



Search for new hadronically decaying resonances with masses down to 20 GeV in pp collisions using the ATLAS Detector at the Large Hadron Collider

Lars Henkelmann
of
Clare College, Cambridge

September 2023



This thesis is submitted for the degree of Doctor of Philosophy

Declaration

This thesis is the result of my own work and includes nothing which is the outcome of work done in collaboration except as declared in the preface and specified in the text.

It is not substantially the same as any work that has already been submitted before for any degree or other qualification except as declared in the preface and specified in the text.

It does not exceed the prescribed word limit for the Physics and Chemistry Degree Committee.

Search for new hadronically decaying resonances with masses down to 20 GeV in pp collisions using the ATLAS Detector at the Large Hadron Collider

Lars Henkelmann

Many proposed models of physics beyond the Standard Model of particle physics (BSM) predict new resonances. Thus, a single experimental search for such new resonances can probe a multitude of BSM models. In particular, searches for resonances decaying into hadrons can probe for any resonances that can be produced at the LHC. However, such dijet searches lose sensitivity at low masses due to transverse momentum trigger thresholds.

This thesis presents a search targeting events where the BSM resonance is produced back-to-back with a photon radiated from the initial state partons. The search achieves sensitivity to BSM resonance masses as low as 20 GeV for the first time using data from the ATLAS detector. This is done by triggering on the initial state radiation (ISR), which decouples the accessible resonance mass from the trigger thresholds. Trigger thresholds instead imply a minimum boost of the resonance, causing the lowest-mass resonances to produce overlapping jets that are best reconstructed as a single large-radius jet.

Precise large-radius jet reconstruction down to very low masses is thus a crucial requirement of the search. However, the standard Run 2 large-radius jet definition is only calibrated down to 50 GeV. To overcome this limitation, in this thesis the resonance candidate jets are instead reconstructed by reclustering calibrated $R = 0.2$ hadronic anti- k_T jets and inner detector tracks into so-called track-assisted reclustered (TAR) jets.

A dataset corresponding to 140.1 fb^{-1} of proton-proton collisions at $\sqrt{s} = 13 \text{ TeV}$ recorded with the ATLAS detector from 2015 through 2018 is analysed. The most significant excess over the SM expectation is localized near 40 GeV of jet mass, but at 2.8σ it is not significant. Thus, upper limits are set on the cross-section for producing Z' vector resonances. The 95 % CL_s limits range from 0.02 pb for a 20 GeV resonance to 0.2 pb for a 225 GeV resonance.

Preface

Information is the resolution of uncertainty.

attributed to Claude Shannon

The joy of research, to me, lies in the kind of surprise one experiences when one tests an assumption and the test fails in an unexpected way. This experience is surely very common. Indeed, from the perspective of information theory, the *information content* of any piece of data is measured by how unlikely it is to occur, in other words, how much one is surprised to encounter it.

This naive joy of surprise is why I decided to pursue experimental research, and it also motivates the kind of research I present in this thesis. The realm of very low jet masses has not so far been investigated much in ATLAS, and thus one may hope to be surprised by what one finds. The use of a novel reconstruction method was of course necessitated by choosing to study jets that are lighter than the usual calibration admits. However, personally the novelty of the method attracted me with the promise of many small surprises along the way, and I am pleased to say, it has fulfilled that promise. Similarly, I hope that should you be so kind as to read some of my thesis, you will encounter a surprise or two as you do, indicating information has been transmitted.

Author's contributions

The work presented here was carried out in the context of the ATLAS collaboration. Thus, it relies heavily on code, calibrations, computing and data processing facilities, datasets, simulations etc. provided by members of the collaboration for members of the collaboration. The major contributions of the author are described in the following.

The author contributed to the calibration of hadronic $R = 0.2$ anti- k_T jets used in the search. (He also contributed to the calibration of $R = 0.6$ anti- k_T jets, but these are not used in this thesis.) Specifically, the author contributed the residual pile-up correction calibration and associated systematic uncertainties for $R = 0.2$ and $R = 0.6$. A second contribution is the statistical combination of in-situ jet energy scale calibrations using different observed final states across different energy scales into a combined calibration continuously parameterized in the η, p_T plane of the calibrated $R = 0.2$, $R = 0.6$ jets. Finally, the author contributed a study evaluating the dependence of the accuracy of the energy scale calibration of $R = 0.2$ jets on the presence of close-by dense hadronic radiation. These contributions to jet calibration are used in this thesis to calibrate $R = 0.2$ jets (see Section 5.2.2), but in the interest of brevity they are not themselves documented.

The $R = 0.2$ and $R = 0.6$ jet calibrations are described in more detail in a paper being prepared for publication at the time of writing, and are used by several teams analysing ATLAS data.

The author is the main analyser and analysis contact for the search for low-mass hadronic resonances presented in this thesis. On the suggestion of his advisor Oleg Brandt, the author synthesized the analysis strategy based on existing approaches (cf. the discussions in Section 2.5 and Chapter 6). He adapted the analysis pre-selections (see Section 5.2.4) for use with lower-mass TAR jets based on great work done by Roy S. Brener* for higher-mass calorimeter jets [1]. The simulation of hypothetical signal resonance decays down to very low masses was investigated by Chris M. Delitzsch†. The remaining background simulations used in this thesis are commonly used across ATLAS, and were provided by the Physics Modelling Group of the ATLAS Collaboration. The author defined the selections for the signal- and control regions used in the analysis, including the kinematic decorrelation of the D_2 observable. The determination of the resonant D_2 scale factor using data observed in the top control region was done by Jad M. Sardain‡. The author designed and implemented the likelihood model described in Chapter 7, and carried the validation and characterization studies thereof, except for the envelope-smoothing of the spurious signal uncertainty, for which he is grateful to Davide Melini§. The author validated the likelihood model in-situ as described in Section 8.1. The studies done in the course of unblinding on the full Run 2 dataset were preformed by the author and Jad M. Sardain in close collaboration. The unblinded results presented here are also part of a separate publication currently in preparation.

The brief description of the ATLAS detector at the LHC in Chapter 3 and the overview of the Standard Model of Particle Physics in Chapter 1 are edited and adapted from the authors previous such descriptions in [2].

Acknowledgements

I would like to express my deepest appreciation to my examiners, Caterina Doglioni and Chris Lester, for taking the time to carefully and critically read my thesis.

I extend my deep gratitude and sincere appreciation to my Supervisor, Oleg Brandt. His unwavering support, guidance, and advice have been invaluable throughout my journey. Always alert to new opportunities, his meticulous feedback and assistance in navigating the — at times intricate — landscape of collaboration politics and communication were truly indispensable.

* Technion (Israel), now Weizmann Institute (Israel)

† T.U. Dortmund (Germany)

‡ U. Arizona (United States)

§ Technion (Israel)

I am sincerely grateful to the members of the wider (boosted- and resolved-) dijet + ISR (-jet or -photon) team, with special mention to Jad M. Sardain, Jennifer Roloff, Roy S. Brener, Davide Melini, and Chris M. Delitzsch. Our collaboration has been incredibly productive, and their abundant advice and feedback have been instrumental in shaping the course of my work.

I offer my heartfelt thanks to all the numerous members of the “TARscan” team, with a special mention to Jona Bossio, Jason Veatch, and Bogdan Malaescu. The TAR and R-scan projects have been exceptionally productive, especially considering the demanding long-term nature of the effort. Both TAR reclustering and the $R = 0.2$ jet calibration are integral to the results presented in this work, and would not be possible without the sustained hard work of all the team. Furthermore, their valuable advice and inspiring examples of good science have been of immeasurable significance to my own development as a researcher.

I want to convey my heartfelt appreciation and special gratitude to two remarkable Cambridge Part III students, Dennis Lindebaum and Gustavs V. Kehris. It was a genuine delight to help supervise their final-year projects linked to my research. I learned a great deal from their perspectives. Although time constraints regrettably prevented me from integrating Dennis’ exceptional work with adversarial neural architectures into the final result described in this thesis, it was fascinating to witness its potential. Gustavs’ efforts resulted in our first independent in-situ validation of the TAR energy scale, providing a substantial boost in confidence regarding the validity of $R=0.2$ calibration in the TAR context.

Special thanks go to all the members of the Cavendish HEP group, including David Rouso, Dominic Jones, Fionn C. R. Bishop, Ben H. Hodkinson, Rupert Tombs, Herschel A. Chawdhry, Andrei Popescu, Holly Pacey, Jon Burr, Sebastian R. Colmenares, Anna J. Mullin, Richard M. Williams, Chris Lester, William Fawcett, Bill K. Balunas, and Tina Potter. Their friendly and welcoming environment provided much inspiration and support. I will fondly remember numerous stimulating discussions over tea and coffee.

I greatly appreciate the powerful computing cluster provided by and to the Cavendish HEP group, which has been impeccably maintained by the incomparable Steve Wotton. This infrastructure has been utterly indispensable, since without it the research presented in this thesis would not have been achievable.

During the initial months of my thesis project, I was with the ATLAS group at the Kirchhoff Institut für Physik in Heidelberg. I am very grateful to the group, among them Falk Bartels, Sebastian Weber, Fabrizio Napolitano, Merve Sahinsoy Thomas M. Spieker, Claire Antel, Monica Dunford, and Pavel Starovoitov. Their creation of a wonderful, friendly, and supportive environment has been truly inspiring, fostering curiosity and providing unwavering encouragement. They have set remarkable examples, and I am thankful for the privilege of utilizing their exceptional computing cluster.

I am thankful for the innumerable ways in which the ATLAS Collaboration enabled and facilitated this work, providing resources and support across many areas, such as computing, data-taking operations, and calibration, to name just a few. I would similarly like to express my appreciation to the LHC collaboration, whose outstanding work at the forefront of accelerator physics is a prerequisite to every result presented in the following.

I wish to extend my sincere gratitude to the Science and Technology Funding Council of UK Research and Innovation, fetch.ai, the Centre for Doctoral Training in Data Intensive Science in Cambridge (UK), and the International Max Planck Research School for Precision Tests of Fundamental Symmetries in Heidelberg for their invaluable financial support, which has been pivotal to the realization of this work.

I would also like to express my heartfelt thanks to the wonderful team at OKRA.ai. My industry placement with them has been both exhilarating and invigorating, providing me with a fresh perspective and a truly enjoyable experience.

I am thankful to the many kind people of Clare College for providing an environment that has been both entertaining and joyous.

I wish to express my heartfelt appreciation to Tom, Alex, Sascha, Anna, and Jan, for their unwavering support and friendship throughout this journey. Your presence has provided a much-needed balance to my life. To my family, especially Christine and Joachim, your enduring encouragement has been my rock, and I extend special thanks for your constant belief in and support of me. I am deeply grateful to Tania. Your patience and support, even in the face of way too many late nights stolen by debugging, has been invaluable. Your remarkable ability to divert my mind from Physics when I needed it most has been a true gift during challenging times. Ni mitz tlazotla, no moyolotzin.

Lars Henkelmann, Cambridge (UK), Thursday 28th September, 2023

Contents

Declaration	iii
Abstract	v
Preface	vii
Contents	xi
TRIVIA	1
1 The Standard Model of Particle Physics	3
1.1 Field content	3
1.2 Kinematics	4
1.3 Local Gauge Symmetries	5
1.3.1 Quantum Chromodynamics	5
1.3.2 The Electroweak Interaction	7
1.4 Higgs Mechanism	9
1.4.1 Scalar Sector Lagrangian	10
1.4.2 Spontaneous Electroweak Symmetry Breaking	11
2 Physics beyond the Standard Model	15
2.1 Motivating new physics	15
2.2 Dark Matter	16
2.2.1 Theory of WIMP Dark Matter	16
2.2.2 Dark Matter Searches	17
2.3 Constraints on low-mass leptophobic Z' models	21
2.4 Unlocking the lowest resonance masses at the LHC	22
2.5 Prior searches for very-low-mass hadronic resonances	23
3 The ATLAS Detector at the LHC	27
3.1 The Large Hadron Collider	27
3.2 The ATLAS Detector	28
3.2.1 Magnets	30
3.2.2 Inner Detector	30
3.2.3 Calorimeters	32
3.2.4 Muon System	35

3.2.5	Trigger	36
3.3	Luminosity	37
ANALYSIS		41
4	Analysis overview	43
5	Samples and selections	45
5.1	Samples	45
5.1.1	Collision data samples	45
5.1.2	Simulated samples	46
5.2	Selections	50
5.2.1	Photon selection	50
5.2.2	Jet reconstruction and calibration	51
5.2.3	Jet selection	54
5.2.4	Event selection	55
6	Background estimate	57
6.1	Strategy	57
6.1.1	Background processes	57
6.1.2	Data-driven estimate of non-resonant background	58
6.1.3	Simultaneous subtraction of resonant background	59
6.2	Fixed-efficiency cut method for a mass-decorrelated two-prong tag	60
6.2.1	Measuring jet mass shaping using the Jensen-Shannon distance	60
6.2.2	Fixed efficiency cut method	62
6.2.3	Measured quantiles and fixed-efficiency cut maps	64
6.2.4	Mass shaping of decorrelated taggers	64
6.2.5	Summary of tagger performance for different observables	66
6.3	Forward photon control region	67
6.3.1	Candidate observables	68
6.3.2	Purity and contamination of control region	70
6.3.3	Bias on jet kinematics in control region	71
6.3.4	Choice of control region selection	73
6.4	Top control region	74
6.4.1	Overview	74
6.4.2	Object reconstruction and selections	74
6.4.3	Jet mass distributions	76
6.4.4	D_2^{DDT} scale factor for two-prong resonances	78
7	Fit studies	79
7.1	Likelihood construction	79
7.1.1	Transfer factor parameterization	81
7.1.2	Modelling systematic uncertainties	82

7.2	Smoothed template of QCD multijet background jet mass spectrum	87
7.2.1	Smoothing method overview	88
7.2.2	Smoothing the QCD multijet template	90
7.2.3	Smoothing the pass-to-fail ratio	91
7.2.4	Smoothing the double ratio	91
7.2.5	Applicability	93
7.3	Choice of transfer factor polynomial degree	94
7.4	Spurious signal	96
7.5	Linearity of fit response	97
7.6	Impact of systematic uncertainties	99
7.7	Expected exclusion estimate using pseudodata	101
7.7.1	The asymptotic approximation for CL_s exclusion intervals	102
7.7.2	Expected exclusion intervals	105
7.7.3	Limitations of interpreting the 10 GeV signal extraction fit	106
8	Results	109
8.1	In-situ validation results	109
8.1.1	Validation in 2015 data	110
8.1.2	Validation in data from the 2018 L4 period	111
8.2	15.8 fb ⁻¹ unblinding results	112
8.3	Fully unblinded results	113
9	Conclusions	121
9.1	Summary	121
9.2	Outlook	123
9.2.1	Extended datasets	123
9.2.2	Extended methodology	125
	APPENDIX	127
A	The Kolmogorov-Smirnov test and measuring residual mass shaping	129
B	Auxiliary material on the smoothed multijet background template	131
B.1	Further illustration of excessive MC noise in simulated samples	131
B.2	Further KDE smoothing variants and examples	133
C	Auxiliary material about fit linearity	137
C.1	Injection up to higher values of μ	138
D	Auxiliary material on systematic uncertainties	141
E	Auxiliary material on the in-situ validation results	145
E.1	Validation in 2015 data	145
E.2	Validation in data from the 2018 L4 period	150

F	Auxiliary material on unblinding 15.8 fb^{-1} of observed data	155
F.1	Limited fit of the transfer factor polynomial to only the forward photon CRs	155
F.1.1	Background-only maximum-likelihood fit in all regions	157
F.2	Example signal-extraction maximum likelihood fit	160
F.3	CLs limits	162
F.3.1	Comparison to ATLAS and CMS results	162
G	Auxiliary material on the unblinded analysis of 2015-2018 data	165
	Bibliography	167

List of Figures

1.1	The particles of the Standard Model of elementary particles.	4
1.2	Interactions of the elementary particles in the Standard Mode	5
1.3	Illustration of the shape of the SM Higgs potential.	11
2.1	Relation of WIMP freeze-out to Dark Matter searches at colliders by CPT symmetry . .	16
2.2	Topologies targeted by LHC searches for Dark Matter	18
2.3	Summary of ATLAS searches for Dark Matter	20
3.1	The LHC and its pre-accelerators.	28
3.2	The components of the ATLAS Detector	29
3.3	The ATLAS magnet system.	30
3.4	The ATLAS inner detector	31
3.5	The ATLAS calorimeters	33
3.6	Cutaway view of the ATLAS muon System	35
5.1	Reconstructed mass of the resonance candidate jet in benchmark signal samples. . . .	49
5.2	Narrow $R = 0.2$ jet energy scale calibration stages.	52
5.3	TAR jet reconstruction algorithm.	53
6.1	Feynman graphs for non-resonant background.	57
6.2	Feynman graphs for resonant backgrounds.	57
6.3	The relation of the signal region and the three main control regions.	58
6.4	Distortion of the jet mass spectrum due to two-prong tagging.	61
6.5	Smoothed map of D_2 cuts at a fixed non-resonant background efficiency.	65
6.6	Distortion of the jet mass spectrum due to a D_2^{DDT} selection in simulated samples. . . .	66
6.7	Distributions of candidate observables for constructing the forward-photon CR	69
6.8	Control region compositions resulting from different candidate observables.	70
6.9	Shaping of jet kinematic distribution by cuts on CR candidate observables.	72
6.10	Semi-leptonic $t\bar{t}$ topology targeted by the top CR.	74
6.11	TAR jet mass distribution in top CR	76
6.12	TAR jet mass distribution in top CR split by D_2^{DDT}	77
7.1	Structure of unconstrained nuisance parameters for the non-resonant background. . .	80
7.2	Degree 5 Bernstein basis polynomials.	81
7.3	Effect of the JES NP JET_EffectiveNP_Modelling1 on the $Z^0\gamma$ background template	83
7.4	Effect of the JER NP JET_JER_EffectiveNP_2 on the $W^\pm\gamma$ background template . . .	84
7.5	Effect of the PDF eigenvariations (ME_PDF260000_var) on the $m'_Z = 20$ GeV signal template	85
7.6	Post-fit mass spectra including non-resonant QCD background simulations.	87
7.7	Toy example of using a KDE to smooth a histogram sampled from a χ^2 -distribution. .	89
7.8	Examples of using a KDE (bandwidth 0.16) to smooth a multijet mass histogram. . . .	90

7.9	Pseudodata and samples in the analysis regions, using the smooth QCD multijet template.	91
7.10	Transfer function polynomials extracted from fits to the smooth QCD multijet template.	92
7.11	Post-fit residuals of fits to the Asimov dataset using basis polynomials of varying degree.	95
7.12	Summary of iterative scans over alternative transfer factor polynomial degree choices.	95
7.13	Distributions of best-fit signal strength $\hat{\mu}$ of fits to pseudodata toys.	96
7.14	Linearity of the fit response to varying injected signal strength μ .	98
7.15	NPs ranked by impact on the POI.	100
7.16	Signal strength scan used to determine the upper end of the one-sided 95% CL_s interval.	104
7.17	95% CL_s limits on the Z' production cross-section	106
7.18	Level scheme of bound $b\bar{b}$ systems.	108
8.1	Mass spectra observed in the 2015 data compared to the expected post-fit distribution.	110
8.2	Mass spectra observed in the 2018 L4 data compared to the post-fit expectation.	111
8.3	Post-fit mass spectra of the signal extraction fit.	112
8.4	95% CL_s limits on Z' production for a range of Z' resonance masses.	113
8.5	Pass-to-fail ratios and transfer function polynomials observed in 2015-2018 data.	114
8.6	Post-fit mass spectra of the background-only fit.	115
8.8	95% CL_s limits on Z' production for a range of Z' resonance masses.	117
8.7	Iterative likelihood-ratio tests for different transfer function polynomial degrees.	117
8.9	Post-fit mass spectra of the $m_{Z'} = 40$ GeV signal extraction fit.	118
B.1	Pre-fit mass spectra including non-resonant QCD background simulations.	132
B.2	Comparison of pre- and postfit <code>shapefactor</code> values of fits to the un-smoothed pseudodata.	132
B.3	Toy example of using a KDE to smooth a histogram sampled from a normal distribution.	134
B.4	Examples of using a KDE (bandwidth 0.04) to smooth a multijet mass histogram.	134
B.5	Examples of using a KDE (bandwidth 0.12) to smooth a multijet mass histogram.	135
B.6	Examples of using a KDE (bandwidth 0.2) to smooth a multijet mass histogram.	135
C.1	Linearity of the fit response to varying injected signal strength (Boxplot).	137
C.2	Linearity of the fit response to injected signal strengths $\mu \geq 1$.	138
C.3	Linearity of the fit response to varying injected signal strength $\mu \geq 1$ (Boxplot).	139
D.1	Effect of the JES NP <code>JET_EffectiveNP_Modelling1</code> on the $m'_Z = 20$ GeV signal template	142
D.2	Effect of the photon energy scale NP <code>EG_SCALE_ALL</code> on the $Z^0\gamma$ background template	143
E.1	Mass spectra observed in the 2015 data compared to the expected pre-fit distribution.	145
E.2	Pulls of NPs with prior constraints (except MC statistics).	146
E.3	Pulls of MC statistical uncertainty NPs.	147
E.4	Change in the floating shapefactor NPs relative to their initial value in the fit.	148
E.5	Transfer function polynomials extracted from fits to the 2015 data validation sample.	149
E.6	Transfer polynomial coefficients extracted from fits to the 2015 data.	149
E.7	Mass spectra observed in the 2018 L4 data compared to the pre-fit expectation.	150
E.8	Pulls of NPs with prior constraints (except MC statistics).	151
E.9	Pulls of MC statistical uncertainty NPs.	152
E.10	Change in the floating shapefactor NPs relative to their initial value in the fit.	153

E.11	Transfer polynomials extracted from fits to the 2018 L4 data sample.	154
E.12	Transfer polynomial coefficients extracted from fits to the 2018 L4 data sample.	154
F.1	Transfer function polynomials extracted from fits to the 15.8 fb ⁻¹ sample.	156
F.2	Pre-fit mass spectra	156
F.3	Post-fit mass spectra of the background-only fit.	158
F.4	Pulls of NPs with prior constraints in the background-only fit.	159
F.5	Pulls of NPs with prior constraints.	161
F.6	Signal strength scans used to determine the upper end of the one-sided 95% CL _s interval.	162
F.7	Comparison of the 95% CL _s limits on the Z' to prior ATLAS and CMS results.	163
G.1	Pass-to-fail ratios and transfer function polynomials observed in 2015-2018 data.	165
G.2	95 % CL _s exclusion limits on the Z' -quark coupling strength	166
G.3	Local p-values of the observed data given the background-only hypothesis.	166

List of Tables

5.1	Triggers used to record top CR data.	45
5.2	Exact versions of common tools used for $t\bar{t}$ and single-top simulation.	47
5.3	Inclusive single-top production cross-section predictions.	48
5.4	Exact versions of common tools used for V +jets simulation.	48
5.5	Z' masses included in the grid of simulated signal templates.	49
5.6	Photon selection criteria.	50
5.7	Narrow jet reconstruction.	52
5.8	Track pre-selection criteria.	53
5.9	Large- R jet reconstruction.	53
5.10	Large- R jet selection criteria.	54
5.11	Narrow jet ($R = 0.2$ anti- k_T) selection criteria.	54
5.12	Event selection criteria.	55
6.1	Selections defining the signal region and control regions.	58
6.2	Comparison of residual mass shaping by various taggers after mass-decorrelation.	67
6.3	On-peak $W^\pm$ identification efficiency of various taggers after mass-decorrelation	67
6.4	Muon selection criteria.	74
6.5	$R = 0.4$ anti- k_T EMPFlow jet selection.	74
6.6	b -tagging selection criteria.	75
6.7	$R = 1$ jet selections applied in the top CR	75
6.8	Event selections for the top control region.	75
7.1	Mass bin edges in GeV.	79
7.2	Design of constrained nuisance parameters modelling systematic uncertainties.	82
7.3	Uncertainty on the best-fit signal strength $\hat{\mu}$ by category.	101

8.1 Measured integrated luminosity, by year of data taking.	114
---	-----

TRIVIA

The Standard Model of Particle Physics

1

This chapter gives a short overview of the Standard Model of Particle Physics, based on the descriptions in refs. [3, 4], adapting a similar prior summary in ref. [2].

Particle physics is concerned with the study of the fundamental particles and forces¹ of nature. The current understanding of those particles and forces is summarized in one theory, called the Standard Model of particle physics (SM). Section 1.1 summarizes the fundamental particles described by the SM. The SM describes the dynamics of these fundamental particles using Quantum Field Theory (QFT). Therefore, the SM describes the kinematics and spins of spin 1/2 and spin 1 or spin 0 particles using the Dirac and Klein-Gordon equations, respectively, as described in Section 1.2. The SM describes the interactions of these particles as arising from a local gauge principle, as discussed in Section 1.3. The masses of the particles in the SM are interpreted as effective masses that arise from interactions of all² massive particles with a constant background field permeating space, as described in Section 1.4.

1.1 Field content

The SM contains 15 different species of particles, shown in Figure 1.1. This includes 12 fermions with spin 1/2, 4 vector bosons with spin 1, and one scalar boson with spin 0, called the Higgs boson.

The 4 vector bosons are the massless uncharged photon γ , the massive electroweak vector bosons W^+ , Z^0 , of which the W^+ carries electric charge, and the massless and colour-charged gluon g . The photon is the particle mediating the electromagnetic interaction as described by Quantum Electrodynamics (QED). The W^+ and Z^0 mediate the weak interaction, and g is the mediator of the strong interaction, as described in Quantum Chromodynamics (QCD).

The 12 fermions consist of 3 generations of 4 fermions. The only difference between the fermion generations are the masses of the fermions³. Each generation of fermions consists of two types of quarks, which are colour charged and thus couple to gluons, and leptons, which do not.

One quark in a fermion generation has an electric charge of 2/3. These are called up-type quarks after the first generation quark with charge 2/3, the up quark u . The up-type quarks of the second and third generation are called the charm quark c and top quark t , respectively. The other quark in any generation, with charge $-1/3$, is correspondingly called a down-type quark after the down quark d in the first generation. The

[3]: Thomson (2013), *Modern particle physics*

[4]: Vergados (2017), *The Standard Model and Beyond*

1: except gravity

2: This may include neutrinos if one considers e.g. three light sterile right-handed Dirac-fermion neutrinos to be part of the SM particle content. However, the origin and hierarchy of the neutrino mass eigenvalues are still open research questions, and not all proposed explanations use the SM Higgs mechanism to generate the neutrino masses.

3: Different mass-eigenstates of fermions do also exhibit different rates of coupling to other mass-eigenstates. However, the overall coupling to fermions of any mass is not different. In this way, one may view these coupling differences as resulting from the presence of mass differences, rather than being revealed by them (cf. Section 1.3.2).

Standard Model of Elementary Particles




Figure 1.1: The particles of the Standard Model of elementary particles. Taken from Ref. [5]

down-type quarks of the second and third generation are called the strange quark s and bottom quark b .

Of the two leptons in a fermion generation, one carries an electric charge, and the other is electrically neutral. The neutral leptons, called neutrinos, have masses that are orders of magnitude smaller than the masses of the charged leptons. The charged leptons, in order of generation, are called electron e , muon μ , and τ -lepton τ . Neutrinos are referred to by the charged lepton that they couple to in charged current interactions, as electron-, muon- or τ -neutrinos: ν_e, ν_μ, ν_τ . These charged current interaction neutrino eigenstates are very misaligned from the mass eigenstates of the neutrinos.

All the electrically or colour-charged fermions in the SM have corresponding anti-particles, i.e. negative-energy solutions to the Dirac Equation propagating backwards in time, interpreted as particles with opposite charges propagating forwards in time. The fundamental bosons are their own anti-particles⁴.

4: in the case of the bosons carrying some conserved charge, the charge-conjugated boson is the anti-particle: e.g., the W^+ is the anti-particle of the W^-

1.2 Kinematics

The kinematics of the 12 spin-1/2 fermions in the SM are described by the Dirac equation, with a free-field Lagrangian density

$$\mathcal{L} = \bar{\psi}(i\partial - m)\psi, \quad (1.1)$$

where ψ is a fermion field operator, m is the mass of the fermion, and $\partial \equiv \gamma^\mu \partial_\mu$. γ^μ are four generating elements⁵ of the Dirac Algebra, which can be defined by

$$\{\gamma^\nu, \gamma^\mu\} = \gamma^\nu \gamma^\mu + \gamma^\mu \gamma^\nu = 2\eta^{\mu\nu}, (\gamma^\mu)^\dagger = \eta^\mu_\nu \gamma^\nu,$$

5: Also called "gamma matrices".

where $\eta^{\mu\nu} = \text{diag}(1, -1, -1, -1)$ is the Minkowski metric and where $(\dots)^\dagger = ((\dots)^*)^T$ denotes the hermitian conjugate. Equation 1.1 fixes the kinematics of the fermion to a momentum four vector with the energy $E = \pm\sqrt{\vec{p}^2 + m^2}$, and gives rise to an intrinsic angular momentum, i.e spin, of $1/2$.

The kinematics of particles with spin 0 or 1 are expressed by the Klein-Gordon equation, with a Lagrangian density

$$\mathcal{L} = \frac{1}{2} [(\partial_\mu \phi)(\partial^\mu \phi) - m^2 \phi^2] , \quad (1.2)$$

with a scalar (vector) field operator ϕ and m the mass of the scalar (vector) particle. While Equation 1.2 specifies the energy-momentum relation of ϕ to $E^2 = m^2 + \vec{p}^2$ as well, it admits bosonic integer-spin solutions. This is unlike Equation 1.1, which is solved only by spinors.

1.3 Local Gauge Symmetries

The SM Lagrangian is constructed as the sum of all Lorentz-invariant terms of mass dimension 4 or lower that are invariant under local field transformations in the gauge symmetry group $SU(3) \otimes SU(2) \otimes U(1)$. The requirement of a mass dimension 4 ensures that the theory is renormalizable, that is, the same set of operators with adjusted prefactors is a valid description at any energy scale. Renormalizability thus ensures the theory is self-consistent up to arbitrarily large energies [7]. Requiring all terms to be Lorentz invariant ensures that the theory is consistent with special relativity. For example, the kinetic terms of the SM Lagrangian comply with this requirement by taking the form of either the Klein-Gordon equation, or the Dirac equation⁶.

The structure of the interactions between the SM particles is given by the local gauge symmetry group $SU(3) \otimes SU(2) \otimes U(1)$. Which particles couple to which other particles, and how, is given by the representation of the corresponding field operators with respect to the gauge group $SU(3) \otimes SU(2) \otimes U(1)$. A qualitative overview of this interaction structure is given in Figure 1.2.




Figure 1.2: Interactions of the elementary particles in the Standard Model. Black ovals represent particles, blue lines connect particles that interact. Taken from Ref. [6]

6: see Section 1.2

1.3.1 Quantum Chromodynamics

The strong interaction is described by QCD, which is included in the SM gauge group $SU(3) \otimes SU(2) \otimes U(1)$ via the $SU(3)$ subgroup of rotations in three complex dimensions. All quarks transform like elements of a triplet under $SU(3)$ transformations:

$$q(x) \rightarrow e^{ig_s \alpha^a(x) T^a} q(x) ,$$

where T^a are the eight 3×3 dimensional generators of SU(3). Under such a transformation, Equation 1.1 is not invariant. Introducing gluons, which couple to quarks in a modified kinetic term

$$\mathcal{L}_{QCD} \supset \bar{q}(i\not{D} + ig_s \not{G}_\mu^a T^a - m)q \equiv \bar{q}(i\not{D}_S - m)q, \quad (1.3)$$

and transform as part of an octet under SU(3)

$$G_\mu^c(x) \rightarrow G_\mu^c(x) - \partial_\mu \alpha^c(x) - g_s f_{abc} \alpha^a(x) G_\mu^b(x),$$

with f_{abc} the SU(3) structure constants, defined by $[T^a, T^b] = if_{abc}T^c$. By defining the covariant derivative D_S of SU(3) in Equation 1.3, the form of Equation 1.1 is recovered, reflecting the fact that in an interacting theory, fermion four momenta may be modified in interactions with the mediator. Defining the gluon field strength tensor [8]

$$\mathcal{G}_{\mu\nu}^a \equiv \partial_\mu G_\nu^a - \partial_\nu G_\mu^a - g_s f_{abc} G_\mu^b G_\nu^c, \quad (1.4)$$

one can then write the full QCD Lagrangian density by adding the kinetic and self-interaction terms of the gluons to Equation 1.3 and summing over all the quark flavours [8]:

$$\mathcal{L}_{QCD} = \sum_q \bar{q}(\not{D}_S)q - \frac{1}{2} \text{tr}(\mathcal{G}_{\mu\nu} \mathcal{G}^{\mu\nu}), \quad (1.5)$$

where the trace gives rise to the $m = 0$ Klein-Gordon term of the gluons and gluon self-interaction terms.

By Noether's theorem, the SU(3) symmetry exhibited by Equation 1.5 yields an associated conserved charge, called colour charge. Since quarks and gluons are triplets and octets under SU(3), they carry colour charge. All other SM particles transform as singlets under SU(3), i.e. they are unchanged. Thus, the other SM particles are colour-neutral.

No free particles carrying colour charge have been observed [8], but only colour neutral systems of bound quarks, i.e., hadrons. This non-existence of free coloured states is called colour confinement. It is hypothesized to be due to the gluon self-interaction giving rise to an attractive force of gluon field lines between two quarks. Thus, the field lines would not spread out as a function of the distance of the quarks, as they would for an electromagnetic field. Hence, the energy stored in the field between two separated colour charges increases with the distance, until it is large enough to pair-produce new quarks, shielding the quark fields in bound states.

1.3.2 The Electroweak Interaction

The electroweak local gauge symmetry group $SU(2) \otimes U(1)$ is the non-QCD part of the SM gauge group $SU(3) \otimes SU(2) \otimes U(1)$. $SU(2)$ corresponds to rotations in two complex dimensions. $U(1)$ corresponds to a free complex phase.

The left-handed components of all SM fermion fields transform as doublets under $SU(2)$:

$$L_\ell = (v_{\ell,L}, \ell_L^-)^T, \quad \ell \in \{e, \mu, \tau\}, \quad (1.6)$$

$$Q = (u_L, \mathcal{D}'_L)^T, \quad (1.7)$$

$$(\mathcal{U}, \mathcal{D}') \in \left\{ \left(u, d' = \sum_{\mathcal{D} \in \{d,s,b\}} V_{u\mathcal{D}} \mathcal{D} \right), \left(c, s' = \sum_{\mathcal{D} \in \{d,s,b\}} V_{c\mathcal{D}} \mathcal{D} \right), \left(t, b' = \sum_{\mathcal{D} \in \{d,s,b\}} V_{t\mathcal{D}} \mathcal{D} \right) \right\},$$

where V_{ij} are the elements of the Cabibbo-Kobayashi-Maskawa (CKM) quark mixing matrix, non-trivial as a result of the misalignment of the quark mass eigenstates with the states forming the $SU(2)$ doublets. The right-handed components are $SU(2)$ singlets. Here “handedness” refers to the chiral projections $\psi_{L,R} = P_{L,R} \psi$ with $P_{L,R} \equiv \frac{1}{2}(1 \mp \gamma^5)$, $\gamma^5 \equiv i\gamma^0\gamma^1\gamma^2\gamma^3$. Since the chirally projected fermion fields change when applying the parity operator $P = \gamma^0$, the interaction described by the $SU(2)$ symmetry group does not conserve parity. Furthermore, the combination of charge and parity CP is not conserved either, due to a complex phase in the CKM matrix.

Similarly to the $SU(3)$ gluon octet of QCD, to allow local gauge transformations of the fermion doublets as

$$\begin{aligned} L_\ell(x) &\rightarrow e^{ig_W \alpha^a(x) t^a} L_\ell(x), \\ Q(x) &\rightarrow e^{ig_W \alpha^a(x) t^a} Q(x), \end{aligned}$$

with t^a as the three generators t^1, t^2, t^3 of $SU(2)$, the coordinate derivative is replaced by a covariant derivative

$$D_{\mu,L} = \partial_\mu - it^a W_\mu^a(x),$$

in the Dirac Equation 1.1, thus introducing a $SU(2)$ triplet of gauge bosons $(W^1, W^2, W^3)^T$. The transformation of the W^a under $SU(2)$ is exactly analogous to that of the gluons under $SU(3)$ in Equation 1.4, with 3 W^a fields instead of 8 gluon fields G^a and the generators and structure constants of $SU(2)$ instead of the $SU(3)$ ones. Like colour in QCD, the $SU(2)$ symmetry gives rise to a conserved charge by Noether’s theorem, which is called weak isospin. The name is due to the fact that the inherent symmetry of a spin-1/2 system can also be described as a $SU(2)$ symmetry. Unlike in the case of colour, particles carrying different isospin can occur in isolation, e.g. as isolated electrons or neutrinos. Consequently, specific

weak isospin projections are assigned to the components of the SU(2) left-handed fermion doublets in Equations 1.6 and 1.7. The t_3 eigenvalue I_3 for charged left-handed leptons and left-handed up-type quarks is $+\frac{1}{2}$. For left-handed neutrinos and left-handed down-type quarks I_3 is $-\frac{1}{2}$. All right-handed fermions are SU(2) singlets, and thus have $I_3 = 0 = I$.

Both left- and right-handed fermions have a non-trivial transformation under the remaining U(1) symmetry group. Its only generator is the identity map, thus the gauge transformation for fermions is to apply a complex phase factor

$$\psi(x) \rightarrow e^{iYg_Y\alpha(x)}\psi(x),$$

where the weak hypercharge Y is the conserved charge associated to the U(1) symmetry. Because there is only one generator, all transformations commute and there are no structure constants. Consequently, the U(1) gauge boson B transforms as

$$B_\mu(x) \rightarrow B_\mu(x) - \partial_\mu\alpha(x),$$

which restores gauge invariance in the Dirac term after replacing the coordinate- with the covariant derivative

$$D_{\mu,Y} = \partial_\mu - iYg_YB_\mu.$$

Left-handed leptons have weak hypercharge $Y = -1$, right-handed leptons have weak hypercharge $Y = -2$, left-handed quarks have $Y = \frac{1}{3}$, right-handed up type quarks have $Y = \frac{4}{3}$, and right-handed down type quarks have $Y = -\frac{2}{3}$. In extensions of the SM that include neutrino masses by the presence of right-handed neutrinos, there are right-handed neutrinos. These right-handed neutrinos would then have $Y = 0$, i.e. they would be sterile.

Collecting the above constructions, and defining field strength tensors $\mathcal{B}, \mathcal{W}$ for the B, W fields in the way of Equation 1.4, the SM Lagrangian thus far reads:

$$\begin{aligned} \mathcal{L}_{\text{massless}} = & -\frac{1}{2}\text{tr}(\mathcal{G}_{\mu\nu}\mathcal{G}^{\mu\nu}) - \frac{1}{2}\text{tr}(\mathcal{W}_{\mu\nu}\mathcal{W}^{\mu\nu}) - \frac{1}{4}\mathcal{B}_{\mu\nu}\mathcal{B}^{\mu\nu} \\ & + \sum_Q \bar{Q}i\not{D}_Q Q + \sum_{u_R} \bar{u}_R i\not{D}_{u_R} u_R + \sum_{d_R} \bar{d}_R i\not{D}_{d_R} d_R \\ & + \sum_\ell \bar{\ell} i\not{D}_\ell \ell \end{aligned} \quad (1.8)$$

with the covariant derivative for the three left-handed quark doublets Q

$$D_Q \equiv \partial - i\frac{1}{3}g_Y B - ig_S T^b G^b - ig_W t^a W^a,$$

and the covariant derivative for the three left-handed lepton doublets

L_ℓ

$$D_{L_\ell} \equiv \partial + i g_Y B - i g_W t^a W^a ,$$

and the covariant derivative for the three right-handed down-type quark singlets $\mathcal{D}_R$

$$D_{\mathcal{D}_R} \equiv \partial + i \frac{2}{3} g_Y B - i g_S T^a G^a ,$$

and the covariant derivative for the three right-handed up-type quark singlets $\mathcal{U}_R$

$$D_{\mathcal{U}_R} \equiv \partial - i \frac{4}{3} g_Y B - i g_S T^a G^a ,$$

and the covariant derivative for the three right-handed charged lepton singlets ℓ_R

$$D_{\ell_R} \equiv \partial + i 2 g_Y B .$$

1.4 Higgs Mechanism

The part of the SM Lagrangian written out in Equation 1.8 contains descriptions of all the SM fermions and vector bosons. For the Lagrangian density written in Equation 1.8, all particles are massless. Mass terms of the forms given in Equations 1.1 and 1.2 would break the $SU(2) \otimes U(1)$ gauge symmetry of the SM. Adding a Klein-Gordon mass term as in Equation 1.2 for e.g., the B boson breaks the underlying $U(1)$ symmetry: $\frac{m_B^2}{2} B^\mu B_\mu \rightarrow \frac{m_B^2}{2} (B^\mu - \partial^\mu \alpha)(B_\mu - \partial_\mu \alpha) \neq \frac{m_B^2}{2} B^\mu B_\mu$. Similarly, since the left-handed fermion fields are $SU(2)$ doublets, but the right-handed fermion fields are $SU(2)$ singlets, a mass term like $m \bar{\psi}_L \psi_R$ breaks $SU(2)$.

To keep the principle of local gauge symmetry intact, the SM interprets the observed masses of the charged fermions and weak gauge bosons as effective masses. These effective masses originate from couplings to a scalar field with a non-trivial vacuum expectation value. This mass generation mechanism is called the Higgs mechanism. Section 1.4.1 describes the construction of this scalar field. Section 1.4.2 discusses how the couplings to the scalar vacuum give rise to effective mass terms.

1.4.1 Scalar Sector Lagrangian

Calling the scalar field ϕ and ignoring all couplings to other fields, its associated Lagrangian density is

$$\begin{aligned}\mathcal{L}_\phi'' &= \frac{1}{2}(\partial_\nu \phi)(\partial^\nu \phi) - V(\phi) \\ &= \frac{1}{2} \left[(\partial_\nu \phi)(\partial^\nu \phi) - \mu^2 \phi^2 - \frac{\lambda}{2} \phi^4 \right],\end{aligned}\quad (1.9)$$

where $V(\phi)$ is called the scalar potential. In order to balance the left-handed fermion doublet appearing in fermion mass terms, the scalar field needs to transform as a SU(2) doublet. This also generates mass terms for the W^a bosons. Furthermore, mass terms for the B boson are needed to obtain the correct chiral coupling structure for the Z^0 boson. Thus, the scalar field is a complex doublet under SU(2), and has a weak hypercharge, $Y = -1$. Introducing the coupling to the SU(2) $\otimes$ U(1) gauge bosons, Equation 1.9 becomes

$$\mathcal{L}_\phi' = (D_{\phi,\nu} \phi)^\dagger (D_\phi^\nu \phi) - \mu^2 \phi^\dagger \phi - \lambda (\phi^\dagger \phi)^2, \quad (1.10)$$

with the covariant derivative

$$D_\phi \equiv \partial + ig_Y B - ig_W t^a W^a.$$

The scalar field of Equation 1.10 has the right quantum numbers and couplings to bosons, but no couplings to fermions. These are included as Yukawa couplings [8]:

$$\begin{aligned}\mathcal{L}_\phi &= \mathcal{L}_\phi' + \sum_\ell y_\ell \left(\bar{L}_\ell \phi \ell_R + \bar{\ell}_R \phi^\dagger L_\ell \right) \\ &\quad + \sum_Q \left[\sum_{\mathcal{U}} \left(\bar{Q} \phi_c Y_{Q\mathcal{U}} \mathcal{U}_R + \bar{\mathcal{U}}_R Y_{Q\mathcal{U}}^\dagger \phi_c^\dagger Q \right) + \sum_{\mathcal{D}} \left(\bar{Q} \phi Y_{Q\mathcal{D}} \mathcal{D}_R + \bar{\mathcal{D}}_R Y_{Q\mathcal{D}}^\dagger \phi^\dagger Q \right) \right],\end{aligned}\quad (1.11)$$

with y_ℓ the strength of the Yukawa coupling to lepton $\ell = e, \mu, \tau$, and $Y_{Q\mathcal{D}}$ and $Y_{Q\mathcal{U}}$ two 3×3 matrices inscribing the Yukawa coupling strengths to down-type quarks $\mathcal{D} = d, s, b$ and up-type quarks $\mathcal{U} = u, c, t$, and Q the doublets of left-handed quarks from Equation 1.7. The conjugate Higgs doublet $\phi_c \equiv \epsilon_2 \phi^\star$ in Equation 1.11 transforms just like ϕ under SU(2) $\otimes$ U(1), with ϵ_2 the fully antisymmetric tensor in 2 dimensions. The quark Yukawa couplings $Y_{Q\mathcal{U}}$ and $Y_{Q\mathcal{D}}$ define the CKM matrix V_{CKM} : $V_{\text{CKM}} \equiv V_{\mathcal{U}}^L (V_{\mathcal{D}}^L)^\dagger$, with $V_{\mathcal{U},\mathcal{D}}^L$ the matrices that diagonalize the quark Yukawa couplings by $V_{\mathcal{U},\mathcal{D}}^L Y_{Q\mathcal{U},Q\mathcal{D}} V_{\mathcal{U},\mathcal{D}}^R$ [8].




Figure 1.3: Illustration of the shape the SM Higgs potential using the potential $V = \mu^2 \phi^\dagger \phi + \lambda (\phi^\dagger \phi)^2$ of a complex scalar ϕ for the case $\lambda > 0, \mu^2 < 0$. The SM Higgs potential has a similar dependence, but since it is a complex doublet, the SM potential is a function of four real parameters, not two. Taken from ref. [9].

1.4.2 Spontaneous Electroweak Symmetry Breaking

The potential of the Higgs doublet ϕ in Equation 1.10, with $\mu^2 < 0, \lambda > 0$, is minimal on a sphere in ϕ -space satisfying

$$\phi^\dagger \phi = -\frac{\mu^2}{2\lambda} \equiv \frac{v^2}{2}, \quad (1.12)$$

where the last step defines the (norm of the) vacuum expectation value v . If the Higgs field is not excited by interactions with other fields, it will be in a vacuum state of minimal potential, i.e., a state fulfilling Equation 1.12. While the set of all these possible vacuum states retains the full $SU(2) \otimes U(1)$ symmetry, any particular vacuum state breaks the symmetry, as illustrated in Figure 1.3. In this way, the non-zero vacuum expectation value (VEV) of the Higgs doublet is said to “spontaneously break” the $SU(2) \otimes U(1)$ symmetry. Despite this nomenclature, the full symmetry is still present, if one takes into account the now non-trivial gauge transformations of the Higgs vacuum.

