be true is if all of the above terms are exactly 0 for all values of x^μ . Thus, the Lagrangian above only has four terms:

$$\mathcal{L} = i(\bar{\Psi}_L \gamma^\mu \mathcal{D}_\mu \Psi_L + \bar{\psi}_R \gamma^\mu \mathcal{D}_\mu \psi_R) - m(\bar{\Psi}_L \psi_R + \bar{\psi}_R \Psi_L) \tag{B.14}$$

But since the left-right coupled terms aren't gauge invariant, they can't appear in the Lagrangian. As such, the final Lagrangian only has two terms, as it appears in Equation 1.40.

$$\mathcal{L} = i(\bar{\Psi}_L \gamma^\mu \mathcal{D}_\mu \Psi_L + \bar{\psi}_R \gamma^\mu \mathcal{D}_\mu \psi_R) \tag{B.15}$$

Appendix C

History of Particle Physics and CERN

The field of particle physics has a history spanning just over two centuries, while CERN just recently celebrated its 60th anniversary in 2014 [160]). This appendix will summarise the history of particle physics over the last 200 years, with special focus on the contributions made by CERN scientists after its founding in 1954 [31].

C.1 Particle Physics in the 19th and 20th Centuries


The study of particles has existed in some form or another since the early 1800's when William Prout hypothesised[161] that all matter was built up from hydrogen, foreshadowing the discovery of the proton by almost a century[162]. In the time between Prout's hypothesis in 1815 and the foundation of CERN in 1954[31] particle physics evolved from a disconnected set of chemists and physicists studying seemingly separate phenomena like “cathode rays” and the photoelectric effect to a unified field with a community of physicists who understood what they were studying, how they were related and how to advance the field.

In 1815 William Prout would hypothesise that all matter was built up from hydrogen[161]. He introduced this idea after observing that the measurements of atomic weights of elements that were available in the day all seemed to be integer multiples of the atomic weight of hydrogen. While the measurements taken in those days were inaccurate and would be replaced with better measurements that would show that the atomic weights of elements weren't *exact* multiples of the atomic weight of hydrogen, Prout's hypothesis is still prescient. The idea that hydrogen atoms can be used to construct all other atoms is a precursor to one of the foundational concepts of modern

atomic theory: that the amount of protons (and neutrons) in a nucleus determine which element (and isotope) the atom is.

In the mid-1800's discoveries would be made primarily through the use of vacuum tubes. In 1838, Michael Faraday would take the first steps towards using vacuum tubes to discover new particles. In his experiment, he applied a high voltage across two electrodes with a glass tube between them. This glass tube had then been partially evacuated of air. What he found was that there began appearing strange arks of light that originated at the cathode and landing in the anode[163]. This effect is now understood, the free electrons in the gas are accelerated by the potential difference between the two electrodes, and they collide with either other atoms in the gas or the anode. These collisions produce more electrons as well as positive ions which continue this process (now called “gas discharge”). The net effect of this process was that the gas in the tube began to glow. At the time of Faradays experiment this effect wasn't understood, with subsequent experiments attempting to discover if the glow was caused by the gas in the tube or some force that was invisible to the human eye. To this end, the following decades saw scientists repeat this experiment with lower and lower pressures in the vacuum tube. This culminated in the 1870's when British physicist William Crookes evacuated a tube so completely that the internal pressure was below 10^{-6} atm. In doing so he almost completely removed the air from the tube, and the glow between the electrodes present in Faraday's experiment was no longer there. Instead, the glass on the anode end of the tube began glowing. Crookes would later modify the experiment, moving the anode beneath the glass tube and putting a metal barrier in the way of the beam. Like before, this device would cause the glass at the anode end of the tube to glow, but this time a section of the glass with the same shape as the metal barrier didn't glow. This showed that whatever the cathode was emitting behaved like light, it travelled in straight lines and cast “shadows”. Thus, “cathode rays” were discovered[164], although it wouldn't be until 1876 when Eugen Goldstein would prove they came from the cathode and named them “kathodenstrahlen” (German for “cathode rays”)[164]. Figure C.1 shows what a modern “Crookes tube” looks like and what happens when it's turned on.

However, unlike light, when a magnetic field is applied cathode rays bend in response to it, meaning that cathode rays carry electric charge. In 1886 German scientist Eugen Goldstein would discover “anode rays”[165] which carry a positive electric charge (the opposite of the cathode rays negative charge), simply by reversing the positions of the anode and cathode in a Crookes tube.



(a) Crookes Tube Apparatus - Off.

(b) Crookes Tube Apparatus - On.

Fig. C.1 The Crookes tube apparatus. The cathode is on the right end of the tube, with the anode at the base of the tube. When the tube is turned on, the green glow caused by the cathode rays is clear, as is the “shadow” cast by the iron cross.

As of the 1890’s the concepts of cathode and anode rays are still somewhat vague, since there isn’t really a good way to measure the properties that they possess. This would change in 1897 when English physicist Joseph John Thomson, along with colleagues John Townsend and Harold Wilson, measured the charge-to-mass ratio ($\frac{q}{m}$) of cathode rays[166]. In doing so, he showed that cathode ray beams contained exactly one kind of particle with a definite charge-to-mass ratio. The cathode rays were then renamed to reflect this fact. The name “electron” was proposed by Irish physicist George Johnstone Stoney, originally as a name for the fundamental unit of electric charge, but was used as a name for the particle itself and has become the accepted name for these particles since.

