Hadronic Jet Physics and the Quest for New Matter in Multi-Jet Final States with the ATLAS Detector at the Large Hadron Collider

Michael E. Nelson

Corpus Christi College,
University of Oxford

A thesis submitted to the University of Oxford in fulfilment of the requirements for the degree of Doctor of Philosophy

Hilary, 2019

Abstract

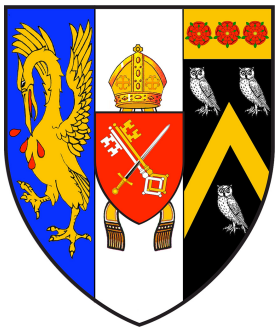
The Standard Model of particle physics is one of the triumphs of contemporary science, but it is well-established that this theory must be incomplete. This thesis focuses on searching for the production of new matter, motivated by beyond-Standard Model framework of *supersymmetry*, with data recorded by the ATLAS experiment during Run 2 of the Large Hadron Collider. One of the crucial physics objects in the search for new matter described in this thesis is the *hadronic jet*. Before describing in detail the 'multi-jets' search for new matter, which uses 36 fb^{-1} of ATLAS data in signal regions where at least 7 hadronic jets are required in the final state, the latest developments in hadronic jet physics at ATLAS shall be motivated and studied. Particular detail will be paid to new approaches in jet reconstruction, jet mass calculations and the corresponding jet mass resolution, and jet calibration at the energy and mass scales. Furthermore the issue of pileup contributions to hadronic jets shall be described, together with recent studies outlining how new developments in inputs to jet reconstruction and jet grooming can be used to improve the performance of hadronic jets in ATLAS and mitigate the effects of pileup therein. The full power of jet reconstruction at ATLAS will then be utilised in the multi-jets analysis. With the 36 fb^{-1} ATLAS data-set, no statistically significant evidence for new matter is reported, and new constraints on the mass of gluinos and stops are pushed up to 1.8 TeV and 1.4 TeV, at the 95 % confidence level, respectively (within the context of particular sets of new physics models). The multi-jets search is sensitive to final states with moderate $E_{\text{T}}^{\text{miss}}$, and consequently is used to assess the LHC sensitivity to both R-parity-conserving (RPC) and R-parity-violating (RPV) supersymmetric scenarios. World-leading constraints across this 'RPC-meets-RPV' spectrum will be assessed at the 95 % confidence level. Finally the ultimate plans for the full Run 2 iteration of the multi-jets analysis, amounting to 140 fb^{-1} of data, will be described.

Hadronic Jet Physics and the Quest for New Matter in Multi-Jet Final States with the ATLAS Detector at the Large Hadron Collider

Michael E. Nelson

Corpus Christi College,

University of Oxford





A thesis submitted to the University of Oxford in fulfilment of the requirements for the degree of Doctor of Philosophy

Hilary, 2019

Acknowledgements

It is exceptionally difficult to condense into a few paragraphs the numerous people who have supported and guided me up to this point. I am hugely indebted to each and every one of you, and there are far more of you who have helped me than I can mention here. To you all, I am grateful to have had your support, guidance, and friendship every step of the way.

First of all, I would like to thank my supervisor Claire Gwenlan and the Oxford SUSY group leader Alan Barr, both of whom have provided an immense amount of support and insight throughout this time, and have not only been incredibly supportive and helpful in my research, but have also encouraged me to pursue the areas of particle physics on ATLAS which I became most interested in. It has been a privilege to be a part of the Oxford SUSY group, and to work with some of the most remarkable scientists I have ever encountered. I've had a lot of fun working with (and learning from) the present and past academics, postdocs, and students in Oxford's particle physics group, including Will Fawcett, Will Kalderon, Jon Burr, Nurfikri Norjoharuddeen, Moritz Backes, Bill Balunas, Todd Huffman, Cigdem Issever, Lydia Beresford, James Frost, Hayden Smith, Luigi Marchese, Gabija Zemaityte, Kathrin Becker, Santiago Paredes, Jesse Liu, and Aaron O'Neill (to name a few!)

I owe immeasureable amounts of thanks to the generations of Oxford SUSY group students who preceded me. In particular I would like to thank Will Fawcett and Will Kalderon for welcoming me to the group, and providing a wealth of advice, patience, and insight from the very beginning. Will Fawcett's infectious enthusiasm and sterling work ethic was a real inspiration during the wonderful days of the 'multi-jets' analysis. I am also grateful to Jon Burr, who could always be relied on to ask the difficult (and important!) physics questions and provide hours and hours of helpful coding and software advice. Finally, I must personally thank Jesse Liu for the intense, thought-provoking, and insightful physics discussions during our studies. Jesse was always there to provide a new take on a problem in an analysis/the field as a whole, and I have grown so much as a physicist throughout my time working with him. I can only hope that I have been even half as helpful to the younger Oxford SUSY generation: Aaron, Gabriel, and Jonas.

It has been a privilege to be a member of the 'multi-jets' analysis, and even more of a privilege to work with such a talented team of scientists. I'm particularly grateful to Will Fawcett and Jon Burr (again!), as well as Teng Jian Khoo and Marco Valente who have been an absolute pleasure to collaborate with, and were always there with the right answers to any questions I had. Many, many thanks also to the members of the wider ATLAS SUSY group for

their help over the years, in particular Till Eifert, Zach Marshall, Max Swiatlowski, and Chiara Rizzi.

In this thesis I talk about jets. A lot. It has been a real highlight of my past few years to work in the ATLAS Jet/$E_{\mathrm{T}}^{\mathrm{miss}}$ group. When I started on ATLAS, I thoughts I would be involved in study of jets as a means of qualifying as an ATLAS author, but the immense enthusiasm for the field and fun I had propelled me to work with this group for the rest of my doctorate (and hopefully a few years beyond it!) I have loved every moment of learning about the intricacies of jet physics and reconstruction, substructure, grooming, calibration, and much more besides. Not to mention the karaoke nights at the Hadronic Calibration workshops. My sincere thanks to David Miller, Piere-Antoine Delsart, Steven Schramm, and Cigdem Issever who really got me hooked on jets and have taught me so much. Cigdem deserves special recognition for supervising (and getting me through!) my qualification task − thanks so much! Within the jet substructure group, I am grateful for the happy memories and wonderful discussions and workshops, and in particular must thank Chris Delitzsch, Jen Roloff, Matt LeBlanc, Chris Young, Christophe Clement, Oleg Brandt, Joe Taenzer, Sam Meehan, Amal Vaidya, Jason Veatch, Ece Akili, Dag Gilberg, Ben Nachman, Francesco Spano, Jacob Rawling, and Peter Loch for being wonderful friends and colleagues in the jet world. I have learned so much from each of you.

I must extend a special thanks to the LTA community for providing a wonderful experience of life at CERN; those LTA dinners, parties, and Wednesday/Friday beer sessions kept me (almost!) sane, and sharing those experiences in particular with Amal Vaidya, Alison Tully, James Cowley, and Matt Tilley has been a blast.

To close, I must thank a few particularly important people who have shaped me as a scientist over the years. During my undergraduate years at Cambridge, Julia Riley was a particularly inspiring teacher, mentor, and Director of Studies, instilling in me not only the foundations of science, but also a belief in myself. Unforgettably, I would not be where I am today without the constant guidance and friendship provided by John Fry and John Dainton at the University of Liverpool. During my late schooldays, John Fry was a real inspiration and introduced me to the world of research, particle physics, and CERN. The days spent working with him on NA62 during my undergraduate summers are some of my happiest, and I owe him so much for his kindness and encouragement to pursue scientific research as a career.

To *really* close, I must thank those who have shaped me as a person over the years, in particular my parents, Beverley and Gary; my sister, Bethany; and my late-grandmother, Mavis. My immense thanks to them for their unending love and support. This thesis is dedicated to all four.

M. E. Nelson (April, 2019)

Preface

The work presented in this thesis was completed by the author as a member of the ATLAS collaboration. In ATLAS no publication is created alone, and the work described in this thesis is of course no exception. What follows is a list in reverse chronological order of publications that the author has made leading contributions to during the 2015 − 2018 Run 2 ATLAS data-taking. Many, but not all, of these publications are covered in this thesis. For completeness, two upcoming journal publications (where the author is serving as an analysis leader and paper editor) are briefly described at the end.

Publications

- *In situ* **calibration of large-R jet energy and mass in 13 TeV proton-proton collisions with the ATLAS detector** (ATLAS Collaboration, 2018, accepted by EPJC, see reference [104])

The author was the main developer of the combined jet mass (Chapter 6 of this thesis) and sole responsible for its creation and continued maintenance in ATLAS, which resulted in substantial improvements in the calculation and measurement of the jet mass resolution, applied to both Monte Carlo and *in situ* studies in ATLAS.

- **Jet Mass Resolution in ATLAS Run 2 Monte Carlo Simulations** (ATLAS Collaboration, 2018, ATL-PHYS-PUB-2018-015, see reference [101])

This work is described in Chapter 6. The author was the analysis contact, contact editor, and chief developer of all studies in this publication, primarily interested in formally quantifying the behaviour of the jet mass resolution in ATLAS Monte Carlo, and its systematic uncertainties. These studies will directly impact and change how jet mass resolution systematics are calculated in ATLAS, allowing for the eventual switch from systematics based on Gaussian smearing to systematics based on the forward-folding of Monte Carlo jet masses. The author's development of jet mass resolution smearing methods has allowed for the resolution systematic uncertainties to be directly applied to many searches and measurements in ATLAS, from exotics searches to top differential cross-section measurements.

- **Search for pair production of higgsinos in the HH in final states with at least 3 b-tagged jets in $\sqrt{s} = 13$ TeV pp collisions using the ATLAS detector** (ATLAS Collaboration, 2018, Phys. Rev. D 98, 092002)

This paper summarises a search for electroweak production of Higgsinos, which decay in final states with b-jets. The author led the $H \to b\bar{b}$ tagging study by comparing b-tagging approaches using track-jets with those using standard small-radius jets, for both resolved and boosted Higgsino production signals, and was also the responsible for evaluating the theory systematics for the key Monte Carlo backgrounds, such as $t\bar{t}$ and single top production. This work is not described in this thesis.

- **Reinterpretation of searches for supersymmetry in models with variable R-parity-violating coupling strength and long-lived R-hadrons** (ATLAS Collaboration, 2018, ATLAS-CONF-2018-003, see reference [161])

The author led the reinterpretation of the multi-jets analysis in large classes of R-parity-violating supersymmetric models with long-lived particles and displaced vertices, resulting in new limits on the R-parity-violating coupling which had never been determined before at ATLAS, and presented this work at the SUSY 2018 conference. This work is also briefly described in Chapter 9 of this thesis.

- **Impact of Alternative Inputs and Grooming Methods on Large-R Jet Reconstruction in ATLAS** (ATLAS Collaboration, 2017, ATL-PHYS-PUB-2017-020, see reference [123])

Much of this work is described in Chapter 7. The author was the sole developer of a generic jet mass and energy calibration framework, allowing for the simultaneous jet energy and jet mass calibration of the shortlisted jet collections in this publication, which are based on Soft Drop grooming and constituent subtraction techniques. These calibrations were key to the final conclusions of this study, which has gone on to change the nature of large-radius jet reconstruction in ATLAS. The author also studied the role of jet reclustering as a jet groomer, which had not previously been considered in ATLAS. Finally the author directly contributed to all decisions and discussions on the final metrics used to determine the shortlisted large-radius jet definitions, and presented this work at the BOOST 2018 conference.

- **Search for new phenomena with large jet multiplicities and missing transverse momentum using large-radius jets and flavour-tagging at ATLAS in 13 TeV pp collisions** (ATLAS Collaboration, 2017, JHEP 12 (2017) 034, see reference [70])

This paper forms the basis of Part III of this thesis, and the author was one of the lead investigators, playing a role in every aspect of this challenging multi-jets analysis: from n-tuple production to the final HEPData entry, to the determination of key systematics such as the flavour-dependence of the data-driven QCD background estimation technique. In addition the author led all studies in applying jet substructure techniques using reclustered jet masses and top-tagging with jet-level variables, and also introduced reinterpretations based on R-parity-violating supersymmetry to this (historically) R-parity-conserving supersymmetry search, allowing for the sensitivity to stop production to be assessed. Finally, the author was the sole responsible for the final analysis exclusion limits on three simplified models and the pMSSM slice of supersymmetric phase-space. Since publication, the author has become an analysis contact for the current iteration of this new physics search, which uses the full Run 2 ATLAS data-set, and has presented the work in this paper at the following conferences: SUSY 2018, LHCC in 2018, and the UK IoP conference in 2018.

- **Jet reclustering and close-by effects in ATLAS run II** (ATLAS Collaboration, 2017, ATLAS-CONF-2017-062)

The author developed and derived the jet mass calibrations for small-radius jets, used to create the reclustered large-radius jets in this study. This had direct implications for the reclustered jet mass resolution, which offers performance gains over the mass resolution of conventional large-radius jets in ATLAS. Reclustered jets are further utilised in the multi-jets search, described in Part III of this thesis. The details of this publication shall not be discussed further in this work.

- **Pursuit of new phenomena in final states with high jet multiplicity, high jet masses and missing transverse momentum with ATLAS at $\sqrt{s} = 13$ TeV** (ATLAS Collaboration, 2016, ATLAS-CONF-2016-095)

This paper is the 18 fb^{-1} iteration of the multi-jets analysis, and highlights the author's first contributions to that search for new physics. For this publication, the author was responsible for devising cross-checks of the analysis region definitions and cut-flow, and also developed and calculated the systematic uncertainty arising from the flavour-dependence of the QCD background estimation technique. This publication is not discussed in this thesis.

- **Jet mass reconstruction with the ATLAS Detector in early Run 2 data** (ATLAS Collaboration, 2016, ATLAS-CONF-2016-035, see reference [98])

The author was responsible for deriving the jet mass calibrations for both the calorimeter and track-assisted jet masses (Chapter 6), and led the software implementation and calculation of the combined jet mass (Chapter 6), which subsequently became the recommended jet mass definition for jet reconstruction in ATLAS. The author also presented an extension of this work at the LHCC in 2017.

Upcoming Publications (working titles)

- **New Approaches in Jet Reconstruction and Grooming in Run 2 of the ATLAS detector**

This will be a journal paper aiming for publication later in 2019. The author is one of the paper editors, responsible for the final publication. In addition the author has worked on jet energy and jet mass calibrations of new jet collections, including those based on Soft Drop, particle flow, and track-calo cluster (TCC) jet definitions. The author has also been working on using unfolding techniques in data with these new jet definitions, and is the responsible for the data and Monte Carlo comparison studies of several jet substructure variables which will be considered in this publication, in order to assess the performance of Monte Carlo modelling of Run 2 ATLAS data. Some of the ongoing work in this analysis is briefly described in Chapter 7.

- **Search for new phenomena with large jet multiplicities and missing transverse momentum using large-radius jets and flavour-tagging with the full run 2 ATLAS data-set**

This is the upcoming paper for the full Run 2 multi-jets analysis, for which the author is the analysis contact and contact editor. In this analysis the author is leading a programme of research and development which will include the study of particle flow jets, object-based $E_{\mathrm{T}}^{\mathrm{miss}}$ significance, quark-gluon tagging, and jet reclustering optimisation into this search for the first time in order to enhance analysis sensitivity to gluino production at the LHC. The author is active in every aspect of this analysis, contributing to the n-tuple production, region definitions, systematic uncertainties, and fitting strategies. Some of the ongoing work in this analysis is briefly described in Chapter 10.

Throughout this thesis, if a diagram, plot or other figure has not been produced by the author or author's analysis, then a reference to the source is stated in the caption. Plots that the author has created that have been published by the ATLAS collaboration are labelled **ATLAS** or **ATLAS Preliminary**.

Abstract

The Standard Model of particle physics is one of the triumphs of contemporary science, but it is well-established that this theory must be incomplete. This thesis focuses on searching for the production of new matter, motivated by beyond-Standard Model framework of *supersymmetry*, with data recorded by the ATLAS experiment during Run 2 of the Large Hadron Collider. One of the crucial physics objects in the search for new matter described in this thesis is the *hadronic jet*. Before describing in detail the 'multi-jets' search for new matter, which uses 36 fb^{-1} of ATLAS data in signal regions where at least 7 hadronic jets are required in the final state, the latest developments in hadronic jet physics at ATLAS shall be motivated and studied. Particular detail will be paid to new approaches in jet reconstruction, jet mass calculations and the corresponding jet mass resolution, and jet calibration at the energy and mass scales. Furthermore the issue of pileup contributions to hadronic jets shall be described, together with recent studies outlining how new developments in inputs to jet reconstruction and jet grooming can be used to improve the performance of hadronic jets in ATLAS and mitigate the effects of pileup therein. The full power of jet reconstruction at ATLAS will then be utilised in the multi-jets analysis. With the 36 fb^{-1} ATLAS data-set, no statistically significant evidence for new matter is reported, and new constraints on the mass of gluinos and stops are pushed up to 1.8 TeV and 1.4 TeV, at the 95 % confidence level, respectively (within the context of particular sets of new physics models). The multi-jets search is sensitive to final states with moderate $E_{\mathrm{T}}^{\mathrm{miss}}$, and consequently is used to assess the LHC sensitivity to both R-parity-conserving (RPC) and R-parity-violating (RPV) supersymmetric scenarios. World-leading constraints across this 'RPC-meets-RPV' spectrum will be assessed at the 95 % confidence level. Finally the ultimate plans for the full Run 2 iteration of the multi-jets analysis, amounting to 140 fb^{-1} of data, will be described.

Contents

1 Introduction **1**

I The Theory of Experimental Particle Physics **4**

2 The Standard Model of Particle Physics **5**
2.1 The Particle Content of the Standard Model 5
2.1.1 Fermions and Matter 6
2.1.2 Bosons and Forces 7
2.2 Quantum Fields and their Interactions 8
2.2.1 Quantum Electrodynamics (QED) 8
2.2.2 Quantum Chromodynamics (QCD) 10
2.2.3 Electroweak Theory 13
2.3 The Higgs Mechanism 17
2.3.1 Electroweak Symmetry Breaking (EWSB) 19
2.3.2 Fermonic Masses from the Higgs 20
2.4 Putting it all together: the SM Lagrangian 23
2.5 Collider Physics 23
2.5.1 Scattering Theory 24
2.5.2 Luminosity 24
2.6 Summary 25

3 *Beyond* the Standard Model of Particle Physics **26**
3.1 Problems with the Standard Model 26
3.1.1 Neutrino Masses 26
3.1.2 Dark Matter 28
3.1.3 Hierarchies 29
3.1.4 Matter-Antimatter Asymmetry and Baryogenesis 29
3.1.5 Grand Unification 30
3.2 Supersymmetry (SUSY) 31
3.2.1 Higgs Mass Hierarchy Problem 31
3.2.2 Superalgebras and Superfields 32
3.2.3 Phenomenology of $N = 1$ SUSY 34
3.3 Supersymmetric Models at Hadron Colliders 39
3.3.1 SUSY Production at the LHC 39
3.3.2 Simplfied Models 39
3.4 Summary 41

4 The ATLAS Experiment **42**
4.1 The Large Hadron Collider (LHC) . . . 42
4.2 The ATLAS Detector . . . 44
4.2.1 Coordinate system . . . 44
4.2.2 Overview . . . 47
4.2.3 Inner Detector . . . 47
4.2.4 Calorimeters . . . 50
4.2.5 Muon System . . . 52
4.3 Trigger . . . 54
4.3.1 L1 Trigger . . . 54
4.3.2 HLT . . . 55
4.4 Physics Object Reconstruction . . . 56
4.4.1 Primary Vertices and Tracks . . . 56
4.4.2 Electrons and Photons . . . 57
4.4.3 Muons . . . 57
4.4.4 Jets . . . 58
4.4.5 Taus . . . 59
4.4.6 Missing Transverse Momentum . . . 59
4.4.7 Overlap Removal . . . 60
4.5 Summary . . . 60

II Hadronic Jet Performance during Run 2 of ATLAS **61**

5 Jet Reconstruction and Calibration in Run 2 **62**
5.1 Overview . . . 63
5.1.1 Building an Interaction . . . 63
5.1.2 What is a Jet? . . . 66
5.2 Jet Reconstruction . . . 67
5.2.1 Jet Inputs . . . 68
5.2.2 Jet Algorithms . . . 70
5.2.3 Jet Grooming . . . 75
5.3 Calibration at the Jet Energy Scale (JES) . . . 77
5.3.1 Small-Radius Jets . . . 77
5.3.2 Large-Radius Jets . . . 90
5.3.3 JES Uncertainties . . . 92

6 The Jet Mass and its Resolution **95**
6.1 The Jet Mass Scale . . . 96
6.1.1 Overview . . . 96
6.1.2 The Calorimeter Mass . . . 96
6.1.3 The Track-Assisted Mass . . . 97
6.1.4 Jet Mass Calibration . . . 98

6.1.5 The Combined Mass . . . 103
6.1.6 Particle Flow . . . 109
6.2 The Jet Mass Resolution . . . 110
6.2.1 Monte Carlo Samples . . . 112
6.2.2 Monte Carlo-based Jet Mass Resolution . . . 113
6.2.3 Investigating non-Gaussian features of the Jet Mass Resolution . . . 119
6.2.4 Jet Mass Resolution from Forward-Folding . . . 121
6.2.5 JMR comparisons: Gaussian Smearing versus Forward-Folding . . . 126

7 Optimising Large-Radius Jets in Run 2 of ATLAS 131
7.1 New Jet Inputs . . . 132
7.1.1 Voronoi Subtraction . . . 133
7.1.2 Constituent Subtraction . . . 135
7.1.3 SoftKiller . . . 136
7.2 New Jet Groomers . . . 137
7.2.1 Pruning . . . 139
7.2.2 Modified Mass Drop Tagger . . . 140
7.2.3 Soft Drop . . . 141
7.3 Jet Substructure Variables . . . 144
7.3.1 N-subjettiness . . . 145
7.3.2 Energy Correlation Functions . . . 146
7.4 The Jet Inputs and Grooming Scan . . . 148
7.4.1 First Run 2 Inputs and Grooming Scan . . . 149
7.4.2 The Updated Inputs and Grooming Scan . . . 174
7.4.3 The Future of Large-Radius Jets at ATLAS . . . 184

III The Hunt for New Matter in Multi-Jet Final States 185

8 Hunting for new phenomena in multi-jet final states 186
8.1 Analysis Overview . . . 187
8.2 Signal Models . . . 188
8.2.1 Two-Step Model . . . 189
8.2.2 pMSSM Slice . . . 190
8.2.3 $\tilde{g}t\bar{t}$ Model . . . 191
8.2.4 RPV $\tilde{t} \to \bar{s}\bar{b}$ Model . . . 191
8.3 Analysis Strategy and Physics Objects . . . 192
8.3.1 Trigger . . . 192
8.3.2 Event Selection and Cleaning . . . 192
8.3.3 Physics Objects Selection . . . 194
8.3.4 Kinematic Variables . . . 196
8.3.5 Region Definitions . . . 198
8.3.6 Jet Reclustering and Jet Substructure for Top-Tagging . . . 201

8.4 Leptonic Background Estimation . . . 205
8.5 QCD Multi-Jet Background Estimation . . . 207
8.6 Treatment of Uncertainties . . . 213
8.6.1 Detector Systematics . . . 213
8.6.2 Theory Systematics . . . 215
8.6.3 Template Systematics . . . 217
8.6.4 Systematics Summary . . . 220

9 Constraining new phenomena in multi-jet final states 223
9.1 Results with 36 fb^{-1} of ATLAS Data . . . 223
9.1.1 Background-Only Fit . . . 223
9.1.2 Unblinded Results . . . 225
9.2 Interpretations with 36 fb^{-1} of ATLAS Data . . . 227
9.2.1 Limit Setting . . . 227
9.2.2 Constraining Gluinos, Neutralinos, and Stops . . . 232
9.2.3 Constraining R-Parity Violating SUSY . . . 233

10 Multi-Jets with Full Run 2 Data 240
10.1 Jet Reclustering Optimisation . . . 240
10.2 Object-Based $E_{\text{T}}^{\text{miss}}$ Significance . . . 242
10.3 Particle Flow Jets . . . 245
10.4 Quark-gluon Tagging . . . 246
10.5 Summary . . . 248

11 Conclusions 249

Appendices

A Appendix: Jet Mass Resolution at the HL-LHC 252

B Appendix: Soft Drop Combined Mass 254

C Appendix: $E_{\text{T}}^{\text{miss}}/\sqrt{H_{\text{T}}}$ Distributions in all Multi-Jets Signal Regions 258

List of Figures 265

List of Tables 266

1 Introduction

The goal of particle physics is to analyse and understand the universe from its most fundamental constituents, and to explain how those most basic 'building blocks' of reality are formed and interact with other fundamental physical objects, or *particles*. This endeavour has been practised by humanity for thousands of years: in 6th century BC, the Ancient Grecian pre-Socratic philosopher Anaximenes of Miletus proposed that all of creation is composed of four fundamental elements − earth, air, fire, and water. Around 4th century BC, Leucippus and his pupil Democritus coined the word *atomos* − meaning 'uncuttable' − to describe the indivisible building blocks of matter. This early atomism was further developed in Greek society by Epicurus, and subsequently (and independently) emerged in Islamic philosophy. Although far from a true reflection of the rich and complex subatomic nature of our observable universe, these theories do illustrate the early human desire to understand and underpin reality in its most basic possible terms.

Modern particle physics has grown from countless theoretical and experimental scientific developments over the past millennia. The 1896 discovery of the electron − the first experimental verification of fundamental matter − propelled the experimental study of particle physics into mainstream scientific endeavour. Over the course of the 20th century, the electron discovery was accompanied by the experimental observation of a plethora of new particles across laboratories around the world. In parallel to this, the field of modern theoretical physics was turned on its head by the developments of relativity, quantum theory, and the eventual unification of

special relativity and quantum physics into *Quantum Field Theory* (QFT). With the theoretical framework of QFT as a guide, the theoretical foundations of a theory which explained the observation of these new particles − fundamental fermions which form bound hadronic states of baryons and mesons, along with leptons; and the bosons which mediate the interactions of reality itself − were synthesised into the *Standard Model* (SM) of particle physics. The SM continues to be tested by experiments around the world, often situated at large national and international laboratories, such as the *Organisation européenne pour la recherche nucléaire*, or CERN, situated at the French-Swiss border outside Geneva, Switzerland. It was at CERN that the discovery of the Higgs boson − a long-sought crucial ingredient to the SM − was announced in July, 2012, by the ATLAS and CMS experiments which operate on the Large Hadron Collider ring. This discovery is the jewel in the diadem of the SM, and provides further support that we are moving ever closer to a more complete description of the matter content of our universe.

This thesis is not about praising verbatim the numerous (and astonishingly accurate!) predictions of the SM, but rather focuses on describing its theoretical shortcomings and outlining how modern particle physics experiments are searching for new physics and new matter *beyond* what the SM predicts. In this thesis, we shall aim to establish connections between theoretical proposals for sources of new physics beyond the SM and experimental searches for that same new physics, using data taken at the ATLAS experiment throughout 2015 and 2016. Physicists use complex reconstruction techniques to measure and analyse the production of particles at the LHC, be they electrons, muons, photons, etc. However, one crucial physics object − the *central* physics object to the work described in this thesis − is the *hadronic jet*. Jets are experimental manifestations of the complex processes of hadronisation and fragmentation of particles, as predicted by QFT, produced in the proton-proton collisions at the LHC. This thesis will present some of the latest studies of jet physics at the ATLAS experiment, including the reconstruction, calibration and substructure of hadronic jets. These jets shall then be utilised in a search for new physics − motivated by the beyond SM theory of *Supersymmetry* − taking place at the ATLAS experiment. We shall see how jet physics can be used to set world-leading constraints on the production of supersymmetric new matter at the LHC.

This story shall be described across three 'Parts'. In the first Part of this thesis (Chapters 2-4) the theory of experimental particle physics shall be briefly outlined, starting with the theoretical

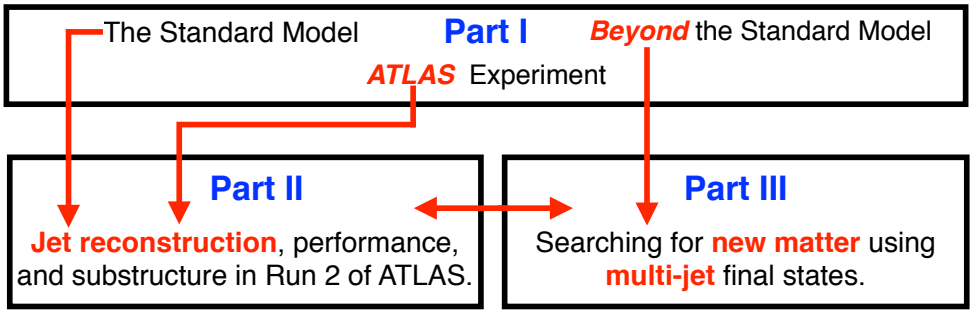


Figure 1.1: The connections between theoretical developments, jet physics, and experimental searches for new physics described within this thesis.

foundations of the SM, the subsequent development and motivation of beyond SM theories through the example of supersymmetry, and finally a brief segway into the hardware, software, and reconstruction techniques comprising the ATLAS experiment. The second Part (Chapters 5-7) describes the developments in hadronic jet physics at ATLAS in which the author has played, and continues to play, a leading role. Finally, the third Part (Chapters 8-10) attempts to connect the theoretical background described in Part I with the jet reconstruction techniques and developments presented in Part II, and motivate a benchmark search for new physics which the author has been leading, inspired by the prospect of the production of new matter at the LHC, which subsequently decays to final states with many hadronic jets. These connections are visually illustrated in Figure 1.1. A summary of this thesis is then provided in Chapter 11.

Part I

The Theory of Experimental Particle Physics

The law is the true embodiment,
Of ev'rything that's excellent,
It has no kind of fault or flaw,
And I, my lords, embody the law.

– W. S. Gilbert,
Iolanthe, Act I

2

The Standard Model of Particle Physics

The modern theory of particle physics is a mathematical description of a collection of particles, which can be characterised by fields described in the framework of Quantum Field Theory, and their associated interactions. Together these particles and interactions comprise the Standard Model of Particle Physics (SM), which utilises the concept of *symmetry* and *symmetry groups* to produce predictive, measurable quantities which are realised in nature to a significant degree of accuracy. The SM, which can be thought of as an *effective field theory* describing the particles and interactions up to large energy scales [1], is certainly one of the greatest theoretical and experimental triumphs of the 20th Century. In this Chapter we shall examine the fundamental concepts which underpin the SM, first describing the constituent fundamental particles and forces of our universe, before discussing the quantum field theories of the *electromagnetic, electroweak, and strong interactions* of subatomic particles. We shall then move to a brief review of the *Brout-Englert-Higgs mechanism* [2–4]; this process provides a deep understanding of where and how the 'building blocks' of our observable universe acquire mass. Finally, the basic principles of collider physics phenomenology shall be reviewed.

2.1 The Particle Content of the Standard Model

The contents of the SM can be broadly classified according to a quantum number characterising a given particle's intrinsic angular momentum, referred to as *spin*. The spin takes integer or half-integer values in units of the reduced Planck constant, $\hbar$; the nature of a particle's spin

value is fundamental to the role it plays in the SM, and its subsequent interactions.

2.1.1 Fermions and Matter

Particles possessing half-integer spin, such as 1/2, are referred to as *fermions*. The fermions of the SM are traditionally thought of as fundamental (or point-like) particles which give rise to matter in the observable universe. In other words, they can be thought of as representing the fundamental 'building blocks' of reality as described by the SM. The beginnings of particle physics are rooted in the discovery of the electron, the first spin-1/2 fundamental particle to be measured, by J. J. Thomson, in 1896 [5]. The electron, and other subsequently discovered fermions, are summarised in Table 2.1, together with their electromagnetic charges (in units of elementary charge, e).

Fermion	Mass	Electric charge [e]	Year of Discovery
up quark, u	2.3 MeV	$+\frac{2}{3}$	1968
down quark, d	4.8 MeV	$-\frac{1}{3}$	1968
charm quark, c	1.275 GeV	$+\frac{2}{3}$	1974
strange quark, s	95 MeV	$-\frac{1}{3}$	1968
bottom quark, b	4.18 GeV	$-\frac{1}{3}$	1977
top quark, t	173 GeV	$+\frac{2}{3}$	1995
electron, e	0.511 MeV	-1	1896
electron neutrino, ν_e	< 2 eV	0	1956
muon, μ	105.7 MeV	-1	1936
muon neutrino, ν_μ	< 0.19 MeV	0	1962
tau, τ	1.777 GeV	-1	1974 − 1977
tau neutrino, ν_τ	< 18.2 MeV	0	2000

Table 2.1: The SM fermions, each possessing a spin of $+\frac{1}{2}$. The top six are the quarks; the bottom six the leptons. There are also corresponding anti-fermions of opposite electric charge. Masses are taken from the 2018 Review of Particle Physics [6].

The fermions can be further subdivided in sets of *quarks* and sets of *leptons*, where the former interact via the strong and electroweak forces of nature (to be described in Section 2.2) while the latter only interact electroweakly. There are six 'flavours' of quark, each described by a different flavour quantum number, and each quark type is further distinguished by three 'colours' (red, green, blue) corresponding to quantum numbers of the strong interaction; anti-quarks carry corresponding anti-colour. There are six leptons, and additionally six anti-leptons,

comprising the electron, muon, and tau particles, together with their nearly massless [7] neutrino counterparts. Putting these together, the SM endows nature with 48 fundamental fermions.

2.1.2 Bosons and Forces

Particles possessing integer spin are referred to as *bosons*, where spin-0 bosons are said to be *scalars* and spin-1 *vectors* (or gauge bosons). The bosons can be thought of as the particles which mediate the interactions between the fermions; as such they can be interpreted as the 'force carriers' of nature. Perhaps most familiar is the spin-1 gauge boson responsible for mediating the electromagnetic interaction between charged fermions (and described by the theory of Quantum Electrodynamics): the photon. Together with the W^+, W^-, and Z particles, we arrive at the gauge bosons responsible for mediating electroweak interactions. Finally, the gluon, g, is the gauge boson responsible for mediating strong interactions in nature, as described by the theory of Quantum Chromodynamics. There are in fact 8 gluons, each with different colour wavefunctions, bringing the total number of bosons necessary for strong and electroweak physics up to 12. The bosons of the SM are summarised in Table 2.2.

Boson	Mass	Electric charge [e]	Spin	Force mediated	Year of Discovery
photon, γ	$< 10^{-18}$ eV	0	1	electromagnetic	1923
W boson	80.4 GeV	$\pm$ 1	1	weak	1983
Z boson	91.2 GeV	0	1	weak	1983
gluon, g	0 GeV	0	1	strong	1979
Higgs boson, H	125.1 GeV	0	0	N.A.	2012

Table 2.2: The SM bosons, together with the forces they mediate. Masses are taken from the Review of Particle Physics [6].

One distinction among the family of bosons is the presence of what is currently the only experimentally detected scalar particle, the spin-0 Higgs boson [8, 9]. The Higgs boson has a special place in the SM, and we shall see in Section 2.3 how it gives rise to the experimentally measured masses of the fermions and bosons comprising our universe. Adding in the 13 bosons described by the SM, we are left with 61 fundamental 'building blocks' of nature.

2.2 Quantum Fields and their Interactions

To understand the theoretical dynamics of the SM's particles and their interactions, we need to introduce the formalism of Quantum Field Theory. In classical physics, it is possible to construct a Lagrangian for some mechanical system or classical field which, when processed through the Euler-Lagrange equations, gives rise to the equations of motion underpining that system or field. The forces and particles of the SM can be understood using the same Lagrangian formalism, where now the fundamental objects of the system are *quantum* fields. We shall see that the interactions of the SM fall out naturally when one requires that the SM Lagrangian (density), $\mathcal{L}_{\text{SM}}$ be invariant under certain special symmetry transformations. This is the group-theoretic essence of the SM. To illustrate this concretely, let us start with the simplest case, *quantum electrodynamics*, before building up the various Lagrangians which will allow us to obtain the final $\mathcal{L}_{\text{SM}}$.

2.2.1 Quantum Electrodynamics (QED)

Let us begin with a Lagrangian for the (special) relativistic motion of a spin-1/2 particle, such as an electron, of mass m, which is described by the fermionic quantum field ψ:

$$\mathcal{L}_{\text{Dirac}} = i\bar{\psi}\gamma^{\mu}\partial_{\mu}\bar{\psi} - m\psi\bar{\psi}, \tag{2.1}$$

where the *adjoint* fermionic field is denoted $\bar{\psi}$, and γ^{μ} are the Dirac γ matrices in some representation. This Lagrangian is invariant under a global $U(1)$ transformation of ψ. However, let us now consider how the Lagrangian changes under a *local* $U(1)$ transformation:

$$\psi(x) \to \psi'(x) = e^{i\alpha(x)}\psi(x). \tag{2.2}$$

It is clear that $\mathcal{L}_{\text{Dirac}}$ is no longer invariant under this local symmetry. However, it is possible to restore the invariance observed under the global $U(1)$ symmetry by first generalising to the *covariant derivative*:

$$\partial_{\mu} \to D_{\mu} = \partial_{\mu} - ieA_{\mu}, \tag{2.3}$$

where it is necessary to introduce a new gauge field, A_{μ}, which we also require to transform under a local $U(1)$ according to the following rule:

$$A_\mu \to A'_\mu = A_\mu + \frac{1}{e}\partial_\mu \alpha(x). \tag{2.4}$$

We now see that, with these new variables, the Dirac Lagrangian is invariant under a local $U(1)$ transformation. Furthermore, by introducing a new gauge field A_μ with a well-defined local $U(1)$ transformation, the invariance of the original Dirac Lagrangian has been restored. Physically, A_μ corresponds to the electromagnetic field four-potential, and it is related to the electromagnetic field tensor, $F_{\mu\nu}$, by

$$F_{\mu\nu} = \partial_\mu A_\nu - \partial_\nu A_\mu. \tag{2.5}$$

.

Recall from classical field theory that the manifestly covariant form of Maxwell's equations follow from the Lagrangian

$$\mathcal{L}_{\text{Maxwell}} = -\frac{1}{4} F_{\mu\nu} F^{\mu\nu}. \tag{2.6}$$

By combining this Lagrangian, which fully captures the classical electromagnetic field *in vacuo*, with the Dirac Lagrangian for a relativistic fermion, we have arrived at a result which describes the field theory of quantum electrodynamics:

$$\mathcal{L}_{\text{QED}} = \mathcal{L}_{\text{Dirac}} + \mathcal{L}_{\text{Maxwell}} = i\bar{\psi}\gamma^\mu D_\mu \bar{\psi} - m\psi\bar{\psi} - \frac{1}{4} F_{\mu\nu} F^{\mu\nu}. \tag{2.7}$$

Expanding out the first term in this new Lagrangian, we now obtain the original first term in $\mathcal{L}_{\text{Dirac}}$, and an additional $-ie\bar{\psi}\gamma^\mu A_\mu \psi$ term. This additional term corresponds to an interaction vertex between two fermion fields and the new A_μ gauge field, and is characterised by the e scale of charge. Such an interaction is shown in the Feynman diagram in Figure 2.1, and it represents the most basic QED vertex. This is perhaps the most profound concept in the Standard Model of particle physics: *by requiring the invariance of Lagrangians under specific continuous symmetry transformations, new gauge interactions arise naturally.* This QED Lagrangian characterises the physics of interactions between charged matter particles, and the physical realisation of the new gauge field is the photon.

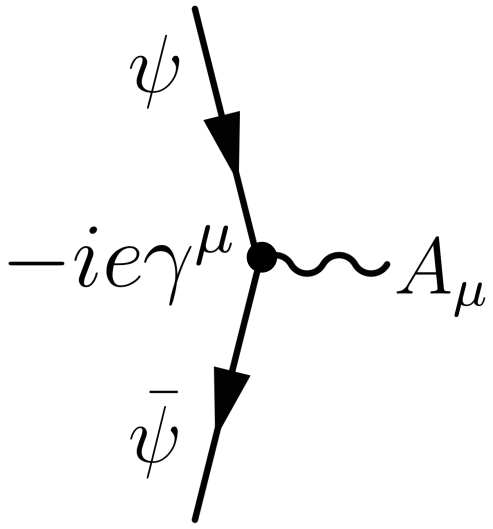


Figure 2.1: Feynman diagram corresponding to the interaction vertex obtained from the $-ie\bar{\psi}\gamma^\mu A_\mu\psi$-term in the QED Lagrangian.

2.2.2 Quantum Chromodynamics (QCD)

QED is a local $U(1)$ gauge theory, where two local $U(1)$ transformations also commute; thus it is said to be an *Abelian* theory, and, as such, is the simplest of the gauge field theories realised in nature. An important physical consequence of Abelian gauge field theories is that the charge of the group does not act on the new gauge fields which are generated by imposing the symmetry requirement. In other words, in QED the photons (the gauge bosons) are not electromagnetically charged and as such do not directly interact with each other electromagnetically; this is of course not the case for e.g. electrons in QED. The theory of the strong interactions is described by Quantum Chromodynamics (QCD), and is similar in structure to QED. However, there are two important differences: 1) the QCD Lagrangian is invariant under a local $SU(3)$ transformation, giving rise to new $SU(3)$ gauge fields; 2) the theory is *non-Abelian*[1]. Consequently, the gauge bosons introduced through the local $SU(3)$ invariance of QCD *carry a group charge*, and as such they interact with each other. Specifically, we say that QCD possesses an $SU(3)_C$ symmetry (or 'colour' symmetry) and the gluon gauge bosons have a colour charge. The fermions of the theory, quarks, also interact through this same colour charge. Akin to the positive and negative e charge of QED, in QCD there are three colour charges: red, green, and blue. Anti-quarks possess anti-colour charges. The quarks are described by fermion fields, ψ, which transform in

[1]Non-Abelian gauge theories belong to a specific class of results referred to as *Yang-Mills theories.*

3-dimensional representations of $SU(3)_C$. The QCD Lagrangian, including a summation over the six quark flavours (up, down, strange, charm, bottom, and top) is thus defined:

$$\mathcal{L}_{\text{QCD}} = -\frac{1}{4}G^a_{\mu\nu}G^{\mu\nu}_a + \sum_{f\in\{u,d,s,c,b,t\}}(\bar{\psi}(i\gamma^\mu D_\mu - m_f)\psi), \tag{2.8}$$

where the multiple gluon fields are labelled by the a index, f corresponds to a particular quark flavour, and $G^a_{\mu\nu}$ is the gluon field strength tensor. The QCD covariant derivative has the form:

$$D_\mu = \partial_\mu + ig_s t_a A^a_\mu, \tag{2.9}$$

and the field strength tensor is described by

$$G^a_{\mu\nu} = \partial_\mu A^a_\nu - \partial_\nu A^a_\mu - g_s f^{abc} A^b_\mu A^c_\nu. \tag{2.10}$$

Here g_s is the strong interaction coupling strength, and the A^a_μ are vector fields introduced to ensure the local $SU(3)_C$ invariance of the Lagrangian. The t_a are the *generators* of the $SU(3)$ group, and these directly enter the local $SU(3)$ symmetry transformation:

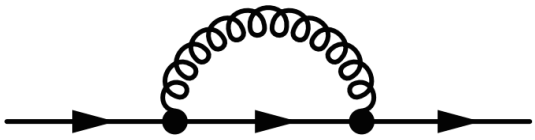
$$\psi(x) \to \psi'(x) = e^{it^a\alpha_a(x)}\psi(x), \tag{2.11}$$

which leaves $\mathcal{L}_{\text{QCD}}$ invariant. In $SU(3)$ groups, there are 8 such generators and consequently 8 gluon fields associated with the theory. Physically these correspond to 8 coloured gluon states, justifying the need for 8 gluons which was taken for granted in Section 2.1.2. These generators also define the $SU(3)$ *Lie algebra*, through the commutation relation:

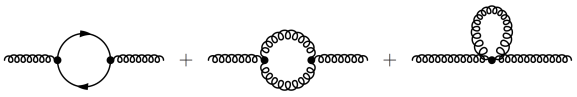
$$[t^a, t^b] = if^{abc}t^c. \tag{2.12}$$

The f^{abc} are the $SU(3)$ structure constants, and manifestly f^{abc} is an antisymmetric tensor. It is worth noting that, in Abelian gauge theories, these structure constants are 0, and consequently the $A^b_\mu A^c_\nu$ cross-term in $G^a_{\mu\nu}$ disappears; what remains is something akin to the QED field strength tensor. There also exists a 3-dimensional representation of these t^a generators, given by $\frac{\lambda^a}{2}$, where the λ^a are the 8 3×3 *Gell-Mann matrices* [10], which extend the original Pauli matrices to dimension-3.

One vital feature of QCD (and QED) for the physics described in this thesis is that the g_s coupling varies with respect to the energy scale of interest (the coupling is said to 'run'). The



(a) 1-loop quark self-energy correction.



(b) 1-loop gluon self-energy correction.

Figure 2.2: The leading order loop corrections for the quark and gluon self-energies, where the effects of the vacuum polarisation are included for the gluon loop. These are important 1-loop corrections to the quark and gluon propagators that contribute to the running of the strong coupling.

g_s characterises the QCD interaction strength at the vertex; consequently the actual coupling strength varies as g_s^2 (at tree-level), and we usually write $\alpha_s = g_s^2/4\pi$. The value of α_s can be calculated at the quark and gluon propagator loop level by considering contributions from both the quark colour interaction and gluon colour interaction loops, shown in Figure 2.2[2].

At the 1-loop level, α_s runs with energy scale Q according to:

$$\alpha_s(Q^2) = \frac{\alpha_s(\mu^2)}{1 + \beta_0 \alpha_s(\mu^2) \ln(Q^2/\mu^2)}, \tag{2.13}$$

where

$$\beta_0 = \frac{11N_C - 2N_f}{12\pi}. \tag{2.14}$$

This relates the strong coupling at the energy scale of interest Q to the value of the coupling at some fixed scale μ, often referred to as the *renormalisation* scale; it is usual practice to take $\mu = m_Z$, the Z boson mass pole. In the SM the number of colours and quark flavours are $N_C = 3$ and $N_f = 6$ respectively, and hence β_0 is positive. Therefore, α_s is predicted to decrease with Q^2; this has been experimentally confirmed in a number of particle physics experiments, as shown in Figure 2.3.

As $Q^2 \to \infty$, $\alpha_s \to 0$. This vanishing coupling gives rise to *asymptotic freedom*, an important physical consequence of QCD, where at sufficiently high energies (such as in the early universe) the QCD coupling is small enough for quark and gluons to behave like free particles. A physical

[2]There are additional contributions from the quark-gluon vertices, but here we just illustrate the leading order quark and gluon self-energy corrections.

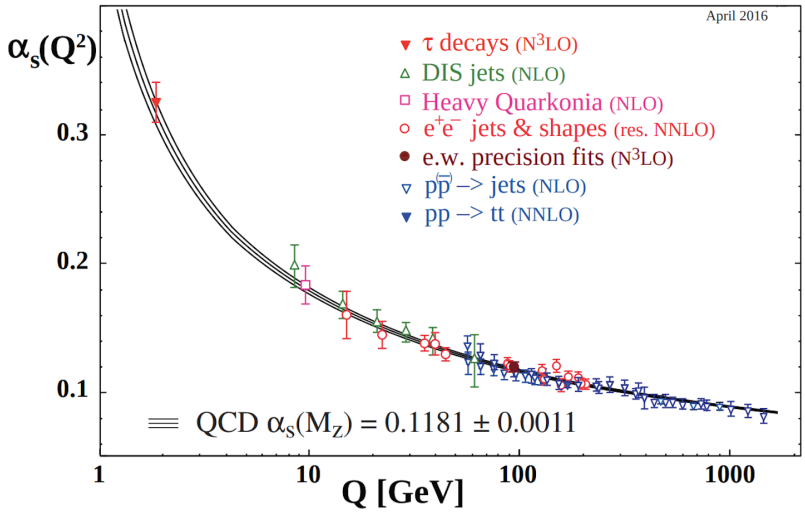


Figure 2.3: A summary of different measurements of the running of the strong coupling with energy scale Q. Results are shown for deep inelastic scattering experiments, electron-positron annihilation experiments, as well as pp and $p\bar{p}$ collider experiments. From [6].

realisation of this is the creation of a plasma of free quarks and gluons, a state of matter which can be created and studied in heavy-ion collision experiments, such as the ALICE experiment [11]. Based on coupling strength, QCD can now be split into perturbative and non-perturbative regimes. At higher energies and lower α_s, QCD theory is perturbative, and thus predictive calculations can be derived analytically. At $Q^2 \sim 1$ GeV^2 perturbative QCD breaks down and we enter a non-perturbative regime (set by some QCD energy scale, $\Lambda_{\text{QCD}} \sim 200$ MeV [12]) where analytical calculation can no longer be performed. This has huge implications for the physics of *hadronic jets*, and the lower-energy process of *hadronisation*, which we will introduce in Chapter 5. For completeness, it is worth noting that the equivalent β_0 arising from 1-loop corrections to the QED photon propagator is negative. Therefore, the QED coupling increases with energy and the theory contains no asymptotic freedom. The fact that different couplings run with energy has important consequences for treatments of *beyond* Standard Model physics, which shall be discussed in the next chapter of this thesis.

2.2.3 Electroweak Theory

The electroweak theory is an example of the unification of two seemingly disparate laws of nature: the electromagnetic theory of QED and the theory of weak interactions, responsible for nuclear processes such as β-decay. Having discussed QED in some detail, we shall first outline some of the key experimental observations of the theory of weak interactions in order to build

up a mathematical formulation of weak physics, before generalising this to the full electroweak theory. The weak force mediates nuclear β-decay, based on the underlying $d \to u + e^- + \bar{\nu}_e$ process. We want to describe this process using a non-Abelian gauge theory, where the gauge boson of interest should turn the down quark to an up quark at one vertex (conserving quark, and consequently baryon, number), and the electron to a neutrino at another vertex (conserving lepton number). Since one particle is being turned into another at each vertex, we clearly require a representation of this theory that contains at least two elements; a representation based on the non-Abelian $SU(2)$ Lie group achieves this, so we shall construct a weak theory using $SU(2)$. To make progress, we need to understand what the Lie algebra of $SU(2)$ is. Fortunately, because $SU(2)$ is the (double) covering group of the $SO(3)$ group describing three-dimensional rotations, the two groups have *identical* Lie algebras. Consequently, the generators of the two-dimensional $SU(2)$ representation are just the Pauli matrices used to generate $SO(3)$:

$$\sigma^1 = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, \sigma^2 = \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix}, \sigma^3 = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}, \tag{2.15}$$

with Lie algebra $[\frac{\sigma^i}{2}, \frac{\sigma^j}{2}] = i\epsilon_{ijk}\frac{\sigma^k}{2}$, where ϵ_{ijk} is the Levi-Civita tensor.

Denoting the $SU(2)$ gauge fields by W^i_μ, we may write the covariant derivative for this two-dimensional representation as:

$$D_\mu = \partial_\mu + i\frac{g}{2}W^i_\mu\sigma^i = \partial_\mu + i\frac{g}{2}\begin{pmatrix} W^3_\mu & W^1_\mu - iW^2_\mu \\ W^1_\mu + iW^2_\mu & -W^3_\mu \end{pmatrix} = \partial_\mu + i\frac{g}{2}\begin{pmatrix} W^3_\mu & \sqrt{2}W^+_\mu \\ \sqrt{2}W^-_\mu & -W^3_\mu \end{pmatrix}, \tag{2.16}$$

where we have now introduced the complex gauge fields $W^\pm_\mu = \frac{1}{\sqrt{2}}(W^1_\mu \mp iW^2_\mu)$. The significance of these new complex fields is apparent when the covariant derivative acts on the $SU(2)$ quark and lepton doublets, which for the first generation particles are denoted $\chi_q = \begin{pmatrix} u \\ d \end{pmatrix}$ and $\chi_l = \begin{pmatrix} \nu \\ e \end{pmatrix}$. Putting these doublets into the covariant derivative subset of the weak Lagrangian,

$$\mathcal{L}_{\text{weak}} \supset i\chi_l\gamma^\mu D_\mu\chi_l + i\chi_q\gamma^\mu D_\mu\chi_q, \tag{2.17}$$

generates interactions of the form $-ig\sqrt{2}\bar{\nu}\gamma^\mu W^+_\mu e^-$. This is a very satisfying result, and we have now recovered the desired weak interactions where the quark and lepton transformations are achieved via the exchange of W^+ and W^- gauge bosons, where the $\pm$ signifies electromagnetic charge. Starting from the foreknowledge of β-decay, and underlying $SU(2)$ symmetry, the

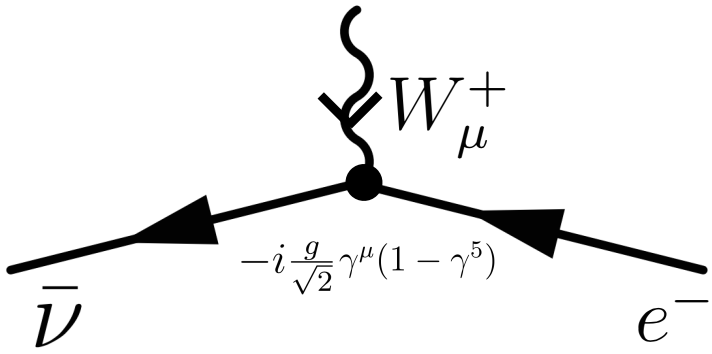


Figure 2.4: The weak interaction vertex for electron and electron-antineutrino transitions. The factor of $(1-\gamma^5)$ arises from the need to project out the left-handed components of the particle wavefunctions, since the weak interaction does not conserve parity.

weak interaction vertices have followed straightforwardly. However, we have neglected one key empirical observation thus far: parity violation. In 1957 it was observed by Madam Wu *et al* [13] that weak interactions do not conserve parity, and therefore $\mathcal{L}_{\text{weak}}$ is not invariant under the spatial inversion $x \to -x$. It is possible to write down Lagrangians which are both Lorentz invariant and parity violating by introducing the matrix:

$$\gamma^5 \equiv i\gamma^0\gamma^1\gamma^2\gamma^3. \tag{2.18}$$

The projection operators, $P_{L,R} = \frac{1}{2}(1 \mp \gamma^5)$, act on fermion fields, ψ, to generate the left- and right-handed components $P_{L,R}\psi = \psi_{L,R}$. Right-handed fermions have directions of spin pointing in the same direction as their motion, whilst left-handed fermions have their direction of spin pointing in the opposite direction to their motion. In the parity-violating weak interaction, only left-handed fermions and right-handed anti-fermions couple to the W gauge bosons. This feature can be added to the theory by using the P_L projection factor in the flavour-changing W vertex. Now the parity-violating weak Lagrangian contains interactions of the form $-i\frac{g}{\sqrt{2}}\bar{\nu}\gamma^\mu(1-\gamma^5)W^+_\mu e^-$, as shown in Figure 2.4 for case of electrons and electron-antineutrinos. We have now formulated an $SU(2)$ theory of weak interactions, where the fermion fields are charged under *weak isospin*, $I_3 = 0, \pm\frac{1}{2}$ (corresponding to the eigenvalues of the third $SU(2)$ generator). So far, we have derived the $W^\pm_\mu$ vertices which give rise to β-decay, but there are

also vertices which involve the W_μ^3. Naively one might assume that this field corresponds to the physical Z boson, however the Z *does* couple to right-handed fermions. Progress can be made by recalling that the photon, a gauge boson arising from a $U(1)$ theory, also couples to both left- and right-handed fermions. It is therefore possible to generate the Z field by considering a mixture of the W_μ^3 field with a second $U(1)$ boson (denoted B_μ) which couples to both left- and right-handed fermions. We have now arrived at the first real indication that the $SU(2)$ weak theory and $U(1)$ electromagnetic theory can be combined into one *electroweak* $SU(2) \times U(1)$ theory.

First, put the left-handed quark and leptons into $SU(2)$ doublets $\chi_{q,L} = \begin{pmatrix} u_L \\ d_L \end{pmatrix}$ and $\chi_{l,L} = \begin{pmatrix} \nu_L \\ e_L \end{pmatrix}$, with coupling g. We will now also endow these doublets with a new charge, called *weak hypercharge*, $Y_{q,l}$ arising from an additional $U(1)$ transformation with a B_μ gauge field and a coupling constant g'. The right-handed fermions are $SU(2)$ singlets (because they do not transform under the symmetry) u_R, d_R, and e_R, with weak hypercharges $Y_{u,d,e}$. The W_μ^3 and B_μ fields now mix to produce the physical Z_μ and A_μ fields of the Z boson and photon in the following manner:

$$\begin{pmatrix} W_\mu^3 \\ B_\mu \end{pmatrix} = \begin{pmatrix} \cos\theta_W & \sin\theta_W \\ -\sin\theta_W & \cos\theta_W \end{pmatrix} \begin{pmatrix} Z_\mu \\ A_\mu \end{pmatrix}, \tag{2.19}$$

where the *Weinberg angle*, θ_W has been measured by, among other experiments, the LEP experiments with value $\sin^2\theta_W \approx 0.23$ [14].

With the mixing of the W_μ^3 and B_μ fields established, the covariant derivatives can be used to deduce the values of the charges in electroweak theory. Since only the B_μ is associated with the right-handed fermions, these particles must contribute the following covariant derivative to the electroweak Lagrangian:

$$\mathcal{L}_{\text{electroweak}} \supset -\bar{\psi}_R g' Y_\psi \gamma^\mu B_\mu \psi_R \supset -\bar{\psi}_R g' \cos\theta_W Y_\psi \gamma^\mu A_\mu \psi_R. \tag{2.20}$$

By comparing this with the QED Lagrangian, it follows immediately that $|e| = g' \cos\theta_W$ and the Y_ψ corresponds to the electric charge: $Y_e = -1$, $Y_u = +\frac{2}{3}$, and $Y_d = -\frac{1}{3}$, again confirming what was taken from granted in Table 2.1. Similarly, the covariant derivative for the left-handed fermions will contain the piece

$$\mathcal{L}_{\text{electroweak}} \supset -\bar{\psi}_L(\frac{g}{2}\sigma^3\gamma^\mu W^3_\mu + g'Y_\psi\gamma^\mu B_\mu)\psi_L \supset -\bar{\psi}_L(\frac{g}{2}\sin\theta_W\sigma^3 + g'\cos\theta_W Y_\psi)\gamma^\mu A_\mu\psi_L. \quad (2.21)$$

This result, and the fact that the $\chi_{q,L}$ and $\chi_{l,L}$ doublets contain states which only differ in electric charge by one unit of e, leads to $|e| = g\sin\theta_W$, $Y_q = +\frac{1}{6}$, and $Y_l = -\frac{1}{2}$. Thus we have now arrived at a unified electroweak theory, satisfying the experimentally observed β-decay of the weak interaction, the parity violation of the weak interaction, the parity conservation of interactions associated with the Z boson and photon, and the associated weak isospins and weak hypercharges. Denoting the left-handed weak interaction symmetry $SU(2)_L$ and the additional B_μ field hypercharge symmetry $U(1)_Y$, we can summarise the group structure of electroweak theory as $SU(2)_L \times U(1)_Y$. Including the QCD colour symmetry, the SM can now be summarised as a gauge field theory operating under $SU(3)_C \times SU(2)_L \times U(1)_Y$ symmetry.