Specifying the VEV of ϕ such that only the neutral real component is non-vanishing

$$\langle 0 | \phi | 0 \rangle = \frac{1}{\sqrt{2}} \begin{pmatrix} 0 \\ v \end{pmatrix},$$

ensures that the photon is massless in the broken phase. Writing ϕ as relative to the VEV one obtains

$$\phi(x) = \frac{1}{\sqrt{2}} \begin{pmatrix} \phi_1(x) + i\phi_2(x) \\ v + H(x) + i\phi_4(x) \end{pmatrix},$$

which can be simplified by picking a gauge in $SU(2) \otimes U(1)$ such that only the real neutral component is non-vanishing

$$\phi(x) = \frac{1}{\sqrt{2}} \begin{pmatrix} 0 \\ v + H(x) \end{pmatrix},$$

which is called the unitary gauge. The degrees of freedom $\phi_{1,2,4}$ reappear via gauge terms of the $SU(2) \otimes U(1)$ gauge bosons W^a, B . They can be absorbed as longitudinal degrees of freedom into the vector bosons. In the unitary gauge, the kinetic term from Equation 1.10 can be written

as

$$(D_{\phi,\nu}\phi)^\dagger(D_\phi^\nu\phi) = \frac{1}{2}(\partial_\mu H)(\partial^\mu H) + \frac{g_W^2}{8}(v+H)^2(W^1+iW^2)_\mu(W^1-iW^2)^\mu + \frac{1}{8}(v+H)^2(g_W W^3 - g_Y B)_\mu(g_W W^3 - g_Y B)^\mu, \quad (1.13)$$

where by defining $W^\pm \equiv \frac{W^1 \mp iW^2}{\sqrt{2}}$ and considering the term

$$(D_{\phi,\nu}\phi)^\dagger(D_\phi^\nu\phi) \supset \frac{1}{2} \left(\frac{g_W v}{2} \right)^2 W_\mu^- W^{+\mu},$$

one can identify the mass

$$M_W = \frac{g_W v}{2}.$$

Writing the term with W^3, B in Equation 1.13 as a matrix multiplication

$$(D_{\phi,\nu}\phi)^\dagger(D_\phi^\nu\phi) \supset \frac{v^2}{8} \begin{pmatrix} W^3 & B \end{pmatrix} \begin{pmatrix} g_W^2 & -g_W g_Y \\ -g_W g_Y & g_Y^2 \end{pmatrix} \begin{pmatrix} W^3 \\ B \end{pmatrix},$$

and rotating $(W^3, B)^T$ into the eigenbasis $(A_\gamma, Z)^T$, thus diagonalizing the matrix

$$(D_{\phi,\nu}\phi)^\dagger(D_\phi^\nu\phi) \supset \frac{1}{2} \begin{pmatrix} A_\gamma & Z \end{pmatrix} \begin{pmatrix} 0 & 0 \\ 0 & (\frac{v}{2})^2(g_Y^2 + g_W^2) \end{pmatrix} \begin{pmatrix} A_\gamma \\ Z \end{pmatrix},$$

one can identify the masses

$$\begin{aligned} M_\gamma &= 0, \\ M_Z &= \frac{v \sqrt{g_W^2 + g_Y^2}}{2}, \end{aligned} \quad (1.14)$$

where A_γ is the photon field. Thus, the photon is the gauge boson of the remaining U(1) unbroken gauge symmetry of complex rotations around the axis defined by the direction of $\langle \phi \rangle$. The angle of the rotation into the mass eigenstates A_γ, Z is the Weinberg angle θ_W with

$$\cos \theta_W = \frac{M_W}{M_Z},$$

or, equivalently

$$\tan \theta_W = \frac{g_Y}{g_W}.$$

The remaining neutral scalar degrees of freedom H is called the Higgs boson. Its couplings to the massive vector bosons $W^\pm, Z$ can also be read

off from Equation 1.13. Thus, e.g., the triple couplings are given by:

$$\begin{aligned} (D_{\phi,\nu}\phi)^\dagger(D_\phi^\nu\phi) &\supset \frac{g_W^2 v}{2} W_\mu^+ W^{-\mu} h H + \frac{g_W^2 v}{2} Z_\mu Z^\mu H \\ &= \frac{2M_W^2}{v} W_\mu^+ W^{-\mu} H + \frac{M_Z^2}{v} Z_\mu Z^\mu H. \end{aligned}$$

The potential term in Equation 1.10, in the unitary gauge, has a term $V \supset v^2 \lambda H^2$ and thus the Higgs boson H has a mass

$$M_H = \sqrt{2\lambda} v.$$

While v can be determined precisely from measurements of the Fermi coupling constant $G_F = \frac{v^2}{\sqrt{2}} = 1.1663788(7) \times 10^{-5} \text{ GeV}^{-2}$, e.g., from measurements of the muon lifetime [8, 10], λ is a free parameter of the theory that affects only M_H . Thus, M_H is not predicted and must be measured. Combining ATLAS and CMS measurements of M_H in the $H \rightarrow \gamma\gamma$ and $H \rightarrow ZZ^* \rightarrow 4\ell$ with $\ell = e, \mu$ channels during the first two runs of the LHC, [11] find $M_H = 125.25 \pm 0.17 \text{ GeV}$ on average.

Finally, the Yukawa couplings in Equation 1.11 give rise to effective fermion masses. For the example of the electron, one gets

$$\mathcal{L}_\phi \supset -\frac{vy_e}{\sqrt{2}}(\bar{e}_L e_R + \bar{e}_R e_L),$$

and thus the fermion masses are

$$M_f = \frac{y_f v}{\sqrt{2}},$$

and a similar way the fermion couplings to the Higgs boson are read off

$$\begin{aligned} \mathcal{L}_\phi &\supset -\frac{y_e}{\sqrt{2}}(\bar{e}_L e_R + \bar{e}_R e_L)H \\ &= -\frac{M_e}{v}(\bar{e}_L e_R + \bar{e}_R e_L)H. \end{aligned}$$

Thus the tree-level couplings of the Higgs boson to both fermions and vector bosons are proportional to the mass of the fermion or boson under consideration.

Physics beyond the Standard Model

2

2.1 Motivating new physics

The Standard Model of Particle Physics (SM) very successfully describes a wide range of phenomena across many orders of magnitude of energy and cross-section in collider experiments, astrophysics, and other fields [12]. While successful in predicting or matching a wide range of observations, there are reasons to anticipate fundamental physics that is outside the scope of the SM. Examples include:

- The observation of neutrino oscillation, which indicates that neutrinos have mass, is by some [11] taken as direct evidence of physics beyond the SM (BSM), by adopting a narrow view of the SM requiring the SM neutrinos to be massless. But even less narrow interpretations suggest the possibility of neutrino masses originating from e.g., a Seesaw mechanism involving BSM heavy neutral leptons.
- The observed matter-antimatter asymmetry of the universe can in current cosmological models be explained as resulting from CP-violating interactions. However, the CP-violating interactions of the SM are not frequent enough to explain the observed asymmetry [13].
- The existence of Dark Matter (DM), a gravitating fluid with at most weak non-gravitational interactions with light and known matter, is confirmed by many astronomical and cosmological observations [14–16]. Many proposed accounts of the nature of DM are not compatible with it being composed of known SM particles.

For all the above reasons, one might expect physics beyond the SM (BSM) to be present at energy scales well below the Planck scale, and potentially even at energies accessible at the LHC. At the same time, proposed BSM models are plentiful [17–28] and so it is useful to pursue widely-interpretable searches that have the sensitivity to discover or rule out as large as possible a class of theories.

To achieve this sensitivity to a large class of theories, dijet resonance searches have long been used [29, 30], and it is one motivation for the dijet search presented in this thesis. Low-mass hadronic resonances are of a particular interest due to a relative paucity of constraints, both from theory and from experiments, as Section 2.3 discusses in more detail later.

A second motivation for hadronic resonance searches arises out of the particular case of searches for DM production at the LHC, as Section 2.2 argues in the following.

[11]: Workman et al. (2022), *Review of Particle Physics*

[14]: Bertone et al. (2005), *Particle dark matter: Evidence, candidates and constraints*

[15]: Bergström (2000), *Nonbaryonic dark matter: Observational evidence and detection methods*

[16]: Clowe et al. (2006), *A direct empirical proof of the existence of dark matter*

2.2 Dark Matter

2.2.1 Theory of WIMP Dark Matter

Dark Matter is indicated to exist by a wide range of astronomical and cosmological evidence [14–16, 31]. There is a plethora of proposed Dark Matter candidates, including hidden baryonic matter such as MACHOs, primordial black holes, sterile neutrinos, and axions [14, 15]. This section focuses on Weakly Interacting Massive Particles (WIMPs), which are conventionally used to interpret many LHC Dark Matter Searches [28, 32, 33].




Figure 2.1: In the early universe, WIMPs need to decay in large numbers through the graph (a) in order to match the observed relic density of dark matter [14, 34, 35]. By CPT symmetry this implies a comparable cross-section for the graph (b), where Dark Matter is produced in association with initial state radiation. Dark Matter cannot be observed by LHC detectors, but the initial state radiation can. This gives rise to $E_T^{\text{miss}} + X$ - or mono- X signatures, where X denotes the observable initial state radiation (e.g., a photon or jet). Taken from [35].

Hypothetical WIMPs are produced and annihilated in thermal equilibrium in the very early Universe. As the universe expands, it cools. Eventually the typical energy of particle collisions drops below the threshold for WIMP production. WIMPs continue to annihilate into lower-mass standard model particles as they collide, which can be expressed through the Boltzmann equation as [34, 36]

$$\dot{n} = -3Hn - \langle \sigma_{\text{ann}} v \rangle (n^2 - n_{\text{eq}}^2), \quad (2.1)$$

where the number density of WIMPs n changes as a result of both the expansion of the universe with rate H , and the pair annihilation with velocity-averaged cross-section $\langle \sigma_{\text{ann}} v \rangle$. Solving equation (2.1) for n and expressing the solution relative to the critical density of a flat universe, one obtains [14]

$$\Omega h^2 \approx \frac{3 \times 10^{-27} \text{ cm}^3 \text{ s}^{-1}}{\langle \sigma_{\text{ann}} v \rangle}. \quad (2.2)$$

The Planck collaboration measured cosmic microwave background anisotropies, and from these inferred the relic density of cold dark matter to be $\Omega h^2 = 0.1187(17)$ [37]. Matching this measurement with Equation 2.2 requires a velocity averaged annihilation cross-section of a size that is expected for WIMP candidates with masses of order 100 GeV – 1 TeV [14, 34]. This range of masses and cross-sections is accessible in LHC searches, which motivates the popularity of WIMP models for the interpretation of such searches. Furthermore, any interaction allowing Dark Matter annihilation into standard model particles implies the existence of scattering processes allowing production of Dark Matter at the LHC, as illustrated in Figure 2.1 [14, 34, 35].

A variety of theories of new Physics beyond the Standard Model incorporate WIMP Dark Matter. For example, in supersymmetric theories that conserve R -parity the Lightest Supersymmetric Particle (LSP) cannot decay, since no Standard Model decay products can match its R -parity [14, 15]. If the LSP only couples weakly to Standard Model fields and has a mass in the 100 GeV – 1 TeV range, it could constitute WIMP Dark Matter.

[14]: Bertone et al. (2005), *Particle dark matter: Evidence, candidates and constraints*

[34]: Taoso et al. (2008), *Dark Matter Candidates: A Ten-Point Test*

[35]: Kahlhoefer (2017), *Review of LHC Dark Matter Searches*

SUSY and other complete theories providing WIMP Dark Matter candidates typically have very large parameter spaces, and most of the parameters and fields in those theories do not impact the interesting interactions of the WIMP [14, 15]. Hence instead of the full theory, Effective Field Theories (EFTs) that approximate complete theories in the low energy limit are often used for searches [38]. In such EFTs, the Dark Matter - SM interaction expressed through a contact interaction term in the Lagrangian density [35, 39], such as

$$\frac{\bar{\chi}\gamma^\mu\chi\bar{q}\gamma_\mu q}{\Lambda^2}. \quad (2.3)$$

EFTs are advantageous when comparing experimental results to theory predictions, because large classes of different theories can be treated simultaneously using few parameters. The parameters distinguishing different EFTs for WIMP searches are the effective suppression scale of the contact interaction and the mass of the Dark Matter candidate. Such contact interaction EFTs are the framework used for discussing experiments where Dark Matter scatters off nuclei that are at rest in the laboratory frame [38]. In this case $Q_{\text{scatter}}^2 \ll m_{\text{DM}}^2, \Lambda$, so the EFT captures all accessible degrees of freedom, and experiments can set limits in the spaces of Dark Matter mass and the spin dependent or spin independent nuclear scattering cross-section.

LHC searches however deal with collisions at energy scale up to the nominal centre-of-mass energy of $\sqrt{s} = 14$ TeV. In this case, the effective suppression scale Λ of EFTs the search is sensitive to is often reached or even surpassed by the momentum transferred in the collision. Consequently, EFTs are not an appropriate framework for the interpretation of LHC searches [35, 39]. So-called simplified models are used instead, consisting of the dark matter candidate and one or more additional particle mediating the interaction between Dark Matter and the Standard Model particles [28, 35, 39]. This mediators propagator terms can be integrated out into EFT coefficients in the low-energy limit, allowing consistent comparison with nuclear scattering experiments. A typical simplified model is specified by the masses of the Dark Matter candidate and mediator, the coupling between the mediator and Dark Matter, the couplings of the mediator to Standard Model particles, and the spin and CP quantum numbers of the mediator and Dark Matter candidate [28].

[35]: Kahlhoefer (2017), *Review of LHC Dark Matter Searches*

[39]: Buchmüller et al. (2014), *Beyond Effective Field Theory for Dark Matter Searches at the LHC*

[38]: Marrodán Undagoitia et al. (2016), *Dark matter direct-detection experiments*

[28]: Abdallah et al. (2015), *Simplified Models for Dark Matter Searches at the LHC*

2.2.2 Dark Matter Searches

The astronomical and cosmological evidence for dark matter is indirect [14–16], hinging on the gravitational effects of Dark Matter at galactic and larger scales. The properties of Dark Matter at length scales below the radius of small galaxies are thus largely unknown [14, 15]. Hence, there is a sustained interest in searches for the observation of Dark Matter

interacting with Standard Model particles. Many of these Dark Matter searches fall under one of three distinct categories [14, 15]:

- indirect detection searches,
- direct detection searches,
- and collider searches.

Indirect detection searches attempt to observe the products of Dark Matter annihilation. Since the mean number density of Dark Matter today would be too low to make such annihilations likely, one typically focuses the search on regions where the Dark Matter is expected to be densely clustered, such as the centres of galaxies [40].

Direct detection searches are experiments trying to observe Dark Matter particles scattering off nuclei on earth as earth moves through the Dark Matter halo of the Milky Way [38]. The scattering imparts momentum on the nucleus, which can be measured using a very sensitive and highly shielded particle physics detector. A variety of detector technologies can be used, such as dual phase time projection chambers [41, 42], or cryogenic solid-state bolometers [43].

Collider searches try to find evidence for the production of Dark Matter particles in high-energy particle collisions, e.g. in high-energy proton-proton collisions at the Large Hadron Collider (LHC) [14, 15, 35]. The models used to interpret such searches typically assume one or more novel boson that mediates the Dark Matter interaction with the Standard Model [28, 32, 33, 44]. However, if such a model can produce Dark Matter in proton-proton collisions, it couples to at least some partons of the proton, either directly or through loops. Due to this coupling to quarks or gluons, the decay of the mediating boson into hadronic jets should also be observable at the LHC [35]. Searches for dijet resonances can thus guide and constrain searches for Dark Matter at the LHC.




Figure 2.2: Example Feynman graphs and final state topologies commonly targeted by LHC Dark Matter searches. The advantages and limitations of these types of topologies are discussed in the text.


2.2.2.1 LHC Searches

Dark Matter searches at particle colliders have the potential to produce WIMPs and similar Dark Matter candidates that undergo thermal freeze-out in the early universe, as argued in Section 2.2.1 and Figure 2.1.


Mono- X searches are searches for events in which Dark Matter is produced, indicated by missing transverse momentum in the event [35]. Missing transverse momentum $E_T^{\text{miss}} = -\sum_i \vec{p}_T^i$ is defined as the negative sum of all transverse momenta of any object observed in the event. Dark matter production results in E_T^{miss} because momentum is conserved and the initial state has no transverse momentum, but the Dark Matter particles are invisible to the detector. Thus, their momenta show up as missing in the balance of final state transverse momentum.

Observing Dark Matter production requires a visible object to be produced alongside it, with different searches targeting different visible objects [44]. This object can always be produced as initial state radiation (ISR), irrespective of the details of the model of Dark Matter production. This motivates the very generic mono-jet [45, 46] and mono-photon searches [47]. In more complete, but less generic models of Dark Matter production, both the Dark Matter and the visible object are produced in the hard interaction [33, 44]. These less simplified models are targeted by searches such as mono- H [48], mono- $Z(l\bar{l})$ [49], mono- $V(\text{had.})$ [50], and others [44]. Figure 2.2(a) illustrates the mono- X topology.

New particles acting as mediators in the s -channel are a feature of most simplified models used to interpret mono- X searches [35, 44]. Since these mediators can be created in proton collisions, they couple to partons, and can also decay into partons. Thus, any Dark Matter model with an s -channel mediator targeted by a mono- X search can also be probed using a dijet resonance search for the decay of the mediator into hadrons [35]. Figure 2.2(b) depicts the resonance topology. Although insufficient to prove Dark Matter production, the resonance searches typically have greater exclusion- or discovery-potential, Figure 2.3(a) illustrates. Because the decay into Dark Matter is not involved in the targeted topology, the sensitivity of a resonance search does not have a strong dependence on the Dark Matter mass, unlike mono- X searches [29, 35, 44, 51]. Thus searches for new hadronically decaying resonances are a core part of the LHC Dark Matter search program, complementing mono- X searches [44].



(a) Summary of resonance and mono-X searches



(b) Comparison of hadronic resonance searches.

Figure 2.3: Summary of 95% CL ATLAS confidence contours for the leptophobic axial vector mediator simplified Dark Matter. The coupling of the mediator to Dark Matter is chosen as $g_\chi = 1$, and the coupling to leptons is chosen as $g_l = 0$. Retrieved from [51].

2.3 Constraints on low-mass leptophobic Z' models

Hadronic resonance searches have long been used to probe a variety of new physics models using a single generic search [29, 30]. However, most extant searches focus on the high-mass frontier. For example, the published ATLAS dijet search programme does not constrain resonances with masses below 100 GeV [51]. While historical results do constrain such models, they only do so for comparatively large coupling strengths. The unprecedented luminosity delivered by the LHC and recorded by experiments makes tests of these as-of-yet unconstrained smaller-coupling scenarios feasible, as will be argued in the following for the more concrete example of low-mass leptophobic Z' models.

Z' bosons are hypothetical spin 1 vector bosons that are singlets under electric and colour charge [12]. They are predicted in a variety of models for new physics beyond the standard model. For example, several simplified WIMP dark matter models include a Z' mediating the interaction between the dark matter and standard model fields [20, 28, 29, 32, 35, 39, 44]. Many such models can be probed over a wide range of model parameters using the results of a single resonance search, as discussed in Section 2.2.2.1.

Constraints from LHC searches for Z' decays into pairs of charged leptons are strong [44, 51], but may be elided by choosing the couplings of the Z' such that it is *leptophobic*, meaning it does not couple to SM leptons. However, constraints for searches for hadronically decaying Z' resonances can not be avoided in this way, since the same $Z' q\bar{q}$ decay vertex is also responsible for Z' production at the LHC for any kind of search. Thus hadronic Z' production searches are a powerful tool to probe even models that are tuned to avoid constraints from other LHC searches.

Very-low mass ($m_{Z'} \lesssim 100$ GeV) leptophobic Z' models are not constrained by dijet searches at previous colliders [30]. Thus, they are mostly constrained indirectly. Mixing with the SM Z^0 can lead to modifications of the Z^0 hadronic decay width [52]. Models with $m_{Z'} \lesssim 10$ GeV are furthermore constrained by measurements of the width of quarkonium states such as $\Upsilon(1S)$, $\Upsilon(2S)$, J/Ψ , $\Psi(2S)$, and Φ [52]. However, these constraints allow small couplings to SM quarks ($g_q \lesssim 0.2$) for Z' with a mass that is far from the Z^0 or quarkonium mass. Hence, limits from LHC searches such as [53] have the potential to be the most stringent limit on low-mass leptophobic Z' models, if the sensitivity is sufficient to probe small couplings and low masses simultaneously.

Leptophobic Z' models are one example of a well motivated class of hypothesized BSM physics models that can be tested by a low-mass hadronic dijet search. However, they are not the only such models, and

[12]: Particle Data Group et al. (2020), *Review of Particle Physics*

[44]: Aaboud et al. (2019), *Constraints on mediator-based dark matter and scalar dark energy models using $\sqrt{s} = 13$ TeV pp collision data collected by the ATLAS detector*

[51]: (2023), *Dark matter summary plots for s-channel, 2HDM+a, Higgs portal and Dark Higgs models*

[30]: Harris et al. (2011), *Searches for Dijet Resonances at Hadron Colliders*

[52]: Dobrescu et al. (2014), *Hidden GeV-scale interactions of quarks*

beyond considering specific models, extant analyses of the Run 2 ATLAS dataset do not cover dijet resonances of any origin with masses below 100 GeV [51].

2.4 Unlocking the lowest resonance masses at the LHC

Dijet resonance searches have limited sensitivity to resonance masses below several hundred GeV. This is because towards low dijet masses the QCD cross-section for dijet production rises quickly. The production rate of background events thus outpaces the bandwidth of the data acquisition system [54]. To make the best use of the available bandwidth, only events where the highest-momentum jet exceeds some transverse momentum threshold are triggered. This transverse momentum threshold results in a dijet mass threshold at approximately $m \approx 2p_T$ [54]. Hence, the conventional dijet resonance search [55], relying on events triggered by the presence of at least one jet with $p_T \geq 450$ GeV, has no sensitivity below $O(1 \text{ TeV})$, as shown in Figure 2.3(b).

1: The overview given here is not exhaustive and other approaches have been proposed, e.g. ref. [56].

[57]: Aaij et al. (2016), *Tesla : an application for real-time data analysis in High Energy Physics*

[58]: Khachatryan et al. (2016), *Search for narrow resonances in dijet final states at $\sqrt{s} = 8 \text{ TeV}$ with the novel CMS technique of data scouting*

[59]: Aaboud et al. (2018), *Search for low-mass dijet resonances using trigger-level jets with the ATLAS detector in pp collisions at $\sqrt{s} = 13 \text{ TeV}$*

2: However, the requirement of sufficiently hard ISR effectively acts like a trigger prescale, reducing the observable event rate.

3: For anti- k_T jets clustering a pair of particles, the reach is the radius parameter R [61].

There are several¹ ways to work around this limitation imposed by the jet trigger threshold. One is to cluster, calibrate, and record jets during online data taking, before or as the final high-level trigger decision is made. This allows recording much fewer data per-jet and per-event. The reduced amount of data can be recorded at the rate of the Level 1 hardware trigger, which has a lower jet p_T threshold even without prescaling. Many collider experiments pursue this strategy, which LHCb calls “real-time data analysis” [57], CMS calls “data scouting” [58], and ATLAS calls “trigger-level-object analysis” [59], or TLA. Due to lower trigger thresholds, TLAs are sensitive to lower dijet masses than the default offline search. However, TLAs still are limited at low masses due to the L1 trigger threshold, which in the case of the ATLAS Run 2 TLA means there is no sensitivity for $m_{Z'} < 450$ GeV [59].

Another way to avoid the limiting trigger thresholds, and achieve sensitivity to even lower masses, is to modify the targeted signature. By searching for ISR produced in association with the dijet resonance, the dijet mass is decoupled from the trigger threshold [53, 54, 60]. The ISR can have a high enough momentum to trigger the recording of the event, no matter how low the mass of the hadronic resonance.² In this way, even the lowest dijet resonance masses can be probed at the LHC.

The interplay of dijet, TLA dijet, and ISR-associated dijet searches is shown in Figure 2.3(b). At very low resonance masses, the jets of hadrons radiated by the decay products of the resonance are so collimated that their angular separation is comparable to or smaller than the angular reach³ of the jet clustering algorithm. Thus, to probe the very lowest

resonance masses, the resonance candidate is reconstructed as a single large- R jet, and the substructure of this jet is used to determine the resonance mass. Figure 2.2(c) illustrates this topology. This is the approach pursued in this thesis, using a single hard ISR photon to trigger events and reconstructing resonance candidates as a large- R jet.

As discussed above, the two leading limiting factors on the discovery potential of low-mass dijet resonance searches are high trigger thresholds, and the precision with which the shapes and four-momenta of jets can be reconstructed [53, 54, 58–60]. The search for Initial State Radiation (ISR) photon and a boosted hadronic resonance tries to address both of these limiting factors to enhance its discovery potential.

2.5 Prior searches for very-low-mass hadronic resonances

Both ATLAS and CMS have previously published searches for boosted dijet resonance produced in association with an ISR jet or photon [53, 54, 60]. This section briefly reviews some common challenges faced by these previous iterations of the search, and how they were addressed.

The boosted hadronic resonance candidate is reconstructed as a large-radius jet. The ATLAS searches use anti- k_T jets [62] with radius parameter 1, the CMS photon channel search uses anti- k_T jets with radius parameter 0.8, and the CMS jet channel search uses Cambridge-Aachen jets [63] with radius parameter 1.5. The large radius parameters make it likely that the jet contains all the decay products of the resonance, however, the increased catchment area increases the jet's susceptibility to ISR and pileup [64, 65]. Pileup biases and deteriorates the determination of the mass and energy of jets. To mitigate the effects of pileup, all searches employ jet grooming techniques. The ATLAS searches use jet trimming [66], where the jet is reclustered into smaller subjets, and those subjets that carry less than 5% of the total jet momentum are discarded. The CMS searches employ the soft-drop algorithm [67, 68] instead, iterating through the jets clustering history while filtering out soft- and wide-angle radiation contributions. The CMS searches also suppress pileup contribution by applying per-particle weights to indicate the probable pileup contribution [69].

Two background processes are considered in all searches. One is non-resonant QCD multijet production, either with or without an associated hard photon. The other common background are hadronically decaying electroweak vector bosons produced in association with a photon or jet. The CMS searches also consider single-top and $t\bar{t}$ production as backgrounds, which are neglected in the ATLAS searches.

[53]: Sirunyan et al. (2019), *Search for Low-Mass Quark-Antiquark Resonances Produced in Association with a Photon at $\sqrt{s} = 13$ TeV*

[54]: Aaboud et al. (2019), *Search for light resonances decaying to boosted quark pairs and produced in association with a photon or a jet in proton-proton collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector*

[60]: Sirunyan et al. (2019), *Search for low mass vector resonances decaying into quark-antiquark pairs in proton-proton collisions at $\sqrt{s} = 13$ TeV*

Jet substructure observables quantify the shape of the distribution of jet constituents. The shape of jets originating from the non-resonant QCD background are determined by the parton shower and hadronization of a single original parton, resulting in a distribution of constituents clustering isotropically around single central hard energy deposit. Jets originating from a hadronically decaying vector boson or BSM (axial-)vector mediator instead derive their shape from the quark pair the resonance decays. Thus, their constituents tend to cluster around two hard prongs.

[70]: Thaler et al. (2011), *Identifying Boosted Objects with N-subjettiness*

[71]: Larkoski et al. (2013), *Energy Correlation Functions for Jet Substructure*

[72]: Moulton et al. (2016), *New Angles on Energy Correlation Functions*

[68]: Dasgupta et al. (2013), *Towards an understanding of jet substructure*

[73]: Dolen et al. (2016), *Thinking outside the ROCs: Designing Decorrelated Taggers (DDT) for jet substructure*

All searches exploit these jet substructure differences to reduce the large contribution of the nonresonant backgrounds. The ATLAS searches use the n -subjettiness ratio τ_{21} [70], comparing how well the jet constituents align with two subjet axes, as compared to a single axis. The CMS searches use a ratio of powers of generalized energy correlation functions N_2 [71, 72], which likewise is sensitive to two-prong structure in the jet constituents.

Both jet substructure observables used are correlated with the mass of the jet, due to their dependence on the opening angles between jet constituents. To avoid shaping of the background mass spectrum by jet substructure selections, all searches use dedicated techniques to decorrelate τ_{21} and N_2 from the jet mass [68, 73]. This decorrelation is achieved by rescaling and shifting these observables as a function of jet mass and momentum. Selections on the decorrelated substructure variants no longer introduce resonance-like structures in the mass distribution of the non-resonant backgrounds.

The effect of the cut on the decorrelated jet substructure observable is not modelled accurately by the Monte-Carlo templates of the nonresonant backgrounds. All searches avoid this mismodelling by using data-driven estimates to describe the main nonresonant background. These estimates are typically based on a data control region where the jet substructure cut is inverted. An in-situ efficiency correction is extracted from a fit to data in phase-space near the SR, and applied to the data-driven estimate to correct for the residual jet-kinematic dependence of the tagger efficiency. The ATLAS search applies a linear fit to decorrelate the tagger, and a non-parametric interpolation of the tagger-efficiency. By contrast, the CMS search applies a non-parametric smoothing technique to decorrelate the tagger, and then fits a bivariate polynomial to determine the tagger efficiency. Because the mass- and p_T scales at which resonant and non-resonant jets are correlated are both determined by the jet's mass resolution and energy resolution, the use non-parametric interpolation to model the non-resonant tagging efficiency implies an unfavourable trade-off between the uncertainty on the interpolated efficiency and signal contamination in the CR.

All searches fit combined signal and background templates to the ob-

served jet mass spectrum and compare these to background-only template fit results to find confidence intervals constraining the Z' production cross-section⁴ None of the searches has found a significant deviation from known backgrounds. ATLAS exclude Z' mediator production cross-sections down to 90.1 through 0.01 pb in a mediator mass range covering 100 – 220 GeV. CMS exclude similar Z' production cross-sections for masses ranging from 10 GeV to 425 GeV.

4: or equivalently, Z' - quark coupling strength

The ATLAS Detector at the LHC

The Large Hadron Collider (LHC) at CERN near Geneva in Switzerland is the most powerful particle collider in operation at the time of writing. The ATLAS detector is one of several detectors recording the LHC hadron collisions. This chapter gives a short overview of both, adapting a previous description in [2]. Section 3.1 briefly describes the LHC, following Refs. [74, 75]. Section 3.2 gives a short overview of the ATLAS Detector, following Ref. [76]. The concept of luminosity and its measurement with the ATLAS detector are introduced in Section 3.3.

[74]: Evans et al. (2008), *LHC Machine*

[75]: Vidal et al. (2018), *Taking a closer look at the LHC*

[76]: Aad et al. (2008), *The ATLAS Experiment at the CERN Large Hadron Collider*

3.1 The Large Hadron Collider

The LHC consists of two counter-rotating storage rings. The storage rings are located in a single circular underground tunnel with a circumference of 26.7 km located near Geneva, Switzerland, at a depth between 45 m and 175 m below the surface. The storage rings can be filled with either protons or lead ions, but this thesis uses only data recording proton collisions.


The protons for the LHC are produced by ionizing hydrogen gas with an electric field. Once the electrons are stripped away by the electric field, the remaining protons are accelerated in a sequence of preaccelerators, displayed in Figure 3.1. These include, in sequence, the linear accelerator LINAC2, followed by the synchrotron accelerators BOOSTER, proton Synchrotron (PS), and Super Proton Synchrotron (SPS). Because the pre-accelerating synchrotrons use radio frequency (RF) cavities to accelerate the protons, the protons travel as *bunches* of $\approx 10^{11}$ protons. The proton bunches from the SPS are injected into the two LHC rings at two injection points, from where they are accelerated to the final beam energy.

One LHC beam can consist of up to 2808 bunches spaced¹ 25 ns apart. Accelerating a proton beam to the final LHC beam energy takes roughly two hours in total, of which the proton beam spends roughly half an hour being accelerated in the LHC. The proton beam is then stored in the LHC for dozens of hours. During this time pairs of bunches from the two beams are made to cross through each other, allowing protons to collide. Due to these collisions, the beams are losing protons over the course of many hours, until the beam is eventually dumped.

The LHC rings consist of a sequence of 8 straight segments and 8 curved segments. The straight segments host RF cavities that accelerate protons in a certain phase of an electromagnetic field which oscillates at 400 MHz and has an effective field strength of 2 MV/m. The straight sections are interspersed with curved sections containing a total of 1232

1: Other bunch spacings are also used (e.g., 50 ns) depending on the target luminosity and experimental program. Furthermore, even given a certain bunch spacing, not every bucket (potential bunch position along the beam) is filled with a bunch. Rather, bunches are carefully structured into densely spaced trains of bunches, with a particular structure of empty and full buckets. The details of the bunch spacing and bunch train structure are chosen based on technical constraints from both the accelerator and the detectors recording collisions.

Figure 3.1: Schematic illustrating the LHC and the associated chain of lower energy accelerators. The names, final beam energy and year of first operation are given for the injectors. For the LHC, the maximum design collision energy is given instead. The position of the four main LHC detectors is indicated by grey dots. Taken from ref. [77].



superconducting NbTi magnet coils. The coils produce dipole fields up to 8.33 T in strength that force the protons on a circular trajectory. Thus, despite gaining only 2 MeV of energy per RF cavity, the protons can revolve many times, and reach kinetic energies up to the nominal design energy of 7 TeV. The maximum energy is limited by the strength of the dipole magnets, given the curvature radius of the tunnel. The Run 2 data used in this thesis correspond to 6.5 TeV beam energy, but this has been superseded in Run 3 with a maximum² beam energy of 6.8 TeV. Furthermore, 858 quadrupole magnets are used to squeeze the beam, counteracting beam dispersion and focusing the beam to a radius of $16 \mu\text{m}$ prior to reaching an interaction point. The LHC magnets are cooled down to 1.9 K, using superfluid and therefore highly heat-conducting helium as coolant.

The two counter-rotating beams are brought to collision at four interaction points. Around these interaction points, detectors are installed to record the particles produced in the collisions. These locations are indicated in Figure 3.1. The four main detectors installed at the LHC are:

- ▶ ALICE [78], specialized towards the study of heavy ion collisions and Quantum Chromo Dynamics (QCD);
- ▶ ATLAS [76], a general purpose detector constructed to allow reconstructing the results of particle collisions as close to completely as possible, to facilitate precise measurements of SM properties and searches for physics not predicted by the SM;
- ▶ CMS [79], a general purpose detector like ATLAS and sharing the same physics goals, but using a different technological implementation;
- ▶ LHCb [80], a single-arm forward spectrometer specialized towards the study of CP violation in hadron collisions and rare processes involving the decays of hadrons containing b quarks.

3.2 The ATLAS Detector

The ATLAS detector, illustrated in Figure 3.2 is a multipurpose detector, surrounding interaction point 1 of the LHC. Its three main components are the Inner Detector (ID), the calorimeters, and the muon spectrometers.

2: at time of writing




Figure 3.2: The components of the ATLAS Detector. Taken from Ref. [81].

These components successively and concentrically surround the beam axis, in a forward-backwards symmetric way. The sub-detectors are immersed in a magnetic field generated by a magnet system consisting of one solenoid magnet surrounding the ID, and a toroid magnet system in the barrel and end-cap regions close to the muon spectrometers. The ATLAS detector is 44 m long in the beam direction, 25 m high, and weighs a total of 7000 tons.

When describing the geometry of the ATLAS detector and events recorded with it, a reference coordinate system is used. The reference coordinate system is a right-handed and orthogonal coordinate system. The coordinate origin is the nominal interaction point, and the direction of the z axis is given by the beam axis. Thus, the $x - y$ plane is transverse to the beam direction. In this transverse plane, the x axis points towards the centre of the idealized LHC ring, and the y -axis points upwards.

An alternative description, often used for describing the momenta of final state particles, uses (ϕ, η, p_T) instead of (x, y, z) . ϕ is the azimuthal angle in the transverse plane, $\vec{p}_T$, $\vec{E}_T^{\text{miss}}$, and similar transverse quantities are defined as projections into the transverse plane, and p_T , E_T^{miss} are the corresponding two-dimensional norms of the projection in the transverse plane. The polar angle θ is defined relative to the z axis, which is the beam axis. The rapidity is defined as $y \equiv \frac{1}{2} \ln \frac{E + p_z}{E - p_z}$, and the pseudorapidity is $\eta \equiv -\ln \tan \frac{\theta}{2}$. A common measure of the angular distance between two measured four-momenta (e.g., the momenta of photon or electron candidates) in η, ϕ is defined as $\Delta R \equiv \sqrt{(\Delta\eta)^2 + (\Delta\phi)^2}$. The benefit of the use of ΔR is that it is approximately³ invariant with respect to Lorentz

3: The invariance is exact if and only if the four-momenta are massless. Then, the pseudorapidity equals longitudinal rapidity, differences of which are invariant under boosts in the z -direction.

boosts along the beam axis.

3.2.1 Magnets

A superconducting solenoid around the ID and a set of superconducting toroids in the outer detector comprise the ATLAS magnet system. Their geometric arrangement is shown in Figure 3.3. The magnets give rise to magnetic fields that curve the trajectories of charged particles travelling through them, enabling track based momentum measurements and charge discrimination. Because the momenta of these particles can be quite high, the magnets are cooled to 4.5 K. At this temperature, the magnets material is superconducting, and thus the field strength is high.




Figure 3.3: The ATLAS magnet system [76]. The magnets are shown in red. The inner solenoid is located inside the tile calorimeter, the outermost layer of which plays the role of a return yoke. The other return yokes and the end-cap shielding are not pictured.

4: Which would reduce the resolution of the energy measurements in the calorimeters.

5: A radiation length is the distance an electron will on average travel until interactions with the material have reduced its energy to $1/e$ of the energy it had before entering the material [8].

The superconducting solenoid is a cylinder aligned with the beam axis. Its length is 5.8 m, with a diameter of 2.56 m. It is made from a NbTi alloy stabilized with aluminium, to achieve high field strengths while avoiding putting a lot of material between the ID and the calorimeters⁴. The thickness of the solenoid corresponds to 0.66 electron radiation lengths⁵. The single layer of solenoid coil generates a 2 T axial magnetic field, which deflects the trajectories of charged particles in the ID in the transverse plane. Instead of a dedicated return yoke, the surrounding calorimeters return the flux of the solenoid field.

There are three superconducting toroid magnets: eight coils envelop the barrel, and interleaved with them are two sets of eight coils near the forward and backward end-caps. Multiple scattering in the toroids, which would reduce the precision of the track reconstruction, is reduced by the toroid coils being filled with air. All toroid coils are composed of an aluminium-stabilized NbTiCu alloy. The resulting field bends the trajectories of charged particles traversing the muon spectrometer in R and z . The central barrel toroid has an outer diameter of 20.1 m with a length of 25.2 m along the beam direction, providing a 0.5 T magnetic field within $|\eta| < 1.4$. The end-cap toroids have outer diameters of 10.7 m and are 5 m long. They yield magnetic fields with a strength of 1 T in the $1.6 < |\eta| < 2.7$ region. The overlap region between $1.4 < |\eta| < 1.6$ the fields of the barrel and end-cap toroids are superposed.

3.2.2 Inner Detector

The Inner Detector (ID) is designed to allow precisely reconstructing charged particle tracks with $p_T > 0.5 \text{ GeV}$, which can be used for momentum measurements, particle identification, and vertex reconstruction. The analysis presented in this thesis relies on heavily on the data measured by the ID. ID tracks are crucial inputs for reconstructing resonance candidate jets (see Section 5.2.2). Vertices identified as the intersection of




Figure 3.4: The ATLAS inner detector. Taken from Ref. [82].

several interpolated ID tracks are used both to identify the primary vertex in the event pre-selection (see Section 5.2.4), and as an input feature used to identify the decays of long-lived B-mesons for b -tagging hadronic jets in the top control region selections. Finally, ID-tracks are used in the reconstruction of photons (to reject prompt electrons and recover converted photons) and some muons (if a muon-system track can be matched to an ID track).

The ID is 6.2 m long and has a diameter of 2.1 m, thus covering the region of $|\eta| < 2.5$. Charged particle trajectories are reconstructed by combining the locations where ionization charges the active ID material, indicating the passage of a charged particle, into tracks. Because the ID is entirely contained in the 2 T magnetic field of the solenoid magnet, the tracks are curved. The radius of curvature is used to measure the momentum of the particle leaving the track.

The ID is composed out of three independent tracking subsystems. From inner- to outermost, these are the silicon pixel detectors, the silicon microstrip layers, and the transition radiation tracker. The design of the ID is illustrated in Figure 3.4. The individual tracking subsystems are each separated into a central barrel region, consisting of concentric cylinders, and an end-cap region, consisting of arrays of wheels positioned transversal to the beam line. This is done to optimize the angle of incidence of charged particles originating from the interaction points. This improves the resolution by reducing the material a passing charged particle is exposed to, reducing the effects of multiple scattering and energy loss in the ID.

The silicon pixel and silicon microstrip detectors are constructed from layers of semiconducting silicon. Externally applied voltages generate depletion zones, which are ionized when a charge particle passes through them. The ionization charges are collected and measured as electronic signal, which is converted into a space point based on the location of the microstrip or pixel. The size of the strip or pixel determines the resolution with which the space point is determined. The pixel and microstrip detectors are cooled to temperatures from -5 to -10 °C in order to reduce thermal noise, which is increased by the radiation damage the pixel and

microstrip detectors sustain due to their proximity to the beam pipe.

The silicon pixel detector is located closest to the beryllium beam pipe, to exploit the fact that it has the best resolution of all ID subsystems. It consists of 3 barrel layers, the innermost of which has a distance to the nominal interaction point of 5 cm, and 3 wheels at each end.

The silicon microstrip detector consists of 4 barrel layers and 18 wheels. The microstrip layers are divided into strips $6.4 \text{ cm} \times 80 \text{ }\mu\text{m}$ large. Half of the barrel strip layers are longitudinally aligned to the beam axis, the other half are at an angle of $40 \text{ }\mu\text{rad}$ w.r.t. the beam axis, and are alternated with the aligned strips. In the end-cap wheels, the microstrips are radially aligned with the beam axis, and half have a $40 \text{ }\mu\text{rad}$ angle, as well.

The Transition Radiation Tracker (TRT) consists of straw tubes filled with a gas mixture of xenon, carbon dioxide, and oxygen. The straw tubes are interleaved with polypropylene, and cover the region $|\eta| < 2$. The diameter of the straw tubes is 4 mm, and a gold-plated tungsten wire runs centrally through each straw tube. When a charged particle traverses the TRT, it locally ionizes the gas mixture in the tubes it passes through. The wires, acting as anodes, collect the ionization charges, generating an electronic signal. In the barrel region, the tubes are 144 mm long, with anodes segmented in two halves, and are aligned with the z -axis. In the end-cap, the tubes have a length of 37 cm. The TRT provides no measurement in z , but can greatly enhance the track momentum resolution by yielding additional hits for more accurate track reconstruction. The TRT contributes ca. 36 hits to a typical track in the transverse $R - \phi$ plane, in which the tracks are bent.

The TRT furthermore provides information for electron identification in addition to that provided by the calorimeters. The polypropylene filling the space between the straw tubes gives rise to transition radiation in the form of a γ ray when a charged particle passes the boundary of the tube material and the polypropylene. Electrons can be distinguished from hadrons because the lower electron mass results in a probability to emit transition radiation compared to a charged pion. The increased ionization due to the additional radiation can then be distinguished by applying a higher signal amplitude threshold, which charged pions are unlikely to be able to achieve. Unlike the silicon tracker, the TRT is operated at room temperature.

3.2.3 Calorimeters

The ATLAS calorimeter is designed to completely stop incoming particles, such that they deposit all of their energy. This allows measuring the energy of the incident particle, most crucially in the case of particles (such as neutrons and photons) that are not charged, and thus cannot be




Figure 3.5: The ATLAS calorimeters. Taken from Ref. [83].

measured in the ID. It can also yield improved momentum resolution if incident particles are charged, but have large momenta [84, 85]. In this case, the curvature of the track is small, and thus the track-based momentum measurement may be inaccurate.

In the analysis presented by this thesis, data recorded with the calorimeter is used to identify and measure the energy of photons, which are key to triggering and selecting the targeted ISR-photon-associated final state. The calorimeter data is also crucial in the reconstruction of the energy of hadronic radiation from candidate signal resonance jets (see Section 5.2.2). Finally, calorimeter measurements are used to reconstruct and identify B-meson decays used in the definition of the top control region.

The calorimeter is constructed as a sampling calorimeter, meaning it consists of alternating layers of absorber material and active material. The absorber materials⁶ are dense and have high atomic numbers Z . This increases the chance for incident radiation to interact with the atomic electric fields or nuclei of the absorber. Said interactions generate showers of secondary particles of lower energies. These secondary particles may then shower further themselves in the absorber, until the radiation either leaves the final calorimeter layer⁷, or, more typically, all the initial energy is converted or absorbed. The secondary radiation that is not absorbed enters the layers of active material. There it generates a signal in the by ionization or scintillation light, which allows measuring a fraction of the initially incident energy. This loss to absorption makes it necessary to carefully calibrate the data measured in the calorimeter to identify the energy scale of the originally incident radiation corresponding to the measured signal amplitude.

The design of the ATLAS calorimeter is shown in Figure 3.5. The ATLAS calorimeter is 12 m long, has an outer diameter of 8.5 m, and covers the pseudorapidity range $|\eta| < 4.9$. It consists of three subsystems. The electromagnetic calorimeter, designed to absorb and measure the energy of photons and electrons, is the innermost part of the calorimeter, and described in Section 3.2.3.1. The hadronic calorimeter, designed to absorb and measure the energy of hadrons in the event, surrounds the electromagnetic calorimeter and is described in Section 3.2.3.2.

[84]: ATLAS (2017), *Improving jet substructure performance in ATLAS using Track-CaloClusters*

[85]: ATLAS (2018), *Track assisted techniques for jet substructure*

6: Lead, tungsten, or copper, depending on the particular subsystem.

7: this phenomenon is also called *punch-through*

3.2.3.1 Electromagnetic Calorimeter

The electromagnetic calorimeter (ECAL) is optimized for the absorption of electrons and photons. Other particles will in general also deposit energy there, but will not lose all of it. It uses liquid argon as the active material, in which 1.113 – 1.54 mm thick lead plates are immersed, acting as absorbers. The lead plates are interleaved with copper electrodes. Two electrodes are used to apply a high voltage, and the third electrode provides readout capabilities. The lead plates are shaped like an accordion, avoiding acceptance losses across the ϕ direction due to gaps.

High energy particles entering the ECAL start losing energy via pair production and bremsstrahlung, primarily in the lead absorber, generating a cascade of increasing numbers of particles with reduced energy. When the particles have sufficiently low energies, typically several MeV, they start losing most of their energy via the photoelectric effect and ionization. The ionization charges in the liquid argon are made to drift by the 2000V voltage applied between the first second copper electrode. Capacitive coupling via the third copper electrode then generates an electronic signal.