In the next few years a slew of new “rays” would be discovered. The α , β and γ rays were discovered in 1899 and 1900 by Ernest Rutherford (α and β rays) and Paul Villard (γ rays). Like so-called “cathode rays” these rays were named “ α ”, “ β ” and “ γ ” rays because the properties of the particles weren’t known at the time. The only thing that was known about them was their penetration power; α rays are stopped by a sheet of paper, β rays are stopped by a sheet of aluminium foil, γ rays can only be stopped by several meters of concrete or lead.

In time each of these particles would be studied in more detail and what they actually are would be uncovered. In 1900, Henri Becquerel would put the β ray under the same test that J. J. Thomson applied to the cathode ray[167]. He would show that β rays were indeed just electrons. Similarly, the α particles would be found to be

helium ions by Ernest Rutherford and Thomas Royds in 1907 the results of which they would publish the next year[168].

In 1905 Einstein would postulate the photon as an explanation for the photoelectric effect[169], although photons would never be “discovered” as a particle since photons were the particles used to demonstrate wave-particle duality in Young’s Double-Slit experiment. The γ rays that Villard discovered ultimately proved to be unaffected by the introduction of a magnetic field, showing that they had no electric charge. In 1914, γ rays were shown to be electromagnetic radiation and their wavelengths were measured soon afterward. Ultimately, the idea of a photon as a particle wouldn’t be “proven” as such, rather, the photon would be used as the example of wave-particle duality.

Up until this point the nature of atoms wasn’t well understood. It wasn’t until the famous Geiger-Marsden experiments (performed by Hans Geiger and Ernest Marsden under the supervision of Ernest Rutherford, also known as the Rutherford gold foil experiment) that atoms would be shown to have a tiny nucleus at the centre where all of the positive charge resided with electrons orbiting around it[170][171][172][173]. After the discovery of the nucleus it wouldn’t be long until the proton was discovered as well. The hydrogen atom had long been thought to be the building block of other atoms, but it wouldn’t be proven until Rutherford successfully demonstrated it by experiment in 1919 that the mechanism was nailed down. Rutherford noticed that when he fired α particles into nitrogen, his scintillators would pick up hydrogen nuclei as a product. He traced the production of these hydrogen nuclei to the nitrogen in the air, and repeated the experiment with pure nitrogen as a target for his α particles. In this case the effect was even larger, leading Rutherford to conclude that the hydrogen nucleus was present inside the nitrogen nucleus. He had in fact just shown the existence of the proton[162] as well as produced the first reported nuclear reaction: $^{14}\text{N} + \alpha \rightarrow ^{17}\text{O} + p$.

This discovery would lead to another problem, however. The properties that the nucleus had couldn’t be achieved if the nucleus contained *only* protons. Everything from the nuclear mass to its spin to the Heisenberg uncertainty principle meant that it was impossible for a nucleus to contain only protons. This culminated with an experiment by Walter Bothe and Herbert Becker in 1931[174][175]. They found that if α particles from polonium fell on beryllium, boron or lithium an exceptionally penetrating radiation was produced. Since it was unaffected by electric fields, they concluded it was γ radiation. The next year Irène and Frédéric Joliot-Curie showed that if this “ γ radiation” fell on paraffin it produced very high energy protons. The γ radiation hypothesis wasn’t popular, and in 1932 James Chadwick of the Cavendish

Laboratory would perform a series of experiment showing that this radiation was made up of particles with a similar mass to the proton[176][177]. These were the neutrons, and their discovery would win James Chadwick the Nobel Prize in 1935[178].

In 1928 Paul Dirac published a paper[179] introducing the “Dirac equation”. This equation fully modelled electron interactions (including the then-new concept of electron spin) while combining quantum mechanics with special relativity. This paper didn’t *explicitly* predict a new particle, but it did allow for electrons with either positive or negative energies as solutions. The negative energy solutions were very much an unintended consequence, and they puzzled the physics community. To explain it Dirac first argued that a “sea” of negative energy states existed throughout the whole universe, while others argued that this negative energy state was in-fact the proton. None of these hypotheses would satisfy, so in 1931 Dirac published another paper explicitly predicting a new particle, the “antielectron”, that had the same mass as the electron, but with the opposite charge and which would mutually annihilate on contact with an electron[180]. This new particle would be observed several times before being formally discovered. In 1929 a Soviet physicist named Dmitri Skobeltsyn would use a cloud chamber to try to detect γ radiation coming from cosmic rays, and he would find particles that had roughly the same mass as electrons but travelling in the opposite direction under a magnetic field. He was puzzled by these results and couldn’t explain them. At CalTech the same year, a graduate student named Chung-Yao Chao noticed a similar phenomenon, where a group of particles behaved like electrons but with a positive charge. However, his results were inconclusive and the phenomenon was not studied further. The final discovery would be made by Carl David Anderson a year later. His apparatus consisted of a cloud chamber surrounded by magnets. When he allowed cosmic rays to pass through the chamber, he discovered several ion trails that had the same charge-to-mass ratio as electrons, but curved in the opposite direction under his magnetic field. It curved in such a way that showed that it’s charge was positive. Thus, he had discovered a particle with the same charge-to-mass ratio as an electron, but with the opposite signed charge[181]. He had discovered the positron, and would receive a Nobel Prize for his work in 1936[182].