We are now almost in a position to formulate the SM Lagrangian, but two crucial details have been neglected. Firstly, because left- and right-handed fermions have been placed into different representations of $SU(2)_L \times U(1)_Y$, we are prevented from writing a mass term for fermions. Since electrons clearly have a mass, this is a problem. The invariance of these Lagrangians under the various symmetries (or gauge invariance) also forbids the gauge bosons in the theory from possessing a mass. This is fine for the case of QED, since the photon is massless and electromagnetism is a long-range force. However, the weak interaction is a short-range force, due to the relatively large masses of the W and Z bosons which have been experimentally measured. Again, we have arrived at the same mass problem for bosons − how can these masses be recovered without breaking the gauge invariance of the theory? The solution is referred to as the *Higgs mechanism*, originally introduced in [2–4].

2.3 The Higgs Mechanism

The issue of generating boson and fermion masses without breaking the gauge invariance of the electroweak theory was solved through an elegant approach introduced by Brout, Englert, Guralnik, Hagen, Higgs, and Kibble in a series of independent papers published around 1964, most commonly referred to as the *Higgs mechanism*. The idea was to spontaneously break the electroweak symmetry $SU(2)_L \times U(1)_Y \to U(1)_{EM}$, where $U(1)_{EM}$ is the local electromagnetic

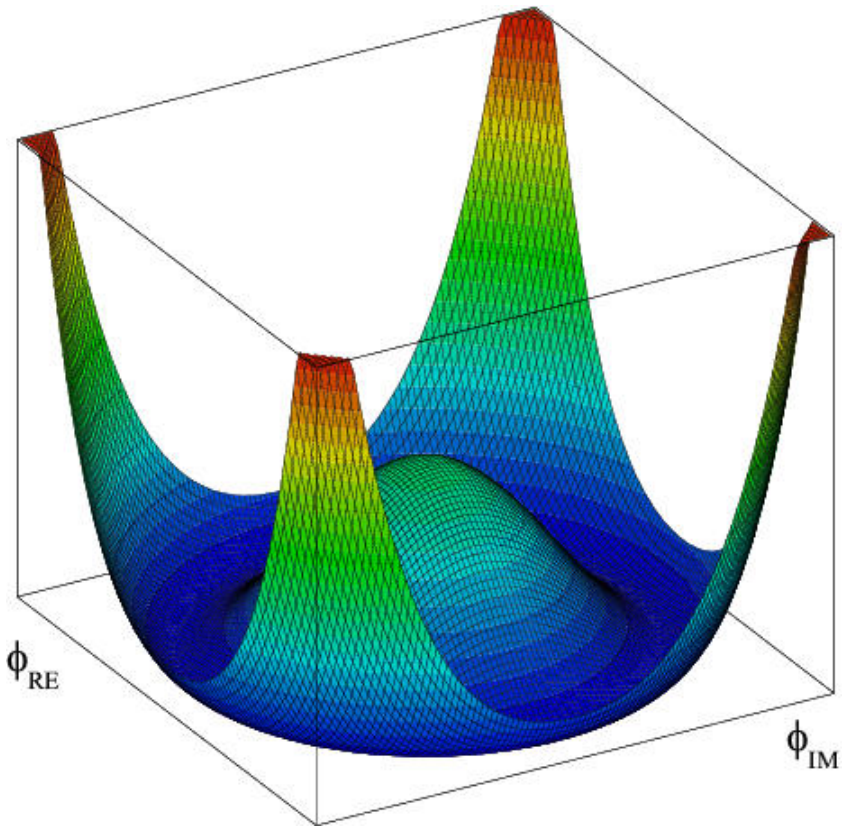


Figure 2.5: The 'Mexican hat' Higgs potential in the complex plane, after spontaneous symmetry breaking where the vacuum expectation value of the Higgs field is generated. The circle traces out the minimum of the potential, with radius related to the magnitude of the Higgs vacuum expectation value. In this complex plane there are both radial and circular degrees of freedom, again generated by spontaneous symmetry breaking.

$U(1)$ symmetry of QED, by introducing a new complex scalar $SU(2)$ doublet, Φ, with weak hypercharge 1 and weak isospin $\frac{1}{2}$. The corresponding Lagrangian is

$$\mathcal{L}_{\text{Higgs}} = (D^{\mu}\Phi)^{\dagger}(D_{\mu}\Phi) + y_f(\bar{\chi}_{l,L}\Phi e_R + \bar{e}_R\Phi^{\dagger}\chi_{l,L}) + V(\Phi^{\dagger}\Phi), \tag{2.22}$$

where the first component contains the electroweak covariant derivative, D_{μ}, and therefore describes the couplings between the new complex scalar and the electroweak gauge bosons. The second term describes the coupling between the new scalar and the fermions. Finally, the third term is the Higgs potential

$$V(\Phi^{\dagger}\Phi) = \mu^2\Phi^{\dagger}\Phi + \lambda(\Phi^{\dagger}\Phi)^2. \tag{2.23}$$

2.3.1 Electroweak Symmetry Breaking (EWSB)

For $\mu^2 < 0$, the Higgs potential will have minima (vacuum expectation values) at $|\Phi| = \frac{v}{2}$, where $v = \sqrt{\frac{-\mu^2}{\lambda}}$. This symmetry can now be spontaneously broken by choosing a vacuum value $\frac{1}{\sqrt{2}}\begin{pmatrix}0\\v\end{pmatrix}$ and expanding the Φ about its radial and (three) circular degrees of freedom. This gives rise to the famous 'Mexican hat' potential, shown in the complex plane of the Φ field in Figure 2.5. We may now write the general expanded Φ as:

$$\Phi(x) = \frac{1}{\sqrt{2}} e^{-iG_a(x)\frac{\sigma^a}{2}} \begin{pmatrix}0\\v+h(x)\end{pmatrix}, \tag{2.24}$$

where the $G_a(x)$ label the three circular degrees of freedom, the σ^a are the usual $SU(2)$ generators, and the $h(x)$ corresponds to the expanded radial degree of freedom. By substituting this Φ field into the covariant derivative term of the Higgs Lagrangian, and assuming that the W^3_μ and B_μ fields mix to produce the correct Z_μ and A_μ fields, we may now write the covariant derivative term of the Higgs Lagrangian in the following expanded form:

$$\begin{aligned}\mathcal{L}_{\text{Higgs}} \supset (D^\mu\Phi)^\dagger(D_\mu\Phi) = \frac{1}{4}g^2(h^2+2vh+v^2)W^+_\mu W^{-\mu} \\ + \frac{1}{8}(g^2+g'^2)(h^2+2vh+v^2)Z_\mu Z^\mu - \mu^2h^2 - \lambda v h^3 - \frac{\lambda}{4}h^4,\end{aligned} \tag{2.25}$$

where we have worked in the *Unitary gauge* which ensures that the G_a circular degrees of freedom, often referred to as Goldstone bosons, are absorbed into the longitudinal components of the W and Z boson gauge fields. We have now generated mass terms for the $W^\pm$ fields, the Z field, and the Higgs itself, which can be read off from the Lagrangian:

$$m_W = \frac{1}{2}gv, \tag{2.26}$$

$$m_Z = \frac{1}{2}\sqrt{g^2+g'^2}v = \frac{m_W}{\cos\theta_W}, \tag{2.27}$$

$$m_H = \sqrt{2}|\mu| = v\sqrt{2\lambda}. \tag{2.28}$$

Notice that there is also no $A_\mu A^\mu$ term generated, and so we have also recovered the massless photon field: $m_\gamma = 0$. A couple of Higgs interaction terms have also been generated: $\lambda v h^3$ describing the *trilinear Higgs vertex* and $\frac{\lambda}{4}h^4$ describing the *quartic Higgs vertex*. Precise measurements of these vertices will be a vital goal of future collider physics programmes.

2.3.2 Fermonic Masses from the Higgs

Again, we may substitute the EWSB Φ field into the second term in $\mathcal{L}_{\text{Higgs}}$ to generate the fermion mass terms. Performing this substitution yields:

$$\mathcal{L}_{\text{Higgs}} \supset y_f(\bar{\chi}_{l,L}\Phi e_R + \bar{e}_R\Phi^\dagger\chi_{l,L}) = -\frac{y_f}{\sqrt{2}}(\bar{\chi}_{l,L}e_R + \bar{e}_R\chi_{l,L})(v+h), \tag{2.29}$$

where it should be recalled that the leptonic fermion mass term is split into left- and right-handed components $\bar{\chi}_{l,L}e_R + \bar{e}_R\chi_{l,L}$. Once again, the fermion masses and fermion-Higgs interactions can be read off directly from the Lagrangian. The masses for leptonic fermions f are given by:

$$m_f = \frac{y_f v}{\sqrt{2}}, \tag{2.30}$$

and the couplings by

$$y_f h, \tag{2.31}$$

where the y_f are the different leptonic fermion *Yukawa couplings*. Clearly, a fermion requires both a left- and right-handed field in order to acquire a mass through these Yukawa couplings; for this reason the Standard Model is unable to endow the neutrinos with mass if we assume that only left-handed neutrinos are present in nature. The fact that several experiments have confirmed the presence of non-zero neutrino masses is therefore the first empirical example of *beyond* Standard Model physics phenomena observed in nature.

In the above treatment, we have focused on how the masses and couplings are obtained for the leptonic fermions by using the $\chi_{l,L}$ doublet and e_R singlet. For quarks, we may obtain the fermion masses using a very similar procedure but we must be aware of one complication: in the quark sector it is possible for weak interactions to take place between different quark generations, and consequently the weak interaction eigenstates of the theory are not identical to the quark mass eigenstates. We shall start from the $SU(2)_L$ quark doublet of hypercharge $Y = \frac{1}{6}$, $\chi^i_{q,L} = \begin{pmatrix} u^i_L \\ d^i_L \end{pmatrix}$, where i labels the three quark generations (previously the i index had been implicitly implied). There are also the usual $SU(2)_L$ right-handed singlets u^i_R and d^i_R, with hypercharges of $\frac{2}{3}$ and $-\frac{1}{3}$ respectively. Working from the fermion term in the Higgs

Lagrangian it is then possible to obtain the general form for the component of the Higgs Lagrangian which generates the required quark-Higgs couplings:

$$\mathcal{L}_{\text{Higgs}} \supset -\sqrt{2}\Big[y_d^{ij}\bar{\chi}_{q,L}^{i}\Phi d_R^j + y_u^{ij}\epsilon^{\alpha,\beta}\bar{\chi}_{q,L}^{i,\alpha}\Phi^{\dagger\beta}u_R^j + \text{h.c.}\Big], \tag{2.32}$$

where i, j = 1,2,3 are generation indices and α, β = 1,2 index the $SU(2)$ representations. Appending 'h.c.' denotes adding the Hermitian conjugate of the preceding expression in the equation. The additional $\epsilon^{\alpha,\beta} = (\sigma^2)^{\alpha,\beta}$ is added to ensure that the Lagrangian is invariant. At this point it is also worth noting that the above component of the Higgs Lagrangian is invariant under transformations of charge and parity (CP invariant) if and only if the Yukawa coupling matrix elements $y_q^{ij} = (y_q^{ij})^*$.

When expanding about the Higgs vacuum, the above Lagrangian should generate the quark mass terms. We have already noted that the weak and mass eigenstates are not identical, and therefore there is some mixing in the quark sector giving rise to non-diagonal y_q^{ij}. The required mass terms will therefore follow from diagonalising these y_q^{ij}:

$$y_u = U_u\mathcal{Y}_u S_u^\dagger, \quad y_d = U_d\mathcal{Y}_d S_d^\dagger, \tag{2.33}$$

where U_u and U_d are 3×3 unitary matrices, and the $\mathcal{Y}$ are diagonal. The quark fields transform according to

$$u_L \to U_u u_L, \quad d_L \to U_d d_L, \quad u_R \to S_u u_R, \quad d_R \to S_d d_R. \tag{2.34}$$

Under these transformations, the terms on the Higgs Lagrangian transform in the following manner:

$$y_d^{ij}\bar{\chi}_{q,L}^{i}\Phi d_R^j \to \bar{\chi}_{q,L}\Phi U_d\mathcal{Y}_d S_d^\dagger S_d d_R = \bar{\chi}_{q,L}\Phi U_d\mathcal{Y}_d d_R. \tag{2.35}$$

The quark masses can now be generated by setting the Φ field to its vacuum expectation value $\frac{1}{\sqrt{2}}\begin{pmatrix}0\\v\end{pmatrix}$. The generates the quark mass term in the Higgs Lagrangian, namely

$$\mathcal{L}_{\text{Higgs}} \supset -v\mathcal{Y}_d^{ij}\bar{d}_L^i d_R^j - v\mathcal{Y}_u^{ij}\bar{u}_L^i u_R^j + \text{h.c.} = -m_d^i\bar{d}_L^i d_R^j - m_u^i\bar{u}_L^i u_R^j + \text{h.c.} \tag{2.36}$$

With the quark masses now defined, the interaction terms between the quarks also transform according to, for example,

$$\bar{u}_L^i \gamma^\mu d_L^i \rightarrow \bar{u}_L^i \gamma^\mu (U_u^\dagger U_d)^{ij} d_L^j. \tag{2.37}$$

Comparing this with the standard vertices (such as those which arise in QED) it is clear that the weak interaction between quarks, which involves a flavour change, can be written in terms of these U matrices. In fact the weak eigenstates can clearly be expressed as a linear combination of the mass eigenstates. The mixing matrix used to achieve this combination is given by

$$U_u^\dagger U_d = V_{\mathrm{CKM}} = \begin{pmatrix} V_{ud} & V_{us} & V_{ub} \\ V_{cd} & V_{cs} & V_{cb} \\ V_{td} & V_{ts} & V_{tb} \end{pmatrix} \tag{2.38}$$

Omitting the complex phase, the inter- and outer-generational V_{CKM} matrix elements have been measured to be [6]:

$$\begin{pmatrix} |V_{ud}| & |V_{us}| & |V_{ub}| \\ |V_{cd}| & |V_{cs}| & |V_{cb}| \\ |V_{td}| & |V_{ts}| & |V_{tb}| \end{pmatrix} = \begin{pmatrix} 0.97420 \pm 0.00021 & 0.2243 \pm 0.0005 & (3.94 \pm 0.36) \times 10^{-3} \\ 0.218 \pm 0.004 & 0.997 \pm 0.017 & (42.2 \pm 0.8) \times 10^{-3} \\ (8.1 \pm 0.5) \times 10^{-3} & (39.4 \pm 2.3) \times 10^{-3} & 1.019 \pm 0.025 \end{pmatrix}. \tag{2.39}$$

V_{CKM} is referred to as the Cabibbo-Kobyashi-Maskawa (CKM) matrix, where the matrix elements must be measured experimentally and the element $(V_{\mathrm{CKM}})_{ij}$ quantifies the probability of a weak flavour changing interaction between the quark flavours i and j. The CKM matrix is approximately diagonal: its diagonal elements have magnitudes close to unity, with small off-diagonal values. The smallest $(V_{\mathrm{CKM}})_{ij}$ values are measured for flavours i and j which are separated by two generations, demonstrating that flavour-changing weak charged current interactions become more suppressed as the generation gap between the quark flavours increases. Another important feature of this matrix is that it contains matrix elements which have a non-zero complex phase, δ, giving rise to CP violation which has been experimentally observed in the quark sector.

The idea of flavour mixing described here is also relevant for the lepton sector as well. Previous measurements of the flux of neutrinos emitted from the Sun, or those measured in the atmosphere, indicate that a neutrino of a particular flavour may transmute into one of another flavour as it propagates. The first indication of such a phenomenon was inferred from measurements of an electron neutrino deficit on Earth with respect to the expected flux from the Sun. Similar flavour mixings have been observed between other neutrino generations, and, as

already mentioned, these *neutrino oscillations* cannot be explained by the SM as described thus far because they lead to the important physical consequence that neutrinos are not massless. However, these oscillations can be quantified by constructing a mixing matrix to account for the difference between the weak eigenstates and mass eigenstates in the neutrino sector. Like the CKM matrix, this is new matrix is also a 3×3 unitary matrix, and contemporary neutrino experiments are working to assess the precise values of its matrix elements. One important outstanding question is whether or not there is CP violation in the neutrino sector; measuring a non-zero complex phase in the matrix elements of this so-called Pontecorvo-Maki-Nakagawa-Sakata (PMNS) neutrino mixing matrix will provide direct evidence that the neutrino sector indeed conspires to violate charge-parity symmetry.

2.4 Putting it all together: the SM Lagrangian

With the Higgs mechanism for bosonic and fermionic masses outlined, it is now possible to summarise the SM Lagrangian as arising from the complete set of particle interactions predicted by the electroweak and QCD Lagrangians, together with the Higgs Lagrangian and Φ field which endows the gauge bosons, quarks, and leptons with their masses without breaking the gauge invariance of the SM. The total SM Lagrangian can now be compactly summarised as:

$$\mathcal{L}_{\mathrm{SM}} = \mathcal{L}_{\mathrm{QCD}} + \mathcal{L}_{\mathrm{electroweak}} + \mathcal{L}_{\mathrm{Higgs}}. \tag{2.40}$$

As previously discussed, the $\mathcal{L}_{\mathrm{SM}}$ can not be the complete story; it does not out-of-the-box generate the necessary neutrino mass terms, nor does it quantify the neutrino oscillations, without the requirement of additional mixing in the neutrino sector. This is a first real hint of *beyond* SM physics (BSM), and indeed the major focus of this thesis is searching for the direct production of beyond SM physics at the Large Hadron Collider. More general BSM theories will be briefly discussed in the next Chapter.

2.5 Collider Physics

Before moving to beyond SM physics, it is instructive to briefly outline the phenomenology of the particle physics expected in proton-proton colliders, such as the Large Hadron Collider (LHC), in order to define a few useful concepts which will be revisited later in this thesis.

2.5.1 Scattering Theory

The process $i \to f$ occurs at some rate $\Gamma_{i\to f}$ which can be calculated to first order using Fermi's 'Golden Rule' [15]:

$$\Gamma_{i\to f} = \frac{2\pi}{\hbar}|\langle\psi_f|\mathcal{H}|\psi_i\rangle|^2\rho, \tag{2.41}$$

where $\mathcal{H}$ is the interaction Hamiltonian for this process of interest, and ρ is the density of final states. This rate is related to the process *cross-section*, σ, where:

$$\sigma_{i\to f} = \frac{\Gamma_{i\to f}}{J}. \tag{2.42}$$

Here J is the initial particle flux in the process. The $\langle\psi_f|\mathcal{H}|\psi_i\rangle$ is the matrix element for the process, to first order in perturbation theory. Generalising this to higher orders, we may define the generalised matrix element to be $\mathcal{M}_{i\to f}$. The density of final states must be determined based on a Lorentz invariant phase-space element. Putting this together, the full Lorentz invariant cross-section is then defined as:

$$\sigma_{i\to f} = \frac{2\pi}{\hbar J}|\mathcal{M}_{i\to f}|^2\rho. \tag{2.43}$$

This formula can be used to compute the cross-section for interactions such as the production of quarks from the initial proton-proton scattering in the LHC beam, such as the process $u\bar{u} \to d\bar{d}$ where down-type quarks are produced from the scatter of the up-quarks in the protons. The cross-section for this process can be calculated to be:

$$\sigma_{u\bar{u}\to d\bar{d}} = \frac{|\mathcal{M}_{u\bar{u}\to d\bar{d}}|^2}{64\pi^2 s}\int d\Omega, \tag{2.44}$$

where s is the square of the centre-of-mass energy of the initial beam of particles and the integral is taken over the solid angle of the interaction.

2.5.2 Luminosity

The *luminosity* is the ratio of the number of events, N (such as $u\bar{u} \to d\bar{d}$ processes), measured in some time interval, and the cross-section for that process. If such a process is measured in a detector, then the measurement will be affected by the detector *efficiency* (ϵ) and *acceptance* (A). As a result, the final number of events measured in an analysis (N_{meas}) is

$$N_{\text{meas}} = \frac{N}{\epsilon A}, \tag{2.45}$$

and the luminosity is given by

$$L = \frac{N_{\text{meas}}}{\sigma}. \tag{2.46}$$

The integrated luminosity (where the integration is over time), given by $\int L dt$, provides a measure of the total number of events recorded over some interval of time, and will be discussed further in the context of the LHC and the ATLAS experiment in Chapter 4.

2.6 Summary

In this Chapter the fundamental theory of the Standard Model of physics has been reviewed, and the ideas of symmetry breaking, gauge invariance, the Higgs mechanism, and the nature of bosonic and fermionic forces and masses have been outlined. In addition the basic concepts of process rate, cross-section, and luminosity have also been presented. The SM is a hugely successful example of an effective field theory in nature; one that accurately predicts the particles and interactions present around the scale of electroweak symmetry breaking. In spite of this, it fails to account for important observations such as the non-zero mass of neutrinos. With this in mind, we shall now briefly discuss other major shortcomings of the Standard Model and how extensions to the Standard Model have been constructed to account for these limitations in the existing theory. The major BSM theory considered in this thesis is that of *supersymmetry*, and we shall briefly summarise the guiding theoretical ideas of this framework.

3

Beyond the Standard Model of Particle Physics

We now move towards a more in depth discussion of some of the major outstanding problems with the current Standard Model of Particle Physics. These include, but are not limited to, the explanation of massive neutrinos, the apparent presence of dark matter in our observable universe, the various mass hierarchies between the SM particles, the clear matter-antimatter asymmetry necessary to give rise to the universe as we know it, and ultimately the need for an eventual Grand Unification of the known particle physics fields, including gravity, into a more complete theory of fundamental physics. We shall explore these fascinating outstanding problems very briefly, before outlining the main phenomenology of the theory of *supersymmetry* [16–23], which is a theoretical extension of the SM which very naturally incorporates solutions to many of these challenging problems.

3.1 Problems with the Standard Model

3.1.1 Neutrino Masses

As reported in Table 2.1 the three experimentally observed generations of neutrino, ν_e, ν_μ, and ν_τ, undergo oscillations due to a mixing between neutrino weak eigenstates and mass eigenstates, and consequently there must exist a mass hierarchy in the neutrino sector. Indeed two possible hierarchies are allowed by nature, and these 'normal' and 'inverted' hierarchies are shown in Figure 3.1. As briefly discussed in the previous Chapter, the presence of a right-handed neutrino, ν_R, naturally allows for mass in the neutrino sector. A class of theories known as *seesaw mechanisms* [25] can also explain how suitably light lepton neutrinos can be produced,

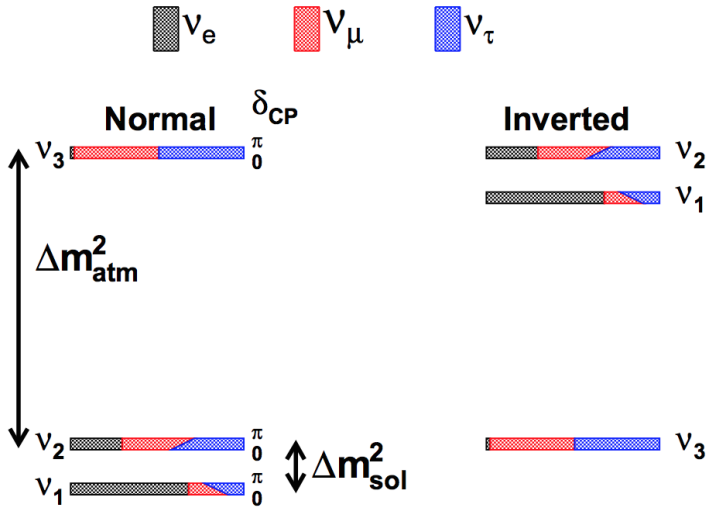


Figure 3.1: The two alternative neutrino mass hierarchies between the mass eigenstates of the neutrinos, ν_1, ν_2, and ν_3. For each mass eigenstate, the relative contributions from the weak eigenstates (ν_e, ν_μ, ν_τ) are also provided as a function of the (currently unknown) CP-violating phase, δ_{CP}. The mass splittings, as determined by solar and atmospheric neutrino experiments, are also indicated. From [24].

together with a much heavier ν_R. In the simplest of these theories, two or more right-handed neutrino fields are introduced which do not interact electroweakly. One may write down a neutrino mass matrix made up of two mass scales: the electroweak scale mass, $m_{\mathrm{electroweak}} \sim 10^2$ GeV, and the Grand Unification Theory (GUT) scale mass $m_{\mathrm{GUT}} \sim 10^{16}$ GeV which is the relevant mass scale for the new physics. For a mass matrix of the form,

$$\begin{pmatrix} 0 & m_{\mathrm{electroweak}} \\ m_{\mathrm{electroweak}} & m_{\mathrm{GUT}} \end{pmatrix}, \tag{3.1}$$

the two eigenvalues are approximately $|\lambda_1| \approx m_{\mathrm{GUT}}$ and $|\lambda_2| \approx \frac{m^2_{\mathrm{electroweak}}}{m_{\mathrm{GUT}}}$, corresponding to the masses of the heavy right-handed neutrinos and the light lepton neutrinos respectively. This provides the necessary lighter neutrino masses around or below 1 eV. The right-handed neutrinos are the new physics particles for the GUT-scale, which corresponds to the energy scale where the running electroweak and strong couplings are comparable in size and might well unify into a single coupling corresponding to a strong-electroweak interaction. A popular GUT theory is the $SO(10)$ extension to the SM, of which $SU(3)_C \times SU(2)_L \times U(1)_Y$ is a subgroup. 15 states of a single SM generation fit nicely into the 16-dimensional representation of $SO(10)$, where the 16th state is an SM gauge singlet, happily playing the role of a right-handed neutrino.

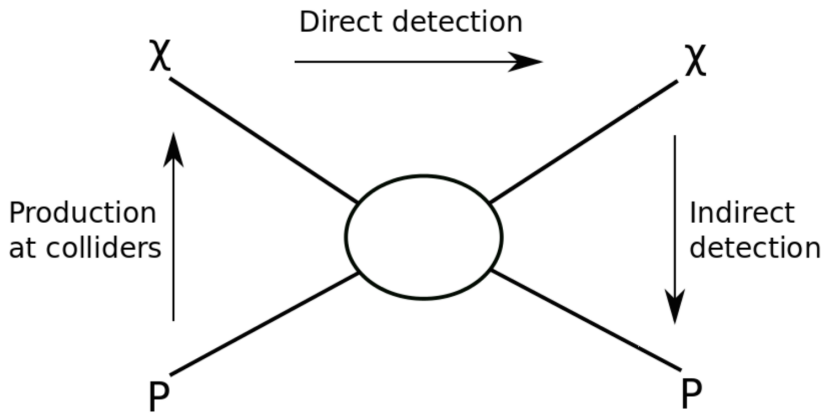


Figure 3.2: The three complementary approaches to dark matter searches. LHC experiments such as ATLAS focus on the production of dark matter from collisions of known particles. From [26].

3.1.2 Dark Matter

Measurements of the rate of the rotation of galactic spirals [27], and the gravitational lensing studies of the fluid motion of gas in the two colliding galaxy clusters forming the Bullet Cluster [28] are examples of the astronomical evidence for the existence of an additional form of matter, referred to as *dark matter*. Dark matter itself must be massive and interact gravitationally but not electromagnetically; as such, we would expect particles of dark matter to be incredibly weakly interacting. The latest projections estimate that only ~ 5 % of the matter-energy content of the observable universe consists of regular SM matter, whilst dark matter alone is responsible for a ~ 25 % slice [29]. Searches for dark matter particles are one of the most active investigations in particle and astrophysics.

A class of dark matter particle of particular theoretical interest are the so-called weakly-interacting massive particles, or *WIMPs*. Such particles are assumed to interact only through gravity and the weak interaction, or other possible new interactions which produce cross-sections no higher than the weak interaction cross-section. The large mass of these WIMPs would cause them to be slow moving enough to be insufficient to overcome their mutual gravitational attraction, and consequently WIMPs would tend to clump together into large clusters of 'cold' dark matter. WIMPs in the ~ 100 GeV mass range also have the correct range of self-annihilation cross-sections to give rise to the correct amount of dark matter required after the thermal freezeout of the early universe (the so-called dark matter relic density) [30]. The three main methods of (typically WIMP) dark matter searches are shown in the schematic in Figure 3.2.

In this thesis, we explore the possibility of creating dark matter through the proton-proton collisions at the LHC. The production of dark matter is described in numerous new physics models, which produce many signatures and search strategies, and it is also predicted by many supersymmetric physics theories.

3.1.3 Hierarchies

With the Higgs mass experimentally measured, the vacuum expectation energy is calculated to be 246 GeV. This value sets the scale for electroweak symmetry breaking. One important outstanding issue throughout the SM is understanding why the electroweak scale is so far from the natural gravitational energy scale, set by the Planck mass $m_{\text{Planck}} \sim 10^{19}$ GeV. This discrepancy lies at the heart of understanding why the weak interaction force is $\sim 10^{24}$ times as strong as gravity. It shall be demonstrated that supersymmetry provides a neat and natural solution to this *Higgs mass hierarchy problem* in Section 3.2.1.

3.1.4 Matter-Antimatter Asymmetry and Baryogenesis

There is no significant experimental evidence to suggest that there exists significant concentrations of antimatter in our observable universe, and there is therefore a clear matter-antimatter asymmetry in reality. The *Sakharov conditions* [31] outline the requirements for a matter-dominated universe to be the need for:

- Baryon number violation.
- CP symmetry violation.
- The presence of particles that do not achieve thermal equilibrium.

Baryon number violation is of course required to produce an excess of baryons over anti-baryons, and the additional CP violation allows for unequal numbers of left-handed baryons and right-handed anti-baryons to be produced. The third condition is easily achieved through a rapidly-expanding universe, causing the rate of particle-antiparticle annihilation to drop. The Standard Model naturally incorporates baryogenesis through certain topological phase-transitions (such as sphaleron production) but this is insufficient to account for the present observed baryon asymmetries in nature. BSM theories, such as supersymmetry, can introduce

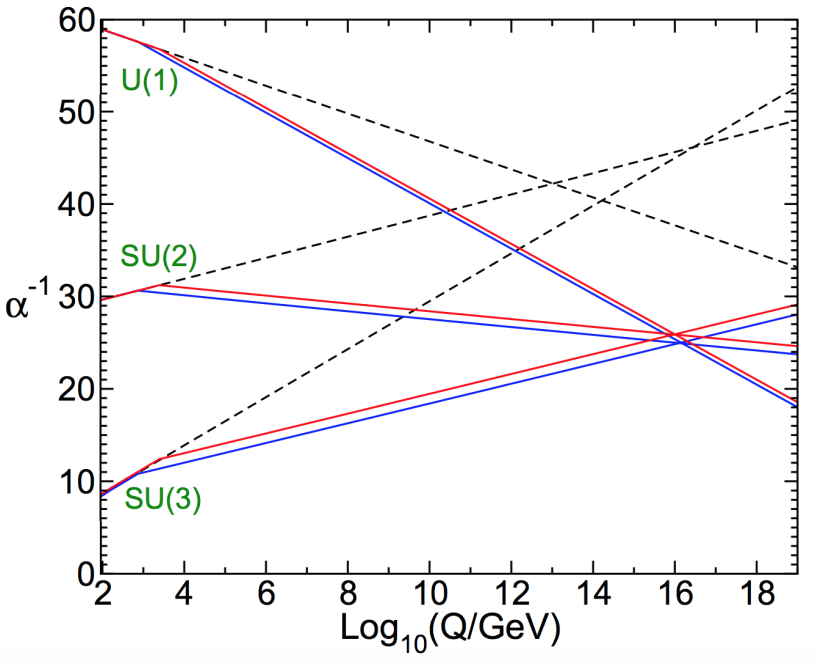
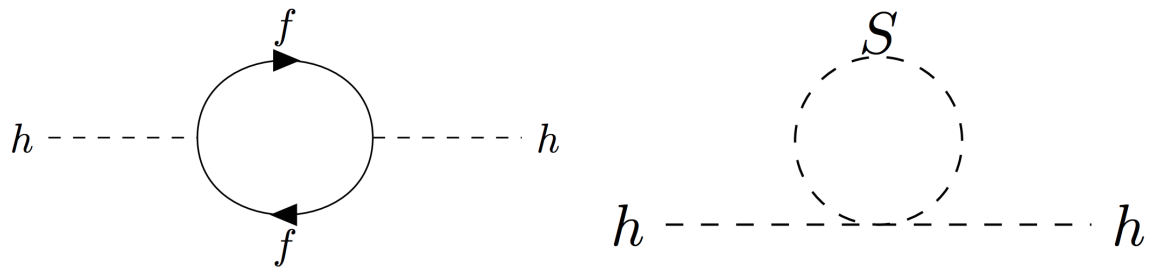


Figure 3.3: The running of the SM couplings (black, dotted lines) and running of the same couplings in a BSM extension based on supersymmetry (solid, coloured lines). In this case, the unification is achieved via the minimal supersymmteric standard model (MSSM), a supersymmetric model which we shall shortly review. The masses of new SUSY particles are varied as a common threshold between 750 GeV and 2.5 TeV, and the $\alpha_S(m_Z)$ between 0.117 and 0.120. From [22].

additional CP violating terms which can drag the predicted baryon asymmetry up to what is experimentally observed through cosmological and astronomical studies.

3.1.5 Grand Unification

The electroweak scale corresponds to the energy scale at which the theories of weak interactions and electromagnetic interactions are successfully unified into a single complete framework. One might imagine that it is also possible to unify the electroweak and strong interactions at some higher energy scale. The Standard Model is unable to provide such a unification out-of-the-box as evidenced by the nature of the running electromagnetic, weak, and strong couplings in Figure 3.3. There is no obvious energy scale where these couplings unify, if we assume the typical running of these couplings as provided by the SM. It is therefore desirable that BSM theories also unify these strong and electroweak couplings. The familiar $SO(10)$ is an example of a GUT which unifies these couplings at the GUT mass scale, $m_{\text{GUT}} \sim 10^{16}$ GeV. Theories of supersymmetry similarly offer a nice unification of the electromagnetic, weak, and strong couplings at the GUT scale. An example of the unification achieved via a SUSY model is shown in Figure 3.3.



(a) Fermion loop correction. **(b)** Scalar loop correction.

Figure 3.4: Feynman diagrams for the first-order loop corrections to the Higgs mass, coming from a fermion loop (left) and a scalar loop (right) respectively.

3.2 Supersymmetry (SUSY)

Having established the incompleteness of the Standard Model, it is natural to ask how one can make some progress in extending its theoretical framework. We have already briefly discussed the idea of extending the *internal symmetries* of the SM, such as introducing an $SO(10)$ GUT symmetry which can be spontaneously broken to yield the original $SU(3)_C \times SU(2)_L \times U(1)_Y$ SM. Another approach is to introduce additional spatial dimensions to the theory, explored for example in so-called *Kaluza-Klein* models. Finally, one may try to extend the external, space-time symmetry of the SM which corresponds to the symmetries of the Lorentz transformation and translations. This latter approach is the essence of supersymmetry (SUSY). Supersymmetry offers solutions to several shortcomings of the SM: it solves the Higgs mass hierarchy problem, unifies the gauge couplings at the GUT scale, provides suitable WIMP dark matter candidates, introduces additional CP violation necessary for baryogenesis in the early universe, and so on.

3.2.1 Higgs Mass Hierarchy Problem

One of the original motivations for SUSY was its elegant solution to the *Higgs mass hierarchy problem*, briefly described in Section 3.1.3. In the Higgs Lagrangian, the $\frac{1}{2}m_H^2 h^2$-term will receive radiative corrections from loops of fermions, f, as shown in Figure 3.4. A fermion with Yukawa coupling y_f will now provide a correction to the Higgs mass [22]:

$$\delta m_H^2 = -\frac{|y_f|^2}{8\pi^2}\Lambda_{\mathrm{UV}}^2 + \mathcal{O}(m_f^2 \ln \Lambda_{\mathrm{UV}}), \tag{3.2}$$

where Λ_{UV} is some cut-off scale which has been introduced the regularise the loop-diagram integral in the mass calculation. This should be interpreted as the scale of new physics, which

could be up to the Planck mass scale where gravity becomes relevant, above which the SM breaks down as an effective field theory. We can now see that the physical Higgs mass gains a quadratic divergence from this loop correction: how can this divergence be removed, and the experimentally verified Higgs mass be recovered? The gauge boson masses, being protected by the symmetries of the SM Lagrangian [32], do not have any quadratic divergences of this nature; can we introduce symmetries to protect the Higgs mass in a similar manner? Indeed we can: consider introducing the second diagram in Figure 3.4, where some new scalar S contributes a term $-y_S h^2 S^\dagger S$ to the Higgs Lagrangian, thus generating a correction to the Higgs mass of size[1]:

$$\delta m_H^2 = \frac{y_S}{16\pi^2}\Lambda_{\mathrm{UV}}^2 + \mathcal{O}(m_S^2 \ln \Lambda_{\mathrm{UV}}). \tag{3.3}$$

It is now possible to cancel the quadratic divergences in the Higgs mass by introducing two new scalar fields, with $y_S = |y_f|^2$. These new scalars can be introduced through an extension of the Poincaré symmetry[2] of the SM.

3.2.2 Superalgebras and Superfields

The Poincaré group is described by generators $M^{\mu\nu}$ (rotations and Lorentz boosts) and P^σ (the momentum operator, describing symmetry under translations) with algebra:

$$\Big[P^\mu, P^\nu\Big] = 0, \tag{3.4}$$

$$\Big[M^{\mu\nu}, P^\sigma\Big] = i(P^\mu \eta^{\nu\sigma} - P^\nu \eta^{\mu\sigma}), \tag{3.5}$$

$$\Big[M^{\mu\nu}, M^{\rho\sigma}\Big] = i(M^{\mu\sigma}\eta^{\nu\rho} + M^{\nu\rho}\eta^{\mu\sigma} - M^{\mu\rho}\eta^{\nu\sigma} - M^{\nu\sigma}\eta^{\mu\rho}). \tag{3.6}$$

We now wish to extend the Poincaré group. However, the famous *no-go* Coleman-Mandula theorem (1967) [33] states that, in 3+1 dimensions, there is no non-trivial mix of Poincaré and internal symmetries except for the direct product of the two. However, this restriction only applies to *bosonic* generators, and indeed it was demonstrated by Golfand and Licktman (1971) [16] that the Poincaré algebra may be extended by introducing the additional *fermionic*

[1] The change of sign for the scalar is vital for the solution to the hierarchy problem, and arises from the spin-statistics theorem.

[2] The Poincaré group describes the Lorentz and translational symmetries of spacetime: $x^\mu \to x'^\mu = \Lambda^\mu_\nu x^\nu$. The Lorentz symmetry alone is described by the *proper orthochronous group*, $SO(1,3)^+$.

spinor generators Q_α (where $\alpha = 1, 2$) which act on bosonic representations and turn them into fermionic ones (and vice-versa):

$$Q|\text{boson}\rangle \propto |\text{fermion}\rangle, \quad Q|\text{fermion}\rangle \propto |\text{boson}\rangle. \tag{3.7}$$

Adding in these additional generators extends the Poincaré group to the so-called *super-Poincaré group*, which now possesses an algebra describing Lorentz and translational symmetries, and additional supersymmetric transformations between bosons and fermions. The algebra is summarised as[3]:

$$\Big[Q_\alpha, M^{\mu\nu}\Big] = (\sigma^{\mu\nu})^\beta_\alpha Q_\beta, \tag{3.8}$$

$$\Big[Q_\alpha, P^\mu\Big] = \Big[\bar{Q}^{\dot{\alpha}}, P^\mu\Big] = 0, \tag{3.9}$$

$$\Big\{Q_\alpha, Q_\beta\Big\} = \Big\{\bar{Q}_{\dot{\alpha}}, \bar{Q}_{\dot{\beta}}\Big\} = 0, \tag{3.10}$$

$$\Big\{Q_\alpha, \bar{Q}_{\dot{\beta}}\Big\} = 2(\sigma^\mu)_{\alpha\dot{\beta}} P_\mu, \tag{3.11}$$

where $\sigma^{\mu\nu}$ are the antisymmetrised products of the Pauli matrices (see [34] for more details):

$$(\sigma^{\mu\nu})^\beta_\alpha = \frac{i}{4}\Big(\sigma^\mu\bar{\sigma}^\nu - \sigma^\nu\bar{\sigma}^\mu\Big)^\beta_\alpha. \tag{3.12}$$

The 'dot' notation we have adopted is a common shorthand for these *Weyl spinors*:

$$(\psi_\alpha)^\dagger := \bar{\psi}_{\dot{\alpha}}, \quad \bar{\psi}^{\dot{\alpha}} := \psi^*_\beta(\sigma^0)^{\beta\dot{\alpha}}. \tag{3.13}$$

In 1975 Haag, Lopuszanski, and Sohnius generalised the Coleman-Mandula theorem in order to demonstrate that the most general quantum field theories are those described by the direct product of the above super-Poincaré group algebra and internal symmetries. Having gone to some pains to describe this algebra, let us now try and understand how its representations help to solve the original hierarchy problem. The effect of acting on a standard model particle representation

[3]This algebra corresponds to $N = 1$ supersymmetry, where only one spinor generator Q_α is added to the theory. SUSY further generalises to $N > 1$ but it has been demonstrated [34] that such theories are not chiral, and therefore cannot correspond to the observable universe as we know it. They do however have interesting applications in pure mathematics and certain classes of string theory, but we shall stop at that in this thesis ...

(i.e. a SM particle wavefunction) with Q is to produce a new particle representation with identical quantum numbers except for the spin, which now differs by $\frac{1}{2}$. These new particles are referred to as *superpartners*. For any fermion f this generates two new scalar fields S_L and S_R, created from Q acting on the left- and right-handed fermion representations respectively[4]. This is precisely what is required to cancel the quadratic divergence in the Higgs mass due to the fermion f.

To write down supersymmetric Lagrangians it is necessary to generalise the notation of space and fields to *superspace* and *superfields*. We shall not discuss this in detail, but simply comment that a superfield is a mathematical object with arguments created from x^μ (regular spacetime tensors) and spinors of Grassmann numbers [34], often denoted θ^α and $\bar{\theta}_{\dot{\alpha}}$. The superfield also satisfies a specific set of transformations of its (generalised) covariant derivatives. More detail can be found in [34]. Rather than dwell on this further, we shall now briefly describe the particle content predicted by the *minimal supersymmetric Standard Model* (MSSM) [22].

3.2.3 Phenomenology of $N = 1$ SUSY

3.2.3.1 Particle Content

New supermultiplets of particles may be formed under the action of Q: these are divided into *chiral* and *vector* supermultiplets. In the language of supersymmetric Lagragians, the quantum fields corresponding to the particles in these supermultiplets are described by chiral and vector superfields respectively, which generalise the SM quantum fields to superspace. The fermionic particles of the SM fall into the chiral supermultiplet, where the left- and right-handed components of each are paired with new scalar particles. The symbols for these new fields match that of the original fermion, but with a tilde, and the naming convention is to prepend an s to the original fermion name. Thus, the electron, which has right- and left-handed components in the SM, is paired to two *selectron* scalar superparticles, $\tilde{e}_R$ and $\tilde{e}_L$ respectively.

The SM gauge bosons are associated with spin-$\frac{1}{2}$ fermions in vector supermultiplets; gauge bosons are of course massless before EWSB, and the superpartners are correspondingly spin-$\frac{1}{2}$ Weyl fermions. These new fermions are named by appending an 'ino' to the name of the original

[4]The L and R in $S_{L,R}$ label the handedness of the fermion which generates the new scalar, since scalars by definition do not have such handedness.

SM boson, and adding a tilde. Thus, the gluon has a superpartner referred to as the gluino, denoted $\tilde{g}$.

Finally, we turn to the supersymmetric extension of the Higgs field. In the SM we have a single Higgs doublet which, based on its scalar nature, should be placed into the chiral supermultiplets. However, in doing this, the SUSY transformation has the effect of generating new fermionic fields for each of the complex scalar fields in the Higgs doublet. This gives rise to a *gauge anomaly* in which new loop diagrams break the gauge invariance of the Lagrangian. Such an issue can be circumvented by imposing that the sum of the hypercharge over all of the fermions in the theory is 0; thus the original gauge anomaly may be cancelled by adding a second Higgs doublet with $Y = -1$ to the theory [22]. In this extended Higgs mechanism we introduce these separate doublets, $(H_u^+,\, H_u^0)$ and $(H_d^0,\, H_d^-)$, which generate the masses of the up- and down-type quarks respectively. Under the SUSY transformation these complex Higgs scalars now generate two additional fermion doublets, $(\tilde{H}_u^+,\, \tilde{H}_u^0)$ and $(\tilde{H}_d^0,\, \tilde{H}_d^-)$, referred to as Higgsinos. Following electroweak symmetry breaking, each of these Higgs doublets acquires a vacuum expectation value, v_u and v_d for the up- and down-type doublet respectively. These are related to the overall Higgs VEV by $v_u^2 + v_d^2 = v^2$. Another important free parameter of the SUSY theory is the ratio of these new VEVs, $\tan\beta = \frac{v_u}{v_d}$. We now have the full particle content of $N = 1$ SUSY; the resultant particles are summarised in Table 3.1.

Particles	spin-0	spin-$\frac{1}{2}$	spin-1
squarks, quarks	$(\tilde{u}_L,\, \tilde{d}_L)$	$(u_L,\, d_L)$	
($\times$ 3 generations)	$\tilde{u}_R$	u_R	
	$\tilde{d}_R$	d_R	
sleptons, leptons	$(\tilde{\nu}_e,\, \tilde{e}_L)$	$(\nu_e,\, e_L)$	
($\times$ 3 generations)	$\tilde{e}_R$	e_R	
Higgs, Higgsinos	$(H_u^+,\, H_u^0)$	$(\tilde{H}_u^+,\, \tilde{H}_u^0)$	
	$(H_d^0,\, H_d^-)$	$(\tilde{H}_d^0,\, \tilde{H}_d^-)$	
gluino, gluon		$\tilde{g}$	g
winos, W bosons		$\tilde{W}^{\pm},\, \tilde{W}^0$	$W^{\pm},\, W^0$
bino, B boson		$\tilde{B}^0$	B^0

Table 3.1: The particle content of the (minimal) supersymmetric Standard Model (MSSM), prior to electroweak symmetry breaking.

Two Higgs doublets, each consisting of two complex fields, gives rise to 8 degrees of freedom: 3 of these degrees of freedom are absorbed into longitudinal modes of the gauge bosons as before. The remaining 5 degrees of freedom give rise to 5 physical Higgs fields: 2 scalars H^0 (corresponding to the original SM Higgs) and h^0, 1 pseudoscalar A^0, and two charged scalars $H^{\pm}$. In addition, the wino, zino, and bino superpartners (collectively referred to as *gauginos*) mix with the Higgsino states: mixing with the neutral Higgsino states generates sets of *neutralinos* χ_i^0 ($i = 1, 2, 3, 4$), and mixing with charged Higgsino states produces sets of *charginos* $\chi_i^{\pm}$, where $i = 1$ corresponds to the lightest mass state, and $i = 4$ the heaviest mass state. A relatively large mixing between the third-generation sfermion mass eigenstates is also permitted by current experimental constraints on flavour-changing neutral currents and mixing in the kaon system [35].

3.2.3.2 SUSY Breaking Scale and Naturalness

We have briefly touched on SUSY particle masses. However, in the current framework, the new SUSY particles will be massless before electroweak symmetry breaking, and mass degenerate with the physical SM particles afterwards. Had such a symmetry been truly realised in nature, SUSY particles would have been experimentally observed by now. This suggests that some additional physics must exist which serves to break the current symmetry, pushing the physical SUSY particle masses to higher scales than their SM partner masses. We therefore propose that there must exist some mass scale, at which SUSY is spontaneously broken in a manner that grants higher masses to the superpartners and still ensures that they have masses in the right range to allow for the cancellation of quadratically divergent terms in the Higgs mass to all orders, thus resolving the hierarchy problem. This form of SUSY breaking is said to be *soft*. Numerous mechanisms for soft SUSY breaking exist [36–39], and are beyond the scope of this thesis, and the addition of SUSY breaking terms to the Lagrangian has the effect of introducing many new free parameters for the theory. In the MSSM-variant of SUSY breaking, the most general gauge invariant, R-parity-conserving (see Section 3.2.3.3) addition to the MSSM Lagrangian which provides soft SUSY breaking adds an additional 105 free parameters to the theory beyond those of the SM. These 105 additional parameters generate a rich and complex phenomenology in the MSSM, however there are 6 particularly important parameters in this theory which are worth highlighting. These are:

- M_1, M_2, M_3: Bino, wino, gluino masses respectively.

- μ: Higgsino mass.
- A_f: Higgs-scalar trilinear coupling for the superpartner, $\tilde{f}$, to the SM fermion f.
- $\tan\beta$: Ratio of the two Higgs doublet VEVs, $\frac{v_u}{v_d}$.

A crucial concept when dealing with models of SUSY breaking is that of *naturalness*, namely, the degree of fine-tuning required by a particular SUSY breaking model to generate the correct physical SM particle masses after radiative corrections. A 'natural' SUSY scenario is one where the fine-tuning is as mild as possible, and as such should e.g. generate the correct physical Higgs mass after cancelling the quadratically divergent terms with a minimal amount of fine-tuning required. Following soft SUSY breaking, there are several important contributions to the physical Higgs mass and the electroweak fine-tuning of the MSSM: the most crucial contributions come from constraints on Higgsinos, stops, and gluinos. At the Higgs mass one-loop level, the stop particle contribution is generated through the mixing $(\tilde{t}_R,\, \tilde{t}_L) \rightarrow (\tilde{t}_1,\, \tilde{t}_2)$. Soft SUSY breaking will then endow these stops with physical masses $m_{\tilde{t}_1}$ and $m_{\tilde{t}_2}$. The correction from these stop and top partners, in terms of the key free parameters of the MSSM, is then:

$$m_H^2 = m_Z^2 \cos^2 2\beta + \frac{3m_t^4}{2\pi^2 v^2}\Big(\ln \frac{m_{\tilde{t}_1} m_{\tilde{t}_2}}{m_t^2} + \frac{(A_f - \mu\cot\beta)^2}{m_{\tilde{t}_1} m_{\tilde{t}_2}}\Big(1 - \frac{(A_f - \mu\cot\beta)^2}{12 m_{\tilde{t}_1} m_{\tilde{t}_2}}\Big)\Big) + ... \quad (3.14)$$

From this equation it seems apparent that experimental constraints on the mass of stops are important for addressing the Higgs hierarchy problem, and to recover the correct physical Higgs mass with minimal fine-tuning it is required that the stops have $m_{\tilde{t}_1,\tilde{t}_2} \sim \mathcal{O}(10^3)$ GeV. However, this is a misleading picture of natural SUSY contraints. The stop mass also depends on the gluino mass, and therefore the gluino mass also provides a correction to the Higgs mass at the two-loop level. It is in fact the case that the most stringent naturalness upper limits, from a phenomenological point-of-view, are provided by the gluino mass. Furthermore, it has been demonstrated [40] that light stops are not a particularly generic requirement of natural SUSY, but light gluinos generally are. Constraints on the stop and gluino masses shall be discussed in much greater detail in Part III of this thesis, and the 'multi-jets' search described therein provides new sensitivities to both of those supersymmetric particles. Constrains on Higgsino masses are not determined in the LHC analysis outlined in this thesis, but it is worth mentioning here that the most robust requirement for low electroweak fine-tuning of the MSSM is that Higgsinos should be relatively light, typically with a mass upper bound around 1.0 TeV, which

is also in the range required for dark matter to be made from Higgsinos [40]. The notion of naturalness is important but should not be the overarching factor when dealing with different SUSY breaking scenarios. There is no 'good' reason for us to exist in a natural, minimally fine-tuned universe, although this notion is an aesthetically and theoretically satisfying one. If indeed our reality is realised as a minimally fine-tuned one, we might well expect to see evidence of gluino, stop, and/or Higgsino production at the LHC.

In addition to solving the hierarchy problem through naturalness, the MSSM also provides corrections to the electromagnetic, weak, and strong coupling such that the three are approximately unified at the m_{GUT} scale.

3.2.3.3 R-Parity-Violating SUSY

One important quantum number we have yet to consider in our review of SUSY phenomenology is R-parity, P_R. In the MSSM baryon and lepton number are no longer conserved by all couplings in the theory. However, experimental studies of the stability of the proton [41] also directly constrain baryon- and lepton-number violating terms in BSM theories. In SUSY we can 'force' such baryon- and lepton-violating couplings to be forbidden by imposing the conservation of the R-parity quantum number, defined as:

$$P_R = (-1)^{3(B-L)+2s}, \tag{3.15}$$

where B, L, and s are the baryon, lepton, and spin quantum numbers of the particle of interest. Under this definition, SUSY particles have $P_R = -1$ and SM particles have $P_R = +1$. In R-parity-conserving SUSY the number of sparticles is conserved at each decay vertex, and consequently there exists a lightest supersymmetric particle (LSP) which is stable. Typically, this LSP corresponds to the neutralino χ_1^0, and can be interpreted as a candidate for WIMP dark matter. Once again, we arrive at another pleasing consequence of SUSY: the RPC MSSM naturally produces a dark matter candidate which may exist in the mass range necessary to be consistent with WIMP dark matter and universal relic density constraints.

However, SUSY could in fact be an R-parity-violating (RPV) theory. RPV theories contain additional lepton- and baryon-number violating terms in the SUSY Lagrangian. Such terms are

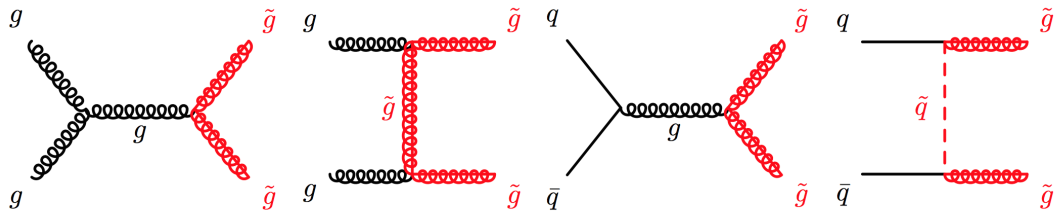


Figure 3.5: Leading-order gluino production diagrams in proton-proton collisions, where the t-channel diagram on the right is suppressed if squark masses are much greater than gluino masses.

characterised by the additional couplings λ'_{ijk} and λ''_{ijk}[5], which characterise the probabilities for lepton- and baryon-number violating decays of the LSP respectively. In such scenarios the LSP may propagate as a long-lived particle, before decaying to SM particles. The size of the decay probability determines the lifetime of the LSP, ranging from prompt decays to more long-lived scenarios. We shall consider specific examples of RPV and RPC SUSY models in Part III of this thesis.

3.3 Supersymmetric Models at Hadron Colliders

3.3.1 SUSY Production at the LHC

Given the motivation of naturalness, it is possible for the proton-proton collisions at the LHC to directly produce gluinos and stops via strong-production gluon interactions. Example leading-order production mechanisms for gluinos at the LHC are highlighted in Figure 3.5. It is also possible to directly produce other gauginos and Higgsinos via electroweak processes, however such production mechanisms are suppressed with respect to the strong-production of gluinos and stops due to their smaller production cross-sections. These cross-sections are shown in Figure 3.6.

3.3.2 Simplfied Models

In direct searches for the production of new physics, it is helpful to motivate (in part) analyses using so-called *simplified models* [44] which capture the preeminent features of the new physics. Such models are constructed from a specific signature (e.g. high multiplicities of hadronic jets produced in SUSY cascade decays), rather than from the full SUSY Lagrangian and SUSY breaking description, and can be thought of as models which capture the essence of SUSY phenomenology rather than the details and predictions of a very specific, complex model capturing all of the features of a small portion of the phase-space. Such models are typically

[5]The ijk label the quark generation indices

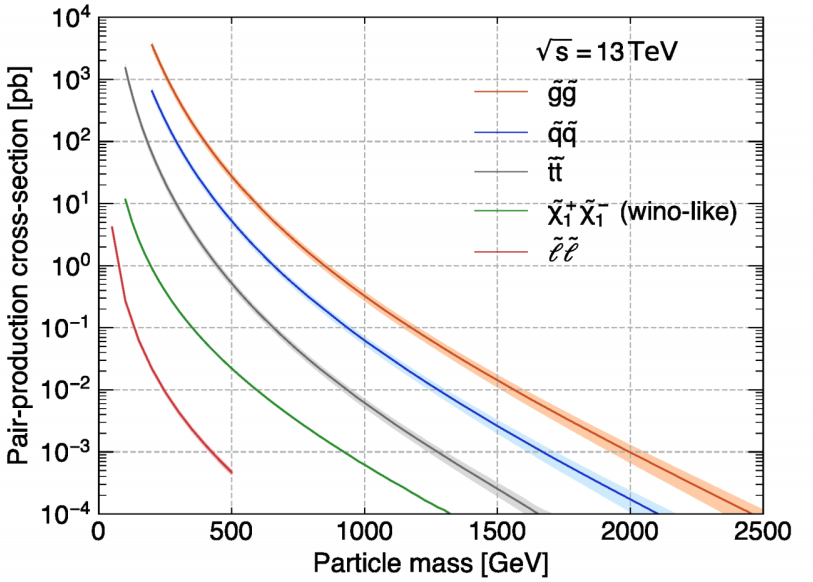


Figure 3.6: Production cross-sections for several different SUSY processes, from [42], where the gluino production cross-section assumes all squarks are decoupled. For a given particle mass, the relatively large gluino production cross-section is evident. These cross-sections were calculated using Reference [43].

parametrised in only a few variables (rather than the 105 of a full MSSM model), such as the gluino mass, LSP mass, stop mass, and branching fraction for certain decays. This simplification makes the models useful for guiding how analyses should optimise their search strategy, without being too constrained to a specific area of the SUSY parameter space. By basing the model on a small set of free parameters, it is also easier to statistically interpret the analysis results in a meaningful way. Since simplified models are, by their nature, not fully calculated models of SUSY and SUSY breaking, their statistical interpretation should be taken with a pinch of salt, but is still useful for gauging the sensitivity of the LHC to gluino and squark production. The task of performing full reinterpretations of SUSY models, such as the MSSM or related scenarios, is an important piece of work and involves input from both the experimental and theory communities. We shall revisit the simplified models considered in search for supersymmetry in final states with many hadronic jets in Part III of this thesis.