The ECAL is segmented into a central barrel wrapped concentrically around the beam the solenoid magnet, and a forward and a backward end-cap. The barrel covers $|\eta| < 1.475$, the end-caps cover $1.375 < |\eta| < 3.2$. The ECAL barrel consist of 3 separately sampled segments at different radii, where the innermost segment has the highest cell granularity in η, ϕ and the outermost segment has the lowest. The ECAL endcaps are similarly axially segmented into two separately sampled segments.

3.2.3.2 Hadronic Calorimeter

The hadronic calorimeter (HCAL) is optimized for the absorption of hadrons. These will typically lose only a fraction of their energy in the ECAL, because even if they are charged, their energy loss due to Bremsstrahlung is suppressed by their mass, which is much larger than the electron mass. Consequently, the HCAL is optimized to induce hadrons such as protons, neutrons, pions and kaons to lose energy via inelastic nuclear interactions with the detector material.

The hadrons lose energy in spallation processes, which at high energies yields a cascade of nuclear fragments and secondary particles, forming a hadronic shower. As the hadron energies decrease and the multiplicity increases, the spallation processes tend to more frequently result in hadrons traversing the nuclei, radiating more hadrons. The traversing hadron can also activate the nucleus into an excited state, that later relaxes via fission or evaporation. Since the average rate of nuclear interactions in a material is much lower than that of electromagnetic interactions, hadronic showers are deeper than electromagnetic ones. Finally, decays $\pi^0 \rightarrow \gamma\gamma$,




Figure 3.6: Cutaway view of the ATLAS muon system. Taken from Ref. [76]

causing electromagnetic showers, direct ionization or scintillation by charged hadrons, and indirect ionization by charged nuclei scattered by slow neutrons, allow measuring part of the deposited energy. However, the kinetic energy of escaping neutrinos and muons, as well as the activation energy of nuclei excited into long-lived higher energy states, are not detected.

The HCAL wraps around the ECAL. It comprises the central barrel covering and the extended barrel, together covering $|\eta| < 1.7$, two end-caps covering $1.5 < |\eta| < 3.2$ and two forward calorimeters, extend the coverage by $3.1 < |\eta| < 4.9$.

The central barrel and extended barrel are located far away from the beam pipe. Thus, they are built from less radiation hard but cost-effective material. They use iron plates as absorbers, alternating with scintillating paraterphenyl-doped polystyrene tiles used as the active material. The 1.97 m thick central and extended barrels are built from 64 modules with a size of $\Delta\Phi$ 0.1. Wavelength-shifting fibres transport the scintillation light to the outside of the calorimeter, where photomultipliers measure the scintillation light.

The HCAL end-caps and the forward calorimeter are closer to the beam line, thus they use liquid argon as the active material, and tungsten and copper plates as absorber. The readout is done by capacitive coupling to drifting ionization charges, like in the ECAL.

3.2.4 Muon System

Muons, being much heavier than electrons, lose only very little energy by bremsstrahlung in the atomic fields of the matter in the calorimeters. Furthermore, they do not couple with the strong nuclear force, so their cross-section for interactions with nuclei are strongly suppressed when compared to hadrons in a similar mass range. Thus, muons typically leave the calorimeters keeping most of their kinetic energy.

To identify muons originating from the interaction point, or secondary muons produced in e.g., B -meson decays, the ATLAS detector features a dedicated muon system. In this thesis, the muon system is used to trigger the recording of data entering the top control region (see Section 6.4), and select offline muons to enhance this regions purity.

The muon system is the outermost active part of the ATLAS detector and is depicted in Figure 3.6. It comprises the Thin Gap Chambers (TGCs), the Resistive Plate Chambers (RPCs), the Cathode Strip Chambers (CSCs), and the Monitored Drift Tubes (MDT). The MDTs and CSCs provide precise tracking information for $2.7 < |\eta|$. The TGCs and the RPCs are used for fast triggers on muons in the region with $2.4 < |\eta|$.

The MDTs are three to eight layers of aluminium drift tubes with a diameter of 30 mm per tube. The drift tubes are filled with a gas mixture consisting of argon and carbon dioxide, at a pressure of 3 bar. A central tungsten-rhenium wire with a diameter of $50\text{ }\mu\text{m}$ functions as the anode. Charged particles passing through the drift tubes ionize the gas that fills them, and the ionization charges are collected at the anode wire, generating an electronic signal that indicates a hit.

The CSCs are multiwire proportional chambers. They are installed at $2.0 < |\eta| < 2.7$ in the innermost layer, where the rate and resolution requirements are highest, as dictated by the large expected rate of incident particles. The diameter of the segmented anode wires is $35\text{ }\mu\text{m}$. The CSCs are filled with a gas mixture consisting of carbon dioxide and argon. The charges from gas that has been ionized by a passing charged particle are collected at the anodes and the timing and wire segment information are combined into localized hit information.

The RPCs are parallel electrode plates spaced 2 mm apart, with a potential difference of 9.8 kV between them. They cover $|\eta| < 1.05$, in the barrel region. Muons passing through an RPC ionize the gas in the gap between the plates. The high field in the gap creates an avalanche of secondary ions and electrons that quickly create a large signal when they hit the electrode plates.

The TGCs are multiwire proportional chambers installed in the end-cap regions, covering the $1.05 < |\eta| < 2.7$ region. The distance between the electrodes is smaller than the distance from one anode wire to another, allowing quasi-saturated operation. The TGCs are filled with a highly quenching gas that is a mixture of carbon dioxide and n-pentane.

3.2.5 Trigger

While running, the LHC delivers up to 4×10^7 bunch crossings per second, and tens of proton-proton interactions per bunch crossing are typical. However, only about a thousand full events per second can be written to

the permanent data storage. Thus, the event rate needs to be reduced, which is the task of the trigger system. The trigger system reduces the event rate by selecting events that are deemed interesting, and rejecting events that are not. Several categories of interesting events are defined, such as events with a photon or muon candidate with a transverse momentum above a certain threshold, called a "single photon trigger" or a "single muon trigger", respectively. The ATLAS trigger system uses 2 consecutive stages, or levels. The stages operate at successively slower rates, allowing increasingly refined selections.

The first trigger stage is called the level 1 trigger. It reduces the rate of incoming events to the design rate of 100 kHz. The level 1 trigger is implemented in custom-built hardware so that it can keep up with the incoming event rate. It uses partial information from the calorimeter and muon subsystems to coarsely reconstruct candidate muons, electrons, photons, jets, and hadronically decaying τ leptons, as well as total energy and missing transverse energy. If the transverse momenta of these objects or the energy sums exceed threshold values, the event is accepted. The decision to accept or reject events is taken within $2.5 \mu\text{s}$. Accepted events are read out and then buffered, while being evaluated further in the second trigger stage.

The second trigger stage, called the high level trigger, evaluates events using the full granularity of the detector, but considers only certain regions of interest. Regions of interest are high occupancy regions, indicated by the level 1 trigger. Due to the reduced rate of incoming events, the level 2 trigger is implemented as software running on commercially available hardware. If an event is accepted by the high level trigger, it is written to permanent storage and is available to be used in later analyses. The design rate at which the high level trigger should accept events is 1 kHz.

There can be trigger items with low momentum or multiplicity thresholds that are physically interesting, but occur so frequently that accepting each such event would make the trigger system exceed the design acceptance rates. In such cases, the corresponding triggers are pre-scaled, meaning that only a fraction of the events passing the trigger requirement are accepted. A trigger that is prescaled with prescale S accepts only one out of S events that satisfy its conditions.

3.3 Luminosity

The instantaneous luminosity of a particle collider is a measure of the rate at which collisions occur. It can be calculated from the number of bunches n_B , the particles per bunch $N_{1,2}$ in the bunches of beams 1, 2, the revolution frequency of the beams f and the cross-sectional area of

the beams A :

$$L = \frac{n_b N_1 N_2 f}{A} . \quad (3.1)$$

The corresponding integrated luminosity $L_{int} \int_{t_0}^{t_1} L dt$ is a measure of the number of collisions that occurred in the time interval $[t_0, t_1]$. The concept of luminosity is useful because it decouples the experiment-specific observed number of events with a certain final state n_{process} from the fundamental and universal probability for a process to occur given that a collision happened:

$$\sigma_{\text{process}} = n_{\text{process}} / L_{int} , \quad (3.2)$$

which is called the production cross-section of the process.

In order to determine production cross-sections using Equation 3.2, the luminosity needs to be measured. This is done using several forward detectors providing relative and absolute measurements of the bunch-wise instantaneous luminosity. The Beam Conditions Monitor (BCM) [86] and the LUCID detector are examples of detectors providing relative luminosity measurements, the ALFA detector is an example providing an absolute measurement.

The BCM is located 1.84 m up- and downstream of the interaction point [86]. At both points it consists of 4 diamond sensors with a cross-sectional detector area of 1 cm^2 , arranged in a cross pattern around the beam pipe at a radius of 5.5 cm. The relative luminosity is measured by the rate at which particles traversing the material induce ionization.

The LUCID detector is located ± 17 m up and down from the interaction point. At each point a module containing 16 gas-filled tubes with diameter of 15 mm and a length of 1.5 m, enclosing the beam axis. Particles passing through the tubes generate Cherenkov light. Photomultiplier tubes measure the Cherenkov light. Again, the rate at which particles are detected serves as an indicator of the relative luminosity.

The ALFA detector is placed ± 240 m from the interaction point. At both locations, scintillating fibre trackers are installed inside Roman Pots that can be moved up to 1 mm close to the beam, inside the beam pipe. For the absolute luminosity measurements, special LHC runs are performed, during which the Roman Pots are deployed, allowing measurements of the rate of elastic proton scattering at angles smaller than the nominal beam divergence. This allows determining the total cross-section from the measured elastic forward scattering cross-section via applying the optical theorem. By then applying Equation 3.2, an absolute luminosity measurement can be performed, allowing the calibration of the relative measurements obtained by LUCID and the BCM during regular LHC runs.

It is not necessary to use ALFA to obtain an absolute luminosity measurement for calibration. Instead, the other detectors can also calibrate their interaction rate measurements to the luminosity using displacement scans, in a technique initially developed by van der Meer [87] to determine the Luminosity at the Intersecting Storage Rings at CERN. During a displacement scan the two beams are vertically (horizontally) displaced at different subsequent displacements. By monitoring the change in the observed interaction rates as a function of the displacement, one can extract the beam profile, either (as done at the LHC) by a simple fit of the beam profiles [88], or (as proposed by van der Meer for the Intersecting Storage Rings), by considering a suitable ratio of displacement-integrated observed rates [87, 89]. Knowing the beam profile, Equation 3.1 can be used to calculate the luminosity. f and n_b are known operational parameters, and the N_i can be determined by measuring the beam currents using pickup electrodes placed in proximity to the beam.

ANALYSIS

The search targets the prompt production of a hadronically decaying two-prong resonance back-to-back with an ISR photon. The resonances being considered have masses between 10 GeV and 225 GeV (see Section 5.1 for details about the signal samples). Thus the resonance is assumed to be collimated into a single large-radius jet. The accurate reconstruction, precise calibration, and selection of these jets is crucial for the mass reach performance of the analysis, so the jet reconstruction is non-standard.¹ Sections 5.2 and 5.2.2 describe the jet reconstruction.

1: compared to other ATLAS analyses using large-radius jets

The dominant background is QCD multijet production in association with a photon. This process, together with QCD multijet production where one jet is misreconstructed as a photon, is labelled as the *non-resonant background*. The non-resonant background is estimated purely from data using a combination of control regions (CRs) and a polynomial fit using a method adapted from Ref. [53]. Chapter 6 describes the details of the control region definition. Beyond the non-resonant background, there are resonant background processes involving hadronic $W^\pm$ and Z^0 boson decays. In order from highest- to lowest rate in the fiducial phase-space, these are

[53]: Sirunyan et al. (2019), *Search for Low-Mass Quark-Antiquark Resonances Produced in Association with a Photon at $\sqrt{s} = 13$ TeV*

- ▶ $W^\pm\gamma$, a $W^\pm$ produced with a photon,
- ▶ $Z^0\gamma$, a Z^0 produced with a photon,
- ▶ $t\bar{t}$ where one t decay involves a hadronic $W^\pm$ decay, and the other a $W^\pm \rightarrow e^\pm\nu_e$ decay where the electron is misreconstructed as a photon.

These resonant backgrounds are modelled using Monte-Carlo (MC) simulation for both their shape and their yield, using the samples described in Section 5.1.

The resonant- and non-resonant background estimates are combined into a shared likelihood-based model, increasing the accuracy of the data-driven estimate by subtracting the resonant backgrounds to account for resonant contamination of the CRs. The definition of this likelihood is described in Section 7.1. Simulated signal samples are used to model hypothetical BSM Z' production for a range of signal masses. This signal extraction fit is studied using simulated pseudodata in Chapter 7. E.g. the fit is found to be free from spurious signal² and respond linearly³ to the presence of injected signal. The anticipated sensitivity of the signal extraction fit to the full run-2 dataset is evaluated using simulated pseudodata in Section 7.7. These fit studies are described in Chapter 7.

2: see Section 7.4

3: see Section 7.5

The analysis of observed ATLAS pp collision data using this likelihood model is discussed in Chapter 8. Several aspects of the analyses motivate a very careful unblinding strategy. The search uses novel jet reconstruction

techniques in a low-mass phasespace that has not previously been studied. Furthermore it covers data samples recorded in 2017 and 2018 that have not been used for low-mass dijet resonance searches before. Thus the dataset was not unblinded all at once, but in several steps. Initially, two 3 fb^{-1} data samples were fit using this likelihood to validate the background model, discussed in Section 8.1. Then a slightly larger data samples corresponding to 15.8 fb^{-1} was analysed, the results of which are discussed in Section 8.2. The results of analysing the full 140 fb^{-1} dataset are described in Section 8.3.

Finally, Chapter 9 summarizes the main points of the previous chapters, and quickly outlines potential avenues of further research, based on what has been presented here.

5.1 Samples

5.1.1 Collision data samples

The search analyses proton-proton collisions produced at the LHC and recorded with the ATLAS detector during LHC run 2, from 2015 through 2018, with a centre-of-mass energy of $\sqrt{s} = 13$ TeV. In the signal region and most control regions, recording of collision events is triggered by the detection of at least one highly energetic isolated photon. The trigger¹ requires at least one photon candidate with at least 140 GeV of transverse momentum.

Collision events recorded in the top control region² are triggered using energetic single muons. The muon trigger menu changed during this period [90, 91], so the isolation requirements and p_T threshold of the single muon trigger used depends on the year of data-taking. Table 5.1 lists the triggers used.

Selections to ensure data quality and good detector status are applied as described in Section 5.2.4. After these selections, the sample used in the top CR corresponds to 139 fb^{-1} .

The fit strategy is validated using two small (each corresponding to ca. 3 fb^{-1} of integrated luminosity) samples of pp collision data. This approach is chosen because the rate at which signal may enter the control regions (other than the top control region) is not small enough to be neglected (cf. Chapter 6). The full 2015 dataset is the first validation sample. In this sample, the median of the average number of interactions per bunch crossing μ is 12.4, and μ does not exceed 25. This is not representative of the full Run 2 dataset. Hence, a second validation dataset corresponding to the L4 data taking period in 2018 was used, with a μ up to 55.5. Thus, it can be used to test the robustness of the background model and fit method to high pileup.

Unblinding proceeds in two phases. In an initial step, 15.8 fb^{-1} of observed data recorded in 2015-2016 was unblinded and analysed. This partial sample is composed by successively adding data recorded during single “periods” from 2016 to the 2015 dataset until the total integrated luminosity exceeds 15 fb^{-1} . Only after this initial unblinding has concluded is the full Run 2 dataset recorded from 2015-2018 analysed. This final observed dataset corresponds to an integrated luminosity of $(140.1 \pm 1.2) \text{ fb}^{-1}$ [92], after applying the data quality selections described in Section 5.2.4.

1: HLT_g140_loose

Table 5.1: Triggers used to record top CR data.

Year(s)	Trigger
2015	HLT_mu20_iloose_L1MU15
2016–2018	HLT_mu26_ivarmediumtrigger

2: see Section 6.4

[90]: ATLAS (2016), *2015 start-up trigger menu and initial performance assessment of the ATLAS trigger using Run-2 data*

[91]: Aad et al. (2020), *Performance of the ATLAS muon triggers in Run 2*

5.1.2 Simulated samples

Signal-and background processes are simulated using Monte-Carlo (MC) techniques. Event generators are used to generate events as sets of particle momenta, distributed and weighted according to the (differential) cross-sections predicted by fixed-order perturbative scattering theory. The response of the ATLAS detector to these particles is simulated using GEANT4 [93] and a detailed model of the detector [94].

Different event generators are used to simulate the matrix element of the various hard-scattering processes.

Parton shower, fragmentation, and hadronization are modelled by the Lund string model [95, 96] using PYTHIA 8 [97] with the A14 set of tuned model parameters [98] and the NNPDF2.3LO PDF set [99] for most generators. The SHERPA 2 event generator instead uses its internal fragmentation modelling, with parameters tuned by the SHERPA authors. The event generation for all processes bar one are carried out in the five-flavour scheme³. The exception is the t -channel contribution to single-top production, which uses the four-flavour scheme⁴ instead. EVTGEN 1 [100] models the decays of hadrons containing b -quarks or c -quarks.

The effect of multiple interactions in the same and neighbouring bunch crossings (pile-up) is modelled by overlaying the simulated hard-scattering event with inelastic proton–proton (pp) events. The inelastic events are generated with PYTHIA 8.186 [101] using the NNPDF2.3LO set of PDFs [99] and the A3 set of tuned parameters [102].

The uncertainties on the generated samples are propagated using variations of the weights of affected simulated events. These uncertainties include terms due to the fits determining α_s and the PDF parameters used, and are treated according to the PDF4LHC15 prescription [103]. To estimate the uncertainty from neglecting higher orders of the perturbative series, the renormalisation and factorisation scale parameters μ_r and μ_f are varied up- and down by a factor of 2 independently⁵. The various sets of weights propagating these uncertainties on the simulated templates can have $O(100)$ components. These component variations are combined into a small number of summary variations, using the LHAPDF 6.2.3 program [104] for the combination of PDF-based components.

5.1.2.1 Resonant background processes

Three resonant background processes are simulated for use in the signal region: $W^\pm\gamma$, $Z^0\gamma$, and $t\bar{t}$. These simulated samples are used to create template distributions used in the background estimate and the signal extraction fit. The composition of backgrounds in the top control region is very different from all other analysis regions. Two additional background

3: In the five-flavour scheme, b -quarks are treated as massless quarks and included in the sea-quarks of the proton PDF, which thus contains five flavours of quarks.

4: In the four-flavour scheme, b -quarks are treated as massive fermions produced in the matrix element as a result of explicitly calculated interactions of lighter initial state partons. As b -quarks are not counted among the sea quarks, the proton PDF "contains" only four flavours of quarks.

5: However, the combinations $(2\mu_r, 2\mu_f)$ and $(\mu_r/2, \mu_f/2)$ are not evaluated.

processes are thus simulated to model the data observed in the top control region: single-top production, and $W^\pm$ -associated top production.

The generation of $W^\pm\gamma$ and $Z^0\gamma$ events targets only the fully-hadronic decays of the $W^\pm$ and Z^0 , all other branching ratios are fixed to 0. It uses the SHERPA 2.2.11 event generator to compute scattering matrix elements up to next-to-leading-order (NLO) accuracy in α_s for one additional parton and leading-order (LO) accuracy for emitting up to three extra partons [105, 106]. These virtual NLO-accurate QCD corrections are computed using OPENLOOPS 2 [107–110] in SHERPA 2.2.11. The matrix elements approximate the $W^\pm$ and Z^0 as strictly on-shell, neglecting γ^* amplitudes and interference with $\gamma \rightarrow q\bar{q}$.

SHERPA 2.2.11 computes the parton shower using Catani–Seymour dipole factorisation [111]. The parton shower is then matched to the matrix element results using the MEPS@NLO prescription [112–115]. The generation uses the NNPDF3.0_{NNLO} PDF set [116].

Generator-level cuts are used to reject events that would fail the analysis selections, increasing the efficiency of simulation. In the $W^\pm\gamma$ sample, the hard photon and the on-shell $W^\pm$ are each required to have p_T greater than 140 GeV. For the $Z^0\gamma$ production the photon selection requires $p_T > 15$ GeV and the on-shell Z^0 is required to have $p_T > 100$ GeV.

The Hessian parameterization is used to evaluate the PDF uncertainty [116, 117]. An uncertainty due to the choice of PDF set is estimated by comparing the central function of the nominal set to the central functions of alternative choices. Both samples use the NNPDF3.1_{NNLO} PDF set [118] as one of the alternatives. The $W^\pm\gamma$ sample furthermore uses the MSHT2020_{NNLO} [119] and CTEQ18_{NNLO} [120] PDF sets as alternatives, whereas the $Z^0\gamma$ sample uses CT14_{NNLO} [121] and MMHT2014_{NNLO} [122] instead.

The generation of $t\bar{t}$ events targets both the fully hadronic and semi-hadronic decay channels⁶. The $t\bar{t}$ matrix elements are calculated at NLO accuracy using the POWHEG Box v2 [123–126] generator and the NNPDF3.0_{NLO} [116] PDF.

The parton shower and matrix element calculations are matched using the POWHEG Box method⁷ with $h_{\text{damp}} = 1.5 m_{\text{top}}$.

The cross-sections of the $t\bar{t}$ samples are scaled to match the total inclusive crosssection $\sigma_{t\bar{t}}^{\text{total}} = 832 \pm 51$ pb. This value is calculated using TOP++ 2.0 [128–134] at NNLO accuracy in α_s and NNLL in soft-gluon contributions, for $\sqrt{s} = 13$ TeV, $m_{\text{top}} = 172.5$ GeV. The uncertainty on this total crosssection is evaluated using the PDF set NNPDF2.3_{LO} [99], MSTW2008_{NNLO} [135, 136], and CT10_{NNLO} [137, 138].

Table 5.2: Exact versions of common tools used for $t\bar{t}$ and single-top simulation.

Tool	version
PYTHIA	8.230
EVTGEN	1.6.0

6: However, the semi-leptonic decay is found to contribute most of the expected $t\bar{t}$ event rate after the analysis selections of the signal- and associated control regions are applied.

7: In the POWHEG Box method, the h_{damp} parameter regulates the kinematic dependence of the transition from matrix element to parton shower as emitted radiation transitions from hard to soft [127].

Uncertainties on the generated samples due to the parton shower modelling of ISR and final-state radiation (FSR) are estimated by varying scale and PDF tune parameters [139]. The impact on ISR is estimated by varying μ_r , μ_f , and h_{damp} simultaneously, and using "Var3c" variations of the A14 tune [98]. The impact of FSR is estimated using variations of μ_r for emission.

An additional uncertainty due to the choice of PDF set is assigned by comparing to the central-value replica of the CT14_{NNLO} [121] and MMHT2014_{NNLO} [122] PDF sets, respectively.

The generation of single-top-quark production includes $W^\pm$ -associated production ($tW^\pm$) and uses POWHEG Box v2 [124–126, 140–142] to calculate matrix elements to NLO accuracy in α_s using the NNPDF3.0_{NLO} [116] PDF set. These matrix elements use $\mu_r = \mu_f = m_{\text{top}} = 172.5$ GeV, except for the t -channel component of the single-top process, which uses $\mu_f = \mu_r = \sqrt{m_b^2 + p_{T,b}^2}$ the transverse b -quark mass as suggested in [141]. The $tW^\pm$ final state partially overlaps the $t\bar{t}$ final state, thus the production channels can interfere. This is accounted for using the diagram removal scheme [143, 144]. The decays of top quarks originating from t -channel single-top production are simulated at LO with MADSPIN [145, 146], preserving spin correlations.

Table 5.3: Inclusive single-top production cross-section predictions. See text for details of tools and inputs.

Production process	σ [pb]
t (s-channel)	$6.35^{+0.23}_{-0.20}$
$\bar{t}$ (s-channel)	$3.97^{+0.19}_{-0.17}$
t (t -channel)	$136.02^{+5.40}_{-4.57}$
$\bar{t}$ (t -channel)	$80.95^{+4.06}_{-3.61}$
tW^- and $\bar{t}W^+$	71.7 ± 3.8

The inclusive cross-sections of the s - and t -channel single-top and $W^\pm$ -associated production contributions are individually corrected to match theoretical predictions. All predictions are evaluated for $\sqrt{s} = 13$ TeV and assuming $m_{\text{top}} = 172.5$ GeV. The cross-section predictions for both single-top channels are calculated to NLO accuracy in α_s using the HATHOR 2.1 program [147, 148]. The prediction for the $tW^\pm$ production cross-section is calculated at NLO in α_s and including NNLL soft-gluon corrections, using the calculation in [149, 150]. Table 5.3 details these cross-section predictions.

Table 5.4: Exact versions of common tools used for V +jets simulation.

Tool	version
PYTHIA	8.186
EVTGEN	1.2.0

The generation of multijet-associated V-boson production uses MADGRAPH5_AMC@NLO 2.2.2 [151] at LO accuracy for up to four partons in the matrix element, with NNPDF3.0_{NLO} PDF set [116]. Overlap between matrix element and parton shower radiation is removed by the CKKW-L merging procedure [152, 153] with a merging scale of 30 GeV and jet radius parameter $R = 0.2$.

5.1.2.2 Non-resonant background processes

Simulated non-resonant background samples are used to define and tune the analysis selections, and to generate pseudodata for studying the sensitivity and bias of the signal extraction fit prior to unblinding. The simulated non-resonant samples are not used when analysing the

10 GeV	20 GeV	30 GeV	40 GeV	50 GeV
60 GeV	70 GeV	80 GeV	90 GeV	100 GeV
125 GeV	150 GeV	175 GeV	200 GeV	225 GeV

Table 5.5: Z' masses included in the grid of simulated signal templates.




Figure 5.1: Reconstructed mass of the resonance candidate jet in benchmark signal samples. The year(s) in the individual captions indicate the period of data taking the pileup- and detector conditions of which the simulation is configured to approximate.

observed data after unblinding, however⁸. Two samples are generated: one of multijet events with at least one prompt, hard, and isolated photon, and one of multijet events containing no such photon. When considered together, the two samples are referred to here and in the following as the non-resonant background. Since the non-resonant simulated samples are not used in the signal extraction fit, no uncertainties are estimated for them.

The generation of events with a prompt photon and multiple jets uses SHERPA 2.1 [105] to simulate matrix elements for up to three additional partons. These matrix elements are matched to parton showers using the MEPS@LO method [115]. The generation uses the CT10NLO PDF set [137]. Hadronization is described using the cluster model [154]. Photons from the matrix elements are required to be isolated from hadronic radiation using a smooth-cone isolation criterion [155] with $\delta_0 = 0.3$, $\epsilon_\gamma = 0.025$ and $n = 2$. To enhance the rate at which high- p_T are sampled⁹, the generation is split into separate photon p_T ranges at the generator level. Each such p_T -sliced sample is generated with a similar number of simulated events, and the event weights are reduced to compensate.

The generation of events with multiple jets but no isolated photon uses PYTHIA 8.230 [97] with leading-order matrix elements to simulate dijet production. To efficiently sample the long tail of the jet p_T spectrum, the sample is generated in slices of jet p_T , similarly to the photon simulation.

8: The nonresonant background is estimated using observed data instead, see Chapter 6.

9: thus reducing the Monte-Carlo statistical uncertainty in the tails of distributions

5.1.2.3 Signal processes

Benchmark signal Z' ($q\bar{q}$) [32, 156, 157] production in association with an ISR photon is simulated using the MADGRAPH5_AMC@NLO v2.9.2 [151] event generator for the hard matrix element and the NNPDF3.0_{NLO} PDF set [116]. The results are interfaced with PYTHIA 8.244 [97, 158] to simulate parton shower and hadronisation. The mass of the hypothesized Z' resonance is varied between 10 GeV and 225 GeV, as listed in Table 5.5. The Z' coupling to Standard Model (SM) leptons is set to 0, and the coupling to SM quarks is set to 1/5. The coupling of the Z' field to the dark matter field is set to 1. The mass of the dark matter fermion is set to 10 TeV to ensure the $Z' \rightarrow f_{\text{DM}}\bar{f}_{\text{DM}}$ decay is kinematically inaccessible.

An uncertainty due to the arbitrary choice of PDF set is estimated by comparing to the central values of the CT14_{NNLO} [121] and the MMHT2014_{NNLO} [122] PDF sets.

Figure 5.1 shows the reconstructed mass distributions of the benchmark signal models.

5.2 Selections

5.2.1 Photon selection

Table 5.6: Photon selection criteria.

Feature	Criterion
Pseudorapidity range	$(\eta < 1.37) \quad \quad (1.52 < \eta < 2.37)$
Transverse momentum	$p_T > 10 \text{ GeV}$
Identification	Tight [159]
Isolation	FixedCutTight ("tight isolation requirement" in [159])

10: prompt highly energetic photons may be *converted* to an e^+e^- pair in the material of the inner tracking detector

11: this correction is sometimes called "fudging"

Photons are identified in candidate events to ensure they conform to the targeted final state and are in a region of phase space where the trigger is efficient. Photon candidates are reconstructed from clusters of energy deposits recorded in the electromagnetic calorimeter and tracks recorded in the inner detector. The reconstruction distinguishes unconverted¹⁰ photons, converted photons and electrons using the number and properties of tracks associated to the calorimeter cluster [159, 160]. Photon candidates located in regions of the electromagnetic calorimeter known to have unusually low efficiency due to malfunctioning components are vetoed. In simulated events, observables describing the transverse shape of photon showers are shifted¹¹ to correct for known differences of the simulated and observed shower shape distributions [159].

To reduce the contamination from neutral hadrons such as π^0 or jets depositing significant energy in the electromagnetic calorimeter layers, photon identification requirements are applied. The photon identification

uses the profile of the first two layers of the electromagnetic calorimeter. This search uses the tight photon identification and tight fixed-cut photon isolation criteria specified in [159, 161].

Photons are rejected if their pseudorapidity falls into the transition region between 1.37 and 1.52 where the acceptances of the central barrel of the electromagnetic calorimeter and the end-cap overlap. Table 5.6 summarizes these photon selection criteria.

5.2.2 Jet reconstruction and calibration

Jets are collimated clusters of radiation produced in highly energetic particle collisions [61, 66, 162–169]. While the intuitive impression of jets may appear obvious when observing the distribution of radiation in such a collision, a concrete definition of the jet concept requires specifying an algorithm to determine the association of radiation into clusters¹² [61]. The phenomenon of jets has a long history in the study of such collisions [162, 172], and a great number of jet clustering algorithms have historically been used [61–63, 164, 169, 173, 174]. The choice of algorithm can impact both the multiplicities and constituents of jets one observes [61, 166, 169, 174]. In the analysis of LHC data, the anti- k_T clustering algorithm [62] has emerged as the conventional choice. Anti- k_T clustering can be conceptually understood¹³ as the following iterative procedure:

1. The initial set of pseudo-jets is the set of all input constituents (measured or simulated four-momentum vectors) in the event. The initial set of output jets is the empty set.
2. the anti- k_T distance metric $d_{ij} = \min(p_{T,i}^{-2}, p_{T,j}^{-2})R_{ij}^2$ is computed between any pair of pseudo-jets i and j
3. $d_{iB} = p_{T,i}^{-2}R^2$ is computed for each pseudo-jet¹⁴.
4. the minimum of the current d_{ij} and d_{iB} is found
 - if it is a d_{iB} , the corresponding pseudo-jet is moved from the set of pseudo-jets into the set of output jets.
 - if it is a d_{ij} , the two pseudo-jets are removed from the set of pseudo-jets. Their four-momenta are added, and the sum is inserted into the set of pseudo-jets.
5. the cardinality of the set of pseudo-jets is determined
 - if it is greater than unity, clustering continues from step 2.
 - otherwise, the final remaining pseudo-jet is an output jet, and clustering terminates.

This anti- k_T clustering has attractive properties from both experimental and phenomenological perspectives. Experimentally, the observation that the expected catchment area [166] of anti- k_T jets is a circle of radius R allows for straightforward calibration of pile-up and similar

12: However, an explicit clustering algorithm is not necessary in order to analyse events involving jets [162, 170, 171].

[61]: Salam (2010), *Towards Jetography*

[62]: Cacciari et al. (2008), *The anti- k_t jet clustering algorithm*

13: Computationally efficient implementations used in practice [165, 167] differ substantially from the algorithm discussed here. However, the resulting assignment of constituents to clusters is identical.

14: $R > 0$ is a parameter of the algorithm, which users must choose to fully specify the clustering.

[166]: Cacciari et al. (2008), *The Catchment Area of Jets*




Figure 5.2: Narrow $R = 0.2$ jet energy scale calibration stages. Taken from [177].

15: Infrared- and collinear (IRC) safety is a property of observables. It holds if neither the addition of radiation of infinitesimally low momentum into any direction, nor the splitting of any input into two collinear inputs of the same total momentum, change the value the observable takes. Thus, IRC-safety ensures that calculations using the safe observable are consistent even when affected by hadronisation, cancelling divergences in perturbative QCD, and experimental effects such as angular granularity and momentum thresholds [61, 169].

Table 5.7: Narrow jet reconstruction.

Parameter	Value
Algorithm	anti- k_T
R -parameter	0.2
Input constituent	LCTopo

16: The resulting jets are also called LCTopo.

area-dependent effects [64]. Furthermore, the simplicity of the definition allows a coherent comparison of jets at the level of particles and the (simulated) detector responses, which is also useful for calibration. Phenomenologically, this ability to apply a single consistent jet definition to both particle simulation and reconstruction is similarly a benefit when comparing predictions to observations. Crucial to the consistency of such comparisons is the IRC-safety¹⁵ of the definition. Together with its simplicity, it furthermore makes the resulting jets amenable to analytic calculation [61, 169, 175, 176]. To fully define an anti- k_T jet clustering, the parameter R must be chosen. The best choice depends on the problem at hand. In the remainder of this thesis, three values are used for different tasks: a large one ($R = 1$), an intermediate one ($R = 0.4$), and a small one ($R = 0.2$).

Large-radius (large- R) jets are reclustered out of narrow jets and inner detector tracks.

Narrow jets are clustered from topoclusters [178] using the anti- k_T algorithm with $R = 0.2$ [62]. The topocluster energy is corrected using the local cell signal weighting (LCW) method¹⁶ [178]. The reconstructed $R = 0.2$ jets are calibrated to the truth jet energy using simulated Monte-Carlo (MC) events [177]. Residual differences of the MC-calibrated jet energy scale to the energy scale observed in collision data are corrected using calibrated $R = 0.4$ EMTopo reference jets [179]. This correction is applied only to $R = 0.2$ jets reconstructed in collision data. Figure 5.2 shows an overview of the calibration stages. The $R = 0.2$ jet energy resolution is measured in data and simulation. In regions of jet $p_T \times \eta$ where the observed resolution is greater than simulated, the momenta of simulated jets are smeared such that the effective resolution matches that observed in data. In regions where the data resolution is smaller than simulated, pseudodata are smeared to propagate the difference, which is taken as an additional systematic uncertainty on the simulated jet energy resolution. The reconstruction and calibration parameters for narrow jets are summarized in Table 5.7.

Inner detector tracks [180, 181] are pre-selected to ensure reconstruction quality, vertex association, and association to any small- R calorimeter jet. These selections are detailed in Table 5.8. The jet association requirement is chosen such that it has a strictly larger acceptance than the jet association

[64]: Cacciari et al. Jet reconstruction using jet area




Figure 5.3: TAR jet reconstruction algorithm. Adapted from [85].

used in the large- R reclustering described below.

Large- R jets are used to reconstruct Z' resonance candidates in the analysis. They are reconstructed using the track-assisted-reclustered (TAR) algorithm [85]. The TAR algorithm combines track-based and calorimeter-based measurements, enabling substructure reconstruction that uses the superior angular resolution of the tracks, and the more precise energy resolution of the clusters. The inputs to the reconstruction are calibrated $R = 0.2$ jets (see Tables 5.7 and 5.11) and tracks of charged particles observed in the inner detector. The $R = 0.2$ jets are reclustered into $R = 1$ anti- k_T -jets using the FASTJET software library [167]. Inner detector tracks are ghost-associated [166] to the $R = 0.2$ subjects. The reclustered $R = 1$ jets are trimmed [66], removing subjects with p_T less than $f_{\text{cut}} = 0.05$ of the reclustered jets p_T . Inner detector tracks associated to trimmed-away subjects are reassociated to the subject nearest the track in $\Delta R(\text{jet}, \text{track})$, unless the nearest subject has a distance greater than $\Delta R_{\text{max}} = 0.3$ to the track. Tracks without nearby subjects after trimming are discarded. The p_T of tracks is then rescaled such that the sum of the p_T of tracks associated to a given $R = 0.2$ jet matches the p_T of the jet. Finally, the TAR jet is taken as the set of rescaled tracks. Figure 5.3 summarizes the algorithm.

[85]: ATLAS (2018), *Track assisted techniques for jet substructure*

Systematic uncertainties on the jet energy scale and resolution are propagated from the calibrated $R = 0.2$ input jets. Uncertainties on track reconstruction are negligible compared to those on the $R = 0.2$ jets, and are

Tracking algorithm	Primary [180]
Quality selection	Loose [180]
p_T	$> 0.5 \text{ GeV}$
Vertex association	$z_0 \sin(\theta) < 3.0 \text{ mm}$ and $d_0 < 2.0 \text{ mm}$
Jet association	$\Delta R < R_{\text{anti-}k_T} + 0.13$

Table 5.8: Track pre-selection criteria. Track-to-jet association accepts a track if it is associated to at least one $R_{\text{anti-}k_T} = 0.2$ anti- k_T LCTopo [177, 178], $R_{\text{anti-}k_T} = 0.4$ anti- k_T EMTopo, or $R_{\text{anti-}k_T} = 0.4$ EMPFlow jet [179].

Parameter	Value
Algorithm	anti- k_T
R-parameter	1.0
Input constituents	$R = 0.2$ LCTopo jets tracks
Grooming algorithm	Trimming
f_{cut}	0.05
R_{trim}	0.2

Table 5.9: Large- R jet reconstruction. The input $R = 0.2$ LCTopo jets are reconstructed as specified in Table 5.7 and selected following the criteria in Table 5.11.

therefore not propagated to the TAR jets. The reconstruction parameters for large- R jets are summarized in Table 5.9.

17: cf. e.g. [182, 183]

By propagating TAR jet uncertainties in this *bottom-up* fashion from sufficiently calibrated inputs, the calibrated jets can be used much more flexibly than with the conventional calibration¹⁷, which is applied to the final reconstructed large- R jet. Conventional late calibration does not usually apply outside the phase space wherein the calibration was derived, adding substantial effort to the exploration of unconventional phase space (such as very low large- R jet masses). By contrast, bottom-up calibrations- and uncertainties are much more likely to be easily transferable, because unconventional phase space for the high-level object (e.g. $R = 1$ TAR jet) may not appear unconventional as far as the observables of the lower-level calibrated objects (e.g. $R = 0.2$ LCTopo jets) are concerned. This is in fact the situation encountered in this analysis for very low TAR-jet masses at moderate-to-high transverse boost: the relevant calibrated observables are the transverse momenta of the constituent $R = 0.2$ LCTopo jets, which at typical values of ca. 70 GeV – 150 GeV are well within their calibrated range. Thus the use of TAR jets to reconstruct signal resonance candidates is crucial to enabling the analysis of resonances lighter than 50 GeV presented in this thesis¹⁸.

18: The conventional ATLAS calibrations [182, 183] only extend to 50 GeV.

5.2.3 Jet selection

Table 5.10: Large- R jet selection criteria. J refers to the candidate large- R jet. γ refers to any photons in the same event selected according to Table 5.6. γ_0 refers to the photon in the event with the largest p_T .

Observable	Requirement
$ \eta $	< 2
$\frac{2m}{p_T}$	$< R = 1$
p_T	$> 150 \text{ GeV}$
$\Delta R(J, \gamma_0)$	> 1
$\Delta \phi(J, \gamma_0)$	$> \frac{\pi}{2}$
Photon isolation	$\Delta R(J, \gamma) \geq 1.2$ $(\gamma \text{ s.t. } \frac{p_T^\gamma}{p_T} > \frac{1}{10})$
$n_{\text{const.track}}$	≥ 3

Table 5.11: Narrow jet ($R = 0.2$ anti- k_T) selection criteria.

Observable	Requirement
p_T	$> 15 \text{ GeV}$
$ \eta $	< 2.2

Narrow $R = 0.2$ jets are reconstructed as per the configuration shown in Table 5.7 and selected according to the criteria in Table 5.11. Requiring a narrow jet p_T of at least 15 GeV ensures the jets are in the calibrated range where systematic uncertainties are available [177].

Large- R jets are reconstructed as described in Section 5.2.2 and Table 5.9. Jets are only accepted in the central pseudorapidity region where at least part of the jet area is covered by the inner tracker acceptance, and the input calorimeter jets are well-measured. Badly collimated jets ($\frac{2m}{p_T} > R$) are rejected to avoid susceptibility to edge effects due to the limited angular catchment of the jets. A minimum p_T cut ensures candidate jets are hard enough to balance the ISR photon candidate p_T . Similarly jets are rejected if they are too aligned with the leading photon (the ISR

candidate) in $\Delta R(J, \gamma_0)$ and $\Delta\phi(J, \gamma_0)$ since the signal resonances are expected to predominantly be produced back-to-back with the ISR photon. Hard photon candidates other than the ISR candidate in proximity to or overlapping with a large- R jet would make the interpretation of ISR photon candidate and resonance candidate jet ambiguous, hence such jets are rejected. Finally at least three constituent rescaled tracks are required to ensure jet shape observables are well defined. Table 5.10 summarizes the $R = 1$ jet selection criteria.

5.2.4 Event selection

Observable	Requirement
Photon multiplicity	≥ 1
$p_T^{\gamma_0}$	$> 150 \text{ GeV}$
Large- R jet multiplicity	≥ 1
Second jet veto	$\nexists j$ such that $\Delta R(J_0, j) > 1.2$ and $p_T^j > p_T^{\gamma_0}$
$\rho^{J_0} = 2 \log \left(\frac{m^{J_0}}{p_T^{J_0}} \right)$	> -5.4

Table 5.12: Event selection criteria. Photons are selected according to Table 5.6. γ_0 refers to the photon with the greatest p_T . Large- R (narrow) jets have been selected according to Table 5.10 (Table 5.11). j refers to any narrow jet in the event, J_0 to the large- R jet with the greatest p_T .

Observed collision events are filtered to ensure data quality [184]. Reasons why events would be vetoed by these filters include:

- noise bursts in the liquid argon calorimeter leading to data corruption [185];
- issues with tile calorimeter operation [186];
- the recovery procedure of the semiconductor tracker for single event upsets [187];
- an event not having been recorded completely.

Both observed and simulated events are required to have a primary vertex with at least two associated tracks. The primary vertex is selected as the one with the largest Σp_T^2 , where the sum is over all tracks with transverse momentum $p_T > 0.4 \text{ GeV}$ that are associated with the vertex.

Table 5.12 details the further event selection criteria. Requiring at least one hard photon and a large- R jet passing the jet criteria in Table 5.10 selects events that conform to the ISR-photon associated jet production topology. Events with one or more additional hard jet(s) are vetoed since they correspond to higher orders in the perturbative prediction of the QCD cross-section and are thus challenging to model well. Very-low-mass highly collimated jets were found to introduce challenging QCD multijet mass spectra, and are thus rejected using a cut on ρ .

[184]: Aad et al. (2020), *ATLAS data quality operations and performance for 2015–2018 data-taking*

[185]: Aad et al. (2014), *Monitoring and data quality assessment of the ATLAS liquid argon calorimeter*

[186]: Aaboud et al. (2018), *Operation and performance of the ATLAS Tile Calorimeter in Run 1*

[187]: Aad et al. (2014), *Operation and performance of the ATLAS semiconductor tracker*

Background estimate

6

6.1 Strategy

6.1.1 Background processes

The dominant background consists of events where a hard photon is produced in association with multiple QCD jets (e.g. Figure 6.1). Furthermore, multijet production without a real associated photon can enter the preselection (see Section 5.2) if a jet is misidentified as a photon. This misidentification rate is small, but due to the overwhelming QCD multijets production cross-section such events nevertheless enter the selection with a non-negligible rate. Multijet production of either kind is referred to as the *non-resonant background*.

Resonant electroweak production of hadronic jets in association with a photon (e.g. Figure 6.2(a)) is a genuine example of the targeted signal topology, but predicted by the SM. It is thus inherently irreducible and must be modelled. The highest-rate of these processes is the photon-associated production of a hadronically decaying $W^\pm$. The $W^\pm\gamma$ cross-section is ca. two orders of magnitude smaller than that of photon-associated multijet events due to the smaller effective electroweak coupling. Thus $W^\pm\gamma$ is a sub-dominant background in all phasespace probed by the search. A further electroweak background is the hadronically decaying Z^0 in association with a photon. $Z^0\gamma$ is a very similar irreducible resonant background, but with a rate even lower than $W^\pm\gamma$.

In $t\bar{t}$ events where one top quark decays hadronically ($t \rightarrow W(\rightarrow qq')b$) and the other top decay involves an electron or positron ($t \rightarrow W(\rightarrow e\nu)b$), the electron can be misreconstructed as a photon (see Figure 6.2(b)). Such events can in rare cases enter the preselection. If the large- R jet reconstruction captures most of the products of the hadronic $W^\pm$ decay, it will show a two-prong structure and likely enter the signal region. If the large- R jet reconstruction contains the full hadronic top decay, it exhibits a three-prong structure. Such events still enter the preselection, but are more likely to end up in a control- rather than the signal region. Thus the $t\bar{t}$ background shape is different from that of $W^\pm\gamma$ or $Z^0\gamma$, and requires dedicated modelling. The effective rate of the $t\bar{t}$ background is even smaller than that of $Z^0\gamma$. It is the lowest-yield background modelled in the search.

Taken together, $W^\pm\gamma$, $Z^0\gamma$, and $t\bar{t}$ are referred to as the *resonant background*.




Figure 6.1: Example feynman graphs contributing to the dominant non-resonant background from photon-associated multijet production.




Figure 6.2: Example leading order feynman graphs contributing to the resonant background from hadronically decaying $W^\pm\gamma$ and $Z^0\gamma$ (a) or semi-leptonically decaying $t\bar{t}$ (b).

6.1.2 Data-driven estimate of non-resonant background

Table 6.1: Selections defining the signal region and control regions. A control region is defined for each combination of inverting a selection (cf. Figure 6.3). All preselections described in Section 5.2 are applied before any analysis region selection. See Sections 6.2 and 6.3 for the definition of D_2^{DDT} and comparative studies of the forward-photon selection, respectively.