In 1930 the neutrino would be postulated by Wolfgang Pauli in an attempt to explain how β -decay could conserve energy, momentum and angular momentum/spin. However, it isn’t until after CERNs foundation that it would be proven to exist.

From the 1920’s all the way to the 1960’s the main tool of discovery in particle physics would cease to be the vacuum tube, instead the cloud chamber would be used.

The cloud chamber has already been used to discover the positron, and will be used to discover several new particles, including the first discovered mesons.

In 1936 Carl D. Anderson and one Seth Neddermeyer were studying cosmic radiation in CalTech when they made a peculiar observation. There were some particles that behaved the same way as electrons when they entered the cloud chamber, in that they seemed to have the same sign of charge, but they curved less sharply than electrons for particles of the same velocity. They measured the charge-to-mass ratio and found that it was smaller than that of electrons but larger than that of protons. The electric charge of these particles was assumed to be the same as that of an electron so it logically must have been heavier. Anderson initially called these particles “mesotrons”[183]. Their existence would be confirmed the next year by J. C. Street and E. C. Stevenson’s cloud chamber experiment[184]. A particle in the same mass range had been predicted by Hideki Yukawa in 1935, so physicists initially suspected that this new particle was the particle Yukawa predicted, called “Yukawa’s meson”[185], but in the end the mesotron was found to have the wrong properties to be Yukawa’s meson. Further, with the discoveries of the π [186] and K [187] mesons in 1947 the mesotron was renamed the μ meson, since it was thought to be a meson at the time. Eventually it would be confirmed that the μ meson is in-fact a lepton but that would also occur after the founding of CERN.

Thus, at the founding of CERN in 1952, the field of particle physics had identified eight particles: protons, neutrons, electrons, positrons, photons, μ mesons, π mesons and K mesons. The atom was understood to be composed of a central nucleus consisting of protons and neutrons with the electrons “orbiting” around it. The concept of an “antiparticle” was understood and the antiparticle of the electron, the positron, has already been discovered. The photon helped physicists understand the concept of wave-particle duality.

At this point there is still work to do. The μ meson will have to wait a few decades before being correctly assigned as the second generation lepton, and the π and K mesons are about to unleash the idea that mesons are composite particles, made of much smaller objects called “quarks”.

C.2 Foundation of CERN

After World War II the European nations found themselves no longer at the forefront of scientific discovery, with many scientists emigrating to other countries and much of the infrastructure destroyed. A handful of scientists across the various European countries

had the idea to create a pan-European atomic physics laboratory, a laboratory that scientists from across Europe could have access to while also sharing the substantial costs of operating nuclear facilities. A laboratory that was unlinked to any military goal that would promote unity among the European states after the cataclysmic world wars. The French physicist Louis de Broglie first put this idea forward during the European Cultural Conference, held in Lausanne on the 9th of December, 1949[188]. A further push came the next year when American physicist Isidor Rabi presented the idea at the fifth UNESCO General Conference, held in Florence in 1950[188]. The first official resolution concerning the establishment of such a laboratory occurred at a UNESCO Intergovernmental Meeting in Paris in December of 1952[188]. Two months after this 11 countries signed an agreement establishing a provisional council, the “European Council for Nuclear Research”. In French, “Conseil Européen pour la Recherche Nucléaire”, or “CERN”. This council would organise the creation of the laboratory, everything from the physical construction of the building and equipment to the organisation of the budget and the institutions that would oversee the laboratories operation. Governments wishing to host the laboratory were given the opportunity to submit proposals before the end of July that year. This council also oversaw the construction of the laboratories first experimental apparatus, a synchrocyclotron with an energy of about 500 MeV[189]. There were proposals submitted by the Danish, Dutch, French and Swiss governments for hosting the laboratory, but in the end Geneva’s central location in Europe, Switzerland’s neutrality during the war and the fact that it already hosted a number of international organisations solidified it as the choice to host the laboratory[190]. Over the next 18 months the CERN Convention would be written before being unanimously approved by the 11 countries who signed the original agreement as well as the UK[191]. After the sixth session of the CERN Council (taking place between 29 June - 1 July 1953) the convention establishing the organisation was signed and would slowly be ratified by the 12 founding states. The ratification was completed on the 29th of September, 1954 when France and Germany ratified it[31]. The provisional council was dissolved but the acronym remained[31].

C.3 Timeline of Modern Particle Physics and CERN Milestones

- 1955: The antiparticle of the proton, the “antiproton”, is discovered at the Bevatron particle accelerator at the University of California, Berkeley[192].