3.4 Summary

We have now completed a review of BSM physics pertinent to the analyses of this thesis, having focused in particular on the theory of supersymmetry. We have seen that SUSY offers an ideal solution to the Higgs fine-tuning problem, albeit constrained by the need for low fine-tuning in a

natural theory. We have also seen that SUSY must be a broken symmetry, and how the various mechanisms of SUSY breaking give rise to a rich phenomenology, in particular focusing on one of the simplest cases of the MSSM. Finally, we have briefly reviewed the idea of SUSY production at the LHC and how the use of simplified models can motivate certain specific signatures for SUSY, allowing LHC searches to reinterpret their results in a simplistic approach based only on a few free variables, such as sparticle masses and branching fractions. In the final section of Part I of this thesis, we shall describe the ATLAS experiment at the LHC, and focus on how common physics objects are reconstructed from detector hardware and firmware, with the end goal of data-production for analysers searching for signatures such as SUSY.

4
The ATLAS Experiment

The ATLAS ("A Toroidal LHC Apparatus") experiment [45], shown schematically in Figure 4.1, is one of the four major experiments at the Large Hadron Collider (LHC) ring [46], located at the European particle physics laboratory, CERN, near Geneva, Switzerland. In this chapter, we shall briefly review the LHC and the ATLAS experiment, including its various detector subsystems, before finishing with a summary of the different physics objects which are reconstructed in ATLAS and correspond to the central objects for all performance studies and physics analyses outlined in this thesis.

4.1 The Large Hadron Collider (LHC)

The LHC is a 27 km circumference proton synchrotron located below the CERN laboratory in the tunnel used for a Large Electron-Positron Collider (LEP), a collider experiment which ran from $1988 - 2000$. The full injector chain of the LHC is illustrated in Figure 4.2: protons are first extracted from a supply of hydrogen gas situated in an electric-field, before being accelerated to energies of 450 GeV. The Super Proton Synchrotron (SPS) [46] acts as the pre-injector for the LHC, and during Run 2 of ATLAS the protons, upon injection into the final synchrotron ring, are accelerated to beam energies of 6.5 TeV. The four major experiments − including the general-purpose ATLAS and CMS detectors [45, 47], and the LHCb detector [48]− are constructed around collision points wherein the two proton beams collide at a centre-of-mass energy of $\sqrt{s} = 13$ TeV. For Run 2 the peak luminosity of the collider has been increased to

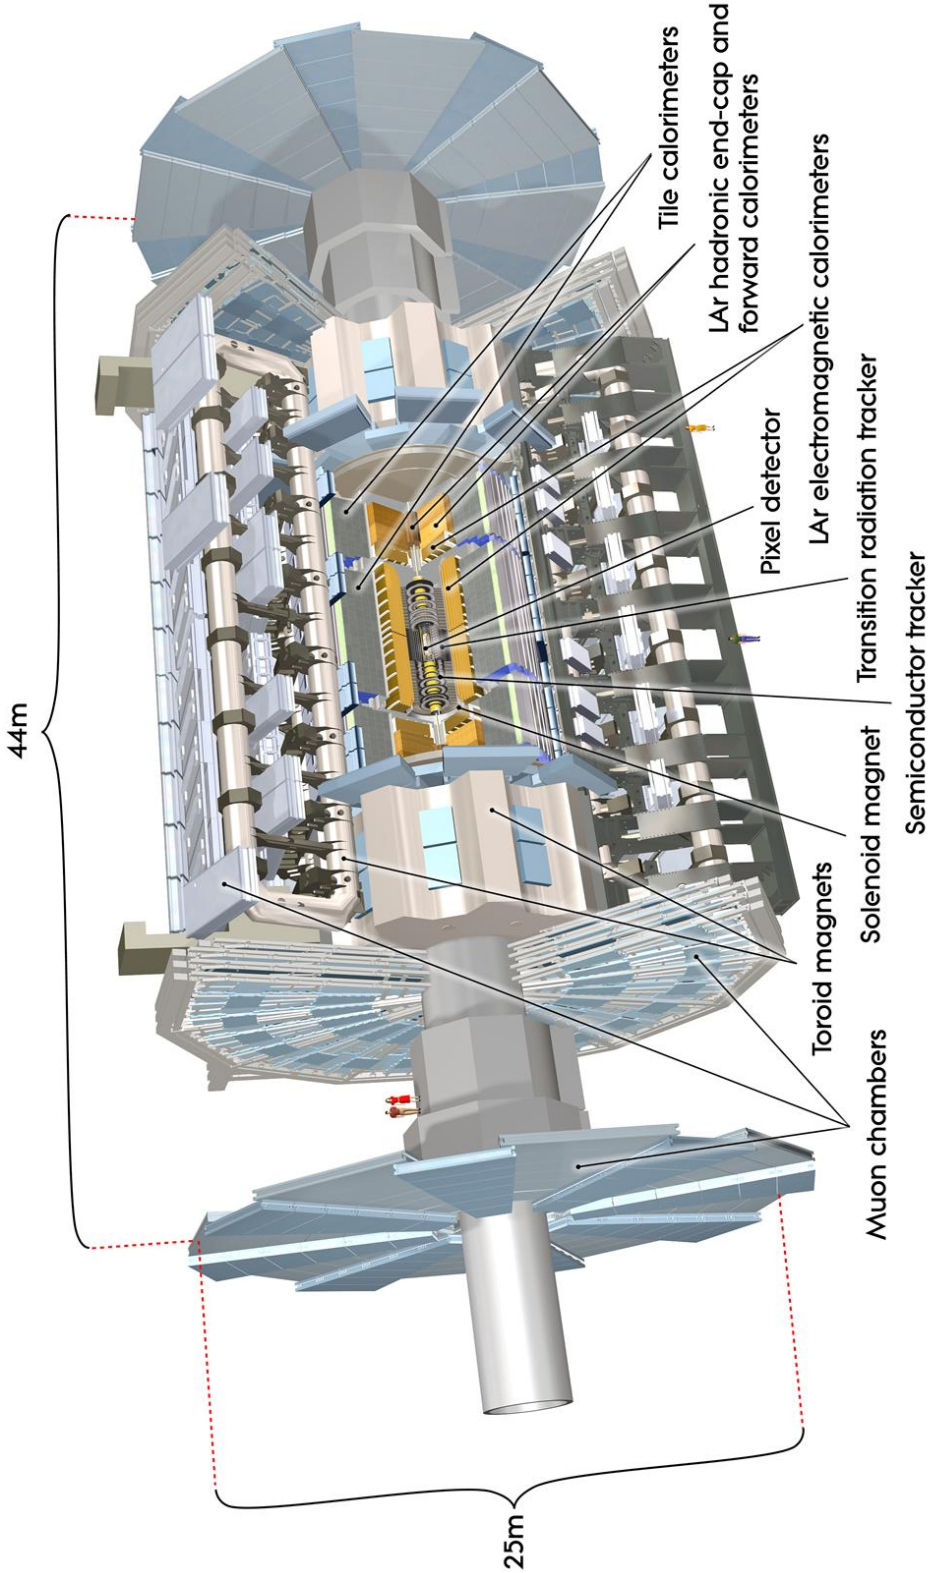


Figure 4.1: A cut-away schematic of the ATLAS detector at the Large Hadron Collider, reproduced from Reference [45], and showing the various detector subsystems. The forward detector subsystems, which are not considered in any analysis in this thesis, are not shown.

21×10^{33} cm^{2}s^{-1} [1], twice the original design luminosity. Protons are accelerated and collided in bunches every 25 ns. In order to achieve these collision rates and luminosities, there must be multiple proton-proton interactions per bunch-crossing. Although such interactions allow us to achieve the desired particle production rates, a consequence of this is the presence of stochastic, QCD-dominated minimum-bias proton-proton collisions which do not correspond to the hard *pp* interaction of interest; such collisions are referred to as *pileup* and will degrade the desired final detector performance unless suitably accounted for. The average pileup and integrated luminosity throughout Run 2 LHC data-taking, which extended from 2015 to the end of 2018, are shown in Figures 4.3 and 4.4 respectively. The mean number of inelastic *pp* collisions per crossing is denoted by μ, where $\langle\mu\rangle$ corresponds to the average pileup over all colliding bunch pairs [49]. Pileup contributions may arise from collidion events in the same bunch-crossing ($in-time$ pileup) or in separate bunch-crossings ($out-of-time$ pileup).

4.2 The ATLAS Detector

4.2.1 Coordinate system

In ATLAS a right-handed coordinate system is adopted: the z-axis points along the beamline, with the x-axis inwards towards the centre of the LHC ring, and the y-axis upwards. In the transverse plane, cylindrical coordinates (r, ϕ) are used, where ϕ is the azimuthal angle around the beam pipe. The polar angle, θ, defines the pseuodrapidilty variable, $\eta = -\ln\tan(\theta/2)$. In the massless limit, this is equivalent to the rapidity $y = \frac{1}{2}\ln\frac{E+p_z}{E-p_z}$, a useful variable which has a nice physical interpretation as the relativistic boost along the beam axis, which takes an observer from the laboratory frame to a frame in which the particle moves only perpendicular to the beam itself (i.e. *transverse* to the beamline). The geometry of the ATLAS cylinder can be unfolded into a flat plane in $\eta-\phi$ coordinates, where the separation between physics objects in this plane is given by $\Delta R = \sqrt{(\Delta\eta)^2+(\Delta\phi)^2}$, and $\eta-\phi$ plane is also transverse to the beamline. The geometry of the ATLAS detector removes the ability for the momentum fraction of the colliding partons in the longitudinal direction of the beam to be easily measured, and therefore transverse quantities are typically used in detector-level analysis. An important example is the transverse momentum vector, $\mathbf{p}_\mathrm{T} = (p_x, p_y)$. Conservation of the transverse momentum plays an important role in the inference of different physics objects in the ATLAS $\eta-\phi$ plane.

[1] This rises to 10^{34} cm^{2}s^{-1} in higher pileup runs.

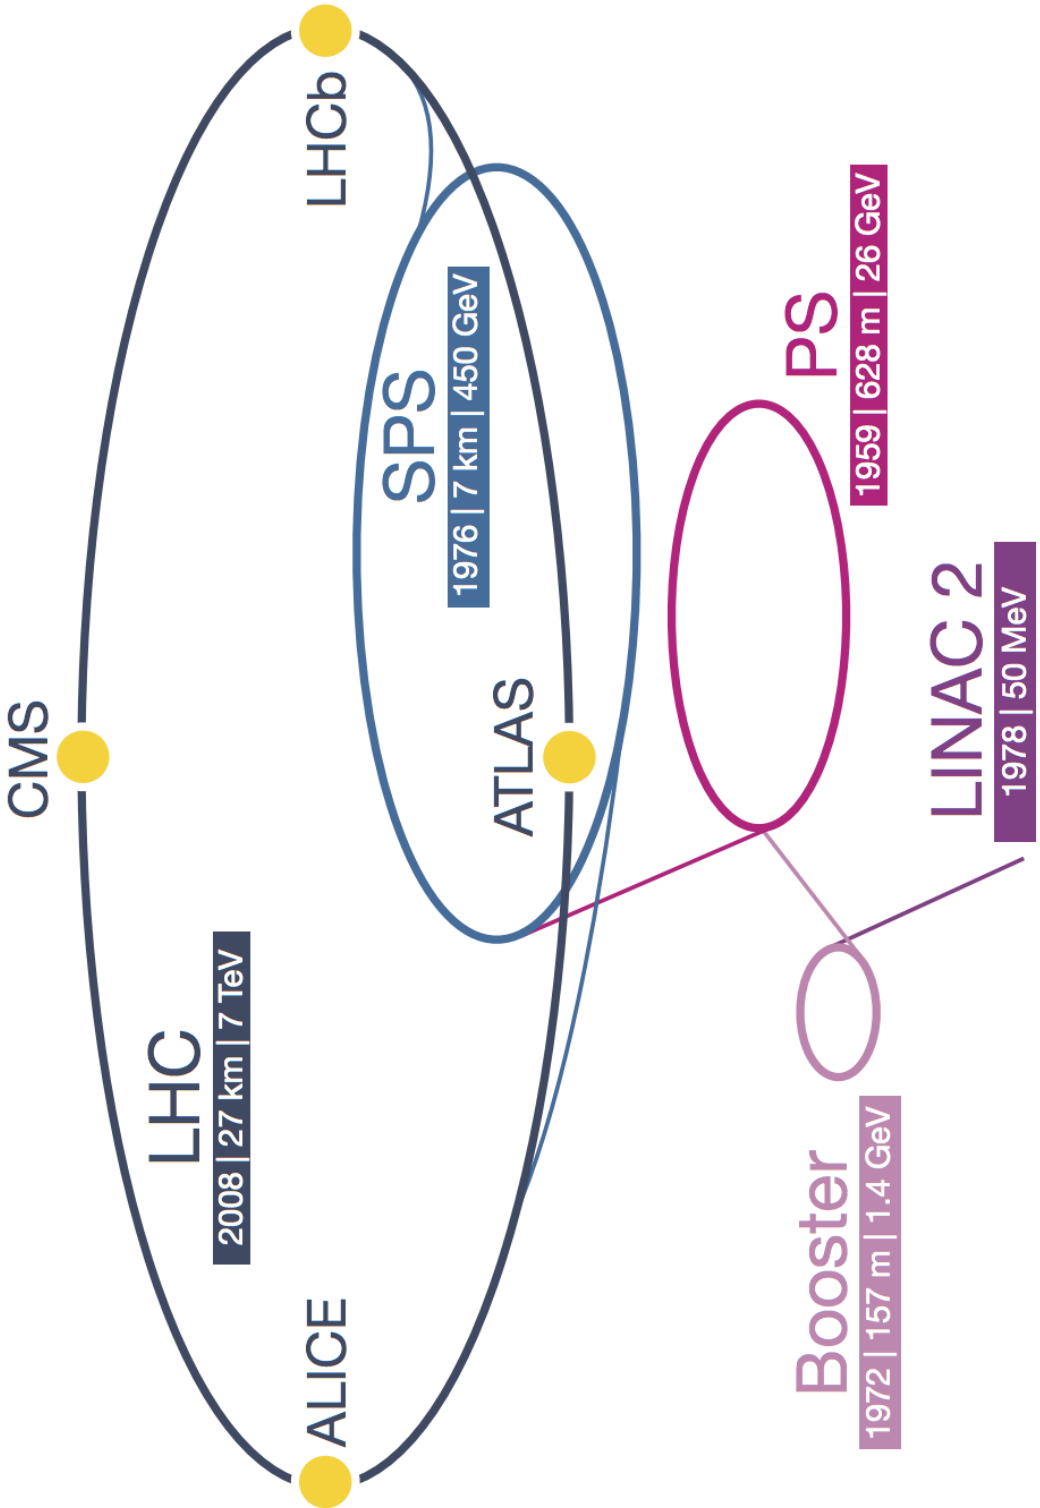


Figure 4.2: The LHC accelerator complex at CERN, including the various pre-accelerators. Shown are the linear accelerator (LINAC 2), the Proton Synchrotron Booster, the Proton Synchrotron (PS), Super Proton Synchrotron (SPS), and the Large Hadron Collider (LHC). The four main experiments on the LHC ring are also shown, although their relative positions on the ring are not to scale. Image taken from Reference [50].

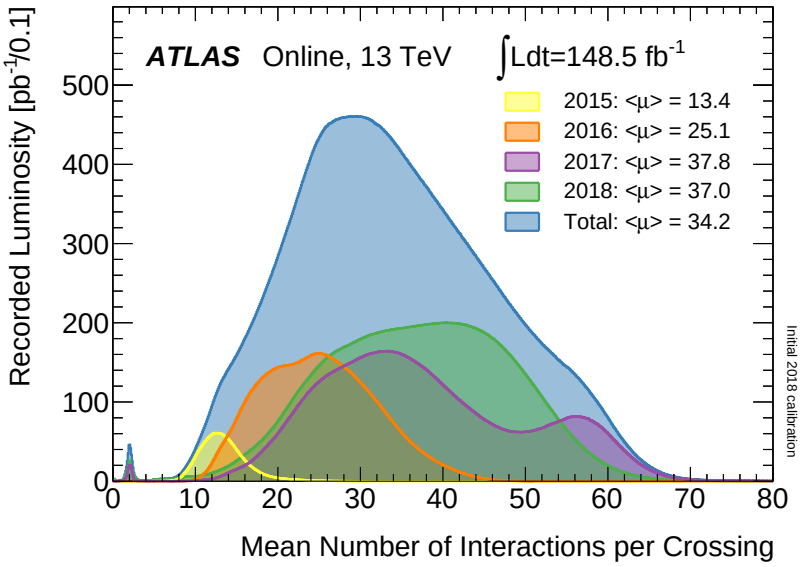


Figure 4.3: The luminosity weighted distribution of the mean number of interactions per bunch-crossing (*pileup*), split by year throughout Run 2 of ATLAS. Taken from [51].

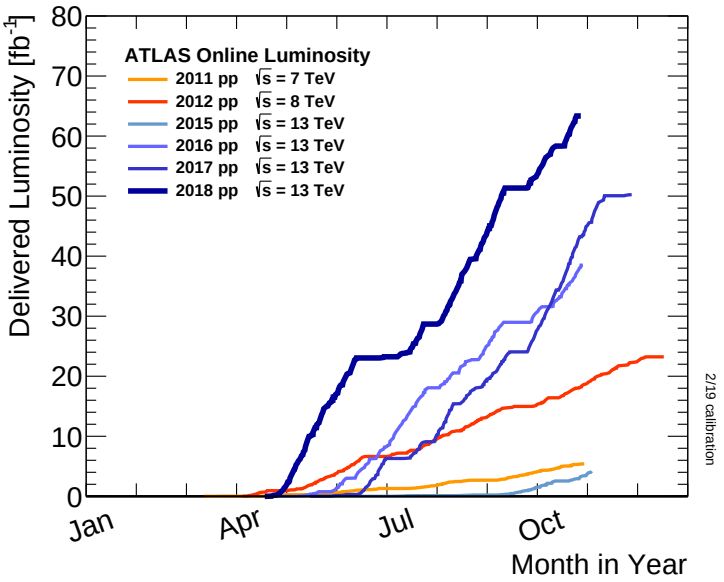


Figure 4.4: The cumulative luminosity vs. day in year delivered to ATLAS during stable beams for high-energy *pp* processes, shown for the various years through Run 1 and Run 2 data-taking. Taken from [51].

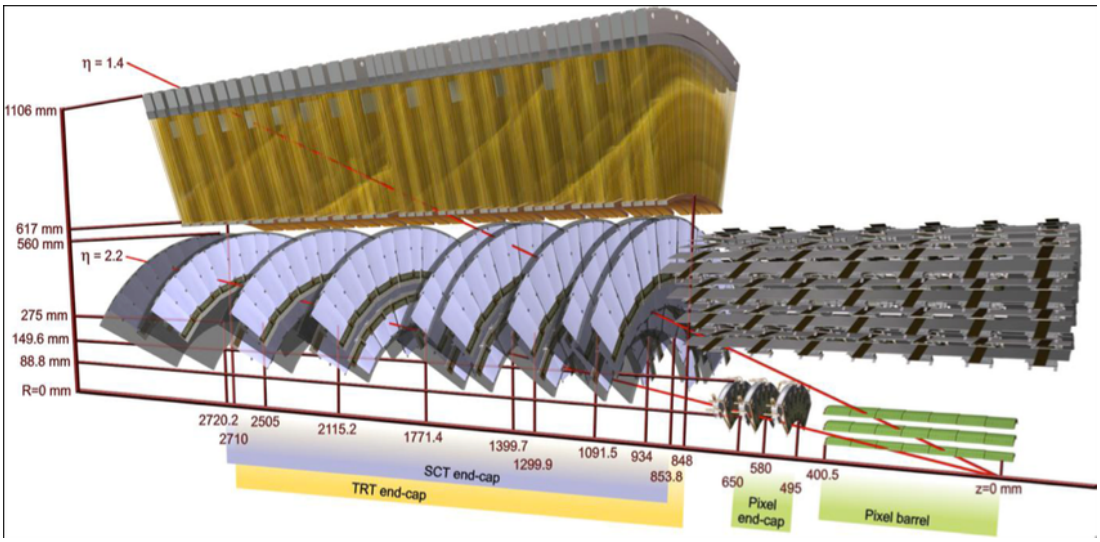
4.2.2 Overview

ATLAS is one of the two large, general purpose particle physics detectors at the LHC and extends to 46 m in length. It has a cylindrical design, and is almost entirely hermetic with most detector subsystems split into barrel and end-cap components (see Figure 4.1). The detector can be divided into several major subsystems. The innermost central inner detector (ID) [52] consists of silicon pixels and silicon strips comprising the semiconductor tracker (SCT), enclosed by the transition radiation tracker (TRT), which together provide tracking information and particle identification up to $|\eta| < 2.5$. The ID is surrounded by a solenoid magnet, which provides a near uniform 2 T axial magnetic field, allowing the transverse momentum of the ID particle tracks to be accurately determined. Moving outwards, we have the calorimeters: these are split into the electromagnetic and hadronic calorimeters respectively, EMCAL and HCAL, which are a series of liquid argon (LAr) detectors [53] in the barrel and end-cap regions, together with plastic scintillator tiles in the barrel region only (this is referred to as the TileCal [54].) A 4 T toroidal magnetic field further surrounds this calorimeter region, which extends to $|\eta| < 4.9$. Beyond this there are a series of muon detectors, comprising the muon spectrometer (MS) subsystem [55], which is responsible for the fast-triggering (in addition to the full ATLAS trigger system, described later) and detection of muons.

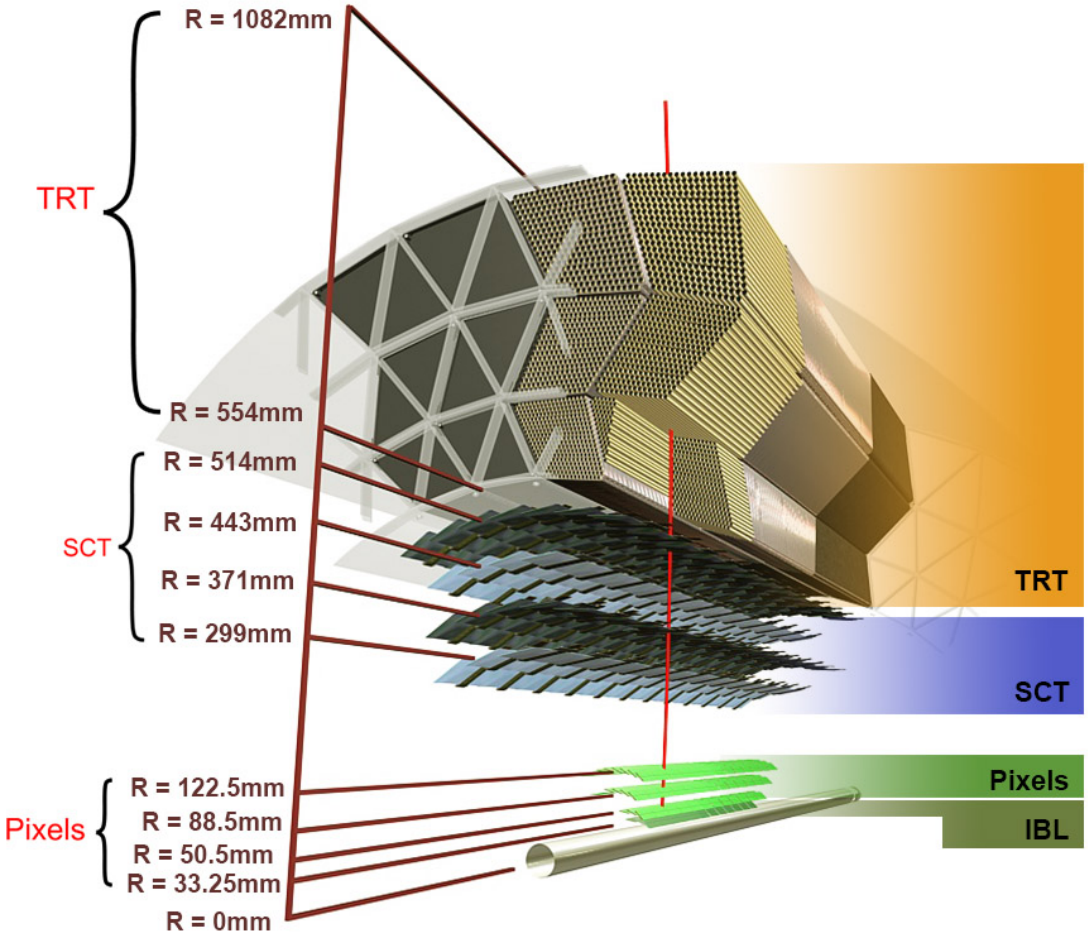
4.2.3 Inner Detector

The ATLAS Inner Detector, and its subsystems, are shown schematically in Figure 4.5. The purpose of the ID is to provide detailed measurements of the paths of the charged particles produced close to the interaction point. By encasing the ID in a solenoid magnetic field, it is possible to measure the curvature of these tracks, from which the transverse momentum of the propagating particle is determined. These tracks may also be associated to the vertices corresponding to the original 'hard-scatter' event[2], or to subsequent secondary vertices produced through the later decay of long-lived particles which propagate some non-negligible distance from the interaction point. The identification of such vertices, associated with the measured tracks, are vital for the identification of e.g. B- and D-mesons close to the original interaction point(see for example techniques utilised in [56, 57]).

[2]This is just the hard process which the partons of the original pp collision are involved in, and it shall be discussed in more detail in the context of jet physics in Chapter 5.



(a) Inner detector end-cap region. Taken from [45].



(b) Inner detector barrel region. Taken from [52].

Figure 4.5: A cut-away view of the ATLAS inner detector system used in Run 2, including the IBL and pixel subsystems, the SCT, and the TRT. The IBL subsystem is closest to the beam and was added for Run 2 data-taking. The end-caps are shown in the upper panel, and the barrel region in the lower panel.

4.2.3.1 Pixel Subsystem

The pixel system [52], which is the innermost ATLAS subsystem, comprises 2 end-caps (each comprising 3 disks) and 3 outer barrel layers made from 1744 biased silicon modules, combined into over 80 million 50 $\times$ 400 μm individual silicon pixel cells. The passage of charged particles through these silicon modules liberates electron-hole pairs, producing a multitude of charge carriers (proportional to the energy of the charged particle). The application of a bias voltage causes these charge carriers to drift, generating a charge separation which is subsequently amplified to produce the recorded signal. These signals are used to accurately measure the position of the charged particle to a precision of approximately 10 μm. For Run 2 an additional innermost pixel layer − the Insertible B-layer (IBL) [52] − has been added, providing an additional 286 modules comprising 13 million 50 $\times$ 250 μm silicon pixels. The IBL allows for an extremely precise determination of tracking parameters, vital to the identification of short-lived mesons. This combined pixel subsystem spans up to 150 mm from the interaction point. Any particle produced within $|\eta| < 2.5$ passes through the 4 pixel layers of the pixel system.

4.2.3.2 Semiconductor Tracker (SCT)

The SCT [45, 52] is another biased silicon detector, and spans from 27.5 to 56 cm from the interaction point. It consists of 4 cylindrical barrel layers and 2 end-caps, each with 9 disks, through which charged particles propagate, allowing for the measurement of tracks to within $|\eta| < 2.5$. The SCT consists of 4088 double-sided silicon modules, which are arranged into strips roughly 13 cm long and 80 μm wide. Additionally there are axial silicon strips at an angle of 40 mrad to the initial strip layers, providing additional tracking information. The result is over 6 million readout channels, which provide measurements on the charged particle ϕ position to within approximately 16 μm. Additional z- and r-measurements are provided, using the barrel and end-cap channels respectively, to a precision of 580 μm in z and approximately 16 μm in $r-\phi$.

4.2.3.3 Transition Radiation Tracker (TRT)

The TRT [45, 52] component is build from straw detectors, 4 mm diameter kapton tubes, filled with a $Xe/CF_4/CO_2$ gas mixture (to a volume of 3 m^3) and 30 μm diameter anode wires running along the detector centreline. Like the SCT, it provides coverage up to $|\eta| < 2.5$. As charged particles propagate through the gas mixture, electrons are liberated and drift towards the central anode wire, liberating additional electrons along the way. Additional

materials of varying dielectric constant surround the straws, and these emit transition radiation when highly-relativistic electrons cross at the boundaries between them. In total, the TRT provides 420 000 readout channels, with the particle drift time measurements providing a charged particle spatial resolution of approximately 170 μm. In this way the TRT provides additional information for electron identification as well as tracking.

4.2.4 Calorimeters

A schematic of the ATLAS calorimeter system is provided in Figure 4.6: these surround the ID solonoid magnets and provide energy measurements of both charged and neutral particles up to $|\eta| < 4.9$. These calorimeters consist of alternative layers of active and passive material. The passive materal acts as an absorber; it is typically very dense in order to produce many interactions between itself and the propagating particles of interest [45, 53]. The interwoven layers of active material react to the particles which pass through them, and the constituents either become ionised or scintilate, generating signal current or emiting light which is detected by photodetectors, respectively. The measurements from these two materials can be combined to measure the shower shapes and energies of the interacting particles in the calorimeter. Since the particles which are measured are also absorbed in the process, calorimeters are destructive detectors.The ATLAS calorimeters are non-compensating, and as such the detector responses at the electromagnetic and hadronic scales differ. Dedicated calibrations are derived to account for such differences, and shall be briefly discussed in the context of jet physics in Chapter 5.

4.2.4.1 Liquid Argon (LAr) Electromagnetic Calorimeter

The LAr calorimeter [45, 53] comprises an electromagnetic barrel (EMB, $|\eta| < 1.475$) and two end-caps which are divided into three parts. The electromagnetic end-caps (EMEC, $1.375 < |\eta| < 3.2$) are situated the closest to the interaction point. Further out are the hadronic end-caps (HEC, $1.5 < |\eta| < 3.2$), and finally at a similar distance, but closer to the beam pipe, lie the forward calorimeters (FCAL, $3.1 < |\eta| < 4.9$). The EMB and EMEC use plates of lead and stainless steel as absorbers, whilst the HEC uses copper plates. The FCAL consists of 3 modules in each end-cap: the first module uses copper as the absorber and is optimised for electromagnetic shower measurements, and the second two use tungsten as the absorber and are optimised for hadronic shower measurements. Liquid argon serves as the active material throughout these calorimeter systems. Copper-kapton tape electrodes are placed between layers of absorber in order to collect

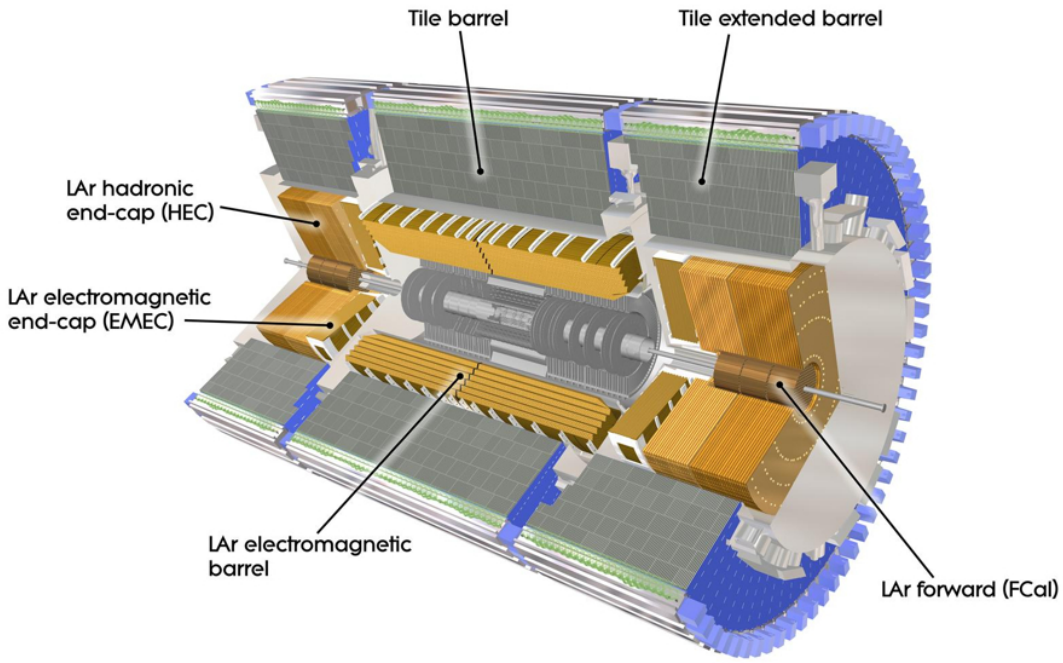


Figure 4.6: A cut-away schematic of the electromagnetic and hadronic calorimeter systems in ATLAS, including the barrel and end-cap regions. Image taken from [45].

electrical signals. In order to measure a uniform energy distribution in ϕ, the absorber plates are arranged in a zig-zag accordion fashion, and the plates are also etched uniformly in order to provide an electrical readout in uniform towers in $\eta - \phi$. The EBM is finely segmented in order to provide precision measurements of the energy depositions of electrons and photons.

4.2.4.2 Tile Hadronic Calorimeter

Surrounding the electromagnetic calorimeter is the simpler Tile hadronic calorimeter (TileCal) [45, 54] subsystem. This comprises a barrel surrounding the EMC and two extended barrels outside the EMECs, with no additional end-cap. Steel plates are used as absorbers, with plastic scintillators providing the active material in a series of tiles, extending radially as a set of wedge-shaped modules. The scintillating active material generates light which is transmitted by blue-to-green wavelength shifting fibre optics, which is then read out using photomultiplier tubes situated outside the module system. These modules are segmented in to three layers in order to provide a segmentation of $\Delta\phi$ and $\Delta\eta$ of 0.1 $\times$ 0.1 for the two inner layers, and 0.1 $\times$ 0.2 for the third outer layer. The TileCal provides a total thickness of $9.7\lambda_I$ in the barrel,

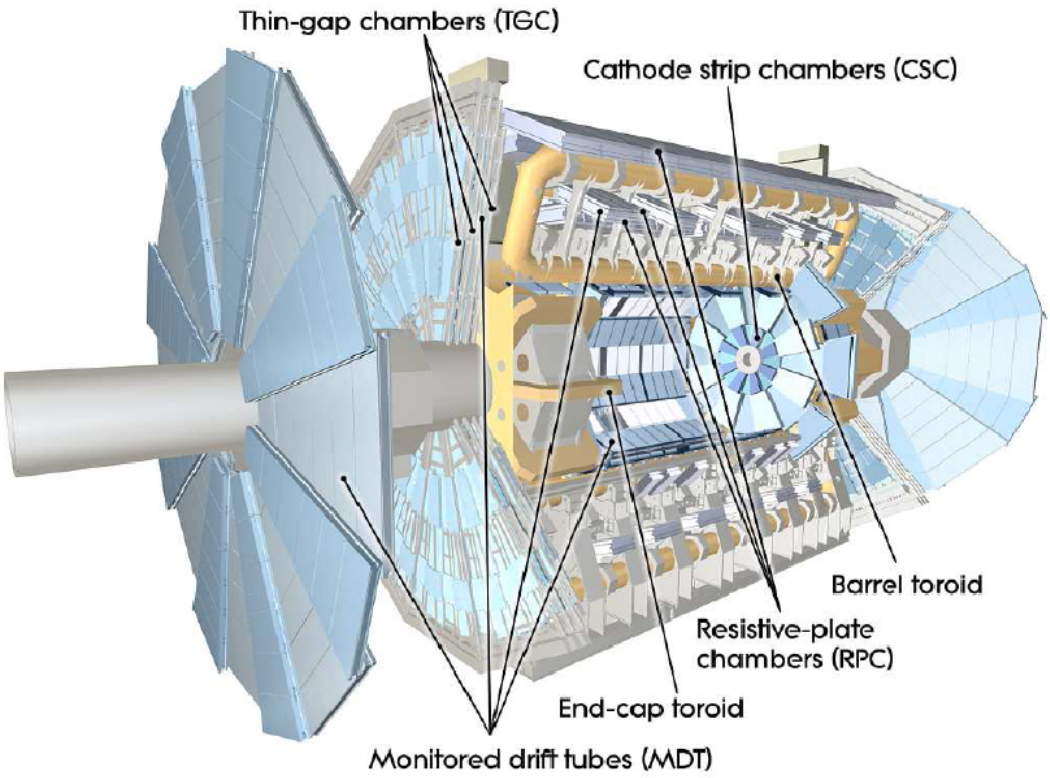


Figure 4.7: A cut-away schematic of the muon spectrometer and its various detector subsystems. Image taken from [45].

and $10\lambda_I$ in the end-caps, where λ_I denotes the *nuclear interaction length*.[3]

4.2.5 Muon System

The outermost muon spectrometer (MS) [45, 55] component (shown schematically in Figure 4.7) comprises several subsystems. The MS provides near-hermetic coverage of the barrel volume of the inner ATLAS subsystems, and also provides measurements of the forward regions using the large lever arm of the muon wheels. Due to access requirements for cabling and services to the detector, the MS is not fully instrumented in the central $|\eta| < 0.1$ region. In total, the MS utilises four different detector technologies [55]: monitored drift tubes, cathode strip chambers, resistive-plate chambers, and thin-gap chambers. Additionally the ATLAS toroid magnets utilised in the MS comprise 8 barrel coils and 2 sets of 8 end-cap coils, generating fields of up to about 4 T.

[3]The nuclear interaction length is the mean free path for nuclear interaction of a high-energy hadronic particle. Shorter interaction lengths result in less calorimeter material required to contain particle showers.

4.2.5.1 Monitored Drift Tubes (MDTs)

The MDTs [55] provide the vast majority of the MS coverage, and operate using similar principles to the TRT straw tubes. The aluminium MDT straws are larger, with diameters of 30 mm and lengths between 0.7 and 6.3 m, and contain a pressurised Ar/CO_2 gas mixture surrounding a central wire. These are separated into sets of 3 or 4 layers, and lie parallel to the field of the toroid magnet in order to provide the best possible measurement of the direction of the bending muons.

4.2.5.2 Cathode Strip Chambers (CSCs)

The CSCs [55] are situated in the forward region of the detector, but closer in than the toroid system, where the particle densities are too high for the MDTs to function reliably. These consist of multiwire proportional chambers with anode wire spacing of 2.5 mm, together with perpendicular cathode strips which provide a second spatial coordinate, and are filled with an Ar/CO_2 gas mixture. The CSCs benefit for a much finer granularity than the MDTs, and are therefore able to measure the muon activity in the forward region, with a radial precision of 40 μm.

4.2.5.3 Resistive-plate Chambers (RPCs)

The typical time for charges to drift in the MDT straws is 500 ns, which is far too slow for accurate muon triggering. To aid triggering, two RPCs [55] are placed in the barrel region. These are made from two resistive Bakelite plates, separated by a 2 mm gap containing a gas mixture of $C_2H_2F_4$/Iso-C_4H_{10}/SF_6 in a uniform electric field. This additional technology provides a 1 ns time resolution and a 10 mm positional resolution (in each case, in the plane of the RPC plates).

4.2.5.4 Thin-gap Chambers (TGCs)

The RPCs, though necessary to provide an exceptional timing resolution, do suffer from poor granularity. In place of the RPCs, TGCs are used in the end-cap regions. TGCs comprise of a plane of 50 μm anode wires, separated by 2 mm, which are held between two grounded plates and immersed in a pentane/CO_2 gas mixture. This allows for much higher granularity in the end-caps (measuring the precision of the propagating particle to within approximately 5 mm), whilst maintaining a sufficiently effective response speed for the fast triggering of muons.

4.3 Trigger

The LHC has a collision rate of approximately 40 MHz, corresponding to a total output rate of around 50 TB per second. This rate of information would completely overwhelm the data-storage capabilities of ATLAS (where the maximum storage size of data-sets is approximately 1.3 MB per event, assuming the full storage capacity is reached). In addition to this, the vast majority of LHC physics collisions give rise to soft QCD processes, which are typically not the physics processes of interest. Finally, there is a limit to the data-storage and information-transfer rate of many of the subdetectors, typically at the level of 100 kHz. It is therefore imperative to limit the actual number of events which the detector reads out so that a more manageable amount of data can be recorded, whilst still ensuring that the physics processes of interest are not lost; this is controlled by a *trigger system* [58].

The ATLAS trigger system is divided into two subsystems for Run 2 data-taking: the Level-1 (L1) trigger, and the Higher-Level Trigger (HLT). The L1 trigger is a hardware-based trigger and is designed to reduce the input 40 MHz rate down to a more manageable 100 kHz, as required by the readout limitations of the hardware of ATLAS. The HLT is a software-based trigger which further reduces this rate down to 1 kHz. Different sets of online cuts and thresholds (trigger strategies) are considered in order to trigger on different physics objects of interest to analysers.

4.3.1 L1 Trigger

The L1 trigger consists of two separate subsystems: the calorimeter trigger (L1Calo) and the muon trigger [59]. Signals from these two systems, together with additional trigger subsystems, are combined in the Central Trigger Processor (CTP). This CTP is responsible for making the final L1 trigger decision, and ultimately deciding whether an event is accepted for further consideration at the HLT, or rejected. The CTP also applies *dead time*, so that there is a sufficient time-gap between two L1 event acceptances, allowing adequate time for detector readouts to register each acceptance without overlap between the two signals; a maximum number of L1 event acceptances is also set in order to prevent inputs to the HLT causing that trigger system to overflow.

The L1Calo receives information from the entire calorimeter, divided into *trigger towers* of varying size in $\Delta\eta \times \Delta\phi$. For example, in the barrel region these towers measure 0.1×0.1 up to $|\eta| < 2.5$, 0.2×0.2 for $2.5 < |\eta| < 3.1$, and 0.1×0.2 for $3.1 < |\eta| < 3.2$. Electron and photon candidates are formed using a sliding-window algorithm [59] across the trigger towers, comparing the energy sizes against a given threshold, and requiring that the shower from which the event is triggered on does not extend to the hadronic calorimeter, remaining localised to the electromagnetic calorimeter. Tau candidates are formed in a similar manner, but the condition of shower localisation to the electromagnetic calorimeter is removed. Similarly, a sliding-window algorithm across trigger towers from both the electromagnetic and hadronic calorimeters is used to form jet and missing transverse momentum candidates (which shall be discussed in more detail shortly).

The L1 muon event acceptance decision within $|\eta| < 1.05$ is determined by the RPCs in the barrel, and for $1.05 < |\eta| < 2.4$ by the TGCs in the end-caps. In the barrel coincident hits in at least 3 of the 4 layers of the inner two planes of the muon system are required, which allows for a low-p_{T} trigger. In order to trigger on high-p_{T} muons in the barrel region the event must pass the acceptance requirements of the low p_{T} trigger and also have coincident hits in the third plane of the muon system. A similar procedure is employed for triggering on events using the end-caps.

4.3.2 HLT

The HLT is able to view ATLAS globally, considering all detectors at once. This is in contrast to the L1 trigger, which makes decisions by considering the various detectors in sections and combining these using a central trigger system including processing from L1Topo [60]. The L1 trigger defines certain *regions of interest* (RoIs), which are used by the HLT to perform CPU-intensive algorithms, such as track reconstruction. The final set of HLT physics objects are designed to be as close to their offline counterparts as possible, prior to any offline physics reconstruction. It is this offline physics reconstruction, using events which pass the final L1 and HLT trigger system, which is employed to generate the final set of physics objects considered by analysers. We shall now briefly turn to these offline physics object definitions.

4.4 Physics Object Reconstruction

Following trigger acceptances at the L1 and HLT, ATLAS software is used to perform an offline reconstruction according to the recommendations of the various *Combined Performance* (CP) groups. The final list of reconstructed physics objects correspond to the objects of interest for most ATLAS analyses[4].

4.4.1 Primary Vertices and Tracks

Tracks are reconstructed from the hits produced by charged particles interacting with the inner detector material, and are required to have $p_{\mathrm{T}} \geq 500$ MeV with reconstruction within $|\eta| < 2.5$. First the raw readout from the pixel and SCT detectors is translated into sets of space-point measurements of the propagating charged particles, giving rise to sets of detector 'hits'. A seed track is then formed from the hits in the pixel layer, and is extended to the SCT. To reduce fake tracks, it is required that the candidate tracks are formed from at least a total of 7 hits in the combination of the pixel and silicon microstrip trackers. In addition, the longitudinal impact parameter relative to the *primary vertex*, $|z_0| \sin\theta$, must be less than 1.5 mm, where the θ angle is the angle between the beam and the measured track. Additionally the track transverse impact parameter $|d_0|$ relative to the primary vertex must be less than 1.5 mm. Finally the candidate tracks are further extended to include information from the TRT. Following this, the entire track is re-fitted using information from all of the ID. An extra track-finding algorithm is then applied backwards, starting from the TRT measurements, in order to improve the efficiency of reconstructing secondary tracks which may arise from the presence of decaying long-lived particles, or photon conversion processes.

Primary vertices are identified using a dedicated vertex-finding algorithm [62], which searches through the vertices of the event. Vertices are reconstructed by requiring the association of at least two charged-particle tracks with $p_{\mathrm{T}} > 400$ MeV, reconstructed within the inner detector. The sum of the squared transverse momentum of all tracks associated to each vertex is calculated, and the vertex with the largest squared sum of p_{T} is deemed to be the primary vertex originating from the hard-scatter event. This provides one primary vertex per pp bunch-crossing.

[4]Notable exceptions are the so-called 'data-scouting', or trigger-level, physics analyses pioneered during Run 2 of ATLAS [61].

4.4.2 Electrons and Photons

Electrons and photons are reconstructed in parallel from the measured calorimeter cell energy clusters. These clusters are formed using the process of *topoclustering*, which shall be discussed in detail in Chapter 5. Complete details can be found in Reference [63]. These two objects are distinguished by matching the clusters to different tracks in the ID, where the electron will be matched to a charged track. The clusters of interest are first identified by starting from seeds with transverse energy, $E_{\mathrm{T}} > 2.5$ GeV, and then applying a sliding-window algorithm in $\Delta\eta \times \Delta\phi$ to fully locate the cluster. The track matching is achieved through extrapolating tracks with $p_{\mathrm{T}} > 400$ MeV from their last-measured point in the ID to the LAr calorimeter, and then comparing the η- and ϕ-coordinates of the extrapolated track to the position of the original seeded cluster: the two are considered matched if the distance between the track impact position in the LAr and the EM calorimeter cluster barycentre has $\Delta\eta < 0.05$. If there is no matched track, then the cluster is identified as an unconverted photon. To then distinguish electrons from the converted photons it is necessary to look for two close-by oppositely charged ID tracks which originate from a vertex displaced with respect to the interaction point. The total reconstructed energy is then calculated from the deposited energy in the material in between the calorimeter and the interaction point and the cluster energy (following corrections for any energy which may fall outside of the LAr).

4.4.3 Muons

Muons are reconstructed by combining tracks formed in the ID and adding in information from the muon spectrometer, giving an improved background rejection from non-LHC muons such as those produced from cosmic ray showers in the atmosphere. The final muon identification procedure results in four different "types" for analyses [64]:

Combined Muon : Here the reconstruction is performed independently in the ID and muon spectrometer and, if the tracks from each system can be successfully combined, then the resulting single track is dubbed a combined track. Following this, the track is re-fitted in order to take into account information from both the ID and the muon spectrometer. This is the most common type of muon used in analyses, and also has the best momentum resolution performance.

Stand Alone : Here the trajectory of the muon is reconstructed from muon spectrometer information only, and then extrapolated back to the point of closest approach with the beam line. No ID information is provided in the reconstruction, and this type of muon is used outside of the ID acceptance, in $2.5 < |\eta| < 2.7$.

Segment − Tagged : It is possible to identify ID tracks as muon tracks if they can be extrapolated from the ID to the muon spectrometer, such that the tracks are associated to at least one track segment in the MDT or CSC chambers. These segment-tagged muons can be used to enhance the final detector muon acceptance.

Calorimeter − Tagged : Finally it is also possible to identify ID tracks as muon tracks if they can be extrapolated to the calorimeter, and from there be associated to calorimeter energy deposits which are consistent with the presence of *minimum ionising particles*. Again this method helps to increase the overall muon acceptance in the detector.

4.4.4 Jets

Jets form the central physics object of this thesis, and the more standard jet definitions are formed from the calorimeter energy deposits created from topoclustering. A jet acts as a proxy for the complex and collimated sprays of hadronic energy produced from QCD, and the theory of jet physics and jet reconstruction is rich, complex, and deeply fascinating. We shall postpone further discussion of jet reconstruction to Chapter 5. It is also possible to *flavour-tag* a jet, where the flavour of a quark inside a hadron can be identified using the jet information. In this thesis b-tagging [65] is employed to identify jets which contain a B-hadron. The relatively long lifetime of a B-hadron (approximately 1.5 ps) gives the particle a longer transverse flight distance. This allows for tracks of charged particles originating from the decay of the B-hadron to be identified as those not originating from the primary vertex, and can be used to construct a secondary vertex. This additional information is used in various approaches to b-tagging in ATLAS.

4.4.5 Taus

Taus present a particular challenge in object reconstruction; they may decay both hadronically and leptonically, and are sufficiently long-lived to be easily mistaken for b-tagged jets. The lep-

tonically decaying taus give rise to muons, electrons, and neutrinos, and cannot be distinguished from prompt electrons or muons in the detector. Consequently, it is only possible to identify hadronically-decaying taus, which correspond to roughly 65 % of all tau decays [6]. Hadronic tau decays are categorised as either *one-prong* or *three-prong*, depending on the number of charged tracks which are subsequently detected. The virtual W boson produced from the tau decay typically decays to a heavy resonance, such as a $a^{\pm}$ or $\rho^{\pm}$. The $\rho^{\pm}$ decays to a charged and neutral pion, giving rise to the single charged track of the one-prong mode, whereas the $a^{\pm}$ decays into three (or four) pions in a mix of three charged pions (and a neutral pion), giving rise to the three-prong mode. Taus are initially identified from the collection of reconstructed jets, where a Boosted Decision Tree is used to classify jets as originating from taus or otherwise [66]. This algorithm uses information from the shape of the jets, as taus typically produce more collimated jets and have different energy fractions in the EM calorimeter due to the presence of neutral pions.

4.4.6 Missing Transverse Momentum

Due to the conservation of momentum in the transverse direction to the LHC beam axis, the sum of the transverse momenta across all final state particle should, under the assumption of a perfect detector reconstruction, sum to zero. Some particles, such as neutrinos or the hypothetical LSP particles in SUSY phenomenology, are invisible to particle detectors and may propagate through the detector without leaving any signal. This would generate a momentum imbalance in the final sum: the negative vector sum of all the momenta in the event is defined as the *missing transverse momentum*, $\mathbf{E}_{\mathrm{T}}^{\mathrm{miss}}$. The magnitude of this vector, $E_{\mathrm{T}}^{\mathrm{miss}}$, is typically used in analyses. In calculating the $E_{\mathrm{T}}^{\mathrm{miss}}$ it is necessary to take in to account the effect of mismeasurement and miscalibration of the physics objects in the event, and the contributions from pileup where particles produced from one interaction point can be confused with those produced in another during the same bunch-crossing. The final missing transverse momentum is calculated from the x- and y-components of the various physics objects in the event which are associated with the primary vertex:

$$E_{x(y)}^{\mathrm{miss}} = E_{x(y)}^{\mathrm{miss,electrons}} + E_{x(y)}^{\mathrm{miss,photons}} + E_{x(y)}^{\mathrm{miss,muons}} + E_{x(y)}^{\mathrm{miss,taus}} + E_{x(y)}^{\mathrm{miss,jets}} + E_{x(y)}^{\mathrm{miss,soft}}. \quad (4.1)$$

The final term, $E_{x(y)}^{\mathrm{miss,soft}}$ is the so-called *soft term* [67] which accounts for signal in the detector which is not associated with any of the well-defined physics objects above. This soft

term can be calculated using either unassociated calorimeter energy deposits (the *calorimeter soft term* (CST)) or tracks originating from the primary vertex which are not associated to any of the physics objects in the event (the *track soft term* (TST)). In Run 2 of ATLAS, the TST is used for the soft contribution as it was found to be more pileup robust than the CST.

4.4.7 Overlap Removal

It is possible for the same reconstruction procedure to double-count a physics object, for example an energy cluster in the calorimeter may be identified as both an electron and a jet. To mitigate this issue in a physics analysis, a procedure called *overlap removal* is applied following the object reconstruction. This is described in detail in Chapter 8.

4.5 Summary

The description of the LHC and the ATLAS detector concludes the first part of this thesis, which aims to provide a theoretical and experimental overview of the techniques used in the various physics analyses at the ATLAS experiment. We shall now turn to the major analysis contributions of this thesis, focusing first on jet reconstruction and performance techniques in Run 2 of the LHC (Part II), before turning to a benchmark search for new physics utilising the Run 2 ATLAS data-set and the jet reconstruction techniques subsequently presented (Part III).

Part II

Hadronic Jet Performance during Run 2 of ATLAS

... I oft admire,
How Nature, wise and frugal,
Could commit such disproportions ...

– John Milton,
Paradise Lost, Book VIII

5

Jet Reconstruction and Calibration in Run 2

The central physics object considered in this thesis is the hadronic jet. Here a jet is *hadronic* in the sense that it is produced through initial hadron-hadron collisions, such as the proton-proton collisions at the LHC, and is generated through the radiation and decay of a hadron (as opposed to a lepton). Broadly speaking, a 'jet' is simply a proxy for capturing the radiation produced by the complex QCD processes taking place after the initial proton-proton hard-scatter event, or additional radiative emissions during the hard-scatter. This radiation is subsequently deposited in the ATLAS calorimeter, where its energy and angular direction can be measured. We will see in future chapters how jet reconstruction [68] and calibration [69] are essential to numerous searches for new physics, and how such searches are also driven by what sort of information (i.e. decay products) can be captured within the jet itself. Indeed the central physics analysis in this thesis, the ATLAS 'multi-jets' search [70, 71], is a search for the production of new physics which can give rise to final states with many hadronic jets reconstructed using calorimeter information. Since hadronic jet physics will play a leading role in the 'multi-jets' search, the second major part of this thesis will therefore aim to bring the reader up-to-date with the state-of-the-art in jet reconstruction and calibration at ATLAS, providing a firm foundation in the jet physics which is utilised in numerous new physics searches motivated by supersymmetric models.

This chapter will present an overview of the concept of a jet in hadron-hadron colliders, focusing on the inputs required to build these jets and how they act as successful proxies for the original hard-scatter collision. The latter requires finding a meaningful way to reconstruct

these jets in the ATLAS detector, and theoretical approaches to a 'meaningful' reconstruction will also be motivated. In addition, the concept of jet calibration, which will be utilised extensively in subsequent chapters, is also presented using the example of calibrating the energy of a jet measured from the ATLAS data-set with the energy determined from Monte Carlo simulation. Such calibration procedures are necessary in order to correct for several effects, including differences between truth and detector-level Monte Carlo, differences between detector-level Monte Carlo and data, and additional corrections due to the flavour-dependence of the reconstructed jet, the non-compensating nature of the ATLAS calorimeter [69] and consequently the loss of energy when photons, leptons, and hadrons traverse both the electromagnetic and hadronic calorimeters.

5.1 Overview

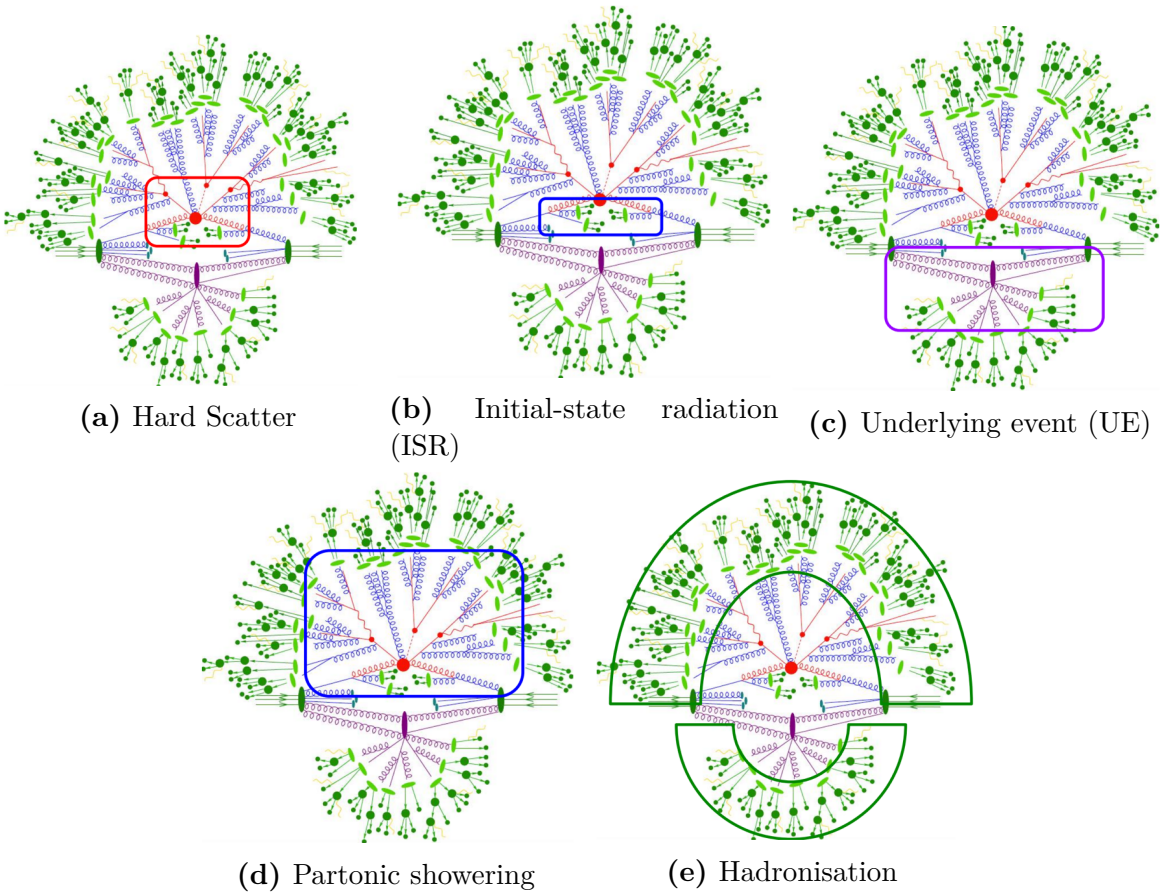
5.1.1 Building an Interaction

Before defining and formalising a jet it is necessary to recall in some detail how interactions occur in proton-proton colliders such as the LHC. In Figure 5.1 five identical cartoons of a toy interaction are presented, with different parts of the interaction highlighted. Such a process can originate from a simple 2-2 scattering of quarks or gluons in the colliding protons, illustrated in the Feynman diagram in Figure 5.2. To be consistent with the literature, these colliding quarks or gluons will be referred to as 'partons' [72, 73] in the following discussion.