Regions	Selection requirements	
	D_2^{DDT}	$ \eta_\gamma $
Signal region (SR)	< 0	< 1.3
Central-photon CR	> 0	< 1.3
Tagged CR	< 0	> 1.3
Forward-photon anti-tagged CR	> 0	> 1.3

The dominant non-resonant background is challenging to model using simulation. The reasons are two-fold:

1. Achieving sensitivity requires selecting events where the resonance candidate jet has a two-prong substructure. However, the non-resonant background efficiency of such a selection depends intricately on details of the parton shower and hadronization steps of the simulation, which are challenging to model.
2. Monte-carlo simulations should use number of Monte-Carlo events that is many times¹ greater than the expected data yield [188]. Otherwise, the uncertainty from noise in the Monte-Carlo sample becomes a limiting systematic uncertainty. For the large multijet cross-section and the full Run 2 luminosity, this requirement makes a sufficiently precise Monte-Carlo estimate unfeasible.

1: O(10-100) as per ref. [188]

[188]: ATLAS (2020), *Recommendations for the Modeling of Smooth Backgrounds*




Figure 6.3: The relation of the signal region and the three control regions used to estimate the non-resonant background, as described in the text. See Table 6.1 for the cuts defining the regions.

To avoid the above challenges, a fully data-driven estimate of the non-resonant background is used.

The non-resonant background estimate hinges on being able to control the (non-resonant) background efficiency of the two-prong cut. It proceeds as follows:

1. The two-prong substructure observable is decorrelated from the mass of nonresonant large- R jet resonance candidates using simulation, as detailed in Section 6.2. This results in a selection with an efficiency that is constant as a function of mass, at least in simulation.
2. A CR is constructed using the same preselections as the SR, but inverting the two-prong substructure requirement.
3. Since the two-prong tag has been mass-decorrelated, it has approximately constant efficiency ϵ_{SR} as a function of mass in simulated samples. Assuming no mismodelling, the observed non-resonant yield in mass bin i in the SR (N_{SR}^i) can thus be predicted by rescaling the observed non-resonant yield in the CR from step 2 (N_{CR}^i):

$$N_{\text{SR}}^i = N_{\text{CR}}^i \frac{\epsilon_{\text{SR}}}{1 - \epsilon_{\text{SR}}} \quad (6.1)$$

4. In practice there is mismodelling, and thus the true efficiency in data likely has a different value $\epsilon'_{\text{SR}} \neq \epsilon_{\text{SR}}$. Furthermore mismodelling

may introduce a mass dependence of the efficiency. The ratio of efficiencies is thus modified by a mass-bin dependent data scale factor τ^i

$$\frac{\epsilon_{\text{SR}}^{'i}}{1 - \epsilon_{\text{SR}}^{'i}} = \frac{\epsilon_{\text{SR}}}{1 - \epsilon_{\text{SR}}} \tau^i \quad (6.2)$$

which is assumed to at most vary slowly as a function of jet mass and be $\mathcal{O}(1)$ due to most of the detailed (mis-) modelling effects on $\epsilon_{\text{SR}}^{'i}$ cancelling in the ratio.

5. The nonresonant efficiency ratio scale factors τ^i can already be determined in the signal extraction fit using a functional form fit². However to achieve as robust as possible a fit, the τ^i parameters³ should be constrained outside the SR. To this end, the forward-photon CR is constructed as described in Section 6.3, and the SR and above CR restricted to the central photon pseudorapidity. The forward-photon selection is split according to passing/failing the two-prong selection of the SR and thus allows an independent constraint of τ^i . Figure 6.3 schematically illustrates this region definition, Table 6.1 lists the selections used.

2: see Equations 7.7 and 7.8

3: or equivalently, the function parameters predicting the τ^i

6.1.3 Simultaneous subtraction of resonant background

The resonant background processes $W^\pm\gamma$, $Z^0\gamma$, and $t\bar{t}$ contribute most strongly to the SR. There, their contribution must be independently constrained, since it is otherwise indistinguishable from a signal. Thus both the shape and normalization of the resonant background mass spectra are modelled using MC simulation (see Section 5.1 for details of the simulations).

However, resonant background also enters the CRs, albeit at a lower rate relative to the non-resonant background. But since the non-resonant background is estimated using the observed mass spectra in the CRs, the background estimate needs to account for this contamination, lest it be biased. Thus the amount and distribution of resonant background contamination in the CRs is simulated, and subtracted⁴ from the observed CR spectra.

4: in the signal extraction fit this is implemented by simultaneously fitting correlated templates (see Section 7.1)

Equation 6.2 has the freedom to fit residual differences between the in-situ pass-to-fail ratio of two-prong tagging non-resonant jets, but similar differences may also exist for resonant jets resulting from genuine two-parton decays. The calibrations applied to observed and simulated jets (see Section 5.2.2) should implicitly account for most of these effects. However, the two-prong tagger itself is not explicitly calibrated. To ensure a sufficiently accurate modelling of the two-prong tag efficiency on resonant two-prong jets in particular, a further CR is constructed. This CR aims to collect a pure sample of two-prong jets from hadronic

decays of top quarks in data, to independently validate and constrain the resonant two-prong-tagging efficiency. It is described in Section 6.4.

6.2 Fixed-efficiency cut method for a mass-decorrelated two-prong tag

Non-resonant QCD jets constitute the majority of the background entering the signal region. To reduce this background, events accepted into the signal region must pass a cut on a two-prong tagging observable. Two-prong tagging observables are observables that are designed to differentiate jets resulting from the decay of a single hard parton (such as a gluon), from jets that result from the decay of a hadronic resonance into a pair of partons (e.g., $Z' \rightarrow q\bar{q}$) [70–72, 189, 190]. By selecting events using a two-prong tagging observable, relatively more of the resonant signal events enter into the signal region. Thus the search is more sensitive to new resonances beyond the SM.

In general, the jet mass spectrum of a sample passing such a two-prong tagging cut is different from that of the inclusive sample of jets with any tagger value [73, 191, 192]. An example of this shaping of the mass distribution is shown in Figure 6.4. This is because two-prong taggers tend to be correlated to the jet mass. However, the tagger-passing and tagger-failing samples are assumed to have the same jet mass distribution in the fit used to extract the strength of the signal process. Thus, it is necessary to modify the tagger observable such that background jets that pass the tagging cut have the same mass distribution as the inclusive background sample.

6.2.1 Measuring jet mass shaping using the Jensen-Shannon distance

The tagger decorrelation method described in this section is repeatedly tuned. When tuning hyperparameters in this way, it is very useful to have a single metric to quantify the effect of the tuning, enabling a quick comparison of alternative parameter choices. The study discussed here aims to minimize the mass shaping of the tagger selection, therefore a metric designed to measure differences between the shapes of distributions, the Jensen-Shannon distance [193], is used⁵.

The Jensen-Shannon distance is used to quantify how much some cut on some two-prong tagging observable shapes the jet mass distribution. It is the square-root of the Jensen-Shannon divergence denoted⁶ JSD , thus

5: A distance is used, rather than a test of compatibility. This is by design: the aim is to measure how different two distributions are, not to test their compatibility. Supposing two equally compatible decorrelated distributions, the smoother one is still preferred in this context. Appendix A discusses this point further.

6: to align with the notation in references [191, 194]




Figure 6.4: Distortion of the jet mass spectrum due to a N_2 [72] selection in simulated samples of non-resonant (a) and resonant (b) jets. In both (a) and (b), the topmost third of the plots shows the mass distribution of the inclusive sample and compares it to the shape of the N_2 -tagged distribution. A well-decorrelated tagger would show the two distributions closely matching in (a) while enhancing the W peak in (b). The middle third shows the bin-wise average of the toys binned in the binning used for the resampling using toys, to illustrate which features the resampled distributions are sensitive to, for both the tagged and resampled cases. The resampling is motivated and described in the text. The lower third shows the efficiency of the N_2 selection as a function of jet mass, for both the finely binned input histograms and the mean of the rebinned toys. A well-decorrelated tagger would show a flat efficiency in (a), and simultaneously an increase in efficiency around the W resonance mass in (b). The distance between the tagged and inclusive distributions is measured using the Jensen-Shannon distance $\sqrt{JSD}$, annotated in the top margins of (a) and (b). $\sqrt{JSD}$ is defined in Equation 6.3 and its use as a metric of jet mass shaping is discussed in Section 6.2.1.

we denote the Jensen-Shannon distance as $\sqrt{JSD}$. It is defined as [193]

$$\sqrt{JSD(p, q)} = \sqrt{\frac{KL(p, \frac{p+q}{2}) + KL(q, \frac{p+q}{2})}{2}}, \quad (6.3)$$

where $p = P(X)$ and $q = Q(X)$ are vectors of probabilities of distributions P and Q over a shared set of events X in general. In this particular application, P is the probability density of the inclusive jet mass distribution, and Q that of the tagged jet mass distribution (or vice versa). p and q correspondingly are the tagged and inclusive jet mass histograms normalized to unity. Figure 6.4 shows examples of such mass histograms. $KL(p, q)$ in Equation 6.3 means the Kullback-Leibler divergence [195] or relative entropy

$$KL(P(X), Q(X)) = - \sum_{x \in X} P(x) (\log_2 Q(x) - \log_2 P(x)). \quad (6.4)$$

Using $\sqrt{JSD}$ is preferred here because, unlike JSD , it fulfils the triangle inequality, allowing for an intuitive interpretation as a metric distance between probability distributions [196], i.e., the distance between tagged and inclusive jet mass spectra.

The lowest value $\sqrt{JSD(q, p)}$ can take is 0, corresponding to $q = p$. However, this study uses Monte-Carlo methods to approximate the probabilities q and p , and so there is a practical limit to the expected $\sqrt{JSD}$ due to the limited size of the Monte-Carlo sample. This limit can be estimated for the MC samples used in this analysis by calculating

7: the mass distributions are rebinned to bins that are larger than the mass resolution of the jets prior to the bootstrap. This is done in order to ensure that neighbouring bins may be approximated as uncorrelated in the bootstrap.

the expected $\sqrt{SD}$ of the inclusive non-resonant background mass distribution computed from 10 000 bootstrapped and rebinned⁷ inclusive distribution which were varied according to the expected bin-wise variance of the number of events entering the bin due to the Monte-Carlo sample size. For the non-resonant background, this limit is estimated to be $\sqrt{SD_{\text{stat.limit}}} = 0.0027 \pm 0.0006$.

The rebinning procedure used in the calculation of the Monte-Carlo statistical lower bound means that the probability space over which the $\sqrt{SD}$ is evaluated is different from the one of the original binning. Hence, the limit can not meaningfully be compared to the $\sqrt{SD}$ evaluated with the original binning. To allow for a meaningful judgement of how close to the statistical limit the $\sqrt{SD}$ is, it is re-computed using 10 000 bootstrapped and rebinned distributions, and the expected value over all bootstraps is quoted in figures such as Figure 6.4 as $\sqrt{SD_{\text{resampled}}}$. Such figures illustrate the rebinned distributions using the average of all 10 000 bootstrapped distributions, which is labelled *resampled* in the lower two panels.

$\sqrt{SD}$ and $\sqrt{SD_{\text{resampled}}}$ are used when deriving the results presented in the following to compare the performance of different decorrelation method tuning choices. They are furthermore used to compare the residual shaping exhibited by different tagging observables after decorrelation, and to judge how close a particular choice of tagger and decorrelation method gets to the theoretical limit on decorrelation performance.

6.2.2 Fixed efficiency cut method

A mass-correlated two-prong tagger selection on the jet mass spectrum can be interpreted as having a cut efficiency that depends on the jet mass, as illustrated in the third panel of Figure 6.4(a). This interpretation suggests the mass correlation can be reduced by choosing tagger working points that depend on the jet mass, such that the background efficiency of the cut is constant with respect to the jet mass. This method, referred to as a *designed decorrelated tagger* (DDT), was first described in Ref. [197], and then used and studied in refs. [53, 60, 191, 194]. As in these references, the tagger working point is chosen in bins of $\rho = 2 \log(m/p_T)$ and p_T instead of just the jet mass m , because to first order, the jet mass is proportional to p_T for QCD jets [61].

Based on this, a two-prong-tagger X may be decorrelated from the mass of background jets for a background tagger efficiency q by these four steps:

1. Measure the distribution of X in bins of $\rho \times p_T$ in the (simulated [60, 191, 194, 197] or observed [53]) background dataset;
2. Identify the q -th percentile of the X distribution, X_q , in each $\rho \times p_T$ bin;

[197]: Sirunyan et al. (2018), *Search for low mass vector resonances decaying into quark-antiquark pairs in proton-proton collisions at $\sqrt{s} = 13$ TeV*

3. Smooth the measured percentiles in the (ρ, p_T) plane to reduce the effects of statistical fluctuations on the cut;
4. Compute the decorrelated observable as $X^{\text{DDT}}(X, \rho, p_T) = X - X_q^{\text{smoothed}}(\rho, p_T)$.

A selection such as $X^{\text{DDT}}(X, \rho, p_T) < 0$ will then have the fixed background efficiency q , irrespective of the mass or p_T of the jet. In the following, this method is applied using a simulated background dataset and a 13 % target background efficiency.

To smooth the measured percentiles in step three, the observed distribution of percentiles is smoothed with a Gaussian filter. The Gaussian filter convolves the binned bivariate distribution with a bivariate Gaussian kernel, with different kernel width σ along the ρ and p_T directions (analogous to the standard deviation of a Gaussian distribution), where the kernel width is measured in units of bins.

The parameters for the binning and smoothing of the percentile maps are manually tuned to achieve good mass decorrelation. The simulated sample used to construct the decorrelated tagger and also to measure its decorrelation performance is affected by noise due to small MC sample size. This presents a challenge to the tuning: care must be taken to avoid tuning to spurious features of the noise. To avoid tuning the parameters of the decorrelated tagger construction to sample noise, the simulated sample is randomly split into three orthogonal subsamples. The three subsamples contain 50%, 25%, and 25% of the full sample, respectively. The larger subsample is used to construct percentile maps, and the cuts from the smoothed maps derived from the larger subsample are applied to the two smaller subsamples, labelled “test sample 1” and “test sample 2” in the following.

By separating the events used in the construction of the percentile map from those used to measure the decorrelation performance, this method ensures that the tuning is not biased to the noise in the sample used to construct the decorrelated tagger. The use of two independent test samples allows comparing the two tagged spectra to check for the presence of common features that might indicate insufficient smoothing of noisy inputs. Furthermore, it reduces the risk of tuning the decorrelated tagger to be optimal for spurious features of one particular test sample.

Based on studies of the smoothing on the two different shards, the following kernel widths are chosen. The kernel width in the p_T direction is $\sigma_{p_T} = 1.8$ bins, and $\sigma_\rho = 1.6$ bins in the ρ direction. The high- p_T region of the percentile distribution is more affected by noise due to limited MC sample size. To counteract this, bins with $p_T > 750$ GeV take their smoothed percentiles from a convolution with a kernel that is more spread out in the ρ direction, using $\sigma_\rho = 2.6$ bins. The kernel is truncated at 4σ in either direction. To avoid bias near the edges of the binned range in ρ and p_T , the bins at the edges are repeated before computing the

kernel convolution, to cover the support of the kernel in the respective direction.

The performance of this mass decorrelation method depends strongly on the choice of binning in the (ρ, p_T) plane. If the bins are too large, some jets with very different kinematics are treated with a common tagger cut, whereas potentially very similar jets on either side of a bin edge are treated very differently. This may lead to significant deviations of the efficiency measured as a function of the mass from the expected constant. These issues are avoided by choosing sufficiently small bins, which is part of the decorrelation parameter tuning. Bins are tuned to be as large as possible to minimize the impact of statistical fluctuations of the measured quantiles due to limited MC sample size, but still small enough to avoid bias due to too-large bins. As with other tuning aspects, this absence of bias is explicitly monitored in the DDT-tagged mass spectrum, and the final binning choice shows no such bias. The final tuned p_T bin edges are 43 geometric-equidistant edges inclusively spanning 150 GeV through 643 GeV, then 7 geometric-equidistant edges that inclusively span 713 GeV through 1200 GeV. The tuned ρ bin edges linear-equidistantly span -8 through -1 inclusively, in steps of 0.05.

6.2.3 Measured quantiles and fixed-efficiency cut maps


The quantiles of D_2 are shown in Figure 6.5 to illustrate the method. To quantify how much the smoothing procedure changes the quantile maps shown in this section, the residual of the smoothing is calculated as

$$\text{residual}(\rho, p_T) = \frac{X_q^{\text{smoothed}}(\rho, p_T) - X_q^{\text{measured}}(\rho, p_T)}{|X_q^{\text{measured}}(\rho, p_T)|}. \quad (6.5)$$


The residual is small in the phase-space region where the background cross-section is high. The residual appears randomly distributed, which is consistent with the expectation that it smoothes out localized fluctuations and discontinuities, but maintains the large-scale quantile map that was measured.

6.2.4 Mass shaping of decorrelated taggers


An example of the mass shaping that remains after the two-prong tagging observables have been decorrelated is shown in Figure 6.6. As seen in the figure, the mass sculpting of the tagged distribution is significantly reduced compared to a fixed cut on D_2 , while still improving the signal shape. Both test samples show similar levels of sculpting, indicating that the fluctuations are statistical, and not induced through the smoothing.



(a) Smoothed cut map at fixed efficiency.



(b) Measured quantiles.



(c) Residual of the smoothed w.r.t. the measured quantiles.

Figure 6.5: Smoothed map of D_2 [198] cuts at a fixed non-resonant background efficiency of $\epsilon_{\text{bkg.}} = 0.13$ (a). White-shaded areas indicate regions of the cutmap that later analysis selections on ρ (> -5.4) and $2m/p_T$ (≤ 1) reject, meaning they are not actually used in practice. The white (grey) lines indicate the exponential curves that are lines of constant mass in the (ρ, p_T) plane. The cutmap is smoothed using a bivariate Gaussian filter (see Section 6.2.2) from a measured distribution of D_2 quantiles (b). The residual (c) of the smoothed cutmap w.r.t. the measured quantiles is computed as in Equation 6.5. The remaining mass shaping due to selecting on D_2^{DDT} decorrelated using this cutmap is shown in Figures 6.6(a) and 6.6(b).




Figure 6.6: Distortion of the jet mass spectrum due to a D_2^{DDT} [198] selection in simulated samples of non-resonant (a), (b) and resonant (c), (d) jets. The two test samples used (test sample 1 in (a), (c) and test sample 2 in (b), (d)) are split using a fully orthogonal random selection that yields two independent samples of the same size, as described in more detail in the text. Section 6.2.2 describes the method of decorrelation by using a map of fixed efficiency cuts. The cutmap used in the selection illustrated in this figure is shown in Figure 6.5(a). $\sqrt{JSD}$ is defined in Equation 6.3 and its use as a metric of jet mass shaping is discussed in Section 6.2.1. Cf. similar visualisations for a not-decorrelated tagger in Figure 6.4.

6.2.5 Summary of tagger performance for different observables

Table 6.2 summarizes the level of residual shape differences achieved by the two-prong observables considered. The mass shaping of all decorrelated taggers is very low. No single tagger stands out as particularly less shaping after decorrelation, and the differences in $\sqrt{JSD}$ between taggers are of a similar size as the differences between test samples for the same tagger. In general, the $\sqrt{JSD}$ of the bootstrapped toys is lower due to the coarser mass binning as mentioned previously. Nevertheless, also for $\sqrt{JSD_{\text{resampled}}}$ the differences in shaping between different taggers are not statistically significant relative to their bootstrapped uncertainty. Thus, all taggers have $\sqrt{JSD_{\text{resampled}}}$ similarly close⁸ to its statistical limit (given the shape of the underlying distribution and the size of the samples) of 0.0027 ± 0.0006 . This indicates that there is some room left for small

8: within 2-3 standard deviations of the difference

Tagger	ref.	$\sqrt{JSD}$		$\sqrt{JSD_{\text{resampled}}}$	
		test sample 1	test sample 2	test sample 1	test sample 2
C_2	[71]	0.0068	0.0066	0.0049(8)	0.0046(7)
$D_2^{(1,2)}$	[72]	0.0063	0.0063	0.0051(8)	0.0044(7)
$D_2^{\beta=1.7}$	[71]	0.0071	0.007	0.0052(8)	0.0053(8)
D_2	[198]	0.0061	0.0064	0.0046(8)	0.0041(7)
M_2	[72]	0.0059	0.0061	0.0046(7)	0.0043(7)
N_2	[72]	0.0067	0.0058	0.0047(8)	0.0044(7)
τ_{21}	[70]	0.0065	0.007	0.0042(7)	0.0047(7)

Table 6.2: Comparison of residual mass shaping by various taggers after mass-decorrelating as described in Section 6.2.2. $\sqrt{JSD}$ and $\sqrt{JSD_{\text{resampled}}}$ are calculated as described in Section 6.2.1 using the definition of $\sqrt{JSD}$ of Equation 6.3.

improvements to the method to achieve decorrelation at the theoretical limit, and this seems to be independent of the tagger used. However since all taggers have residual shaping within a few standard deviations of the limit, large gains in decorrelation performance are not possible without substantially increasing the size of the underlying simulated sample.

Tagger	ref.	$\epsilon_{W^\pm}^{\text{peak}}$	
		test sample 1	test sample 2
C_2	[71]	0.29 ± 0.01	0.30 ± 0.01
$D_2^{(1,2)}$	[72]	0.22 ± 0.01	0.22 ± 0.01
$D_2^{\beta=1.7}$	[71]	0.29 ± 0.01	0.29 ± 0.01
D_2	[198]	0.32 ± 0.01	0.32 ± 0.01
M_2	[72]	0.20 ± 0.01	0.20 ± 0.01
N_2	[72]	0.25 ± 0.01	0.26 ± 0.01
τ_{21}	[70]	0.27 ± 0.01	0.29 ± 0.01

Table 6.3: On-peak $W^\pm$ identification efficiency of various taggers after mass-decorrelation as described in Section 6.2.2. The highest efficiency, highlighted in **bold**, is attained by D_2 . The indicated uncertainty is due to the finite size of the MC samples.

The $W^\pm$ -tagging performance of the different candidate decorrelated taggers is evaluated by considering the efficiency of tagging the $W^\pm$ in samples of simulated $W^\pm\gamma$ events. The efficiency is evaluated in bins of the reconstructed jet mass⁹. There is an evident mass dependence of the efficiency, which is maximal near 90 GeV (the $W^\pm$ mass pole) and smoothly falls to both lower- and higher masses until plateauing close to the efficiency of 13 % targeted by the decorrelation for non-resonant jets. Table 6.3 shows the highest efficiency reached by every decorrelated tagger in any mass bin, for two independent simulated samples. Unlike for the residual mass shaping, different taggers show significantly different performance differences, which are consistent between the two samples. Hence, the highest-efficiency tagger is chosen for the construction of the SR and CRs (cf. Section 6.1 and Figure 6.3), which is D_2^{DDT} with a constant 13 % simulated non-resonant background efficiency, and up to 32 % simulated $W^\pm$ efficiency.

9: See Figures 6.6(c) and 6.6(d) for an example.

6.3 Forward photon control region

In order to estimate the non-resonant QCD background to the ISR photon and boosted large-radius jet channel, a control region (CR) is designed

to contain a signal-region like sample of events where the relative rate of resonant jets is reduced. This CR is defined using a fixed selection independent of the signal hypothesis considered, as described in [53]. This section compares candidate observables to select on, as well as possible values to select at. The candidate observables are chosen which are expected to not depend on the shape or structure of the resonance candidate jet. Their dependence on jet p_T and mass is explicitly studied, and a good final choice should be independent of the jet mass in particular. A well-designed final control region thus should impact different resonant processes similarly. Furthermore the signal contribution in the CR should be reduced, and the non-resonant mass spectrum not shaped by the selection.

6.3.1 Candidate observables

Several potential observables are considered, including the pseudorapidity of the large- R jet (η_J), the pseudorapidity of the ISR photon (η_{ISR}), the pseudorapidity difference between the large- R jet and the ISR photon ($|\Delta\eta_{\text{ISR},J}|$), $y_{\text{ISR},J}^* = |y_{\text{ISR}} - y_J|/2$, where y is the rapidity, $\chi_{\text{ISR},J} = \exp(|y_{\text{ISR}} - y_J|)$, and the ΔR between the large- R jet and the ISR photon ($|\Delta\eta_{\text{ISR},J}|$). The comparisons of the distribution of these observables for the non-resonant background, the resonant backgrounds, and the signal are shown in Figure 6.7.

The η_J distribution in Figure 6.7(a) shows inconsistent trends of the shape difference of the different resonant processes to the dominant non-resonant background. The $Z'\gamma$ process has a relatively higher rate at larger $|\eta_J|$, but the $W\gamma$ process has a relatively reduced rate, and the $Z\gamma$ process has a differential cross-section in η_J consistent with that of the non-resonant process. Similarly, the $\Delta R_{\text{ISR},J}$, shown in Figure 6.7(f), has inconsistent trends of the relative shapes. Thus η_J and $\Delta R_{\text{ISR},J}$ are not good observables to use in the construction of the control region.

The η_{ISR} distribution in Figure 6.7(b) shows a consistent trend of all resonant processes having enhanced rates at low absolute values, and reduced rates at high absolute values, relative to the non-resonant process. This consistent trend makes a cut on $|\eta_{\text{ISR}}|$ a viable candidate for the construction of the control region.

Similarly, $|\Delta\eta_{\text{ISR},J}|$ distribution, shown in Figure 6.7(c), exhibits a clear trend where the resonant processes have a higher rate at low values relative to at high values, when compared to the non-resonant process. The detailed distributions are different, but this trend is the same, when considering the distributions of the related observables $y_{\text{ISR},J}^*$ shown in Figure 6.7(d), and $\chi_{\text{ISR},J}$, shown in Figure 6.7(e). Thus, functions of the absolute (pseudo-) rapidity difference are also considered viable candidates to use to construct the control region.




Figure 6.7: Distributions of candidate observables considered in the construction of a control region with minimal correlation with the kinematics and shape of the resonance candidate jet.




Figure 6.8: Control region composition resulting from constructing the CR using a single selection on several candidate observables.

6.3.2 Purity and contamination of control region

In addition to showing consistent trends across the different processes, it is also important that the signal contamination is lower in the CR than in the signal region (SR). This separation can be quantified by the CR purity and resonant contamination in the CR. This separation is shown in Figure Figure 6.8 for the four observables where the kinematic distributions are considered viable (η_{ISR} , $|\Delta\eta_{\text{ISR},J}|$, $y_{\text{ISR},J}^*$, and $\chi_{\text{ISR},J}$).

Selecting events for the CR using η_{ISR} can decrease the contamination by resonant $W\gamma$ and $Z\gamma$ events from ca. 0.012 to about 0.009, as shown in Figure 6.8(a). In the central barrel region the selection has a stronger relative effect on purity and contamination than in the more end-cap region at higher values.

The effect of using $\Delta\eta_{\text{ISR},J}$ and similar observables to select events for the CR is shown in Figures 6.8(b) to 6.8(d). While the dependence of purity and contamination on exact cut values depends on the observable used, all cover a similar range. Using $\Delta\eta_{\text{ISR},J}$ -like observables, the resonant contamination may be reduced down to 0.002 - 0.001. This is significantly smaller than the smallest contamination (0.009) achievable by selections on η_{ISR} . Thus, as far as CR purity is concerned, a cut on $\Delta\eta_{\text{ISR},J}$, $y_{\text{ISR},J}^*$, or $\chi_{\text{ISR},J}$ would be best.

6.3.3 Bias on jet kinematics in control region

Control region purity is not the only concern in designing the forward-photon CR. The shaping of the jet kinematic distributions by the CR selection must also be considered. Figure 6.9 illustrates the shaping effects of the absolute pseudorapidity of the ISR photon and the pseudorapidity difference between the ISR photon and the resonance candidate jet. The non-resonant processes are by far the dominant contribution, especially in the CR. Thus the concern about shaping is really one of shaping of the non-resonant mass spectrum, which the CR is used to help estimate. Thus in the study of kinematic shaping due to the CR selections, only the non-resonant processes are considered.

The shaping effects of the candidate cut observables on the mass distribution of non-resonant jets is shown in Figure 6.9. In general the binned mass spectra are very similar across all bins in either cut observable. However, the $\Delta\eta_{\text{ISR},J}$ binned distributions in Figure 6.9(b) exhibit stronger shaping than the η_{ISR} binned ones in Figure 6.9(a), especially near the edges of the studied ranges. Thus the η_{ISR} cut is the preferred way to construct the control region when minimal mass shaping is prioritised over other concerns.

Figure 6.9 illustrates the shaping effects of the candidate cut observables on the jet p_T distribution. Whereas the p_T shaping when binning in $\Delta\eta_{\text{ISR},J}$, shown in Figure 6.9(d) is comparable to the mass shaping shown in Figure 6.9(b), the p_T shaping when binning in η_{ISR} is much more pronounced, as Figure 6.9(c) illustrates.

$\rho = \log(m^2/p_T^2)$ is computed from both the jet mass, which is barely shaped by selections on $|\eta_{\text{ISR}}|$, and the jet p_T , which is strongly shaped. However, this combination, outside of the upper end of the distribution, shows very little shaping, as Figure 6.9(f) illustrates. As in the case of the jet mass, this shaping is somewhat stronger when selecting on the rapidity difference instead, as Figure 6.9(e) shows.




Figure 6.9: Shape of the jet mass distribution in bins of η_{ISR} (a), and in bins of $\Delta\eta_{\text{ISR},J}$ (b). Shape of the transverse jet momentum distribution in bins of η_{ISR} (c), and in bins of $\Delta\eta_{\text{ISR},J}$ (d). Shape of the jet $\rho = \log(m_j^2/p_{T,J}^2)$ distribution in bins of η_{ISR} (e), and in bins of $\Delta\eta_{\text{ISR},J}$ (f). The lower panels depict the binned distributions relative to the distribution summed over all bins of η_{ISR} or $\Delta\eta_{\text{ISR},J}$, respectively.

6.3.4 Choice of control region selection

To summarize, wide intervals of possible selections on several observables are considered for constructing the second set of CRs. Two observables are able to consistently separate out events with non-resonant jets, and are studied in more detail: $|\eta_{\text{ISR}}|$ and $|\Delta\eta_{\text{ISR},J}|$. Two other options, $y_{\text{ISR},J}^*$ and $\chi_{\text{ISR},J}$, have very similar behavior to $|\Delta\eta_{\text{ISR},J}|$, and so are not studied further. Both of $|\eta_{\text{ISR}}|$ and $|\Delta\eta_{\text{ISR},J}|$ distinct advantages: the $|\eta_{\text{ISR}}|$ selection results in minimal shaping on the jet mass distribution, while the $|\Delta\eta_{\text{ISR},J}|$ selection yields much reduced jet p_T shaping, and allowing for a more purely non-resonant CR.

The observable chosen to define the forward set of CRs is η_{ISR} . While both mass shaping and CR purity are expected to impact the sensitivity of signal extraction fits, mass shaping of the CR selection does so by potentially invalidating comparisons of distributions observed in the affected regions. Thus, the worst case of mass shaping in the forward CRs may risk the validity of the analysis. An overly signal-contaminated CR on the other hand is in the worst case misnamed, acting in effect like second SR, so long as the likelihood is able to model the contamination consistently. The risk of invalidating the search is judged to be worse than that of mislabelling regions, therefore η_{ISR} is chosen to construct the so-called forward-photon CR.

To choose the η_{ISR} selection separating the SR from the forward-photon CR, two aspects are considered.

1. Relative expected event yield in the SR as opposed to the CR: To be able to provide a useful constraint under likelihood maximization, the CR should contain many more events than the SR.
2. Control region purity: a less signal contaminated CR carries less risk of biasing a fit of a likelihood that fails to model signal contamination perfectly.

The resonant vs. non-resonant separation of η_{ISR} is insufficient to satisfy both considerations with one selection (cf. Figures 6.7(b) and 6.8(a)). As a compromise between contamination and statistical power, the CR selection is chosen at¹⁰ $|\eta_{\text{ISR}}| > 1.3$ (cf. Section 6.1 and Figure 6.3).

10: The proximity of this cut to the barrel/end-cap transition region beginning at $|\eta| \approx 1.37$ is incidental.

6.4 Top control region

6.4.1 Overview




Figure 6.10: Semi-leptonic $t\bar{t}$ topology targeted by the top CR. The illustration shows a conceptual view in the transverse plane through the interaction point, indicated by the concave polygon. The green curve dashed arrow represent the muon track and neutrino (escaping the detector without interacting) resulting from one $W^\pm$ decay, the thin orange circle sectors indicate jets containing B-meson decays, and the wide circle sector represents the two-prong jet resulting from the W decaying into light hadrons.

In order to measure the efficiency scale factor of the D_2^{DDT} selection on two-prong jets, a control region targeting semi-leptonic $t\bar{t}$ decays is defined, as sketched in Figure 6.10. This top control region allows studying the properties of the tagger on the known resonant hadronic decay of a $W^\pm$ to a pair of quarks. The leptonically decaying top allows selecting a sample of pp collision data that has a high purity of such hadronic $W^\pm$ decays. By comparing the observed resonant $W^\pm$ yield to the expected yield, the D_2^{DDT} selection efficiency on two-prong-like jets can be constrained in-situ and in simulation. Furthermore, the agreement of the peak location in the observed and expected mass distributions validates the reconstruction and calibration of the $R = 1$ TAR jets.

6.4.2 Object reconstruction and selections

Muons are reconstructed by combining tracks identified independently in the muon spectrometer and the inner tracking detector [199, 200]. Muon momenta are corrected for effects from detector misalignment [201] and calibrated using observed muon pairs produced at the Z^0 and J/Ψ resonances [199, 200, 202].

Table 6.4: Muon selection criteria.

Criterion	Selection	Identification requirements are applied to suppress fake muon candidates
Identification	Medium [199, 200]	originating from the decays of hadrons such as pions and kaons [199, 200].
Isolation	PFlowLoose [200]	Isolation requirements reduce the number of muons overlapping with
p_T	$> 30 \text{ GeV}$	significant energy deposited in the calorimeter from neutral radiation not
$ \eta $	< 2.5	associated with any tracks, using the variable-cone-radius variant of the
$d_0/\sigma(d_0)$	< 3	PFlowLoose requirement [200]. Muon tracks are required to be consistent
z_0	$< 0.5 \text{ mm}$	with originating from the primary vertex by selecting candidates with

11: The distance from the primary vertex to the point of closest approach of the extrapolated trajectory to the beam line, projected onto the direction of the beam line.

12: The distance in the transverse plane from the beam position to the point where the extrapolated trajectory most closely approaches the beam line.

Table 6.5: $R = 0.4$ anti- k_T EMPFlow jet selection.

Observable	Requirement
p_T	$> 25 \text{ GeV}$
$ \eta $	< 2.5
JVT score [203]	> 0.59

longitudinal impact parameter¹¹ $z_0 < 0.5 \text{ mm}$. A further selection is applied using the longitudinal impact parameter¹² d_0 , but it is applied in terms of the d_0 significance: $d_0/\sigma(d_0) < 3$. The minimum muon p_T requirement of 30 GeV ensures either of the muon triggers used to record the observed data is fully efficient (cf. Table 5.1), and suppresses multijet QCD background. This multijet suppression is strong enough that the multijet background can be neglected entirely, simplifying the modelling of the top CR using simulated data. The muon selections are summarized in Table 6.4.

$R = 0.4$ EMPFlow jets are clustered using the anti- k_T algorithm [62] from inner-detector tracks and topological clusters [178] of energy deposited in the calorimeter. A particle flow algorithm [204] is used to subtract the energy accounted for by tracks from corresponding calorimeter

clusters prior to anti- k_T clustering [179, 205, 206]. The reconstructed momenta of the jets are calibrated using MC simulated data such that they better match the true jet momenta. An additional calibration is applied to jets in observed collision data. It corrects for differences of the MC-calibrated jet energy scale to the energy of calibrated reference objects¹³ using a momentum-balance method. Differences between simulated and observed jet resolutions are mitigated by smearing the energy of jets in simulated samples. Jets are required to have $|\eta| < 2.5$ and $p_T > 25$ GeV, ensuring they are in the phase-space covered by the calibration and reducing the rate of pile-up jets. Pile-up jets are further suppressed by applying jet vertex tagging (JVT) [203], requiring a JVT score > 0.59 for jets with $p_T < 60$ GeV, $|\eta| < 0.4$.

b -tagged $R = 0.4$ EMPFlow jets are $R = 0.4$ jets identified as containing a B-hadron decay by a neural network-based method (DL1r) [207, 208]. This method makes use of differences in the distribution of tracks, presence of secondary vertices, and energy distribution associated with the jet due to the longer lifetimes and higher masses of such hadrons. A fixed cut is placed on the final score, resulting in a typical efficiency of 85 %.

Large- R TAR jets are reconstructed as described in Section 5.2.2 and selected similarly to the pre-selection applied to the $R = 1$ jets entering the SR (see Tables 5.10 and 5.12), to ensure a comparable sample of jets is selected. Additional isolation requirements are applied to the TAR jet selection to avoid overlap with muons or the radiation from the bottom quark originating from the muonic top decay. TAR jets are rejected if any (including soft and non-isolated candidates) muon is found within $\Delta R \leq R_{\text{TAR}} + R_{\text{isolation gap}} = 1.05$. Similarly, the isolation requirement rejects TAR jets if there is a b -tagged $R = 0.4$ jet within $\Delta R \leq R_{\text{TAR}} + R_{\text{PFlow}} + R_{\text{isolation gap}} = 1.45$ of the jet. This is done because the aim of the study is to measure the D_2^{DDT} efficiency in a sample of jets originating from hadronic $W^\pm$ decays. TAR jet overlapping the b -tagged jet are likely to either contain the full top decay, or have missed part of the $W^\pm$ decay. Neither of these cases matches the targeted decay topology. Table 6.7 lists the combined selections.

6.4.2.1 Event selection

Observable	Requirement
isolated μ multiplicity	$= 1$
$R = 0.4$ EMPFlow jet multiplicity	≥ 3
b -tagging multiplicity	≥ 2
$R = 1$ jet multiplicity	≥ 1
$R_{\text{isolation gap}}$	0.05

13: Photons in $\gamma + j$ events or leptons in $Z^0 + j$ events for a lower-energy jet, or multiple lower-energy jets for a higher-energy jet.

Table 6.6: b -tagging selection criteria. See Table 6.5 for the input jet criteria.

Algorithm	DL1r
Efficiency	85 %

Table 6.7: $R = 1$ jet selections applied in the top CR

Observable	Requirement
p_T	> 150 GeV
$ \eta $	< 2
$\frac{2m}{p_T}$	$< R = 1$
$\rho = 2 \log \left(\frac{m}{p_T} \right)$	> -5.4
$n_{\text{const.track}}$	≥ 3
$\Delta R(J, \gamma)$	$> R = 1$
$\Delta R(J, \mu)$	> 1.05
$\Delta R(J, b - \text{jet})$	> 1.45

Table 6.8: Event selections for the top control region.

The top CR selection targets $t\bar{t}$ events with one muonic top decay, and aims to suppress backgrounds from multiple non-resonant jets with or without associated photons. Table 6.8 summarizes the top CR selection. Exactly one isolated muon is required to be reconstructed offline, reducing both the multijets background and the fully hadronic and fully leptonic $t\bar{t}$ decay channels. Three or more $R = 0.4$ anti- k_T EMPFlow jets are required in the event. Of these, at least two b -tagged jets are required, enhancing the proportion of events where top quarks have been produced. One of the $R = 0.4$ jets is required to be located within $\Delta R < 1.5$ of the muon, in order to select events with a somewhat boosted and therefore collimated $t \rightarrow bW(\rightarrow \mu\bar{\nu})$ decay. At least one $R = 1$ TAR jet is required in the event. The highest- p_T TAR jet is taken as the $W^\pm \rightarrow q\bar{q}$ candidate, and used to study mass distributions and D_2^{DDT} efficiency.

6.4.3 Jet mass distributions




Figure 6.11: Mass distribution of the p_T -leading $R = 1$ TAR jets in events passing the top CR selection described in Section 6.4.2. Visualization courtesy of Jad M. Sardain.

Figure 6.11 shows the TAR mass distribution of the $W^\pm$ candidate jets after the top CR pre-selections. The resonance peak at the $W^\pm$ mass is clearly visible, both in the observed full run-2 data and the stacked simulated distribution templates. The location of the $W^\pm$ peak is where it is expected from previous measurements, at about 80 GeV [11]. The location and shape of the peak on the mass scale are the same in the observed data sample as in the sum of simulated templates. This agreement between data and simulation validates the TAR jet reconstruction in-situ. In particular, the input $R = 0.2$ LCTopo jet calibration is sufficient, and no additional calibration at the level of TAR jets is required.

[11]: Workman et al. (2022), *Review of Particle Physics*

The total predicted number of events in Figure 6.11 is roughly 25 % greater than the observed number. This difference is comparable to previous studies of similar simulated samples targeting semi-leptonic $t\bar{t}$ production in the run 2 dataset collected by ATLAS in Ref. [209]. Ref. [209] found that simulation overestimates the $t\bar{t}$ production cross-section in this channel by roughly 12 %. This is less extreme than what is observed here, but this top CR selection selects top-quarks with substantially lower p_T , and Ref. [209] report that the mismodelling they observe is strongly p_T dependent, and worse at lower p_T . Thus, the observed discrepancy likely is due to this previously reported cross-section mismodelling.

[209]: Aad et al. (2022), *Measurements of differential cross-sections in top-quark pair events with a high transverse momentum top quark and limits on beyond the Standard Model contributions to top-quark pair production with the ATLAS detector at $\sqrt{s} = 13$ TeV*




Figure 6.12: Mass distribution of the p_T -leading $R = 1$ TAR jets in events passing (a) or failing (b) the D_2^{DDT} selection. All events in either distribution are pre-selected according to the top CR selection described in Section 6.4.2. Cf. the D_2^{DDT} -inclusive distribution in Figure 6.11. Visualizations courtesy of Jad M. Sardain.

The effect of the D_2^{DDT} selection /used to define the SR and related control regions (cf. Table 6.1) on the jet mass distribution in the top CR is illustrated by Figure 6.12. The processes involving the hadronic decay of a $W^\pm$, $t\bar{t}$ and single-top, have their $W^\pm$ peak enhanced by the SR-like selection and suppressed by the CR-like selection. Conversely, the $W(\mu\nu)$ +jets and $Z(\mu^+\mu^-)$ +jets processes, where the TAR jet is composed of non-resonant radiation, have the same mass distribution in either selection, which is the same as in the inclusive distribution in Figure 6.11. This confirms the expectation of mass-decorrelation based on the construction of D_2^{DDT} in Section 6.2.

Finally, the enhanced $W^\pm$ peak in Figure 6.12(a) exhibits a clear mass-dependence of the disagreement between simulated and observed cross-sections. This makes clear that there is a need to measure the scale-factor

of the D_2^{DDT} efficiency in two-prong resonances, such that it can be corrected in the SR.

6.4.4 D_2^{DDT} scale factor for two-prong resonances

To extract a D_2^{DDT} scale factor for only the resonant processes, a shape parameterization is fit to the D_2^{DDT} -passing and inclusive mass distributions. The parameterization combines a polynomial shape for non-resonant backgrounds and a normal-distributed component peaking at the $W^\pm$ mass. Concretely,

$$v_i(N, m_{W^\pm}, \sigma, A, B, C) = \frac{N}{\sigma\sqrt{2\pi}} \exp\left(-\frac{(m_i - m_{W^\pm})^2}{2\sigma^2}\right) + A(m_i)^2 + Bm_i + C, \quad (6.6)$$

where v_i is the number of predicted events in mass bin i , m_i is the central value of mass bin i , and the number of resonant events is N . The parameterization Equation 6.6 is fit to the observed and simulated mass distributions of the D_2^{DDT} -tagged and inclusive samples by minimizing the mean squared error of the parametrized prediction w.r.t. each distribution.

The D_2^{DDT} efficiencies for both data and Monte Carlo are calculated as $\varepsilon = N^{\text{tagged}}/N^{\text{untagged}}$, where N^{tagged} is the number of resonant jets that passed the tagger and N^{untagged} is the number of resonant jets before the tagger was applied, both determined using the fit of Equation 6.6. Hence, the scale factor is extracted as $\text{SF} = \varepsilon^{\text{data}}/\varepsilon^{\text{MC}}$ where $\varepsilon^{\text{data}}$ is the efficiency in observed data and ε^{MC} in simulation. The final scale factor is $\text{SF} = 0.999 \pm 0.025$.

This extracted scale factor is applied as a multiplicative yield correction to all simulated resonant signal and background templates in the SR and forward-photon tagged CR defined in Section 6.1. To propagate the impact of the scale factor coherently to the anti-tagged CRs, the relation

$$\text{SF}^{\text{CR}} = \frac{1 - \varepsilon_{\text{SR}}^{\text{MC}} \text{SF}^{\text{SR}}}{1 - \varepsilon_{\text{SR}}^{\text{MC}}} \quad (6.7)$$

is applied, with $\varepsilon_{\text{SR}}^{\text{MC}}$ the D_2^{DDT} efficiency of the corresponding analysis region selecting D_2^{DDT} -passing events determined from simulation. Note the relative sign between the SR and CR scale factors, expressing the correct anti-correlated behaviour for its modifications to resonant event rates in D_2^{DDT} passing and failing regions. This parameterization allows the signal extraction fit to treat the scale factor as a single nuisance parameter (SF^{SR}).

7.1 Likelihood construction

The presence of hypothetical signal is extracted using a maximum-likelihood fit. The same fit simultaneously estimates the non-resonant background¹, and adjusts the simulated templates of the resonant backgrounds. The ability of the fit to perform these tasks results from the form of the likelihood it maximizes. The likelihood presented in this section is numerically implemented and evaluated using the `pyhf` software library [210, 211].

1: cf. Section 6.1

The likelihood is of the usual form of a product of bin-wise poisson distribution probabilities

$$\begin{aligned}
p(N, \mu, \theta) &= \prod_{r \in \text{regions}} \prod_{i \in \text{mass bins}} p(N_r^i, \mu, \theta) \\
&= \prod_{r \in \text{regions}} \prod_{i \in \text{mass bins}} \frac{e^{-v_r^i(\mu, \theta)} (v_r^i(\mu, \theta))^{N_r^i}}{N_r^i!}
\end{aligned} \tag{7.1}$$

with N the vector of observed event yields across analysis regions (see Table 6.1), and mass bins, μ the signal strength ² parameter of interest (POI), θ the vector of nuisance parameters (NPs), and $v_r^i(\mu, \theta)$ the expected event yield given parameter values μ, θ . Here and in the following, *bins* refer to the mass bins defined in Table 7.1.