- 1956: A team of 5 physicists at Los Alamos confirm the existence of the (electron) neutrino[193].
- 1957: [CERN](#): The 600 MeV Synchrocyclotron (SC), CERN's first particle accelerator is constructed[194].
- 1959: [CERN](#): The Proton-Synchrotron (PS) starts operating. With an energy of 28 GeV, it was the most powerful accelerator in the world for a short time after starting up. It would replace the SC in providing proton beams for CERN's experiments, with the former accelerator being used for heavy ion physics[195].
- 1962: A team of 7 physicists at the Brookhaven National Laboratory proves the existence of 2 kinds of neutrinos, discovering what would become known as the "muon neutrino"[196].
- 1965: [CERN](#): The antiparticles corresponding to the three particles that make up atoms have all been observed as free particles. Since these three particles come together and form atoms, it is a natural hypothesis that there corresponding antiparticles would come together to form "antiatoms". This was eventually shown to be true when the antideuteron (a stable antiparticle made from an antiproton and an antineutron) was observed separately and simultaneously by two teams of physicists (one at CERN, led by Antonio Zichichi using the Proton Synchrotron and another at Brookhaven National Laboratory, led by Leon Lederman using the Alternating Gradient Synchrotron)[197].
- 1968: [CERN](#): Georges Charpak revolutionises detection. In the decades prior to his invention, the main apparatus used to discover new particles was the cloud chamber. As more and more particles were discovered, later searches had to try to discover particles that were produced by rarer and rarer processes. This meant that the cloud chamber method needed to be automated to some extent. This was done by setting up a camera to take a photo of the cloud chamber whenever it sensed a spark (something that is usually produced when a particle leaves a streak through the cloud chamber). A downside of this method was that it was cumbersome, with on the order of millions of photos being produced and needing to be examined by hand in order to make new particle discoveries. The invention of the "multi-wire proportional chamber" by Georges Charpak changed this. This chamber was a gas-filled box (much like a cloud chamber) which contained a large number of parallel detector wires, with each wire connected to an individual amplifier. Linked to a computer, it could achieve a counting rate 1000 times faster

than existing detectors. Charpak's invention moved experimental particle physics into the electronic age, and today practically every experiment uses technology based on his multi-wire proportional chamber[198].

- 1971: [CERN](#): The Intersecting Storage Rings perform their first proton-proton collisions. In the 1950's and 1960's physicists knew that colliding two proton beams head-on would produce much more energetic processes than firing a proton beam at a stationary target (which was the main method used at the time). To this end, the Intersecting Storage Rings (ISR) were created. The ISR would take two beams from the Proton-Synchrotron and keep them circulating around the ring in opposite directions until they were ready to collide[199]. This year would also see the approval of the construction of the Super Proton-Synchrotron (SPS)[200].
- 1974: [CERN](#): The tunnel for the Super Proton-Synchrotron is completed[201]. The J/Ψ particle is discovered, and would eventually be found to be made up of a charm quark and an charm antiquark[202].
- 1976: The Intersecting Storage Rings laid the foundation for larger particle accelerators to be built. The first of these would be the Super Proton-Synchrotron. Operating at an energy of 450 GeV and at a circumference of 7 km, the SPS would serve as the main accelerator used for CERN's physics program for roughly the next 15 years. An important note about it's construction is that the Proton-Synchrotron served as a sort-of "pre-accelerator" to the SPS. The Proton-Synchrotron would accelerate protons to an energy of 26 GeV before injecting them into the SPS where they're further accelerated to an energy of 450 GeV[203].
- 1977: Fermilab discovers the Upsilon (Υ) particle, confirming the existence of the bottom quark in the process[204]. Additionally, the τ -lepton is discovered this year after a series of experiments leading from 1974 to 1977[205][206].
- 1979: The PLUTO Detector at the DESY laboratory indirectly observes gluons[207].
- 1981: [CERN](#): The ISR produce the worlds first proton-antiproton collisions on April the 4th, 1981[208].
- 1983: [CERN](#): In 1979, CERN decided to convert the SPS into a proton-antiproton collider. In 1983, two experiments on the SPS, called "UA1" and "UA2", discovered the $W^\pm$ and Z^0 bosons after searching the debris of collisions for them[209].

- 1986: [CERN](#): Heavy ion collisions begin[210].
- 1988: [CERN](#): The tunnel for the next accelerator, the **L**arge **E**lectron-**P**ositron (“LEP”) Collider, is completed. At 27 km in circumference, it will house the equally large Large Electron-Positron Collider and later the Large Hadron Collider as well. The two ends of the ring meet each other with only 1 *cm* of error [211].
- 1989: [CERN](#): The Large Electron-Positron Collider collides its first proton beams. At 27 km in circumference, it was - and still is - the largest electron-positron accelerator ever built. During its first runs it accelerated electrons and positrons to 100 GeV, while later on in its operational history it accelerated them to a maximum of 208 GeV[212].
- 1990: [CERN](#): The worlds first browser/editor, website and server go live at CERN. By Christmas in 1990, Sir Tim Berners-Lee would invent the most basic building blocks of the internet: http, html and the URL. Additionally, he would also write the first internet browser/editor and server software. The worlds first web server: info.cern.ch, would be run on a NeXT computer at CERN, with the worlds first web page address providing information about the world wide web[213].
- 1995: Physicists at Fermilab discover the top quark[214][215]. Additionally, the first antiatoms are produced at CERN. Nine atoms of antihydrogen (a positron orbiting and antiproton) are produced, exist for about a 40 billionth of a second, travel a distance of 10 metres at nearly the speed of light before annihilating with regular matter[216].
- 1997: [CERN](#): With LEP reaching the end of its lifetime, a set of new experiments are designed with the goal of exploiting the newly designed Large Hadron Collider. This year, three of the four largest: ATLAS, CMS and ALICE would be approved for construction[217][218]. Additionally, with CERN’s antimatter machines (the Antiproton Accumulator, the Antiproton Collector and the Low Energy Antiproton Ring) being shut down to free up resources for the construction of the LHC, several scientists who wanted to continue studying antimatter applied for a new machine to be constructed that would allow them to do this. The result of these applications would be the Antiproton Decelerator, a machine which would also be approved that year[219].
- 1998: [CERN](#): The LHCb Experiment is also approved for construction[220].