Figure 5.1a shows the initial *hard-scatter* between the two partons originating from the colliding protons produced in the LHC beam, and is further simplified in Figure 5.2a which shows the Feynman diagram for the s-channel [74] hard-scatter event. The physics is encoded in a matrix element, $\mathcal{M}$, containing the wavefunction information for incoming and outgoing partons (denoted ψ_u, where u labels a parton of arbitrary flavour):

$$\mathcal{M} = -4\alpha_S \pi \mathcal{C} \bar{\psi}_u \gamma^\mu \psi_u \frac{g_{\mu\nu}}{Q^2} \bar{\psi}_u \gamma^\nu \psi_u, \tag{5.1}$$

where α_S is the strong coupling, Q is the energy scale of the scattering, γ^μ are the usual Dirac γ matrices [75], $g_{\mu\nu}$ is the Minkowski metric, and the colour interaction of QCD is encapsulated in the $\mathcal{C}$-term, or *colour factor*, obtained by summing over a traceless Hermitian matrix



(a) Hard Scatter **(b)** Initial-state radiation (ISR) **(c)** Underlying event (UE)

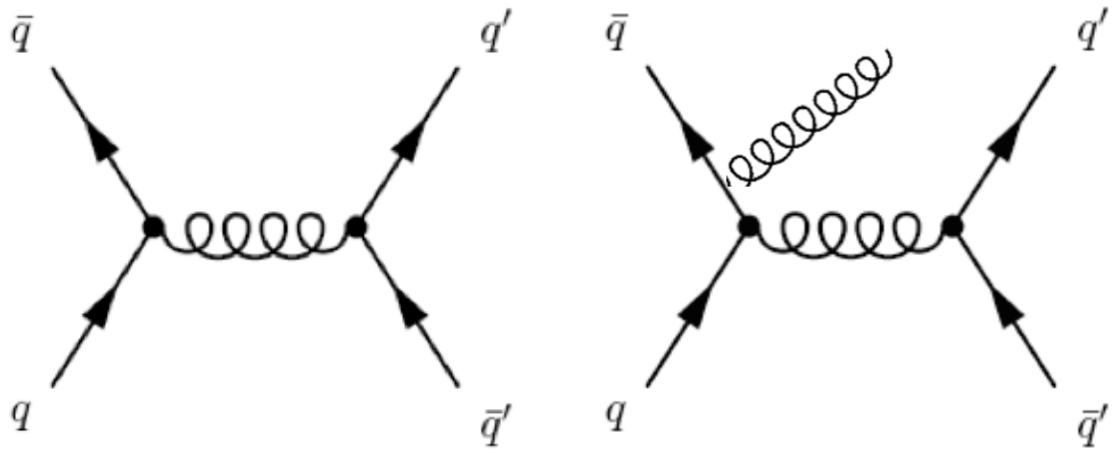
(d) Partonic showering **(e)** Hadronisation

Figure 5.1: A toy model of simple two-two scattering, where the additional processes of multi-partonic interactions (underlying event) and showering of the partons produced from the initial scatter give rise to complex final states consisting of multiple hadrons (commonly pions), leptons from hadronic decay, and radiation. Reconstructing and reinterpreting such processes measured using calorimetry is one of the major challenges in contemporary hadron-hadron collider experiments.

representation of the QCD $SU(3)_C$ generators (the Gell-Mann matrices [10], see Chapter 2).

In addition to the hard-scatter, there is also often *initial state radiation*, (ISR) where one of the incoming partons emits a quark or gluon via the strong interaction before the hard-scatter takes place. This is highlighted in Figure 5.1b and a simplified Feynman diagram, where an ISR gluon is emitted from an incoming quark, is shown in Figure 5.2b.

The proton has a remarkably complex substructure consisting of both valence and 'sea' quarks, as well as virtual gluons produced in strong interactions between those quarks [76]. Consequently additional interactions between other partons which form the proton, but are not the same partons involved in the hard-scatter, also take place. This is shown in Figure 5.1c and is referred to as the *multiparton interaction* (MPI). More generally, the interaction incorporating



(a) $2 \to 2$ hard-scattering **(b)** $2 \to 2$ ISR

Figure 5.2: Feynman diagrams for the s-channel $q\bar{q} \to q'\bar{q}'$ process, showing the hard-scatter and ISR gluon emission.

MPI, ISR, final state radiation (FSR) of quarks or gluons after hard-scatter, and the soft remant of the proton is referred to as the *underlying event* (UE). The partons produced in both the hard-scatter and UE processes can then undergo (non-perturbative) *partonic showering* [77], where they generate cascades of radiation, typically producing copious numbers of soft gluons and light quarks in addition to the hard-scatter particles. Such soft radiation production is demonstrated in 5.1d.

The final, and perhaps most complex, process is that of *hadronisation*, shown in Figure 5.1e. Here the numerous partons which are produced from both the hard-scatter and the partonic showering form colour-singlet bound states of hadrons [10]. These hadrons then propagate through ATLAS and their propagation and decay products can be experimentally measured e.g. as tracks in the ATLAS inner detector and/or energy deposits in the calorimeter.

Having built up this toy model of an interaction, it is clear that even simple 2-2 scattering processes can give rise to final states of many bound hadrons, as well as additional hadrons coming from UE, ISR, and partonic showering. An experimental realisation of such complex QCD-dominated final states is provided in Figure 5.3 which shows an ATLAS Run 1 event where two top quark candidates have been reconstructed in the calorimeters. There are large collimated sprays of energy produced by the partonic showering and hadronisation described above, and the problem now becomes one of reconstructing top quark candidates from these

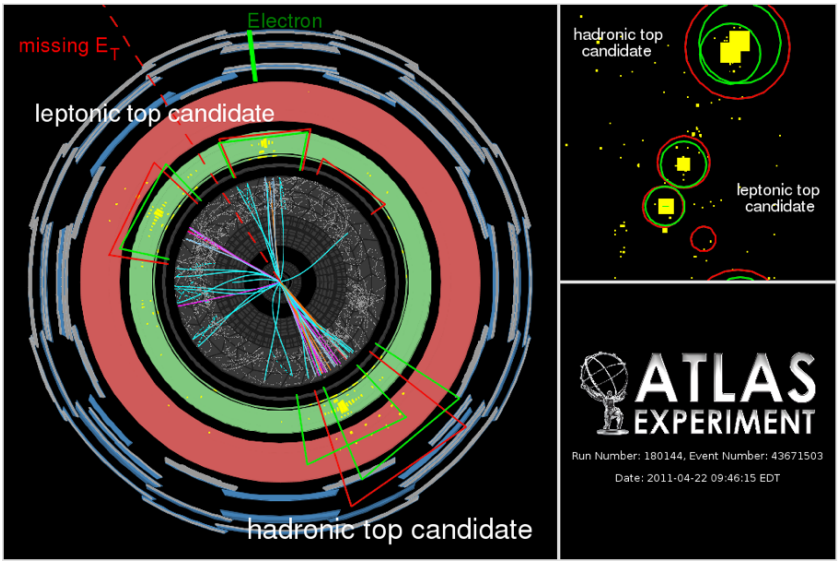


Figure 5.3: A candidate top event observed during the 2011 ATLAS data run, where a leptonically-decaying top candidate is inferred from lepton identification in the electromagnetic calorimeter, and a hadronically-decaying top candidate is inferred using information from both the electromagnetic and hadronic calorimeters. The yellow dots correspond to energy deposits in the calorimeters, and the coloured lines correspond to tracks reconstructed in the inner detectors. Jet reconstruction is a complex process which can utilise information from both tracking and energy deposition. From [78].

deposited sprays of energy and matter. An algorithm which is accurate at capturing the energy and matter and inferring the original hard-scatter process is essential, and this is where hadronic jet reconstruction can now be employed.

5.1.2 What is a Jet?

A jet is simply a proxy for particles (and their decay products) arising from the partonic hard-scatter, showering, and hadronisation described in the previous section. The purpose of jet reconstruction is to find a way of grouping collimated sprays of matter (as measured in e.g. the calorimeters) into a single physics object, with a well-defined four-momentum. Geometrically, this can be thought of as a cone in $r - \eta - \phi$ of the ATLAS coordinate system, with some radius given by $\Delta R = \sqrt{(\Delta\eta)^2 + (\Delta\phi)^2)}$. The radius of the jet can be tuned to capture an entire spray of radiation and matter in one region of the detector.

Jets can be formed from a variety of inputs, in both data and Monte Carlo simulations [79]. Monte Carlo hadrons simulated without any detector interactions are referred to as truth-level 'stable' particles (excluding muons and neutrinos), and can be clustered into a truth-level jet;

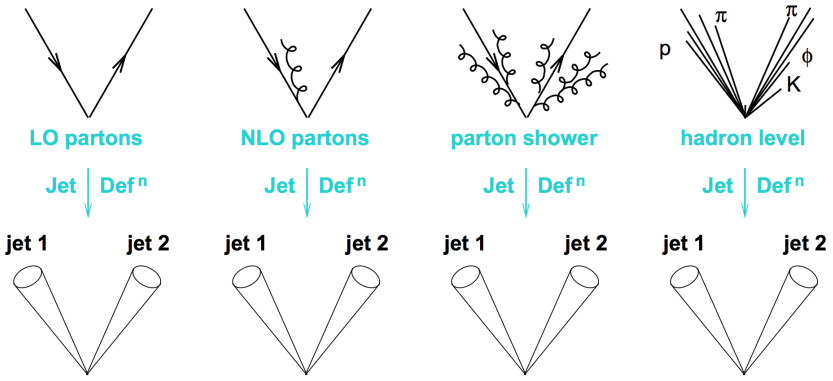


Figure 5.4: An illustration of some of the different Monte Carlo inputs which can be given to a jet reconstruction algorithm: partons, leading-order and beyond, can form parton jets; jets can also be formed from the partonic shower, and finally the full chain of hadrons at truth-level. This latter example gives rise to truth-level jets. From [80].

simulated detector readings can be clustered into a reconstruction-level jet; and of course actual data measurements in the tracker or calorimeter can be clustered to form so-called track or calorimeter jets respectively. The relationship between various classes of inputs and the resultant jets are summarised in Figure 5.4, where the inputs and the final jet are connected via the *jet definition* (next section). As already eluded to, any jets measured in data, such as calorimeter jets, are calibrated by comparing their four-vector properties with those of associated jets which are measured in Monte Carlo, such as these truth-level and reconstruction-level jet objects.

5.2 Jet Reconstruction

A *jet definition* is required to carry out jet reconstruction: this consists of a *jet algorithm* used to build a jet from some class of inputs, together with a *recombination scheme* [81] which takes the jet constituents and combines their energies and momenta in a meaningful way in order to determine the jet four-momentum, from which quantities such as the mass and energy of the jet can be extracted.

5.2.1 Jet Inputs

5.2.1.1 Monte Carlo Inputs

As previously mentioned, truth-(reconstruction-)level jets are formed from the application of a jet definition to truth-(reconstruction-)level hadrons/stable particles, both of which come

entirely from simulation.

Specifically reconstructed ('reco', or detector-level) jets are those obtained from signals that simulate the entire chain of the ATLAS detection and data storage process, from detection through interactions of particles with matter, to signal collection, the final electronic read-out and calibration of the inputs. Truth (also referred to as generator-level) jets are obtained by applying the particular jet reconstruction algorithm to the stable particles in simulated events that have a lifetime τ such that they propagate distances $c\tau > 10$ mm in the ATLAS detector.

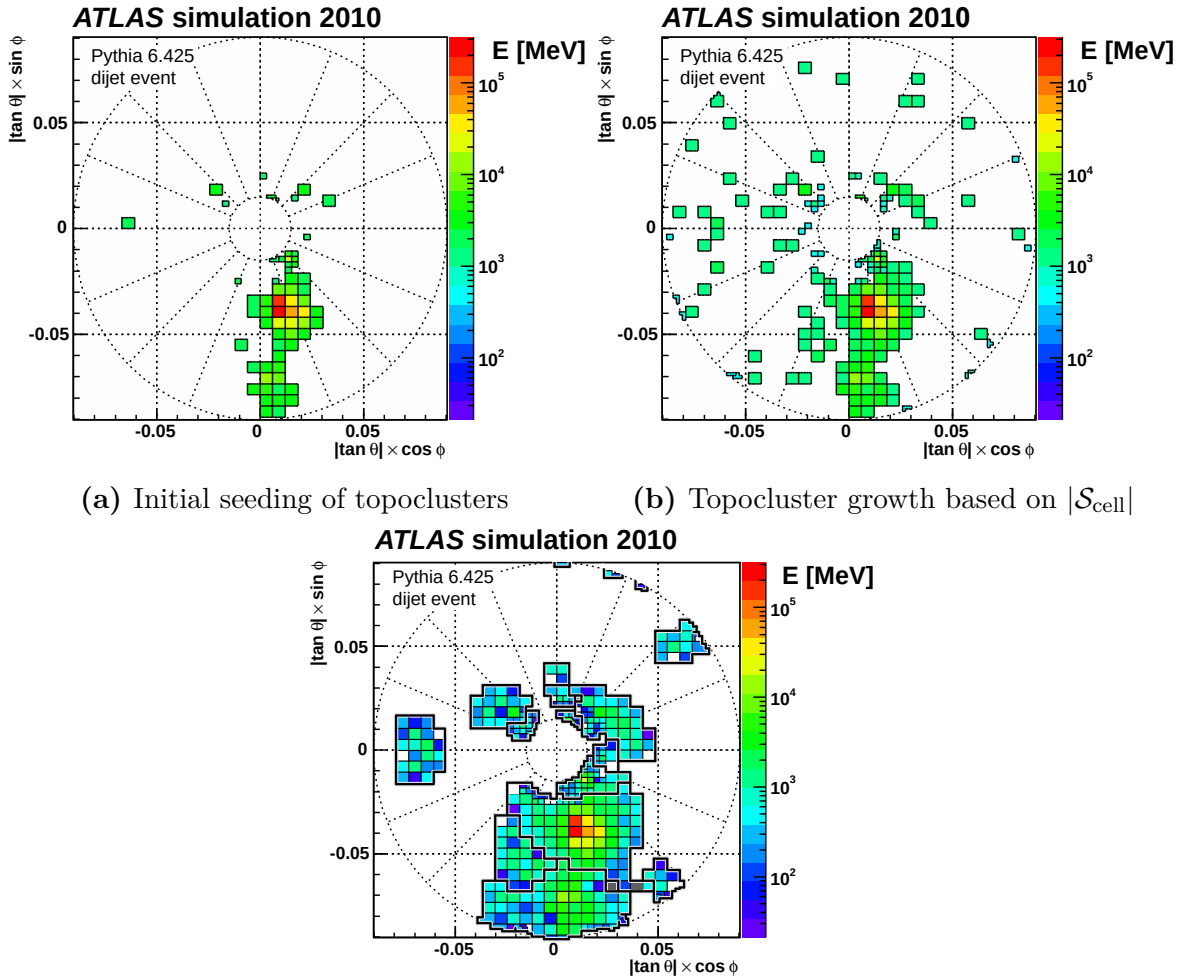
5.2.1.2 Topoclusters

Calorimeter jets are constructed from calorimeter energy deposits in the ATLAS $\eta - \phi$ plane, referred to as topological clusters, or *topoclusters* [82]. Alternative inputs to jets will be considered in subsequent chapters, but historically topoclusters have been the baseline calorimeter input to jet reconstruction throughout Runs 1 and 2 of ATLAS.

Topoclusters are formed from calorimeter cells [82], which are topologically connected in $\eta - \phi$ or on each calorimeter layer, in order to eventually extract an energy signal and suppress electronic noise from the calorimeter. Cells are grouped according to their *cell significance*, $\mathcal{S}_{\text{cell}}$, which is the ratio of the cell signal energy measured at the electromagnetic scale ($E_{\text{cell}}^{\text{EM}}$) and the energy arising from the background noise due to pileup events and detector electronics, denoted $\sigma_{\text{cell}}^{\text{pileup}}$ and $\sigma_{\text{cell}}^{\text{noise}}$ respectively. The cell significance is then defined to be:

$$\mathcal{S}_{\text{cell}} = \frac{E_{\text{cell}}^{\text{EM}}}{\sqrt{(\sigma_{\text{cell}}^{\text{pileup}})^2 + (\sigma_{\text{cell}}^{\text{noise}})^2}}. \tag{5.2}$$

Cells are clustered together based on their cell significance evaluated with respect to three parameters: S, P, and N. The topoclustering algorithm begins by isolating the seed cell which has the largest magnitude of cell significance $|\mathcal{S}_{\text{cell}}| > S$, where $S = 4$ is the default. These seed cells form the initial proto-cluster (illustrated in Figure 5.5a). All cells which are in the same calorimeter layer as the seeds, or which have significant overlap in $\eta - \phi$, and have a significance $|\mathcal{S}_{\text{cell}}| > N$, where $N = 2$ is the default, are added to the proto-cluster, causing it to grow (Figure 5.5b). Once there are no remaining cells satisfying this growth criterion, all additional neighbouring cells satisfying $|\mathcal{S}_{\text{cell}}| > P$, where $P = 0$ is the default, are grouped



(a) Initial seeding of topoclusters **(b)** Topocluster growth based on $|\mathcal{S}_{\text{cell}}|$

(c) Final topoclustering output

Figure 5.5: The topoclustering process, from initial seeding and growth, to the final topoclusters, in spherical polar coordinates based on the ϕ and θ position of the initial energy deposits in the calorimeter. This clustering is based on dijet events in ATLAS 2010 data-taking conditions simulated with Pythia 6 Monte Carlo. The thick black lines indicate the topocluster boundaries. From [82].

into the proto-cluster. Following this, the algorithm moves to the next seed cell and iterates again until there are no remaining cells to cluster. The above procedure does not prevent the formation of proto-clusters with several local energy maxima. In order to minimise such an effect, topoclusters that have at least two local energy maxima with $E_{\text{cell}}^{\text{EM}} > 500$ MeV are split between the two energy maxima and the cells at the boundary are shared based on the distance between each local maximum and the centre and energy of those boundary cells. The final topoclustering output is shown in Figure 5.5c, where the boundaries are indicated in dark lines.

The final product should be a set of topoclusters, where each topocluster is defined to be massless and have a total energy given by the weighted sum of the energies of the separate

constituent calorimeter cells. Each cluster has an associated η and ϕ coordinate given by the sum of the separate η and ϕ of each cell, weighted by the absolute value of the energy of each cell. The resulting topoclusters which are positive in energy are the inputs to the calorimeter jets. However, before these inputs can be provided, it is first necessary to calibrate these topoclusters in order to account for differences between the behaviour of the cells at the electromagnetic and hadronic calorimeter scales. Such differences are a consequence of the non-compensating nature of the calorimeter, which results in smaller measured signals for hadrons than for electrons or photons of the same energy upon entering the calorimeter. The correction is referred to as *Local Cluster Weighting* (LCW) [69, 82] calibration. Topoclusters which are subsequently corrected using this LCW calibration are referred to as LCTopoclusters (denoted `LCTopo`), whilst those which are not corrected are simply used for reconstruction at the electromagnetic scale only, and are referred to as EMTopoclusters (denoted `EMTopo`). Both `LCTopo` and `EMTopo` objects have been the standard ATLAS jet inputs until recently (see Chapter 7).

5.2.2 Jet Algorithms

Having defined the jet inputs for Monte Carlo and data, the next issue is to understand how those inputs can be combined in a theoretically sound manner. The kinematics of the resultant jet should try to match the kinematics of the underlying physics objects or particles that the same jet captures. The basic physics and computational requirements of a jet definition were laid out in the 1990 *Snowmass Accord* [83]. This is summarised by the requirement that a jet definition should be:

- Simple to implement in an experimental analysis;
- Simple to implement in a theoretical calculation;
- Defined at any order in perturbation theory;
- Yields finite cross-sections at any order in perturbation theory;
- Yields a cross-section that is relatively insensitive to hadronisation.

These requirements are automatically satisfied by requiring that the jet algorithm be *infrared and collinear safe*. The concept of *collinear safety* is shown in Figure 5.6. Each cartoon contains

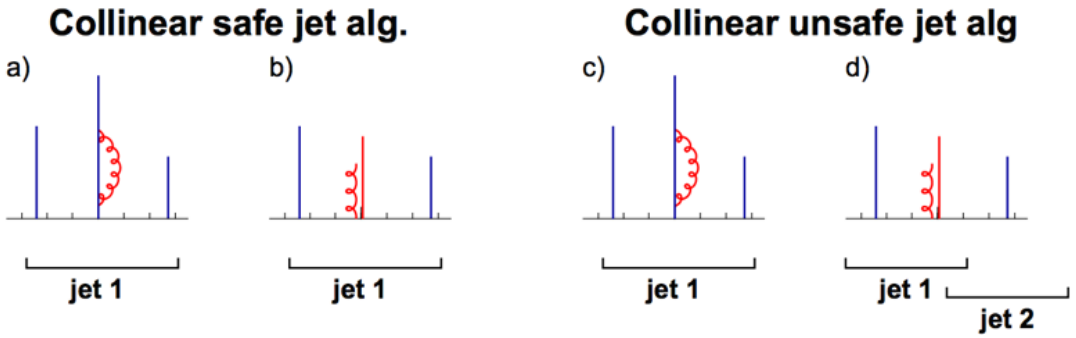


Figure 5.6: An illustration of collinear safe and unsafe jet reconstruction algorithms. From [68].

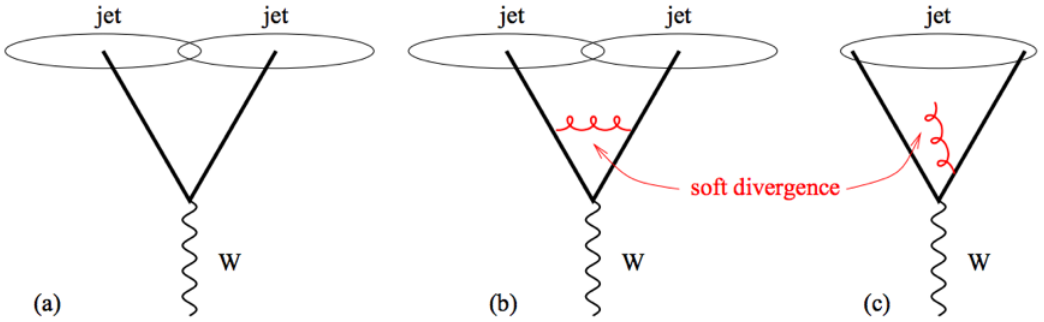


Figure 5.7: An illustration of infrared safe and unsafe jet reconstruction algorithms. From [68].

three lines, where the height of each line corresponds to some p_{T}. These straight lines represent jet inputs, such as truth-level particles. The red looping lines represent additional collinear[1] splittings that have been added to the event, such as gluons radiated off the propagating particle. Case a) is the initial state, and in b) a collinear splitting is added to the event. The resulting jets formed are unmodified by the presence of this extra splitting, and therefore the jet algorithm used is said to be collinear safe because the additional splitting doesn't change the final number of jets reconstructed when compared with a). In c) and d) the introduction of a collinear splitting changes the final jet multiplicity; the jet algorithm now reconstructs two jets in d) due to the extra splitting. In this case, the jet algorithm is collinear unsafe.

Figure 5.7 provides an example of infrared safety in jet algorithms. In a) a W boson's decay products are combined into jets with some algorithm, giving rise to a final state with two jets. In the cases of b) and c) an additional low energy emission, such as a gluon emission, is added to the event. In b) the additional emission doesn't change the final number of reconstructed

[1]Collinear in the sense that the angle of the emitted radiation is along (or nearly along) the direction of the original particle.

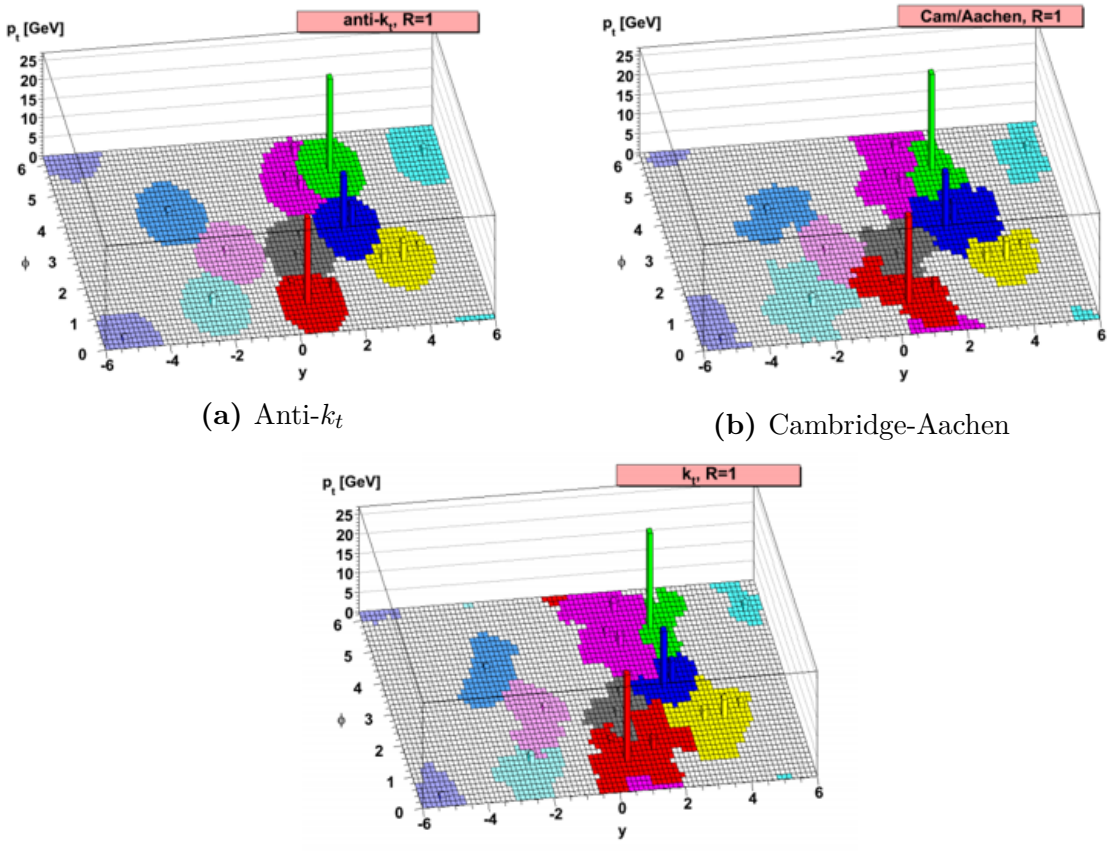
jets but in the case of c) the extra soft radiative emission has the effect of causing the decay products in the event to be combined into a single, larger jet. In b) the jet algorithm is infrared safe whilst in c) is it infrared unsafe.

Many early jet algorithms suffered from being collinear and/or infrared (IRC) unsafe. Promising early IRC safe algorithms, such as the MidPoint algorithm [84], involved sorting nearby particles or calorimeter energy deposits in to geometrical cones of some fixed size, based on some distance measure between those particles or energy deposits. Modern IRC safe jet algorithms used in ATLAS and CMS, referred to as *sequential recombination algorithms* [68, 85], combine different jet inputs i and j, with detector-frame coordinates (η_i, ϕ_i) and (η_j, ϕ_j) respectively. The combination is based on some distance measure d_{ij}, and produces a jet of some radius R. In these recombination algorithms the distance measure is of the form:

$$d_{ij} = \min(p_{\mathrm{T},i}^{2p}, p_{\mathrm{T},j}^{2p}) \frac{\Delta R_{ij}^2}{R}, \tag{5.3}$$

where $\Delta R_{ij} = \sqrt{(\phi_j - \phi_i)^2 + (\eta_i - \eta_j)^2}$ and the distance between the i^{th} input and the beam is, by construction, $d_{iB} = p_{\mathrm{T},i}^{2p}$. The p parameter takes values of -1, 0, 1, corresponding to three different jet algorithms. These are referred to as the anti-k_t [85], Cambridge-Aachen [86], and k_t [87] algorithms respectively. The anti-k_t algorithm combines the hardest (highest p_{T}) jet inputs first, building the core of a jet containing the highest momentum density and then working outwards in momentum space to build the remains of the jet from softer inputs. This gives rise to highly conical jets in $\eta - \phi$. The Cambridge-Aachen approach uses a purely geometric distance measure, based only on the ΔR_{ij} between the inputs, and the k_t works in a similar way to the anti-k_t, but builds up the jet from the softest inputs first. An example of the jets formed from these three algorithms is provided in Figure 5.8.

In ATLAS the anti-k_t algorithm is used to build jets of size $R = 0.4$ from electromagnetic scale topoclusters (`EMTopo`), giving rise to so-called `AntiKt4EMTopo` jets, or small-radius jets. Truth-level particles from Monte Carlo simulation are also used to construct truth jets, and in ATLAS the standard $R = 0.4$ jets made from truth particles are referred to as `AntiKt4EMTruth` jets. These have been used in ATLAS throughout Run 1 and Run 2. Anti-k_t was chosen because it provides geometrically conical jets. It has been implemented in the FastJet [88] jet reconstruction software package, and is computationally fast when compared with many other



(a) Anti-k_t

(b) Cambridge-Aachen

(c) k_t

Figure 5.8: The three different classes of sequential recombination algorithms, used to construct $R = 1.0$ jets from inputs defined in the ϕ and rapidity (y) frames for some simulation of events. The final jets are signified by the coloured objects in the $y - \phi$ plane, and the inputs have heights proportional to their transverse momenta. From [68].

approaches to jet reconstruction, (see Figure 5.9), with a computation time for N inputs scaling like $\mathcal{O}(N \log N)$ [68, 88].

The anti-k_t algorithm proceeds as follows:

1. From the list of all jet inputs (topoclusters, or truth-level particles, ...) determine all possible d_{ij} between each input and all d_{iB} between an input and the beam (see Figure 5.10).
2. If the $\min(d_{ij}) < d_{iB}$, combine the i and j inputs and return to the start.
3. Otherwise, declare input i to be a jet, remove it from the list of inputs, and return to the start.
4. Iterate until there are no remaining inputs in the list, such that all of the original inputs have been either classified as a jet or as part of a jet.

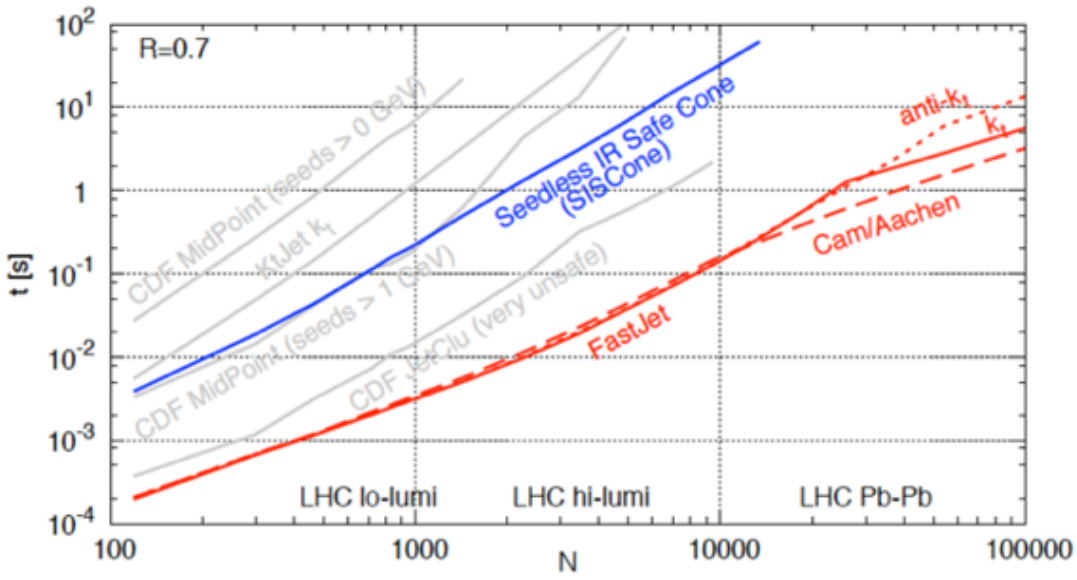


Figure 5.9: The number of input clusters, N, against the computational time, t, in seconds, for various different jet reconstruction algorithms. The sequential recombination algorithms are the most computationally efficient for a given number of input clusters, often by several orders of magnitude. From [68].

Once the jet has been constructed using the jet algorithm, the final step in the jet definition is to apply a recombination scheme which provides the jet with a meaningful four-momentum. In ATLAS the so-called E-scheme [81, 88] is adopted, whereby the total jet four-momentum is calculated by summing the four-momenta of each of the jet inputs. How one can extract important quantities from the jet four-momentum forms an important part of this thesis, and in particular forms the basis of the jet mass discussion in the next chapter.

5.2.3 Jet Grooming

So-far the reconstruction of $R = 0.4$ jets has been discussed, as these are the standard jets adopted for many analyses in ATLAS. It is possible to construct jets of any radius, and in ATLAS the second type of jet considered is the so-called large-radius jet. Typically a radius of size $R = 1.0$ is used.

The need for large-radius jets in ATLAS can be understood by considering the decays of heavy particles, such as W/Z bosons or top quarks. Consider the hadronically-decaying W boson, $W \rightarrow q\bar{q}$. It is reasonable to imagine that each of the quarks produced in this decay would undergo hadronisation and the resultant bound states would be reconstructed as a single `AntiKt4EMTopo`

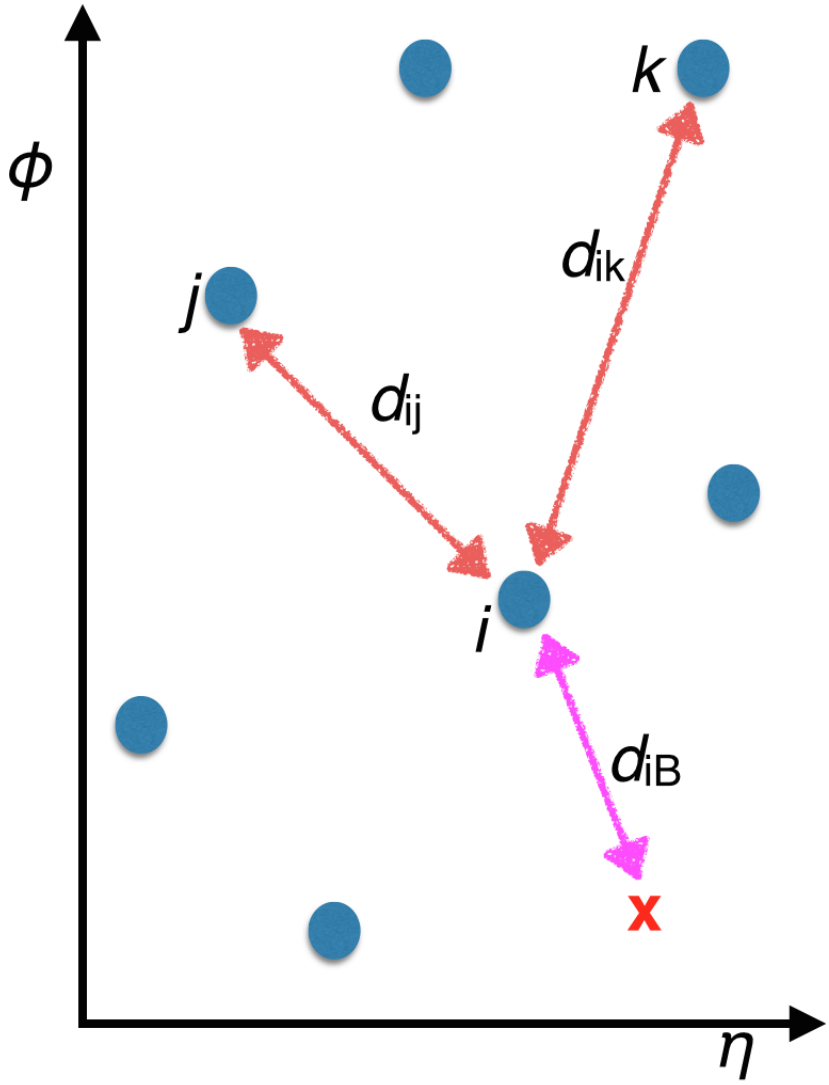


Figure 5.10: A cartoon summarising the basic arrangement of clusters in the anti-k_t algorithm adopted by ATLAS. Clusters are grouped in $\eta - \phi$ space based their separation from other clusters and from the beam itself. This distance measure depends on the geometrical separation of clusters in $\eta - \phi$, the inverse square of the cluster transverse momenta, and the desired radius of the final jets.

jet per initiator quark. However, the very high centre-of-mass energy of the pp collisions in the LHC often give rise to vector bosons with a much higher transverse momentum than their mass. In this case, the decay products of such a W would be more collimated along the direction of the W's momentum. In a two-body decay such as $W \to q\bar{q}$ the angular separation $\Delta R_{q\bar{q}}$ between the decay products depends on the W mass and transverse momentum according (approximately) to:

$$\Delta R_{q\bar{q}} \approx \frac{1}{\sqrt{x(1-x)}} \frac{m_W}{p_{\mathrm{T}}}, \tag{5.4}$$

where x is the momentum fraction carried by each of the two quarks. The decaying quarks will have approximately the same momentum, and so the above expression reduces to $\Delta R_{q\bar{q}} \approx$

$2m_W/p_{\mathrm{T}}$. For a sufficiently boosted W boson, the angle between the decay products can be small enough for those decaying quarks to be exceptionally collimated. Ws produced at a p_{T} in the $400 - 500$ GeV range will decay in to quarks with an angular separation around 0.4, and so resulting parton showers of those quarks will overlap such that they can no longer be separately resolved in a jet of size $R = 0.4$. Instead a jet of size $R = 1.0$ can be used to capture the entire hadronic decay of the W. It is for this reason that large-radius jet reconstruction is crucial when dealing with heavy particles propagating through the detector at high transverse momenta, and particles of this nature are said to have a high Lorentz boost (often deemed to be 'boosted'). Since these $R = 1.0$ jets are typically used to reconstruct the hadronic decay of heavy particles, they need to be built from topoclusters which are calibrated to the hadronic scale, and consequently are build from LCW-calibrated topoclusters. We refer to these as `AntiKt10LCTopo` jets in the ATLAS nomenclature.

One implication of using larger jets is that they are more susceptible to contamination from several sources: radiation from out-of-time pileup in the high-luminosity conditions of the LHC; multiple parton interactions from the underlying event; and the presence of softer particles from showering effects. These soft contaminations can diminish the measured resolution of the jet, and need to be combated. One popular approach for dealing with this is *jet grooming* [68], where an algorithm is applied which attempts to remove the soft particle contamination in the $R = 1.0$ jet, without removing the component from the hard-scatter which is endowed with the meaningful physics. In Chapter 7 the concept of jet grooming and its implications for jet performance will be discussed in detail, and several different jet groomers will be presented and analysed. Section 5.3.2.1 will focus on one example of a jet groomer, *jet trimming* [89], which has been used as the *de-facto* jet grooming algorithm in ATLAS [90] until very recently.

5.3 Calibration at the Jet Energy Scale (JES)

Having presented the method of reconstructing both small and large-radius jets in ATLAS, the next important step is to calibrate these objects beyond the EM and LCW schemes. This section briefly summarises the *jet calibration chain* for the *jet energy scale* (JES) [91] calibration. The JES derives calibrations which first corrects the energy of a reconstruction-level jet in

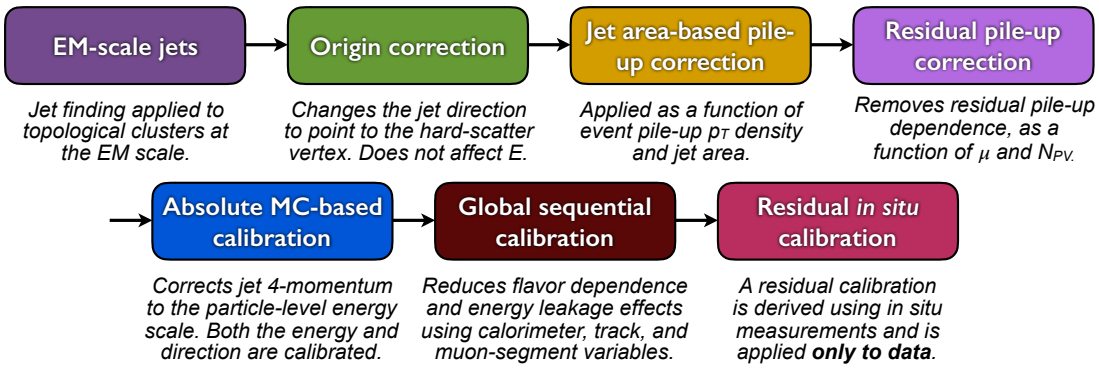


Figure 5.11: A summary of the complete jet energy scale calibration chain for `AntiKt4EMTopo` jets in ATLAS. The initial origin correction is followed by the Monte Carlo JES calibration. After this a correction to the flavour-dependence of the jet is applied using Monte Carlo, before a final *in situ* calibration is determined using events measured in ATLAS data. From [91].

Monte Carlo to the original truth-level jet, followed by correcting the jets measured in data with the energy of that reconstruction-level jet.

5.3.1 Small-Radius Jets

The complete chain for small-R jets is shown in Figure 5.11, and each of these components will now be briefly summarised.

5.3.1.1 Origin Correction

Jets are reconstructed from inputs defined in the detector reference frame, and are therefore pointing in some direction with respect to the centre of the ATLAS detector. The purpose of the origin correction is to recalculate the four-momentum of all jets such that they are pointing in some direction now with respect to the primary vertex associated with the hard-scatter event, while keeping the energy of the jet constant. This has the effect of improving the η resolution in Monte Carlo, defined to be the residual of reco-level jet η and the truth-level jet η, $\eta^{\text{reco}} - \eta^{\text{true}}$, where η^{reco} is the origin-corrected, reconstruction-level jet η. This should be distinguished from the η in the detector frame of reference, which is also used in jet calibrations and shall be denoted η^{det}. The origin correction is applied in this way to jets in the 2015 and 2016 ATLAS data-set, but from 2017 onwards this origin correction was automatically applied to the clusters before any jet reconstruction takes place (and would therefore not form part of the jet calibration chain).

5.3.1.2 Jet Area and Residual Correction

The purpose of the jet area and residual corrections is to remove any additional energy which comes from particles and radiation produced from in-time and out-of-time pileup, and would contribute significant background noise to the jet reconstruction in the event.

The *jet area* [92, 93] (also known as the catchment area), A, is a way of defining the region of $\eta - \phi$ space where contributions from pileup and the underlying event affects the kinematics of the jet. In order to form an area out of the point-like topocluster jet inputs, the *active area* method [92] is employed in ATLAS jet reconstruction. [2] In this approach, a large set of randomly distributed particles (referred to as *ghosts*), each with an infinitesimal transverse momentum, is added to the list of topocluster inputs. The resulting four-vector of the final jet is completely unmodified due to the use of the infrared-safe anti-k_t algorithm, and the active area of the jet is the ratio of the number of ghosts clustered into the jet to the number of ghosts per unit area in the $\eta - \phi$ grid of ATLAS, ν_g. A dense coverage of ghosts is injected into the jet $\eta - \phi$ grid of topoclusters, and the final area is determined by averaging over many measurements with different sets of ghosts used. Putting this together, the full four-vector jet area is then:

$$A_\mu = \lim_{\nu_g \to \infty} \left\langle \frac{1}{\nu_g \langle g_t \rangle} \sum_{g_i \in J} g_{\mu i} \right\rangle, \tag{5.5}$$

where the summation runs over all ghosts g_i clustered within a jet J, and $\nu_g \langle g_t \rangle$ is the ghost transverse momentum density. The topoclusters and ghosts will all be clustered to form the jet. As a result, due to pileup contributions in an event being (approximately) randomly-distributed, the resulting jet area will provide a measure of how susceptible the jet is to pileup contributions.

The final ingredient to the area calculation is the p_{T} density of pileup events, which is the transverse momentum per unit area added to the event by pileup. The density, ρ, is calculated from the median of the p_{T}/A distribution of all the i reconstructed $R = 0.4$ jets in the event:

$$\rho = \mathrm{median} \left(\frac{p_{\mathrm{T},i}}{A_i} \right). \tag{5.6}$$

[2] A simplification of the jet area comes from visualising the reconstructed jet as a geometrical cone, of radius R, such that jet area is then simply πR^2.

For simplicity, the ρ was taken as a constant with respect to the jet η until recent developments which incorporated the jet η-dependence into the final ρ calculation.

The pileup correction is now applied to the jet four-vector in its entirety, producing the new corrected jet four-momentum, p_μ^{corr}:

$$p_\mu^{\text{corr}} = p_\mu - \rho A_\mu, \tag{5.7}$$

and this reduces to a single correction on the reconstruction-level jet p_{T},

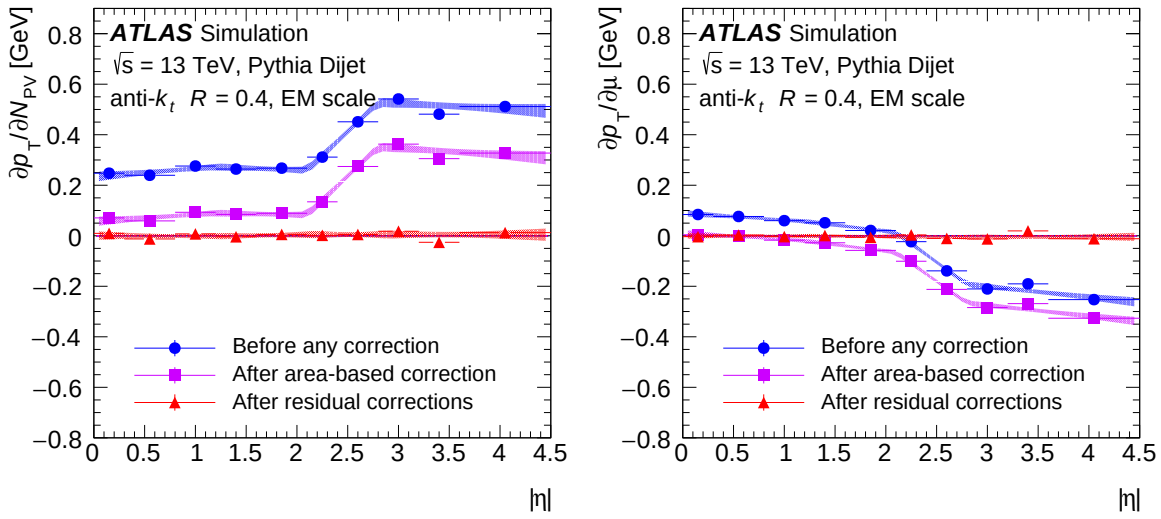
$$p_{\text{T}}^{\text{reco,corr}} = p_{\text{T}}^{\text{reco}} - \rho \times A. \tag{5.8}$$

Following the jet area subtraction, two residual corrections to the p_{T} are applied to remove any additional dependence on the in- and out-of-time pileup contributions. These additional contributions can arise from high-occupancy regions of the calorimeter which are not fully considered in the ρ calculation. The residual p_{T}-dependence on the in-time pileup is investigated by considering the number of primary vertices in per event, N_{PV}, whilst out-of-time pileup dependence is inferred from the number of interactions per bunch-crossing, μ. The residual p_{T} is defined as the difference between the reconstructed jet p_{T} and the p_{T} of a geometrically matched truth jet (which is by definition insensitive to pileup). The truth jet is geometrically matched to be within a $\Delta R < 0.3$, where $\Delta R = \sqrt{(\eta^{\text{reco}} - \eta^{\text{true}})^2 + (\phi^{\text{reco}} - \phi^{\text{true}})^2}$.

The residual p_{T} dependencies on N_{PV} and μ are observed to be roughly linear in nature. The linear fits are parametrised by the gradients α and γ respectively, where different α and γ terms are determined for different values of the jet η. Once the α and γ coefficients have been smoothly extracted for different jet η they can be used to derive the fully-corrected jet $p_{\text{T}}^{\text{reco,corr}}$, which now takes into account the area subtraction and the N_{PV} and μ-dependence:

$$p_{\text{T}}^{\text{reco,corr}} = p_{\text{T}}^{\text{reco}} - \rho \times A - \alpha \times (N_{\text{PV}} - 1) - \gamma \times \mu. \tag{5.9}$$

The results of applying first the jet area correction, and then the full residual correction, are shown in Figure 5.12. The jet p_{T} dependence on N_{PV} and μ is removed by applying these pileup and residual corrections, producing new corrected jets which are ready to be calibrated



(a) Jet p_{T} N_{PV}-dependence. **(b)** Jet p_{T} μ-dependence.

Figure 5.12: The N_{PV} and μ dependencies before and after applying the full jet area and residual corrections. Applying the full residual correction, where the correction is derived in bins of the jet η^{det}, almost entirely removes the pileup dependence. These corrections are derived for `AntiKt4EMTopo` jets using Monte Carlo dijet events produced in PYTHIA, at reconstruction and truth-level. From [91].

at the jet energy scale.

5.3.1.3 Monte Carlo η-JES Correction

Having performed the jet area subtraction on the $R = 0.4$ jets, we are now in a position to derive corrections to the energy and η (beyond the initial origin correction) of those same jets. For these $R = 0.4$ jets calculated at the Monte Carlo reconstruction-level, the purpose of the *jet energy scale* (JES) calibration is to correct the jet energy determined at the reconstruction-level such that it is consistent with energy determined from geometrically matched truth-level jets, to within a few percent of the true value.

The *jet energy response*, $\mathcal{R}_E$, is defined as the ratio of the jet energy determined at the detector-level (i.e. as determined through the Monte Carlo-simulation of the ATLAS calorimeters) to the energy calculated from truth-level jets which are geometrically matched to be within $\Delta R < 0.3$ of the reconstruction-level jets. The response function is defined as:

$$\mathcal{R}_E = \frac{E^{\mathrm{reco}}}{E^{\mathrm{true}}}, \tag{5.10}$$

where 'reco' and 'true' label the MC reconstruction-level and truth-level quantities respectively. The jet responses are different to (and typically smaller than) unity due to the non-compensating nature of the calorimeter, the presence of inactive material in the detectors, detector noise thresholds giving rise to some signal losses, the 'leakage' of energy due to particles falling outside the calorimeter, and 'out-of-cone' effects where particles which should fall into the detector-level jet cone do not fall into the truth-level cone. The simulated topoclusters, from which the detector-level jets are reconstructed, may also be split and this has a direct effect on the mass of the reconstructed jet. Implications for the jet mass scale shall be discussed in the next chapter.

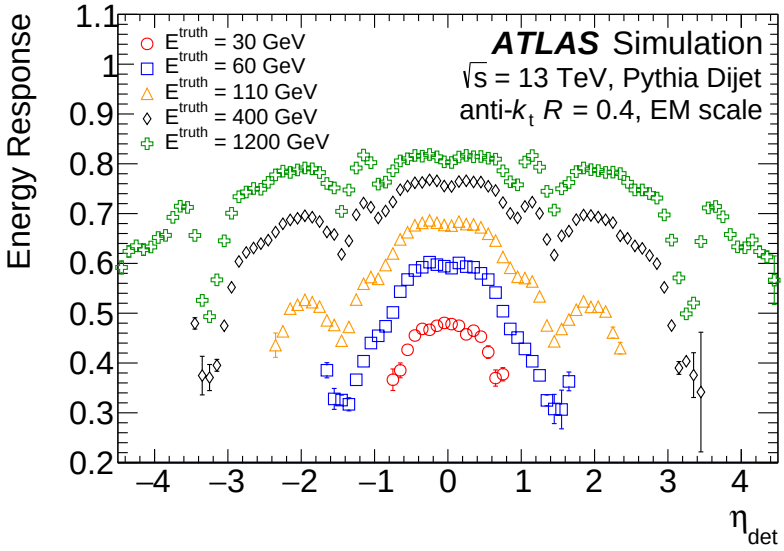
The calibrations determined for $R = 0.4$ jets are derived using samples of QCD dijets produced in simulated LHC pp conditions at collision energies of 13 TeV. The jet energy is calibrated to truth-level using the technique of *numerical inversion*. An energy calibration is determined by first calculating the jet energy response in bins of the η^{det} and E^{true} and extracting the mean of a Gaussian fit to each response. The average response as function of η^{det}, for different bins in the truth jet energy, is depicted in Figure 5.13a for `AntiKt4EMTopo` jets.

The mean jet energy responses are parametrised in E^{true}, but numerical inversion transforms them into a function of E^{reco}, for each jet η^{det} bin . This resulting energy response is then fitted according to a multi-logarithmic 'groom' function, again for each η^{det} bin. This groom function is a parameterisation of the energy loss for particles showering in a calorimeter. This new fitted quantity is referred to as the *calibration function*, $\mathcal{F}^{\text{calib}}_{m,\eta}(E^{\text{reco}})$. There is one calibration function per η^{det} bin. The JES-calibrated jet energy is then defined:

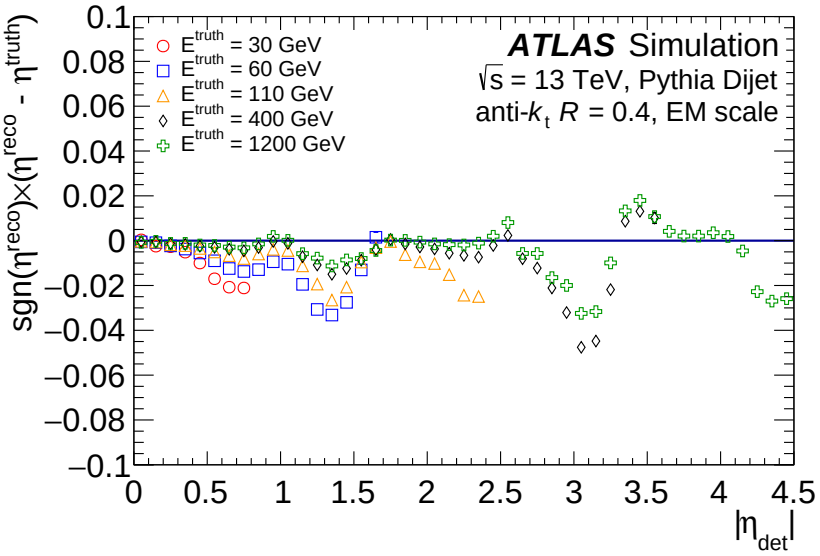
$$E^{\text{reco}}_{\text{JES}} = \frac{E^{\text{reco}}}{\mathcal{F}^{\text{calib}}_{m,\eta}}(E^{\text{reco}}). \tag{5.11}$$

The second stage of the η-JES calibration corrects for the differences between the reconstruction and truth-level jet η. As the η is a dimensionless quantity, a calibration is derived based on the residual η and not the η response. The residual considered is:

$$\Delta\eta = \eta^{\text{reco}} - \eta^{\text{true}}. \tag{5.12}$$



(a) Average jet energy response, before any η-JES correction.



(b) Average residual jet η, before any η-JES correction.

Figure 5.13: The average jet energy response and residual jet η values, before any η-JES correction has been applied, as function of the detector η. The JES correction is obtained from the inverse of the energy response. From [91].

The final $\Delta\eta$ residuals are determined as a function of E^{true}, and are applied to the JES-calibrated jets in bins of the same E^{true}. The typical residual correction to the jet η is small. Putting all of this together, the jets have now been origin and pileup corrected, and calibrated at the Monte Carlo reconstruction-level to truth-level, in both energy and η. After correcting for the η-JES, the closure[3] of the `AntiKt4EMTopo` jets is well within 5 % of unity, and within 2 % for the majority of the jet energy and η phase-space, particularly at higher jet energies.

[3] A non-closure is defined to be the deviation of the jet response *after* calibration from unity, thus a small degree of non-closure means that the jet calibration method closes well.

5.3.1.4 Global Sequential Calibration

Following the η-JES calibration, an additional *Global Sequential Calibration* [69, 91] (GSC) is applied to correct for the additional differences in the energy distribution and calorimeter shower shape of the jets, and to also improve the jet resolution. These effects arise from the dependence of the jet on different quark flavours, including the dependence on the quarks and gluons initiating the jet, as well as the effect of jets passing through the calorimeter subsystems and being misidentified as muons (*punch-through*). A quark-initiated jet will typically include a few hadrons contributing a large fraction of the total jet p_{T}, and will therefore penetrate deeper into the calorimeter. However, a gluon-initiated jet will typically carry more soft radiation and lower p_{T} constituents, leading to a smaller measurable response in the calorimeter and a shower shape with a larger transverse profile, i.e. a larger 'jet width'. For `AntiKt4EMTopo` jets, five observables have been identified, based on Run 1 ATLAS studies [69], which help to reduce the JES resolution and flavour-dependence of the jets when applied as a correction through this GSC.