2: signal yield relative to the nominal
yield of the signal template

In the SR, the expected number of events in mass bin i is

$$v_{\text{SR}}^i(\mu, \boldsymbol{\theta}) = \mu v_{\text{SR, signal}}^i(\boldsymbol{\theta}_{\text{signal}}) + \frac{\epsilon_{\text{SR}}}{1 - \epsilon_{\text{SR}}} \tau^i v_{\text{central}}^i + \sum_{s \in \{W^\pm \gamma, Z^0 \gamma, t\bar{t}\}} v_{\text{SR}, s}(\boldsymbol{\theta}_s), \quad (7.2)$$

i.e. the sum of the nominal SR signal template $v_{\text{SR, signal}}^i$ scaled by the POI μ , the nonresonant background contribution parameterized according to the outline in 6.1, and the sum of the nonresonant SR background templates. Nuisance parameters in θ distribute across these component samples as follows: Common experimental (and some theoretical) uncertainties are shared between the sub-vectors θ_s where $s \in \{\text{signal}, W^\pm\gamma, Z^0\gamma, t\bar{t}\}$, but some (e.g. scale variations to estimate uncertainty due to missing higher

Table 7.1: Mass bin edges in GeV.

[illegible]




Figure 7.1: Structure of unconstrained nuisance parameters for the non-resonant background estimate. Two independent sets of shapefactors are shared among the regions with a central (forward) photon requirement, respectively (a). Cf. $v_{\text{central } \gamma}^i$ and $v_{\text{forward } \gamma}^i$ in Equations 7.2, 7.3, 7.5 and 7.6. One shared set of transfer factors is shared by only the two-prong tagged analysis regions (b). Cf. τ^i in Equations 7.2 and 7.5, also equivalently $f(m_j^i)$ in Equation 7.8. See Table 6.1 for the cuts defining the regions.

orders in the perturbative expansion) are specific to a given process. Finally $v_{\text{central } \gamma}^i$ and τ^i are unconstrained nuisance parameters (and thus elements of θ on the left-hand side) but completely disjunct from the nuisance parameters θ_s that affect the simulated templates.

In the central-photon anti-tagged CR, the expected yield per mass bin i similarly is:

$$v_{\text{central } \gamma \text{ CR}}^i(\mu, \theta) = \mu v_{\text{central } \gamma \text{ CR, signal}}^i(\theta_{\text{signal}}) + v_{\text{central } \gamma}^i + \sum_{s \in \{W^\pm \gamma, Z^0 \gamma, t\bar{t}\}} v_{\text{central } \gamma \text{ CR, } s}(\theta_s). \quad (7.3)$$

The relevant difference between Equation 7.3 and Equation 7.2 is the treatment of the non-resonant background. In the CR, only $v_{\text{central } \gamma}^i$ enter — it is not affected by the τ^i . Furthermore, the same $v_{\text{central } \gamma}^i$ are shared between Equation 7.3 and Equation 7.2. This reflects the constructive assumptions Equations 6.1 and 6.2.

The CR-related Poisson factors of the likelihood are more constraining in maximization because roughly 6.7 as much nonresonant background yield enters the CR relative to the SR. Hence, in the large-sample-limit, the total expected best-fit CR yield approaches the observed yield $\hat{v}_{\text{central } \gamma \text{ CR}}^i(\mu, \theta) \rightarrow N_{\text{central } \gamma \text{ CR}}^i$. Then Equation 7.3 yields

$$\hat{v}_{\text{central } \gamma}^i \rightarrow N_{\text{central } \gamma \text{ CR}}^i - \hat{\mu} v_{\text{central } \gamma \text{ CR, signal}}^i(\hat{\theta}_{\text{signal}}) - \sum_{s \in \{W^\pm \gamma, Z^0 \gamma, t\bar{t}\}} v_{\text{central } \gamma \text{ CR, } s}(\hat{\theta}_s), \quad (7.4)$$

recovering the intuition of subtracting resonant contamination from the observed CR yield to estimate the non-resonant background yield in the SR (cf. Section 6.1).

The event yield expectations for the forward-photon-selecting CRs are constructed in full analogy with Equations 7.2 and 7.3. The expected event yield in the forward-photon tagged-jet CR is:

$$v_{\text{fwd. } \gamma \text{ tagged CR}}^i(\mu, \theta) = \mu v_{\text{fwd. } \gamma \text{ tagged CR, signal}}^i(\theta_{\text{signal}}) + \frac{\epsilon_{\text{SR}}}{1 - \epsilon_{\text{SR}}} \tau^i v_{\text{forward } \gamma}^i + \sum_{s \in \{W^\pm \gamma, Z^0 \gamma, t\bar{t}\}} v_{\text{fwd. } \gamma \text{ tagged CR, } s}(\theta_s), \quad (7.5)$$

and in the forward-photon anti-tagged CR:

$$v_{\text{fwd. } \gamma \text{ CR}}^i(\mu, \theta) = \mu v_{\text{fwd. } \gamma \text{ CR, signal}}^i(\theta_{\text{signal}}) + v_{\text{forward } \gamma}^i + \sum_{s \in \{W^\pm \gamma, Z^0 \gamma, t\bar{t}\}} v_{\text{fwd. } \gamma \text{ CR, } s}(\theta_s). \quad (7.6)$$

The use of a separate set of non-resonant background yield parameters $v_{\text{forward } \gamma}^i$ in Equations 7.5 and 7.6 allows the likelihood to fit potential residual differences between the central-photon and forward-photon

mass spectra in-situ. Note that the same background efficiency ϵ_{SR} and (ratio) scale factor τ^i are used in Equations 7.2 and 7.5, thus allowing the forward-photon CRs to constrain the τ^i in a region with reduced resonant contamination.

7.1.1 Transfer factor parameterization

The efficiency ratio scale factors are expected to be $\mathcal{O}(1)$ and to vary over mass scales that are large compared to the jet mass resolution (which determines the anticipated width of signal excesses). To express this expectation in the form of the likelihood, the bin-wise independent τ^i in Equations 7.2 and 7.5 are replaced by a function. This function is evaluated in every mass bin but depends on a globally shared set of coefficients, i.e.,

$$\frac{\epsilon_{\text{SR}}}{1 - \epsilon_{\text{SR}}} \tau^i = f(m_{\text{J}}^i), \quad (7.7)$$

where m_{J}^i is the mid-point of the mass bin i .

The transfer function f is parameterized using the Bernstein polynomial basis of degree d

$$f(m_{\text{J}}^i) = \sum_{k=0}^d c_k \binom{d}{k} (m_{\text{J}}^i)^k (1 - m_{\text{J}}^i)^{(d-k)}, \quad (7.8)$$

where the $d + 1$ polynomial coefficients c_k are parameters without prior constraints whose values are determined in the maximum-likelihood fit. Figure 7.2 shows an example set of basis polynomials.

The Bernstein basis has attractive properties for maximum-likelihood fits when evaluated in the interval $(0, 1)$, namely:

- Linear combinations of Bernstein basis polynomials can approximate arbitrary functions $(0, 1) \rightarrow \mathbb{R}$.
- All basis polynomials take only positive values over the open interval. Thus, by a simple constraint $c_k \geq 0 \forall k$ the transfer function is guaranteed to take only non-negative values. This is beneficial since it ensures real and non-negative expected likelihoods, which are assumed by the `iminuit` likelihood minimization package[212, 213].
- The Bernstein basis forms a partition of unity, i.e., $f(x) = 1$ when $c_k = 1 \forall k$, allowing deviations of the c_k from a shared $\frac{\epsilon_{\text{SR}}}{1 - \epsilon_{\text{SR}}}$ to easily be interpreted as deviations from flatness of the in-situ efficiency ratio.

To benefit from these above properties, the centres of the mass bins³ are linearly mapped to the interval $[0, 1]$ as

$$m_{\text{J}}^i \rightarrow \frac{m_{\text{J}}^i - 11 \text{ GeV}}{264 \text{ GeV}} \quad (7.9)$$




Figure 7.2: Degree 5 Bernstein basis polynomials, cf. Equation 7.8.

3: cf. Table 7.1

Table 7.2: Design of constrained nuisance parameters modelling systematic uncertainties. The PDF sets refer to the respective LHAPDF ID [104].

Type of systematic	corr. across processes	affected processes	template modification
Experimental sources			
JES/JER	✓	$W^\pm\gamma, Z^0\gamma, t\bar{t}$, signal	shape and norm.
Photon energy scale	✓	$W^\pm\gamma, Z^0\gamma, t\bar{t}$, signal	normalization
Luminosity	✓	$W^\pm\gamma, Z^0\gamma, t\bar{t}$, signal	normalization
D_2^{DDT} SF uncertainty	✓	$W^\pm\gamma, Z^0\gamma, t\bar{t}$, signal	normalization
spurious signal		signal	normalization
Theoretical sources			
MC stat. uncertainty	✓	$W^\pm\gamma, Z^0\gamma, t\bar{t}$, signal	bin-wise uncorr.
PDF eigenvariations (PDF set 303200)	✓	$W^\pm\gamma, Z^0\gamma$	normalization
PDF eigenvariations (PDF set 260000)	✓	$t\bar{t}$, signal	normalization
PDF choice		$W^\pm\gamma, Z^0\gamma, t\bar{t}$	normalization
α_S (PDF set 303200)	✓	$W^\pm\gamma, Z^0\gamma$	normalization
α_S (PDF set 260000)	✓	$t\bar{t}$, signal	normalization
Parton Shower/ Hadr.		$W^\pm\gamma, Z^0\gamma, t\bar{t}$	normalization
ME scale variations		$W^\pm\gamma, Z^0\gamma, t\bar{t}$, signal	normalization

before evaluating the polynomial.

The likelihood for the final signal extraction fit uses a degree $d = 4$ polynomial transfer function parameterization. This is chosen as the lowest-degree polynomial that satisfies successive signal injection tests using pseudodata (see Section 7.3).

An alternative parameterization using the Chebyshev-basis polynomials of the first kind is studied. No significant difference in the quality of successful fits to pseudodata is observed. However, a slightly greater fraction of pseudodata fits is observed to fail to converge. This is due to the inability to prevent the polynomial from taking negative values by simple parameter constraints, because the Chebyshev-basis polynomials take negative value in part of the fiducial domain. Thus, no simple parameter constraint can prevent negative expected bin yields while maintaining the universal approximation properties required for sufficient fit. As a result, fits may enter parameter values generating negative yields during maximization, resulting in an ill-defined likelihood, and a failure of the fit.

7.1.2 Modelling systematic uncertainties

Systematic uncertainty about the normalization and shape of the simulated signal- and background templates is expressed as parameterized modifications of the nominal templates. Template modifications are parameterized as a function of nuisance parameters (NPs). Each NP parameterizing systematic uncertainty has an associated prior constraint term in the full likelihood, in the form of a gaussian probability density. The




Figure 7.3: Effect of the JES NP `JET_EffectiveNP_Modelling1` on the $Z^0\gamma$ background template across the four analysis regions (see Table 6.1 for their definition).

constraint terms are maximal at the value of the NP which corresponds to the nominal (not modified) simulated template.

Table 7.2 summarizes the different categories of systematic uncertainty sources included in the likelihood. Only resonant backgrounds and the signal are affected.

Some systematic uncertainties are correlated across processes, because the uncertainty they represent originates from a source that is common to all affected processes (e.g. a prior calibration measurement). Other systematic uncertainties are process-specific, meaning they are allocated separate NPs per process, allowing them to vary independently in the fit. Most theory uncertainty terms are modelled as independent between processes. The two exceptions to this are the uncertainty due to PDF eigenvariations, and the uncertainty on the QCD coupling strength parameter α_S . Both




Figure 7.4: Effect of the JER NP `JET_JER_EffectiveNP_2` on the $W^\pm\gamma$ background template across the four analysis regions (see Table 6.1 for their definition).

are the result of (global fits to) prior measurements, and therefore treated as correlated among processes. But different processes are modelled using different PDF sets, hence a different nuisance parameter is assigned per PDF set, and templates generated using different PDF sets are treated as uncorrelated.

Several uncertainties predominantly affect the normalization of the simulated templates. These are parameterized such that they may only scale the templates by an overall normalization to speed up the numerical evaluation of the likelihood. For each such uncertainty source, scaling of the normalizations is controlled by a single global nuisance parameter, expressing full correlation across the four regions. The proportion by which a particular template in a particular region is scaled, given the same change in the nuisance parameter, depends on both the region and the template. This is the parameterization of all uncertainties except




Figure 7.5: Effect of the PDF eigenvariations (ME_PDF260000_var) on the $m'_Z = 20$ GeV signal template across the four analysis regions (see Table 6.1 for their definition).

for those due to MC statistical uncertainty and those relating to the jet measurements.

The uncertainty due to the finite size of the MC samples is not correlated across mass bins. Hence, the impact of two separate processes pulling this uncertainty in the same bin is degenerate w.r.t. the impact on the likelihood. Thus, the evaluation of the likelihood is sped up by combining all MC statistical uncertainty effects into a single overall bin-wise modification and NP [211].

The jet energy scale (JES) and jet energy resolution (JER) uncertainties are modelled according to the CategoryReduction_FullJER set of reduced NPs⁴. These NPs are reduced using eigenvalue decomposition of the covariance matrix within physically meaningful categories of uncertainty sources, using only the eigenvectors with the greatest eigenvalues to reduce the number of nuisance parameters while preserving the most

4: This category comprises 62 NPs, 18 for JER and 44 for JES.

impactful correlations [177].

The main effect of JES and JER NPs is to distort the shape of the mass spectra. In principle, these NPs could also impact the template normalizations, (e.g. by moving jets close to a minimum p_T selection below / above the threshold), but in practice such effects are found to be minor. Figure 7.3 shows an example of a JES NP affecting the mass spectrum expected for the $Z^0\gamma$ background. The largest effect is near the mode of the distribution, which is shifted towards higher masses by the $+1\sigma$ variation, and vice versa for the opposite -1σ variation.

Similarly, JER NPs mainly impact the shape of the mass distribution templates. However, rather than shifting masses up- and down- across the spectrum like the JES, JER variations tend to modify the width of peaks. In the example shown in Figure 7.4, the down-variation narrows the peak (the rate of bins closer to the peak is increased, that of bins farther from the peak decreased), and the up-variation broadens it (fewer entries close to the peak, more entries in the flanks).

The photon energy scale uses the 1NP_v1 decorrelation model to reduce the number of nuisance parameters. This reduction drops the covariance structure of the input NPs. Hence, significant template modifications due to the decorrelated NPs being pulled are likely unphysical. However, the analysis has no strong dependence on the photon energy, so such pulls are unlikely. In the event that they were observed, a more accurate description would have been substituted.

The sources of theoretical uncertainty are described in Section 5.1.2. They are propagated to the likelihood as nuisance parameters. Because they are found to have a strong impact on the overall cross-sections, but no or only insignificant impact on the shape of the predicted jet mass template distributions⁵, the likelihood models their effect as modifying the template normalizations only.

Table 7.2 shows how different theory uncertainty sources are assumed to be correlated. In particular, the effects of choosing a different PDF set and scale uncertainties on the matrix element (ME) and parton shower (PS) are specific to particular processes, even if they are estimated using similar methods. E.g., the ME scale uncertainty aims to estimate the impact of higher orders on the cross-section calculation, which depends both on the process being calculated and the details of the calculation. Thus, even if the method⁶ of estimating these uncertainties is similar between processes, the effect being estimated can not be assumed to correlate, as e.g., different higher-order terms apply to different processes at the next order of perturbative approximation.

The D_2^{DDT} tag scale factor correcting the simulated tagging efficiency to match that observed in the $t\bar{t}$ CR is determined as 0.999 ± 0.025 in Section 6.4. This uncertainty is modelled as a Gaussian-constrained NP.

5: cf. Figure 7.5

6: variations of factorization and renormalization scales

The effect of the NP is to modify the overall yields of the templates modelling resonant processes. This modification is fully anti-correlated between the D_2^{DDT} -passing and -failing analysis regions (cf. Equation 6.7), and fully correlated between the forward-photon and central-photon analysis regions.


An uncertainty accounting for the effect of potential spurious signal extracted from the fit is modelled as a constrained multiplicative signal yield modification. The impact of the modification is chosen based on the (negligible) amount of spurious signal found in the studies described in Section 7.4.

The measured luminosity of the analysed data sample is used to normalize the resonant background and signal yields. The luminosity measurement uses primarily data from LUCID-2 [214]. The uncertainty on the luminosity measurement is⁷ 0.83 %, with the absolute luminosity calibration using van der Meer scans being the largest contribution [92]. This uncertainty is modelled as a globally correlated multiplicative modification of the normalization of all simulated templates across all regions, controlled by a single parameter with a Gaussian constraint. The non-resonant background, being estimated using observed data, is not scaled by the measured luminosity, and thus not modified by the luminosity uncertainty term.

7.2 Smoothed template of QCD multijet background jet mass spectrum

A *smooth template* of the non-resonant background mass spectra is required for several cross-checks performed on the signal extraction fit. These cross-checks include

- testing for the presence of *spurious signal* (or *bias*, see Section 7.4) ,
- testing the *linearity* of the fit response to artificially injected signal (see Section 7.5),



7: The value quoted in the text is the uncertainty on the 140 fb^{-1} dataset. The uncertainties on the 3 fb^{-1} 2015 and 2018 (L4 period) datasets used to validate the fit (see Section 8.1) are larger at 1.13 % and 1.10 % respectively. Similarly, the 15.8 fb^{-1} dataset used in partial unblinding (Section 8.2) has a slightly larger luminosity uncertainty of 0.87 %.

[92]: ATLAS Collaboration (2022), *Luminosity determination in pp collisions at $\sqrt{s} = 13 \text{ TeV}$ using the ATLAS detector at the LHC*

Figure 7.6: Post-fit mass spectra including non-resonant QCD background simulations in the D_2^{DDT} -passing regions. Both the central-photon selection in (a) and the forward-photon selection in (b) are shown. The pseudodata distributions are compared to the background-only (μ fixed at 0) maximum-likelihood expected yield. The expected sample yields of the likelihood and the pseudodata distribution are scaled to an integrated luminosity of 140 fb^{-1} , and consist of the combination of mc16a, mc16d, and mc16e samples. Error bars indicate the standard deviation of a Poisson distribution with expectation value equal to the bins pseudodata content.

8: This is why the in-situ validation samples used in Section 8.1 are not suitable here.

9: The bin-by-bin variance exhibited by the background template and pseudodata is far greater than the variance expected for observed data, the error bars indicate. The observed variance of the pseudodata is due to the size of the simulated sample, which is thus too small.

10: cf. Figure B.2

- and determining the *expected sensitivity* of the search before unblinding (see Section 7.7).

All of these tests should be performed assuming the full integrated luminosity of the Run 2 dataset, which will be used for the unblinded result.⁸ Furthermore, the template should be constructed in a way that ensures freedom from signal contributions, otherwise the spurious signal test becomes invalid. Thus, a dedicated smooth QCD multijet template needs to be constructed.

7.2.1 Smoothing method overview

Standard methods for constructing smooth background templates include MC simulation, data observed in a control region, and the fit of a smooth function to data [188]. Applied in the usual way, none of these standard approaches are sufficient in this case:

- The available simulated nonresonant background samples have an effective statistical uncertainty much larger than that of data, and the statistical uncertainty is too large to be usable for spurious signal tests. In order to avoid bias being obscured by Poisson noise in the background template, the background template should contain 10-100 times more simulated MC events than the dataset the search will be applied to. This requirement increases further if the search covers a wide range of signal hypotheses, as is the case here [188]. The available QCD-associated single-photon sample is far too small, as Figure 7.6 illustrates⁹.
- The existing search strategy already makes use of most discriminating features (jet mass, jet substructure decorrelated from mass, photon kinematics) that might be used to define control regions. Further, the existing control regions have too much signal contamination to be usable in signal injection tests, and new control regions would be confined to using even less discriminating observables.
- The functional form used to transfer the tagger efficiency measured in the control region to the signal region only smooths a ratio of efficiencies. Thus, it cannot be used on its own to generate smooth background templates. In other words, the respective shape-factors specific to the central-photon and forward-photon selections are sensitive to non-smoothness, and hence require additional smoothing of the inputs¹⁰.

Since the above methods are not suited to generating a smooth QCD multijet template, a dedicated smoothing procedure is used, which extends the method of fitting a smooth function to this analysis.

The smooth QCD multijet template is a set of four univariate template histograms of the jet mass, one for each of the three control regions and the




Figure 7.7: Toy example of using a KDE to smooth a noisy histogram sampled from a χ^2 -distribution for 3 degrees of freedom (green curve). The number of samples used to fill the noisy example histogram is indicated as `n_samples`. The size of the toy dataset the KDE is fit to is $10 * n_samples$. The number of bins of the histograms is chosen as the maximum of the number of optimal bins according to Sturges' rule and the Freedman-Diaconis rule [215], so it increases as the number of toys increases. The bandwidth of the KDE is chosen as twice the width of a histogram bin.

signal region defined in Table 6.1. The four histograms are filled with non-resonant MC-simulated samples, using the analysis selections outlined in 5.2.4. The smoothing procedure is guided by several assumptions about the expected behaviour of the non-resonant background:


- The background should be smooth bin-by-bin, i.e., there should be no discontinuous jumps between neighbouring mass bins due to Monte-Carlo sampling noise;
- The ratios of the D_2^{DDT} -pass-to-fail ratios should also be smooth;
- The double ratio of the D_2^{DDT} -pass-to-fail in the central and forward regions should be smooth.

Based on these assumptions, the four histograms for the different regions are smoothed in three consecutive steps, every step taking the previous steps output histograms as its input:

1. Every histogram is smoothed using a kernel-density estimate (KDE) [215–217].
2. The two D_2^{DDT} -pass-fail ratios are smoothed using a Savitzky-Golay filter [215, 218].
3. A global pass-fail ratio is fit to all four histograms at once, using a polynomial parameterization of an order once higher than in the signal extraction fit. The expected non-resonant yields of the best-fit workspace in this final step are then taken as the smooth QCD multijet template.

Each of these steps is described in greater detail below.

Figure 7.8: Examples of using a KDE to smooth a noisy multijet background mass histogram of the signal region. The original noisy simulated template is shown as the green line. A toy histogram is sampled from it, shown as the blue filled bar chart, to induce additional artificial noise for illustration. (This is not done when smoothing histograms to create the smooth QCD multijet template.) The number of samples used to fill this noisy example histogram is indicated as `n_samples`. The size of the toy dataset the KDE is fit to is $10 * n_samples$. The width of the mass bins is fixed at 2 GeV. The bandwidth of the KDE is 0.16. Refer to Appendix B.2 for illustrations of the effect of bandwidth changes.



7.2.2 Smoothing the QCD multijet template

In order to ensure that the smooth QCD multijet template is smooth bin-by-bin, the first smoothing step uses KDEs to smooth the four histograms corresponding to the three control regions and the signal region independently. Toy mass values are randomly sampled from the histogram. Within any histogram bin a uniform probability distribution is assumed by the sampling. The resulting sample of toy mass values is fit with a KDE. The KDE is then integrated over intervals determined by the bin edges. These bin-wise integrals of the KDE are taken as the smoothed histogram.

The KDE smoothing method is studied in the case of artificial and known underlying probability densities, from which a noisy toy histogram is drawn, which the KDE smoothing should approximately recover. One example check is illustrated in Figure 7.7. The qualitative shape of the generating probability density is approximated well by the KDE, and the bin-wise integral is an appropriate smoothing of the noisy toy histograms.

The KDE smoothing method has two free parameters: the number of toys drawn from the input histogram, and the bandwidth parameter of the KDE. Figure 7.8 illustrates the KDE smoothing applied to the simulated signal region jet mass spectrum, for a range of toy sample sizes¹¹ used to fit the KDE. Increasing the number of toys is computationally intensive, but a too low number of toys may lead to sampling noise biasing the resulting smoothed histogram. 10^7 toys are used, which is more than ten times larger than the unitary-weight sample with a relative Poisson standard deviation equivalent to the relative uncertainty in the weighted simulated sample underlying the input histograms [219]. The bandwidth of the KDE is chosen as 0.16 (cf. Figure 7.8) to achieve a strong smoothing while

11: In Appendix B.2, a range of alternative KDE bandwidths is illustrated.




Figure 7.9: Pseudodata and constituent samples in the four analysis regions, using the smooth QCD multijet template for the non-resonant background. The templates are the combination of mc16a, mc16d, and mc16e scaled to 140 fb^{-1} . The pseudodata distributions are compared to constituent samples from a simple pre-fit where shape factors are set such that the expected yield of the D_2^{DDT} -failing regions exactly matches the observations, the resonant background yields are set to the nominal values, the signal yield set to 0, and using transfer polynomial coefficients taken from a simple least-squares fit to only the forward-photon D_2^{DDT} -pass-to-fail ratio. The transfer function polynomial used here has degree 4, unlike the one used in the smoothing procedure.

balancing this against the increasingly biased shape of the smoothed histogram at the lowest masses. Said bias arises due to the inherent size of the underlying kernel¹², which limits the steepness of the slope the final KDE can approximate.

7.2.3 Smoothing the pass-to-fail ratio

Since the KDE smoothing is performed separately for each histogram, the individual histograms are locally smooth, but the two ratios of the D_2^{DDT} -passing to the D_2^{DDT} -failing histograms are not. This non-smoothness is unphysical, and so the second smoothing step corrects this by applying a Savitzky-Golay filter to the two pass-to-fail ratios of the two respective histogram pairs in the forward-photon and central-photon selections.

The Savitzky-Golay filter consists of a sliding window polynomial fit. For each mass bin, the ratio value is replaced by the value at the bin centre of a quadratic function fit to the 15 nearest mass bins. Then, to produce the smoothed histograms for each region, the smoothed ratios are multiplied by the D_2^{DDT} -failing smoothed histograms¹³. The two D_2^{DDT} -passing histograms are left unchanged in this step.


7.2.4 Smoothing the double ratio

Having separately smoothed the central-photon and forward-photon D_2^{DDT} -pass-to-fail ratios with the Savitzky-Golay filter, non-smoothness

12: This is confirmed by comparison to the toy example in Figure 7.7. It shows that the steep rising slope of the generating χ^2 distribution is not quite reproduced by the KDE. This is due to the KDE bandwidth which is systematically too-large and implies a limit on the slope that can be approximated. By contrast the less steep falling slope of the generating distribution is reproduced well by the KDE.

13: The choice to smooth the ratio by replacing the D_2^{DDT} -passing histograms is motivated by the larger overall yield in the D_2^{DDT} -failing histograms.

Figure 7.10: Transfer function polynomials extracted from fits to the smooth QCD multijet template (cf. Figure 7.9) compared to the D_2^{DDT} -pass-to-fail ratios observed in *unsmoothed* multijet QCD Monte-Carlo simulation (scaled to 140 fb^{-1}). The top two pads show the unsmoothed multijet QCD templates, with different line colors for the spectrum of the D_2^{DDT} passing and failing selection, respectively. The topmost plot shows the two regions where the photon is forward, the middle plot shows the selection where the photon is central. In the bottom plot, the markers show the D_2^{DDT} pass-to-fail ratios of the two photon selections, with error bars indicating the propagated uncertainty due to Monte-Carlo sampling noise. The light blue line indicates the transfer polynomial fit to the D_2^{DDT} -pass-fail ratio of the smooth QCD multijet template in the forward-photon selection only (cf. Figure 7.9), whereas the dark blue line shows the transfer polynomial that is implicitly fit by a maximization of the likelihood defined using the full smooth QCD multijet template.



can remain in the double ratio of the central-to-forward ratio of D_2^{DDT} -pass-to-fail ratios. Furthermore, the two preceding independent smoothing steps may have biased the two D_2^{DDT} -pass-to-fail ratios to disagree systematically. Such systematically induced disagreement would break the underlying assumption of the fit strategy that the D_2^{DDT} background efficiency can be measured in the forward-photon region and applied to the central-photon region. In order to re-align the smoothed histograms with this assumption, the final smoothing step performs a special background-only maximum-likelihood fit.

The final smoothing fit uses a workspace composed of the smoothed histograms of the previous smoothing steps as the non-resonant contribution to the “observed” pseudodata, in addition to the usual resonant background samples¹⁴. The resonant background samples are also used to define the expected yields of the likelihood, but their yields are fixed to the nominal. The non-resonant background is estimated in the data-driven way from the pseudodata “observations”, but the transfer function is parameterized by a degree 5 polynomial, rather than the degree 4 polynomial used in the signal extraction fit proper. Only the coefficients of the transfer function polynomial are free to be varied in the fit, all other nuisance parameters are held fixed at their nominal values. The set of polynomial coefficients that maximize this restricted likelihood are then used to extract expected non-resonant background yields simultaneously in all four analysis regions, and these are taken as the smooth QCD multijet template.

Figure 7.9 shows the final smooth QCD multijet template. The Monte-

14: cf. Section 6.1

Carlo sampling noise of the input templates has been fully removed¹⁵. The shape of the smooth QCD multijet template is qualitatively consistent with the shape of the input template at medium and high masses. The low mass end of the smoothed spectra is biased towards a less steep slope, which is an artefact of the KDE smoothing (cf. Figure 7.8) as previously discussed. However, this bias can be fully absorbed by shapefactors, since it is consistent between the D_2^{DDT} -passing and D_2^{DDT} -failing regions. Thus fitting a likelihood using a degree 4 transfer function polynomial to the smoothed template results in a fitted polynomial that is consistent with the D_2^{DDT} -pass-to-fail ratios observed in the simulated input histograms before any smoothing, as Figure 7.10 shows. Hence, the smooth QCD multijet template described in this section is used for blind checks of the fit strategy that require a smooth expected pseudodata yield and cannot be done using observed data.

15: Cf. Figure 7.6

7.2.5 Applicability

The smooth QCD multijet template described in this section is only used for a restricted set of fit studies prior to unblinding. In particular the smooth QCD multijet template it is *not* used to determine the expected sensitivity of the search, neither for propagating detector and generator uncertainties, nor for studying the modelling of the non-resonant background.

- The smooth QCD multijet template is not used for the expected sensitivity calculation after unblinding. Doing so would risk introducing systematic differences in the likelihood used for the observed limit with respect to the expected limit, and thus break the interpretation as a sensitivity for the search. Instead, an expected sensitivity using the data-driven background estimate is derived using the Asimov-dataset [220] of the unblinded dataset.
- The smooth QCD multijet template is not used for propagating detector-based or theoretical uncertainties because the non-resonant background estimate is fully data-driven, and thus not affected by these uncertainties¹⁶
- The smooth QCD multijet template is not used for studying the simulated shape of the non-resonant background mass spectra. To do so would open up the possibility of shape modifications due to the smoothing procedure, which would be a confounding effect with respect to true mismodelling of the spectrum. The use of a data-driven background estimate elides the question of mismodelling of the background, and is validated in-situ (see Section 8.1).

16: Only the templates describing the resonant backgrounds and the signal processes are affected by these uncertainties, cf. Sections 7.1 and 7.1.2.

In short, the smooth QCD multijet template is only used in those studies that can not be done using data validation samples and the data-driven background estimate. These are the fit linearity test with full statistics

in Section 7.5, the spurious signal test Section 7.4, and the expected sensitivity *prior* to unblinding in Section 7.7.

7.3 Choice of transfer factor polynomial degree

The degree of the transfer factor polynomial used to model the ratio of scale factors of the non-resonant D_2^{DDT} efficiency in the likelihood (see Section 7.1) is determined using iterative fits to pseudodata, adopting the recommendation of [188]. Input pseudodata is constructed by combining the simulated resonant background templates and the smooth QCD multijet template described in Section 7.2. Pseudodata are fitted repeatedly, incrementing the degree of the transfer factor polynomial by one at each iteration. The smallest polynomial degree yielding a satisfactory fit is chosen for the final likelihood definition.

The smooth QCD multijet template construction in Section 7.2 includes a step of fitting a likelihood to partially smoothed pseudodata. When iterating the polynomial degree in the studies described here, the degree of the polynomial used to smooth the QCD pseudodata is always one greater than the degree of the polynomial parameterization being tested.




Figure 7.11: Post-fit residuals of fits to the Asimov dataset in the SR, using basis polynomials of varying degree to parameterize the mass dependence of the non-resonant transfer factor. The residual of the fit with respect to the fitted Asimov dataset is shown by the black points in the bottom subplots. To ensure a negligible spurious signal uncertainty, this residual should be less than a third of the uncertainty on the fit, which is indicated by the hatched area.

Satisfactory fit quality for a given polynomial degree choice is evaluated using the χ^2 test and the supremum of spurious signal across the mass range. The χ^2 test probability is the probability of observing pseudodata at least as discrepant as the actual pseudodata assuming Gaussian per-bin likelihoods. The resulting χ^2 test probabilities are shown in Figure 7.12. Except for $d = 0$ (i.e. a transfer factor that is constant vs. jet mass) $p(\chi^2)$ is always larger than the acceptance criterion of $p(\chi^2) > 0.01$ suggested in [188]. Thus, the final choice is made based on the more stringent spurious signal criterion.

During initial spurious signal scans the amount of spurious signal found using a sufficiently large sample of toy distributions drawn from pseudodata agrees well with the residual disagreement between the fitted distribution and the Asimov dataset¹⁷. Hence, this residual of the fit to the Asimov dataset is used to estimate the spurious signal while iterating through alternative polynomial degrees, rather than repeated fits to toy distributions. This greatly reduces the computational cost of




Figure 7.12: Summary of iterative scans over alternative transfer factor polynomial degree choices.

[188]: ATLAS (2020), *Recommendations for the Modeling of Smooth Backgrounds*

17: The Asimov dataset [220] is the set of expected distributions corresponding to the parameter values that maximize the likelihood of the (pseudo-)data, given the parameter of interest is fixed at 0. It is more commonly used in hypothesis testing, see Section 7.7.1.

the study. The likelihood used to construct the Asimov dataset uses the next-higher-degree polynomial basis to parameterize the transfer factor.

Figure 7.11 shows selected examples of fits to the Asimov datasets. The fits using degree one (see Figure 7.11(a)) and degree three (see Figure 7.11(b)) polynomials exhibit clear bias, even though the χ^2 criterion is satisfied. The first parameterization showing negligible bias uses the basis polynomials of the fourth degree, shown in Figure 7.11(c).

The absence of significant spurious signal given the final choice of the fourth-degree polynomial basis is validated using toy distributions drawn from pseudodata, as described in Section 7.4.

7.4 Spurious signal


The bias – or spurious signal – inherent in the signal extraction fit is checked by drawing toy pseudodata distributions and fitting them. The likelihoods used in the study of fit bias include all other systematic uncertainties. Independent fits are done for every signal mass hypothesis. For each signal mass point, 1080 independent toys are drawn and fitted. Figure 7.13 shows the distributions of the best-fit signal strength $\hat{\mu}$. The average $\hat{\mu}$ is found to be consistent with zero within $\frac{1}{3}$ of the standard deviation of the observed $\hat{\mu}$ sample¹⁸ for every signal mass tested. This indicates the absence of spurious signal, i.e., the signal extraction fit is not overly biased.

18: this is the criterion suggested in ref. [188].

The small amount of spurious signal that is observed is taken as an additional constrained nuisance parameter when constructing the final likelihood. The amount of spurious signal extracted from the ensemble of fits to toys is parameterized as a function of the signal resonance mass. To avoid assigning smaller-than-intuitively credible uncertainty to some masses, an envelope of the mass-parametrized spurious signal is used as recommended in ref. [188]. To construct the envelope, support

[188]: ATLAS (2020), *Recommendations for the Modeling of Smooth Backgrounds*

Figure 7.13: Distributions of best-fit signal strength $\hat{\mu}$ of fits to pseudodata toys. The toys are drawn from a bin-wise Poisson distribution with expectation value equal to the sum of the resonant background templates, and the smooth QCD multijet template described in Section 7.2, fixing the resonant signal yield to zero. 1080 independent toys are drawn for every mass point.



points are chosen at arbitrarily selected local extrema of the absolute value of the parameterized spurious signal. The envelope is then taken as the piecewise linear interpolation between the support points. The value of this interpolated envelope at the pole mass of a given signal mass hypothesis is taken as the spurious signal uncertainty to apply to that given signal template. In the final likelihood parameterization, it is expressed as a constrained multiplicative¹⁹ modification (see also Section 7.1.2).

7.5 Linearity of fit response

The response of the signal extraction fit to varying levels of true signal presence is tested using signal injection into pseudodata. The pseudodata is sampled from the sum of the resonant background templates, one signal template scaled to $\mu \in \{0, 0.2, 0.4, 0.6, 0.8, 1\}$ times its nominal yield, and the smooth QCD multijet template described in Section 7.2. The constituent templates are scaled to 140 fb^{-1} and 1080 independent toy distributions are sampled for every combination of signal mass and signal strength, according to a binwise Poisson distribution. A full signal extraction fit is performed on each toy.

The rate of failing fits is monitored, and in total 1.4 % of fits failed in the scan presented here. Fit failures are not independent of the signal hypothesis. The hypothesis for which the largest number of independent fits failed was $m_{Z'} = 225 \text{ GeV}$ and $\mu = 0.2$, with 45 failed fits out of 1080 independent attempts.

The sample average and standard deviation of the extracted signal strengths are shown in Figure 7.14. The fit is found to have a mostly linear response across the whole range of tested signal strengths. The sample standard deviation appears to remain constant when varying the injected (and extracted) signal, as does the fit bias, which is negligible by comparison. A slight deviation from linearity is observable for the highest Z' mass points due to lack of sample statistics at the high end of the mass distribution to perfectly constrain the transfer factor polynomial. However this effect is small compared to the uncertainty on the extracted signal strength, and in addition well covered by the spurious signal uncertainty after envelope smoothing (cf. Section 7.4). The results of further signal injection and linearity studies are described in Appendix C.

19: This deviates from the suggestion in ref. [188], where an additive modifier is recommended instead. This choice here is due to the lack of an appropriately correlated additive modifier implementation in the version of the `pyhf` library used to describe the likelihood [210, 211]. The deviation is deemed acceptable because the impact of the uncertainty is negligible.




Figure 7.14: Linearity of the fit response to varying injected signal strength μ . For every mass point, six equidistant signal strength values from 0 through 1 are injected. For every combination of signal mass and signal strength, 1080 independent toys are drawn, and a signal extraction fit is performed on each to find the best-fit signal strength $\hat{\mu}$. The grey line corresponds to $\hat{\mu} = \mu$. The sample average best-fit signal strength $\langle \hat{\mu} \rangle$ is indicated by the coloured central line. The sample standard deviation of the $\hat{\mu}$ distribution is indicated by the shaded area.

7.6 Impact of systematic uncertainties

The expected impact of the different sources of systematic uncertainty²⁰ on the best-fit signal strength $\hat{\mu}$ is evaluated using fits to toy mass distributions sampled from simulated signal and resonant background templates and the smooth QCD multijet template.

Signal is injected in the pseudodata generation for this study because in the absence of injected signal, the best-fit signal strength is²¹ at or close to 0, and thus the signal-associated uncertainties do not contribute, but, far from their respective mass peaks, the resonant background template contributions are also negligible. Thus, without injecting some signal, the determined impacts would be almost 0, without allowing conclusions to be drawn about the relative impact of different NPs in the presence of signal. The signal yield being injected is chosen to be near the edge of sensitivity of the fit²².

To determine the impact of a NP, the NP is usually fixed to the best-fit value incremented (decremented) by one standard deviation on the best-fit value, then the fit is repeated. The difference of the best-fit POI value of this fit with respect to the best-fit POI is the impact. Pre-fit impact is evaluated by instead using the width of the prior constraint to increment (decrement) the best-fit POI.

With $O(100)$ NPs in the likelihood, this usual method to estimate the impact of systematic uncertainty would require $O(1000)$ fits. To avoid this significant computational cost, the asymptotic approximation²³ [220, 222, 223] is used instead, requiring only one fit. In the limit of an infinitely large observed sample, the likelihood approximates a normal distribution centred on the parameter values that maximize it. In terms of the negative logarithm of the likelihood²⁴:

$$\mathcal{L}(\theta) = -\log(p(\theta)) \approx \frac{1}{2}(\theta - \hat{\theta})^T V^{-1}(\theta - \hat{\theta}) + \mathcal{L}_{\min} + O((\theta - \hat{\theta})^3) \quad (7.10)$$

where θ is the vector of parameters of the likelihood, $\hat{\theta}$ the vector of parameters minimizing $\mathcal{L}$, V the covariance matrix of the likelihood²⁵. While $\mathcal{L}$ also depends on the observed data through $\mathcal{L}_{\min}$ and V , for this analysis the dependence is not crucial, and so left out of the notation. Similarly, μ is taken to be the first element θ_0 of θ for brevity.

Using the asymptotic approximation, the result of the usual method can be calculated so long as V^{-1} has been estimated²⁶. First, V^{-1} and θ are restricted to a subspace of parameter space orthogonal to the direction of the NP θ_n for which the impact is to be determined. As in the usual method, the two²⁷ cases examined are $\theta_n = \hat{\theta}_n \pm \delta\theta_n$. Then the parameter vector θ' which minimizes $\mathcal{L}$ in the given subspace is found. Because Equation 7.10 is quadratic in θ , the unique global minimum can be found

20: see Section 7.1.2

21: cf. Section 7.4

22: see Section 7.7

23: cf. similar approximations used in Sections 7.3 and 7.7.1

[220]: Cowan et al. (2011), *Asymptotic formulae for likelihood-based tests of new physics*

[222]: Wald (1943), *Tests of Statistical Hypotheses Concerning Several Parameters When the Number of Observations is Large*

[223]: Wilks (1938), *The Large-Sample Distribution of the Likelihood Ratio for Testing Composite Hypotheses*

24: which is more convenient to write and evaluate, and otherwise equivalent

25: as well as the negative inverse of the Fisher information matrix and the inverse of the Hessian of $\mathcal{L}$.

26: e.g. using a finite-element method such as in `iminuit`'s `HESSE` implementation [212, 213], or an automatic differentiation program such as JAX [224].

27: Since Equation 7.10 is invariant under sign-changes of $(\theta - \hat{\theta})$, the impacts thus determined are necessarily symmetric, so it is sufficient to only evaluate one case.




Figure 7.15: NPs ranked by impact on the POI μ for the 40 GeV Z' signal hypothesis. Only the 15 most impactful individual NPs are shown. See Table 7.3 for NP combinations and other mass hypotheses. Refer to Section 7.1.2 for an overview of the modelling of systematic uncertainty. The ranked impact visualization is implemented using `cabinetry` [221].

by solving

$$0 = \frac{\partial \mathcal{L}}{\partial \theta_k} = \left[\sum_{i \neq n} \left(V_{ki}^{-1} (\theta'_i - \hat{\theta}_i) \right) \right] \pm V_{kn}^{-1} \delta \theta_n, \quad (7.11)$$

where k goes over all parameter indices except n . Equation 7.11 is a system of equations linear in θ' and otherwise contains quantities that can be determined using a single fit. Thus, Equation 7.11 can be solved for θ' straightforwardly using linear algebra routines of common scientific programming libraries²⁸. Then, the impact of the given NP on the POI θ_0 is $\Delta\mu = \hat{\theta}_0 - \theta'_0$.

Figure 7.15 shows the most impactful individual NPs for one signal mass hypothesis. In general, systematic uncertainties with prior constraints are not frequent among the most relevant NPs to the overall uncertainty. However, due to the large number of individual NPs, it is more informative to consider combinations of impacts instead.

Table 7.3 lists the combined impacts of related categories of NPs for selected representative signal mass hypotheses. When combining impacts, linear correlations between the different NPs are accounted for. As a result, combined categories may have lower total impact than a naive sum-of-squares combination would suggest. The total post-fit uncertainty on $\hat{\mu}$ for a fit without any systematic uncertainty terms in the likelihood is also given for comparison. Beyond the pure statistical uncertainty due to the limited size of the observed sample in the SR, by far the most relevant uncertainties are those relating to the data-driven estimate of the non-resonant background contribution. In that sense, it is the most

28: For example, the impacts shown in this section are determined using `numpy.linalg` [225].

Table 7.3: Uncertainty on the best-fit signal strength $\hat{\mu}$ by category. The "Statistical uncertainty" row lists the uncertainty on $\hat{\mu}$ from a fit using a likelihood without constrained nuisance parameters.

Uncertainty source	$\Delta\mu/\hat{\mu}$			
	$m'_Z = 20 \text{ GeV}$	$m'_Z = 40 \text{ GeV}$	$m'_Z = 80 \text{ GeV}$	$m'_Z = 100 \text{ GeV}$
Bottom-up JES	0.061 %	0.019 %	0.7 %	0.11 %
Bottom-up JER	0.0077 %	0.016 %	0.061 %	0.63 %
resonant D_2^{DDT} scale-factor	0.056 %	0.0058 %	0.78 %	0.18 %
Spurious signal	0.0043 %	0.055 %	0.13 %	0.089 %
Total exp. uncertainty (excl. non-res.)	0.12 %	0.095 %	1.8 %	1.1 %
Total theory uncertainty	0.021 %	0.011 %	0.029 %	0.11 %
Non-resonant bkg. est. (data-driven)	22 %	14 %	25 %	29 %
Total syst. uncertainty (excl. non-res.)	0.14 %	0.11 %	1.8 %	1.2 %
Statistical uncertainty (incl. non-res.)	32 %	28 %	3.6×10^2 %	34 %
Total uncertainty	32 %	28 %	26 %	26 %

(only) relevant systematic uncertainty.

However, the question of what the most relevant systematic uncertainty is often aims to e.g., determine what aspect of the search to direct effort towards in order to efficiently improve its sensitivity. In such a context, the data driven estimate having the most impact may not be of great interest, since like the statistical uncertainty in the above sense, it will be reduced as the size of the observed data sample is increased. This is a consequence of this background estimate being determined by nearby regions of observed data. Thus, Table 7.3 also lists categories excluding this contribution. The remaining sources of uncertainty have an impact that is roughly one tenth of that of the non-resonant background estimate. The largest contributing category depends on the model hypothesis, but it is typically related to the calibration of the resonance candidate jet: jet energy resolution (JER), jet energy scale (JES), or the resonant jet tagging efficiency scale factor²⁹.

29: cf. Section 6.4

Near the $W^\pm$ and Z^0 mass peaks, the uncertainty of the fit without systematic uncertainty terms is many times larger than that of the fit using the full likelihood. This is because in the former likelihood, the effect of the $Z'\gamma$ signal with masses near 80 GeV or 90 GeV is not distinguishable from that of the $V\gamma$ backgrounds. This is reflected in the large fit uncertainty, and demonstrates why it is important to include the systematic uncertainties due to the theoretical modelling of resonant background processes, despite their low impact on the overall uncertainty in the full likelihood.

7.7 Expected exclusion estimate using pseudodata

The expected exclusion capability of the search is studied using pseudodata before unblinding. This section describes the statistical method for con-

structing the exclusion intervals based on pseudodata, but the exact same method is applied also when analysing observed data after unblinding.