- 2000: [CERN](#): The LEP Collider see its final shutdown[39]. Additionally, the τ neutrino is proven to be a distinct particle from the electron and μ neutrinos[221].

Appendix D

Jet Reconstruction with Sequential Recombination Algorithms

While sequential recombination algorithms have their roots in the e^+e^- experiments of the 1960s and 1970s [79], the lack of available computing power meant that they needed to incorporate complex techniques like preclustering in order to be usable on the computers of the time [79]. This meant that later algorithms would be much simpler due to access to more powerful computers. The first “simple” sequential recombination algorithm was the “JADE” algorithm, named for the JADE Collaboration that introduced it and used it in its e^+e^- experiments in the 1980s [222]. The JADE algorithm introduced the innovation that would define all sequential recombination algorithms: the use of a generalised distance measure between particles in the detector to group them into jets.

The first sequential recombination algorithm to be used in the context of hadron collisions would be the “ k_t algorithm”. The generalised distance measures are defined in equations D.1 and D.2.

$$d_{i,j} = \min(k_{T,i}^2, k_{T,j}^2) \frac{\Delta R_{i,j}^2}{R^2} \quad (\text{D.1})$$

$$d_{i,B} = k_{T,i}^2 \quad (\text{D.2})$$

where:

- y_i , ϕ_i and $k_{T,i}$ are the rapidity coordinate, azimuthal angle coordinate and transverse energy of the i^{th} particle.

- $\Delta R_{i,j}^2 = (y_i - y_j)^2 + (\phi_i - \phi_j)^2$, the distance in rapidity-azimuth space between particles i and j .
- $d_{i,j}$ is the “distance” between particles i and j in the detector.
- $d_{i,B}$ is the “distance” between the particle i and the beam, represented by B
- R is the radius of the cone defined by the jet. It’s value is chosen by the analysis team and usually takes the value of 0.4.

The k_t algorithm is then defined as follows:

1. Calculate $d_{i,j}$ and $d_{i,B}$ for all possible particles
2. Find the minimum of all the calculated
3. If the minimum is a $d_{i,j}$, then combine the i^{th} and j^{th} particles into a jet and repeat.
4. If the minimum is a $d_{i,B}$, then declare i to be a final state jet and remove it from the list of particles before continuing.
5. Repeat until no particles remain.

As a consequence of clustering particles the way it does, using distances in both position and momentum space, this algorithm has the following properties:

- The limits of what is defined as a jet is determined by R . If a particle i has no other particles within a distance of R , then $d_{i,B} > d_{i,j} \forall$ particles j [79].
- A corollary to the above, and because the k_t algorithm continues until all particles in an event are assigned to jets, is that arbitrarily soft particles that are isolated can be designated jets on their own. This is a problem because it can lead to events with many more reconstructed jets than there reasonably should be. A solution to this problem (and indeed what many analyses do) is to require that particles that go into reconstructing jets have a minimum transverse momentum.
- While the cones that define a jet have a limit in their size (radius R), the jets created by the k_t algorithm have very irregular shapes. This is shown more explicitly in Figure D.1a.

In chapter 5 was introduced the idea that any jet reconstruction algorithm must have one specific property: it must have the ability to form a jet around a single, hard isolated particle. In addition to these, it is desirable for jet reconstruction algorithms to have two more properties: infra-red safety and collinear safety.

When a parton emerges from a collision and enters the detector, it emits a variety of soft particles that land in a cone around it. These soft particles should then be reconstructed as part of the jet. The reason for this comes from the original motivation to use jets as physics objects: to develop a way to measure the energy/momentum of particles that hadronise without being dependent on *how* they hadronise, a process governed by complicated QCD interactions that is difficult to model perturbatively at best and impossible at worst. Therefore, when the addition of a soft particle doesn't change the amount of hard jets reconstructed by the jet reconstruction algorithm, that algorithm is said to be “infra-red safe”.

When a hard parton emerges from a collision it has a chance of decaying into two hard particles with similar energies. This is called “collinear splitting” and represents a problem for jet reconstruction algorithms, since their guiding philosophy is to reconstruct a jet around a single, hard particle. If a hard parton decays into two hard particles with similar energies and are close to each-other, then the algorithm might try to reconstruct them as two jets with very strange shapes, which is something to be avoided. As such, all jet reconstruction algorithms should be able to take a pair of hard particles close to each-other and reconstruct them as one jet. Algorithms that are able to do this are called “collinear safe”.