For each observable a jet p_{T} response is constructed from QCD dijets Monte Carlo events, and this response is numerically inverted in order to derive a correction to the jet four-momentum as a function of the jet $p_{\mathrm{T}}^{\mathrm{true}}$ and η^{det}. The effect of each correction is to remove the dependence of the jet response on that particular observable, whilst ensuring that the jet energy at the electromagnetic scale remains unmodified. The corrections from these responses are applied sequentially:

- f_{Tile0}: the fraction of jet energy measured in the first layer of the hadronic tile calorimeter, extending up to $|\eta^{\mathrm{det}}| < 1.7$.
- f_{LAr3}: the fraction of jet energy measured in the third layer of the electromagnetic LAr calorimeter, extending up to $|\eta^{\mathrm{det}}| < 3.5$.
- n_{track}: the inner detector track multiplicity ($|\eta^{\mathrm{det}}| < 2.5$), where the tracks must have p_{T} > 1 GeV, and be associated with a jet.

- $\mathcal{W}_{\text{track}}$: the average p_{T}-weighted transverse difference in the $\eta - \phi$ plane between the axis (defining the geometrical centre) of the jet and all associated inner detector tracks with $p_{\text{T}} > 1$ GeV.

- n_{segments}: the multiplicity of muon tracks associated with the jet, extending up to $|\eta^{\text{det}}| < 2.7$.

The Monte Carlo jet p_{T} response as a function of each of the five GSC observables is shown in Figure 5.14. Features such as the spikes at $f_{\text{Tile0}} = 0$ come from the leakage of jet energy outside of the electromagnetic calorimeter. In spite of such features, the closure is observed to be within 2 % of unity after the full GSC is applied.

5.3.1.5 *In situ* JES

The *in situ* calibration is now performed in order to correct the jet energy scale as measured in data to the reconstruction-level Monte Carlo jet energy scale. This correction is necessary so that the jet energy as measured in real ATLAS data can be faithfully compared with the simulated jet energy, both at reconstruction-level and truth-level. The *in situ* calibration consists of a chain of sub-calibrations:

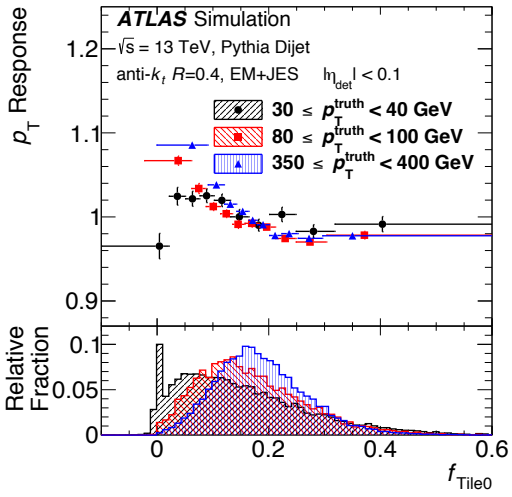
- Dijet η-intercalibration.

- Z + jet balance JES calibration (20 GeV $< p_{\text{T}} <$ 500 GeV).

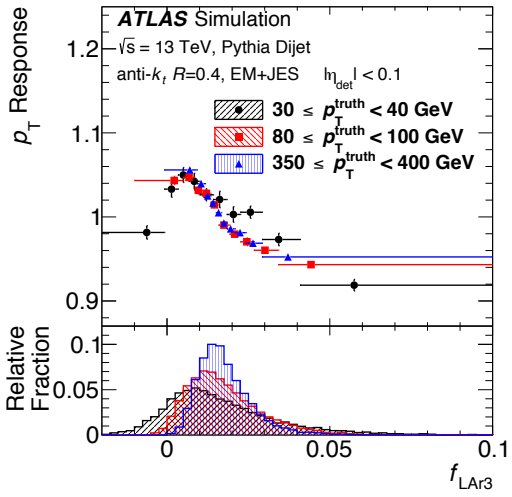
- γ + jet balance JES calibration (36 GeV $< p_{\text{T}} <$ 950 GeV).

- Multi-jet balance JES calibration (300 GeV $< p_{\text{T}} <$ 2000 GeV).

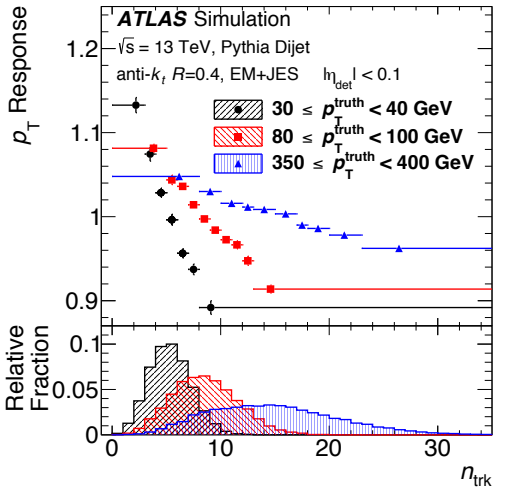
The η-intercalibration is derived using both data and Monte Carlo; it uses dijet events with jets that are calibrated up to the GSC, selecting central jets ($|\eta^{\text{det}}| < 0.8$) in order to calibrate forward jets ($0.8 < |\eta^{\text{det}}| < 4.5$). In each selected dijet event a well-calibrated *reference jet* is isolated in the central region and is balanced in its transverse momentum against a *probe jet* in the forward region, where these two jets have transverse momenta $p_{\text{T}}^{\text{ref}}$ and $p_{\text{T}}^{\text{probe}}$ respectively. Defining $p_{\text{T}}^{\text{avg}}$ as the average of these reference and probe transverse momenta allows one to define the asymmetry, $\mathcal{A}$, and relative jet response $\mathcal{R}_{\text{int}}$:



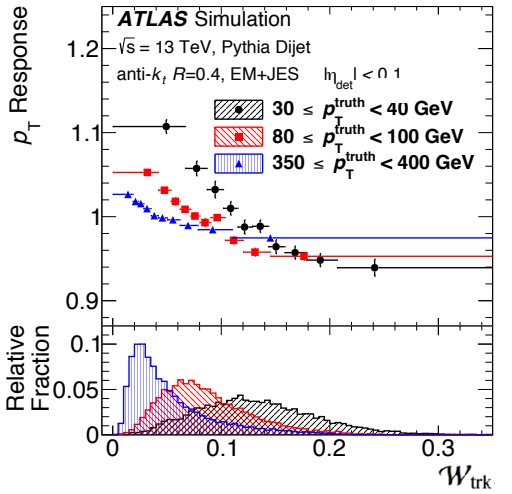
(a) The jet p_{T} response as a function of f_{Tile0}.



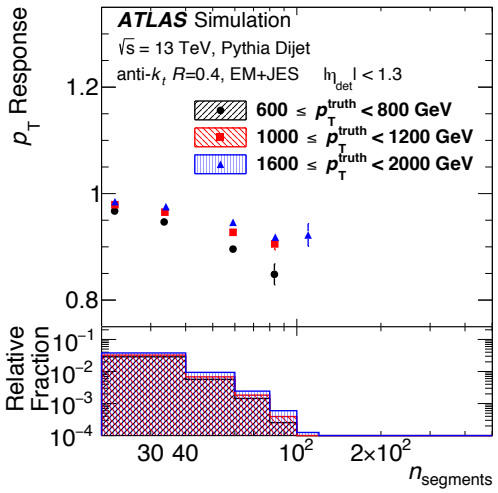
(b) The jet p_{T} response as a function of f_{LAr3}.



(c) The jet p_{T} response as a function of n_{track}.



(d) The jet p_{T} response as a function of $\mathcal{W}_{\mathrm{track}}$.



(e) The jet p_{T} response as a function of n_{segments}.

Figure 5.14: The five observables used to derive the GSC correction for `AntiKt4EMTopo` jets, using dijets events simulated in PYTHIA, at both the truth and reconstruction-levels. These corrections are applied sequentially, with the Tile0 correction applied first and the muon segments correction apply at the end. The responses shown are before any GSC-correction is applied. From [91].

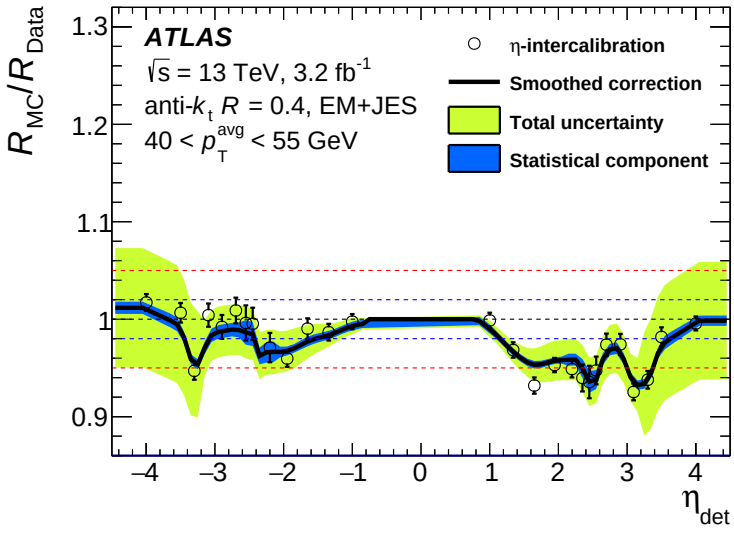


Figure 5.15: The η-intercalibration correction as a function of the detector η, for `AntiKt4EMTopo` jets using the 2015 ATLAS data-set. The smallest corrections are in the central region, which is dominated by statistical uncertainties. The systematic uncertainty components, such as the difference between the POWHEG + PYTHIA 8 (used as the nominal sample) and SHERPA dijets Monte Carlo, are discussed in Section 5.3.3. The final correction (thick black line) is obtained from a smoothed fit to the intercalibration points (open circles). These corrections are obtained in different intervals of $p_{\mathrm{T}}^{\mathrm{avg}}$: 40 − 55 GeV in this case. From [91].

$$\mathcal{A} = \frac{p_{\mathrm{T}}^{\mathrm{probe}} - p_{\mathrm{T}}^{\mathrm{ref}}}{p_{\mathrm{T}}^{\mathrm{avg}}}, \quad \mathcal{R}_{\mathrm{int}} = \frac{2 + \langle \mathcal{A} \rangle}{2 - \langle \mathcal{A} \rangle} = \left\langle \frac{p_{\mathrm{T}}^{\mathrm{probe}}}{p_{\mathrm{T}}^{\mathrm{ref}}} \right\rangle. \tag{5.13}$$

The relative jet response, studied in data, POWHEG + PYTHIA 8 [94], and SHERPA [95] dijet events is shown in Figure 5.15 for different detector η, using the 2015 ATLAS data-set amounting to 3.2 fb^{-1} of integrated luminosity. Imbalances in the jet response are attributed to differences between the calorimeter regions, and the final calibration factors, c_{int} are obtained from the ratio of the response studied in the Monte Carlo dijet events to the dijet events measured in data:

$$\frac{1}{c_{\mathrm{int}}} = \frac{\mathcal{R}_{\mathrm{int}}^{\mathrm{data}}}{\mathcal{R}_{\mathrm{int}}^{\mathrm{MC}}}. \tag{5.14}$$

Jets which have then been calibrared up to the η-intercalibration stage are used in the Z + jet and γ + jet balance, the next stage of the *in situ* calibration. There are two approaches to deriving each of these calibrations: *direct balance* and *missing projection fraction* (MPF). In the direct balance method the transverse momentum balance, $\mathcal{R}_{\mathrm{DB}}$, is calculated from the ratio of the transverse momentum of the leading jet, which is back-to-back with some angle $\Delta\phi$ to a photon or leptonically-decaying Z boson reference object, to the transverse momentum of that same reference object (see the illustration in Figure 5.16):

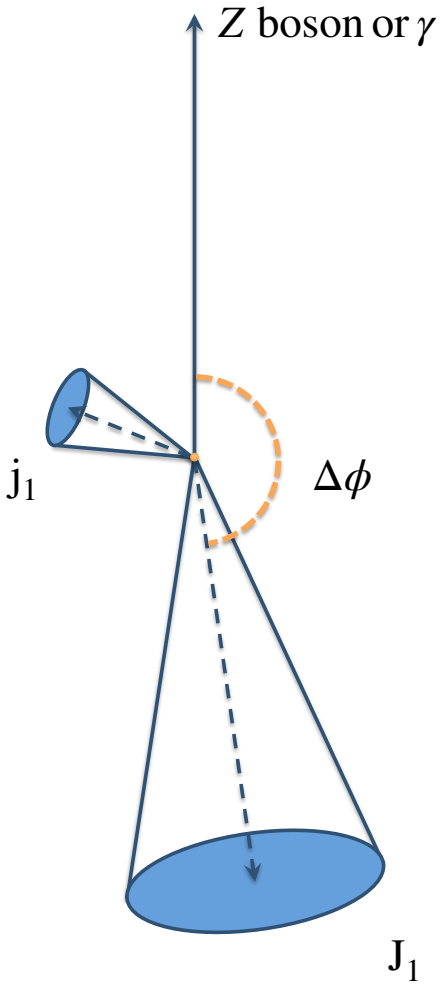


Figure 5.16: The geometry of an event considered in the direct balance method. The balance is calculated from the transverse momentum ratio of the leading jet J_1 and some well-defined reference Z boson or photon, at some angle $\Delta\phi$ to J_1. From [91].

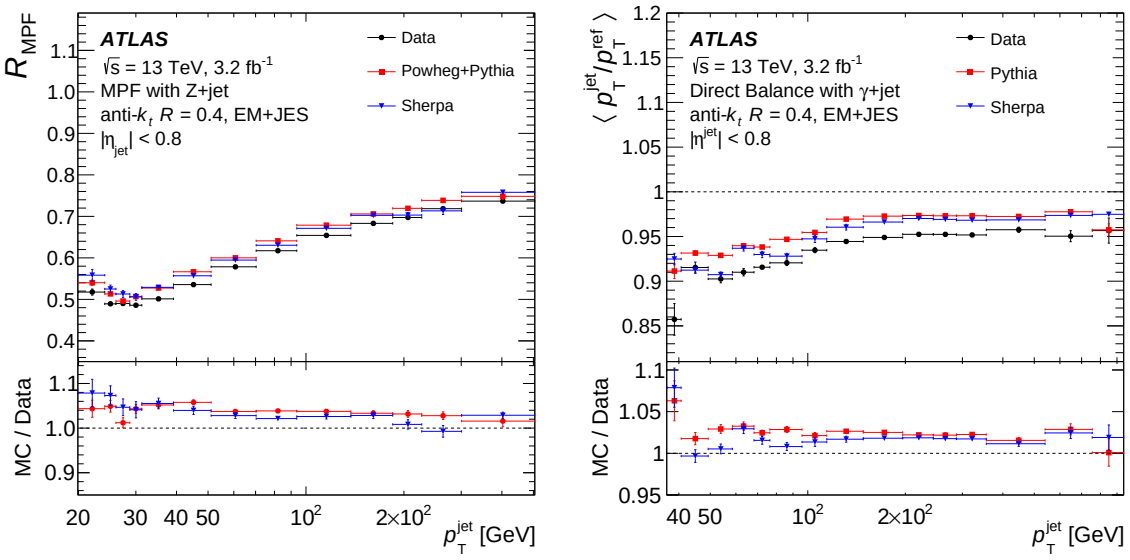
$$\mathcal{R}_{\mathrm{DB}} = \frac{p_{\mathrm{T}}^{\mathrm{jet_1}}}{p_{\mathrm{T}}^{\mathrm{ref}}}, \tag{5.15}$$

where $p_{\mathrm{T}}^{\mathrm{ref}} = p_{\mathrm{T}}^{V} |\cos(\Delta\phi(V, \mathrm{jet}_1))|$ and V refers to the reference photon or Z boson.

In the MPF approach the well-defined reference Z or photon is instead balanced against the entire hadronic recoil of the event[4]. In this way the MPF response represents the calorimeter transverse momentum response to the total hadronic recoil per event:

$$\mathcal{R}_{\mathrm{MPF}} = \left\langle 1 + \frac{\hat{n} \cdot \mathbf{E}_{\mathrm{T}}^{\mathrm{miss}}}{E_{\mathrm{T}}^{\mathrm{ref}}} \right\rangle, \tag{5.16}$$

[4]The hadronic recoil corresponds to the set of all other particles produced in the event, other than the reference object, typically from the hadronisation of proton remnants and gluon ISR.



(a) MPF, Z + jet **(b)** Direct balance, γ + jet

Figure 5.17: The MPF and direct balance methods, using a Z and γ reference respectively, as a function of the jet transverse momentum, measured in the central region of the detector, $|\eta^{\text{det}}| < 0.8$. The balances are calculated using events in 2015 data (black), Monte Carlo POWHEG + PYTHIA (red), and Monte Carlo SHERPA (blue). The lower pad ratio gives the size of the correction to apply to data, where the POWHEG + PYTHIA Monte Carlo is the nominal sample used for the correction and the differences between this and the SHERPA sample provides a handle on the generator uncertainty. From [91].

where $\hat{n}$ is a unit vector along the direction of the reference object, $E_{\text{T}}^{\text{ref}}$ is the missing transverse momentum of the reference object, and $\mathbf{E}_{\text{T}}^{\text{miss}}$ is the missing transverse momentum vector calculated directly from the negative sum of the transverse momenta of the topoclusters in the event. The direct balance and MPF methods are used in tandem as a cross-validation of the Z/γ + jet correction. However, the MPF method has the distinct advantage of being nearly independent of the choice of jet algorithm in the reconstruction since jets are not used directly in its calculation. The MPF is also less sensitive to contributions from pileup and the underlying event [91]. Figure 5.17 shows the balances calculated from MPF, using a Z as a reference, and the direct balance, with a γ reference object. Again this balance is derived for data, POWHEG + PYTHIA 8, and SHERPA events, where the POWHEG + PYTHIA8 is the nominal choice of Monte Carlo. The size of the correction is largest at lower jet transverse momenta.

The final stage in the *in situ* calibration is the *multi-jet balance* (MJB) where a high p_{T} jet is balanced against a system of many lower p_{T} jets which are already fully-calibrated up to and including the Z/γ + jet balances. Using the transverse momentum of the leading

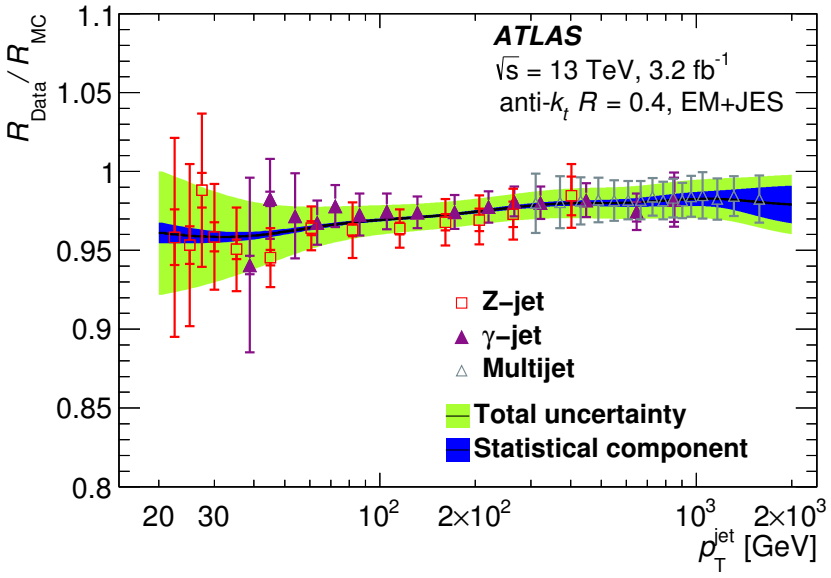


Figure 5.18: The final combined *in situ* correction for `AntiKt4EMTopo` jets measured in 2015 ATLAS data, as a function of the jet transverse momentum. The relative size of the correction is within a few percent of unity. The major uncertainties are discussed briefly in Section 5.3.3. From [91].

jet and the transverse momentum of the vector sum of all the remaining sub-leading jets in the event ($p_{\mathrm{T}}^{\mathrm{recoil}}$), the MJB response is:

$$\mathcal{R}_{\mathrm{MJB}} = \left\langle \frac{p_{\mathrm{T}}^{\mathrm{leading}}}{p_{\mathrm{T}}^{\mathrm{recoil}}} \right\rangle. \tag{5.17}$$

Following the derivation of the calibration factors for each of these *in situ* balances, the final closure in the jet response between data and Monte Carlo is obtained through a statistical combination and interpolation of second-order polynomial splines [91], shown in Figure 5.18. Here the total uncertainty is given by the green band and is consistent with the data/MC response closure from each of the separate balances. The central value of the combination closes to within 2-3 % across a wide range of jet p_{T} up to over 2 TeV.

5.3.2 Large-Radius Jets

We have now considered the full jet energy scale calibration chain for standard `AntiKt4EMTopo` jets. The $R = 1.0$ jets formed from `LCTopo` inputs are calibrated in a similar manner, but without the pileup and residual corrections or any GSC applied. In place of these, the pileup contributions are removed by the application of a jet grooming procedure. Throughout Run 2 of ATLAS the standard $R = 1.0$ jet grooming procedure has been jet trimming [89].

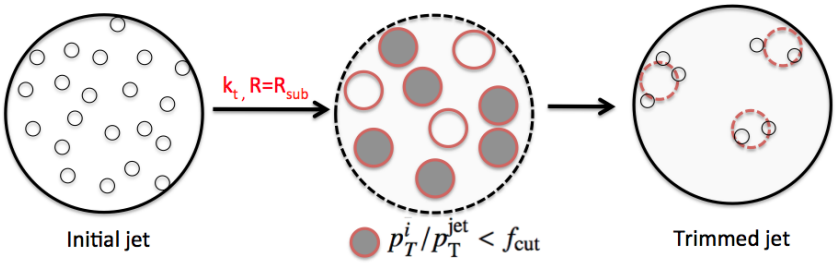


Figure 5.19: A cartoon illustrating the trimming a circular jet. From [90].

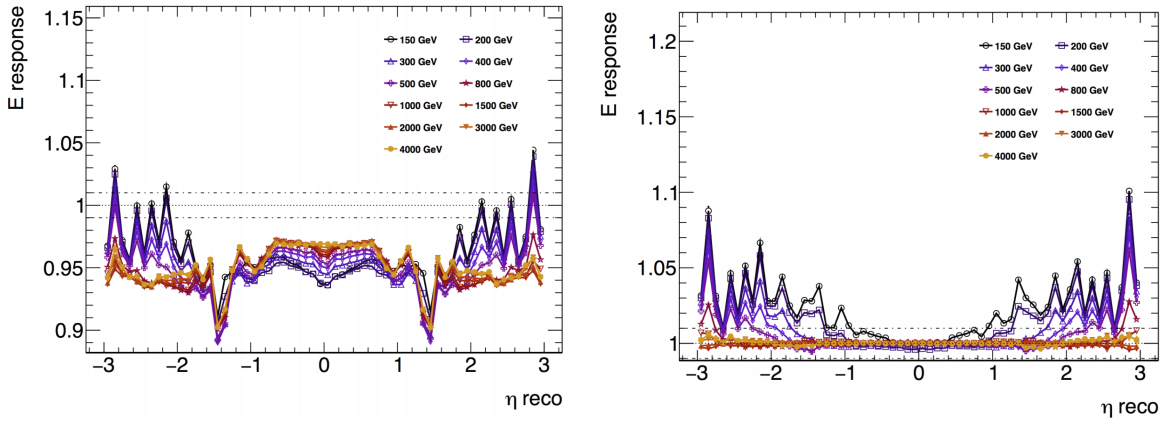
5.3.2.1 $R = 1.0$ Trimmed Jets

In the jet trimming procedure, shown schematically in Figure 5.19 the initial large-radius jet's constituents are first of all clustered into subjets of size R_{sub} using the k_t algorithm to recluster the softest constituents first. After the k_t-clustering, the algorithm proceeds by looping through each of the subjets and calculating the ratio of the transverse momentum of the ith subjet to the total transverse momentum of the jet, $p_{\text{T}}^i/p_{\text{T}}^{\text{jet}}$. If $p_{\text{T}}^i/p_{\text{T}}^{\text{jet}} < f_{\text{cut}}$ (where f_{cut} is a tunable parameter) the subjet is deemed to be sufficiently soft and is removed from the jet. During Run 1 of ATLAS, extensive optimisation studies were carried out to determine a sensible choice of R_{sub} and f_{cut} [90]: $R_{\text{sub}} = 0.2$ and $f_{\text{cut}} = 0.05$ were adopted and the resulting large-radius jets have until recently been the standard $R = 1.0$ jet definition in ATLAS: `AntiKt10LCTopoTrimmedPtFrac5SmallR20` jets.

5.3.2.2 Reclustered Jets

There is a second approach to producing large-radius jets, referred to as *jet reclustering* [96], or 'jets from jets', which will prove vital to the 'multi-jets' search for new physics presented in the third part of this thesis.

The basic idea of jet reclustering is to create a large-radius jet from smaller jets, and thus in ATLAS the standard reclustered jet definition is an $R = 1.0$ jet made from `AntiKt4EMTopo` jets acting as inputs to the sequential recombination algorithm used in the reclustering. The major advantage of reclustered jets over the standard trimmed `LCTopo` jets is that they are built from well-defined and calibrated `AntiKt4EMTopo` inputs. The reclustering process therefore 'inherits' the calibrations and systematic uncertainties associated with the `AntiKt4EMTopo` jets, and consequently the reclustered jets do not require an additional set of calibrations



(a) Before η-JES correction **(b)** After η-JES correction

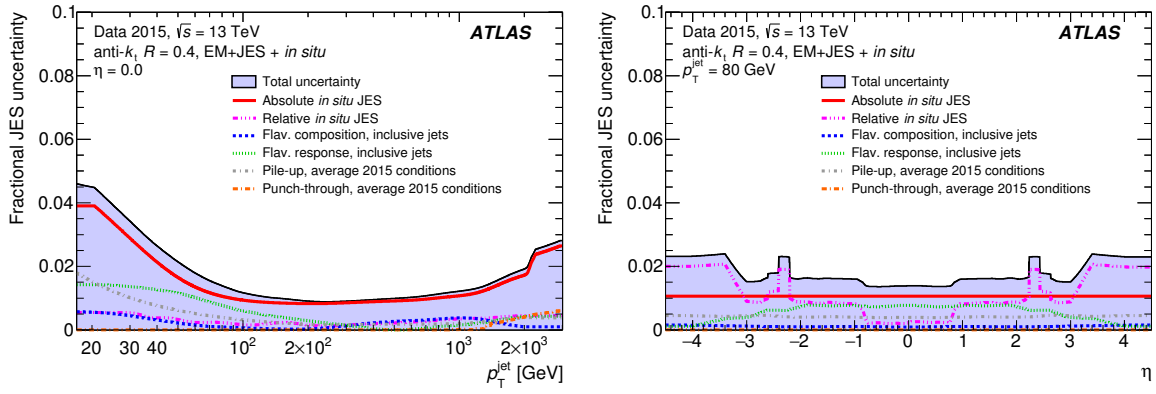
Figure 5.20: The `AntiKt10LCTopoTrimmedPtFrac5SmallR20` POWHEG + PYTHIA dijet Monte Carlo jet energy response as a function of η^{reco} for different bins of E^{true} given by different colours and patterns of points. The closure for jets with energies above 300 GeV is well below even the 1 % level.

or systematics. Another advantage is that jet reclustering can be easily re-run on the same inputs, and therefore turns the jet radius itself into a tunable parameter which the user can modify to make a new jet collection depending on the particular analysis or final state of interest. It is also possible to apply jet grooming to reclustered jets in the same way as the `AntiKt10LCTopoTrimmedPtFrac5SmallR20` jets. These features of the reclustering shall have interesting implications for the 'multi-jets' analysis.

5.3.2.3 $R = 1.0$ Jet Calibration

The `AntiKt10LCTopoTrimmedPtFrac5SmallR20` jets are first origin-corrected, and then η-JES corrected by geometrically matching reconstruction-level and truth-level jets within $\Delta R < 0.75$, and determining the fitted jet energy response in Monte Carlo using dijet events (as outined in Section 5.3.1.3). For completeness, `AntiKt10LCTopoTrimmedPtFrac5SmallR20` jets before and after the η-JES correction is applied are shown in Figure 5.20, as a function of the jet η^{reco}. The level of closure is similar to what is observed for `AntiKt4EMTopo` jets, and is within 5 % for the majority of phase-space. Different bins of the truth-level energy are considered, with the best closure observed at energies above 300 GeV.

After η-JES, no GSC is applied, but another set of *in situ* corrections (Like those described in 5.3.1.5) are applied.



(a) JES uncertainty vs. jet p_{T} **(b)** JES uncertainty vs. jet η

Figure 5.21: The fractional JES uncertainties for `AntiKt4EMTopo` jets as a function of the jet transverse momentum and η. The total uncertainty comes from all of the separate contributions of the Monte Carlo JES and *in situ* calibration uncertainties. From [91].

5.3.3 JES Uncertainties

Briefly, we shall now consider the dominating systematic uncertainties affected the JES calibration of jets. A feeling for the relative magnitude of the different components of the JES systematic can be gleaned from Figure 5.21, which shows the JES uncertainties for `AntiKt4EMTopo` jets, as a function of the jet p_{T} and η, measured in ATLAS 2015 data.

In general the largest systematic uncertainties come from the *in situ* measurements of the JES. For the Monte Carlo-driven residual, η-JES, and GSC corrections the largest uncertainties are due to flavour-dependence of the response and jet composition, and the pileup. The *in situ* uncertainties which dominate originate from:

1. Mismodelling in Monte Carlo (estimated by comparing different Monte Carlo generators and partonic showering frameworks).
2. Topology dependence in the different balances (estimate by varying the cuts applied to select the reference object).
3. Uncertainties inherited from existing calibrations, such as the Monte Carlo-only jet calibration.

The majority of these *in situ* uncertainties come from the Z + jets and multi-jet balance stages of the calibration.

For the MC JES and GSC, an estimate of the flavour response uncertainty is obtained from the difference between the jet energy responses of gluon jets (labelled as gluon jets using the truth record parton identification) in Pythia 8 and Herwig 7 [97] generators:

$$\sigma_{\text{flavour}}^{\text{response}} = f_{\text{gluon}} \times (\mathcal{R}_{E,\text{gluon}}^{\text{Pythia 8}} - \mathcal{R}_{E,\text{gluon}}^{\text{Herwig 7}}), \tag{5.18}$$

where the fraction of gluons in the sample is controlled by $f_{\text{gluon}} = 0.5$, assuming a 50:50 split between gluon and quark jets.

The flavour composition uncertainty (assuming a 100 % uncertainty on the gluon fraction f_{gluon}) is defined:

$$\sigma_{\text{flavour}}^{\text{composition}} = \sigma_{\text{gluon}}^{\text{fraction}} \times \frac{|\mathcal{R}_{E,\text{light-quark}}^{\text{Pythia 8}} - \mathcal{R}_{E,\text{gluon}}^{\text{Pythia 8}}|}{f_{\text{gluon}} \times \mathcal{R}_{E,\text{gluon}}^{\text{Pythia 8}} + (1 - f_{\text{gluon}}) \times \mathcal{R}_{E,\text{light-quark}}^{\text{Pythia 8}}}, \tag{5.19}$$

where 'light-quark' refers to jets associated with highest energy truth-labelled parent u-, d-, and c-quarks. This uncertainty could clearly be reduced by better constraining the gluon fraction estmate in Monte Carlo.

The pileup uncertainties are obtained from the uncertainties on the fits (α and γ) to the residual response dependencies on μ and N_{PV} as a function of the jet $p_{\text{T}}^{\text{true}}$. There is also an additional uncertainty coming from the topology-dependence of ρ, calculated using the difference in the magnitude of ρ between Z + jets, γ + jets, and dijet Monte Carlo, different Monte Carlo generators, and also with data.

6
The Jet Mass and its Resolution

Having presented the basic concept of hadronic jets, their reconstruction in simulated and measured ATLAS data, and the calibration of their energy scale, we shall now focus in some detail on one of the most rudimentary but powerful properties of a jet: the *jet mass*. Such a quantity encodes information about the masses of the hadrons which form the jet, and is therefore the central variable in numerous particle physics analyses which rely on information from reconstructed jets. Indeed we shall see in Part III of this thesis how the total jet mass of reclustered jets is utilised in the multi-jets search as a proxy for the presence of heavy particles which could be contained in those jets, making the jet mass one of the most important quantities in that analysis. More specifically, we can think of the jet mass as the most simple variable that describes the *substructure* [68] of a jet.

In this chapter we shall formally quantify the jet mass and describe how it can be calculated from the constituents of the reconstructed jet. During Run 2 of ATLAS the jet mass definition underwent substantial redevelopment [98–100], and the need for such progress shall also be motivated and described in detail, before the process of jet mass calibration is briefly outlined. Following a discussion of the jet mass magnitude and scale, we shall then investigate the degradation of the jet mass arising from the interaction of hadronic jets with the ATLAS detector. To quantify this degradation, the *jet mass resolution* [101] is introduced and two approaches of calculating this resolution, and its associated systematic uncertainty in physics analyses, are presented and compared [99, 101]. Having invested much time in the most primitive

of *jet substructure* variables, more complex substructure variables will be outlined in Chapter 7, and will prove invaluable to the studies therein described.

6.1 The Jet Mass Scale

6.1.1 Overview

Like the jet energy scale, the *jet mass scale* is a dimensionless quantity and it evaluates the relative size of the jet mass calculated at reconstruction-level to the jet mass at Monte Carlo truth-level, or the relative size of the jet mass as measured in data to the Monte Carlo reconstruction-level mass. Both a Monte Carlo and *in situ* jet mass calibration [98, 99] is derived in order to correct for differences between each of these different jet masses, using techniques in the spirit of jet energy scale calibration which has already been detailed. Before very briefly outlining this jet mass scale calibration, the jet mass definitions of interest must be formally described.

6.1.2 The Calorimeter Mass

Throughout Run 1 of ATLAS, and during early Run 2, the standard jet mass definition was the *calorimeter mass* [98, 102], m^{calo}. This jet mass is calculated solely from information provided by the calorimeter cells. Recall that the two standard ATLAS jet definitions, `AntiKt4EMTopo` and `AntiKt10LCTopoTrimmedPtFrac5SmallR20` jets, are reconstructed from electromagnetic and hadronic scale calorimeter topoclusters respectively. For brevity, we'll refer to such jets as calorimeter jets. For a calorimeter jet, J, reconstructed from calorimeter-cell (i.e. topocluster) constituents, labelled i, which each have an energy E_i and three-momentum $\mathbf{p_i}$, the calorimeter mass is simply the invariant mass obtained from the total four-momentum of all i constituents of the jet:

$$m^{\text{calo}} = \sqrt{\left(\sum_{i\in J} E_i\right)^2 - \left(\sum_{i\in J} \mathbf{p_i}\right)^{\mathbf{2}}}. \tag{6.1}$$

A difficulty with the calorimeter jet mass is that it is more susceptible to soft, low-energy, wide-angle contributions such as pileup. This is because this mass has an explicit dependence on the angular separation between its constituents. To see this, consider a calorimeter mass

calculated from summing over N constituents making up the jet. The mass can be written:

$$\begin{aligned}(m^{\text{calo}})^2 &= \Big(\sum_i^N E_i\Big)^2 - \Big(\sum_i^N \mathbf{p_i}\Big)^{\mathbf{2}} \\ &= \sum_i^N m_i^2 + \sum_{i,j}^N 2E_iE_j(1-\cos\theta_{ij}) \\ &\approx \sum_{i,j}^N E_iE_j(\theta_{ij})^2 \\ &\approx \sum_{i,j}^N p_{\text{T,i}}p_{\text{T,j}}(\Delta R_{ij})^2,\end{aligned} \tag{6.2}$$

where in the final two lines the approximations of $m^{\text{calo}} \gg m_i$ and small θ_{ij} are applied, where θ_{ij} is the angle between the i and j constituents. In the final line the result is re-expressed in typical detector coordinates, where the $\Delta R = \sqrt{(\eta_i-\eta_j)^2+(\phi_i-\phi_j)^2}$ in ATLAS $\eta-\phi$ coordinates. Thus the calorimeter jet mass explicitly depends on the angular separation of its own constituents.

6.1.3 The Track-Assisted Mass

Another limitation of the calorimeter mass is its dependence on the finite granularity (thought of as a 'resolving power') of the ATLAS calorimeters. Recalling that the opening angle between the decay products of a heavy hadron captured within a sufficiently large-radius jet of some p_{T} scales like $m_{\text{hadron}}/p_{\text{T}}$, it is clear that as the transverse momentum of the jet increases the angular separation between the decay products contained in the jet decreases, causing the jet to become more collimated (i.e. it has a larger *Lorentz boost*). For a sufficiently large Lorentz boost, the angular separation between the captured decay products will be comparable in size to the calorimeter cells used in the jet reconstruction. Consequently the ability to resolve the energy and momentum of these decay products, and hence reconstructed mass of the jet, will degrade substantially beyond some high p_{T}. One simple approach to reduce the impact of such degradation is to use more information from the ATLAS detector, such as the four-vector information measured from the tracks reconstructed in its Inner Detector.

Tracking and calorimetry are combined in the *track-assisted mass* [98], m^{TA}, defined as:

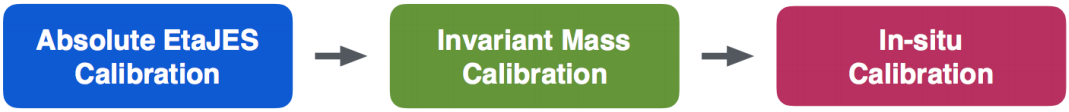


Figure 6.1: A summary of the complete jet mass scale calibration chain for `AntiKt4EMTopo` and `AntiKt10LCTopoTrimmedPtFrac5SmallR20` jets in ATLAS. Currently only the larger-radius jets are calibrated at the *in situ* level. Both the calorimeter and track-assisted jet mass definitions need to be separately calibrated.

$$m^{\mathrm{TA}} = \frac{p_{\mathrm{T}}^{\mathrm{calo}}}{p_{\mathrm{T}}^{\mathrm{track}}} \times m^{\mathrm{track}}. \tag{6.3}$$

Here $p_{\mathrm{T}}^{\mathrm{calo}}$ is the transverse momentum of the calorimeter jet, while $p_{\mathrm{T}}^{\mathrm{track}}$ is the transverse momentum of the four-vector sum of tracks which are associated to the same calorimeter jet. The invariant mass of this track four-vector sum is m^{track}, and this is set to the pion mass (due to the large multiplicity of pions contributing to the reconstructed tracks). The $p_{\mathrm{T}}^{\mathrm{calo}}/p_{\mathrm{T}}^{\mathrm{track}}$ ratio provides a measure of the charged-to-neutral component of the jet, and therefore helps to correct for charged-to-neutral fluctuations in the jet mass, improving the resolving power of the jet mass compared to a track-only definition of m^{track}. As a slightly technical aside, the tracks are associated to the calorimeter jet using the method of *ghost-association* [92]. Ghost-association is a technique which matches tracks to calorimeter jets by treating the tracks as infinitesimally soft, low-p_{T} particles (the track p_{T} is set to some very low sub-MeV value). Such tracks are then added as inputs to the anti-k_t algorithm along with the calorimeter clusters. Following reconstruction it is then possible to identify which tracks were clustered into which jets and thus isolate the list of tracks associated to each particular jet of interest. This can be used to construct track jets, or to match the tracking information to the clustering information in the jet reconstruction.

6.1.4 Jet Mass Calibration

The jet mass is highly sensitive to contributions from additional interactions arising from pileup and the underlying event, as well as soft wide-angle radiation captured in the cone of the jet. Approaches to pileup mitigation, such as the jet area method already discussed, or jet grooming, help to reduce these sensitivities. However, it is necessary to derive a full set of calibrations for the jet mass scale using a similar approach to the jet energy scale calibration chain. The corresponding jet mass calibration chain is shown in Figure 6.1. The jet mass calibration takes jets which have been calibrated at the jet energy scale, both in Monte Carlo and *in situ*, as inputs.

6.1.4.1 Monte Carlo-based Calibration

Jet collection	Requirements			
	on $p_{\mathrm{T}}^{\mathrm{calo}}$ [GeV]	on $\lvert\eta^{\mathrm{det}}\rvert$	Truth-matching	Jet isolation
$R = 1.0$, trimmed	> 200.0	< 2.0	ΔR(reco-truth) < 0.5	ΔR(reco-reco) > 2.5, ΔR(truth-truth) > 1.5
$R = 0.4$	> 100.0	< 2.8	ΔR(reco-truth) < 0.3	ΔR(reco-reco) > 0.6, ΔR(truth-truth) > 1.0

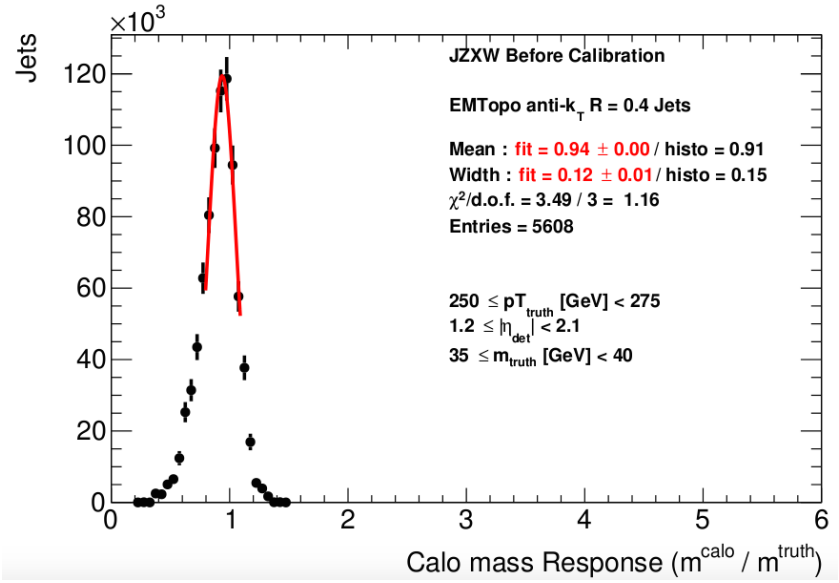
Table 6.1: The basic jet selections, truth-matching, and isolation requirements for the jets considered in the Monte Carlo jet mass calibration.

The basic selection criteria applied to the small and large-radius jets considered in the Monte Carlo jet mass calibration are summarised in Table 6.1. These include the minimum separation in $\eta - \phi$ between two truth and two reconstruction-level jets, as well as the maximum separation in $\eta - \phi$ between a reconstruction-level and truth-level jet considered for those two jets to be geometrically matched. From geometrically matched 'reco' and 'truth' jets, the *jet mass response* is defined in a similar fashion to a jet energy response:

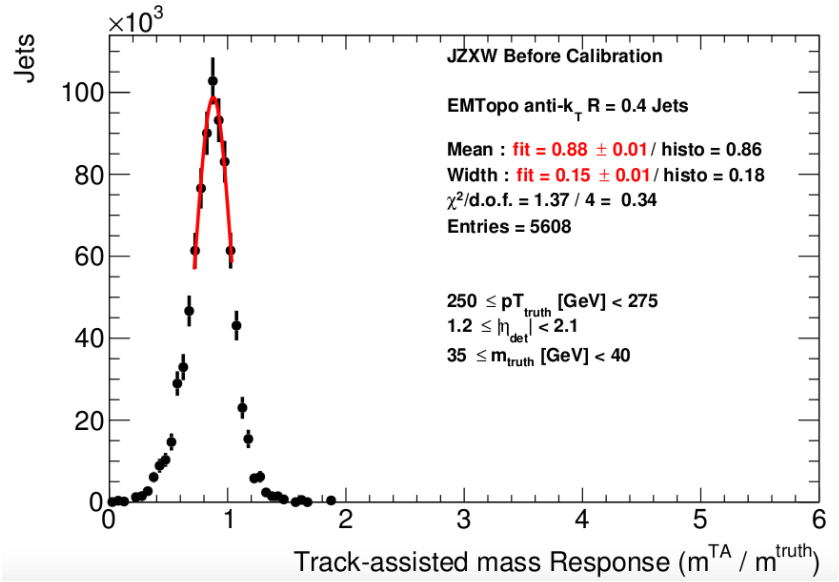
$$\mathcal{R}_m = \frac{m^{\mathrm{reco}}}{m^{\mathrm{true}}}, \tag{6.4}$$

where m^{reco} is the reconstruction-level jet mass, which will be either m^{calo} or m^{TA} obtained from Monte Carlo detector-level events depending on the mass definition considered in the calibration. The jet mass response is calculated in bins of $p_{\mathrm{T}}^{\mathrm{true}}$, $|\eta^{\mathrm{det}}|$, and m^{true}. The number and choice of variables was determined in order to provide distributions of the response which were sufficiently Gaussian, so that the average response could be extracted using the mean of a Gaussian fit. Several examples of such responses and fits, in different bins, for both mass definitions, are shown in Figures 6.2 and 6.3 for `AntiKt4EMTopo` and `AntiKt10LCTopoTrimmedPtFrac5SmallR20` jets respectively. The JMS calibration is derived from the same Pythia 8 QCD dijet samples used in the JES calibration. The small- and large-radius reconstruction-level jets are calibrated relative to the truth-level `AntiKt4Truth` and `AntiKt10TruthTrimmedPtFrac5SmallR20` jet collections respectively.

The calibration is derived bin-by-bin using the same numerical inversion technique procedure in the JES calibration. The calibration function, $\mathcal{F}_{m,\eta}$, is again obtained for different jet $|\eta^{\mathrm{det}}|$ and is constructed from a two-dimensional *Gaussian kernel* [103] of the mean fitted response, $\langle\mathcal{R}_m\rangle$. This has the functional form:

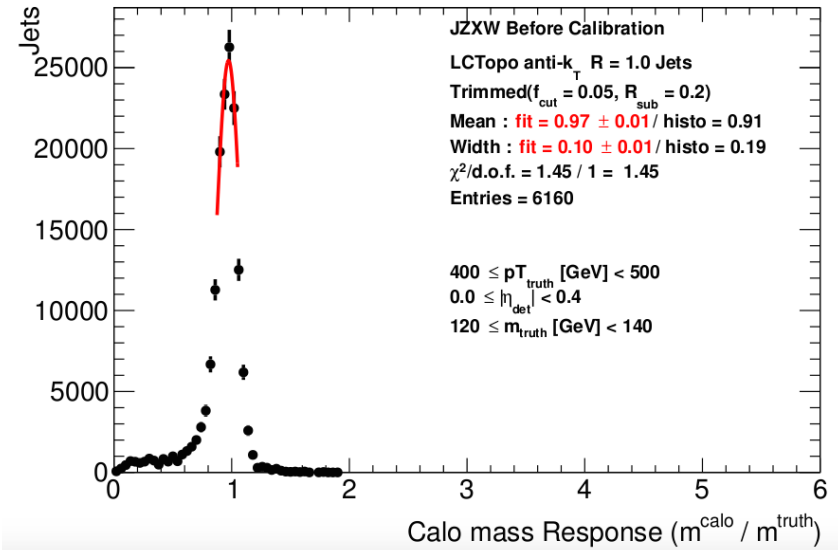


(a) Calorimeter jet mass response, $m^{\text{calo}}/m^{\text{true}}$, before JMS calibration.

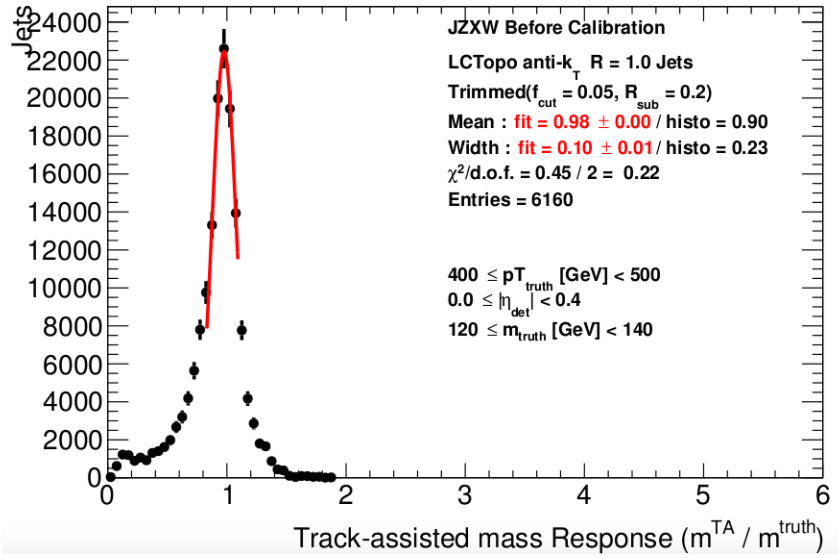


(b) Track-assisted jet mass response, $m^{\text{TA}}/m^{\text{true}}$, before JMS calibration.

Figure 6.2: Example calorimeter and track-assisted jet mass responses, fitted with a Gaussian (red line), for `AntiKt4EMTopo` jets. These jets have been calibrated up to the jet energy scale and are shown for a particular m^{true}, $p_{\text{T}}^{\text{true}}$, and $|\eta^{\text{det}}|$ bin, with the ranges given in each plot. The mean and width are provided, and they depend on the $p_{\text{T}}^{\text{true}}$.



(a) $m^{\text{calo}}/m^{\text{true}}$, before JMS calibration, $R = 1.0$ jets



(b) $m^{\text{TA}}/m^{\text{true}}$, before JMS calibration, $R = 1.0$ jets

Figure 6.3: Example calorimeter and track-assisted jet mass responses, fitted with a Gaussian (red line), for `AntiKt10LCTopoTrimmedPtFrac5SmallR20` jets. These jets have been calibrated up to the jet energy scale and are shown for a particular m^{true}, $p_{\text{T}}^{\text{true}}$, $|\eta^{\text{det}}|$ bin, with the ranges given in each plot. The mean and width are provided, and they depend on the $p_{\text{T}}^{\text{true}}$.

$$\mathcal{F}_{m,\eta}(p_{\mathrm{T}}^{\mathrm{calo}}, m^{\mathrm{reco}}) = \frac{\sum_{i=1}^{N_{\mathrm{bins}}} \langle \mathcal{R}_m \rangle_i w_i}{\sum_{i=1}^{N_{\mathrm{bins}}} w_i}, \tag{6.5}$$

where the response is weighted by the kernel:

$$w_i = \frac{1}{(\Delta\langle \mathcal{R}_m \rangle_i)^2} \times \mathcal{G}\Big(\frac{p_{\mathrm{T}}^{\mathrm{calo}} - \langle p_{\mathrm{T}}^{\mathrm{calo}} \rangle_i}{\sigma_{p_{\mathrm{T}}^{\mathrm{calo}}}} \oplus \frac{m^{\mathrm{reco}} - \langle m^{\mathrm{reco}} \rangle_i}{\sigma_{m^{\mathrm{reco}}}}\Big). \tag{6.6}$$

Here i corresponds to the index of a ($p_{\mathrm{T}}^{\mathrm{calo}}$, m^{calo})-bin, $\Delta\langle \mathcal{R}_m \rangle_i$ is the statistical uncertainty on $\langle \mathcal{R}_m \rangle_i$, and the $\langle p_{\mathrm{T}}^{\mathrm{calo}} \rangle_i$ and $\langle m^{\mathrm{reco}} \rangle_i$ are the average transverse momentum and mass of the jets in the ith bin. These have widths which are extracted from the fit to give the $\sigma_{p_{\mathrm{T}}^{\mathrm{calo}}}$ and $\sigma_{m^{\mathrm{reco}}}$ values. The $\mathcal{G}$ denotes a Gaussian function with a kernel width provided in brackets, and $\oplus$ denotes addition in quadrature.

Following numerical inversion, the extracted calibration function is used to determine the calibrated jet mass, $m_{\mathrm{corr}}^{\mathrm{reco}}$.

$$m_{\mathrm{corr}}^{\mathrm{reco}} = \frac{m^{\mathrm{reco}}}{\mathcal{F}_{m,\eta}}(m^{\mathrm{reco}}, p_{\mathrm{T}}^{\mathrm{calo}}). \tag{6.7}$$

Example calorimeter and track-assisted jet masses, before and after Monte Carlo jet mass calibration, are shown in Figure 6.4, for a sample enriched in W/Z jets. The application of the QCD dijet-derived calibration function moves the position of the jet mass closer to the W and Z pole masses, as would be expected. In addition, the average jet mass response is shifted towards unity following the application of these calibration functions. The calibration 'closure' is shown in Figure 6.5 for the calorimeter and track-assisted jet masses.

6.1.4.2 *In situ* Calibration

Following the Monte Carlo calibration, the JMS is corrected for in data through an *in situ* calibration [104]. This *in situ* procedure uses the same Z+jets balance, γ + jet balance, and multi-jet balance techniques already described in Section 5.3.1.5 to calibrate the jet mass response in data vs. Monte Carlo. These techniques have already been described in detail. For completeness, the double ratio of the jet mass response following the combination of all *in situ* calibrations is shown in Figure 6.6, where the calibration has been derived using the 2015+2016 ATLAS data-set. The *in situ* component of the calibration chain is applied to

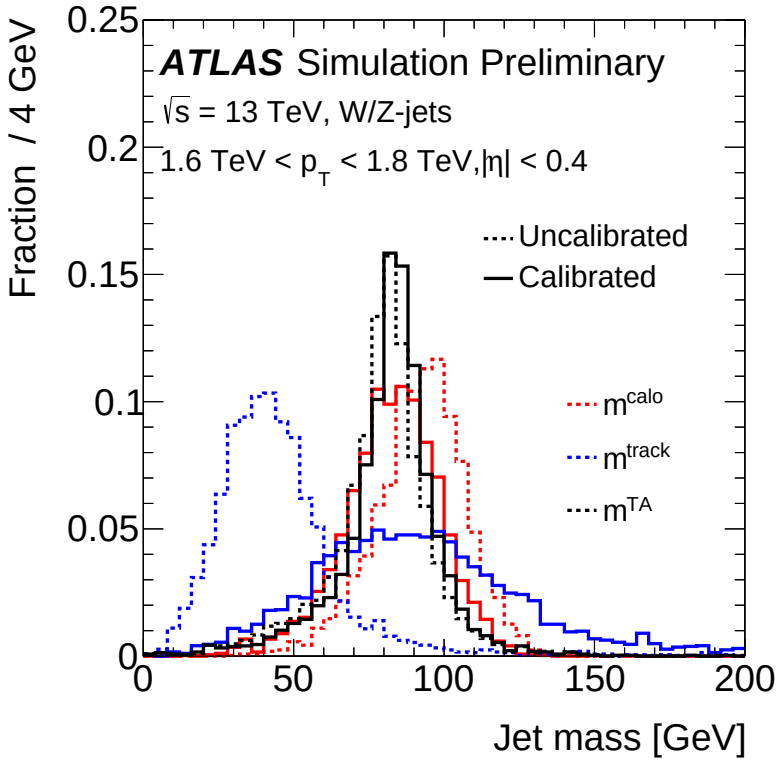
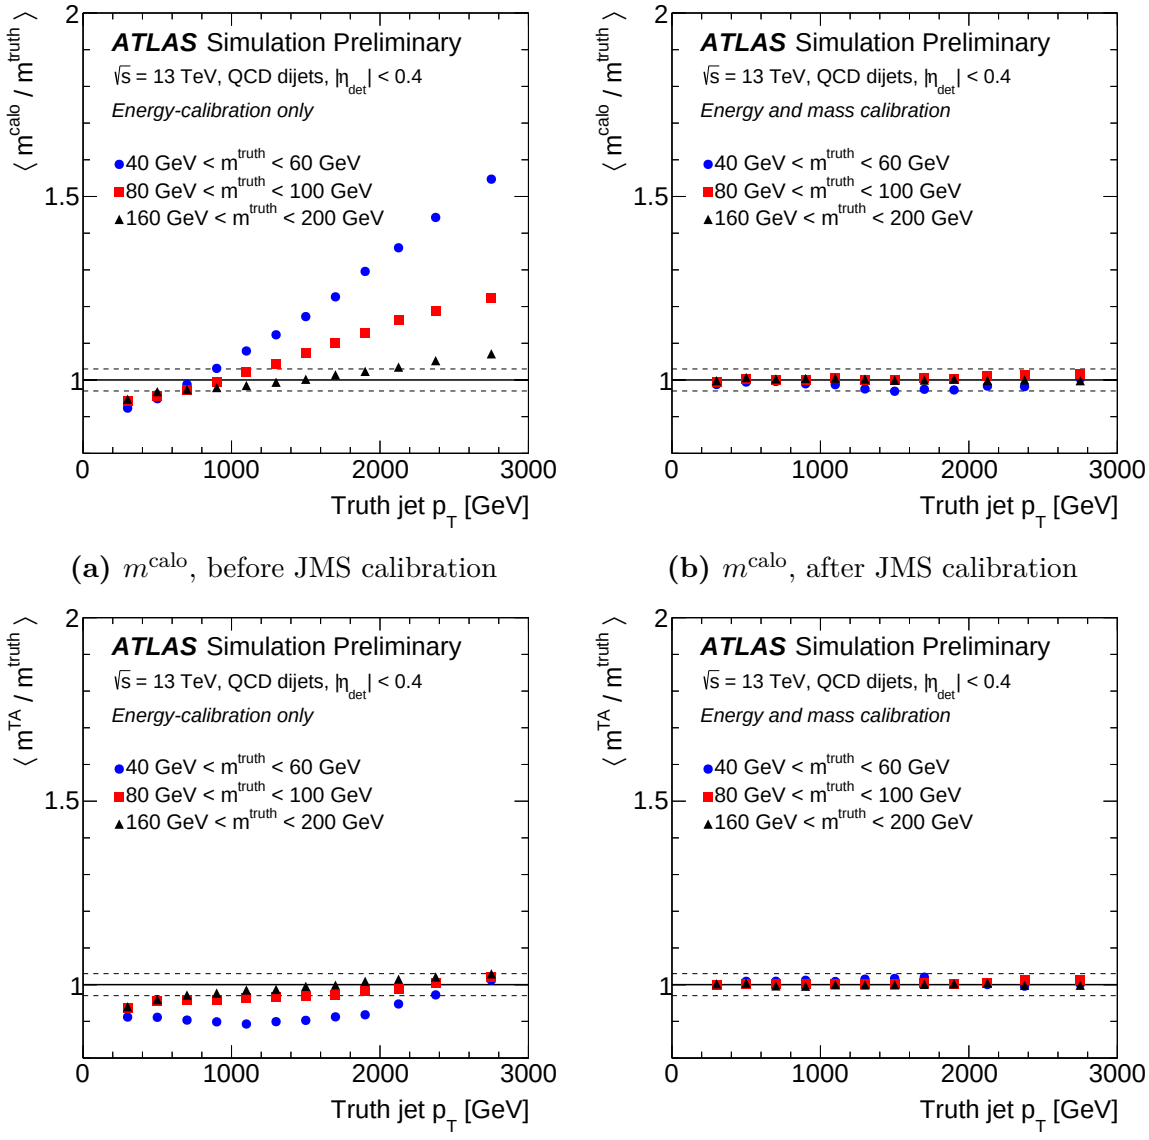


Figure 6.4: The uncalibrated (dashed line) and calibrated (solid line) reconstruction-level jet masses for the calorimeter, track-assisted, and track-only jet mass definitions. The calibrations are obtained from QCD dijets Monte Carlo, and applied to a simulated sample enriched in W/Z jets. `AntiKt10LCTopoTrimmedPtFrac5SmallR20` jets are considered. From [98].