An expected pseudodata distribution is constructed as the sum of nominal simulated resonant background templates and the smooth QCD multijet template described in Section 7.2. In each bin, the Poisson distribution given that bins expected pseudodata yield is sampled from to generate the toy distribution. This sampled toy distribution is kept fixed across all fits of the likelihood, like an actually observed distribution would be.

7.7.1 The asymptotic approximation for CL_s exclusion intervals

30: μ here refers to the factor by which the nominal signal template is scaled (cf. Equations 7.2, 7.3, 7.5 and 7.6), i.e. $\mu = 0$ corresponds to the background-only hypothesis, and $\mu = 1$ to the nominal coupling $g'_q = 1/5$ (cf. Section 5.1.2.3).

31: Table 5.5 lists the signal masses tested

95% CL_s limits [226] are calculated based on the pseudodata distribution across Z' mass and signal strength³⁰ μ . For each of the Z' mass hypotheses³¹, a range of signal strength hypotheses is tested. At each $m_{Z'}$, $\mu_{\text{hypo.}}$ combination, the best-fit signal strength $\hat{\mu}^D$ and NPs $\hat{\theta}^D$ are determined in a fit to the observed binned event yields N^D alongside the best-fit set of NPs $\hat{\theta}^D$ of a fit with constrained $\mu = \mu_{\text{hypo.}}$. In a similarly constrained fit for $\mu = 0$, NP values $\hat{\theta}_0$ called "Asimov parameters" are determined:

$$\hat{\theta}_0 = \underset{\theta}{\operatorname{argmax}} \left(p(N^D, 0, \theta) \right), \quad (7.12)$$

where p is the full likelihood as defined in Equation 7.1. The expected distribution according to p given these $\hat{\theta}_0$ is taken as the asymptotically expected background distribution N^A , also called "Asimov data" [220]

[220]: Cowan et al. (2011), *Asymptotic formulae for likelihood-based tests of new physics*

$$N^A = \underset{N}{\operatorname{argmax}} \left(p(N, 0, \hat{\theta}_0) \right). \quad (7.13)$$

A second set of best-fit parameters are estimated by fitting the likelihood to the Asimov data, i.e. $\hat{\mu}^A$ and $\hat{\theta}^A$ the unconstrained best-fit signal-strength and NPs (which are typically close to 0 and $\hat{\theta}^0$, respectively) and $\hat{\theta}^A$ the best-fit NPs of a likelihood constrained to $\mu = \mu_{\text{hypo.}}$. These fitted parameter estimates and Asimov data are used to estimate the frequentist p-value CL_{s+b} and CL_b , as well as the "modified frequentist" CL_s observable, for each signal strength hypothesis, simplified³² as:

$$\text{CL}_b(\mu_{\text{hypo.}}) = \int_{q_\mu^D(\mu_{\text{hypo.}})}^{\infty} f(q_\mu|0) dq_\mu \quad (7.14)$$

$$\approx 1 - \Phi \left(\sqrt{q_\mu^D(\mu_{\text{hypo.}})} - \sqrt{q_\mu^A(\mu_{\text{hypo.}})} \right) \quad (7.15)$$

$$= 1 - \Phi \left(-2 \ln \left(\frac{p(N^D, \mu_{\text{hypo.}}, \hat{\theta}^D)}{p(N^D, \hat{\mu}^D, \hat{\theta}^D)} \frac{p(N^A, \hat{\mu}^A, \hat{\theta}^A)}{p(N^A, \mu_{\text{hypo.}}, \hat{\theta}^A)} \right) \right), \quad (7.16)$$

$$\text{CL}_{s+b}(\mu_{\text{hypo.}}) = \int_{q_\mu^D(\mu_{\text{hypo.}})}^{\infty} f(q_\mu|\mu_{\text{hypo.}}) dq_\mu \quad (7.17)$$

32: Section 7.7.1.1 discusses the simplifications.

$$\approx 1 - \Phi\left(\sqrt{q_\mu^D(\mu_{\text{hypo.}})}\right) \quad (7.18)$$

$$= 1 - \Phi\left(-2 \ln\left(\frac{p(N^D, \mu_{\text{hypo.}}, \hat{\theta}^D)}{p(N^D, \hat{\mu}^D, \hat{\theta}^D)}\right)\right), \quad (7.19)$$

$$\text{CL}_s(\mu_{\text{hypo.}}) = \frac{\text{CL}_{s+b}(\mu_{\text{hypo.}})}{\text{CL}_b(\mu_{\text{hypo.}})}, \quad (7.20)$$

with Φ the cumulative distribution function of the standard normal distribution (for zero mean and unit variance), the test statistic $q_\mu = -2 \ln(p(N, \mu, \hat{\theta})/p(N, \hat{\mu}, \hat{\theta}))$, which is distributed according to a signal strength-dependent density labeled $f(q_\mu|\mu)$.

Equations 7.14 and 7.17 express the usual definitions of CL_s and CL_{s+b} as frequentist p-values. The test statistic q_μ is constructed as the ratio of the likelihood given the best-fit and tested model respectively, because such test-statistics achieve optimal power to distinguish two hypotheses, as ref. [227] showed.

To evaluate Equations 7.14 and 7.17 as written, the density f of the test statistic q_μ would need to be evaluated, for example by MC-sampling based on data distributions. This is extremely computationally costly because the definition of q_μ requires several numerical likelihood optimisations for each sample. Instead, the step from Equation 7.14 to Equation 7.15 and Equation 7.17 to Equation 7.18 makes two assumptions. Firstly, it assumes sufficient numbers of events observed in all bins, such that the error rate incurred by applying the large-sample limit is negligible. Secondly, it assumes that the true distribution of observed data is contained in the set of functions the parameterized by Equation 7.1, or can be approached by elements of the set with negligible error. Then, results by refs. [222, 223] allow the ratio of likelihoods in q_μ to be approximated as normal-distributed. Ref. [220] shows how this asymptotic distribution of the test statistic can be expressed using the Asimov data. From these expressions, the significance of the (dis-)agreement of the data and expectation can be found to be $\sqrt{q_\mu^D(\mu_{\text{hypo.}})} - \sqrt{q_\mu^A(\mu_{\text{hypo.}})}$ ($\sqrt{q_\mu^D(\mu_{\text{hypo.}})}$ for CL_{s+b}), resulting in the corresponding asymptotic p-value estimates in Equations 7.15 and 7.18. Intuitively, one can think of the Asimov dataset as representing the expected distributions. By fitting it, the implied expected signal strength and the uncertainty on this expectation are extracted, which for a normal-distributed estimate yields a significance $(\hat{\mu} - \mu)/\sigma$ and corresponding p-value $1 - \Phi((\hat{\mu} - \mu)/\sigma)$.

The step from Equations 7.15 and 7.18 to Equations 7.16 and 7.19 inserts the definition of q_μ . The resulting expressions make explicit that to find these p-values, only the above five fits to find the maximum likelihood estimates of the parameters to the likelihood and the Asimov data, are needed.

Finally, Equation 7.20 defines CL_s , a modification of CL_{s+b} [226, 228].

[227]: Neyman et al. (1933), *On the Problem of the Most Efficient Tests of Statistical Hypotheses*

[222]: Wald (1943), *Tests of Statistical Hypotheses Concerning Several Parameters When the Number of Observations is Large*


[223]: Wilks (1938), *The Large-Sample Distribution of the Likelihood Ratio for Testing Composite Hypotheses*

[220]: Cowan et al. (2011), *Asymptotic formulae for likelihood-based tests of new physics*

[226]: Read (2002), *Presentation of search results: The CL(s) technique*

[228]: (2011), *Procedure for the LHC Higgs boson search combination in Summer 2011*

Figure 7.16: Signal strength scan used to determine the upper end of the one-sided 95% CL_s interval. Here the $m_{Z'} = 40$ GeV signal model is shown as an example. A range of signal strength hypotheses is scanned to find the point where the adjusted p-value CL_s intersects the level α of the test, which is the limit shown in Figure 7.17. μ on the x-axis indicates $\mu_{\text{hypo.}}$ in Equation 7.14 ff. The top axis illustrates the coupling strength of the Z' to quarks corresponding to the tested μ .



33: $1 - \alpha$, with α the value of CL_{s+b} below which the combined signal and background hypothesis is rejected

An interval based on CL_{s+b} covers the true signal strength at the rate expected from the confidence level³³ of the test, but this can lead to unintuitive consequences. E.g. in the case of no observed candidate events, a worse search (with larger backgrounds) will find a stronger exclusion to maintain the design coverage. To avoid such counterintuitive effects, CL_{s+b} is divided by CL_b . CL_b measures the compatibility of the observation with the background-only model, and so the modified CL_s will e.g. be increased in the case of an observation incompatible with the background expectation, or conversely decreased if the background-only and signal-inclusive expectations are similarly compatible with an observation. As a result of the modification, CL_s is not a frequentist p-value, and confidence intervals constructed using it tend to over-cover the true value.

The sensitivity curve is constructed by scanning across signal strength hypotheses, as illustrated in Figure 7.16, to find the point where CL_s drops below 0.05. To estimate the uncertainty on the observed CL_s , and visualize the expected exclusion in the absence of signal, Equation 7.15 are evaluated substituting $Z \in \{-2, -1, 0, 1, 2\}$ for the observed significance, i.e:

$$\text{CL}_b^{\text{exp.,}Z} \approx 1 - \Phi(Z), \quad (7.21)$$

$$\text{CL}_{s+b}^{\text{exp.,}Z}(\mu_{\text{hypo.}}) \approx 1 - \Phi\left(Z + \sqrt{q_\mu^A(\mu_{\text{hypo.}})}\right) \quad (7.22)$$

$$= 1 - \Phi\left(Z + 2 \ln\left(\frac{p(N^A, \mu_{\text{hypo.}}, \hat{\theta}^A)}{p(N^A, \hat{\mu}^A, \hat{\theta}^A)}\right)\right), \quad (7.23)$$

to find the expected CL_s and its uncertainty bands. Note that the expected CL_s and its variations do depend on the observed data, but only indirectly, via the Asimov data. Figure 7.16 shows the resulting median expectation ($Z = 0$) as a dashed line, and the uncertainty as colored bands. As for the

observed limit, the intersection with 0.05 determines the expected limit on μ and its associated uncertainty.

7.7.1.1 Test statistic simplifications

The simplifications in Equation 7.14 ff. are immaterial to the resulting exclusion intervals. The full expressions without simplifications treat two edge cases of observed/hypothesized μ more consistently by substituting differently for q_μ in two cases:

- If the best-fit $\hat{\mu} > \mu_{\text{hypo}}$, the test statistic is kept minimal ($q_\mu = 0$), since for greater-than-hypothesized μ one cannot reject the hypothesis. This has the effect of limiting the allowed significances Z in Equations 7.21 and 7.22 to be no lower than $-\sqrt{q_\mu^A}$. Thus, the expected CL_s is at low μ restricted to values in principle attainable (given sufficiently distinctly signal-like observations) by the observed CL_s .
- negative signal strengths are unphysical because negative event rate is unphysical. Thus for cases where $\hat{\mu} < 0$, we substitute the nearest physical value $\hat{\mu} = 0$ instead, resulting in $q_\mu = -2 \ln(p(N, \mu, \hat{\theta})/p(N, 0, \hat{\theta}_0))$.

In the numerical evaluation of the limits presented here, the test statistic without simplifications is used, corresponding to $\tilde{q}_\mu$ in the notation of ref. [220]. Equation 7.14 ff. are evaluated numerically using the `pyhf` software library[210, 211]. The numerical optimisation of the likelihood in order to find $\hat{\theta}_0$, $\hat{\mu}$ etc. uses `MINUIT` [212] and associated algorithms as implemented by the `iminuit` software package [213].

7.7.2 Expected exclusion intervals

The expected sensitivity of the search is evaluated using a toy pseudodata sample instead of observations. The expected exclusion in terms of the Z' production cross-section is shown in Figure 7.17.

Several features are visible in the limit curve shown in Figure 7.17. The smoothing applied to the multijet contribution of the pseudodata (see Section 7.2) biases the multijet contribution to drop less abruptly than in un-smoothed simulation or observed collision samples. Hence, the pseudodata-based estimates at $m_{Z'} = 20 \text{ GeV}$ and even more so 10 GeV are expected not to accurately anticipate the exclusion potential of the unblinded search, which is not impacted by the smoothing.

Around $m_{Z'} = 90 \text{ GeV}$ sensitivity is reduced compared to a smooth interpolation from 60 GeV through to 125 GeV . This is expected due to the presence of resonant $W^\pm\gamma$, $Z^0\gamma$, and $t\bar{t}$ backgrounds. These irreducible backgrounds are modelled in the likelihood. However, they increase the




Figure 7.17: 95% CL_s limits on the Z' production cross-section for a range of Z' resonance masses, based on simulated pseudodata. Values greater than the shown limit are excluded.

background expectation and systematic uncertainty from simulating the resonant processes. These effects decrease the sensitivity slightly compared to the overall trend. This explanation is tested by confirming the absence of this feature when excluding resonant backgrounds from both the likelihood and pseudodata.

In some models of Z' -like new resonances, the Z' may interfere with e.g. the SM Z^0 [229, 230]. Such interference effects would lead to mass-dependent reductions or enhancements of the combined Z^0/Z' cross-section. The interpretation here does not consider any such interference effects³⁴, instead treating the Z' contribution as only adding to the event rate.

[229]: Accomando et al. (2015), *Exclusion and discovery via Drell-Yan in the 4DCHM*
 [230]: Das et al. (2016), *Classically conformal $U(1)'$ extended standard model, electroweak vacuum stability, and LHC Run-2 bounds*

34: Cf. Section 5.1.2.

7.7.3 Limitations of interpreting the 10 GeV signal extraction fit

The reader may have noticed that the title and abstract of this document describe limits down to 20 GeV in resonance mass, but Figure 7.17 shows the lowest-mass limit at 10 GeV. The reason is that there are several non-trivial aspects to a consistent interpretation of the 10 GeV signal point, from both the experimental and phenomenological perspectives. While the analysis method developed here is able to arrive at a result by simply including the 10 GeV point in the mass grid, the interpretation of the result requires more careful thought and discussion than the higher mass points.

Experimentally, there are several limitations that apply in particular to the $\lesssim 10$ GeV range of signal hypotheses:

- Non-smoothness of the simulated non-resonant background,

- ▶ minimum ρ requirement,
- ▶ transfer-factor extrapolation to lowest jet masses.

Even without a dedicated simulation of QCD resonances, the simulated jet mass spectrum (and $\rho = 2 \log(m/p_T)$ spectrum) exhibits rapidly rising/falling features in this mass range³⁵. This presents a challenge to the D_2 quantile smoothing, increases the risk of accidentally shaping the background into a peak-like shape by placing cuts, and is very susceptible to exacerbating mismodelling of the energy scale and resolution.

35: Cf. Figure 6.5(a).

In the present analysis, the above difficulties are avoided by rejecting resonance candidate jets below some minimum $\rho = 2 \log(m/p_T)$ value³⁶. However, such a ρ selection implies a jet- p_T -dependent mass selection, which at the minimum accepted jet p_T of 150 GeV is equal to requiring $m > 10$ GeV. Thus, the acceptance for jets produced according to a symmetric mass distribution centred on 10 GeV is less than $1/2$, severely limiting sensitivity.

36: See Section 5.2.3 and Table 5.10.

When searching for a signal on the edge of the fitted mass spectrum, the mass-dependent transfer factor parameterization is inherently unable to interpolate from higher and lower masses. Instead, it must extrapolate from one direction, reducing sensitivity and increasing the risk of bias³⁷.

37: Though in practice, it is found that the sensitivity reduction is so large that the increased bias is negligible relative to the lost sensitivity.

To confidently interpret and have sensitivity to signals $\lesssim 10$ GeV, the search strategy would need to be amended to address the above limitations.

More serious, however, are the phenomenological limitations around and below interpreting the mass spectrum at and below 10 GeV:

- ▶ The low-mass single-prong jet mass peak,
- ▶ bottomonium QCD resonances.

Even in the absence of massive resonances, a jet originating from a single light parton acquires some mass as a result of parton shower evolution, fragmentation and hadronization. These processes result in hadrons exhibiting a broad distribution of momenta transverse to the momentum of the originating parton, and thus the overall jet has some mass [61, 68]. The exact mass distribution depends on many factors, such as the distribution of jet momenta, the presence of pileup, and the jet grooming technique³⁸, but in the context of this search, the resulting distribution has a peak near 10 GeV. This breaks the assumption that the absence of resonances implies an absence of peaks in the mass-spectrum. The structure of these jets is nevertheless expected to be single-prong-like, and thus the non-resonant jet mass peak might be constrained in a D_2^{DDT} -failing CR. But any bias in the estimated non-resonant jet tagging efficiency is much more likely to induce a peak-like excess when the background shape is inherently peaked.

38: cf. Section 5.2.2

39: As far as the author is aware, this production mode has not been experimentally demonstrated at the LHC [234]. [235]: Gay Ducati et al. (2010), *Diffraction quarkonium production in association with a photon at the LHC*

[11]: Workman et al. (2022), *Review of Particle Physics*




Figure 7.18: Level scheme of bound $b\bar{b}$ systems. Adapted from ref. [11].

There are excited $b\bar{b}$ bound states with masses $\gtrsim 10$ GeV, such as $\Upsilon(1S)$ and its excitations [231–233], illustrated in Figure 7.18. Prompt-photon-associated production of such bottomonia at the LHC is at least theoretically³⁹ feasible [235]. These states predominantly decay in cascades of lower-mass hadronic states[11], and can thus give rise to genuine two-prong resonances, breaking the assumption that resonant backgrounds arise only due to production of electroweak bosons, directly or in top-quark decays. Due to the high multiplicity of possible hadronic decay chains, the D_2^{DDT} -failing selection is not guaranteed to be void of resonances either. As in the case of the fully captured hadronic top decay at high masses, high-prong multiplicity decays of bottomonia will enter the D_2^{DDT} -failing CRs if the additional prongs are resolved. To correctly account for the impact of this effect requires either subtracting an independently-constrained template of the bottomonium backgrounds, or demonstrating that the expected rate of (partially-) hadronic bottomonia decays entering the fiducial phase-space is negligible.

Thus, to confidently separate out the various contributions to peak-like structures observed in the jet mass spectrum near 10 GeV would require much additional effort beyond the scope of this present work. Without having done said separation, the results of this search being applied using a signal template derived using a 10 GeV Z' boson should not be interpreted as indicating the presence, or absence, of a BSM resonance near said mass.

8.1 In-situ validation results

The dominant non-resonant multijet background to the search presented in this thesis is modelled using a fully data-driven method. This allows for some studies of the fit performance using only simulation-based pseudodata (cf. Chapter 7). However, studies based on pseudodata are insufficient tests of the validity of the background model, since any mistaken assumptions that the simulation shares with the background model are not found in such tests. One example could be hypothetical in-situ differences in the jet response large enough to break the applicability of the tagger decorrelation, but many more may be possible.

To rule out or find such broken modelling assumptions, small samples of observed proton-proton collision data are fit using a background-only likelihood. Initially, the fit is validated in the 2015 dataset (see Section 8.1.1), which had already been studied by a previous iteration of the search, and has small enough luminosity to not be sensitive to BSM resonances that are not already excluded by prior results. As the 2015 sample is not representative of the high-pileup conditions that are typical for most of the Run 2 dataset, a second high-pileup validation sample is also studied (see Section 8.1.2), after having observed satisfactory agreement in the fit to the 2015 sample. The two validation samples are not combined during validation, and only fit separately.

The likelihood used in the validation studies incorporates fewer systematic uncertainties than the full likelihood described in Section 7.1. While the full set of JES systematics are used, the JER parameterization is incomplete and preliminary. The uncertainty on photon reconstruction is as described in Section 7.1.2, as is the uncertainty due to sampling noise in MC templates. Uncertainties on the modelling of the generated backgrounds are incomplete and preliminary. For the $t\bar{t}$ background, these uncertainties are similar to the final, full likelihood, but for the $W^\pm\gamma$ and $Z^0\gamma$ templates, no such uncertainties are included. Uncertainties due to the luminosity measurement, spurious signal, and D_2^{DDT} efficiency scale factor on resonant jets are absent.

Common to all these sources of uncertainty not expressed in the likelihood is that they only impact the resonant backgrounds. Due to their low cross-section relative to the nonresonant background, and the small size of the validation data samples, the impact of the resonant backgrounds on the fit quality is minor. Thus, their exclusion does not deteriorate the ability




Figure 8.1: Mass spectra observed in the 2015 data compared to the expected post-fit distribution. The POI is fixed at 0. The hatched area indicates the uncertainty on the expected yield.

of the validation studies to test the applicability of the non-resonant background estimate.

8.1.1 Validation in 2015 data

Figure 8.1 shows the jet mass spectra in the SR and CRs observed in the validation sample recorded in 2015, and the background-only model expectation maximizing the likelihood. The background-only expectation matches the observed distributions well. In particular, the residual distributions in Figures 8.1(a) and 8.1(b) do not exhibit significant shifts or trends, which would indicate an inability of the non-resonant background model to capture the mass dependence of the non-resonant background. In particular, the fourth-degree Bernstein polynomial basis chosen based on simulated data¹ is sufficient to describe the observed data too. In Figures 8.1(c) and 8.1(d) the model expectation follows the data by construction (cf. Sections 6.1 and 7.1), thus agreement is expected here even if the assumptions of the background model are broken.

1: cf. Section 7.3

Appendix E.1 contains additional figures illustrating the result of the




Figure 8.2: Mass spectra observed in the 2018 L4 data compared to the post-fit expectation. The POI is fixed at 0. The hatched area indicates the uncertainty on the expected yield, to which both prior-constrained and floating NPs contribute according to the full likelihood.

validation fit.

8.1.2 Validation in data from the 2018 L4 period

Figure 8.1 shows the mass distributions observed in the validation sample that was taken from the larger 2018 dataset. The number of events in this second validation sample is similar to that of the 2015 sample described in Section 8.1.1. However, the run conditions are different. In particular, the bunch crossings of the events in this sample contain on average up to ca. 50 additional lower-energy pp collisions, allowing to check the sensitivity of the non-resonant background estimate to pileup.

As before, a background-only likelihood is maximised and the corresponding model expectation compared to the observations. Figures 8.1(b) and 8.2(a) show agreement of this expectation with the data. The best-fit transfer factor polynomial has changed w.r.t. Section 8.1.1, now preferring a lower D_2^{DDT} efficiency in the SR, and a reduced mass-dependence of the efficiency. A similar effect can be recreated using simulated data

by changing how many pileup events are overlaid on one generated hard-scatter event. Thus, the effect in observed data is consistent with the known increased presence of pileup. Hence, the nonresonant background estimate is valid also in high-pileup data.

Appendix E.2 contains additional figures illustrating the result of the validation fit.

8.2 15.8 fb^{-1} unblinding results

Before analysing the full dataset, a search is done using only part of the observed data. These partial data correspond to 15.8 fb^{-1} or roughly 10 % to the full Run 2 data, recorded in 2015 and 2016. This partial unblinding is done to be better able to find any potential issues relating to the applicability of the data-driven background model before analysing the full observed data, which may bias analysers.




Figure 8.3: Post-fit mass spectra of the signal extraction fit.

For each signal model (i.e. each Z' mass) under consideration a separate maximum-likelihood fit is done to determine the best-fit signal-strength

for the given Z' mass. In this section, one such fit for $m_{Z'} = 50$ GeV is examined as a representative example. Due to an up-fluctuation in the observed data near $m_J = 50$ GeV in the forward-photon CR, the best-fit model has a non-zero signal strength (see Figure 8.3). However, since the tension in the data is comparable to the statistical uncertainty on the counts in that bin, the extracted signal strength is not large relative to the fit uncertainty.




Figure 8.4: 95% CL_s limits on Z' production for a range of Z' resonance masses.

Since the background-only fit describes the observed distributions well, limits are set on the signal strengths of the signal models. For each Z' mass hypothesis, a range of signal strength hypotheses is tested for compatibility with the observed data, as well as the expected distribution given the best-fit background-only hypothesis. This scan over all masses is summarized by Figure 8.4, where the limit is expressed in terms of observed Z' production cross-section rather than signal strength. The limit grows weaker towards higher masses, which is due to the reduced constraints on the background model in this mass region. At the high-mass end of the spectrum, the background sample observed in the CRs contains fewer events than at lower masses. This is exacerbated by the fixed resonance candidate jet R -parameter value of 1.0 and jet collimation cut², which for e.g. $m = 200$ GeV implicitly require $p_T \geq 400$ GeV, suppressing both SR and CR acceptance.

2: Cf. Table 5.10 and Section 5.2.3.

8.3 Fully unblinded results

Early during the unblinding on the full dataset recorded during 2015-2018, a significant discrepancy between the central-photon and the forward-photon pass-to-fail ratios was noticed. Figure 8.5 shows these ratios. A




Figure 8.5: Pass-to-fail ratios and transfer function polynomials observed in 2015-2018 data. The ratios are shown after subtracting the prior expected distributions due to the $W^\pm\gamma$, $Z^0\gamma$, and $t\bar{t}$ production processes.

3: In high-mass bins, the small sample sizes give rise to bin-by-bin fluctuations greater than the difference seen at lower values. These fluctuations may obscure a constant difference between ratios extending to higher jet masses.

Table 8.1: Measured integrated luminosity, by year of data taking [92].

Year	$\int L$
2015-2016	36.6 fb^{-1}
2017	44.6 fb^{-1}
2018	58.8 fb^{-1}

[92]: ATLAS Collaboration (2022), *Luminosity determination in pp collisions at $\sqrt{s} = 13 \text{ TeV}$ using the ATLAS detector at the LHC*

4: It is found that this separate scaling is necessary to achieve a consistent subtraction that does not generate an artificial deficit in the pass-fail ratios of Figure 8.5.

clear difference is visible starting from ca. 100 GeV, which extends to at least 150 GeV, but possibly to greater values³.

To distinguish the expected difference due to resonant backgrounds from the difference observed here, the nominal simulated resonant background templates are subtracted from the observed jet mass distributions before taking the ratios. Data recorded in different years is subject to different collider- and detector conditions, and versions of the simulated templates attempting to match these changed conditions are combined. Prior to the combination, each template version is normalized to the measured luminosity [92] of the corresponding year⁴, as listed in Table 8.1. A clear difference remains visible after subtracting the resonant background expectation.

The assumption that the in-situ pass-to-fail ratio in the forward-photon selection predicts the same ratio in the central-photon selection was used to motivate the construction of the analysis regions, and the form of the likelihood. The observed difference of the D_2^{DDT} pass-to-fail ratio between SR and CR explicitly breaks this assumption. Thus, before proceeding




Figure 8.6: Post-fit mass spectra of the background-only fit.

further with analysing the unblinded dataset, a suite of studies is carried out aiming to isolate the cause of the observed difference, or a particular phase-space region contributing to it. However, no such cause or phase-space region is found.

In detail, the following explanations are tested:

- additional resonant SM processes:
 - photon-associated $t\bar{t}$ production,
 - $W^\pm$ -associated single-top-quark production,
 - single-top-quark production.
- contamination with pile-up jets,
- special LHC fills with unusual bunch schemes (8b4e),
- photon trigger inefficiency at low p_T ,
- insufficient photon isolation requirements,
- residual correlations with the resonance-candidate jet p_T ,
- angular acceptance of the jet selection,
- angular acceptance of the photon selection,

- the value of the forward-photon/central-photon region split being close to the barrel-endcap transition region,
- a need to apply and derive decorrelation quantile maps separately per year of data taking,
- the difference being particular to using the D_2^{DDT} observable to tag two-prong-like jet shapes.

However, none of the tested hypotheses are able to remove or explain the observed difference.

One remaining tentative account of the origin of the observed ratio difference is a correlation between photon (pseudo-)rapidity and final state parton flavour. Such a correlation may come about as a result of different dependence on photon rapidity of the leading order terms contributing to the photon-associated jet production. The Compton-scattering-like process initiated from a gluon and quark state⁵ has LO contributions from both time-like (s-channel) and space-like (t- and u-channel) amplitudes, and results in a final state quark. The annihilation of a pair of initial state quarks⁶, however, has only space-like contributions, and results in a final state gluon. Thus, if the quark-annihilation process has a slightly lower overall contribution to the low-rapidity part of the differential cross-section, the difference in final state parton flavour may account for different D_2^{DDT} efficiencies.

This account of the gap originating from simple properties of the non-resonant background scattering process is consistent with the available sample of simulated events for said process. However, the size of this simulated sample is too small to resolve the observed difference of pass-fail ratios. Hence, it is not possible to perform a satisfactory test of the explanation unless a significantly larger simulated sample is obtained.

The observed forward-central difference is not consistent with the narrow-width Z' hypotheses that the search targets. This is concluded from two observations:

1. The observed mass-width of the difference is far greater than the jet-mass-resolution-width peaks predicted for narrow resonances;
2. The observed difference in the pass-fail ratio is larger than both the difference expected in the presence of a simulated narrow resonance, and the difference observed at the $W^\pm$ (and Z^0) peak.

The degree of the polynomial parametrizing the transfer function⁷ was chosen to be four in Section 7.3 based on simulated pseudo-data. The compatibility of this background model with the observed collision data is re-examined, because the observed background may exhibit a shape the simulation does not predict. For the test of compatibility to observed data, the test statistic used is the ratio of two likelihoods. When testing a given polynomial degree d , the likelihood⁸ L_d of the best-fit background only model using a degree d polynomial is compared to the likelihood L_{d+1} of

5: cf. Figures 6.1(b) and 6.1(c)

6: cf. Figure 6.1(a)

7: see Section 7.1

8: cf. Equation 7.1




Figure 8.8: 95% CL_s limits on Z' production for a range of Z' resonance masses.

the best-fit background-only model using a degree $d + 1$ polynomial. In the asymptotic limit of many observed events, $-2 \ln(L_d/L_{d+1})$ follows a χ^2 distribution [220, 222, 223], allowing a straightforward computation of p -values for the likelihood ratio. Following the recommendation of ref. [188], the initial background model is accepted as sufficient for $p(L_d/L_{d+1}) \geq 0.05$. If the initial model is not accepted, d is incremented once and the test is repeated.

The likelihood-ratio test described above rejects the default background model using the fourth-degree polynomial, as Figure 8.7 shows. The lowest polynomial degree that passes the criterion is of degree five, so this transfer polynomial parametrization is adopted for the remaining interpretation of the data.

The difference in pass-to-fail ratios observed between the central-photon SR and forward-photon CR may bias the result of the likelihood ratio test. Thus, the test is repeated masking the mass range between 110 GeV and 140 GeV, but the test on masked data also rejects $d = 4$ and accepts $d = 5$.

Similarly, there is a narrow range of bins at low masses where the background-only fit exhibits a significant residual to the observed data (see Figure 8.6). If due to the presence of a signal-like feature in the data, the inclusion of these bins may bias the likelihood ratio test. The test is thus repeated while masking data with m_l between 34 GeV and 64 GeV. Finding $p(L_4/L_5) = 2.3 \times 10^{-5}$ and $p(L_5/L_6) = 0.41$, the original conclusion is upheld.

95% CL_s exclusion limits are calculated as described in Section 7.7.1, illustrated in Figure 8.8. Over most of the tested mass range, Z' production is excluded, with excluded visible cross-sections ranging from 0.01 pb to 0.2 pb, depending on the signal mass hypothesis. For $m_{Z'} \geq 100$ GeV,

[220]: Cowan et al. (2011), *Asymptotic formulae for likelihood-based tests of new physics*

[222]: Wald (1943), *Tests of Statistical Hypotheses Concerning Several Parameters When the Number of Observations is Large*

[223]: Wilks (1938), *The Large-Sample Distribution of the Likelihood Ratio for Testing Composite Hypotheses*

[188]: ATLAS (2020), *Recommendations for the Modeling of Smooth Backgrounds*




Figure 8.7: Iterative likelihood-ratio tests for different transfer function polynomial degrees. The default polynomial degree is four, as determined in Section 7.3. The criterion to accept or reject a given background model is explained in the text.




Figure 8.9: Post-fit mass spectra of the $m_{Z'} = 40$ GeV signal extraction fit.

9: cf. Figure 7.17

10: cf. Figure 7.17 and Sections 7.7 and 7.7.3.

the expected limit is weaker than expected based on simulation⁹. Further, the observed limits are weaker than the expected limits based on Asimov data for this range. Both effects are due to the unexpected central-forward pass-to-fail ratio difference discussed earlier in this section, shown in Figure 8.5. The weakening of the expected limit near 80 GeV – 90 GeV and at 10 GeV is expected, due to the resonant backgrounds and analysis acceptance, respectively¹⁰. The observed limits at 40 GeV and 50 GeV exceed the expected limit, at more than two standard deviations. These two signal templates are separated by a distance of order of the jet mass resolution, thus the excess in both signal hypotheses may be caused by a common feature in data.

Figure 8.9 shows the best-fit of the $m_{Z'} = 40$ GeV signal model. The fitted excess seems to be driven by contributions in both the SR and the forward CR. The respective pass-fail ratios in Figure 8.5 exhibit a narrow jump in the same bins. The inclusion of the signal contribution induces a downward shift in the values attained by the transfer function polynomial in the bins exhibiting the excess. However, the downward shift is covered by the respective post-fit uncertainties on the polynomial

coefficients of the background-only and signal-inclusive fits. Thus, the nominal uncertainty bands of the exclusion plot in Figure 8.8, which includes these uncertainties on the nonresonant background estimate, should cover this effect.

Many nuisance parameters related to jet energy and σ -resolution are constrained in the fit. This may be expected, since the full 2015-2018 dataset is used in the fits shown here, but the in-situ anti- k_T $R = 0.2$ jet calibration is based on only the 2015-2017 subset. The most-impacted such NPs are the energy scale uncertainties relating to parton flavour composition and response. The constraint on NPs is consistent between the background-only and full fits, as well as when masking the bins near 40 GeV.

The degree of (dis-)agreement with the background-only expectation is quantified by finding the p-values of the “discovery” test statistic q_0 of ref. [220]. This p-value is the probability of an excess as large as or larger than the observed one, according to the likelihood of the background-only model. The local significance of the excess near 40 GeV is equivalent to 3.9 standard deviations. The global significance corrects the local significance to account for the increased probability of encountering an excess in a search testing multiple independent signal hypotheses. Using an asymptotic correction [236], the global significance of the same excess is equivalent to ca. 2.8 standard deviations.

[220]: Cowan et al. (2011), *Asymptotic formulae for likelihood-based tests of new physics*

[236]: Gross et al. (2010), *Trial factors for the look elsewhere effect in high energy physics*

9.1 Summary

This thesis presents a search for hadronically decaying new resonances. For the first time, LHC pp collision data recorded with the ATLAS experiment is searched for such resonances with masses down to 20 GeV. To access such low masses¹ the recording of 140 fb^{-1} of analysed data is triggered by the incidence of a highly energetic photon². Such a photon may be produced as ISR in events where the new resonance decays in the opposite side of the detector. The transverse momentum of the ISR photon is balanced by the momentum of the hadronic radiation³. Hence, the momentum of light resonances is greater than their mass, and they are highly boosted and collimated in the rest-frame of the detector. This collimated hadronic radiation is reconstructed as a large- R jet, enabling one unified reconstruction to cover a range of masses and boosts.

The usual ATLAS large- R jet calibration is limited to masses greater than 50 GeV. To achieve calibrated large- R jets with lower masses, large- R jets are instead reclustered from calibrated small- R jets and charged tracks⁴. Because the inputs are already well-calibrated, the calibration and related uncertainties can be propagated bottom-up from the inputs. For best granularity, a novel $R = 0.2$ jet calibration is applied, which is particularly beneficial at low masses where the $R = 0.4$ jets is too coarse to resolve a two-prong decay as two clusters.

The great majority of background events are inelastic pp collisions that do not involve any resonances, where an energetic photon is emitted in the hard scatter⁵. In such background events, the hadronic radiation with high transverse momentum is most commonly due to the production of a single energetic gluon or quark. The energy flow of this radiation thus tends to be concentrated in a single prong, surrounded by a sparse, isotropic pattern of secondary emissions. This pattern of non-resonant radiation is different from that of a resonance decaying into a quark-antiquark pair, which yields a characteristic pattern of two distinct prongs where most of the energy flow is concentrated.

This shape difference between non-resonant background and resonant signal events is exploited to select a sample where resonances are more frequent, and the background rate reduced. This is achieved by requiring the ratio of energy-correlation functions D_2 [198] to be small in the SR. However, a sample of D_2 -selected non-resonant jets has a different mass distribution compared to the inclusive sample. To avoid biasing the search by this effect, the value of D_2 required to enter the SR is chosen as

1: Searches for light hadronic resonances are challenging due to trigger thresholds, see Section 2.4.

2: See Sections 5.1.1 and 5.2.1

3: Cf. Sections 5.2.3 and 5.2.4

4: See Section 5.2.2.

5: See Section 6.1.

[198]: Larkoski et al. (2014), *Power Counting to Better Jet Observables*

6: See Section 6.2.

a function of the mass and p_T of the jet, such that the fraction of selected jets is constant as a function of p_T and mass⁶.

7: See Section 7.3.

Modelling the dominant nonresonant background using simulation is challenging due to high expected yields and a complex dependence on intricate details of hadronic radiation patterns. Instead, its shape and rate are modelled using observed data. The observed events failing the D_2 selection are used to predict the expected mass distribution in the SR. Deviations from the constant expected D_2 efficiency are modelled using a polynomial parametrization. The number of coefficients parametrizing the polynomial is chosen based on absence of bias in simulated events and compatibility with simulated and observed data⁷. To better distinguish the impact of resonant backgrounds or signals on the fitted polynomial, it is simultaneously fitted in the SR and a second sample of events with high-pseudorapidity ISR photons, where resonant yields are reduced⁸.

8: See Section 6.3.

9: See Section 5.1.2.

Resonant background processes, producing $Z^0\gamma$, $W^\pm\gamma$, and $t\bar{t}$, result in signal-like contributions to the SR at the $W^\pm$ and Z^0 masses. These resonant backgrounds are modelled using simulation⁹, and included in the signal extraction fit as mass templates. The efficiency of the D_2 selection when applied to resonant backgrounds and the signal is measured using a CR targeting semi-muonically decaying $t\bar{t}$ pairs¹⁰, and the resulting scale factor is applied to the signal and resonant background regions.

10: See Section 6.4.

11: See Section 7.1.

A combined likelihood is constructed, modelling the mass distributions of the signal- and background processes in the SR and related CRs¹¹. The compatibility of the data with the background-only hypothesis and various signal hypotheses is examined by comparing ratios of the maximal values the likelihood can attain while constrained to the respective hypotheses¹².

12: See Section 7.7.1.

13: See Section 8.2

Several overlapping subsets of the pp collision data recorded by ATLAS from 2015 - 2018 are analysed. The smaller subset corresponds to 15.8 fb^{-1} . It is found to be consistent with the assumptions underlying the likelihood¹³. In particular, the mass dependence of the observed D_2^{DDT} efficiency of resonance candidate jets appears to be consistent between events with a forward ($|\eta| > 1.3$) as compared to a central ($|\eta| < 1.3$) hard photon. In the small dataset, no significant deviations from the SM expectation are observed, and so limits are set on the visible production crosssection of BSM resonances¹⁴.

14: See Figure 8.4

15: See Section 8.3

The full dataset corresponding to 140.1 fb^{-1} is similarly analysed¹⁵. However, with the additional statistical power enabled by the larger dataset, an unexpected feature is resolved: the observed D_2^{DDT} efficiency of resonance candidate jets with masses $\geq 100 \text{ GeV}$ appears to depend on the pseudorapidity of the associated hard photon. The efficiency ratio $\varepsilon/(1 - \varepsilon)$ in this sample of jets is a few percentage points higher when the photon $|\eta| < 1.3$ compared to the sample where $|\eta| > 1.3$. This breaks

the assumption that the forward-photon sample may be used to constrain or predict the efficiency (-ratio) in the central photon sample, first used in this context in ref. [53]. This assumption is a core part of the form of the likelihood used for the search presented here. Since it is not upheld in data at jet masses $\gtrsim 100$ GeV, the search is less sensitive than anticipated in that mass range.

The full dataset is analysed to identify potential new resonances. A hint of an excess, with a global significance of ca. 2.8σ , is found in jets with masses near 40 GeV. In the remainder of the studied mass interval, 95 % CL_s upper limits on the visible cross-section for production of new resonances are set. These range from 0.02 pb for a Z' resonance with a mass of 20 GeV to 0.2 pb for a 225 GeV resonance, representing the most sensitive constraints on the leptophobic Z' model in the mass range 20 GeV through 60 GeV at time of writing. Furthermore, these are the first such constraints below 100 GeV set using ATLAS data.

9.2 Outlook

Future searches using similar approaches as the one presented here will need to account for the unanticipated dependence of the resonance candidate jets D_2^{DDT} efficiency on the ISR-photons pseudorapidity. The simplest way to do so is to abandon the separation into forward-photon and central-photon (or similar) control regions. Instead, the influence of spurious signal-like fluctuations in observed data can be reduced by constructing separate control- and validation regions in mass side-bands around the SR, as e.g. ref. [59] does. This does however necessitate the use of separate likelihoods, being separately tested and validated, with separate background estimates, for each of the signal hypotheses being considered. Furthermore, it limits the potential benefit of non-parametric background estimates (e.g. Gaussian kernel regression), because the mass-width of the SR controls both CR contamination¹⁶ and the uncertainty on the background interpolation into the SR¹⁷. But parametric estimates, e.g., using a polynomial series as in this thesis, are less affected, since they can interpolate beyond the auto-correlation length of the background distribution.

9.2.1 Extended datasets

The search presented here, evaluated using 140 fb^{-1} of pp collision data, is statistically limited. For most signal mass hypotheses, this true in the usual sense — the most important contribution to the uncertainties limiting the sensitivity is the expected variance of observed event counts in the bins of the SR. However, it is true also in a less conventional sense: the most important uncertainties on the expected number of background events,

[53]: Sirunyan et al. (2019), *Search for Low-Mass Quark-Antiquark Resonances Produced in Association with a Photon at $\sqrt{s} = 13$ TeV*

[59]: Aaboud et al. (2018), *Search for low-mass dijet resonances using trigger-level jets with the ATLAS detector in pp collisions at $\sqrt{s} = 13$ TeV*

16: and thus spurious signal, which is smaller with a wider SR

17: which is smaller with a narrower SR

18: See Section 7.6.

19: expected to deliver ca. 250 fb⁻¹ of pp collision data according to ref. [237].

[60]: Sirunyan et al. (2019), *Search for low mass vector resonances decaying into quark-antiquark pairs in proton-proton collisions at $\sqrt{s} = 13$ TeV*

[197]: Sirunyan et al. (2018), *Search for low mass vector resonances decaying into quark-antiquark pairs in proton-proton collisions at $\sqrt{s} = 13$ TeV*

20: trigger-level analysis, see Section 2.4

which would usually be counted among the systematic uncertainties, are reduced as the size of the dataset is increased¹⁸. This reduction is due to the data-driven determination of the background, the precision of which increases with the number of events observed in the D_2^{DDT} -failing CRs.

Thus, future datasets collected during the currently ongoing run 3 of the LHC¹⁹, and future runs 4 through 6, anticipated to deliver a total high-luminosity LHC dataset on the order of 3 ab⁻¹ [238], can be expected to enable significantly more sensitive iterations on the present search.

The additional sensitivity will likely make it necessary to model additional resonant backgrounds to avoid biasing the non-resonant background estimate. Examples might be photon-associated top or $t\bar{t}$ production, which the current iteration can neglect due to their low rate, but which may become more relevant with increased luminosity. Becoming more sensitive to processes involving top decays may motivate extending the use of the $t\bar{t}$ CR, to e.g., independently constrain the yields or shapes of some of these backgrounds in-situ.

Furthermore, a 3 ab⁻¹ dataset will likely necessitate re-examining the two-prong tagger mass decorrelation method. The smoothing of tagger quantiles derived from simulated data is currently able to alleviate the insufficient size of simulated data. However, it is not a given that fully simulated sample sizes will be (able to be) increased at a rate proportional to that of the increase in observed events. Similarly, inaccuracies of the simulation photon-associated multijets production, such as neglecting higher orders in the perturbative expansion of the matrix element, may end up biasing the mass-decorrelation, which is constructed using simulated samples. Additional freedom of the fitted transfer factor estimate, e.g., using higher order polynomials, additional dependent variables such as jet p_T as in refs. [60, 197], or a non-parametric estimate, may be needed to keep the background estimate accurate.

Finally, it is important to note that the high luminosity of later LHC runs will be achieved by substantially increased instantaneous luminosity. This implies corresponding increases in pile-up. To be able to record data under higher-pile-up conditions, trigger thresholds are likely to be increased. This has two implications for future iterations of this search. Firstly, it partially counteracts the sensitivity gained from a larger dataset, since the cross-section for the signal $Z'\gamma$ production above some ISR- γ p_T threshold falls as the threshold increases. Secondly, the increased boost of recordable resonances might make the substructure of the resonance candidate jet more difficult to resolve, thereby deteriorating the jet mass resolution, and two-prong tagging performance. These effects of increased trigger thresholds, if found to be significant, may be mitigated to some degree by performing the search as a TLA²⁰, reducing the trigger threshold.

A second effect of increased pile-up is the direct impact on jet reconstruction. In general, mass and substructure reconstruction suffer increased bias and decreased resolution under high-pile-up conditions. Hence, improvements to the robustness of jets to pile-up, e.g., by means of better jet grooming, improved reconstruction techniques (e.g., UFO jets [182]), or pile-up tagging methods, are likely to be very relevant to future iterations on this search.

[182]: Aad et al. (2021), *Optimisation of large-radius jet reconstruction for the ATLAS detector in 13 TeV proton–proton collisions*

9.2.2 Extended methodology

Some modifications may be required to adapt the search to a larger dataset and changing data-taking conditions, but others can be of inherent interest. One promising example is the use of machine learning methods to obtain a more performant jet tag. Such taggers have previously been found to outperform analytical taggers such as D_2 [191, 239, 240] when tagging $W^\pm \rightarrow q\bar{q}$ decays. By making mass-decorrelation part of the loss being minimized when training such a classifier, the reduction in signal efficiency associated with quantile-based decorrelation may be mitigated to an extent [191, 241]. While the aforementioned examples use high-level jet features (such as D_2, η, p_T) as inputs, further signal efficiency may be gained by using networks operating on graph representations based on constituent-level jet observables [163]. Such constituent-based classifiers may be of particular interest when using TAR jets, as the constituents are already calibrated.

Another extension enabled by calibrated TAR constituents jets are variable- R reclustered jets [174]. By letting the iterative anti- k_T clustering termination criterion depend on the p_T or mass of the pseudo-jet clustered up to the latest iteration, a smaller²¹ effective cone-size for higher- p_T jets can be achieved. This may be beneficial for both the low-mass and high-mass regime. Low mass resonances may benefit from a reduced impact of pile-up²² on highly collimated high- p_T resonance candidate jets. Meanwhile, the sensitivity to high-mass jets may benefit from an ability to capture all decay products of the resonance, even if the angle ΔR between them would be too great to allow clustering into a single fixed- R anti- k_T jet.