Ideally, jets reconstruction algorithms should be both infra-red and collinear safe. While the k_t algorithm fulfils both of these requirements, the irregular shape of the boundary of the jet means that while the number of hard jets reconstructed by the algorithm is unchanged by the addition of soft radiation, it does change the shape of the boundary. Algorithms other than the k_t algorithm also have this property, and those that do are said to have a “soft-adaptable” boundary.

The k_t algorithm was one of the first sequential recombination algorithms to be used by hadron collider experiments, and since its inception a few more have been introduced. These other algorithms came about as generalisations of the k_t algorithm, which materialised as changes to equations D.1 and D.2. The new measures are given in equations D.3 and D.4.

$$d_{i,j} = \min(k_{T,i}^{2p}, k_{T,j}^{2p}) \frac{\Delta R_{i,j}^2}{R^2} \quad (\text{D.3})$$

$$d_{i,B} = k_{T,i}^{2p} \quad (\text{D.4})$$

Choosing $p = 1$ immediately reduces the algorithm back to the k_t algorithm, and more generally $p > 0$ gives algorithms with similar behaviours to the k_t algorithm.

Alternatively, choosing $p = 0$ yields another sequential recombination algorithm called the “Cambridge-Aachen” algorithm [223].

The anti- k_t algorithm is created by taking the choice of $p = -1$, resulting in equations D.5 and D.6.

$$d_{i,j} = \min\left(\frac{1}{k_{T,i}^2}, \frac{1}{k_{T,j}^2}\right) \frac{\Delta R_{i,j}^2}{R^2} \quad (\text{D.5})$$

$$d_{i,B} = \frac{1}{k_{T,i}^2} \quad (\text{D.6})$$

Given this choice, it isn’t surprising that this algorithm has significantly different properties from the k_t algorithm. Consider a hard particle surrounded by many softer particles. Then $d_{1,i} = \min(1/k_{T,1}^2, k_{T,i}^2) \Delta R_{1,i}^2 / R^2$, the “distance” between hard particle “1” and soft particle “i” is exclusively dependent on the transverse momentum of the hard particle $k_{T,1}$ (since if “1” is harder than all of the soft particles “i”, ie $k_{T,1} > k_{T,i}$, then $1/k_{T,1}^2 < 1/k_{T,i}^2 \forall i$) and the $\Delta R_{1,i}$ separation. Conversely, the $d_{i,j}$ between similarly separated soft particles will be much larger. This leads to harder particles clustering soft particles to themselves before the soft particles cluster with each-other. If a hard particle has no similarly hard particles within a distance of $2R$, then it will simply cluster all of the softer particles with $\Delta R_{1,i} < R$ to itself and designate itself as a jet that is conical in shape.

If another hard particle “2” is present such that $R < \Delta R_{1,2} < 2R$ then two hard jets will be reconstructed, but neither will be conical. Instead, the degree to which each jets shape approaches a cone depends on the momentum of the hard particles. If $k_{T,1} \gg k_{T,2}$, then “1” will be conical while “2” will take a semi-crescent shape, defined by what would be it’s cone if “2” were isolated but with the any particle belonging to the cone of “1” removed. Alternatively, if $k_{T,1} = k_{T,2}$ then both particles will be partly conical, but with the boundary between them defined by a flat plane perpendicular to the line joining the two hard particles. In general if $k_{T,1} \approx k_{T,2}$ then the two cones won’t be perfectly conical and the boundary between them, b , will be curved in a way defined by $\Delta R_{1,b}/k_{T,1} = \Delta R_{2,b}/k_{T,2}$.

Finally, if two hard particles are close to each-other ($\Delta R_{1,2} < R$) then they will be reconstructed as one jet. If $k_{T,1} \gg k_{T,2}$ then the final jet will be roughly conical in shape, centred on “1”. If $k_{T,1} \approx k_{T,2}$ then the shape will be more complex, resembling a union of the two cones that “1” and “2” would define if they were isolated in the detector.

Knowing the properties of anti- k_t reconstructed jets, there are two key conclusions:

- anti- k_t jets are both infra-red and collinear safe, since it clusters soft particles to harder particles rather than hard particles to hard particles and it reconstructs two hard particles (like the ones that come from a collinear splitting) as a single jet.
- The anti- k_t algorithm produces jets that have shapes almost exclusively determined by the presence of hard particles, not soft ones. In this way, the anti- k_t algorithm is said to be “soft-resilient (in contrast with the “soft-adaptable” k_t algorithm).

These properties can be clearly seen in Figure D.1b.




Fig. D.1 A pair of 2-D plots comparing the way that the k_t and anti- k_t algorithms reconstruct the same event. The x, y -plane shows the distribution in rapidity-azimuthal angle space of cells in one layer of the calorimeter and the z -axis shows the energy in each cell. The cells are coloured based on which jet they are reconstructed as.

Appendix E

Systematic Uncertainty Contributions by Region

The systematic uncertainty contributions in each region are summarised in detail in this Appendix. The contributions are listed in the following tables:

- Table E.1: $\text{SR}2\ell_{\text{Low}}$
- Table E.2: $\text{SR}2\ell_{\text{Int}}$
- Table E.3: $\text{SR}2\ell_{\text{High}}$
- Table E.4: $\text{SR}2\ell_{\text{ISR}}$
- Table E.5: $\text{SR}3\ell_{\text{Low}}$
- Table E.6: $\text{SR}3\ell_{\text{Int}}$
- Table E.7: $\text{SR}3\ell_{\text{High}}$
- Table E.8: $\text{SR}3\ell_{\text{ISR}}$

In these tables, the contributions are listed in terms of their absolute, rather than relative, values. Further, the contributions needn't sum to the total systematic uncertainty in each region because of correlations between the sources.