`AntiKt10LCTopoTrimmedPtFrac5SmallR20` jets only.

6.1.5 The Combined Mass

The primary concern when interested in a precision measurement of the jet mass is how well this quantity can be resolved in the ATLAS detector, based on the detector limitations of reconstructing the initial jet. This is quantified by the *jet mass resolution* (JMR) [101]. Lower values of the JMR correspond to better resolving power in the measurement of the jet mass. As stated earlier, at higher transverse momenta the finite granularity of the calorimeter causes the calorimeter jet mass resolution to degrade; in this case JMR improvements are obtained by using the track-assisted mass. However, at lower transverse momenta the track-assisted mass resolution degrades due to the appreciable charged-to-neutral fluctuations in the $p_{\mathrm{T}}^{\mathrm{calo}}/p_{\mathrm{T}}^{\mathrm{track}}$-term. At these lower transverse momenta it is therefore the calorimeter mass resolution which is smaller. Ideally what a user needs is a jet mass definition with the smallest possible resolution at any given jet p_{T} in their particular measurement or search. The logical step in achieving this



(a) m^{calo}, before JMS calibration

(b) m^{calo}, after JMS calibration

(c) m^{TA}, before JMS calibration

(d) m^{TA}, after JMS calibration

Figure 6.5: The average calorimeter and track-assisted jet mass response from QCD dijets as a function of the truth-level transverse momentum, before and after JMS calibration. Three different truth-level jet mass bins are shown, and the closure following calibration is within 2 % in all cases. `AntiKt10LCTopoTrimmedPtFrac5SmallR20` jets are considered. From [98].

is to find a meaningful way of combining m^{TA} and m^{calo} such that JMR improvements from m^{calo} are emphasised at lower p_{T}, while the improvements from using m^{TA} are still expressed at higher p_{T}. The combination of these two mass definitions, such that the resolution on the mass is minimised across the jet p_{T}, is referred to as the *combined jet mass*, m^{comb}.

The combined jet mass is a linear sum of the calorimeter and track-assisted masses, weighted such that, across the range of the jet p_{T}, it has a lower resolution than the separate calorimeter

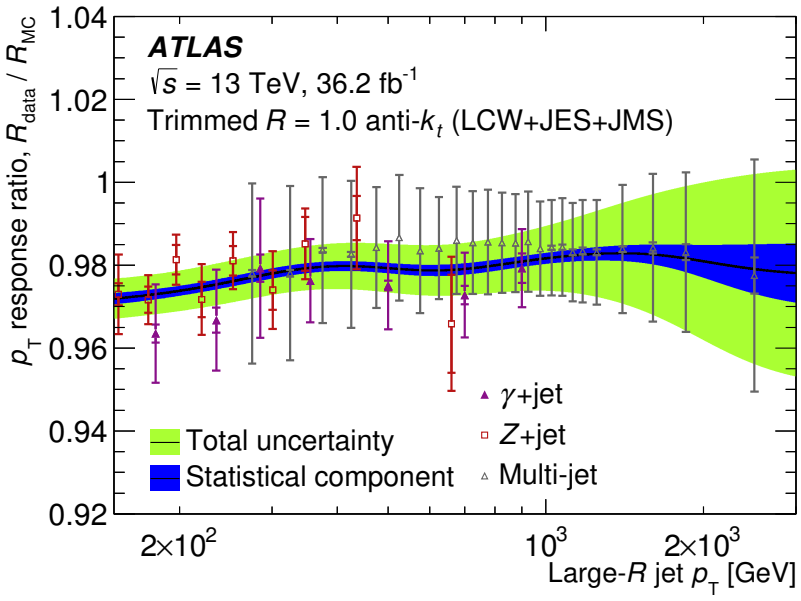


Figure 6.6: The final combined *in situ* correction for `AntiKt10LCTopoTrimmedPtFrac5SmallR20` jets measured in 2015+2016 ATLAS data, as a function of the jet transverse momentum. The relative size of the correction is within a few percent of unity. From [104].

mass and track-assisted mass. m^{comb} is a simple linear combination, and is defined as:

$$m^{\text{comb}} = w^{\text{calo}} \times m^{\text{calo}} + w^{\text{TA}} \times m^{\text{TA}}, \tag{6.8}$$

where the two weights are constrained to sum to unity: $w^{\text{calo}} + w^{\text{TA}} = 1$. The optimal weights are calculated by minimising the combined jet mass resolution, starting from its formula. The jet mass resolution is calculated from the same jet mass response employed in the JMS calibration. The combined mass response can be written in terms of (approximately) Gaussian jet mass responses by dividing throught the combined mass by the truth jet mass, giving :

$$\mathcal{R}_m^{\text{comb}} = w^{\text{calo}} \times \mathcal{R}_m^{\text{calo}} + w^{\text{TA}} \times \mathcal{R}_m^{\text{TA}}. \tag{6.9}$$

Assuming the correlations between the calorimeter and track-assisted masses are negligible, the mathematical variance of the above mass response can be written

$$\text{var}(\mathcal{R}_m^{\text{comb}}) = (w^{\text{calo}})^2 \times \text{var}(\mathcal{R}_m^{\text{calo}}) + (w^{\text{TA}})^2 \times \text{var}(\mathcal{R}_m^{\text{TA}}). \tag{6.10}$$

Each jet mass response distribution will have some width σ, coming from the square root of the variance ($\text{var}(x) = \sigma_x^2$). Using this, and re-writing w^{TA} in terms of w^{calo}, we obtain an expression for the width of the combined jet mass response:

$$(\sigma^{\text{comb}})^2 = (w^{\text{calo}})^2 \times (\sigma^{\text{calo}})^2 + (1 - w^{\text{calo}})^2 \times (\sigma^{\text{TA}})^2. \tag{6.11}$$

Recall that the final goal is to determine a mass definition that minimises the jet mass resolution. The jet mass resolution is obtained from the width of the jet mass responses, corresponding to these σ. Consequently the problem now reduces to finding the w^{calo} that minimises σ^{comb}. Performing this minimisation with respect to w^{calo}, we arrive at the combined mass weights:

$$w^{\text{calo}} = \frac{(\sigma^{\text{calo}})^{-2}}{(\sigma^{\text{calo}})^{-2} + (\sigma^{\text{TA}})^{-2}}, \quad w^{\text{TA}} = \frac{(\sigma^{\text{TA}})^{-2}}{(\sigma^{\text{calo}})^{-2} + (\sigma^{\text{TA}})^{-2}}. \tag{6.12}$$

In the presence of non-negligible correlations between the calorimeter and track-assisted mass responses, the combined mass variance expression acquires an additional covariance term:

$$\text{var}(\mathcal{R}_m^{\text{comb}}) = (w^{\text{calo}})^2 \times \text{var}(\mathcal{R}_m^{\text{calo}}) + (w^{\text{TA}})^2 \times \text{var}(\mathcal{R}_m^{\text{TA}}) + 2w^{\text{calo}}w^{\text{TA}} \times \text{cov}(\mathcal{R}_m^{\text{calo}}, \mathcal{R}_m^{\text{TA}}). \tag{6.13}$$

The new combined mass weights follow from an identical minimisation of the above equation, yielding, for the calorimeter mass-term:

$$w^{\text{calo}} = \frac{(\sigma^{\text{TA}})^2 - \rho\sigma^{\text{calo}}\sigma^{\text{TA}}}{(\sigma^{\text{calo}})^2 + (\sigma^{\text{TA}})^2 - 2\rho\sigma^{\text{calo}}\sigma^{\text{TA}}}, \tag{6.14}$$

where ρ is the correlation coefficient of $\mathcal{R}_m^{\text{calo}}$ vs. $\mathcal{R}_m^{\text{TA}}$ and is related to the covariance through the expression

$$\rho = \frac{\text{cov}(\mathcal{R}_m^{\text{calo}}, \mathcal{R}_m^{\text{TA}})}{\sigma^{\text{calo}}\sigma^{\text{TA}}}. \tag{6.15}$$

The calculation of the combined jet mass therefore depends directly on the resolutions of the calorimeter and track-assisted masses. The combined jet mass is designed to be used with the `AntiKt10LCTopoTrimmedPtFrac5SmallR20` jets and the weights are obtained from calorimeter and track-assisted mass resolutions studied in Monte Carlo QCD dijets events, in *pp* collisions at 13 TeV. The basic selections for the jets, including truth-matching and isolation, are summarised in Table 6.1 i.e. they follow identical selection criteria for the large-radius jets used in the original Monte Carlo jet mass calibration. The jets used to obtain the weights are calibrated at both the jet energy and jet mass scales. The jet mass resolution is calculated from the interquantile range of the jet mass response of Monte Carlo reconstruction-level jets to truth-level jets:

$$\sigma = \frac{IQnR(\mathcal{R}_m)}{2 \times \mathcal{M}(\mathcal{R}_m)}, \tag{6.16}$$

where $\mathcal{M}(\mathcal{R})$ denotes the *median* of the jet mass response distribution and $IQnR(\mathcal{R})$ is the 68.0 % coverage interval of the interquantile range[1] of the jet mass response, and is a measure of the *absolute* size of the resolution. The quantity that has been referred to as the JMR is in fact the $fractional$ JMR, where the absolute value is quantified by $IQnR(\mathcal{R})/2$, and the median acts as a central value of the jet mass response. In this chapter, when the JMR is calculated and compared it is always a fractional JMR, but the 'fractional' word in the text description will be dropped for simplicity.

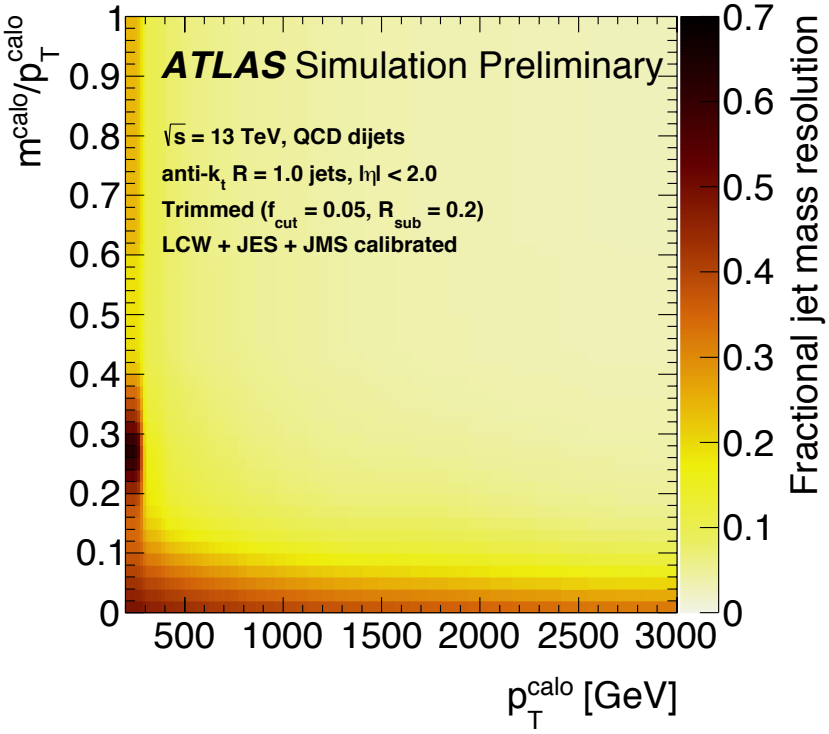
The jet mass resolutions for the calorimeter and track-assisted masses which are used to derive a combined mass are shown in Figure 6.7. The JMR is parametrised in terms of the reconstruction-level mass and the ratio of that mass to the calorimeter jet transverse momentum. Figure 6.7a shows the general behaviour for the calorimeter mass resolution: at lower $m^{\text{calo}}/p_{\text{T}}^{\text{calo}}$ the jets have a larger Lorentz boost and are therefore more collimated. Due to finite granularity of the calorimeter, the JMR consequently degrades.

The final performance of the combined mass will be assessed in Section 6.2. Figure 6.8 summarises the first application of the combined jet mass, comparing the performance of the combined jet mass resolution with respect to the mass resolutions obtained from the calorimeter and track-assisted masses respectively. The combined jet mass resolution is significantly smaller than the calorimeter mass resolution at higher truth jet transverse momenta, due to the additional tracking information provided by the track-assisted mass component. The reduced resolution in this boosted regime has been utilised by numerous searches for exotic particle production on ATLAS.

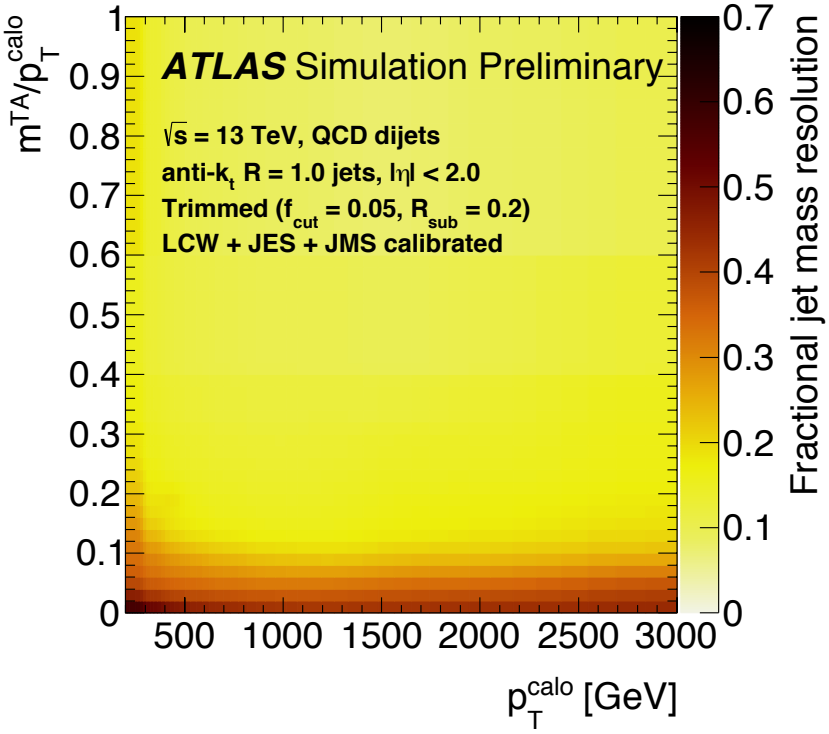
6.1.6 Particle Flow

We need to now briefly summarise the concept of *Particle Flow* and its use in jet reconstruction, as particle flow jets will be directly compared with topocluster-based jets both later in this chapter, and in subsequent chapters. Particle flow jets offer improvements to the jet performance

[1]Such range is defined as $q_{84\%} - q_{16\%}$, where $q_{84\%}$ and $q_{16\%}$ are the 16^{th} and 84^{th} percentiles of a given distribution.



(a) JMR for the calorimeter mass from QCD dijets



(b) JMR for the track-assisted mass from QCD dijets

Figure 6.7: The fractional jet mass resolution for the `AntiKt10LCTopoTrimmedPtFrac5SmallR20` jets used to determine the weights for the combined mass.

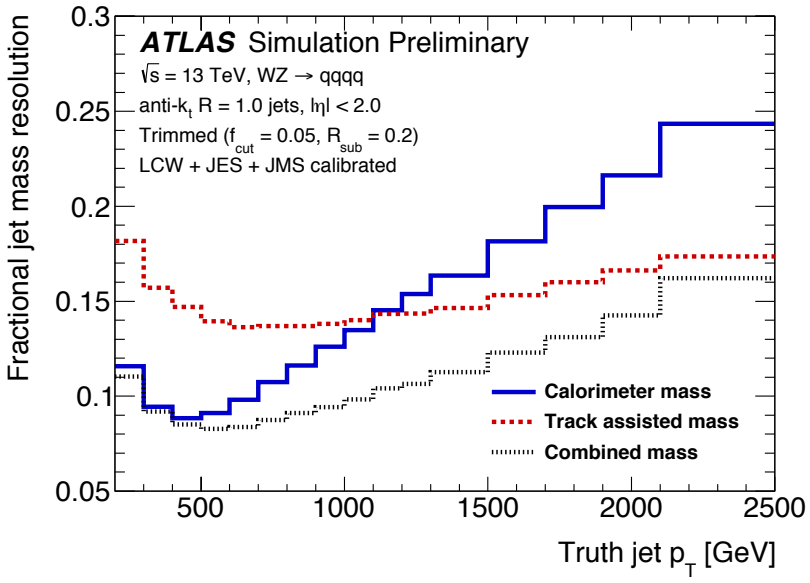


Figure 6.8: The three different jet mass definitions for `AntiKt10LCTopoTrimmedPtFrac5SmallR20` jets for a simulated sample of W/Z jets. The jet mass resolution is provided as a function of the jet transverse momentum, and the resolution improvement from the combined jet mass is clearly seen.

(e.g. a better jet energy resolution) by adding tracking information explicitly into the inputs of the jet algorithm.

Particle flow objects (PFOs) [105] are constructed by matching tracks, which have passed stringent selection criteria, to electromagnetic scale topoclusters in the calorimeter. The tracks used in the matching must be reconstructed from at least nine hits in the silicon detectors, where no missing hits in the pixel tracker are observed. To avoid double-counting energy, the charged component from the track is compared with its expected deposition in the calorimeter, which is then removed from the corresponding matched topocluster. Following this subtraction, the algorithm further handles the case of energy being split between multiple adjacent topoclusters. The final PFOs contain tracking information from charged tracks, as well as calorimeter information from the corrected topoclusters. One may think of particle flow as a generalisation of the track-assisted jet mass, where the matching of the track has been extended to each calorimeter cell, rather than just a matched reconstructed calorimeter jet. The anti-k_t algorithm is then applied to produce small-radius jets from these PFO inputs, where tracks satisfying $|z_0| \sin\theta < 3.0$ mm are considered. The resulting particle flow jet energy and mass are then fully calibrated in precisely the same manner as the topocluster-based jet definitions which have already been described in some detail. We will refer to the resulting $R = 0.4$ jets as

`AntiKt4EMPFlow` jets. A schematic illustration of the process of track-matching and cluster splitting in ATLAS PFO reconstruction is provided by Figure 6.9.

6.2 The Jet Mass Resolution

From an analysis perspective, the jet mass is crucial to differentiating a number of different final state topologies. Resonant peaks are absent in the invariant mass distribution derived from jets originating from gluon, b-, c-, or light-quark jets. We will refer to such jets collectively as 'QCD jets'. Such jets have a mass distribution which features only a low mass logarithmic enhancement[2] followed by a decline arising from phase-space limitations. Understanding the form of different jet mass distributions allows for the selection and reconstruction of event samples of final states containing W/Z, top, and $H \to b\bar{b}$ jets. Consequently, one may measure their differential kinematic properties and/or search for resonant and non-resonant deviations from the Standard Model (SM) predictions. Estimating the mass resolution for such analyses, and developing a reliable and testable methodology for the propagation of its systematic uncertainties, forms the basis of this Section.

The jet mass plays an important role in identifying (and tagging [106]) the hadronic decays of top quarks and heavy bosons (W, Z, Higgs) and in differentiating between them. Having discussed the jet mass and its calibration in some detail, together with formally defining how the jet mass resolution is obtained, we now focus on studying the impact of determining the jet mass resolution in various Monte Carlo final states, for both small- and large-radius jets. All jet JMRs will be presented for several topologies: W/Z, top, $H \to b\bar{b}$, and QCD jets.

In these studies, two different JMRs will be discussed and evaluated. A Monte Carlo-based JMR is evaluated from the spread of the difference between reconstruction-level and (geometrically matched - see Table 6.1) truth jet masses in simulated events. Such a Monte Carlo-based JMR has already been defined in Section 6.1.5. In turn, a second JMR is obtained by using the JMR relative scale between data and simulation derived using the *in situ* technique

[2]The leading order mass differential cross-section $\frac{d\sigma}{dm^2_{\text{jet}}} \approx \frac{\alpha_S}{m^2_{\text{jet}}} \log(\frac{R^2 p^2_{\text{jet,T}}}{m^2_{\text{jet}}})$ where R, $p_{\text{jet,T}}$ and m_{jet} are the jet radius, transverse momentum and mass respectively [68].

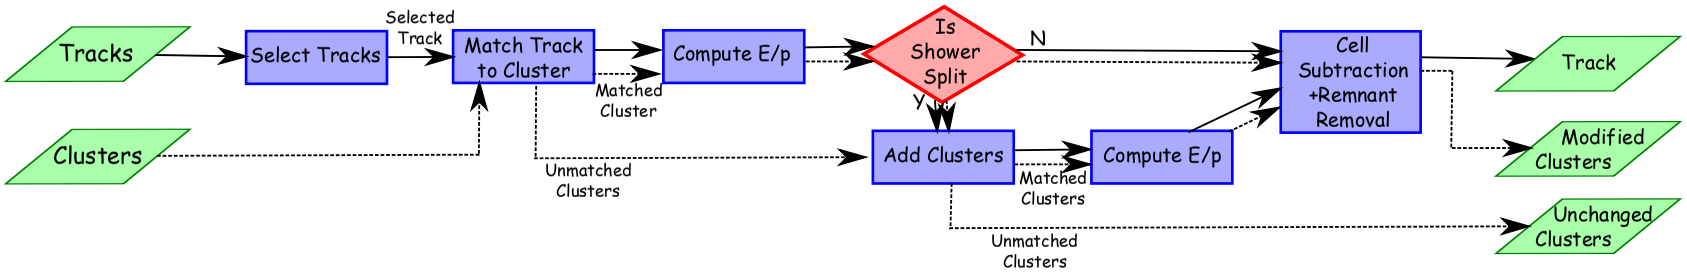


Figure 6.9: A flow chart summarising the particle flow object reconstruction algorithm. Tracks are first matched to clusters. The charged track component's energy of then subtracted to avoid double-counting the reconstructed energy. Finally the shower calorimeter shower splitting is assessed and additional calorimeter clusters are added in appropriately. From [105].

of *forward-folding* [98, 102] (discussed in Section 6.2.4). This second JMR has the advantage of being corrected using Run 2 data; it does not just rely on Monte Carlo information alone. Finally the propagation of JMR systematic uncertainties is illustrated first for the MC-based JMR [107], and then for this second data-corrected JMR (which is often referred to as a forward-folded JMR in the literature [99, 104]). This is presented in Section 6.2.5.

6.2.1 Monte Carlo Samples

Samples of simulated events for proton-proton (pp) collisions at a center-of-mass energy of 13 TeV are used to study the ATLAS detector response and JMR of various reconstructed jet definitions. The samples used for these studies are summarised in Table 6.2.

The exotic production of W' and Z' bosons [108, 109] provides a copious source of simulated boosted W and Z bosons, and of top quarks, respectively. The graviton [110] production provides a source of di-Higgs states, which decay to $b\bar{b}$. The W' and Z' signals are motivated by several BSM theories, such as Little Higgs Models [108], whilst the graviton is motivated by Randall-Sundrum models which generate gravitons and describe reality as a higher-dimensional universe with warped geometry between different 'branes' [110]. These three 'signal' samples are produced with W' and Z' boson masses ranging from 400 to 5000 GeV and graviton masses ranging from 400 to 3000 GeV, providing a spectrum of final state jets with a wide range of Lorentz boosts. The simulated events are re-weighted so that the distribution of the mean number of interactions per bunch crossing in simulation matched what has been observed in the data collected by the ATLAS detector in the years 2015 and 2016. These same signal-samples will be used in the studies in Chapter 7.

The samples are processed using the full ATLAS detector simulation [111], based on GEANT 4 [112].

6.2.2 Monte Carlo-based Jet Mass Resolution

The MC-based JMR is evaluated for the four simulated event samples described in Section 6.2.1: $W' \to WZ \to q\bar{q}q\bar{q}$, $Z' \to t\bar{t}$, $G \to HH \to b\bar{b}b\bar{b}$, and $2 \to 2$ strong interactions processes, a

Process	Matrix Element Generator	PDF	Parton Shower	Underlying Event Tune
$W'(\rightarrow W/Z)$	Pythia 8 [113]	NNPDF23LO [114]	Pythia 8	A14 [115]
$G(\rightarrow HH)$	MadGraph 5 [116]	NNPDF23LO	Pythia 8	A14
$Z'(\rightarrow t\bar{t})$	Pythia 8	NNPDF23LO	Pythia 8	A14
QCD jets	Pythia 8	NNPDF23LO	Pythia 8	A14

Table 6.2: Overview of the simulated samples used in these jet mass resolution studies.

common background in ATLAS. These samples give rise to large-radius W/Z jets, top jets, $H \rightarrow b\bar{b}$ jets, and QCD jets, respectively. In addition to the basic requirements outlined in Table 6.1, such jets are matched, using ghost-association [92], to a truth-level W/Z, top or H particle respectively, with the exception of QCD jets, where no matching requirement to a truth particle is imposed.

6.2.2.1 Monte Carlo JMR Estimate for Large-Radius Jets

Let's first describe JMR as a function of large-radius jet $p_{\mathrm{T}}^{\mathrm{true}}$ in some detail. The behaviour is shown in $m^{\mathrm{true}}/p_{\mathrm{T}}^{\mathrm{true}}$ bins for QCD jets, and m^{true} bins for W/Z, top and $H \rightarrow b\bar{b}$ jets ('signal jets'), for the calorimeter mass m^{calo} resolutions in Figure 6.10. For the jets arising from the signal samples, the JMR is binned in m^{true}. This is motivated by the presence of an intrinsic mass scale expectation for such jets, namely at the W, Z, and Higgs mass poles. However there is no intrinsic mass scale for QCD jets, and so the JMR is shown as a function of $m^{\mathrm{true}}/p_{\mathrm{T}}^{\mathrm{true}}$ in that case, allowing one to proxy the particle-level angular spread of the decay products produced by the parent particle generating the jet.

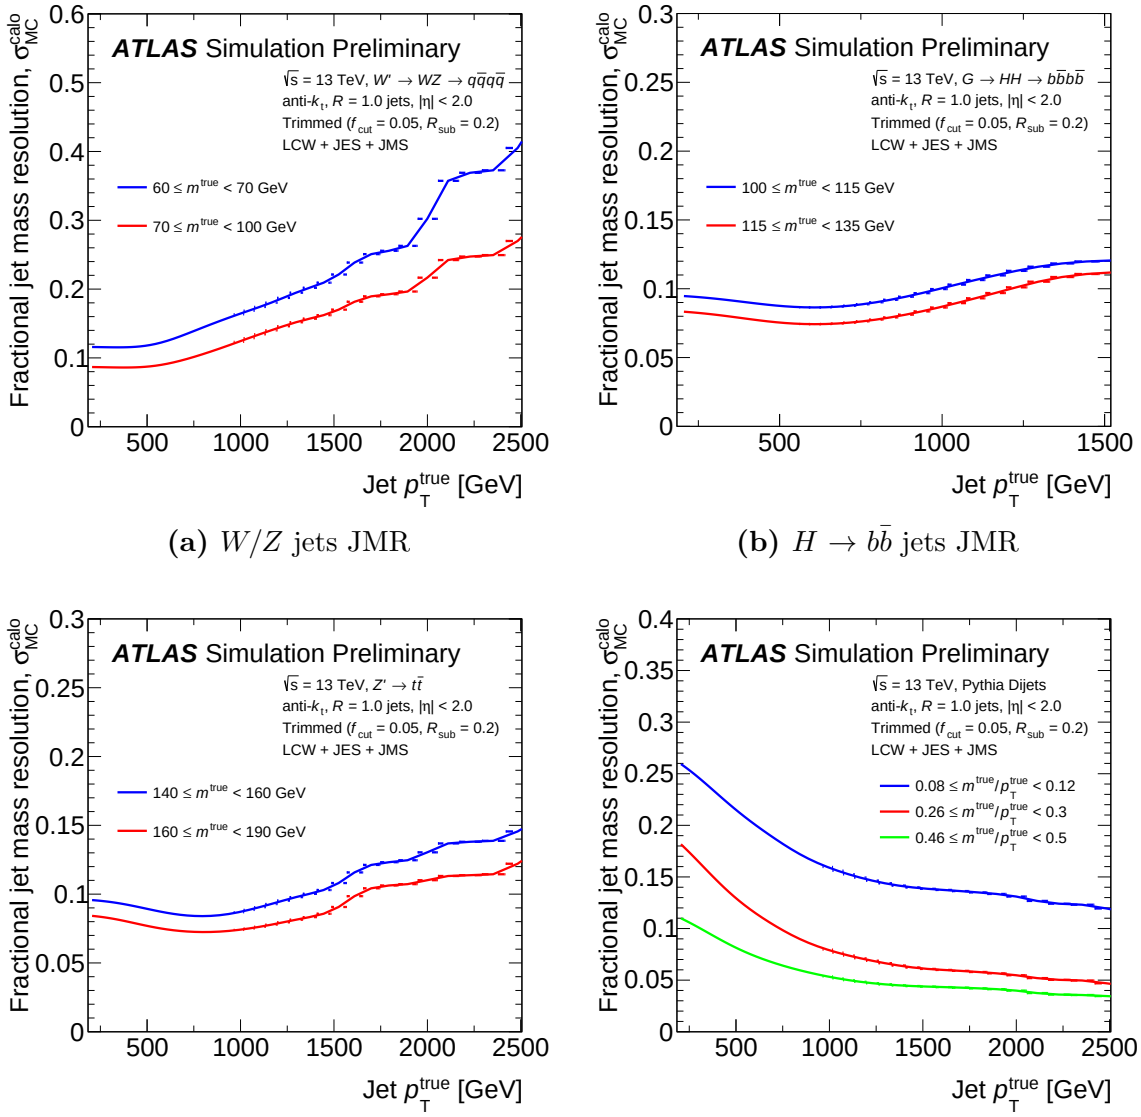
Binning in $m^{\mathrm{true}}/p_{\mathrm{T}}^{\mathrm{true}}$ provides a complete overview of the general JMR behaviour in the QCD jets distribution, where the JMR is largest at low $m^{\mathrm{true}}/p_{\mathrm{T}}^{\mathrm{true}}$ values. Fixing this ratio selects jets in which the energy distribution has approximately constant angular spread. For this fixed ratio, the energy content of the jet increases as $p_{\mathrm{T}}^{\mathrm{true}}$ increases. The calorimeter energy resolution is reduced with increasing jet energy and, as the angular resolution contribution is roughly constant, the overall JMR resolution is expected to be reduced as a function of $p_{\mathrm{T}}^{\mathrm{true}}$ at fixed $m^{\mathrm{true}}/p_{\mathrm{T}}^{\mathrm{true}}$. On the other hand, if instead $p_{\mathrm{T}}^{\mathrm{true}}$ is fixed, then as $m^{\mathrm{true}}/p_{\mathrm{T}}^{\mathrm{true}}$ decreases in size, the angular separation between the decay products becomes comparable to the finite granularity of the calorimeter, causing the jet mass resolution to degrade. The largest JMR values are

observed in the curve corresponding to the smallest $m^{\text{true}}/p_{\text{T}}^{\text{true}}$: here the subjet angular spread becomes comparable in size to a cell in the hadronic calorimeter and the limitations due to the spatial granularity of the detector are more significant.

For W/Z, top and $H \to b\bar{b}$ jets, each JMR curve corresponds to an ensemble of jets with an approximately fixed mass (distributed about the upper and lower bounds of each m^{true} bin). Fixing the mass gives rise to a decreasing angular spread of the decay products as $p_{\text{T}}^{\text{true}}$ increases. It is therefore expected that the reduction in energy resolution of the calorimeter at higher $p_{\text{T}}^{\text{true}}$ is offset by the increase in angular resolution of the jets. In addition, at a fixed $p_{\text{T}}^{\text{true}}$, jets with smaller masses are expected to have smaller angular spread and similar energy content to the larger mass jets, thus generally resulting in a larger JMR for a large fraction of the explored $p_{\text{T}}^{\text{true}}$ range.

Having described the calorimeter mass resolution in some detail we now move towards quantifying the combined mass performance. In order to make a quantitative comparison of the MC-based JMR for m^{calo} and m^{comb}, the two different resolutions have been overlaid, as a function of transverse momentum of the truth jet matched to the detector-level jet. The JMR for each of the two jet mass definitions is compared in Figure 6.11, where now large mass windows have been placed on all four sample topologies. In addition to the truth jet matching, the truth jet mass is required to be in an interval containing the mass peak corresponding to the reconstructed W/Z, $H \to b\bar{b}$, or top jets. For the QCD jets, an m^{true} interval ranging up to 200 GeV is imposed. For top, W/Z, and $H \to b\bar{b}$ jets there is an absolute 2−7 % improvement in the resolution for m^{comb}, particularly at higher momenta where the jet is more boosted and the m^{comb} definition contains substantial component from the tracking information. The m^{comb} weights are derived using a sample of QCD jets and so, in topologies different from QCD jets, the expected improvement in performance for m^{comb} can be reduced in the very low $p_{\text{T}}^{\text{true}}$ regions where the m^{calo} contribution to m^{comb} is dominant. This is most evident in the very first $p_{\text{T}}^{\text{true}}$ bin, for example in the case of W/Z jets. Other important sources of under-performance in m^{comb} include the breakdown in the assumption of a Gaussian detector response and neglecting the correlations between the calorimeter and track-assisted masses[3].

[3]The effect of including these correlations has been previously studied, and has been found to be small in comparison to e.g. non-Gaussian detector effects.



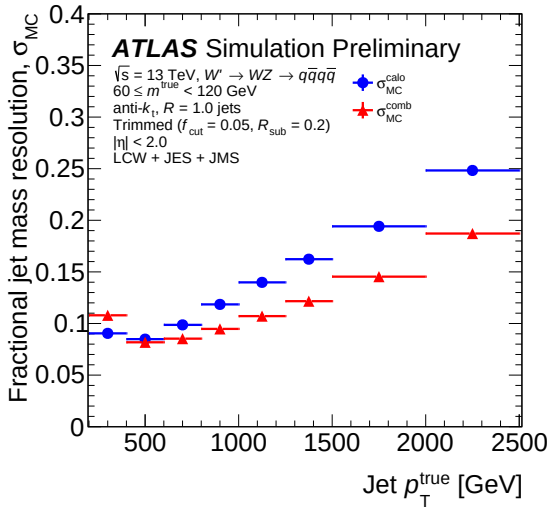
(a) W/Z jets JMR **(b)** $H \to b\bar{b}$ jets JMR

(c) Top jets JMR **(d)** QCD jets JMR

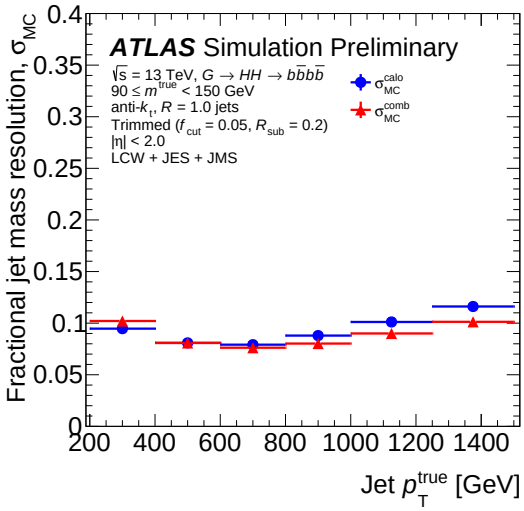
Figure 6.10: The MC-based calorimeter JMR for m^{true} of large-radius jets as a function of jet $p_{\text{T}}^{\text{true}}$, for four jet topologies: W/Z jets, top jets, $H \to b\bar{b}$ jets, and QCD jets. Representative ranges of jet m^{true} or $m^{\text{true}}/p_{\text{T}}^{\text{true}}$ are shown.

6.2.2.2 Monte Carlo JMR Estimate for Small-Radius Jets

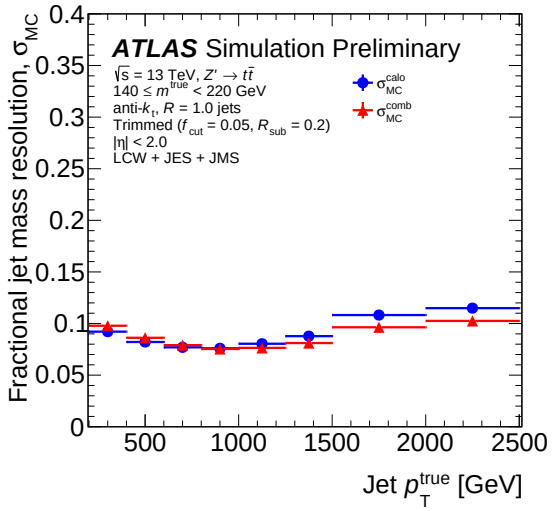
In Figure 6.12 the MC-based JMRs of three types of small-radius jets are compared: those reconstructed from electromagnetic scale topoclusters, from electromagnetic scale particle flow objects, and from hadronic scale/LCW topoclusters. As with the comparison of the calorimeter and combined mass JMRs, here a large mass window is taken in each sample. The additional tracking information utilised in particle flow provides a superior resolution across the majority of the jet $p_{\text{T}}^{\text{true}}$ range, most clearly for the QCD jets where there is up to an absolute 6 % improvement at higher $p_{\text{T}}^{\text{true}}$. The use of particle flow has previously been shown to improve the



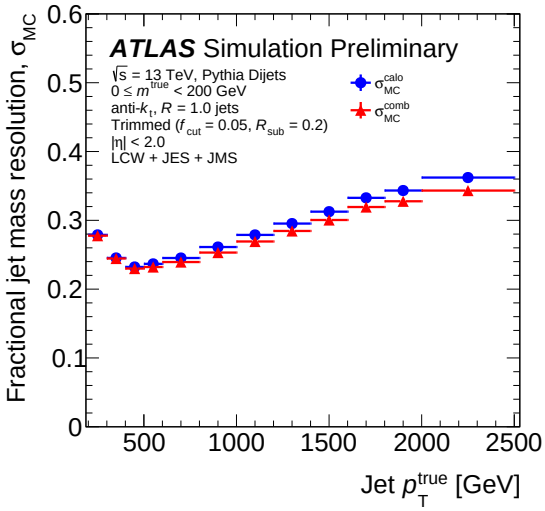
(a) W/Z jets JMR



(b) $H \to b\bar{b}$ jets JMR



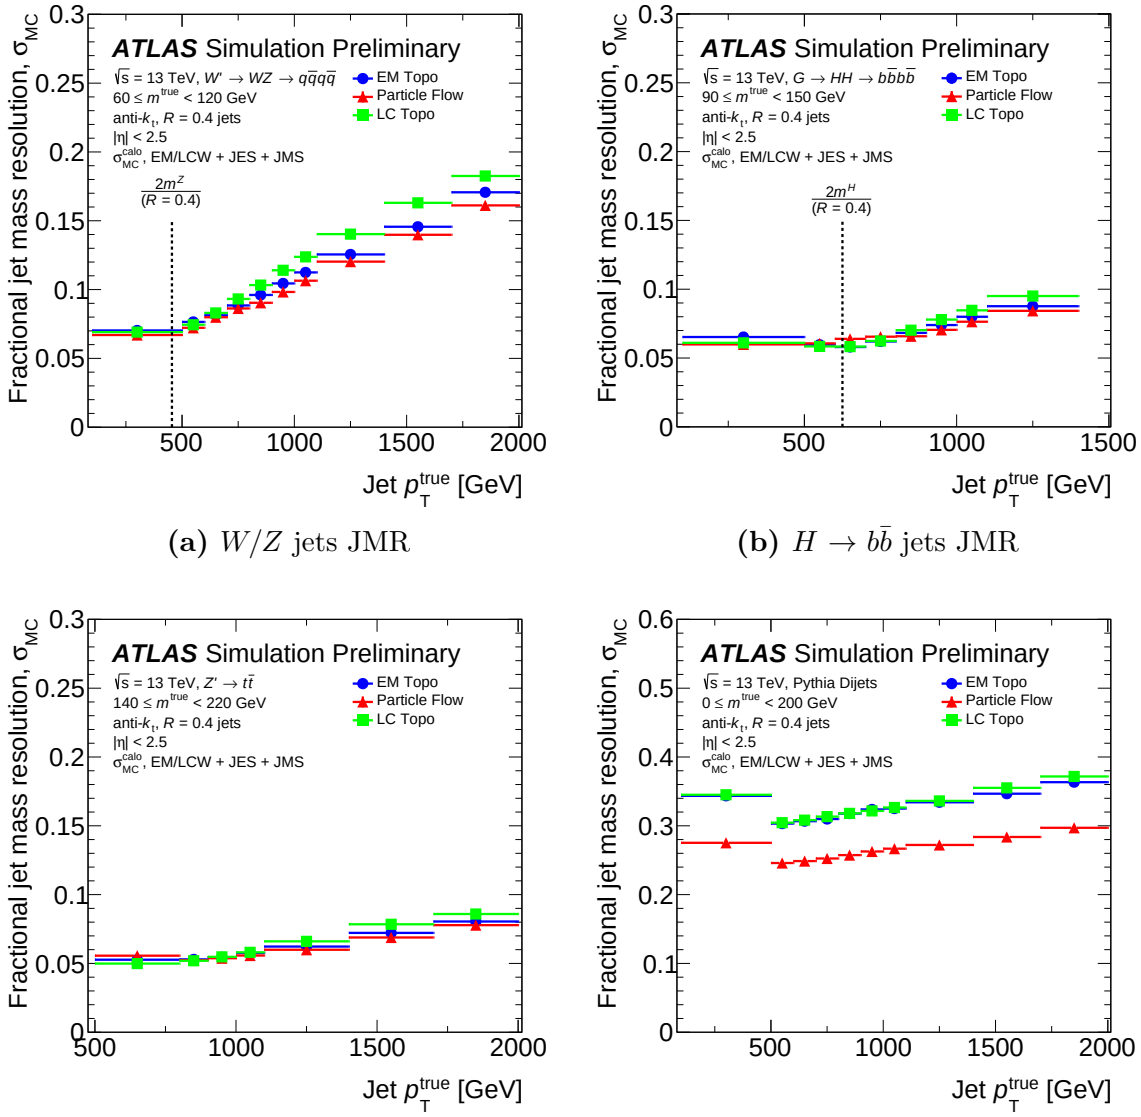
(c) Top jets JMR



(d) QCD jets JMR

Figure 6.11: The MC-based JMR for m^{calo} and m^{comb} of large-radius jets as a function of truth jet transverse momentum $p_{\text{T}}^{\text{true}}$ in jet topology-dependent m^{true} intervals. Results for four jet topologies are shown: W/Z jets, top jets, $H \to b\bar{b}$ jets, and QCD jets.

jet energy resolution at low $p_{\text{T}}^{\text{true}}$, because in that regime the tracking resolution is superior to the calorimeter jet energy resolution. This feature serves as one of the most important motivations for using particle flow jets, and the improvement in the JMR seen here provides further support for their wider usage. JMR performance of LCW-topocluster-based jets is generally similar to EM-topocluster-based jets, as both approaches rely on calorimeter-level information only.



(a) W/Z jets JMR **(b)** $H \to b\bar{b}$ jets JMR

(c) Top jets JMR **(d)** QCD jets JMR

Figure 6.12: The MC-based JMR for m^{calo} of small-radius jets as a function of truth jet transverse momentum $p_{\text{T}}^{\text{true}}$ in jet topology-dependent m^{true} intervals. Results for four jet topologies are shown: W/Z jets, top jets, $H \to b\bar{b}$ jets, and QCD jets. For each topology three distributions are shown for jets reconstructed from EM topoclusters, particle flow objects (PFO) and LCW topoclusters respectively. The black lines indicate the $p_{\text{T}}^{\text{true}}$ scale with $2m^{\text{true}}/(R = 0.4)$, corresponding to the boundaries where the particle of interest can be reasonably contained within the jet.

6.2.2.3 Systematic Uncertainties on the Monte Carlo-based Jet Mass Resolution

The MC-based JMR systematic uncertainties result from a multitude of sources encompassing effects from the detector geometry, higher-order corrections in the simulation of the final state of interest, the calculation of the non-perturbative QCD parton shower, hadronisation, and the underlying event, and the modelling of the interaction of the simulated particles with matter. During Run 1 these were estimated by varying the assumptions on each effect separately, re-calculating the JMR for the varied scenarios and quantifying the spread in the new resolutions from the nominal resolution where none of these assumptions are modified. Based on these studies [107] the resulting relative systematic uncertainty on the JMR was estimated to be 20 %. This value is also supported by comparisons of Run 1 data with simulations [117]. To understand this uncertainty, we need to build up a picture of resolution *smearing*. Resolution smearing is a way of modifying a simulated kinematic quantity in order to obtain new values which are closer to what is truly measured in data, and there will be some uncertainty associated with this process. This 20 % can be thought of as an uncertainty in over-smearing the jet mass, and it is applied directly to the MC-based JMR using the method of *Gaussian smearing* (see below), in order to account for additional detector effects and mismodellings in the original MC simulation. The MC statistical uncertainties for large-radius jets are comparable in size for both mass definitions, and are at the percent level in the bulk of the JMR phase-space. At higher transverse momentum and boost, they are slightly larger due to the combination of the falling nature of the involved distributions with the size of the simulated samples. The MC statistical uncertainty is therefore a subdominant contribution compared to the much larger detector and modelling effects, and thus it is not considered further.

To mimic the variation induced by the JMR uncertainty, the JMR is increased by multiplying the jet mass by a random number drawn from a Gaussian distribution. A way to increase the standard deviation of a Gaussian variable by a factor α is to apply a transformation on the variable that results in a scaled enhancement. This can be applied to the Monte Carlo JMR in the following way:

$$(\alpha\sigma_{\mathrm{MC}})^2 = \sigma_{\mathrm{MC}}^2 + (x\sigma_{\mathrm{MC}})^2. \tag{6.17}$$

For example, for the 20 % increase obtained from Run 1 studies, α is set to 1.2 and Equation 6.17 is solved to give $x \approx 0.66$. The Gaussian random variable used for smearing, $\mathcal{G}$, is then constructed as a Gaussian distribution of unitary mean and a standard deviation of $x\sigma_{\mathrm{MC}}$. The smearing transformation is then applied to the jet mass in the following way:

$$m^{\mathrm{reco,smeared}} = m^{\mathrm{reco}} \times \mathcal{G}(1, x\sigma_{\mathrm{MC}}). \tag{6.18}$$

Following smearing, a new JMR is calculated for $m^{\mathrm{reco,smeared}}$. If the response distribution of the smeared mass is Gaussian, this formalism guarantees that the following condition is satisfied:

$$\sigma_{\mathrm{MC}}^{\mathrm{reco,smeared}} \equiv \sigma_{\mathrm{MC}}(m^{\mathrm{reco,smeared}}/m^{\mathrm{true}}) = \alpha\sigma_{\mathrm{MC}}(m^{\mathrm{reco}}/m^{\mathrm{true}}) \tag{6.19}$$

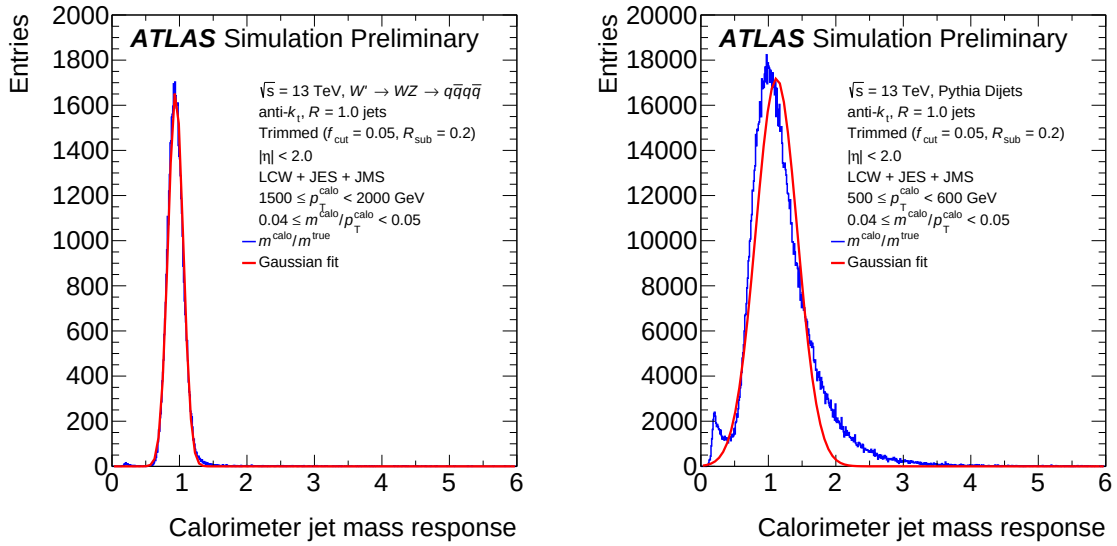
where the JMR σ_{MC} for a given response $\mathcal{R}$ is defined in Equation 6.16.

6.2.3 Investigating non-Gaussian features of the Jet Mass Resolution

The ideal case of a perfectly Gaussian detector mass response is not realised in reality. In general the jet mass response is affected by non-Gaussian features arising from e.g. the non-compensating nature of the calorimetry, and therefore deviations from $\alpha = 1.2$ are expected when comparing the smeared and original resolutions. This leads to a definition of *non-Gaussianity* that is obtained purely from MC:

$$\frac{\sigma_{\mathrm{MC}}(m^{\mathrm{reco,smeared}}/m^{\mathrm{true}})}{\sigma_{\mathrm{MC}}(m^{\mathrm{reco}}/m^{\mathrm{true}})} - \left(1 + \frac{s}{100}\right), \tag{6.20}$$

if a smearing of s % is applied. For a Gaussian detector response, this non-Gaussianity will have a value of 0. An immediate consequence of this result is that any non-Gaussianities could be inferred by looking at the jet mass response. Example jet mass responses illustrating some non-Gaussian behaviour in QCD jets, and an example of a more Gaussian response in W/Z jets, are shown in Figure 6.13. Such response features give rise to the non-Gaussianity for a 20 % smearing shown in Figures 6.14 and 6.15. These non-Gaussianities are quantified for m^{calo} and m^{comb} of large-radius jets, respectively, and are obtained relative to the truth-level transverse momentum, selecting a suitable window in truth mass for a given sample topology.



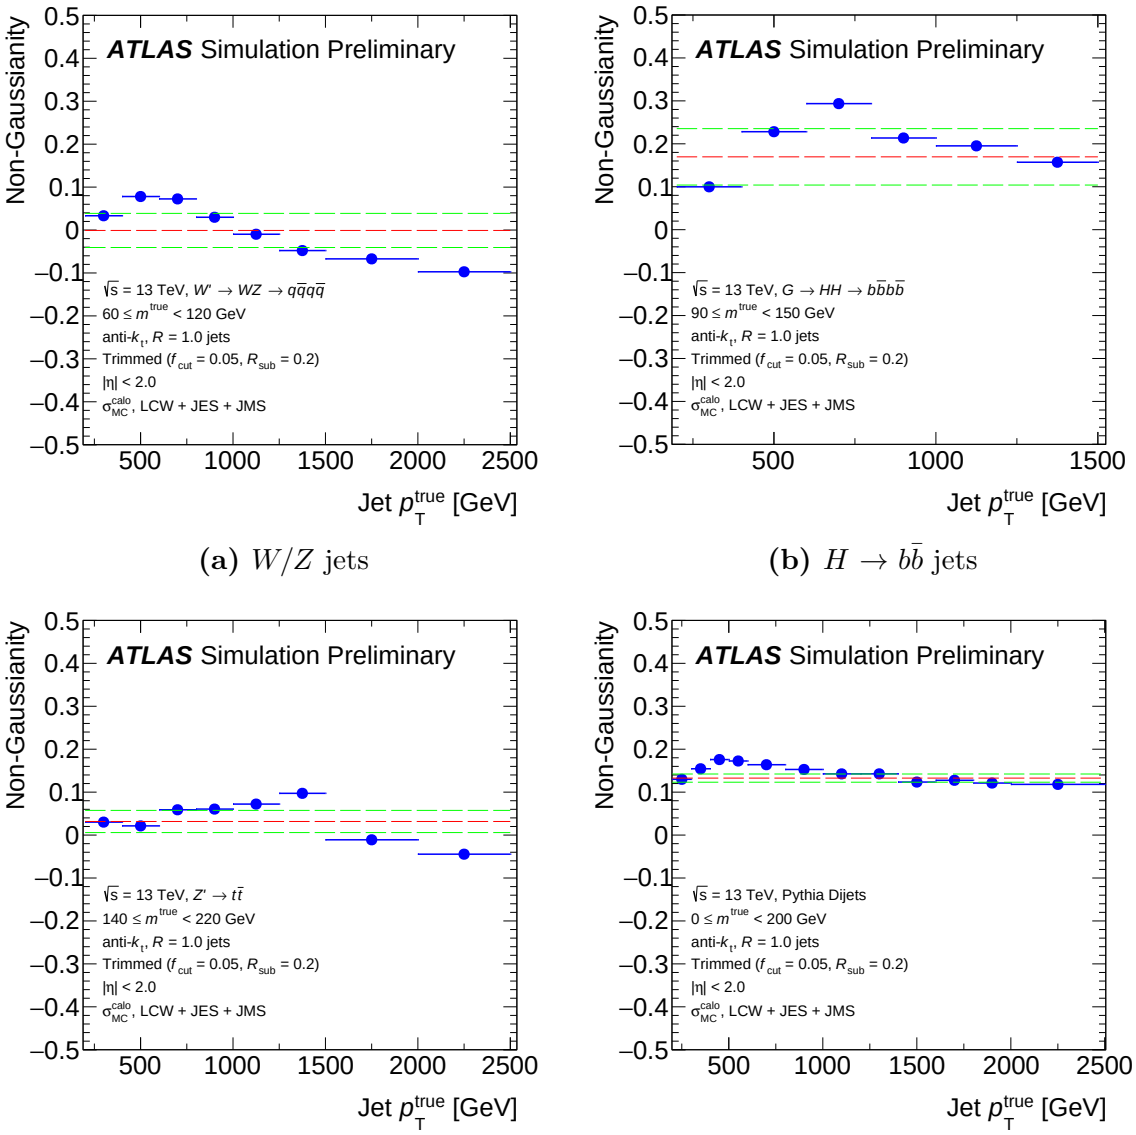
(a) W/Z jets **(b)** QCD jets

Figure 6.13: The jet mass response for large-radius jets. The calorimeter jet mass definition is used. A Gaussian fit has been overlaid for illustrative purposes. The $p_{\mathrm{T}}^{\mathrm{calo}}$ bins have been chosen for demonstrative purposes as examples of distributions which are or are not reasonably Gaussian.

Clearly the ideal smearing of $\alpha = 1.2$ is not recovered in any of the samples or jet mass definitions, i.e. a non-zero value of non-Gaussianity is obtained in all cases. The combined mass closure appears to be similar to the closure from the calorimeter mass resolution, with some improvements from the combined mass (depending on the bin of interest), such as for the $H \to b\bar{b}$ jets.

Figure 6.16 shows the impact of the non-Gaussianity for the four different sample topologies considered for the small-radius jets based on EM topoclusters (`AntiKt4EMTopo` jets). Similar levels of non-closure are observed. Differences in non-Gaussianity between small and large-radius jets are expected due to the different jet reconstruction, use of grooming for large-radius jets, and jet-area pileup correction used for small-radius jets. At low $p_{\mathrm{T}}^{\mathrm{true}}$, the $b\bar{b}$, top, and W/Z jets generally have larger statistical uncertainties for small-radius jets.

The results of this section indicate one of the important limitations of using Gaussian smearing to estimate the JMR systematic uncertainty: that the underlying responses are not Gaussian. Section 6.2.5 will explore this further, using a JMR based on information from data vs. MC. This is a second, data-driven, JMR approach. However, first the idea of *forward-folding* needs to be presented.