[174]: Krohn et al. (2009), *Jets with Variable R*

21: And vice versa, the effective cone-size can be increased for softer jets.

22: which is proportional to the catchment area of the jet, due to the isotropically distributed transverse energy density of pile-up

It would be interesting to extend the search and interpretation to even lower jet masses. Section 7.7.3 discusses many of the challenges such an extension is likely to face. The most pressing one is the suppression and accurate modelling of the low-jet-mass peak. The jet grooming method used for the TAR jet reconstruction in the current iteration of the analysis, trimming [66], is challenging to use when measuring very light jets because non-global logarithms induce said peak and limit the accuracy of fixed-order simulation in its modeling [68]. More recent approaches to grooming, such as the (modified-) Mass Drop or Soft Drop methods [67, 68], can avoid the issue of non-global logarithms, and may thus be

[66]: Krohn et al. (2010), *Jet Trimming*

[67]: Larkoski et al. (2014), *Soft Drop*

[68]: Dasgupta et al. (2013), *Towards an understanding of jet substructure*

essential to accessing hadronic resonances at 10 GeV and below.

APPENDIX

The Kolmogorov-Smirnov test and measuring residual mass shaping

A

In the context of the fixed-efficiency mass decorrelation method described in Section 6.2, the choice of the Jensen-Shannon distance to measure the extent of residual mass shaping after decorrelation sometimes triggers a question along the lines of "Why not use the Kolmogorov-Smirnov test (-statistic) instead?". To discuss this question in great detail in Section 6.2 would give undue weight to a tangential point, so this appendix covers it instead.

The two-sample Kolmogorov-Smirnov test (KS test) is sometimes used to test the compatibility of distributions. It is not an appropriate tool here since we need a metric of the distance between distributions, which the KS test with a binary result is not. The KS test statistic is a distance metric, but it still has unsuitable properties. It is constructed as the supremum of the difference at any point between the two empirical cumulative distribution functions. This makes it less sensitive to small differences spread out over many points of the compared distributions. These will tend to average out to a more stable value as difference are accumulated, which is unsuitable for the use case described here, which requires sensitivity to such deviations.

Two histograms can be statistically compatible by the KS-test but different in detail, which would not be an optimal solution to the problem of mass decorrelation. When tuning, we still prefer the tagger that shapes the mass less, even if our sample is not sufficient to confidently rule out differences due to statistical fluctuations. This is possible with the JS distance, but not with a test of compatibility, or the less sensitive KS distance.

This intuition of the insensitivity of the KS-test statistic to small deviations is one consequence of the more general formal result that the KL divergence (on which the JS distance is based) can be shown to upper-bound the KS-test statistic. That is [242],

$$KS(P, Q) \leq \sqrt{\frac{KL(P, Q)}{2}}, \quad (\text{A.1})$$

where KS is the Kolmogorov-Smirnov distance, KL the Kullback-Leibler divergence, and P and Q are two probability distributions. By Equation A.1 the KL divergence (and thus, the JS distance, cf. Equation 6.3) between any two distributions implies an upper bound on the greatest value the KS distance may take, but the reverse does not hold. In this sense, the JS distance is generally at least as sensitive to differences as the KS distance, and in the context of mass-decorrelating a two-prong tagger, the case where the JS distance is more sensitive can matter.

[242]: Gibbs et al. (2002), *On Choosing and Bounding Probability Metrics*

Auxiliary material on the smoothed multijet background template

B

This appendix collects additional material auxiliary to the discussion of the nonresonant template smoothing discussed in Section 7.2.

B.1 Further illustration of excessive MC noise in simulated samples

A background template base on a MC sample of sufficient size would show the variance in the shapefactors' pre- and postfit values to be negligible relative to the expected yield in the bin (i.e., the expected variance of observations). Figures B.1 and B.2 show that the available sample does not meet this condition. The pre-fit shapefactor values are the anti-tagged CRs by construction. The post-fit values attempt to find a consistent average of anti-tagged and tagged yields as a result of the fitting. However, the found values disagree with the pre-fit values by much more than the expected data variance. Since there is no visible pattern to the disagreement, it is interpreted as being due to MC noise, and it is found to be consistent with the known level of MC statistical uncertainty of the simulated templates. Thus, the size of the simulated sample is insufficient for the targeted observed luminosity.

Figure B.1: Pre-fit mass spectra including non-resonant QCD background simulations in the D_2^{DDT} -passing regions. Both the central-photon selection in (a) and the forward-photon selection in (b) are shown. The pseudodata are compared to a simple pre-fit where shapefactors are set such that the expected yield of the D_2^{DDT} -failing regions exactly matches the observations, and with transfer polynomial coefficients taken from a simple least-squares fit to only the forward-photon D_2^{DDT} -pass-to-fail ratio. The expected sample yields of the likelihood and the pseudodata distribution are scaled to an integrated luminosity of 140fb^{-1} , and consist of the combination of mc16a, mc16d, and mc16e samples. Error bars indicate the standard deviation of a Poisson distribution with expectation value equal to the bins pseudodata content. Figure 7.6 shows the corresponding maximum-likelihood expectation.





Figure B.2: Comparison of pre- and postfit shapefactor values of the fits to the un-smoothed pseudodata shown in Figures 7.6 and B.1. The shapefactors are multiplicative unconstrained modifiers that are applied to unitary nominal yields. Their initial values are the observed yields in the D_2^{DDT} -failing control regions less the resonant backgrounds. In a consistent fit, and assuming the non-resonant background efficiency of D_2^{DDT} is treated correctly in the likelihood, the post-fit parameter values should vary according to the binwise Poisson error of the D_2^{DDT} -passing and D_2^{DDT} -failing observed yields, which are dominated by the much larger D_2^{DDT} -failing observed yields. This motivates choosing $\sqrt{\theta_{\text{initial}}}$ as the uncertainty to scale the postfit-to-prefit parameter difference by.

B.2 Further KDE smoothing variants and examples

Figure B.3 shows the result of applying KDE smoothing¹ to a normal distribution. Because the kernel used here is itself normal, this example illustrates an "ideal" case where even with low sample statistics or large bandwidths, the bias due to the kernel shape is minimal.

1: see Section 7.2.2

Figures B.4 to B.6 illustrate kernel smoothing of toys sampled from simulated mass distributions, for a range of kernel bandwidths. The smoothing used for the studies in Chapter 7 uses a bandwidth value of 0.16, which is shown in Figure 7.8 and not repeated here.

Figure B.3: Toy example of using a KDE to smooth a noisy histogram sampled from a normal distribution (green curve). The number of samples used to fill the noisy example histogram is indicated as `n_samples`. The size of the toy dataset the KDE is fit to is $10 * n_samples$. The number of bins of the histograms is chosen as the maximum of the number of optimal bins according to Sturges' rule and the Freedman-Diaconis rule [215], so it increases as the number of toys increases. The bandwidth of the KDE is chosen as twice the width of a histogram bin.





Figure B.4: Examples of using a KDE to smooth a noisy multijet background mass histogram of the signal region. The original noisy simulated template is shown as the green line. A toy histogram is sampled from it, shown as the blue filled bar chart, to induce additional artificial noise for illustration. (This is not done when smoothing histograms to create the smooth QCD multijet template.) The number of samples used to fill this noisy example histogram is indicated as `n_samples`. The size of the toy dataset the KDE is fit to is $10 * n_samples$. The width of the mass bins is fixed at 2 GeV. The bandwidth of the KDE is 0.04.






Figure B.5: Examples of using a KDE to smooth a noisy multijet background mass histogram of the signal region. The original noisy simulated template is shown as the green line. A toy histogram is sampled from it, shown as the blue filled bar chart, to induce additional artificial noise for illustration. (This is not done when smoothing histograms to create the smooth QCD multijet template.) The number of samples used to fill this noisy example histogram is indicated as `n_samples`. The size of the toy dataset the KDE is fit to is $10 * n_samples$. The width of the mass bins is fixed at 2 GeV. The bandwidth of the KDE is 0.12.




Figure B.6: Examples of using a KDE to smooth a noisy multijet background mass histogram of the signal region. The original noisy simulated template is shown as the green line. A toy histogram is sampled from it, shown as the blue filled bar chart, to induce additional artificial noise for illustration. (This is not done when smoothing histograms to create the smooth QCD multijet template.) The number of samples used to fill this noisy example histogram is indicated as `n_samples`. The size of the toy dataset the KDE is fit to is $10 * n_samples$. The width of the mass bins is fixed at 2 GeV. The bandwidth of the KDE is 0.2.

Auxiliary material about fit linearity

C

Figure C.1 shows more details of the distributions underlying the linearity tests described in Section 7.5.




Figure C.1: Linearity of the fit response to varying injected signal strength (Boxplot). For each injected μ , the distribution of best-fit $\hat{\mu}$ values is indicated by an extended box plot. The central line indicates the median of the distribution, the central quartiles are indicated by the boxes above/below. Successively halved quantiles of the remaining sample are shown by successively smaller boxes, and outlying data are indicated as a scatterplot beyond the final box.

C.1 Injection up to higher values of μ

This section shows the results of testing the linearity of the signal extraction fit response for large signal strengths ($\mu \geq 1$), complementing the scans over smaller values documented in Section 7.5. For every mass point, six equidistant signal strength values from 0 through 5 are injected. For every combination of signal mass and signal strength, 1000 independent toys are drawn. The toys are drawn from a bin-wise Poisson distribution with expectation value equal to the sum of the resonant background templates, and the smooth QCD multijet template described in Section 7.2.




Figure C.2: Linearity of the fit response to injected signal strengths $\mu \geq 1$. The grey line corresponds to $\hat{\mu} = \mu$. The sample average best-fit signal strength $\langle \hat{\mu} \rangle$ is indicated by the coloured central line. The sample standard deviation of $\hat{\mu}$ is indicated by the shaded area.




Figure C.3: Linearity of the fit response to varying injected signal strength $\mu \geq 1$ (Boxplot). For each injected μ , the distribution of best-fit $\hat{\mu}$ values is indicated by an extended box plot. The central line indicates the median of the distribution, the central quartiles are indicated by the boxes above/below. Successively halved quantiles of the remaining sample are shown by successively smaller boxes, and outlying data are indicated as a scatterplot beyond the final box.

Auxiliary material on systematic uncertainties

D

Figures D.1 and D.2 show further examples of template variations beyond those discussed in Section 7.1.2.




Figure D.1: Effect of the JES NP JET_EffectiveNP_Modelling1 on the $m'_Z = 20$ GeV signal template across the four analysis regions (see Table 6.1 for their definition).




Figure D.2: Effect of the photon energy scale NP EG_SCALE_ALL on the $Z^0\gamma$ background template across the four analysis regions (see Table 6.1 for their definition).

Auxiliary material on the in-situ validation results

E

E.1 Validation in 2015 data

This appendix collects additional visualizations of the background-only fit to the data observed in the 2015 validation sample described in Section 8.1.1.




Figure E.1: Mass spectra observed in the 2015 data compared to the expected pre-fit distribution. NPs without prior constraints are set to values that achieve an exact match to the observed distribution in the anti-tagged control regions, and transfer factor polynomial coefficients are chosen to minimize the squared error between the polynomial and the ratio of tagger-passing and -failing resonant-background-subtracted observed yields in the forward-photon control regions. NPs with prior constraints are set to their nominal values (maximizing their probability according to the constraint term). Data from the SR is not used to choose NP values. The POI is fixed at 0. The hatched area indicates uncertainty on the expected yield, which is purely due to the width of the prior constraint terms of constrained NPs.




Figure E.2: Pulls of NPs with prior constraints (except MC statistics). See Figure E.3 for pulls of MC statistical uncertainty NPs.




Figure E.3: Pulls of MC statistical uncertainty NPs.




Figure E.4: Change in the floating shapefactor NPs relative to their initial value in the fit. The shapefactors are multiplicative unconstrained modifiers that are applied to unitary nominal yields. Their initial values are the observed yields in the D_2^{DDT} -failing control regions less the resonant backgrounds. In a consistent fit, and assuming the non-resonant background efficiency of D_2^{DDT} is treated correctly in the likelihood, the post-fit parameter values should vary according to the binwise Poisson error of the D_2^{DDT} -passing and D_2^{DDT} -failing observed yields, which are dominated by the much larger D_2^{DDT} -failing observed yields. This motivates choosing $\sqrt{\theta_{\text{initial}}}$ as the uncertainty to scale the postfit-to-prefit parameter difference by.




Figure E.5: Transfer function polynomials extracted from fits to the 2015 data validation sample compared to the D_2^{DDT} -pass-to-fail ratios observed in the same data. The top two pads show the data, with different line colors for the spectrum of the D_2^{DDT} passing and failing selection, respectively. The topmost plot shows the two regions where the photon is forward, the middle plot shows the selection where the photon is central. In the bottom plot, the markers show the D_2^{DDT} pass-to-fail ratios of the two photon selections, with error bars indicating the expected uncertainty due to statistical fluctuation of per-bin poisson distributions. The light blue line indicates the transfer polynomial fit to the D_2^{DDT} -pass-fail ratio of the data less expected pre-fit resonant backgrounds in the forward-photon selection only (cf. Figure E.1), whereas the dark blue line shows the transfer polynomial that is implicitly fit by a maximisation of the full likelihood using 2015 Data. See Figure E.6 for the polynomial coefficients.




Figure E.6: Transfer polynomial coefficients extracted from fits to the 2015 data validation sample. See Figure E.5 for the evaluated polynomials.

E.2 Validation in data from the 2018 L4 period

This appendix collects additional visualizations of the background-only fit to the data observed in the 2018 L4 validation sample described in Section 8.1.2.




Figure E.7: Mass spectra observed in the 2018 L4 data compared to the pre-fit expectation. NPs without prior constraints are set to values that achieve an exact match to the observed distribution in the anti-tagged control regions, and transfer factor polynomial coefficients are chosen to minimise the squared error between the polynomial and the ratio of tagger-passing and -failing resonant-background-subtracted observed yields in the forward-photon control regions. NPs with prior constraints are set to their nominal values (maximising their probability according to the constraint term). Data from the SR is not used to choose NP values. The POI is fixed at 0. The hatched area indicates uncertainty on the expected yield, which is purely due to the width of the prior constraint terms of constrained NPs.




Figure E.8: Pulls of NPs with prior constraints (except MC statistics). See Figure E.9 for pulls of MC statistical uncertainty NPs.



Figure E.9: Pulls of MC statistical uncertainty NPs.




Figure E.10: Change in the floating shapefactor NPs relative to their initial value in the fit. The shapefactors are multiplicative unconstrained modifiers that are applied to unitary nominal yields. Their initial values are the observed yields in the D_2^{DDT} -failing control regions less the resonant backgrounds. In a consistent fit, and assuming the non-resonant background efficiency of D_2^{DDT} is treated correctly in the likelihood, the post-fit parameter values should vary according to the binwise Poisson error of the D_2^{DDT} -passing and D_2^{DDT} -failing observed yields, which are dominated by the much larger D_2^{DDT} -failing observed yields. This motivates choosing $\sqrt{\theta_{\text{initial}}}$ as the uncertainty to scale the postfit-to-prefit parameter difference by.

Figure E.11: Transfer polynomials extracted from fits to the 2018 L4 data sample compared to the D_2^{DDT} -pass-to-fail ratios observed in the same data. The top two pads show the data, with different line colors for the spectrum of the D_2^{DDT} passing and failing selection, respectively. The topmost plot shows the two regions where the photon is forward, the middle plot shows the selection where the photon is central. In the bottom plot, the markers show the D_2^{DDT} pass-to-fail ratios of the two photon selections, with error bars indicating the expected uncertainty due to statistical fluctuation of per-bin poisson distributions. The light blue line indicates the transfer polynomial fit to the D_2^{DDT} -pass-fail ratio of the data less expected pre-fit resonant backgrounds in the forward-photon selection only (cf. Figure E.7), whereas the dark blue line shows the transfer polynomial that is implicitly fit by a maximisation of the full likelihood using 2018 Data. See Figure E.12 for the polynomial coefficients.





Figure E.12: Transfer polynomial coefficients extracted from fits to the 2018 L4 data sample. See Figure E.11 for the evaluated polynomials.





Auxiliary material on unblinding 15.8 fb⁻¹ of observed data

F


F.1 Limited fit of the transfer factor polynomial to only the forward photon CRs

The data in the SR and CRs are compared prior to fitting the likelihood, to spot potential discrepancies before fully unblinding. The nominal expected resonant background templates are subtracted from the observed mass spectra in all regions. The D_2^{DDT} pass-to-fail ratio is computed independently in the central-photon and forward-photon selection. Figure F.1 illustrates the resulting ratios. They are compared to each other, and the result of fitting a Bernstein polynomial to the forward photon selection only, by minimizing the mean square error to the observed ratio. The pre-fit transfer factor matches the pass-to-fail ratio in the forward selection well, and the bin-wise observed ratios in both region are consistent with each other.


Figure F.1: Transfer function polynomials extracted from fits to the 15.8 fb⁻¹ sample after subtracting the nominal resonant background templates compared to the D_2^{DDT} -pass-to-fail ratios observed in the same data. The top two pads show the subtracted data, with different line colors for the spectrum of the D_2^{DDT} passing and failing selection, respectively. The topmost plot shows the two regions where the photon is forward, the middle plot shows the selection where the photon is central. In the bottom plot, the markers show the D_2^{DDT} pass-to-fail ratios of the two photon selections, with error bars indicating the expected uncertainty due to statistical fluctuation of per-bin Poisson distributions. The light blue line indicates the transfer polynomial fit to the D_2^{DDT} -pass-fail ratio of the data less expected pre-fit resonant backgrounds in the forward-photon selection only, whereas the dark blue line shows the transfer polynomial that is implicitly fit by a maximization of the full likelihood.




(a) SR



(b) Forward-photon tagged CR



(c) Central-photon anti-tagged CR



(d) Forward-photon anti-tagged CR

Figure F.2: Pre-fit mass spectra

F.1.1 Background-only maximum-likelihood fit in all regions

The full likelihood is maximized, while keeping the signal strength fixed at 0. This background-only fit matches the observed data well, as Figure F.3 shows. No nuisance parameters with prior constraints are pulled far from their prior-optimizing value, as Figure F.4 illustrates.




Figure F.3: Post-fit mass spectra of the background-only fit.




Figure F.4: Pulls of NPs with prior constraints in the background-only fit.

F.2 Example signal-extraction maximum likelihood fit

The pulls on constrained nuisance parameters are shown in Figure F.5. While there are some small differences due to the observed fluctuation preferring a narrower peak shape than the template provides, these shifts are small and the best-fit values are all compatible with those seen in the background-only fit.




Figure F.5: Pulls of NPs with prior constraints.

F.3 CLs limits




Figure F.6: Signal strength scans used to determine the upper end of the one-sided 95% CL_s interval. Here the $m_{Z'} = 50$ GeV signal model is shown as an example. A range of signal strength hypotheses is scanned to find the point where the adjusted p-value CL_s intersects the level α of the test.


The uncertainty on the resulting limits is illustrated by also providing the scan results for the background-only expected distribution shifted up-and-down by one and two standard deviations of the signal strength extraction. Figure F.6 illustrates this scan over signal strengths for the example of the $m_{Z'} = 50$ GeV model.

F.3.1 Comparison to ATLAS and CMS results


Figure F.7 shows the comparison between the results of the partial unblinding and the early-run two results previously published by the ATLAS [54] and CMS [53] collaborations. These references represent the latest results targeting the photon-associated boosted hadronic resonance production published by the respective collaborations. Despite analysing less than half the data of the published results, the limits from the partial unblinding shown here are competitive with these prior results. Furthermore, the interpretation down to 10 GeV in the hadronic resonance mass is the first time this part of the phase space has been constrained using ATLAS data.

[54]: Aaboud et al. (2019), *Search for light resonances decaying to boosted quark pairs and produced in association with a photon or a jet in proton-proton collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector*

[53]: Sirunyan et al. (2019), *Search for Low-Mass Quark-Antiquark Resonances Produced in Association with a Photon at $\sqrt{s} = 13$ TeV*



(a) Comparison to early run 2 ATLAS result



(b) Comparison to early run 2 CMS result

Figure F.7: Comparison of the 95% CL_s limits on Z' production to (a) ATLAS [54] and (b) CMS [53].

Auxiliary material on the unblinded analysis of 2015-2018 data

G




Figure G.1: Pass-to-fail ratios and transfer function polynomials observed in 2015-2018 data. The ratios are shown as they are observed in data, without subtracting any resonant backgrounds. Cf. Figure 8.5 where the resonant backgrounds are subtracted.

Figure G.2: 95 % CL_s exclusion limits on the Z' -quark coupling strength. These correspond to the same signal limits in terms of signal strength as Figure 8.8. Casting the limits in terms of the coupling parameter allows easy comparison to the summary plots maintained by the ATLAS and CMS collaborations. The CMS limit shown for comparison is the one published in ref. [53].





Figure G.3: Local p-values of the observed data given the background-only hypothesis.

Bibliography

Here are the references in citation order.

- [1] Roy Schimmel Brener. ‘Estimating the Sensitivity to New Resonances Decaying to Boosted Quark Pairs and Produced in Association with a Photon in the ATLAS Experiment.’ Presented 29 Dec 2020. 2020 (cit. on p. viii).
- [2] Lars Henkelmann. ‘Search for Dark Matter using the $h(b\bar{b}) + E_T^{\text{miss}}$ Signature of the New "Two-Higgs-Doublet Model with Pseudoscalar Mediator" Benchmark Model with the ATLAS Detector’. M.Sc. thesis. Universität Heidelberg, May 2018 (cit. on pp. viii, 3, 27).
- [3] Mark Thomson. *Modern particle physics*. New York: Cambridge University Press, 2013 (cit. on p. 3).
- [4] J. D. Vergados. *The Standard Model and Beyond*. Singapore: World Scientific, 2017 (cit. on p. 3).
- [5] Fermilab Particle Data Group PBS NOVA. *Standard Model Particle Content*. 2008. URL: https://en.wikipedia.org/wiki/File:Standard_Model_of_Elementary_Particles.svg (cit. on p. 4).
- [6] Eric Drexler. *Elementary Particle Interactions in the Standard Model*. 2014. URL: https://en.wikipedia.org/wiki/File:Elementary_particle_interactions_in_the_Standard_Model.png (cit. on p. 5).
- [7] Michael E. Peskin and Daniel V. Schroeder. *An Introduction to quantum field theory*. Reading, USA: Addison-Wesley, 1995 (cit. on p. 5).
- [8] C. Patrignani et al. ‘Review of Particle Physics’. In: *Chin. Phys.* C40.10 (2016), p. 100001. DOI: [10.1088/1674-1137/40/10/100001](https://doi.org/10.1088/1674-1137/40/10/100001) (cit. on pp. 6, 10, 13, 30).
- [9] John Ellis. *Higgs Physics*. Tech. rep. arXiv:1312.5672. KCL-PH-TH-2013-49. LCTS-2013-36. CERN-PH-TH-2013-315. Dec. 2013, 117–168. 52 p (cit. on p. 11).
- [10] D. M. Webber et al. ‘Measurement of the Positive Muon Lifetime and Determination of the Fermi Constant to Part-per-Million Precision’. In: *Phys. Rev. Lett.* 106 (2011). [Phys. Rev. Lett.106,079901(2011)], p. 041803. DOI: [10.1103/PhysRevLett.106.041803](https://doi.org/10.1103/PhysRevLett.106.041803), [10.1103/PhysRevLett.106.079901](https://doi.org/10.1103/PhysRevLett.106.079901) (cit. on p. 13).
- [11] R. L. Workman et al. ‘Review of Particle Physics’. In: *PTEP* 2022 (2022), p. 083C01. DOI: [10.1093/ptep/ptac097](https://doi.org/10.1093/ptep/ptac097) (cit. on pp. 13, 15, 76, 108).
- [12] Particle Data Group, P.A. Zyla et al. ‘Review of Particle Physics’. In: *PTEP* 2020.8 (2020), p. 083C01. DOI: [10.1093/ptep/ptaa104](https://doi.org/10.1093/ptep/ptaa104) (cit. on pp. 15, 21).
- [13] Laurent Canetti, Marco Drewes and Mikhail Shaposhnikov. ‘Matter and Antimatter in the Universe’. In: *New J. Phys.* 14 (2012), p. 095012. DOI: [10.1088/1367-2630/14/9/095012](https://doi.org/10.1088/1367-2630/14/9/095012) (cit. on p. 15).
- [14] Gianfranco Bertone, Dan Hooper and Joseph Silk. ‘Particle dark matter: Evidence, candidates and constraints’. In: *Phys. Rept.* 405 (2005), pp. 279–390. DOI: [10.1016/j.physrep.2004.08.031](https://doi.org/10.1016/j.physrep.2004.08.031) (cit. on pp. 15–18).
- [15] Lars Bergström. ‘Nonbaryonic dark matter: Observational evidence and detection methods’. In: *Rept. Prog. Phys.* 63 (2000), p. 793. DOI: [10.1088/0034-4885/63/5/2r3](https://doi.org/10.1088/0034-4885/63/5/2r3) (cit. on pp. 15–18).
- [16] Douglas Clowe et al. ‘A direct empirical proof of the existence of dark matter’. In: *Astrophys. J. Lett.* 648 (2006), pp. L109–L113. DOI: [10.1086/508162](https://doi.org/10.1086/508162) (cit. on pp. 15–17).

- [17] Stephen P. Martin. ‘A Supersymmetry primer’. In: *Adv. Ser. Direct. High Energy Phys.* 18 (1998). Ed. by Gordon L. Kane, pp. 1–98. doi: [10.1142/9789812839657_0001](https://doi.org/10.1142/9789812839657_0001) (cit. on p. 15).
- [18] Nicolás Bernal et al. ‘The Dawn of FIMP Dark Matter: A Review of Models and Constraints’. In: *Int. J. Mod. Phys. A* 32.27 (2017), p. 1730023. doi: [10.1142/S0217751X1730023X](https://doi.org/10.1142/S0217751X1730023X) (cit. on p. 15).
- [19] Paola Arias et al. ‘WISPy Cold Dark Matter’. In: *JCAP* 06 (2012), p. 013. doi: [10.1088/1475-7516/2012/06/013](https://doi.org/10.1088/1475-7516/2012/06/013) (cit. on p. 15).
- [20] Felix Kahlhoefer et al. ‘Implications of unitarity and gauge invariance for simplified dark matter models’. In: *JHEP* 02 (2016), p. 016. doi: [10.1007/JHEP02\(2016\)016](https://doi.org/10.1007/JHEP02(2016)016) (cit. on pp. 15, 21).
- [21] J. A. Aguilar-Saavedra. ‘Identifying top partners at LHC’. In: *JHEP* 11 (2009), p. 030. doi: [10.1088/1126-6708/2009/11/030](https://doi.org/10.1088/1126-6708/2009/11/030) (cit. on p. 15).
- [22] Gian F. Giudice et al. ‘Clockwork/linear dilaton: structure and phenomenology’. In: *JHEP* 06 (2018), p. 009. doi: [10.1007/JHEP06\(2018\)009](https://doi.org/10.1007/JHEP06(2018)009) (cit. on p. 15).
- [23] Matthew J. Strassler and Kathryn M. Zurek. ‘Echoes of a hidden valley at hadron colliders’. In: *Phys. Lett. B* 651 (2007), pp. 374–379. doi: [10.1016/j.physletb.2007.06.055](https://doi.org/10.1016/j.physletb.2007.06.055) (cit. on p. 15).
- [24] JiJi Fan, Matthew Reece and Joshua T. Ruderman. ‘Stealth Supersymmetry’. In: *JHEP* 11 (2011), p. 012. doi: [10.1007/JHEP11\(2011\)012](https://doi.org/10.1007/JHEP11(2011)012) (cit. on p. 15).
- [25] Diptimoy Ghosh et al. ‘Looking for an Invisible Higgs Signal at the LHC’. In: *Phys. Lett. B* 725 (2013), pp. 344–351. doi: [10.1016/j.physletb.2013.07.042](https://doi.org/10.1016/j.physletb.2013.07.042) (cit. on p. 15).
- [26] Partha Konar et al. ‘How to look for supersymmetry under the lamppost at the LHC’. In: *Phys. Rev. Lett.* 105 (2010), p. 221801. doi: [10.1103/PhysRevLett.105.221801](https://doi.org/10.1103/PhysRevLett.105.221801) (cit. on p. 15).
- [27] Andrés Flórez et al. ‘Anapole Dark Matter via Vector Boson Fusion Processes at the LHC’. In: *Phys. Rev. D* 100, 016017 (2019) (4th Feb. 2019). doi: [10.1103/PhysRevD.100.016017](https://doi.org/10.1103/PhysRevD.100.016017) (cit. on p. 15).
- [28] Jalal Abdallah et al. ‘Simplified Models for Dark Matter Searches at the LHC’. In: *Phys. Dark Univ.* 9–10 (2015), pp. 8–23. doi: [10.1016/j.dark.2015.08.001](https://doi.org/10.1016/j.dark.2015.08.001) (cit. on pp. 15–18, 21).
- [29] Malcolm Fairbairn et al. ‘Constraints on Z’ models from LHC dijet searches and implications for dark matter’. In: *JHEP* 09 (2016), p. 018. doi: [10.1007/JHEP09\(2016\)018](https://doi.org/10.1007/JHEP09(2016)018) (cit. on pp. 15, 19, 21).
- [30] Robert M. Harris and Konstantinos Kousouris. ‘Searches for Dijet Resonances at Hadron Colliders’. In: *Int. J. Mod. Phys. A* 26 (2011), pp. 5005–5055. doi: [10.1142/S0217751X11054905](https://doi.org/10.1142/S0217751X11054905) (cit. on pp. 15, 21).
- [31] Peter Schneider. *Extragalactic Astronomy and Cosmology*. 2nd ed. Springer-Verlag BerlinHeidelberg, 2015 (cit. on p. 16).
- [32] Daniel Abercrombie et al. ‘Dark Matter Benchmark Models for Early LHC Run-2 Searches: Report of the ATLAS/CMS Dark Matter Forum’. In: *Phys. Dark Univ.* 27 (2020). Ed. by Antonio Boveia et al., p. 100371. doi: [10.1016/j.dark.2019.100371](https://doi.org/10.1016/j.dark.2019.100371) (cit. on pp. 16, 18, 21, 50).
- [33] Tomohiro Abe et al. ‘LHC Dark Matter Working Group: Next-generation spin-0 dark matter models’. In: *Phys. Dark Univ.* 27 (2020), p. 100351. doi: [10.1016/j.dark.2019.100351](https://doi.org/10.1016/j.dark.2019.100351) (cit. on pp. 16, 18, 19).
- [34] Marco Taoso, Gianfranco Bertone and Antonio Masiero. ‘Dark Matter Candidates: A Ten-Point Test’. In: *JCAP* 0803 (2008), p. 022. doi: [10.1088/1475-7516/2008/03/022](https://doi.org/10.1088/1475-7516/2008/03/022) (cit. on p. 16).
- [35] Felix Kahlhoefer. ‘Review of LHC Dark Matter Searches’. In: *Int. J. Mod. Phys. A* 32.13 (2017), p. 1730006. doi: [10.1142/S0217751X1730006X](https://doi.org/10.1142/S0217751X1730006X) (cit. on pp. 16–19, 21).
- [36] Robert J. Scherrer and Michael S. Turner. ‘On the Relic, Cosmic Abundance of Stable Weakly Interacting Massive Particles’. In: *Phys. Rev. D* 33 (1986). [Erratum: *Phys. Rev. D* 34, 3263 (1986)], p. 1585. doi: [10.1103/PhysRevD.33.1585](https://doi.org/10.1103/PhysRevD.33.1585), [10.1103/PhysRevD.34.3263](https://doi.org/10.1103/PhysRevD.34.3263) (cit. on p. 16).

- [37] Planck Collaboration et al. ‘Planck 2013 results. I. Overview of products and scientific results’. In: *A&A* 571 (2014), A1. doi: [10.1051/0004-6361/201321529](https://doi.org/10.1051/0004-6361/201321529) (cit. on p. 16).
- [38] Teresa Marrodán Undagoitia and Ludwig Rauch. ‘Dark matter direct-detection experiments’. In: *J. Phys.* G43.1 (2016), p. 013001. doi: [10.1088/0954-3899/43/1/013001](https://doi.org/10.1088/0954-3899/43/1/013001) (cit. on pp. 17, 18).
- [39] O. Buchmueller, Matthew J. Dolan and Christopher McCabe. ‘Beyond Effective Field Theory for Dark Matter Searches at the LHC’. In: *JHEP* 01 (2014), p. 025. doi: [10.1007/JHEP01\(2014\)025](https://doi.org/10.1007/JHEP01(2014)025) (cit. on pp. 17, 21).
- [40] Gianfranco Bertone and David Merritt. ‘Dark matter dynamics and indirect detection’. In: *Mod. Phys. Lett.* A20 (2005), p. 1021. doi: [10.1142/S0217732305017391](https://doi.org/10.1142/S0217732305017391) (cit. on p. 18).
- [41] D.S. Akerib et al. ‘The Large Underground Xenon (LUX) Experiment’. In: *Nucl. Instrum. Meth. A* 704 (2013), pp. 111–126. doi: [10.1016/j.nima.2012.11.135](https://doi.org/10.1016/j.nima.2012.11.135) (cit. on p. 18).
- [42] E. Aprile et al. ‘Design and Performance of the XENON10 Dark Matter Experiment’. In: *Astropart. Phys.* 34 (2011), pp. 679–698. doi: [10.1016/j.astropartphys.2011.01.006](https://doi.org/10.1016/j.astropartphys.2011.01.006) (cit. on p. 18).
- [43] E. Armengaud et al. ‘Performance of the EDELWEISS-III experiment for direct dark matter searches’. In: *JINST* 12.08 (2017), P08010. doi: [10.1088/1748-0221/12/08/P08010](https://doi.org/10.1088/1748-0221/12/08/P08010) (cit. on p. 18).
- [44] Morad Aaboud et al. ‘Constraints on mediator-based dark matter and scalar dark energy models using $\sqrt{s} = 13$ TeV pp collision data collected by the ATLAS detector’. In: *JHEP* 05 (2019), p. 142. doi: [10.1007/JHEP05\(2019\)142](https://doi.org/10.1007/JHEP05(2019)142) (cit. on pp. 18, 19, 21).
- [45] Morad Aaboud et al. ‘Search for dark matter and other new phenomena in events with an energetic jet and large missing transverse momentum using the ATLAS detector’. In: *JHEP* 01 (2018), p. 126. doi: [10.1007/JHEP01\(2018\)126](https://doi.org/10.1007/JHEP01(2018)126) (cit. on p. 19).
- [46] A. M. Sirunyan et al. ‘Search for new physics in final states with an energetic jet or a hadronically decaying W or Z boson and transverse momentum imbalance at $\sqrt{s} = 13$ TeV’. In: *Phys. Rev. D* 97.9 (2018), p. 092005. doi: [10.1103/PhysRevD.97.092005](https://doi.org/10.1103/PhysRevD.97.092005) (cit. on p. 19).
- [47] Morad Aaboud et al. ‘Search for dark matter at $\sqrt{s} = 13$ TeV in final states containing an energetic photon and large missing transverse momentum with the ATLAS detector’. In: *Eur. Phys. J. C* 77.6 (2017), p. 393. doi: [10.1140/epjc/s10052-017-4965-8](https://doi.org/10.1140/epjc/s10052-017-4965-8) (cit. on p. 19).
- [48] Morad Aaboud et al. ‘Search for Dark Matter Produced in Association with a Higgs Boson Decaying to $b\bar{b}$ using 36 fb^{-1} of pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS Detector’. In: *Phys. Rev. Lett.* 119.18 (2017), p. 181804. doi: [10.1103/PhysRevLett.119.181804](https://doi.org/10.1103/PhysRevLett.119.181804) (cit. on p. 19).
- [49] Georges Aad et al. ‘Search for dark matter in events with a Z boson and missing transverse momentum in pp collisions at $\sqrt{s}=8$ TeV with the ATLAS detector’. In: *Phys. Rev. D* 90.1 (2014), p. 012004. doi: [10.1103/PhysRevD.90.012004](https://doi.org/10.1103/PhysRevD.90.012004) (cit. on p. 19).
- [50] M. Aaboud et al. ‘Search for dark matter in events with a hadronically decaying vector boson and missing transverse momentum in pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector’. In: *JHEP* 10 (2018), p. 180. doi: [10.1007/JHEP10\(2018\)180](https://doi.org/10.1007/JHEP10(2018)180) (cit. on p. 19).
- [51] *Dark matter summary plots for s-channel, 2HDM+a, Higgs portal and Dark Higgs models*. Tech. rep. All figures including auxiliary figures are available at <https://atlas.web.cern.ch/Atlas/GROUPS/PHYSICS/PUBNOTES/ATL-PHYS-PUB-2023-018>. Geneva: CERN, 2023 (cit. on pp. 19–22).
- [52] Bogdan A. Dobrescu and Claudia Frugieule. ‘Hidden GeV-scale interactions of quarks’. In: *Phys. Rev. Lett.* 113 (2014), p. 061801. doi: [10.1103/PhysRevLett.113.061801](https://doi.org/10.1103/PhysRevLett.113.061801) (cit. on p. 21).

- [53] Albert M Sirunyan et al. ‘Search for Low-Mass Quark-Antiquark Resonances Produced in Association with a Photon at $\sqrt{s} = 13$ TeV’. In: *Phys. Rev. Lett.* 123.23 (2019), p. 231803. doi: [10.1103/PhysRevLett.123.231803](#) (cit. on pp. 21–23, 43, 62, 68, 123, 162, 163, 166).
- [54] Morad Aaboud et al. ‘Search for light resonances decaying to boosted quark pairs and produced in association with a photon or a jet in proton-proton collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector’. In: *Phys. Lett. B* 788 (2019), pp. 316–335. doi: [10.1016/j.physletb.2018.09.062](#) (cit. on pp. 22, 23, 162, 163).
- [55] Georges Aad et al. ‘Search for new resonances in mass distributions of jet pairs using 139 fb^{-1} of pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector’. In: *JHEP* 03 (2020), p. 145. doi: [10.1007/JHEP03\(2020\)145](#) (cit. on p. 22).
- [56] Benjamin Nachman and Francesco Rubbo. ‘Search strategy using LHC pileup interactions as a zero bias sample’. In: *Phys. Rev. D* 97.9 (2018), p. 092002. doi: [10.1103/PhysRevD.97.092002](#) (cit. on p. 22).
- [57] R. Aaij et al. ‘Tesla : an application for real-time data analysis in High Energy Physics’. In: *Comput. Phys. Commun.* 208 (2016), pp. 35–42. doi: [10.1016/j.cpc.2016.07.022](#) (cit. on p. 22).
- [58] Vardan Khachatryan et al. ‘Search for narrow resonances in dijet final states at $\sqrt{s} = 8$ TeV with the novel CMS technique of data scouting’. In: *Phys. Rev. Lett.* 117.3 (2016), p. 031802. doi: [10.1103/PhysRevLett.117.031802](#) (cit. on pp. 22, 23).
- [59] M. Aaboud et al. ‘Search for low-mass dijet resonances using trigger-level jets with the ATLAS detector in pp collisions at $\sqrt{s} = 13$ TeV’. In: *Phys. Rev. Lett.* 121.8 (2018), p. 081801. doi: [10.1103/PhysRevLett.121.081801](#) (cit. on pp. 22, 23, 123).
- [60] Albert M Sirunyan et al. ‘Search for low mass vector resonances decaying into quark-antiquark pairs in proton-proton collisions at $\sqrt{s} = 13$ TeV’. In: *Phys. Rev. D* 100.11 (2019), p. 112007. doi: [10.1103/PhysRevD.100.112007](#) (cit. on pp. 22, 23, 62, 124).
- [61] Gavin P. Salam. ‘Towards Jetography’. In: *Eur. Phys. J. C* 67 (2010), pp. 637–686. doi: [10.1140/epjc/s10052-010-1314-6](#) (cit. on pp. 22, 51, 52, 62, 107).
- [62] Matteo Cacciari, Gavin P. Salam and Gregory Soyez. ‘The anti- k_t jet clustering algorithm’. In: *JHEP* 04 (2008), p. 063. doi: [10.1088/1126-6708/2008/04/063](#) (cit. on pp. 23, 51, 52, 74).
- [63] Yuri L. Dokshitzer et al. ‘Better jet clustering algorithms’. In: *JHEP* 08 (1997), p. 001. doi: [10.1088/1126-6708/1997/08/001](#) (cit. on pp. 23, 51).
- [64] Matteo Cacciari and Gavin P. Salam. ‘Pileup subtraction using jet areas’. In: *Phys. Lett. B* 659 (2008), pp. 119–126. doi: [10.1016/j.physletb.2007.09.077](#) (cit. on pp. 23, 52).
- [65] Georges Aad et al. ‘Performance of pile-up mitigation techniques for jets in pp collisions at $\sqrt{s} = 8$ TeV using the ATLAS detector’. In: *Eur. Phys. J. C* 76.11 (2016), p. 581. doi: [10.1140/epjc/s10052-016-4395-z](#) (cit. on p. 23).
- [66] David Krohn, Jesse Thaler and Lian-Tao Wang. ‘Jet Trimming’. In: *JHEP* 02 (2010), p. 084. doi: [10.1007/JHEP02\(2010\)084](#) (cit. on pp. 23, 51, 53, 125).
- [67] Andrew J. Larkoski et al. ‘Soft Drop’. In: *JHEP* 05 (2014), p. 146. doi: [10.1007/JHEP05\(2014\)146](#) (cit. on pp. 23, 125).
- [68] Mrinal Dasgupta et al. ‘Towards an understanding of jet substructure’. In: *JHEP* 09 (2013), p. 029. doi: [10.1007/JHEP09\(2013\)029](#) (cit. on pp. 23, 24, 107, 125).
- [69] Daniele Bertolini et al. ‘Pileup Per Particle Identification’. In: *JHEP* 10 (2014), p. 059. doi: [10.1007/JHEP10\(2014\)059](#) (cit. on p. 23).

- [70] Jesse Thaler and Ken Van Tilburg. ‘Identifying Boosted Objects with N-subjettiness’. In: *JHEP* 03 (2011), p. 015. doi: [10.1007/JHEP03\(2011\)015](https://doi.org/10.1007/JHEP03(2011)015) (cit. on pp. 24, 60, 67).
- [71] Andrew J. Larkoski, Gavin P. Salam and Jesse Thaler. ‘Energy Correlation Functions for Jet Substructure’. In: *JHEP* 06 (2013), p. 108. doi: [10.1007/JHEP06\(2013\)108](https://doi.org/10.1007/JHEP06(2013)108) (cit. on pp. 24, 60, 67).
- [72] Ian Moult, Lina Necib and Jesse Thaler. ‘New Angles on Energy Correlation Functions’. In: *JHEP* 12 (2016), p. 153. doi: [10.1007/JHEP12\(2016\)153](https://doi.org/10.1007/JHEP12(2016)153) (cit. on pp. 24, 60, 61, 67).
- [73] James Dolen et al. ‘Thinking outside the ROCs: Designing Decorrelated Taggers (DDT) for jet substructure’. In: *JHEP* 05 (2016), p. 156. doi: [10.1007/JHEP05\(2016\)156](https://doi.org/10.1007/JHEP05(2016)156) (cit. on pp. 24, 60).
- [74] Lyndon Evans and Philip Bryant. ‘LHC Machine’. In: *JINST* 3 (2008), S08001. doi: [10.1088/1748-0221/3/08/S08001](https://doi.org/10.1088/1748-0221/3/08/S08001) (cit. on p. 27).
- [75] Xabier Cid Vidal and Ramon Cid Mandano. *Taking a closer look at the LHC*. 2018. URL: https://www.lhc-closer.es/taking_a_closer_look_at_lhc (cit. on p. 27).
- [76] G. Aad et al. ‘The ATLAS Experiment at the CERN Large Hadron Collider’. In: *JINST* 3 (2008), S08003. doi: [10.1088/1748-0221/3/08/S08003](https://doi.org/10.1088/1748-0221/3/08/S08003) (cit. on pp. 27, 28, 30, 35).
- [77] Felix Müller. ‘Jet Production Measurements at the ATLAS Experiment’. PhD thesis. Heidelberg University, 2013 (cit. on p. 28).
- [78] K. Aamodt et al. ‘The ALICE experiment at the CERN LHC’. In: *JINST* 3 (2008), S08002. doi: [10.1088/1748-0221/3/08/S08002](https://doi.org/10.1088/1748-0221/3/08/S08002) (cit. on p. 28).
- [79] S. Chatrchyan et al. ‘The CMS Experiment at the CERN LHC’. In: *JINST* 3 (2008), S08004. doi: [10.1088/1748-0221/3/08/S08004](https://doi.org/10.1088/1748-0221/3/08/S08004) (cit. on p. 28).
- [80] A. Augusto Alves Jr. et al. ‘The LHCb Detector at the LHC’. In: *JINST* 3 (2008), S08005. doi: [10.1088/1748-0221/3/08/S08005](https://doi.org/10.1088/1748-0221/3/08/S08005) (cit. on p. 28).
- [81] J. Pequeano CERN PhotoLab. *Computer generated image of the whole ATLAS detector*. 2008. URL: <https://cds.cern.ch/record/1095924> (cit. on p. 29).
- [82] G. Aad et al. ‘The ATLAS Inner Detector commissioning and calibration’. In: *Eur. Phys. J. C* 70 (2010), pp. 787–821. doi: [10.1140/epjc/s10052-010-1366-7](https://doi.org/10.1140/epjc/s10052-010-1366-7) (cit. on p. 31).
- [83] J. Pequeano and CERN PhotoLab. *Computer generated image of the ATLAS calorimeter*. 2018. URL: <https://cds.cern.ch/record/1095927> (cit. on p. 33).
- [84] ATLAS. *Improving jet substructure performance in ATLAS using Track-CaloClusters*. Tech. rep. ATL-PHYS-PUB-2017-015. Geneva: CERN, July 2017 (cit. on p. 33).
- [85] ATLAS. *Track assisted techniques for jet substructure*. Tech. rep. ATL-PHYS-PUB-2018-012. 2018 (cit. on pp. 33, 53).
- [86] V. Cindro et al. ‘The ATLAS beam conditions monitor’. In: *JINST* 3 (2008), P02004. doi: [10.1088/1748-0221/3/02/P02004](https://doi.org/10.1088/1748-0221/3/02/P02004) (cit. on p. 38).
- [87] S. van der Meer. *Calibration of the Effective Beam Height in the ISR*. Tech. rep. CERN, 1968 (cit. on p. 39).
- [88] H Burkhardt and P Grafström. *Absolute Luminosity from Machine Parameters*. Tech. rep. LHC-PROJECT-Report-1019. CERN-LHC-PROJECT-Report-1019. CERN, 2007 (cit. on p. 39).
- [89] A. K. Barlow et al. ‘Instrumentation and beam diagnostics in the ISR’. In: *Proceedings of the 8th International Conference on High- Energy Accelerators*. Ed. by M. Hildred Blewett. CERN. Jan. 1971, pp 426–30 (cit. on p. 39).