Table E.1 A summary of the systematic uncertainty contributions in the SR2 ℓ _Low region.

Uncertainty	SR2 ℓ _Low
Total Background Expectation	8.38
Total Statistical Uncertainty ($\sqrt{N}$)	± 2.89
Total Background Systematic	± 5.83 [69.57%]
alpha_syst_zjets_LowMass_2L	± 5.77
gamma_stat_RJS2L_SR_LowMass_cuts_bin_0	± 2.57
alpha_syst_MM_2L_Low	± 0.54
alpha_syst_theoryVV_LowMass_2L	± 0.51
alpha_syst_jer_NF	± 0.35
mu_VV2L_RJS	± 0.21
Lumi	± 0.16
mu_TOP2L_RJS	± 0.15
alpha_syst_MET_SoftTrk_ResoPerp_NF	± 0.03
alpha_syst_MET_SoftTrk_ResoPara_NF	± 0.02
alpha_syst_ElecSF_ID_NF	± 0.01
alpha_syst_Muon_ID_NF	± 0.01
alpha_syst_jes_1	± 0.01
alpha_syst_jes_2_NF	± 0.01
alpha_syst_ElecIsoSF_NF	± 0.01
alpha_syst_btagC_NF	± 0.01

Table E.2 A summary of the systematic uncertainty contributions in the SR2 ℓ _Int region.

Uncertainty	SR2 ℓ _Int
Total Background Expectation	2.42
Total Statistical Uncertainty ($\sqrt{N}$)	± 1.56
Total Background Systematic	± 0.91 [37.60%]
alpha_syst_theoryVV_IntMass_2L	± 0.65
alpha_syst_zjets_IntMass_2L	± 0.37
mu_VV2L_RJS	± 0.33
gamma_stat_RJS2L_SR_IntMass_cuts_bin_0	± 0.27
alpha_syst_jer_NF	± 0.24
alpha_syst_MET_SoftTrk_ResoPerp_NF	± 0.12
alpha_syst_MET_SoftTrk_ResoPara_NF	± 0.10
alpha_syst_MM_2L_Int	± 0.06
alpha_syst_jes_3_NF	± 0.06
alpha_syst_jes_2_NF	± 0.04
alpha_syst_Muon_ID_NF	± 0.03
alpha_syst_Muon_MS_NF	± 0.01
alpha_syst_ElecSF_ID_NF	± 0.01
alpha_syst_jes_1	± 0.01
alpha_syst_EG_res_NF	± 0.01

Table E.3 A summary of the systematic uncertainty contributions in the SR2 ℓ _High region.

Uncertainty	SR2 ℓ _High
Total Background Expectation	1.85
Total Statistical Uncertainty ($\sqrt{N}$)	± 1.36
Total Background Systematic	± 0.78 [42.16%]
alpha_syst_theoryVV_HighMass_2L	± 0.51
alpha_syst_zjets_HighMass_2L	± 0.38
gamma_stat_RJS2L_SR_HighMass_cuts_bin_0	± 0.30
mu_VV2L_RJS	± 0.25
alpha_syst_jer_NF	± 0.10
alpha_syst_MET_SoftTrk_ResoPerp_NF	± 0.07
alpha_syst_MET_SoftTrk_ResoPara_NF	± 0.05
alpha_syst_jes_3_NF	± 0.02
alpha_syst_EG_scale_NF	± 0.01
alpha_syst_Muon_Scale_NF	± 0.01
alpha_syst_Muon_ID_NF	± 0.01
alpha_syst_Muon_MS_NF	± 0.01
alpha_syst_ElecSF_ID_NF	± 0.01
alpha_syst_btagB	± 0.00

Table E.4 A summary of the systematic uncertainty contributions in the SR2 ℓ _ISR region.

Uncertainty	SR2 ℓ _ISR
Total Background Expectation	2.71
Total Statistical Uncertainty ($\sqrt{N}$)	± 1.64
Total Background Systematic	± 2.79 [102.95%]
alpha_syst_zjets_Compressed_2L	± 2.58
alpha_syst_theoryVV_Compressed_2L	± 0.96
mu_VV2L_RJC	± 0.44
gamma_stat_RJC2L_SR_Compressed_cuts_bin_0	± 0.25
alpha_syst_MM_2L_Compressed	± 0.17
alpha_syst_jer_NF	± 0.10
alpha_syst_jes_2_NF	± 0.08
alpha_syst_btagB_NF	± 0.05
alpha_syst_jes_3_NF	± 0.04
mu_TOP2L_RJC	± 0.04
alpha_syst_MET_SoftTrk_ResoPerp_NF	± 0.03
alpha_syst_MET_SoftTrk_ResoPara_NF	± 0.02
alpha_syst_EG_scale_NF	± 0.01
Lumi	± 0.01
alpha_syst_btagL_NF	± 0.01
alpha_syst_ElecSF_ID_NF	± 0.01
alpha_syst_Muon_ID_NF	± 0.01
alpha_syst_btagB	± 0.01
alpha_syst_ElecIsoSF_NF	± 0.01
alpha_syst_Muon_MS_NF	± 0.01

Table E.5 A summary of the systematic uncertainty contributions in the SR3 ℓ _Low region.