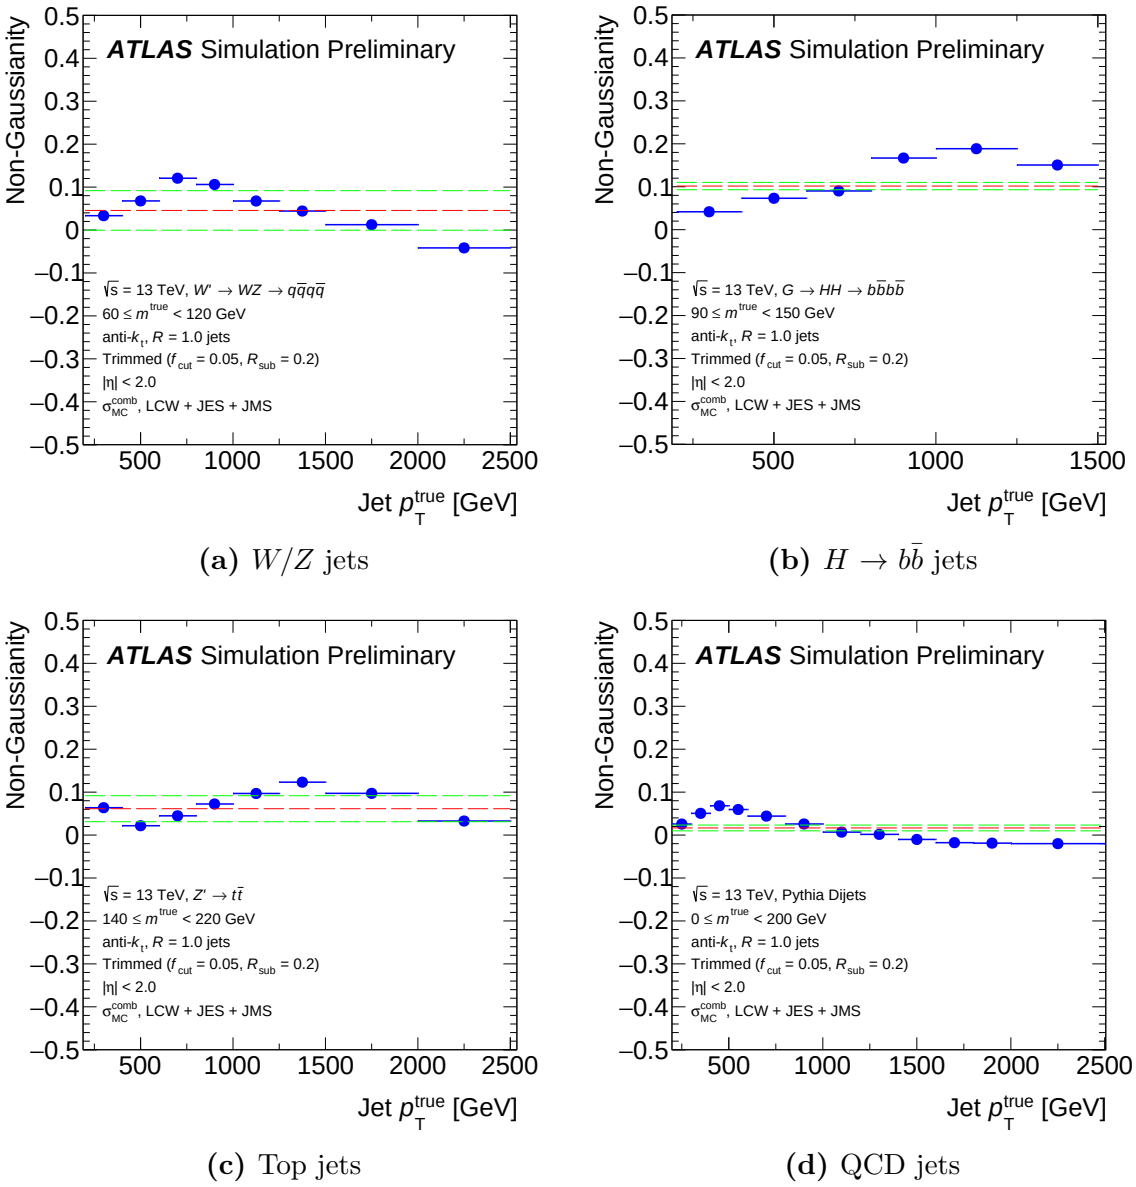
(a) W/Z jets **(b)** $H \to b\bar{b}$ jets

(c) Top jets **(d)** QCD jets

Figure 6.14: The non-Gaussianity, $(\sigma_{\mathrm{MC}}^{\mathrm{calo,smeared}}/\sigma_{\mathrm{MC}}^{\mathrm{calo}}) - 1.2$, for large-radius jets, as a function of truth jet transverse momentum. The calorimeter jet mass definition is used. A 20.0 % Gaussian smearing has been used. The red line gives the mean value, averaged across the points. The green lines provide the standard deviation across the points.

6.2.4 Jet Mass Resolution from Forward-Folding

We shall now describe the approach of forward-folding, which is an alternative method to Gaussian smearing, and is also used to modify the simulated jet mass so that it better agrees with data. Forward-folding has an immediate advantage over Gaussian smearing in that it uses Run 2 *in situ* information in the transformation it applies to the jet mass. As we saw previously,



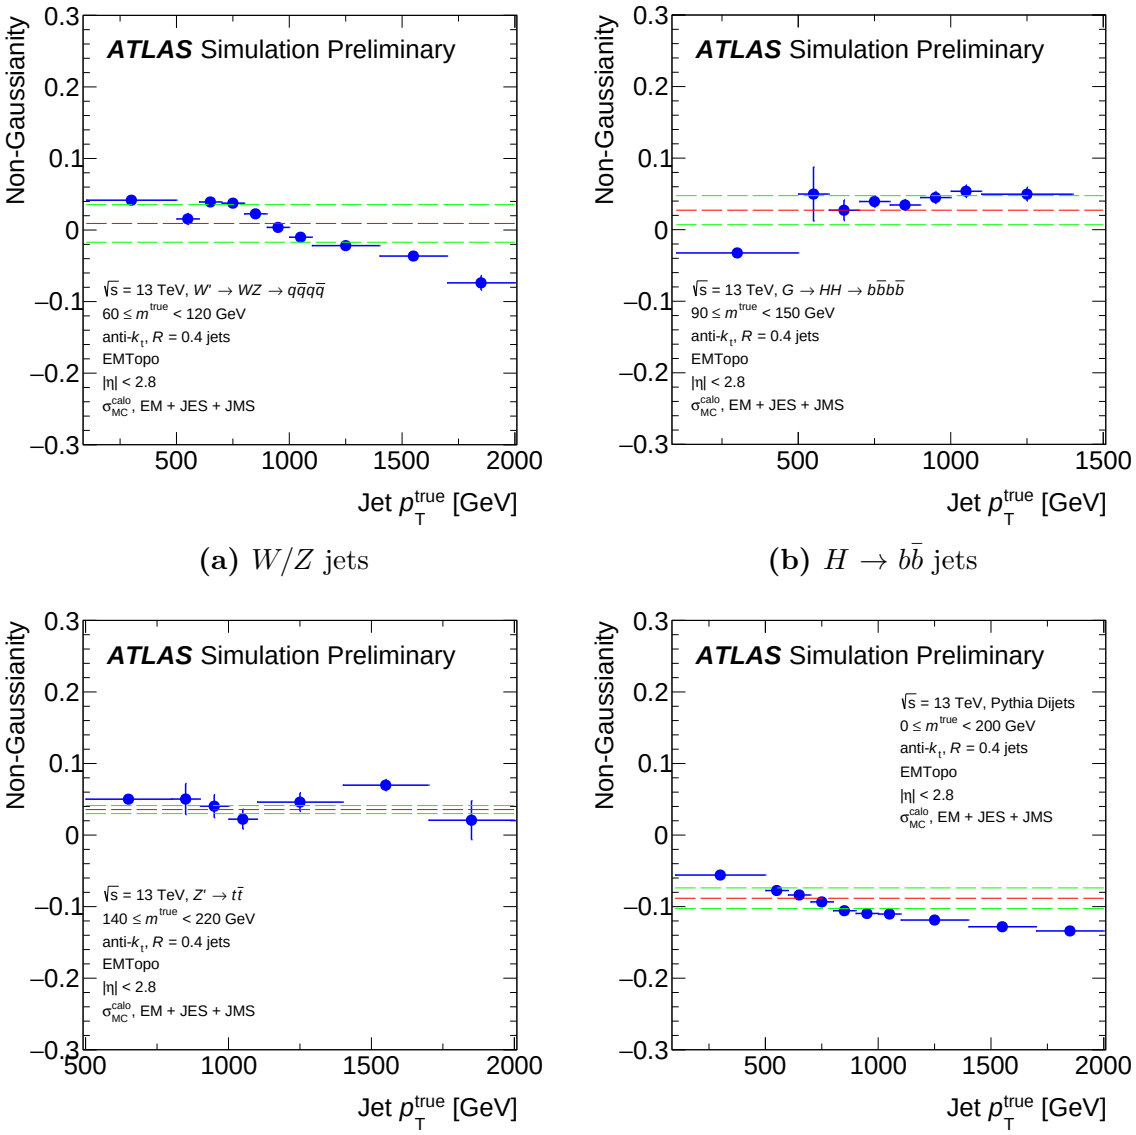
(a) W/Z jets **(b)** $H \to b\bar{b}$ jets

(c) Top jets **(d)** QCD jets

Figure 6.15: The non-Gaussianity, $(\sigma_{\mathrm{MC}}^{\mathrm{comb,smeared}}/\sigma_{\mathrm{MC}}^{\mathrm{comb}}) - 1.2$, for large-radius jets, as a function of truth jet transverse momentum. The combined jet mass definition is used. A 20.0 % Gaussian smearing has been used. The red line gives the mean value, averaged across the points. The green lines provide the standard deviation across the points.

Gaussian smearing only relies on a JMR estimated from Monte Carlo. The studies outlined here provide the first ever detailed comparison between Gaussian smearing and forward-folding, aiming to establish forward-folding as the *de-facto* method to use for all future ATLAS analyses which rely the jet mass and its systematic uncertainty.

Forward-folding measures the JMR in data relative to what is expected from Monte Carlo, where a selection has been applied in data to isolate final state jets with similar properties



(a) W/Z jets **(b)** $H \to b\bar{b}$ jets

(c) Top jets **(d)** QCD jets

Figure 6.16: The non-Gaussianity, $(\sigma_{\mathrm{MC}}^{\mathrm{calo,smeared}}/\sigma_{\mathrm{MC}}^{\mathrm{calo}}) - 1.2$, for small-radius EM-topocluster-based jets, as a function of truth jet transverse momentum. The calorimeter jet mass definition is used. A 20.0 % Gaussian smearing has been used. The red line gives the mean value, averaged across the points. The green lines provide the standard deviation across the points.

to the simulation. The associated systematic uncertainties on such a measurement are also extracted. This information is captured by a set of *in situ*-derived parameters defined below. In the present studies such parameters are then used to correct the Monte Carlo JMR to match data.

In the forward-folding transformation, the detector-level jet mass is modified in the following way (assuming the relative jet mass scale is unity [99]):

$$m_{\mathrm{FF}}^{\mathrm{reco}} = m^{\mathrm{reco}}|(m^{\mathrm{true}}, p_{\mathrm{T}}^{\mathrm{calo}}) + \Theta\left(m^{\mathrm{reco}}|(m^{\mathrm{true}}, p_{\mathrm{T}}^{\mathrm{calo}}) - \langle m^{\mathrm{reco}}|(m^{\mathrm{true}}, p_{\mathrm{T}}^{\mathrm{calo}})\rangle\right), \tag{6.21}$$

where $\langle m^{\mathrm{reco}}|(m^{\mathrm{true}}, p_{\mathrm{T}}^{\mathrm{calo}})\rangle$ is the average detector-level jet mass, given a truth mass m^{true} and a detector-level $p_{\mathrm{T}}^{\mathrm{calo}}$. The Θ parameter is derived from *in situ* studies [99] involving samples enhanced in top jets and W jets with $p_{\mathrm{T}}^{\mathrm{calo}} < 1.0$ TeV.

For the studies outlined in this chapter, the forward-folding transformation is applied to large-radius jets in simulated samples of W and top jets, where the Θ depends on the mass and transverse momentum of the jet. To perform the transformation and extract a JMR, information about the jet mass response is required. This is obtained by dividing both terms of Equation 6.21 by the truth mass, *which is fixed,* thus obtaining:

$$\mathcal{R}_{m,\mathrm{FF}} = \mathcal{R}_m|(m^{\mathrm{true}}, p_{\mathrm{T}}^{\mathrm{calo}}) + \Theta\left(\mathcal{R}_m|(m^{\mathrm{true}}, p_{\mathrm{T}}^{\mathrm{calo}}) - \langle \mathcal{R}_m|(m^{\mathrm{true}}, p_{\mathrm{T}}^{\mathrm{calo}})\rangle\right), \tag{6.22}$$

where $\langle \mathcal{R}_m|(m^{\mathrm{true}}, p_{\mathrm{T}}^{\mathrm{calo}})\rangle$ is the average jet mass response, given m^{true} and $p_{\mathrm{T}}^{\mathrm{calo}}$. In the forward-folding transformation, the value of Θ depends on the mass and on the momentum of the jet. It is more usual to quote the variable κ, defined as $\kappa \equiv 1 + \Theta$. The measured values for κ, and the associated uncertainties $\delta\kappa$, are summarised in Table 6.3. These have been obtained from the detailed studies in [99], where the relative JMR and JMS were extracted using 2016 ATLAS data, amounting to 33 fb^{-1} of integrated luminosity.

Consequently, given $m_{\mathrm{FF}}^{\mathrm{reco}}$ and $\mathcal{R}_{m,\mathrm{FF}}$, the forward-folding JMR definition is obtained from Equation 6.16 as:

$$\sigma_{\mathrm{FF}} = \frac{IQnR(\mathcal{R}_{m,\mathrm{FF}})}{2 \times \mathcal{M}(\mathcal{R}_{m,\mathrm{FF}})}, \tag{6.23}$$

where $\mathcal{R}_{m,\mathrm{FF}}$ is parametrised as a function of m^{true} and $p_{\mathrm{T}}^{\mathrm{calo}}$.

When $IQnR(\mathcal{R}_{m,\mathrm{FF}})/2$ coincides with the standard deviation of $\mathcal{R}_{m,\mathrm{FF}}$, which we will denote by $\Sigma(\mathcal{R}_{m,\mathrm{FF}})$, using the definition of $\mathcal{R}_{m,\mathrm{FF}}$ in Equation 6.22 allows one to write $IQnR(\mathcal{R}_{m,\mathrm{FF}})/2$ $= \kappa\ \Sigma(\mathcal{R}_m|(m^{\mathrm{true}}, p_{\mathrm{T}}^{\mathrm{calo}}))$. In addition, given that taking the average of Equation 6.22 results in $\langle \mathcal{R}_{m,\mathrm{FF}}\rangle = \langle \mathcal{R}_m|(m^{\mathrm{true}}, p_{\mathrm{T}}^{\mathrm{calo}})\rangle$, when $\mathcal{M}(\mathcal{R}_{m,\mathrm{FF}})$ coincides with $\langle \mathcal{R}_{m,\mathrm{FF}}\rangle$ and $IQnR(\mathcal{R}_{m,\mathrm{FF}})/2$ coincides with $\Sigma(\mathcal{R}_{m,\mathrm{FF}})$, Equation 6.23 becomes:

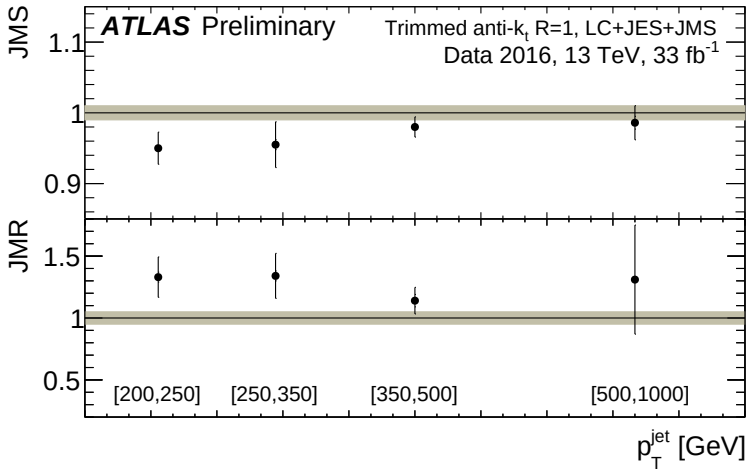


Figure 6.17: Summary of the large-radius jet *in situ* jet mass response and relative jet mass resolution, parametrised in several bins of the jet transverse momentum. These values were extracted by selecting $t\bar{t}$ events with final states containing a lepton and jets. The first two bins are to be used in W-enriched samples, and the last two bins in top-enriched samples. From [99].

$$\sigma_{\mathrm{FF}} = \frac{IQnR(\mathcal{R}_{m,\mathrm{FF}})}{2 \times \mathcal{M}(\mathcal{R}_{m,\mathrm{FF}})} = \frac{\kappa\Sigma(\mathcal{R}_m|(m^{\mathrm{true}}, p_{\mathrm{T}}^{\mathrm{calo}}))}{\langle\mathcal{R}_m|(m^{\mathrm{true}}, p_{\mathrm{T}}^{\mathrm{calo}})\rangle}. \tag{6.24}$$

As a consequence, Equation 6.24 provides the following approximation for the JMR derived from forward-folding:

$$\sigma_{\mathrm{FF}} = \kappa\sigma(\mathcal{R}_m|(m^{\mathrm{true}}, p_{\mathrm{T}}^{\mathrm{calo}})). \tag{6.25}$$

This approximation is adopted in order to estimate σ_{FF} in these studies. In this light it is relevant to emphasise the physical interpretation of the κ-value: it is the *relative* JMR, representing the ratio of the forward-folding (data-driven) JMR to the MC-based JMR, and therefore the natural parameter for forward-folding. As previously highlighted, the studies in [99] established the relative JMR from 2016 data, summarised in the lower ratio of Figure 6.17. The κ adopted in these studies are taken directly from this Figure.

The uncertainty on κ is used to derive the corresponding up (+) and down (−) systematic JMR variations $\sigma_{\mathrm{FF}\pm}$ according to:

$$\sigma_{\mathrm{FF}\pm} = (\kappa \pm \delta\kappa)\sigma(\mathcal{R}_m|(m^{\mathrm{true}}, p_{\mathrm{T}}^{\mathrm{calo}})). \tag{6.26}$$

6.2.5 JMR comparisons: Gaussian Smearing versus Forward-Folding

With forward-folding and its JMR defined, two comparisons between this new JMR and the purely MC-based, Gaussian smeared JMR can now be made: first the relative magnitudes of each JMR, and second the relative uncertainties ascribed to them. To ensure a fair comparison, we must apply a Gaussian smearing to the MC-based JMR which should be equivalent to a forward-folding scaled by κ, and in general we shall see that this smearing does not correspond to the proposed 20 % based on Run 1 studies. The comparison of the JMR magnitude for equivalent-sized Gaussian smearing and forward-folding corrections provides a further assessment of the size of non-Gaussian effects on the jet mass reconstruction.

6.2.5.1 Comparing JMR values

In the ideal case of a Gaussian detector response, the ratio of forward-folding to an MC-based JMR suitably modified with a Gaussian smearing is expected to yield unity.[4]

Jet m^{true} [GeV]	Jet $p_{\text{T}}^{\text{true}}$[GeV]	Nominal		Up-variation		Down-variation	
		κ	Equivalent x_κ	$\kappa + \delta\kappa$	Equivalent $x_\kappa + \delta x_\kappa$	$\kappa - \delta\kappa$	Equivalent $x_\kappa - \delta x_\kappa$
50 - 120	200 − 250	1.32	0.86	1.48	1.09	1.16	0.59
50 - 120	250 − 350	1.33	0.88	1.52	1.14	1.14	0.55
120 - 300	350 − 500	1.15	0.58	1.27	0.78	1.03	0.25
120 - 300	500 − 1000	1.31	0.85	1.82	1.52	0.62	0.00

Table 6.3: The different κ-values considered in the forward-folding procedure, extracted from Figure 6.17. The top two rows are the κ values derived from W jets, and the bottom two rows are derived from top jets. The corresponding $p_{\text{T}}^{\text{true}}$-values are given. The equivalent x_κ-values for a Gaussian smearing, assuming a Gaussian detector response, are also given. In the final row, the $\kappa - \delta\kappa$ of 0.62 would produce a non-real $x_\kappa - \delta x_\kappa$, and so this lower bound is set to 0 (no smearing).

Given a κ value used in forward-folding, let's first define $\sigma_{\text{MC,corr}}$ to be the MC-based JMR, which has been corrected to the value of the JMR determined *in situ* from data using forward-folding. This is achieved by applying a Gaussian smearing by x to the jet mass such that, assuming a Gaussian detector response, its JMR is the same as the forward-folding JMR derived using κ. We will denote the x value used in this Gaussian smearing to mimic the effect of a κ forward-folding x_κ, to make it explicit that it has been calculated to be equivalent to a given

[4]This property has been confirmed by performing a simulated experiment, in which a Gaussian-distributed random variable is both smeared and forward-folded on an event-by-event basis, and the ratio of the two resulting distributions is consistent with unity.

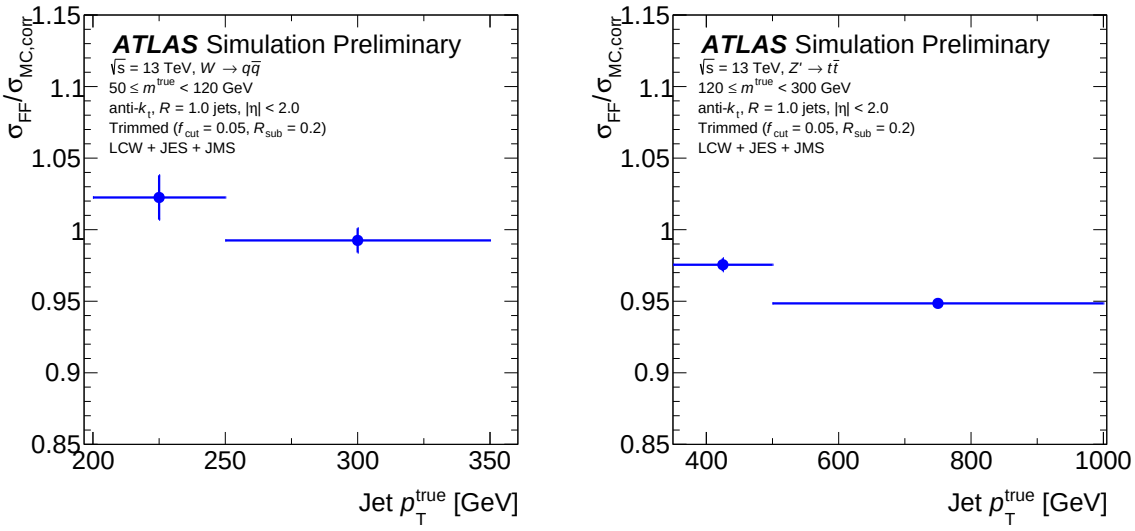
κ value assuming a Gaussian distribution. The values of σ_{FF} and the corresponding $\sigma_{\mathrm{MC,corr}}$ are referred to as the *nominal data-corrected JMR*. Now, recall that in the forward-folding procedure the κ are a function of jet m^{true} and jet $p_{\mathrm{T}}^{\mathrm{true}}$ (see Table 6.3). Therefore, in order to determine the appropriate equivalent x_κ-values, a variable smearing approach is adopted, where the Gaussian smearing also depends on the m^{true} and $p_{\mathrm{T}}^{\mathrm{true}}$ of the jet. This allows for a consistent comparison to be made with the forward-folding. The ratio $\sigma_{\mathrm{FF}}/\sigma_{\mathrm{MC,corr}}$ then directly tests the initial assumption of a Gaussian detector response necessary to match x_κ to κ, and therefore deviations from unity provide another handle on the magnitude of the non-Gaussianity in reconstructing the jet mass.

The x_κ are calculated by substituting the forward-folding coefficient κ in Equation 6.25 with the α coefficient in Equation 6.19, and by then using Equation 6.17 to derive the corresponding x_κ-value. The corresponding uncertainties on x_κ are obtained using the same approach, but now the κ-value is replaced by its uncertainty, $\delta\kappa$. The results of this procedure are reported in Table 6.3. These new x_κ-values are then used to construct a Gaussian smearing function where the x_κ-value changes as κ changes. The $p_{\mathrm{T}}^{\mathrm{true}}$ ranges considered are different for the two different sample topologies for which the κ parameter is available (as shown in Figure 6.17): first a W jets topology where 200 GeV $< p_{\mathrm{T}}^{\mathrm{true}} <$ 350 GeV, and then a top jets topology where 350 GeV $\leq p_{\mathrm{T}}^{\mathrm{true}} <$ 1000 GeV.

The new jet mass- and jet $p_{\mathrm{T}}^{\mathrm{true}}$-dependent Gaussian smearing scenario is compared with the forward-folding scenario in Figure 6.18. The magnitude of the ratio is within about 5 % of unity, thus quantifying the effects from the non-Gaussianity to be at the 5 % level.

6.2.5.2 Comparing JMR Uncertainties

Given the systematic uncertainty on forward-folding JMR defined in Section 6.2.4, let's now denote the up (+) and down (−) variations of the nominal $\sigma_{\mathrm{MC,corr}}$, by $\sigma_{\mathrm{MC,corr+}}$ and $\sigma_{\mathrm{MC,corr-}}$. These JMRs depend on $x_\kappa + \delta x_\kappa$ and $x_\kappa - \delta x_\kappa$, respectively. Recalling the discussion of the previous section, it should now be clear that these variables correspond to an MC-based JMR with Gaussian smearing of $x_\kappa \pm \delta x_\kappa$ applied, such that it has an equal magnitude to a JMR



(a) W jets **(b)**

Figure 6.18: The ratio $\sigma_{\mathrm{FF}}/\sigma_{\mathrm{MC,corr}}$ for W and top jets using the κ- factors derived in previous studies, and a variable Gaussian smearing, with m^{true} windows as specified in the plots. The jets are smeared in a p_{T}-dependent fashion. The error bars represent the uncertainty due to the finite sample size.

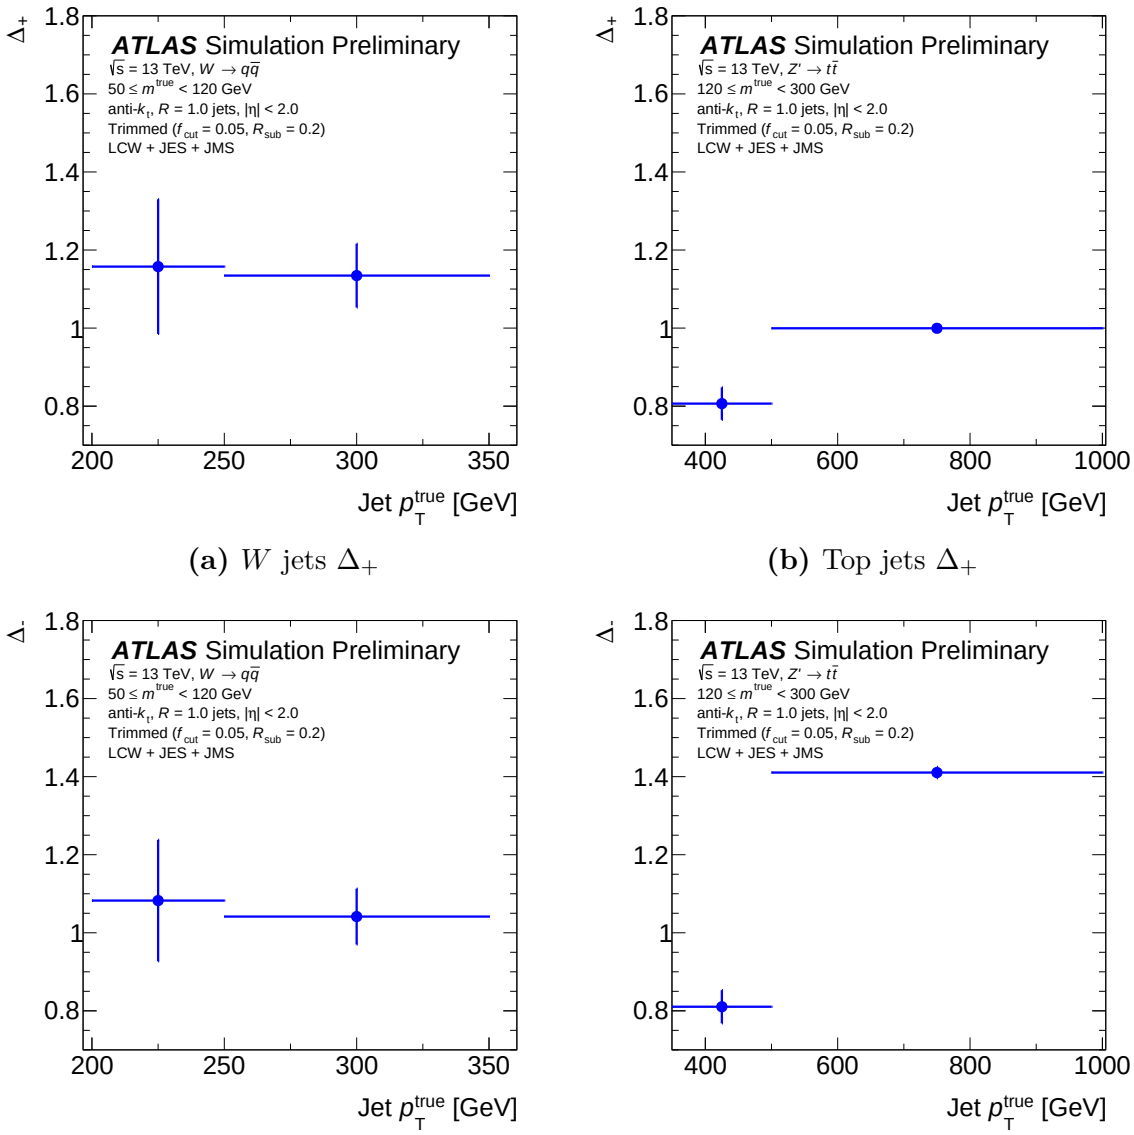
obtained from forward-folding with $\kappa \pm \delta\kappa$, assuming a Gaussian detector response.

The relative size of the propagated JMR systematic uncertainty is then compared by defining two ratios: $\Delta_{\pm}$, corresponding to the up $(+)$ and down $(-)$ systematic variations respectively:

$$\Delta_{\pm} = \frac{(\sigma_{\mathrm{FF}\pm} - \sigma_{\mathrm{FF}})/\sigma_{\mathrm{FF}}}{(\sigma_{\mathrm{MC,corr}\pm} - \sigma_{\mathrm{MC,corr}})/\sigma_{\mathrm{MC,corr}}}. \tag{6.27}$$

Table 6.4 gives the overview of the 'nominal' JMRs and of their systematic uncertainties. The numerator and the denominator of Δ measure the difference between the nominal JMR and the JMR including systematic uncertainties (this is just the difference 'nominal + uncertainty' − 'nominal') for the forward-folding and the corresponding Gaussian smearing approach respectively, normalised by the nominal value.

The values of the $\Delta_{\pm}$ are shown in Figure 6.19 and demonstrate that there is a difference between the relative sizes of the two systematic uncertainties propagated in each of these JMR approaches. For W jets with $p_{\mathrm{T}}^{\mathrm{true}} < 350$ GeV, the sizes of the systematic uncertainties associated with forward-folding are consistent with what is observed in the case of the Gaussian smearing. However, for top jets with 350 GeV $< p_{\mathrm{T}}^{\mathrm{true}} < 500$ GeV the forward-folding systematic uncertainties are consistent with being as much as 20 % smaller than the size of the uncertainties



(a) W jets Δ_+ **(b)** Top jets Δ_+

(c) W jets Δ_- **(d)** Top jets Δ_-

Figure 6.19: The Δ ratios for W and top quark jets. Mass and p_{T} windows, as specified on the plots, reflect the mass and p_{T} where forward-folding κ parameters have been derived. The jets are smeared in a p_{T}-dependent fashion. The up variations are calculated by considering a smearing of $x_\kappa + \delta x_\kappa$. The down variations are calculated by considering a smearing of $x_\kappa - \delta x_\kappa$. The error bars represent the uncertainty due to the finite sample size. The large difference in the final bin of Δ_- is due to a large $\delta\kappa$ and an undersmeared $x_\kappa - \delta x_\kappa$ (see Table 6.3).

	Nominal	Up-variation	Down-variation
MC-based JMR	σ_{MC}	N. A.	N. A.
MC-based JMR corrected to data using Gaussian smearing	$\sigma_{\mathrm{MC,corr}}$	$\sigma_{\mathrm{MC,corr+}}$	$\sigma_{\mathrm{MC,corr-}}$
Data-corrected (forward-folding) JMR	σ_{FF}	$\sigma_{\mathrm{FF+}}$	$\sigma_{\mathrm{FF-}}$

Table 6.4: The relationship between the different JMR definitions. The nominal JMRs are presented in the first column. The second and third column represent the nominal JMRs which have been modified by additional systematic variations, coming from δx_κ in the Gaussian smearing approach and $\delta\kappa$ in the forward-folding approach. Systematic variations are not named (N. A. = Not Applicable) for the nominal MC-based JMR.

associated with Gaussian smearing for for Δ_+. Furthermore, the forward-folding uncertainties are then observed to be up to 40 % larger than the Gaussian smearing uncertainties in Δ_-, for $p_{\mathrm{T}}^{\mathrm{true}} > 500$ GeV. The explanation behind this sudden jump in the size of the forward-folding uncertainty may have multiple contributions, including differences in the origin of the jet (W vs. top), as well as the large increase in $\delta\kappa$ for $p_{\mathrm{T}}^{\mathrm{true}} > 500$ GeV, as shown in Table 6.3. There is also an impact from the variable smearing approach itself: the $x_\kappa - \delta x_\kappa$ variation for $p_{\mathrm{T}}^{\mathrm{true}} > 500$ GeV requires a Gaussian under-smearing which would not be a real number (i.e. $(x_\kappa - \delta x_\kappa)^2 < 0$), and therefore no smearing is applied as an approximation. It is therefore reasonable to partially attribute this 40 % difference between the two methods to the inability of the Gaussian smearing approach to correctly modify the JMR. This demonstrates a fundamental advantage of the forward-folding procedure in the treatment of JMR systematic uncertainties. It is clear that the previous Run 1 procedure of assessing JMR sytematic uncertainties with Gaussian smearing needs to be reworked to an approach inspired by the *in situ* forward-folding methodology. This is now an ongoing and active area of research within ATLAS, and new recommendations for deriving the JMR systematic uncertainty based on forward-folding with Run 2 data are expected in 2019/2020, inspired by the detailed studies summarised in this chapter.

7

Optimising Large-Radius Jets in Run 2 of ATLAS

During Run 1 of the LHC, large-radius jet reconstruction was studied extensively [90] and the optimal reconstruction algorithm for such large-radius jets was found to be based on anti-k_t reconstruction with a radius parameter of $R = 1.0$. These jets were constructed from locally calibrated topological clusters, and were groomed with the trimming algorithm. Different values of the trimming parameters (subjet radius, R_{sub}, and subjet transverse momentum fraction, f_{cut}) were studied and an optimal set of values obtained. Trimming and `LCTopo` inputs have already been described in some detail, as has the 'standard' large-radius jet definition in ATLAS: `AntiKt10LCTopoTrimmedPtFrac5SmallR20`. This final choice of jet algorithm, input, groomer, and grooming parameters was based on the ability of the reconstructed jet to discriminate some signal from background, as well as the robustness of the jet reconstruction in the presence of pileup.

Since these `AntiKt10LCTopoTrimmedPtFrac5SmallR20` jets were determined to be an optimal large-radius jet definition in Run 1, there has been a renaissance in the field of jet grooming and jet reconstruction. Developments include a series of detailed papers proposing new jet grooming algorithms [118–120], new methods of pileup mitigation [121, 122] beyond the more standard approach of jet area subtraction, and consequently the creation of new classes of inputs to the jet reconstruction. With a myriad of new jet groomers and inputs on the market, as well as the significant gain in luminosity and pileup during Run 2 of ATLAS, it is natural to

wonder if the original 'optimal' Run 1 jets are still performing the best out of a broad suite of jet definitions, or can we do better and arrive at a new large-radius jet definition that is more robust against pileup and offers improved signal and background discrimination? It is this question that drives the studies summarised in this Chapter. The end goal is to determine new, more sophisticated $R = 1.0$ jet definitions that will ensure that future analyses using the full ATLAS Run 2 data-set utilise the best jet reconstruction techniques, in order to enhance that final measurement or sensitivity to New Physics. Another important consideration (not described in this Chapter, but discussed briefly in Appendix A) is the performance of both current and new large-R jets in future, upgraded incarnations of the ATLAS detector.

This Chapter outlines the Run 2 jet reconstruction studies which we shall refer to as the *jet inputs and grooming scan*. In these studies, hundreds of different $R = 1.0$ jet definitions are considered, which are constructed from different combinations of new jet groomers, grooming parameters, and new jet inputs. Several metrics of 'good' jet performance are investigated: the ability to discriminate W reconstruction from QCD background, the jet mass scale and resolution, the pileup dependence of *jet substructure variables* (we shall describe these in more detail shortly), and the ability of these jets to identify, or 'tag', bosons and top quarks. These initial studies [123] are based entirely on Monte Carlo simulation, but ultimately will be followed by dedicated data-driven assessments of the jet performance. Of the hundreds of new jet definitions, the most promising jets from this scan were put in to a 'shortlist', and new jet energy and mass calibrations were obtained in order to confirm that calibration at the jet substructure level doesn't modify the conclusions of the initial longlist inputs and grooming scan. Before describing the studies of the scan in detail, it is necessary to summarise the new inputs and groomers that have been scrutinised. Therefore these shall first be discussed in some detail.

7.1 New Jet Inputs

The initial set of new jet inputs are modifications to the traditional locally calibrated topoclusters by including pileup subtraction at the topocluster (or jet constituents) level using the techniques of *Voronoi subtraction* and *constituent subtraction* [121], and afterward applying an additional pileup reduction procedure which cuts on the jet p_{T}, known as *SoftKiller* [122]. The very latest inputs and grooming scan now includes a class of inputs which exploit additional detector

information from tracking in the jet reconstruction. These include modifying clusters by associating tracks and subtracting off the matched cluster energy, which is the *particle flow* [105] procedure already outlined in Section 6.1.6, as well as the association of tracks in order to reduce the angular resolution of the final jet; these new objects are referred to as *track-calo clusters* and are an ATLAS innovation which have been shown to greatly enhance sensitivity to searches for new physics in specific dijets channels [124]. We shall not focus on these track-based inputs in this thesis, but they are under active study and shall follow in additional publications in 2019.

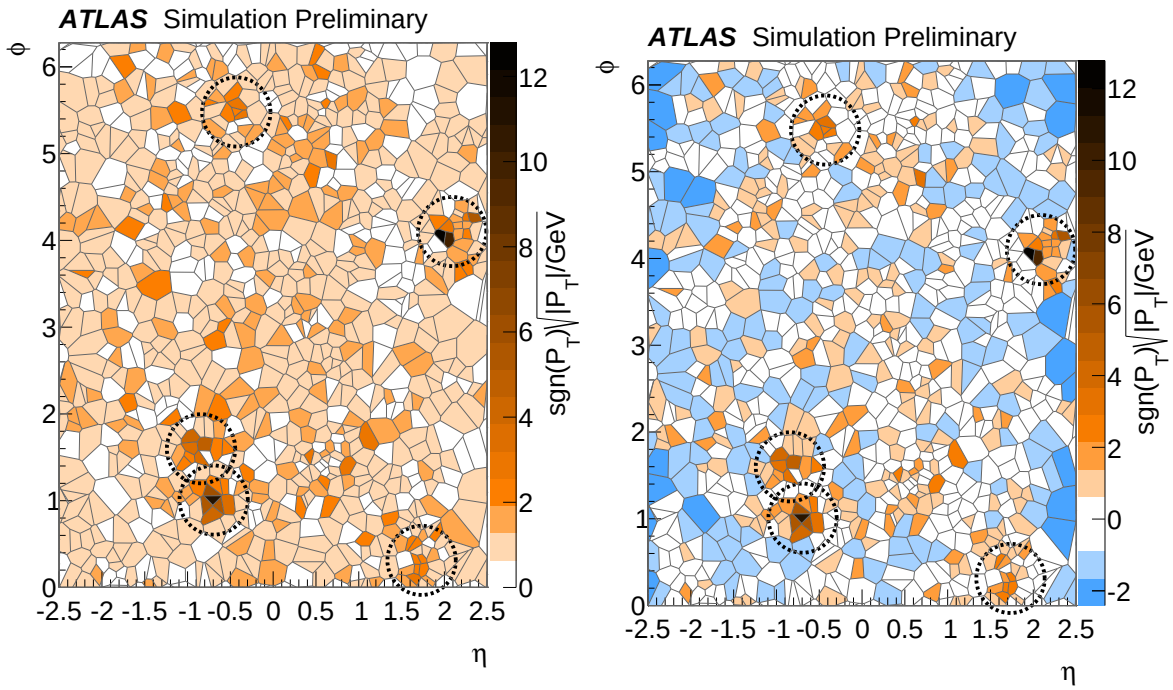
7.1.1 Voronoi Subtraction

The method of Voronoi subtraction extends the original jet area ρ-subtraction [92, 93], which is applied at the jet reconstruction level, to the jet constituents level. First, let's define the *Voronoi area*, A_{Voronoi}, of the i^{th} jet constituent (e.g. a topocluster in $\eta - \phi$) to be the area of all the points in $\eta - \phi$ that are closer to this i^{th} constituent than to any other constituent in the event. As discussed in Section 5.3.1.2, the ρ is the median ratio of the jet transverse momentum to jet area for the radius R jet reconstructed in the event. The Voronoi area is then subtracted from *each constituent* (again, labelling constituents with i) in the following way:

$$p_{\text{T},i}^{\text{corr}} = p_{\text{T},i} - \rho \times A_{\text{Voronoi}}. \tag{7.1}$$

The effects of applying Voronoi subtraction are summarised in Figure 7.1, which shows the Voronoi areas in a simulated event in $\eta - \phi$, before and after the Voronoi area subtraction. Clearly the subtraction produces many constituents with negative $p_{\text{T}}^{\text{corr}}$; we need an approach to deal with such situations, and two common procedures adopted are *negative suppression* and *Voronoi Spreading*. For negative suppression, the clusters with negative $p_{\text{T}}^{\text{corr}}$ are simply removed from the event; they are not considered in the subsequent jet clustering or calculation of the energy in the event arising from hard-scatter. Alternatively, the procedure of Voronoi Spreading seeks to suppress the negative energy clusters, rather than remove them entirely. The negative $p_{\text{T}}^{\text{corr}}$ in a given cluster is now spread across several nearby clusters. First consider some i^{th} cluster with negative $p_{\text{T}}^{\text{corr}}$, and for each cluster i:

1. Find all $j \in S$ clusters in the event with an angular separation between i, ΔR_{ij}, less than 0.4.
2. Calculate a weight for each cluster j, with respect to the target cluster, i:



(a) Simulated event, before Voronoi subtraction.

(b) Simulated event, after Voronoi subtraction.

Figure 7.1: A sample simulated event display with an average of 20 interactions per pp bunch-crossing. The z-axis describes the relative sign of the transverse momentum of each Voronoi cell, relative to its square root. The same event is shown before and after Voronoi subtraction, where in the latter case several negative cells are generated through the subtraction process. From [125].

$$w_{ij} = \frac{1}{\Delta R_{ij}^2} \times \frac{1}{\sum_{k \in S} \frac{1}{\Delta R_{ik}^2}}, \tag{7.2}$$

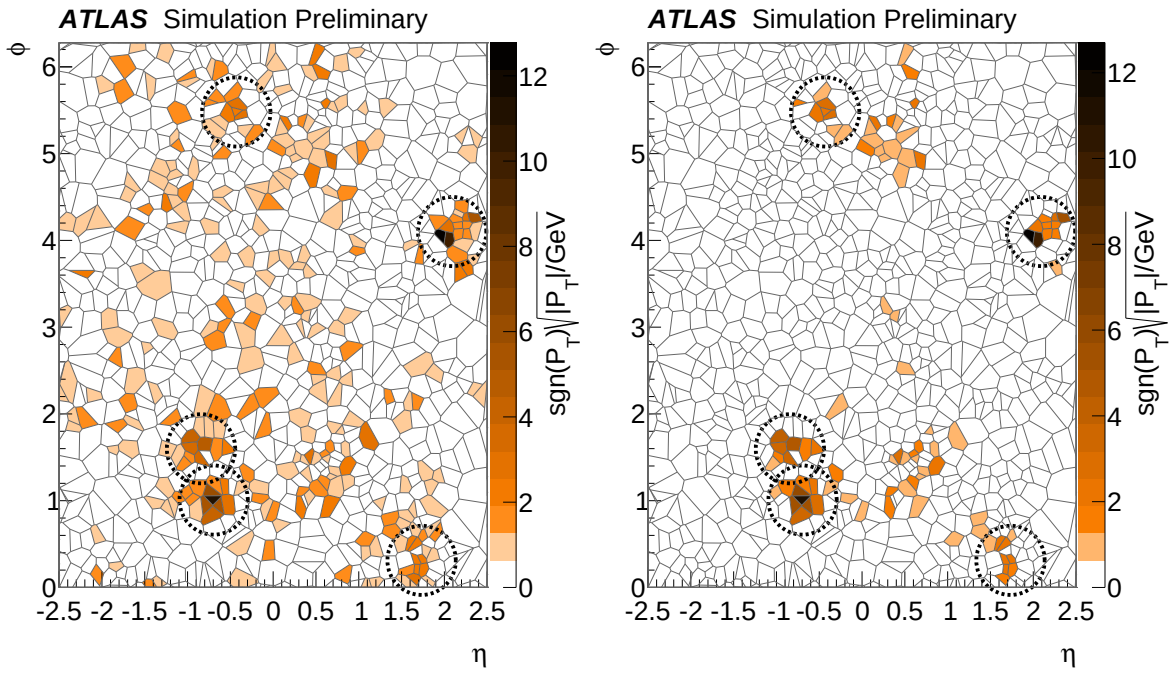
such that the sum of all w_{ij} across the j clusters is unity. This weight corresponds physically to the relative influence of a given cluster on its neighbours.

3. Following this, the negative transverse momentum of the target i can now be spread across all $j \in S$ according to:

$$p_{\mathrm{T},j}^{\mathrm{corr}} \to \max(0, p_{\mathrm{T},j}^{\mathrm{corr}} + w_{ij} \times p_{\mathrm{T},i}^{\mathrm{corr}}). \tag{7.3}$$

4. Once the target cluster's transverse momentum has been spread, set its transverse momentum value to zero.

Figure 7.2 shows the effects of the application of negative suppression and Voronoi spreading of the target clusters, after the initial Voronoi subtraction. The black dotted circles indicate truth-level jets with a transverse momentum over 20 GeV. What can be seen is that the application



(a) Simulated event, after Voronoi subtraction and suppression.

(b) Simulated event, after Voronoi subtraction and spreading.

Figure 7.2: A sample simulated event display with an average of 20 interactions per pp bunch-crossing. The z-axis describes the relative sign of the transverse momentum of each Voronoi cell, relative to its square root. The same event is shown after Voronoi subtraction, followed by either an additional Voronoi suppression or Voronoi spreading. From [125].

of the suppression and spreading preserves such jets. Both the Voronoi-suppressed and Voronoi-spread inputs are considered in these studies, and will be referred to as `VorSpreadLCTopo` and `VorSuppLCTopo` inputs respectively.

7.1.2 Constituent Subtraction

A second approach to pileup subtraction at the jet constituents level is the aptly-named *Constituent subtraction* [121]. The constituent subtraction approach uses an area calculated by populating the event, in $\eta - \phi$, with 'ghost' particles (the same 'ghosts' used in the method of ghost-association), rather than the purely geometrical area calculation of Voronoi subtraction. The ρ is computed from the median transverse momentum density in the same manner as in Voronoi and jet area subtraction. The subtraction algorithm begins by populating the event with these virtual ghost particles, which have a transverse momentum satisfying:

$$p_{\mathrm{T}}^{g} = A_g \times \rho. \tag{7.4}$$

Here A_g is the area of the ghost particle, which is set to 0.01, and in this way the p_{T}^{g} mimics the typical pileup observed in a small $\eta-\phi$ slice of 0.1×0.1. Once these ghosts have been added to the event, the next step is to calculate the angular separation between the i^{th} cluster and k^{th} ghost:

$$\Delta R_{ik} = \sqrt{(\eta_i - \eta_k^g)^2 + (\phi_i - \phi_k^g)^2}. \tag{7.5}$$

Once all cluster-ghost combinations of angular separations have been calculated, the associated cluster-ghost pairs are then sorted in order of ascending ΔR_{ik}. The algorithm now proceeds iteratively through each of these cluster-ghost, i-k, pairs, modifying each cluster according to the information from the associated ghost.

$$\text{If} : p_{\mathrm{T},i} \geq p_{\mathrm{T},k}^{g} \text{:}\ p_{\mathrm{T},i} \to p_{\mathrm{T},i} - p_{\mathrm{T},k}^{g};\ p_{\mathrm{T},k}^{g} \to 0. \tag{7.6}$$

$$\text{Otherwise} : p_{\mathrm{T},k}^{g} \to p_{\mathrm{T},k}^{g} - p_{\mathrm{T},i},\ p_{\mathrm{T},i} \to 0. \tag{7.7}$$

This algorithm is continued until some $\Delta R_{ik} > \Delta R_{\mathrm{max}}$ is achieved. This ΔR_{max} is just a tunable parameter in the algorithm, and for the grooming and inputs scan a value of 0.25 was chosen based on dedicated ATLAS studies described in [125]. Inputs based on topoclusters with this additional constituent subtraction applied shall be denoted `CSLCTopo`.

7.1.3 SoftKiller

Finally, the *SoftKiller* (SK) [122] approach simply uses the p_{T} of a given particle to classify whether it originates from a hard-scatter event, or from pileup. The SoftKiller algorithm makes a hard transverse momentum cut on an event-by-event basis by first splitting the entire event in to $\eta - \phi$ grids of some scale l, and applying a cut on the p_{T} such that afterwards half of the grids in the event are empty. This scale l, which is just the length of one of the side of an $\eta - \phi$ grid, is the tunable parameter in SoftKiller, and different values have been investigated in the ATLAS studies summarised in [125]. For the inputs and grooming scan, a grid size of 0.6×0.6 has been used. The area-based pileup mitigation techniques are applied to the clusters *before they are provided to the SoftKiller algorithm*, because doing so means that the cluster four-vectors can be appropriately corrected before a final decision is made to remove them entirely from the event. The inputs based on the combination of constituent subtraction and SoftKiller,

or Voronoi subtraction and SoftKiller, we shall denote `CSSKLCTopo`, and `VorSuppSKLCTopo` or `VorSpreadSKLCTopo` respectively.

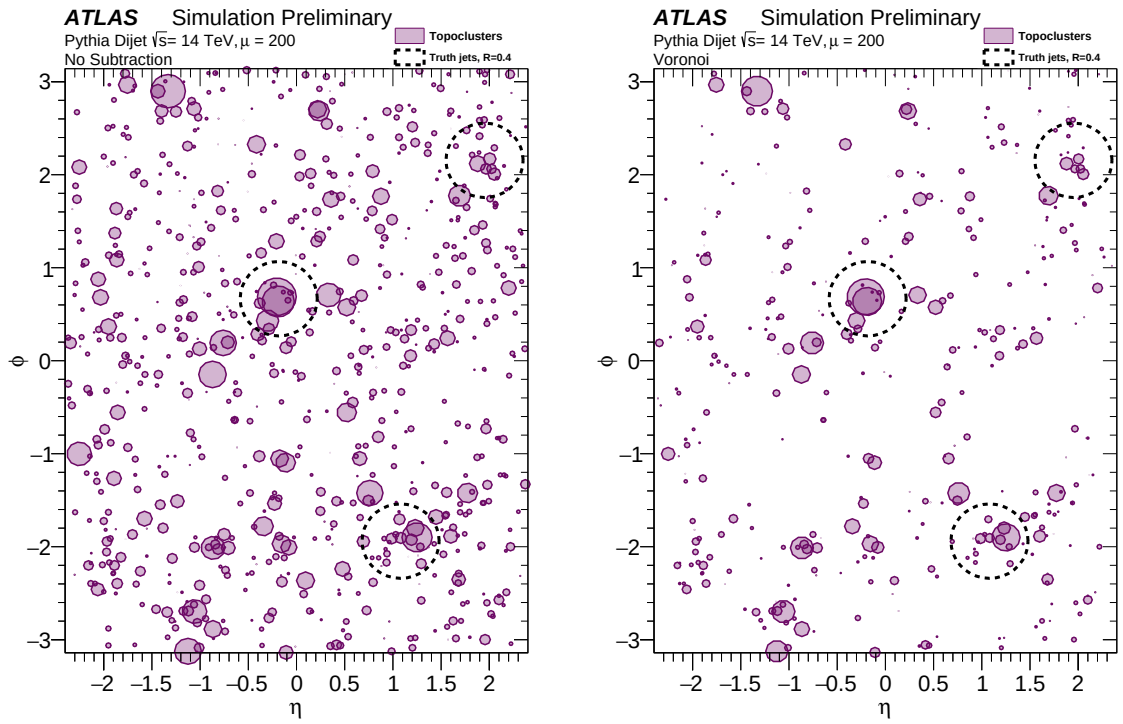
The sequential effects of applying Voronoi subtraction (and suppression), followed by SoftKiller, are shown in Figure 7.3 for a high pileup event (average interactions per bunch-crossing, $\langle\mu\rangle = 200$) of Pythia 8 QCD dijets with a pp collision energy of 14 TeV, expected for the High-Luminosity upgrade of the LHC. It can be seen that the combination of the subtraction and SoftKiller leaves an event which is dominated by the hardest jet constituents. Following this, jet reconstruction can now be applied in order to generate new `VorSuppSKLCTopo` jets.

7.2 New Jet Groomers

Jet trimming has been the ATLAS workhorse of jet grooming since its implementation following the extensive Run 1 jet definition optimisation studies [90], and has already been described in some detail (Section 5.3.2) following the discussion of `AntiKt10LCTopoTrimmedPtFrac5SmallR20` jet reconstruction. Another popular jet grooming procedure is *pruning* [126], which was proposed by the theory community shortly after trimming. Since the implementation of `AntiKt10LCTopoTrimmedPtFrac5SmallR20` jets in ATLAS there have been two major developments in jet grooming: first the introduction of the *Modified Mass Drop Tagger* (mMDT) [118] in 2013 and, crucially the 2014 generalisation of this mMDT algorithm, referred to as *Soft Drop* [119]. Soft Drop is a highly favoured groomer for theorists, and has been employed in CMS large-radius jet reconstruction for some time, whereas ATLAS lacked any dedicated study of its performance in jet reconstruction. Even in the past two years, Soft Drop has been further extended, spawning new algorithms of *recursive Soft Drop* and *bottom-up Soft Drop* [120] which shall be studied in some detail. There is also a relatively new approach to quark-gluon discrimination using a new technique of *iterative Soft Drop* [127], but we shall not explore that particular facet of Soft Drop declustering any further.

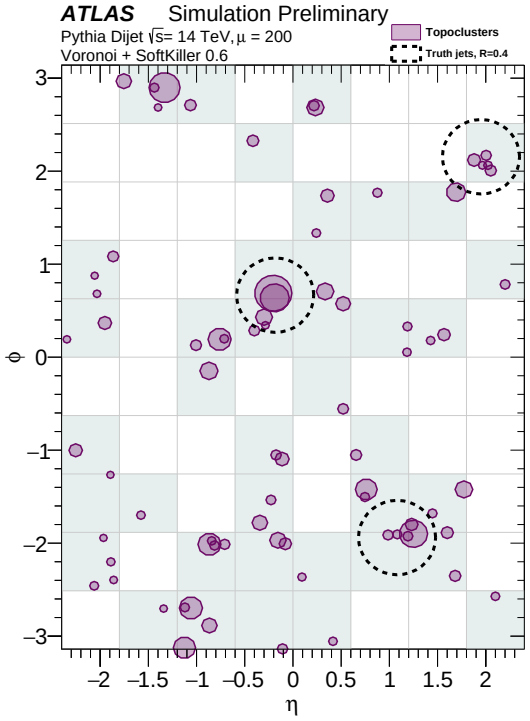
7.2.1 Pruning

The motivation of *pruning* [126] is to find a way to algorithmically modify the clustering sequence of a reconstructed jet, such that specific splittings which make up that clustering sequence, and which pull in contributions from pileup and the underlying event, are discarded. To achieve this



(a) Simulated event, before Voronoi subtraction.

(b) Simulated event, after Voronoi subtraction and suppression.



(c) Simulated event, after Voronoi subtraction and suppression, and SoftKiller with 0.6×0.6 grid spacing.

Figure 7.3: A simulated Pythia 8 QCD dijets event-display using pileup conditions typical for the HL-LHC. The effects of Voronoi subtraction and suppression, followed by SoftKiller, are shown. From [125].

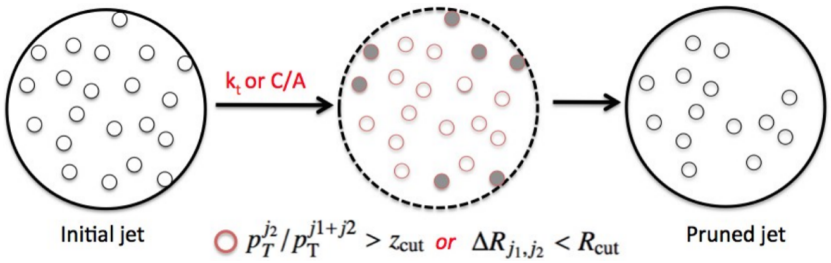


Figure 7.4: A cartoon of the pruning algorithm, including wide-angle and soft conditions. From [90].

a pruning radius must be determined, such that the hardest prongs/splittings of the jet fall in to separate subjets, allowing softer radiation falling outside of these prongs to be removed. The algorithm commences by taking the constituents of the initial large-radius jet and reclustering them using the Cambridge-Aachen algorithm[1]. This forms a clustering sequence, which is supposed to represent the intuitive parton shower branching of the jet based on the angular ordering of the constituents.

As the name suggests, the sequence is split into a series of clusters, j_1, j_2, ... etc. At each merging step in the cluster sequence, $j_1 + j_2 \rightarrow j_3$, a check is conducted to see if the clustering is either soft or wide-angled using the following conditions:

$$\text{Wide} - \text{angle} : \Delta R(j_1, j_2) \geq R_{\text{cut}} \times \frac{2m(j_1, j_2)}{p_{\text{T}}(j_1, j_2)} \tag{7.8}$$

$$\text{Soft} : z \leq z_{\text{cut}}, \tag{7.9}$$

where $\Delta R(j_1, j_2)$, $m(j_1, j_2)$, and $p_{\text{T}}(j_1, j_2)$ are the angular distance, mass, and transverse momentum from the potential cluster combination j_1 and j_2, and $z = \frac{\min(p_{\text{T}}(j_1), p_{\text{T}}(j_2))}{p_{\text{T}}(j_1, j_2)}$, assuming the j_1 cluster is softest. The z can be thought of as the energy fraction of the softest cluster at the merging step. The R_{cut} and z_{cut} are tunable parameters in the algorithm. If the clustering is either soft or wide-angled (i.e. if one of the above conditions is met), then the cluster with the smaller p_{T} is discarded. A cartoon of the pruning procedure is given in Figure 7.4.

Pruning was used to groom jets in CMS during Run 1 [90], before Soft Drop (see Section 7.2.3) was adopted. CMS pruned with parameters $R_{\text{cut}} = 0.5$, $z_{\text{cut}} = 0.1$: smaller R_{cut} and larger z_{cut}

[1]Recall that this is a sequential recombination algorithm with $p = 0$, and thus combines constituents i and j based solely on their angular separation in η and ϕ, $\frac{\Delta R_{ij}}{R}$, where $R = 1.0$ for ATLAS large-radius jets.

correspond to more aggressive pruning.