- [90] ATLAS. *2015 start-up trigger menu and initial performance assessment of the ATLAS trigger using Run-2 data*. Tech. rep. All figures including auxiliary figures are available at <https://atlas.web.cern.ch/Atlas/GROUPS/PHYSICS/PUBS/DAQ-PUB-2016-001>. Geneva: CERN, 2016 (cit. on p. 45).
- [91] Georges Aad et al. ‘Performance of the ATLAS muon triggers in Run 2’. In: *JINST* 15.09 (2020), P09015. doi: [10.1088/1748-0221/15/09/p09015](https://doi.org/10.1088/1748-0221/15/09/p09015) (cit. on p. 45).
- [92] ATLAS Collaboration. ‘Luminosity determination in pp collisions at $\sqrt{s} = 13$ TeV using the ATLAS detector at the LHC’. In: (2022). doi: [10.48550/ARXIV.2212.09379](https://doi.org/10.48550/ARXIV.2212.09379) (cit. on pp. 45, 87, 114).
- [93] GEANT4 Collaboration, S. Agostinelli et al. ‘GEANT4 – a simulation toolkit’. In: *Nucl. Instrum. Meth. A* 506 (2003), p. 250. doi: [10.1016/S0168-9002\(03\)01368-8](https://doi.org/10.1016/S0168-9002(03)01368-8) (cit. on p. 46).
- [94] ATLAS Collaboration. ‘The ATLAS Simulation Infrastructure’. In: *Eur. Phys. J. C* 70 (2010), p. 823. doi: [10.1140/epjc/s10052-010-1429-9](https://doi.org/10.1140/epjc/s10052-010-1429-9) (cit. on p. 46).
- [95] Bo Andersson et al. ‘Parton fragmentation and string dynamics’. In: *Phys. Rept.* 97 (1983), pp. 31–145. doi: [10.1016/0370-1573\(83\)90080-7](https://doi.org/10.1016/0370-1573(83)90080-7) (cit. on p. 46).
- [96] Torbjorn Sjöstrand. ‘Jet fragmentation of multiparton configurations in a string framework’. In: *Nucl. Phys. B* 248 (1984), pp. 469–502. doi: [10.1016/0550-3213\(84\)90607-2](https://doi.org/10.1016/0550-3213(84)90607-2) (cit. on p. 46).
- [97] Torbjörn Sjöstrand et al. ‘An introduction to PYTHIA 8.2’. In: *Comput. Phys. Commun.* 191 (2015), p. 159. doi: [10.1016/j.cpc.2015.01.024](https://doi.org/10.1016/j.cpc.2015.01.024) (cit. on pp. 46, 49, 50).
- [98] ATLAS Collaboration. *ATLAS Pythia 8 tunes to 7 TeV data*. ATL-PHYS-PUB-2014-021. 2014 (cit. on pp. 46, 48).
- [99] Richard D. Ball et al. ‘Parton distributions with LHC data’. In: *Nucl. Phys. B* 867 (2013), pp. 244–289. doi: [10.1016/j.nuclphysb.2012.10.003](https://doi.org/10.1016/j.nuclphysb.2012.10.003) (cit. on pp. 46, 47).
- [100] D. J. Lange. ‘The EvtGen particle decay simulation package’. In: *Nucl. Instrum. Meth. A* 462 (2001), p. 152. doi: [10.1016/S0168-9002\(01\)00089-4](https://doi.org/10.1016/S0168-9002(01)00089-4) (cit. on p. 46).
- [101] T. Sjöstrand, S. Mrenna and P. Skands. ‘A brief introduction to PYTHIA 8.1’. In: *Comput. Phys. Commun.* 178 (2008), pp. 852–867. doi: [10.1016/j.cpc.2008.01.036](https://doi.org/10.1016/j.cpc.2008.01.036) (cit. on p. 46).
- [102] ATLAS Collaboration. *The Pythia 8 A3 tune description of ATLAS minimum bias and inelastic measurements incorporating the Donnachie–Landshoff diffractive model*. ATL-PHYS-PUB-2016-017. 2016 (cit. on p. 46).
- [103] Jon Butterworth et al. ‘PDF4LHC recommendations for LHC Run II’. In: *J. Phys. G* 43 (2016), p. 023001. doi: [10.1088/0954-3899/43/2/023001](https://doi.org/10.1088/0954-3899/43/2/023001) (cit. on p. 46).
- [104] Andy Buckley et al. ‘LHAPDF6: parton density access in the LHC precision era’. In: *Eur. Phys. J. C* 75 (2015), p. 132. doi: [10.1140/epjc/s10052-015-3318-8](https://doi.org/10.1140/epjc/s10052-015-3318-8) (cit. on pp. 46, 82).
- [105] Enrico Bothmann et al. ‘Event generation with Sherpa 2.2’. In: *SciPost Phys.* 7.3 (2019), p. 034. doi: [10.21468/SciPostPhys.7.3.034](https://doi.org/10.21468/SciPostPhys.7.3.034) (cit. on pp. 47, 49).
- [106] Tanju Gleisberg and Stefan Höche. ‘Comix, a new matrix element generator’. In: *JHEP* 12 (2008), p. 039. doi: [10.1088/1126-6708/2008/12/039](https://doi.org/10.1088/1126-6708/2008/12/039) (cit. on p. 47).
- [107] Federico Buccioni et al. ‘OpenLoops 2’. In: *Eur. Phys. J. C* 79.10 (2019), p. 866. doi: [10.1140/epjc/s10052-019-7306-2](https://doi.org/10.1140/epjc/s10052-019-7306-2) (cit. on p. 47).
- [108] Fabio Cascioli, Philipp Maierhöfer and Stefano Pozzorini. ‘Scattering Amplitudes with Open Loops’. In: *Phys. Rev. Lett.* 108 (2012), p. 111601. doi: [10.1103/PhysRevLett.108.111601](https://doi.org/10.1103/PhysRevLett.108.111601) (cit. on p. 47).
- [109] Federico Buccioni, Stefano Pozzorini and Max Zoller. ‘On-the-fly reduction of open loops’. In: *Eur. Phys. J. C* 78.1 (2018), p. 70. doi: [10.1140/epjc/s10052-018-5562-1](https://doi.org/10.1140/epjc/s10052-018-5562-1) (cit. on p. 47).

- [110] Ansgar Denner, Stefan Dittmaier and Lars Hofer. ‘COLLIER: A fortran-based complex one-loop library in extended regularizations’. In: *Comput. Phys. Commun.* 212 (2017), pp. 220–238. doi: [10.1016/j.cpc.2016.10.013](#) (cit. on p. 47).
- [111] Steffen Schumann and Frank Krauss. ‘A parton shower algorithm based on Catani–Seymour dipole factorisation’. In: *JHEP* 03 (2008), p. 038. doi: [10.1088/1126-6708/2008/03/038](#) (cit. on p. 47).
- [112] Stefan Höche et al. ‘A critical appraisal of NLO+PS matching methods’. In: *JHEP* 09 (2012), p. 049. doi: [10.1007/JHEP09\(2012\)049](#) (cit. on p. 47).
- [113] Stefan Höche et al. ‘QCD matrix elements + parton showers. The NLO case’. In: *JHEP* 04 (2013), p. 027. doi: [10.1007/JHEP04\(2013\)027](#) (cit. on p. 47).
- [114] S. Catani et al. ‘QCD Matrix Elements + Parton Showers’. In: *JHEP* 11 (2001), p. 063. doi: [10.1088/1126-6708/2001/11/063](#) (cit. on p. 47).
- [115] Stefan Höche et al. ‘QCD matrix elements and truncated showers’. In: *JHEP* 05 (2009), p. 053. doi: [10.1088/1126-6708/2009/05/053](#) (cit. on pp. 47, 49).
- [116] Richard D. Ball et al. ‘Parton distributions for the LHC run II’. In: *JHEP* 04 (2015), p. 040. doi: [10.1007/JHEP04\(2015\)040](#) (cit. on pp. 47, 48, 50).
- [117] Stefano Carrazza et al. ‘An Unbiased Hessian Representation for Monte Carlo PDFs’. In: *Eur. Phys. J. C* 75.8 (2015), p. 369. doi: [10.1140/epjc/s10052-015-3590-7](#) (cit. on p. 47).
- [118] Richard D. Ball et al. ‘Parton distributions from high-precision collider data’. In: *Eur. Phys. J. C* 77.10 (2017), p. 663. doi: [10.1140/epjc/s10052-017-5199-5](#) (cit. on p. 47).
- [119] S. Bailey et al. ‘Parton distributions from LHC, HERA, Tevatron and fixed target data: MSHT20 PDFs’. In: *Eur. Phys. J. C* 81.4 (2021), p. 341. doi: [10.1140/epjc/s10052-021-09057-0](#) (cit. on p. 47).
- [120] Tie-Jiun Hou et al. ‘New CTEQ global analysis of quantum chromodynamics with high-precision data from the LHC’. In: *Phys. Rev. D* 103.1 (2021), p. 014013. doi: [10.1103/PhysRevD.103.014013](#) (cit. on p. 47).
- [121] Sayipjamal Dulat et al. ‘New parton distribution functions from a global analysis of quantum chromodynamics’. In: *Phys. Rev. D* 93.3 (2016), p. 033006. doi: [10.1103/PhysRevD.93.033006](#) (cit. on pp. 47, 48, 50).
- [122] L. A. Harland-Lang et al. ‘Parton distributions in the LHC era: MMHT 2014 PDFs’. In: *Eur. Phys. J. C* 75.5 (2015), p. 204. doi: [10.1140/epjc/s10052-015-3397-6](#) (cit. on pp. 47, 48, 50).
- [123] Stefano Frixione, Giovanni Ridolfi and Paolo Nason. ‘A positive-weight next-to-leading-order Monte Carlo for heavy flavour hadroproduction’. In: *JHEP* 09 (2007), p. 126. doi: [10.1088/1126-6708/2007/09/126](#) (cit. on p. 47).
- [124] Paolo Nason. ‘A new method for combining NLO QCD with shower Monte Carlo algorithms’. In: *JHEP* 11 (2004), p. 040. doi: [10.1088/1126-6708/2004/11/040](#) (cit. on pp. 47, 48).
- [125] Stefano Frixione, Paolo Nason and Carlo Oleari. ‘Matching NLO QCD computations with parton shower simulations: the POWHEG method’. In: *JHEP* 11 (2007), p. 070. doi: [10.1088/1126-6708/2007/11/070](#) (cit. on pp. 47, 48).
- [126] Simone Alioli et al. ‘A general framework for implementing NLO calculations in shower Monte Carlo programs: the POWHEG BOX’. In: *JHEP* 06 (2010), p. 043. doi: [10.1007/JHEP06\(2010\)043](#) (cit. on pp. 47, 48).

- [127] F. Febres Cordero, M. Kraus and L. Reina. ‘Top-quark pair production in association with a $W^\pm$ gauge boson in the POWHEG-BOX’. In: *Phys. Rev. D* 103.9 (2021), p. 094014. doi: [10.1103/PhysRevD.103.094014](#) (cit. on p. 47).
- [128] M. Beneke et al. ‘Hadronic top-quark pair production with NNLL threshold resummation’. In: *Nucl. Phys. B* 855 (2012), pp. 695–741. doi: [10.1016/j.nuclphysb.2011.10.021](#) (cit. on p. 47).
- [129] Matteo Cacciari et al. ‘Top-pair production at hadron colliders with next-to-next-to-leading logarithmic soft-gluon resummation’. In: *Phys. Lett. B* 710 (2012), pp. 612–622. doi: [10.1016/j.physletb.2012.03.013](#) (cit. on p. 47).
- [130] Peter Bärnreuther, Michal Czakon and Alexander Mitov. ‘Percent-Level-Precision Physics at the Tevatron: Next-to-Next-to-Leading Order QCD Corrections to $q\bar{q} \rightarrow t\bar{t} + X$ ’. In: *Phys. Rev. Lett.* 109 (2012), p. 132001. doi: [10.1103/PhysRevLett.109.132001](#) (cit. on p. 47).
- [131] Michal Czakon and Alexander Mitov. ‘NNLO corrections to top-pair production at hadron colliders: the all-fermionic scattering channels’. In: *JHEP* 12 (2012), p. 054. doi: [10.1007/JHEP12\(2012\)054](#) (cit. on p. 47).
- [132] Michal Czakon and Alexander Mitov. ‘NNLO corrections to top pair production at hadron colliders: the quark-gluon reaction’. In: *JHEP* 01 (2013), p. 080. doi: [10.1007/JHEP01\(2013\)080](#) (cit. on p. 47).
- [133] Michal Czakon, Paul Fiedler and Alexander Mitov. ‘Total Top-Quark Pair-Production Cross Section at Hadron Colliders Through $O(\alpha_s^4)$ ’. In: *Phys. Rev. Lett.* 110 (2013), p. 252004. doi: [10.1103/PhysRevLett.110.252004](#) (cit. on p. 47).
- [134] Michal Czakon and Alexander Mitov. ‘Top++: A program for the calculation of the top-pair cross-section at hadron colliders’. In: *Comput. Phys. Commun.* 185 (2014), p. 2930. doi: [10.1016/j.cpc.2014.06.021](#) (cit. on p. 47).
- [135] A. D. Martin et al. ‘Parton distributions for the LHC’. In: *Eur. Phys. J. C* 63 (2009), p. 189. doi: [10.1140/epjc/s10052-009-1072-5](#) (cit. on p. 47).
- [136] A. D. Martin et al. ‘Uncertainties on α_s in global PDF analyses and implications for predicted hadronic cross sections’. In: *Eur. Phys. J. C* 64 (2009), pp. 653–680. doi: [10.1140/epjc/s10052-009-1164-2](#) (cit. on p. 47).
- [137] H.-L. Lai et al. ‘New parton distributions for collider physics’. In: *Phys. Rev. D* 82 (2010), p. 074024. doi: [10.1103/PhysRevD.82.074024](#) (cit. on p. 47, 49).
- [138] J. Gao et al. ‘CT10 next-to-next-to-leading order global analysis of QCD’. In: *Phys. Rev. D* 89 (2014), p. 033009. doi: [10.1103/PhysRevD.89.033009](#) (cit. on p. 47).
- [139] ATLAS Collaboration. *Studies on top-quark Monte Carlo modelling with Sherpa and MG5_aMC@NLO*. ATL-PHYS-PUB-2017-007. 2017 (cit. on p. 48).
- [140] Simone Alioli et al. ‘NLO single-top production matched with shower in POWHEG: s- and t-channel contributions’. In: *JHEP* 09 (2009), p. 111. doi: [10.1088/1126-6708/2009/09/111](#) (cit. on p. 48). Erratum: in: *JHEP* 02 (2010), p. 011. doi: [10.1007/JHEP02\(2010\)011](#).
- [141] Rikkert Frederix, Emanuele Re and Paolo Torrielli. ‘Single-top t-channel hadroproduction in the four-flavour scheme with POWHEG and aMC@NLO’. In: *JHEP* 09 (2012), p. 130. doi: [10.1007/JHEP09\(2012\)130](#) (cit. on p. 48).
- [142] Emanuele Re. ‘Single-top Wt -channel production matched with parton showers using the POWHEG method’. In: *Eur. Phys. J. C* 71 (2011), p. 1547. doi: [10.1140/epjc/s10052-011-1547-z](#) (cit. on p. 48).

- [143] Stefano Frixione et al. ‘Single-top hadroproduction in association with a W boson’. In: *JHEP* 07 (2008), p. 029. doi: [10.1088/1126-6708/2008/07/029](https://doi.org/10.1088/1126-6708/2008/07/029) (cit. on p. 48).
- [144] ATLAS Collaboration. *Studies on top-quark Monte Carlo modelling for Top2016*. ATL-PHYS-PUB-2016-020. 2016 (cit. on p. 48).
- [145] Stefano Frixione et al. ‘Angular correlations of lepton pairs from vector boson and top quark decays in Monte Carlo simulations’. In: *JHEP* 04 (2007), p. 081. doi: [10.1088/1126-6708/2007/04/081](https://doi.org/10.1088/1126-6708/2007/04/081) (cit. on p. 48).
- [146] Pierre Artoisenet et al. ‘Automatic spin-entangled decays of heavy resonances in Monte Carlo simulations’. In: *JHEP* 03 (2013), p. 015. doi: [10.1007/JHEP03\(2013\)015](https://doi.org/10.1007/JHEP03(2013)015) (cit. on p. 48).
- [147] M. Aliev et al. ‘HATHOR – HAdronic Top and Heavy quarks crOss section calculator’. In: *Comput. Phys. Commun.* 182 (2011), pp. 1034–1046. doi: [10.1016/j.cpc.2010.12.040](https://doi.org/10.1016/j.cpc.2010.12.040) (cit. on p. 48).
- [148] P. Kant et al. ‘HatHor for single top-quark production: Updated predictions and uncertainty estimates for single top-quark production in hadronic collisions’. In: *Comput. Phys. Commun.* 191 (2015), pp. 74–89. doi: [10.1016/j.cpc.2015.02.001](https://doi.org/10.1016/j.cpc.2015.02.001) (cit. on p. 48).
- [149] Nikolaos Kidonakis. ‘Two-loop soft anomalous dimensions for single top quark associated production with a W^- or H^- ’. In: *Phys. Rev. D* 82 (2010), p. 054018. doi: [10.1103/PhysRevD.82.054018](https://doi.org/10.1103/PhysRevD.82.054018) (cit. on p. 48).
- [150] Nikolaos Kidonakis. ‘Top Quark Production’. In: *Proceedings, Helmholtz International Summer School on Physics of Heavy Quarks and Hadrons (HQ 2013)* (JINR, Dubna, Russia, 15th–28th July 2013), pp. 139–168. doi: [10.3204/DESY-PROC-2013-03/Kidonakis](https://doi.org/10.3204/DESY-PROC-2013-03/Kidonakis) (cit. on p. 48).
- [151] 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: *JHEP* 07 (2014), p. 079. doi: [10.1007/JHEP07\(2014\)079](https://doi.org/10.1007/JHEP07(2014)079) (cit. on pp. 48, 50).
- [152] Leif Lönnblad. ‘Correcting the Colour-Dipole Cascade Model with Fixed Order Matrix Elements’. In: *JHEP* 05 (2002), p. 046. doi: [10.1088/1126-6708/2002/05/046](https://doi.org/10.1088/1126-6708/2002/05/046) (cit. on p. 48).
- [153] Leif Lönnblad and Stefan Prestel. ‘Matching tree-level matrix elements with interleaved showers’. In: *JHEP* 03 (2012), p. 019. doi: [10.1007/JHEP03\(2012\)019](https://doi.org/10.1007/JHEP03(2012)019) (cit. on p. 48).
- [154] Jan-Christopher Winter, Frank Krauss and Gerhard Soff. ‘A modified cluster-hadronization model’. In: *Eur. Phys. J. C* 36 (2004), pp. 381–395. doi: [10.1140/epjc/s2004-01960-8](https://doi.org/10.1140/epjc/s2004-01960-8) (cit. on p. 49).
- [155] Stefano Frixione. ‘Isolated photons in perturbative QCD’. In: *Phys. Lett. B* 429 (1998), pp. 369–374. doi: [10.1016/S0370-2693\(98\)00454-7](https://doi.org/10.1016/S0370-2693(98)00454-7) (cit. on p. 49).
- [156] Bogdan A. Dobrescu and Felix Yu. ‘Coupling-Mass Mapping of Dijet Peak Searches’. In: *Phys. Rev. D* 88.3 (2013). [Erratum: *Phys.Rev.D* 90, 079901 (2014)], p. 035021. doi: [10.1103/PhysRevD.88.035021](https://doi.org/10.1103/PhysRevD.88.035021) (cit. on p. 50).
- [157] Mihailo Backović et al. ‘Higher-order QCD predictions for dark matter production at the LHC in simplified models with s-channel mediators’. In: *Eur. Phys. J. C* 75.10 (2015), p. 482. doi: [10.1140/epjc/s10052-015-3700-6](https://doi.org/10.1140/epjc/s10052-015-3700-6) (cit. on p. 50).
- [158] S. Mrenna and P. Skands. ‘Automated parton-shower variations in PYTHIA 8’. In: *Phys. Rev. D* 94 (2016), p. 074005. doi: [10.1103/PhysRevD.94.074005](https://doi.org/10.1103/PhysRevD.94.074005) (cit. on p. 50).
- [159] ATLAS Collaboration. ‘Measurement of the photon identification efficiencies with the ATLAS detector using LHC Run 2 data collected in 2015 and 2016’. In: *Eur. Phys. J. C* 79 (2019), p. 205. doi: [10.1140/epjc/s10052-019-6650-6](https://doi.org/10.1140/epjc/s10052-019-6650-6) (cit. on pp. 50, 51).

- [160] Morad Aaboud et al. ‘Measurement of the photon identification efficiencies with the ATLAS detector using LHC Run-1 data’. In: *Eur. Phys. J. C* 76.12 (2016), p. 666. doi: [10.1140/epjc/s10052-016-4507-9](https://doi.org/10.1140/epjc/s10052-016-4507-9) (cit. on p. 50).
- [161] ATLAS Collaboration. ‘Electron and photon performance measurements with the ATLAS detector using the 2015–2017 LHC proton-proton collision data’. In: *JINST* 14.12 (2019), P12006. doi: [10.1088/1748-0221/14/12/P12006](https://doi.org/10.1088/1748-0221/14/12/P12006) (cit. on p. 51).
- [162] S. Brandt et al. ‘The principal axis of jets — an attempt to analyse high-energy collisions as two-body processes’. In: *Physics Letters* 12.1 (1964), pp. 57–61. doi: [https://doi.org/10.1016/0031-9163\(64\)91176-X](https://doi.org/10.1016/0031-9163(64)91176-X) (cit. on p. 51).
- [163] Huilin Qu and Loukas Gouskos. ‘ParticleNet: Jet Tagging via Particle Clouds’. In: *Phys. Rev. D* 101.5 (2020), p. 056019. doi: [10.1103/PhysRevD.101.056019](https://doi.org/10.1103/PhysRevD.101.056019) (cit. on pp. 51, 125).
- [164] S. Catani et al. ‘Longitudinally invariant K_t clustering algorithms for hadron hadron collisions’. In: *Nucl. Phys. B* 406 (1993), pp. 187–224. doi: [10.1016/0550-3213\(93\)90166-M](https://doi.org/10.1016/0550-3213(93)90166-M) (cit. on p. 51).
- [165] Matteo Cacciari and Gavin P. Salam. ‘Dispelling the N^3 myth for the k_t jet-finder’. In: *Phys. Lett. B* 641 (2006), pp. 57–61. doi: [10.1016/j.physletb.2006.08.037](https://doi.org/10.1016/j.physletb.2006.08.037) (cit. on p. 51).
- [166] Matteo Cacciari, Gavin P. Salam and Gregory Soyez. ‘The Catchment Area of Jets’. In: *JHEP* 04 (2008), p. 005. doi: [10.1088/1126-6708/2008/04/005](https://doi.org/10.1088/1126-6708/2008/04/005) (cit. on pp. 51, 53).
- [167] Matteo Cacciari, Gavin P. Salam and Gregory Soyez. ‘FastJet User Manual’. In: *Eur. Phys. J. C* 72 (2012), p. 1896. doi: [10.1140/epjc/s10052-012-1896-2](https://doi.org/10.1140/epjc/s10052-012-1896-2) (cit. on pp. 51, 53).
- [168] Georges Aad et al. ‘Performance of jet substructure techniques for large- R jets in proton-proton collisions at $\sqrt{s} = 7$ TeV using the ATLAS detector’. In: *JHEP* 09 (2013), p. 076. doi: [10.1007/JHEP09\(2013\)076](https://doi.org/10.1007/JHEP09(2013)076) (cit. on p. 51).
- [169] Simone Marzani, Gregory Soyez and Michael Spannowsky. *Looking inside jets: an introduction to jet substructure and boosted-object phenomenology*. Vol. 958. Springer, 2019 (cit. on pp. 51, 52).
- [170] S. Catani et al. ‘Resummation of large logarithms in e^+e^- event shape distributions’. In: *Nucl. Phys. B* 407 (1993), pp. 3–42. doi: [10.1016/0550-3213\(93\)90271-P](https://doi.org/10.1016/0550-3213(93)90271-P) (cit. on p. 51).
- [171] Mrinal Dasgupta and Gavin P. Salam. ‘Event shapes in e^+e^- annihilation and deep inelastic scattering’. In: *J. Phys. G* 30 (2004), R143. doi: [10.1088/0954-3899/30/5/R01](https://doi.org/10.1088/0954-3899/30/5/R01) (cit. on p. 51).
- [172] George F. Sterman and Steven Weinberg. ‘Jets from Quantum Chromodynamics’. In: *Phys. Rev. Lett.* 39 (1977), p. 1436. doi: [10.1103/PhysRevLett.39.1436](https://doi.org/10.1103/PhysRevLett.39.1436) (cit. on p. 51).
- [173] Stephen D. Ellis and Davison E. Soper. ‘Successive combination jet algorithm for hadron collisions’. In: *Phys. Rev. D* 48 (1993), pp. 3160–3166. doi: [10.1103/PhysRevD.48.3160](https://doi.org/10.1103/PhysRevD.48.3160) (cit. on p. 51).
- [174] David Krohn, Jesse Thaler and Lian-Tao Wang. ‘Jets with Variable R ’. In: *JHEP* 06 (2009), p. 059. doi: [10.1088/1126-6708/2009/06/059](https://doi.org/10.1088/1126-6708/2009/06/059) (cit. on pp. 51, 125).
- [175] Mrinal Dasgupta et al. ‘Jet substructure with analytical methods’. In: *Eur. Phys. J. C* 73.11 (2013), p. 2623. doi: [10.1140/epjc/s10052-013-2623-3](https://doi.org/10.1140/epjc/s10052-013-2623-3) (cit. on p. 52).
- [176] Ian Moulton, Benjamin Nachman and Duff Neill. ‘Convolved Substructure: Analytically Decorrelating Jet Substructure Observables’. In: *JHEP* 05 (2018), p. 002. doi: [10.1007/JHEP05\(2018\)002](https://doi.org/10.1007/JHEP05(2018)002) (cit. on p. 52).
- [177] Bogdan Malaescu et al. *In situ calibrations for R-scan jets (In preparation)*. Tech. rep. ATL-COM-PHYS-2018-407. Geneva: CERN, Apr. 2019 (cit. on pp. 52–54, 86).

- [178] Georges Aad et al. ‘Topological cell clustering in the ATLAS calorimeters and its performance in LHC Run 1’. In: *Eur. Phys. J. C* 77 (2017), p. 490. doi: [10.1140/epjc/s10052-017-5004-5](https://doi.org/10.1140/epjc/s10052-017-5004-5) (cit. on pp. 52, 53, 74).
- [179] Georges Aad et al. ‘Jet energy scale and resolution measured in proton–proton collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector’. In: *Eur. Phys. J. C* 81.8 (2021), p. 689. doi: [10.1140/epjc/s10052-021-09402-3](https://doi.org/10.1140/epjc/s10052-021-09402-3) (cit. on pp. 52, 53, 75).
- [180] M. Aaboud et al. ‘Performance of the ATLAS Track Reconstruction Algorithms in Dense Environments in LHC Run 2’. In: *Eur. Phys. J. C* 77.10 (2017), p. 673. doi: [10.1140/epjc/s10052-017-5225-7](https://doi.org/10.1140/epjc/s10052-017-5225-7) (cit. on pp. 52, 53).
- [181] T. Cornelissen et al. ‘The new ATLAS track reconstruction (NEWT)’. In: *J. Phys. Conf. Ser.* 119 (2008). Ed. by Randall Sobie, Reda Tafirout and Jana Thomson, p. 032014. doi: [10.1088/1742-6596/119/3/032014](https://doi.org/10.1088/1742-6596/119/3/032014) (cit. on p. 52).
- [182] Georges Aad et al. ‘Optimisation of large-radius jet reconstruction for the ATLAS detector in 13 TeV proton–proton collisions’. In: *Eur. Phys. J. C* 81.4 (2021), p. 334. doi: [10.1140/epjc/s10052-021-09054-3](https://doi.org/10.1140/epjc/s10052-021-09054-3) (cit. on pp. 54, 125).
- [183] Morad Aaboud et al. ‘In situ calibration of large-radius jet energy and mass in 13 TeV proton–proton collisions with the ATLAS detector’. In: *Eur. Phys. J. C* 79.2 (2019), p. 135. doi: [10.1140/epjc/s10052-019-6632-8](https://doi.org/10.1140/epjc/s10052-019-6632-8) (cit. on p. 54).
- [184] Georges Aad et al. ‘ATLAS data quality operations and performance for 2015–2018 data-taking’. In: *JINST* 15.04 (2020), P04003. doi: [10.1088/1748-0221/15/04/P04003](https://doi.org/10.1088/1748-0221/15/04/P04003) (cit. on p. 55).
- [185] Georges Aad et al. ‘Monitoring and data quality assessment of the ATLAS liquid argon calorimeter’. In: *JINST* 9 (2014), P07024. doi: [10.1088/1748-0221/9/07/P07024](https://doi.org/10.1088/1748-0221/9/07/P07024) (cit. on p. 55).
- [186] Morad Aaboud et al. ‘Operation and performance of the ATLAS Tile Calorimeter in Run 1’. In: *Eur. Phys. J. C* 78.12 (2018), p. 987. doi: [10.1140/epjc/s10052-018-6374-z](https://doi.org/10.1140/epjc/s10052-018-6374-z) (cit. on p. 55).
- [187] Georges Aad et al. ‘Operation and performance of the ATLAS semiconductor tracker’. In: *JINST* 9 (2014), P08009. doi: [10.1088/1748-0221/9/08/P08009](https://doi.org/10.1088/1748-0221/9/08/P08009) (cit. on p. 55).
- [188] ATLAS. *Recommendations for the Modeling of Smooth Backgrounds*. Tech. rep. All figures including auxiliary figures are available at <https://atlas.web.cern.ch/Atlas/GROUPS/PHYSICS/PUBNOTES/ATL-PHYS-PUB-2020-028>. Geneva: CERN, 2020 (cit. on pp. 58, 88, 94–97, 117).
- [189] ‘Identification of boosted, hadronically-decaying W and Z bosons in $\sqrt{s} = 13$ TeV Monte Carlo Simulations for ATLAS’. In: (2015) (cit. on p. 60).
- [190] Vardan Khachatryan et al. ‘Identification techniques for highly boosted W bosons that decay into hadrons’. In: *JHEP* 12 (2014), p. 017. doi: [10.1007/JHEP12\(2014\)017](https://doi.org/10.1007/JHEP12(2014)017) (cit. on p. 60).
- [191] *Performance of mass-decorrelated jet substructure observables for hadronic two-body decay tagging in ATLAS*. Tech. rep. 2018 (cit. on pp. 60, 62, 125).
- [192] ‘Identification of Hadronically-Decaying W Bosons and Top Quarks Using High-Level Features as Input to Boosted Decision Trees and Deep Neural Networks in ATLAS at $\sqrt{s} = 13$ TeV’. In: (Apr. 2017) (cit. on p. 60).
- [193] J. Lin. ‘Divergence measures based on the Shannon entropy’. In: *IEEE Transactions on Information Theory* 37.1 (1991), pp. 145–151. doi: [10.1109/18.61115](https://doi.org/10.1109/18.61115) (cit. on pp. 60, 61).
- [194] *Performance of W/Z taggers using UFO jets in ATLAS*. Tech. rep. All figures including auxiliary figures are available at <https://atlas.web.cern.ch/Atlas/GROUPS/PHYSICS/PUBNOTES/ATL-PHYS-PUB-2021-029>. Geneva: CERN, July 2021 (cit. on pp. 60, 62).

- [195] S. Kullback and R. A. Leibler. ‘On Information and Sufficiency’. In: *The Annals of Mathematical Statistics* 22.1 (1951), pp. 79–86. doi: [10.1214/aoms/1177729694](https://doi.org/10.1214/aoms/1177729694) (cit. on p. 61).
- [196] D.M. Endres and J.E. Schindelin. ‘A new metric for probability distributions’. In: *IEEE Transactions on Information Theory* 49.7 (2003), pp. 1858–1860. doi: [10.1109/TIT.2003.813506](https://doi.org/10.1109/TIT.2003.813506) (cit. on p. 61).
- [197] Albert M Sirunyan et al. ‘Search for low mass vector resonances decaying into quark-antiquark pairs in proton-proton collisions at $\sqrt{s} = 13$ TeV’. In: *JHEP* 01 (2018), p. 097. doi: [10.1007/JHEP01\(2018\)097](https://doi.org/10.1007/JHEP01(2018)097) (cit. on pp. 62, 124).
- [198] Andrew J. Larkoski, Ian Moult and Duff Neill. ‘Power Counting to Better Jet Observables’. In: *JHEP* 12 (2014), p. 009. doi: [10.1007/JHEP12\(2014\)009](https://doi.org/10.1007/JHEP12(2014)009) (cit. on pp. 65–67, 121).
- [199] Georges Aad et al. ‘Muon reconstruction performance of the ATLAS detector in proton–proton collision data at $\sqrt{s} = 13$ TeV’. In: *Eur. Phys. J. C* 76.5 (2016), p. 292. doi: [10.1140/epjc/s10052-016-4120-y](https://doi.org/10.1140/epjc/s10052-016-4120-y) (cit. on p. 74).
- [200] Georges Aad et al. ‘Muon reconstruction and identification efficiency in ATLAS using the full Run 2 pp collision data set at $\sqrt{s} = 13$ TeV’. In: *Eur. Phys. J. C* 81.7 (2021), p. 578. doi: [10.1140/epjc/s10052-021-09233-2](https://doi.org/10.1140/epjc/s10052-021-09233-2) (cit. on p. 74).
- [201] Georges Aad et al. ‘Alignment of the ATLAS Inner Detector in Run-2’. In: *Eur. Phys. J. C* 80.12 (2020), p. 1194. doi: [10.1140/epjc/s10052-020-08700-6](https://doi.org/10.1140/epjc/s10052-020-08700-6) (cit. on p. 74).
- [202] G. Aad et al. ‘Studies of the muon momentum calibration and performance of the ATLAS detector with pp collisions at $\sqrt{s} = 13$ TeV’. In: *Eur. Phys. J. C* 83.8 (2023), p. 686. doi: [10.1140/epjc/s10052-023-11584-x](https://doi.org/10.1140/epjc/s10052-023-11584-x) (cit. on p. 74).
- [203] *Tagging and suppression of pileup jets with the ATLAS detector*. Tech. rep. All figures including auxiliary figures are available at <https://atlas.web.cern.ch/Atlas/GROUPS/PHYSICS/CONFNOTES/ATLAS-CONF-2014-018>. Geneva: CERN, 2014 (cit. on pp. 74, 75).
- [204] D. Buskulic et al. ‘Performance of the ALEPH detector at LEP’. In: *Nucl. Instrum. Meth. A* 360 (1995), pp. 481–506. doi: [10.1016/0168-9002\(95\)00138-7](https://doi.org/10.1016/0168-9002(95)00138-7) (cit. on p. 74).
- [205] Georges Aad et al. ‘New techniques for jet calibration with the ATLAS detector’. In: (Mar. 2023) (cit. on p. 75).
- [206] Morad Aaboud et al. ‘Jet reconstruction and performance using particle flow with the ATLAS Detector’. In: *Eur. Phys. J. C* 77.7 (2017), p. 466. doi: [10.1140/epjc/s10052-017-5031-2](https://doi.org/10.1140/epjc/s10052-017-5031-2) (cit. on p. 75).
- [207] Georges Aad et al. ‘ATLAS b-jet identification performance and efficiency measurement with $t\bar{t}$ events in pp collisions at $\sqrt{s} = 13$ TeV’. In: *Eur. Phys. J. C* 79.11 (2019), p. 970. doi: [10.1140/epjc/s10052-019-7450-8](https://doi.org/10.1140/epjc/s10052-019-7450-8) (cit. on p. 75).
- [208] G. Aad et al. ‘ATLAS flavour-tagging algorithms for the LHC Run 2 pp collision dataset’. In: *Eur. Phys. J. C* 83.7 (2023), p. 681. doi: [10.1140/epjc/s10052-023-11699-1](https://doi.org/10.1140/epjc/s10052-023-11699-1) (cit. on p. 75).
- [209] Georges Aad et al. ‘Measurements of differential cross-sections in top-quark pair events with a high transverse momentum top quark and limits on beyond the Standard Model contributions to top-quark pair production with the ATLAS detector at $\sqrt{s} = 13$ TeV’. In: *JHEP* 06 (2022), p. 063. doi: [10.1007/JHEP06\(2022\)063](https://doi.org/10.1007/JHEP06(2022)063) (cit. on p. 77).
- [210] Lukas Heinrich et al. ‘pyhf: pure-Python implementation of HistFactory statistical models’. In: *Journal of Open Source Software* 6.58 (2021), p. 2823. doi: [10.21105/joss.02823](https://doi.org/10.21105/joss.02823) (cit. on pp. 79, 97, 105).
- [211] Lukas Heinrich, Matthew Feickert and Giordon Stark. *scikit-hep/pyhf: v0.7.0*. Version v0.7.0. Sept. 2022. doi: [10.5281/zenodo.7110486](https://doi.org/10.5281/zenodo.7110486) (cit. on pp. 79, 85, 97, 105).

- [212] F. James and M. Roos. ‘Minuit - a system for function minimization and analysis of the parameter errors and correlations’. In: *Computer Physics Communications* 10.6 (1975), pp. 343–367. doi: [https://doi.org/10.1016/0010-4655\(75\)90039-9](https://doi.org/10.1016/0010-4655(75)90039-9) (cit. on pp. 81, 99, 105).
- [213] Hans Dembinski et al. *scikit-hep/iminuit: v2.17.0*. Version v2.17.0. Sept. 2022. doi: [10.5281/zenodo.7115916](https://doi.org/10.5281/zenodo.7115916) (cit. on pp. 81, 99, 105).
- [214] G. Avoni et al. ‘The new LUCID-2 detector for luminosity measurement and monitoring in ATLAS’. In: *JINST* 13.07 (2018), P07017. doi: [10.1088/1748-0221/13/07/P07017](https://doi.org/10.1088/1748-0221/13/07/P07017) (cit. on p. 87).
- [215] Pauli Virtanen et al. ‘SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python’. In: *Nature Methods* 17 (2020), pp. 261–272. doi: [10.1038/s41592-019-0686-2](https://doi.org/10.1038/s41592-019-0686-2) (cit. on pp. 89, 134).
- [216] Murray Rosenblatt. ‘Remarks on Some Nonparametric Estimates of a Density Function’. In: *The Annals of Mathematical Statistics* 27.3 (1956), pp. 832–837. doi: [10.1214/aoms/1177728190](https://doi.org/10.1214/aoms/1177728190) (cit. on p. 89).
- [217] Emanuel Parzen. ‘On Estimation of a Probability Density Function and Mode’. In: *The Annals of Mathematical Statistics* 33.3 (1962), pp. 1065–1076. doi: [10.1214/aoms/1177704472](https://doi.org/10.1214/aoms/1177704472) (cit. on p. 89).
- [218] Abraham. Savitzky and M Golay. ‘Smoothing and Differentiation of Data by Simplified Least Squares Procedures.’ In: *Analytical Chemistry* 36.8 (1st July 1964), pp. 1627–1639. doi: [10.1021/ac60214a047](https://doi.org/10.1021/ac60214a047) (cit. on p. 89).
- [219] G. Bohm and G. Zech. ‘Statistics of weighted Poisson events and its applications’. In: *Nucl. Instrum. Meth. A* 748 (2014), pp. 1–6. doi: [10.1016/j.nima.2014.02.021](https://doi.org/10.1016/j.nima.2014.02.021) (cit. on p. 90).
- [220] Glen Cowan et al. ‘Asymptotic formulae for likelihood-based tests of new physics’. In: *Eur. Phys. J. C* 71 (2011). [Erratum: *Eur.Phys.J.C* 73, 2501 (2013)], p. 1554. doi: [10.1140/epjc/s10052-011-1554-0](https://doi.org/10.1140/epjc/s10052-011-1554-0) (cit. on pp. 93, 95, 99, 102, 103, 105, 117, 119).
- [221] Alexander Held et al. *scikit-hep/cabinetry: v0.5.1*. Version v0.5.1. Oct. 2022. doi: [10.5281/zenodo.7267629](https://doi.org/10.5281/zenodo.7267629) (cit. on p. 100).
- [222] Abraham Wald. ‘Tests of Statistical Hypotheses Concerning Several Parameters When the Number of Observations is Large’. In: *Transactions of the American Mathematical Society* 54.3 (1943), pp. 426–482 (cit. on pp. 99, 103, 117).
- [223] S. S. Wilks. ‘The Large-Sample Distribution of the Likelihood Ratio for Testing Composite Hypotheses’. In: *The Annals of Mathematical Statistics* 9.1 (1938), pp. 60–62. doi: [10.1214/aoms/1177732360](https://doi.org/10.1214/aoms/1177732360) (cit. on pp. 99, 103, 117).
- [224] James Bradbury et al. *JAX: composable transformations of Python+NumPy programs*. Version 0.3.13. 2018 (cit. on p. 99).
- [225] Charles R. Harris et al. ‘Array programming with NumPy’. In: *Nature* 585.7825 (Sept. 2020), pp. 357–362. doi: [10.1038/s41586-020-2649-2](https://doi.org/10.1038/s41586-020-2649-2) (cit. on p. 100).
- [226] Alexander L. Read. ‘Presentation of search results: The CL(s) technique’. In: *J. Phys. G* 28 (2002). Ed. by M. R. Whalley and L. Lyons, pp. 2693–2704. doi: [10.1088/0954-3899/28/10/313](https://doi.org/10.1088/0954-3899/28/10/313) (cit. on pp. 102, 103).
- [227] Jerzy Neyman and Egon Sharpe Pearson. ‘On the Problem of the Most Efficient Tests of Statistical Hypotheses’. In: *Phil. Trans. Roy. Soc. Lond. A* 231.694-706 (1933), pp. 289–337. doi: [10.1098/rsta.1933.0009](https://doi.org/10.1098/rsta.1933.0009) (cit. on p. 103).
- [228] *Procedure for the LHC Higgs boson search combination in Summer 2011*. Tech. rep. Geneva: CERN, 2011 (cit. on p. 103).

- [229] Elena Accomando et al. ‘Exclusion and discovery via Drell-Yan in the 4DCHM’. In: *PoS DIS2015* (2015), p. 105. doi: [10.22323/1.247.0105](https://doi.org/10.22323/1.247.0105) (cit. on p. 106).
- [230] Arindam Das et al. ‘Classically conformal U(1)’ extended standard model, electroweak vacuum stability, and LHC Run-2 bounds’. In: *Phys. Rev. D* 93.11 (2016), p. 115038. doi: [10.1103/PhysRevD.93.115038](https://doi.org/10.1103/PhysRevD.93.115038) (cit. on p. 106).
- [231] Georges Aad et al. ‘Measurement of the $\Upsilon(1S)$ production cross-section in pp collisions at $\sqrt{s} = 7$ TeV in ATLAS’. In: *Phys. Lett. B* 705 (2011), pp. 9–27. doi: [10.1016/j.physletb.2011.09.092](https://doi.org/10.1016/j.physletb.2011.09.092) (cit. on p. 108).
- [232] Georges Aad et al. ‘Measurement of Upsilon production in 7 TeV pp collisions at ATLAS’. In: *Phys. Rev. D* 87.5 (2013), p. 052004. doi: [10.1103/PhysRevD.87.052004](https://doi.org/10.1103/PhysRevD.87.052004) (cit. on p. 108).
- [233] A. M. Sirunyan et al. ‘Measurement of quarkonium production cross sections in pp collisions at $\sqrt{s} = 13$ TeV’. In: *Phys. Lett. B* 780 (2018), pp. 251–272. doi: [10.1016/j.physletb.2018.02.033](https://doi.org/10.1016/j.physletb.2018.02.033) (cit. on p. 108).
- [234] Sen Jia, Xingyu Zhou and Chengping Shen. ‘Experimental review of the $\Upsilon(1S, 2S, 3S)$ physics at e^+e^- colliders and the LHC’. In: *Front. Phys. (Beijing)* 15.6 (2020), p. 64301. doi: [10.1007/s11467-020-0978-0](https://doi.org/10.1007/s11467-020-0978-0) (cit. on p. 108).
- [235] M. B. Gay Ducati, M. M. Machado and M. V. T. Machado. ‘Diffractive quarkonium production in association with a photon at the LHC’. In: *Phys. Lett. B* 683 (2010), pp. 150–153. doi: [10.1016/j.physletb.2009.12.025](https://doi.org/10.1016/j.physletb.2009.12.025) (cit. on p. 108).
- [236] Eilam Gross and Ofer Vitells. ‘Trial factors for the look elsewhere effect in high energy physics’. In: *Eur. Phys. J. C* 70 (2010), pp. 525–530. doi: [10.1140/epjc/s10052-010-1470-8](https://doi.org/10.1140/epjc/s10052-010-1470-8) (cit. on p. 119).
- [237] *Run 3 luminosity targets*. URL: <https://lhc-commissioning.web.cern.ch/performance/Run-3-targets.htm> (cit. on p. 124).
- [238] *High-Lumi LHC luminosity targets*. URL: https://lhc-commissioning.web.cern.ch/schedule/images/2023/rampup_2022_YETS19weeks_2p.png (cit. on p. 124).
- [239] Morad Aaboud et al. ‘Performance of top-quark and W-boson tagging with ATLAS in Run 2 of the LHC’. In: *Eur. Phys. J. C* 79.5 (2019), p. 375. doi: [10.1140/epjc/s10052-019-6847-8](https://doi.org/10.1140/epjc/s10052-019-6847-8) (cit. on p. 125).
- [240] *Boosted hadronic vector boson and top quark tagging with ATLAS using Run 2 data*. Tech. rep. 2020 (cit. on p. 125).
- [241] Chase Shimmin et al. ‘Decorrelated Jet Substructure Tagging using Adversarial Neural Networks’. In: *Phys. Rev. D* 96.7 (2017), p. 074034. doi: [10.1103/PhysRevD.96.074034](https://doi.org/10.1103/PhysRevD.96.074034) (cit. on p. 125).
- [242] Alison L. Gibbs and Francis Edward Su. ‘On Choosing and Bounding Probability Metrics’. In: *International Statistical Review / Revue Internationale de Statistique* 70.3 (2002), pp. 419–435. (Visited on 29/08/2023) (cit. on p. 129).

Colophon

This document was typeset with the help of KOMA-Script and L^AT_EX using the kaobook class.