Uncertainty	SR3 ℓ _Low
Total Background Expectation	10.31
Total Statistical Uncertainty ($\sqrt{N}$)	± 3.21
Total Background Systematic	± 1.96 [19.01%]
alpha_syst_theoryVV_LowMass_3L	± 1.25
mu_VV3L_RJS	± 0.91
gamma_stat_RJS3L_SR_LowMass_cuts_bin_0	± 0.85
alpha_syst_jer_NF	± 0.74
alpha_syst_jes_2_NF	± 0.25
alpha_syst_EG_scale_NF	± 0.19
alpha_syst_MM_3L_Low	± 0.17
alpha_syst_EG_res_NF	± 0.14
alpha_syst_MET_SoftTrk_ResoPerp_NF	± 0.10
alpha_syst_MET_SoftTrk_ResoPara_NF	± 0.07
alpha_syst_Muon_MS_NF	± 0.06
alpha_syst_jes_3_NF	± 0.05
alpha_syst_jet_jvt_NF	± 0.04
alpha_syst_Muon_ID_NF	± 0.04
alpha_syst_ElecIsoSF_NF	± 0.03
alpha_syst_btagL_NF	± 0.03
alpha_syst_btagC_NF	± 0.02
alpha_syst_ElecSF_ID_NF	± 0.02
Lumi	± 0.01

Table E.6 A summary of the systematic uncertainty contributions in the SR3 ℓ _Int region.

Uncertainty	SR3 ℓ _Int
Total Background Expectation	2.28
Total Statistical Uncertainty ($\sqrt{N}$)	± 1.51
Total Background Systematic	± 0.50 [21.93%]
gamma_stat_RJS3L_SR_IntMass_cuts_bin_0	± 0.39
alpha_syst_theoryVV_IntMass_3L	± 0.20
mu_VV3L_RJS	± 0.17
alpha_syst_EG_scale_NF	± 0.10
alpha_syst_Muon_MS_NF	± 0.10
alpha_syst_Muon_ID_NF	± 0.07
alpha_syst_EG_res_NF	± 0.07
alpha_syst_jes_2_NF	± 0.02
alpha_syst_jer_NF	± 0.02
alpha_syst_jes_3_NF	± 0.02
alpha_syst_ElecIsoSF_NF	± 0.02
alpha_syst_ElecSF_ID_NF	± 0.01
Lumi	± 0.01
alpha_syst_jet_jvt_NF	± 0.01
alpha_syst_MM_3L_Int	± 0.01

Table E.7 A summary of the systematic uncertainty contributions in the SR3 ℓ _High region.

Uncertainty	SR3 ℓ _High
Total Background Expectation	1.05
Total Statistical Uncertainty ($\sqrt{N}$)	± 1.03
Total Background Systematic	± 0.46 [43.81%]
gamma_stat_RJS3L_SR_HighMass_cuts_bin_0	± 0.39
alpha_syst_theoryVV_HighMass_3L	± 0.19
mu_VV3L_RJS	± 0.08
alpha_syst_jes_2_NF	± 0.07
alpha_syst_EG_res_NF	± 0.06
alpha_syst_jes_3_NF	± 0.06
alpha_syst_jer_NF	± 0.04
alpha_syst_MM_3L_High	± 0.04
alpha_syst_Muon_MS_NF	± 0.03
alpha_syst_Muon_ID_NF	± 0.02
alpha_syst_MET_SoftTrk_ResoPerp_NF	± 0.02
alpha_syst_MET_SoftTrk_ResoPara_NF	± 0.01
Lumi	± 0.01
alpha_syst_ElecSF_ID_NF	± 0.01

Table E.8 A summary of the systematic uncertainty contributions in the SR3 ℓ _ISR region.

Uncertainty	SR3 ℓ _ISR
Total Background Expectation	3.95
Total Statistical Uncertainty ($\sqrt{N}$)	± 1.99
Total Background Systematic	± 1.02 [25.82%]
alpha_syst_theoryVV_Compressed_3L	± 0.76
mu_VV3L_RJC	± 0.42
gamma_stat_RJC3L_SR_Compressed_cuts_bin_0	± 0.41
alpha_syst_MET_SoftTrk_ResoPerp_NF	± 0.16
alpha_syst_MET_SoftTrk_ResoPara_NF	± 0.14
alpha_syst_jer_NF	± 0.13
alpha_syst_Muon_MS_NF	± 0.10
alpha_syst_jes_3_NF	± 0.10
alpha_syst_EG_scale_NF	± 0.09
alpha_syst_MM_3L_Compressed	± 0.08
alpha_syst_jes_2_NF	± 0.07
alpha_syst_Muon_ID_NF	± 0.06
alpha_syst_EG_res_NF	± 0.02
alpha_syst_jet_jvt_NF	± 0.02
alpha_syst_btagC_NF	± 0.01
alpha_syst_btagL_NF	± 0.01
alpha_syst_Muon_Scale	± 0.01
alpha_syst_btagB	± 0.01