7.2.2 Modified Mass Drop Tagger

One of the groundbreaking publications which helped to establish jet substructure as a coherent phenomenological and experimental field was that of the *BDRS algorithm* [128], used to remove background jets (which will have more collinear splittings and soft radiation) in searches for the Higgs boson. An extension of this procedure, and a small modification thereafter, led to the *modified mass drop tagger* (mMDT) algorithm [118]. The creation of the mMDT tagger provided a nice solution to an issue with the calculation of the jet mass: namely, it provided a groomed jet which had a radiation pattern that gave an analytical mass calculation free of so-called *non-global logarithms* [68, 118]. This has nice theoretical consequences of making the final mass simpler to calculate analytically, as well as potentially reducing its associated theoretical uncertainties. For this reason mMDT, and generalisations such as Soft Drop, are favoured by the theory community as the *de-facto* grooming tools for experimentalists.

The mMDT procedure begins in a similar manner to pruning: first the constituents of the initial large-radius jet are clustered using the Cambridge-Aachen algorithm, producing a Cambridge-Aachen jet, labelled j. Then the last step of the clustering is undone, splitting j into two subjets, j_1 and j_2, with $p_{\mathrm{T}}(j_1) > p_{\mathrm{T}}(j_2)$. Following this, the mass drop and balance splitting conditions are tested:

$$\text{Mass Drop} : m(j_1) < \mu \times m(j_2) \tag{7.10}$$

$$\text{Balance Splitting} : \frac{\min(p_{\mathrm{T}}(j_1), p_{\mathrm{T}}(j_2))}{p_{\mathrm{T}}(j_1, j_2)} > z_{\mathrm{cut}}. \tag{7.11}$$

If both of these conditions are passed, the algorithm immediately stops and the original jet, j, is kept. If either condition is failed, the softer j_2 is discarded and j_1 becomes the new j in the algorithm. The new j is then once again undone, and its constituents combined into a new Cambridge-Aachen jet. Again the last clustering step is undone, creating a new j_1, and the conditions repeated. For the studies described in this Chapter, the mMDT is used in 'grooming mode': μ is set to an infinitely large number to ensure the mass drop condition is always satisfied.

Once this procedure has been applied to the clustering history of the jet, the remaining constituents are reclustered using the Cambridge-Aachen algorithm, with a radius parameter $R = 1.0$, and a subjet radius parameter of $R_{\text{filt}} = \min(0.3, \frac{\Delta R(j_1,j_2)}{2})$. Any constituents which fall outside of the three highest p_{T} subjets are discarded.

7.2.3 Soft Drop

Soft Drop [119] is an extension of the modified mass drop tagger by modifying the original z_{cut} criterion to include an explicit angular distance dependence on the jet. We shall first briefly discuss the original 'vanilla' Soft Drop before highlighting more recent developments in Soft Drop grooming.

7.2.3.1 'Vanilla' Soft Drop

The original idea of Soft Drop is to simply extend the requirement on the cluster energy fraction, z. The algorithm proceeds by reclustering the jet constituents into a Cambridge-Aachen jet of size R (as always, $R = 1.0$ in these studies). The first step of clustering is undone, giving the two j_1 and j_2 branches just as in the mMDT case, and then if the condition:

$$\frac{\min(p_{\text{T}}(j_1), p_{\text{T}}(j_2))}{p_{\text{T}}(j_1, j_2)} < z_{\text{cut}} \frac{\Delta R^{\beta}}{R^{\beta}} \tag{7.12}$$

holds the lowest p_{T} branch is removed and the process is repeated by going deeper into the clustering sequence of the jet i.e. by then undoing the last clustering step on the remaining branch which has survived this first Soft Drop condition. Otherwise, the algorithm immediately terminates and the jet is unchanged. The angular component of the jet is controlled by the value of β: $\beta = 0$ reduces to the original mMDT case which is then purely a cut on the soft radiation entering the jet. Increasing β permits increasing amounts of soft and wide-angle radiation to enter the jet, making the grooming less aggressive. Like the mMDT algorithm, Soft Drop also removes non-global logarithms in the analytical jet mass; unlike mMDT, there is no filtering with R_{filt} applied to the jets once the clustering history of the jet has been scanned. A cartoon of this 'vanilla' Soft Drop declustering is provided in Figure 7.5.

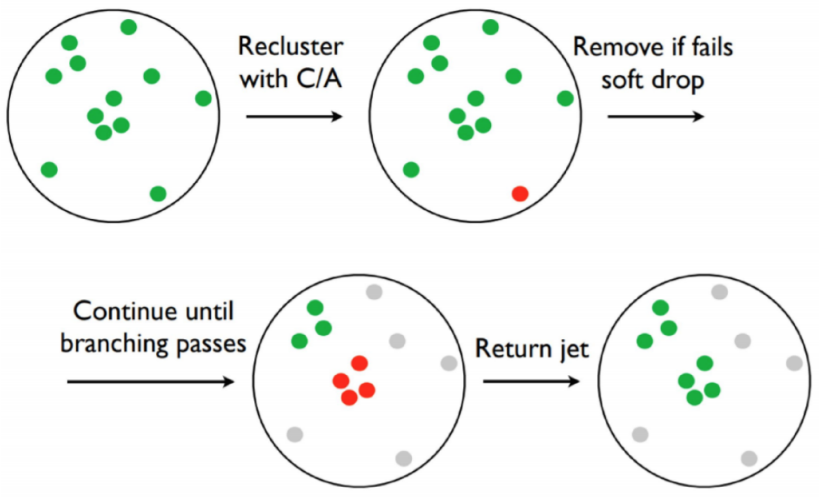


Figure 7.5: A cartoon of the 'vanilla' variant of Soft Drop, based on those in [90].

7.2.3.2 Recursive Soft Drop

Recursive Soft Drop [120] offers a simple extension to 'vanilla' Soft Drop. In the original Soft Drop grooming, the algorithm finishes an iteration once the soft drop condition is satisfied, and consequently it finds the first hard splitting in the clustering history, and takes this as the jet. With a little extra thought it is clear that this procedure is limited. Consider the top quark – a top quark decay can be captured in a large-radius jet and will typically have a three-prong decay pattern: one prong (or energy cluster) from the b-quark, and two from the hadronic decay of the W. Here, the three-prong decay may mean that standard Soft Drop is not sufficient to find the actual jet as it proceeds until only one hard-splitting is found. In recursive Soft Drop, the algorithm instead proceeds until N hard-splittings have been found i.e. the clustering is undone once, giving j_1 and j_2, but the first iteration of the algorithm is not complete until the Soft Drop condition has been satisfied N times, where the clustering history of the hardest subjet is followed at each step where the Soft Drop condition is tested. N is an additional tunable parameter, which takes integer values and has an 'infinity' setting: for the case $N = \infty$ the entire clustering history of the jet is explored before the first iteration of the algorithm terminates. The algorithm therefore progresses until $N + 1$ subjets have been found, or until all subjets are made up of single particles. A schematic depiction of recursive Soft Drop is provided in Figure 7.6. From the above description, it is clear that $N = 1$ recursive Soft Drop reduces to the original 'vanilla' variant, and $N = 0$ simply returns the original jet.

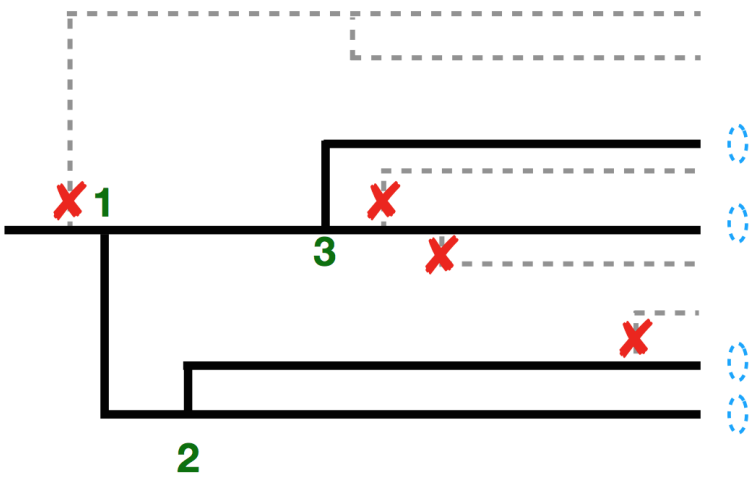


Figure 7.6: A cartoon of the recursive variant of Soft Drop. From [120].

7.2.3.3 Bottom-up Soft Drop

Another variant is *Bottom-up Soft Drop* (BUSD) [120] which has been developed as an alternative to $N = \infty$ recursive Soft Drop. Recursive Soft Drop provides a top-down grooming approach, where the jet is analysed and groomed starting from its complete clustering sequence, thought of as a tree of partonic branches. In contrast, the BUSD approach starts from the 'leaves' of the jet's clustering history, reclustering those leaves based on whether or not the Soft Drop condition is passed. It can be applied locally to groom the jet, or globally to groom the entire event.

The local variant of BUSD starts from the leaves of the clustering history of the jet, and proceeds by applying the Soft Drop condition to each pairwise combination of particles in the initial Cambridge-Aachen clustering history. At each step in the large-radius jet Cambridge-Aachen reclustering, a pseudojet is obtained from combining particles i and j with smallest distance $d_{ij} = \Delta R_{ij}/R$, with a transverse momentum given by:

$$p_{\mathrm{T}}^{\mathrm{pseudojet}} = \begin{cases} \max(p_{\mathrm{T},i}, p_{\mathrm{T},j}) & \text{if the Soft Drop condition fails,} \\ p_{\mathrm{T},i} + p_{\mathrm{T},j} & \text{otherwise.} \end{cases} \tag{7.13}$$

Since one begins from the full clustering history, BUSD effectively inverts the procedure of $N = \infty$ recursive Soft Drop. In the case of global BUSD, the algorithm starts from the entire clustering history of the event itself. In this way, global BUSD offers a way to groom at the

event-level and therefore does not require any specific jet definition. The groomed event will afterwards then contain a subset of the original particles from the ungroomed event, and these can then be combined using any jet algorithm. The global BUSD is an interesting first approach to event-level grooming, and is certainly something which could be applied to ATLAS and CMS data in the future. However, for the studies described in this Chapter, the local variant of BUSD will be studied only.

7.3 Jet Substructure Variables

The previous chapter was devoted to the definition, scale, and resolution of the jet mass, the simplest variable which provides information on the substructure of a jet itself. The field of *jet substructure* is rich and fascinating and in this Chapter we will use several additional jet substructure variables to construct metrics which assess the performance of different large-radius jet definitions in ATLAS. As has been previously discussed, when a particle has a transverse momentum comparable to (or greater than) its total mass, the angular separation of the daughter particles is small enough for the full decay to be captured within a suitably large-radius jet. Understanding the substructure of such jets allows us to differentiate those containing a top decay from those containing a hadronic W decay, and so forth. This discrimination can be achieved through analysing the mass of the jet, or, to take things a step further, the distribution of energy within that jet. With this in mind, let us turn to two incredibly successful approaches at scrutinising the energy distribution of particles captured within boosted jets: *N-subjettiness* [129] and *energy correlation functions* [130]. These quantities belong to a broad class of variables which quantify the overall *jet shape*.

7.3.1 N-subjettiness

Understanding the distribution of energy within a jet, and the direction of the energy distribution with respect to some geometrical axis which defines the jet direction, is very important when one is interested in differentiating large-radius jets which originate from top quarks from those which originate from bosons, or even background QCD processes. A jet originating from hadronic top decay will typically contain a 'three-prong' energy distribution i.e. three well-defined hard lobes of energy within the jet: in the dominant $t \to Wb$ decay, the b-quark will give rise to one energy

spike, and the $W \to qq$, where q generically labels a quark, will give rise to the additional two energy lobes. In contrast, for a jet originating from the hadronic decay of a W boson, there will clearly only be two energy lobes, or 'prongs'. A QCD jet would typically have its energy smeared across the $\eta - \phi$ space it occupies, and not display obvious prongy energy features. It is therefore reasonable to consider the 'pronginess' of the energy distribution of the jet as an assessment of the jet's substructure in differentiating jets originating from different heavy particles and QCD. One popular approach to quantifying this 'pronginess' is *N-subjettiness* [129].

The jet shape N-subjettiness variable, donated τ_N, is calculated by first identifying N candidate subjets from the jet of interest. This is achieved by taking the constituents of the original jet, and applying the k_t clustering algorithm, forcing it to return N subjets. Once the N subjets and their directions have been obtained, the τ_N can now be calculated according to:

$$\tau_N = \frac{1}{d_0} \sum_k p_{\mathrm{T},k} \min(\Delta R_{1,k}, \Delta R_{2,k}, ..., \Delta R_{N,k}), \tag{7.14}$$

where the summation runs over the k jet inputs (particles or energy clusters) with transverse momentum $p_{\mathrm{T},k}$ and $\Delta R_{i,k}$ is the angular separation in $\eta - \phi$ between the i^{th} candidate subjet and the k^{th} cluster. There is an additional factor out at the front, depending on some d_0:

$$d_0 = \sum_k p_{\mathrm{T},k} R_0, \tag{7.15}$$

where R_0 is the radius of the jet of interest, which is of course 1.0 in these studies. From the definition above, it is clear that jets with $\tau_N \approx 0$ have their radiation strongly aligned with the directions of the candidate subjets, and thus for such jets the N subjet hypothesis is a good one. These jets will therefore have N or fewer actual subjets. Conversely, jets with $\tau_N \gg 0$ have radiation distributed away from the subjet directions, and therefore such jets will in fact contain at least $N + 1$ actual subjets.

To utilise N-subjettiness in identifying boosted heavy particles, it is not sufficient to only require a small τ_N. Consider the case of identifying boosted top jets. Top jets should have three hard energy lobes, and hence three subjets. Therefore $N = 3$ is a natural choice, and indeed τ_3 is small for true top jets. However, clearly τ_3 will also be small for jets with only 2 subjets, such as those originating from W bosons. However, we can still use τ_2 usefully here:

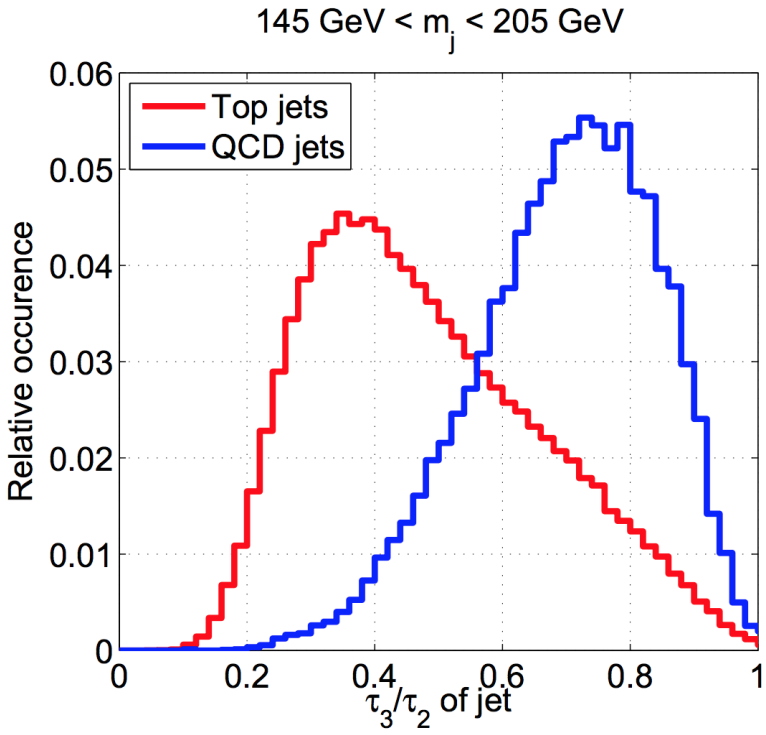


Figure 7.7: The τ_{32} distribution for a truth-level sample of top jets and QCD jets. The top jets have a smaller average τ_{32} as expected from their strongly three-prong energy structure. From [129].

this variable will remain small for true W boson jets but will be larger for top jets, therefore a powerful variable to discriminate 3-prong from 2-prong jets is in fact the ratio:

$$\tau_{32} = \frac{\tau_3}{\tau_2}. \tag{7.16}$$

Now, we have a variable τ_{32} which will be smaller for top jets than for W boson jets, or other jets with 2 energy lobes/prongs/subjets. Similarly the ratio $\tau_{21} = \tau_2/\tau_1$ will be smaller for jets with two-prong structure than one-prong structure, and as such is a good variable for discriminating jets originating from bosons from jets originating from background QCD processes. The separation power of the τ_{32} for a truth-level simulation of top and QCD jets is illustrated in Figure 7.7.

One disadvantage to using N-subjettiness is the need to rely on the direction, and hence axis, of the jet itself. It is possible to quantify energy distributions within a jet without relying on the jet axis by using a second class of substructure variables, referred to as energy correlation functions.

7.3.2 Energy Correlation Functions

The N-point *energy correlation function* (ECF) [130], e_N^β, is defined:

$$e_N^\beta = \sum_{i_1,i_2,\ldots,i_N \in J} \prod_{a=1}^{N} p_{\mathrm{T},i_a} \Big(\prod_{b}^{N-1} \prod_{c=b+1}^{N} \Delta R_{i_a} \Delta R_{i_b} \Big)^\beta, \tag{7.17}$$

where the summation is over all particles within some system J, which can be the entire event or a single jet, and is a function of the transverse momenta of each constituent and the angular separation in $\eta - \phi$ between each pairwise combination of constituents. The β parameter controls the amount of angular information which enters the ECF (in a similar manner to the β parameter used in the Soft Drop condition): $\beta = 0$ removes all angular information in the ECF.

In this way, the two- and three-point energy correlations functions are defined:

$$e_2^\beta = \frac{1}{p_{\mathrm{T}}^2} \sum_{i<j \in J} p_{\mathrm{T},i} p_{\mathrm{T},j} \Delta R_{ij}^\beta, \tag{7.18}$$

$$e_3^\beta = \frac{1}{p_{\mathrm{T}}^3} \sum_{i<j<k \in J} p_{\mathrm{T},i} p_{\mathrm{T},j} p_{\mathrm{T},k} \Delta R_{ij}^\beta \Delta R_{ik}^\beta \Delta R_{jk}^\beta, \tag{7.19}$$

with additional transverse momentum normalisations added at the front of each ECF. A feature of this system of equations is that $(N+1)$-point correlates tend to zero if there are only N actual particles in the system J. In the context of jets, this would mean that a jet with only N well-defined subjets will have a value of e_{N+1}^β significantly smaller than e_N^β. In this way, it can be seen immediately that ratios of energy correlation functions behave much like ratios of N-subjettiness, but without requiring an initial N candidate subjet hypothesis with defined jet axes.

Previous ATLAS studies have established that a variable which is very effective at identifying jets which two hard energy lobes, i.e. those originating from bosons, is the so-called $D_2^{\beta=1}$ variable [107], defined:

$$D_2^{\beta=1} = \frac{e_3^{\beta=1}}{(e_2^{\beta=1})^3}. \tag{7.20}$$

The choice of $\beta = 1$ is used throughout the energy correlation definitions considered in ATLAS. We shall return to this variable in the context of devising metrics for the jet inputs and grooming scan shortly.

7.4 The Jet Inputs and Grooming Scan

Having described the various new jet inputs, jet groomers, and jet substructure variables, we are now in a position to outline the *jet inputs and grooming scan* which has been covered in several public results [123, 131], with a further publication to follow later in 2019. This is an ambitious project which seeks to understand if their is a more robust large-radius jet definition which could be adopted in ATLAS, utilising the very latest jet inputs and grooming technology. The selection is based on important metrics of pileup mitigation, behaviour of the jet mass scale and resolution, and the performance of the jets in conventional W boson and top taggers [132]. We shall first describe the initial ATLAS Run 2 grooming scan study, which studied jet definition based on calorimeter information only, before briefly discussing future and ongoing work towards a second Run 2 scan. This second scan is even more ambitious and considers jet definitions where additional tracking information is used.

The scan uses variables that are directly affected by the jet energy and jet mass calibration. Since the initial combination of inputs and groomers resulted in almost 500 different candidate jet collections, it was impossible to derive dedicated JES and JMS calibrations for all such jet collections within any reasonable timescale. Consequently, a two-step approach is adopted: a set of sensible metrics are devised in order to build up a shortlist of about 5 'winning' jet collections from the initial $\mathcal{O}(500)$. The shortlistsed jets are then calibrated at the JES and JMS, and, since this study is purely Monte Carlo-based, there is no need for an *in situ* JES or JMS to be determined. Once these shortlisted jets have been calibrated, their performance is re-examined in order to confirm that the overall jet performance in the various metrics is not modified by the additional calibration. Three Monte Carlo samples of jets are considered: two are signals, where a $W' \to WZ$ sample is filtered to provide a source of $W-$jets, and has top jets sourced from a $Z' \to t\bar{t}$ signal; the final is a QCD background from a standard sample of Pythia 8 QCD dijets. These Monte Carlo samples have already been described in Chapter 6, and their details are there summarised in Table 6.2.

7.4.1 First Run 2 Inputs and Grooming Scan

Following the application of constituent subtraction and pileup mitigation techniques to the inputs for each jet, the anti-k_t algorithm is used to cluster the jet constituents into large-radius jets of size $R = 1.0$. In addition, jet grooming is then applied to remove additional energy flow from pileup and the underlying event which could still contaminate the large jet. In the initial scan modified Mass Drop, 'vanilla' Soft Drop, trimming, and pruning are used to groom; jet reclustering is also used in order to assess its power as a groomer. The full suite of groomers and grooming parameters are summarised in Table 7.1. These different jet groomers are applied to unmodified `LCTopo` inputs, as well as new inputs based on `CSLCTopo`, `SKLCTopo`, `CSSKLCTopo`, `VorSpreadLCTopo`, `VorSuppLCTopo`, and `VorSuppSKLCTopo`.

Grooming Algorithm	Name	Parameters Tested
Modified Mass Drop Tagger	mMDT	$(z_{\text{cut}}, R_{\text{filt}}) \in [0.05, 0.10, 0.15, 0.20] \times [3]$
Soft Drop	SD	$(z_{\text{cut}}, \beta) \in [0.05, 0.10, 0.15, 0.20, 0.15] \times [0.0, 0.5, 1.0, 1.5, 2.0]$
Trimming	Trimmed	$(f_{\text{cut}}, R_{\text{sub}}) \in [5, 7, 9, 11, 13, 15]$ % $\times [0.1, 0.2, 0.3]$
Pruning	Pruned	$(z_{\text{cut}}, R_{\text{cut}}) \in [0.10, 0.15, 0.20] \times [0.12, 0.25, 0.50]$
Reclustered	RC	$(R_{\text{small}}, f_{\text{cut}}) \in [0.2, 0.3, 0.4] \times [5]$ %

Table 7.1: A different jet groomers, and grooming parameteters, considered in the initial jet inputs and grooming scan.

In order to ensure a fair comparison between each of these different (uncalibrated) jet definitions (recall that we are considering $\mathcal{O}(500)$ different collections in this scan), truth jets of $R = 1.0$ are reconstructed with the anti-k_t algorithm and are then geometrically matched in $\eta - \phi$ to the reconstruction-level jets. Two sample topologies are considered in these initial studies: first a QCD background topology, where the jet pair corresponding to the leading truth jet with $p_{\text{T}} > 300$ GeV, and within $\Delta R < 0.75$ of the corresponding reconstructed jet, is kept. The p_{T} requirements ensure the jet is sufficiently boosted. Following this selection, grooming is applied to both truth and reconstruction-level jets. The second topology is signal $W-$jet sample, where the truth jet originating from the W boson is kept. Such a jet is identified as the truth jet which is closest to the W boson in ΔR, and is also within $\Delta R < 0.75$ of the same particle. Once the truth jet which is to be kept has been identified, the same geometrical matching is employed to find the appropriate reconstructed jet. Again, the truth jet must have a transverse momentum above 300 GeV, to ensure that the W boson is fully-contained such that its substructure can be faithfully studied. To ensure that jets are reconstructed well

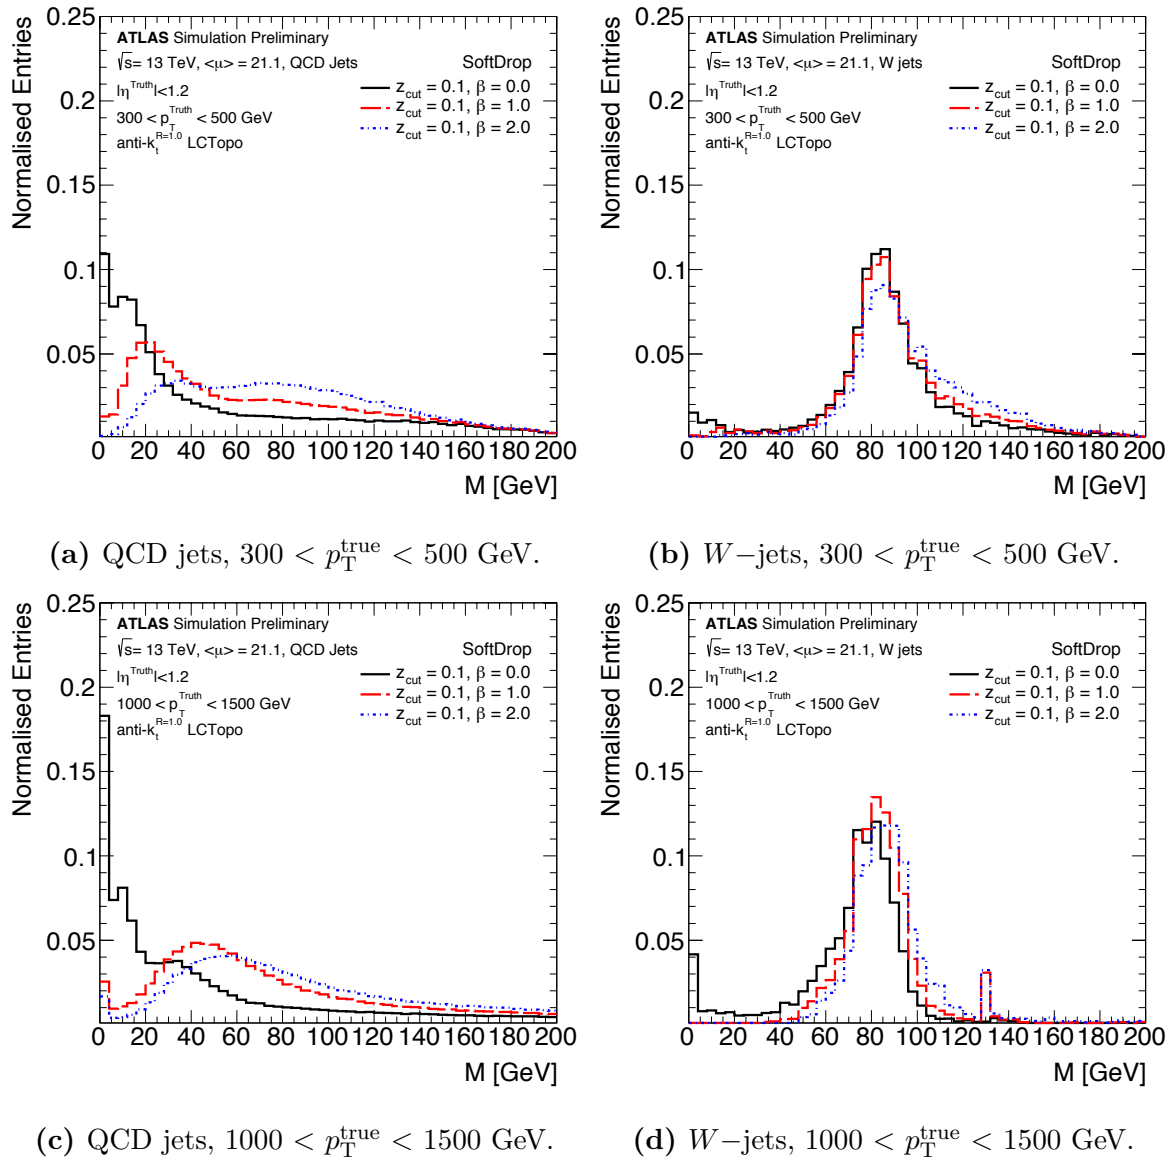
within the region of ATLAS where charged particle tracks may be reconstructed, and additional requirement that the truth jet $|\eta| < 1.2$ is also imposed.

7.4.1.1 Background Rejection Performance

Due to the practical impossibilities of calibrating hundreds of jet collections (where the initial inputs are at the LCW scale), the first sets of metrics considered have to be ones which allow for meaningful conclusions to be drawn in the absence of a calibration. The first such metric attempts to gauge which jet definitions are best at distinguishing the W-jet signal from a QCD jets background using the jet mass as a discriminating variable.

The jet mass (for these studies we only consider the calorimeter jet mass) should provide reasonable discrimination between QCD and $W-$jets: the former has a falling mass spectrum arising from gluon emissions and splittings in the parton shower, the latter has a clear resonant peak around the W boson pole. As previously discussed, the jet mass is highly dependent on pileup, and therefore should also be dependent on the grooming technique and strategy applied. Consequently so is the amount of QCD background expected within the phase-space containing the W boson mass peak. Figure 7.8 shows the jet mass distributions for signal and background, using $R = 1.0$ jets reconstructed from `LCTopo` inputs with Soft Drop grooming. A $z_{\text{cut}} = 0.1$ is applied, and β parameter values of 0.0, 1.0, and 2.0 are considered. The Soft Drop grooming shifts the QCD jets to lower masses, whilst keeping the W mass signal relatively fixed. The most aggressive grooming comes from $\beta = 0.0$, where no wide-angle information entering the jet is kept and therefore more of the mass of the jet originating from soft interactions is groomed away. In the case of $W-$jets, the mass originates from a clear hard splitting and thus the mass peak is preserved when the wide-angle information is removed from the jet. This behaviour is observed at both lower and higher $p_{\text{T}}^{\text{true}}$. Similar behaviour is observed when mMDT is used in the grooming.

As expected, smaller values of z_{cut} lead to a larger fraction of higher mass QCD jets, and therefore the background contamination around the W mass peak (hence the background efficiency, to be defined shortly) also increases. This is expected because decreasing the z_{cut} will result in a larger fraction of the softer components of the jet being kept during the grooming



(a) QCD jets, $300 < p_{\mathrm{T}}^{\mathrm{true}} < 500$ GeV.

(b) $W-$jets, $300 < p_{\mathrm{T}}^{\mathrm{true}} < 500$ GeV.

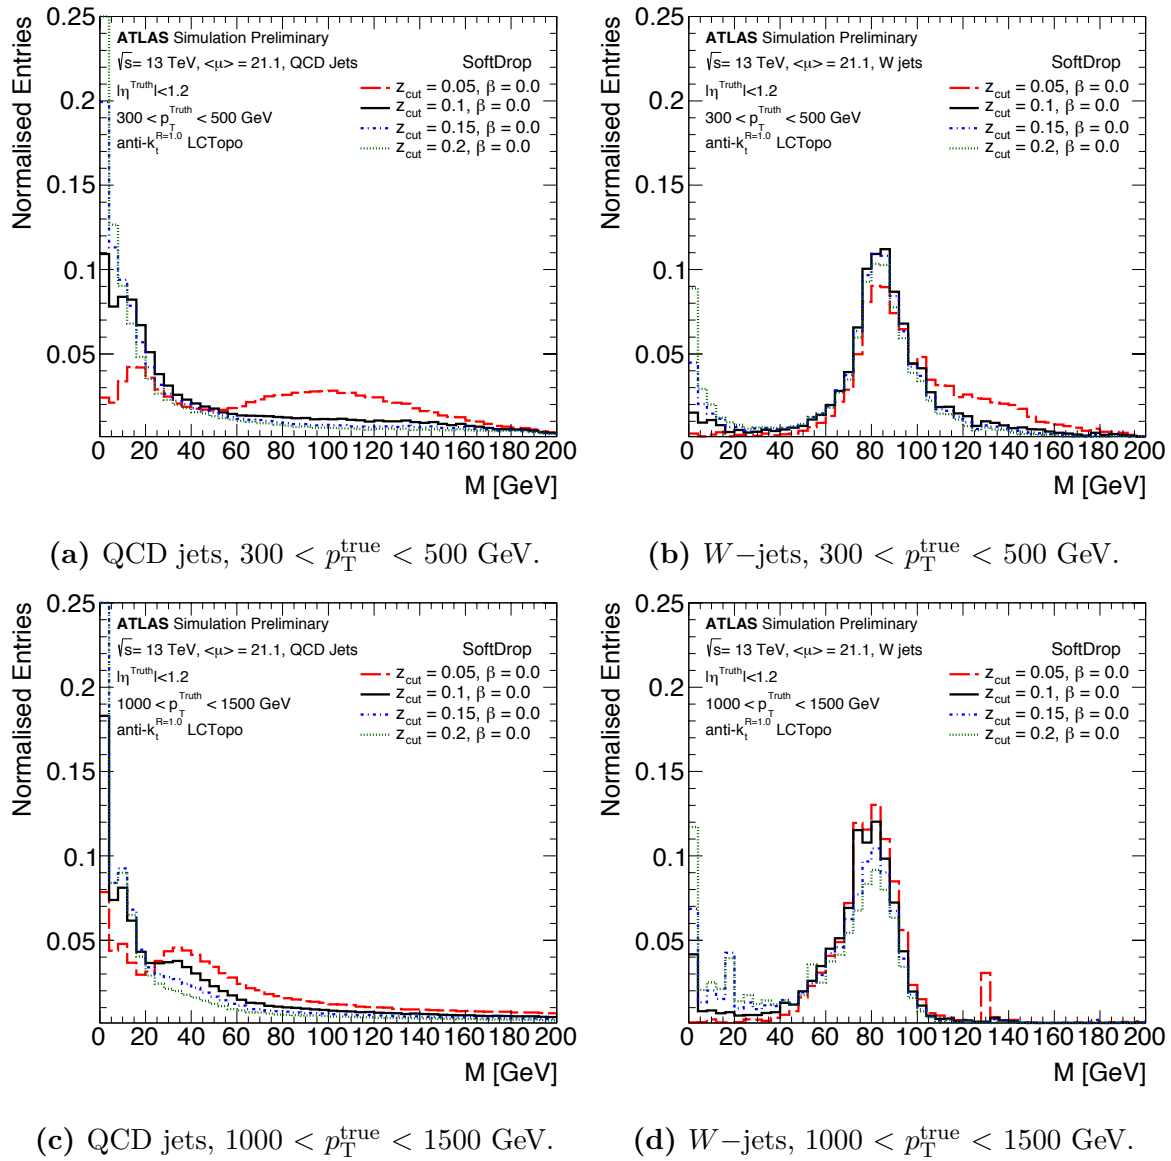
(c) QCD jets, $1000 < p_{\mathrm{T}}^{\mathrm{true}} < 1500$ GeV.

(d) $W-$jets, $1000 < p_{\mathrm{T}}^{\mathrm{true}} < 1500$ GeV.

Figure 7.8: The jet mass distributions for $W-$jet and QCD jet topologies, based on applying Soft Drop grooming to `LCTopo` inputs, for both high and low ranges of $p_{\mathrm{T}}^{\mathrm{true}}$. The z_{cut} is fixed at 0.1, and β-values of 0.0, 1.0, and 2.0 are considered.

process. These conclusions are illustrated in Figure 7.9.

In order to quantify the background rejection power of each jet collection, the smallest mass window around the W peak is constructed such that 68 % of the signal jets are contained within the window mass range. The total fraction of QCD jets which enter this same mass window is determined, and the background efficiency for this fixed 68 % signal acceptance is defined as $\epsilon_{\mathrm{Bkg}}^{68\ \%} = N_{\mathrm{QCD}}/(N_{\mathrm{QCD}} + N_{W-\mathrm{jets}})$. An example of such mass window is illustrated for the standard `AntiKt10LCTopoTrimmedPtFrac5SmallR20` jets in Figure 7.10. For different jet $p_{\mathrm{T}}^{\mathrm{true}}$ ranges, the background efficiency is calculated for every jet collection; lower efficiencies indicate



(a) QCD jets, $300 < p_{\mathrm{T}}^{\mathrm{true}} < 500$ GeV.

(b) $W-$jets, $300 < p_{\mathrm{T}}^{\mathrm{true}} < 500$ GeV.

(c) QCD jets, $1000 < p_{\mathrm{T}}^{\mathrm{true}} < 1500$ GeV.

(d) $W-$jets, $1000 < p_{\mathrm{T}}^{\mathrm{true}} < 1500$ GeV.

Figure 7.9: The jet mass distributions for $W-$jet and QCD jet topologies, based on applying Soft Drop grooming to `LCTopo` inputs, for both high and low ranges of $p_{\mathrm{T}}^{\mathrm{true}}$. The β is fixed at 0.0, and z_{cut}-values of 0.05, 0.1, 0.15, and 0.2 are considered.

that less QCD is contaminating the region of phase-space around the W mass peak. These background efficiencies, for the different jet collections, are summarised in Figure 7.11 for two different transverse momentum ranges.

The large 'heatmap' of background efficiency, although complicated, comes with a few crucial take-home messages. At lower $p_{\mathrm{T}}^{\mathrm{true}}$, background contamination can clearly be reduced by modifying the inputs to include a form of pileup mitigation i.e. constituent subtraction. At larger jet $p_{\mathrm{T}}^{\mathrm{true}}$, this distinction between different jets inputs is less obvious. Now, the amount of background contamination is more dependent on the jet groomer itself; harsher jet grooming

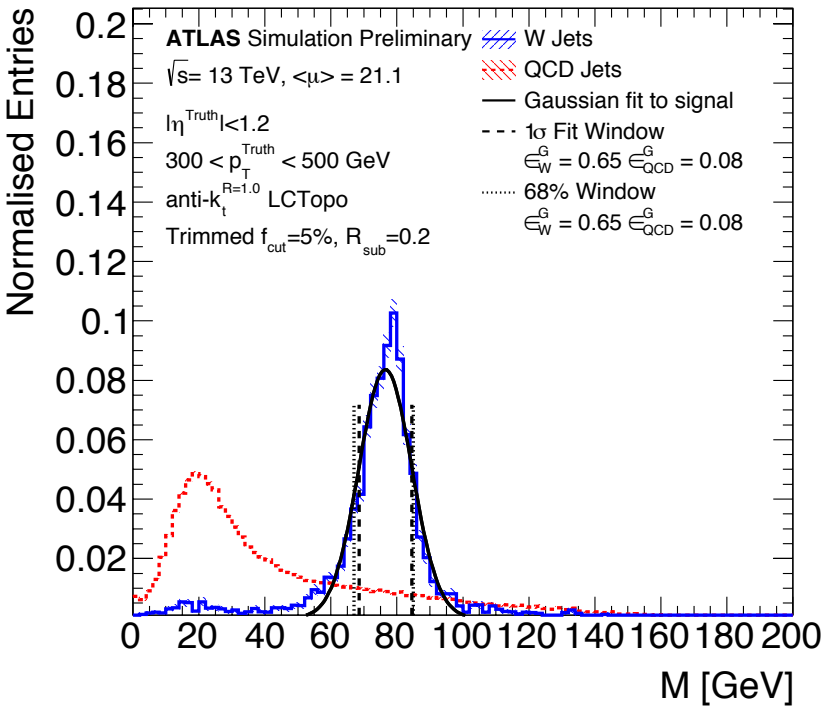


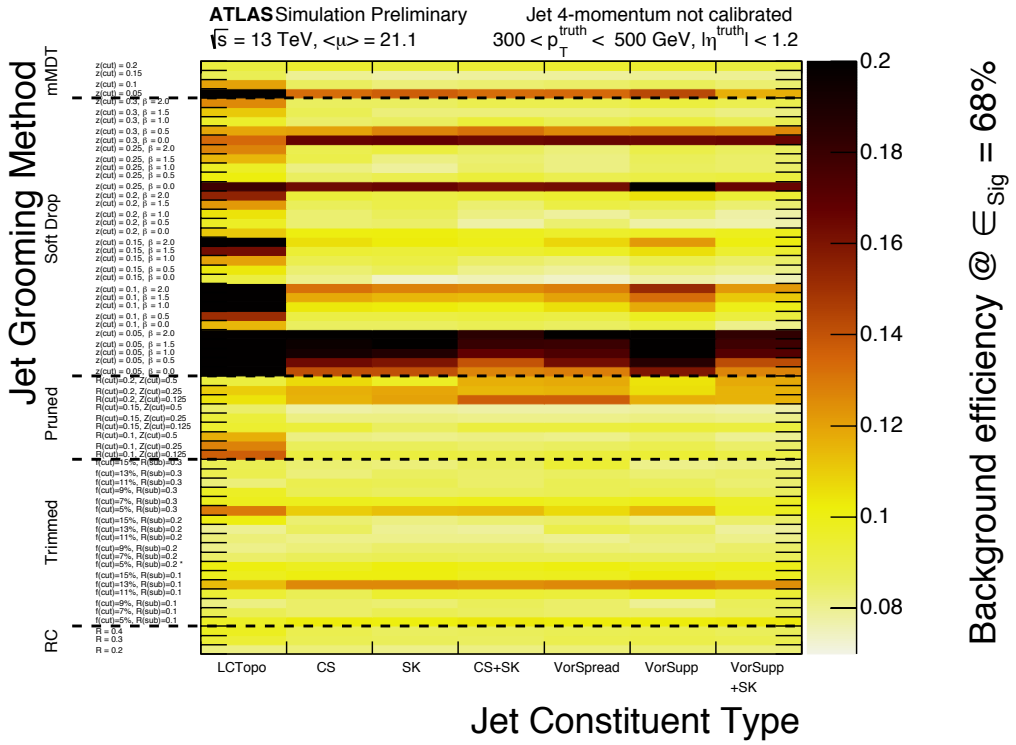
Figure 7.10: The jet mass distribution for `AntiKt10LCTopoTrimmedPtFrac5SmallR20` jets, with the QCD background in red and the $W-$jet signal in blue. A Gaussian fit (black) is applied to the W signal, and the dotted lines indicate the mass window containing 68 % of the desired signal.

pushes the QCD mass distribution to lower values and therefore increases the separating power between the QCD background and the signal jets.

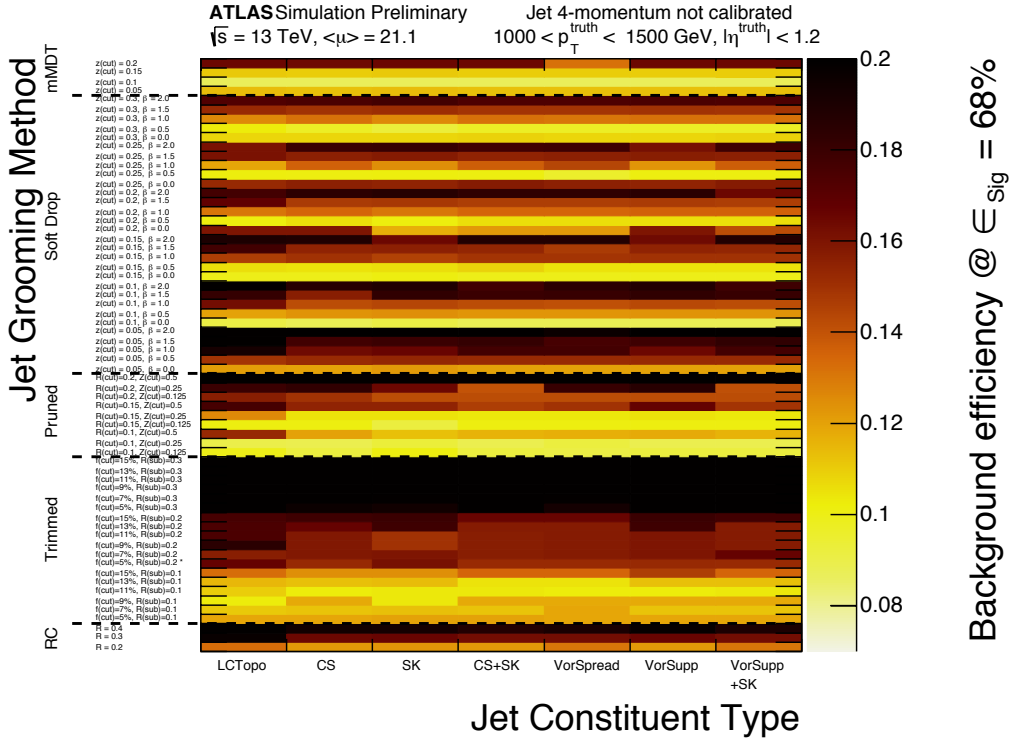
7.4.1.2 Jet Mass Pileup Dependence

The jet mass itself depends on both the energy scale of the jet and the angular spread of that energy. Consequently, the effects of wide-angle contributions from pileup are even more apparent in the jet mass than the jet energy alone. This should result in large changes in the mass distribution with respect to the pileup, and therefore the performance of both the jet mass scale, and its resolution, with respect to the in-time pileup (quantified by the number of primary vertices per event, N_{PV}) makes for an important metric. The variation in the jet mass distribution for different ranges of N_{PV} is illustrated in Figure 7.12.

In reality, it is not viable to study the pileup dependence of the jet mass scale (and its resolution) because the jets are not yet calibrated at the energy and mass scale, and the pileup dependence would be modified following any such calibration. In order to study pileup



(a) Background efficiency, $300 < p_{\mathrm{T}}^{\mathrm{true}} < 500$ GeV.



(b) Background efficiency, $1000 < p_{\mathrm{T}}^{\mathrm{true}} < 1500$ GeV.

Figure 7.11: Heatmaps showing the efficiency for QCD jets to fall into the smallest $W-$jet mass window containing 68 % of the signal, both at low and higher $p_{\mathrm{T}}^{\mathrm{true}}$.

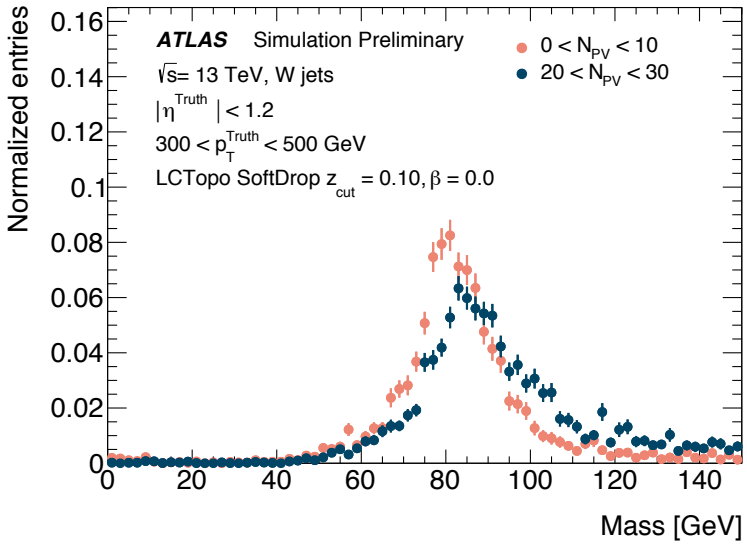
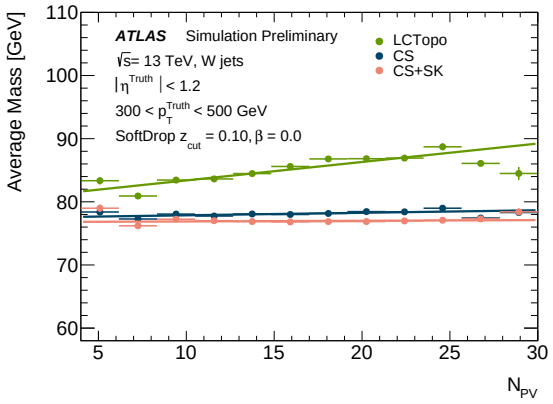
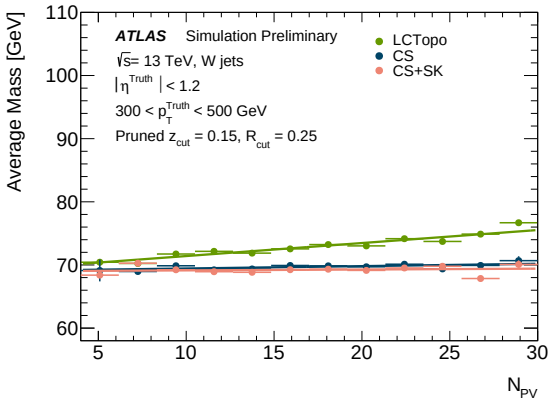


Figure 7.12: The jet mass of a reconstructed W boson, in two exclusive bins of N_{PV}, for Soft Drop jets made from unmodified `LCTopo` inputs.

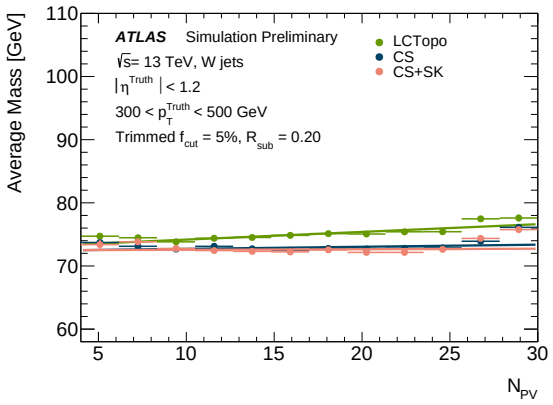
dependence in a meaningful (uncalibrated) manner, instead the reconstructed W boson peak is fitted iteratively with a Gaussian, in different bins of N_{PV}, and the central value and 68 % interquantile range mass window are used as values of the average W boson reconstructed mass and width respectively. Figures 7.13 and 7.14 show the average reconstructed W mass and its width, σ_W, for several different jet collections which were observed to provide some of the lowest background efficiencies. The jets considered lie in the lower $p_{\mathrm{T}}^{\mathrm{true}}$ range between 300 GeV and 500 GeV, which is more susceptible to pileup. Most grooming algorithms which use unmodified `LCTopo` inputs observe an increase in the W jet mass as N_{PV} increases, due to the increase in uncorrelated energy flow entering the jet expected from conditions of higher in-time pileup. Perhaps unsurprisingly, the standard ATLAS trimming configuration ($R_{\mathrm{sub}} = 0.2$, $f_{\mathrm{cut}} = 0.05$) outperforms both the Soft Drop and pruning algorithms, as this trimming algorithm was



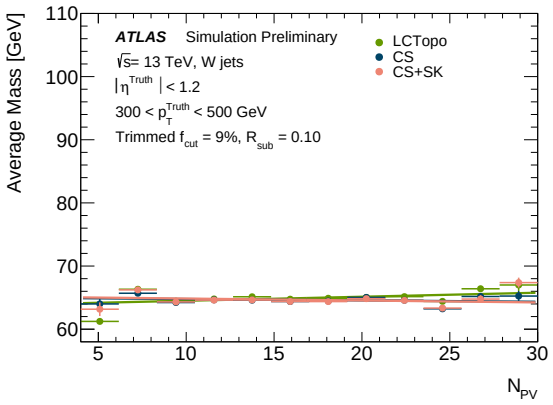
(a) Soft Drop, $z_{\text{cut}} = 0.10$, $\beta = 0.0$.



(b) Pruned, $z_{\text{cut}} = 0.15$, $R_{\text{cut}} = 0.25$.



(c) Standard trimming, $f_{\text{cut}} = 5.0$ %, $R_{\text{sub}} = 0.20$.

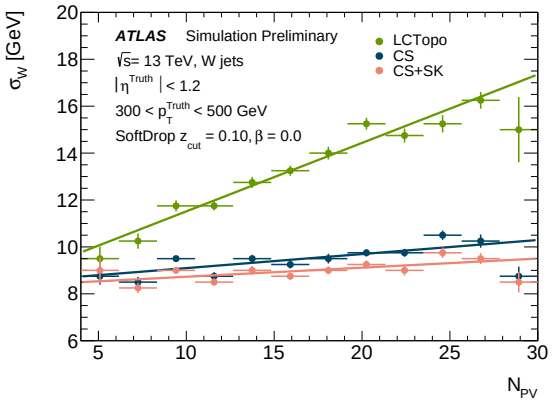


(d) Trimming, $f_{\text{cut}} = 9.0$ %, $R_{\text{sub}} = 0.10$.

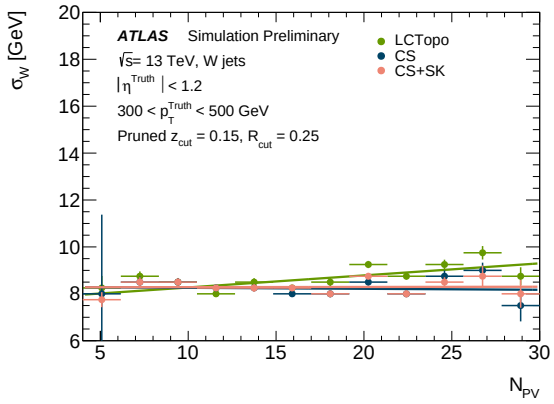
Figure 7.13: The pileup dependence on the average $W-$jet mass, for different jet inputs and grooming algorithms, in the range $300 < p_{\text{T}}^{\text{true}} < 500$ GeV. The values of the different slopes from the linear fits are provided in Figure 7.15

optimised based on Run 1 studies of the pileup dependence of different jet definitions [90] and has been reported in numerous cases to be particularly pileup robust. In all cases, independent of the grooming procedure adopted, the application of the constituent subtraction significantly reduces the pileup dependence. Similar conclusions may also be drawn for the width of the W mass peak.

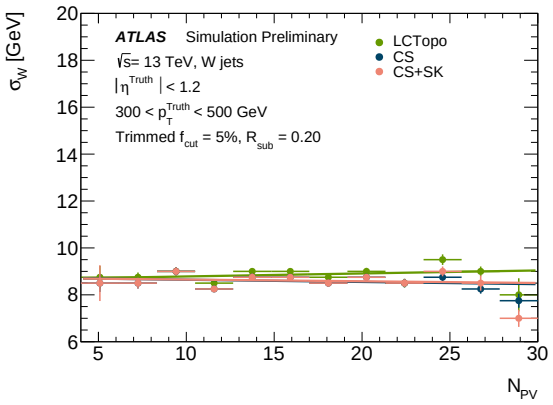
The values of the linear fit to the pileup dependencies of the W mass and its width are used as proxies for the two metrics quantifying the pileup dependence of the jet mass scale and resolution respectively. These values are summarised for all jet collections in the heatmaps shown in Figures 7.15 and 7.16, which attempt to classify the pileup dependence of the jet mass and jet mass resolution respectively. It is now evident that the pileup dependence of both mass and resolution is significantly reduced when the `LCTopo` inputs are modified using constituent subtraction, or Voronoi subtraction, with SoftKiller. The `CSSKLCTopo` and `VorSuppSKLCTopo`



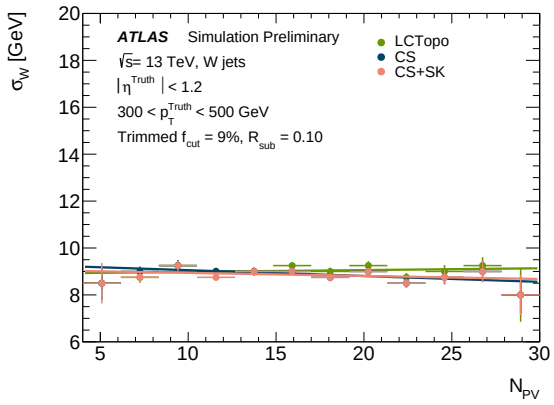
(a) Soft Drop, $z_{\text{cut}} = 0.10$, $\beta = 0.0$.



(b) Pruned, $z_{\text{cut}} = 0.15$, $R_{\text{cut}} = 0.25$.



(c) Standard trimming, $f_{\text{cut}} = 5.0$ %, $R_{\text{sub}} = 0.20$.



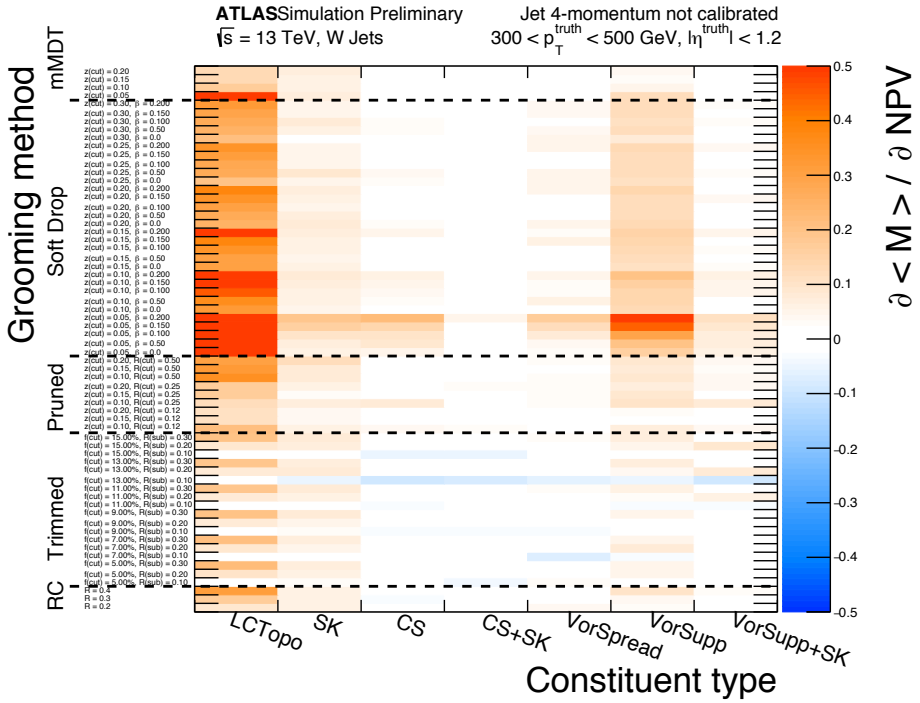
(d) Trimming, $f_{\text{cut}} = 9.0$ %, $R_{\text{sub}} = 0.10$.

Figure 7.14: The pileup dependence on the $W-$jet mass width, for different jet inputs and grooming algorithms, in the range $300 < p_{\text{T}}^{\text{true}} < 500$ GeV. The values of the different slopes from the linear fits are provided in Figure 7.16.

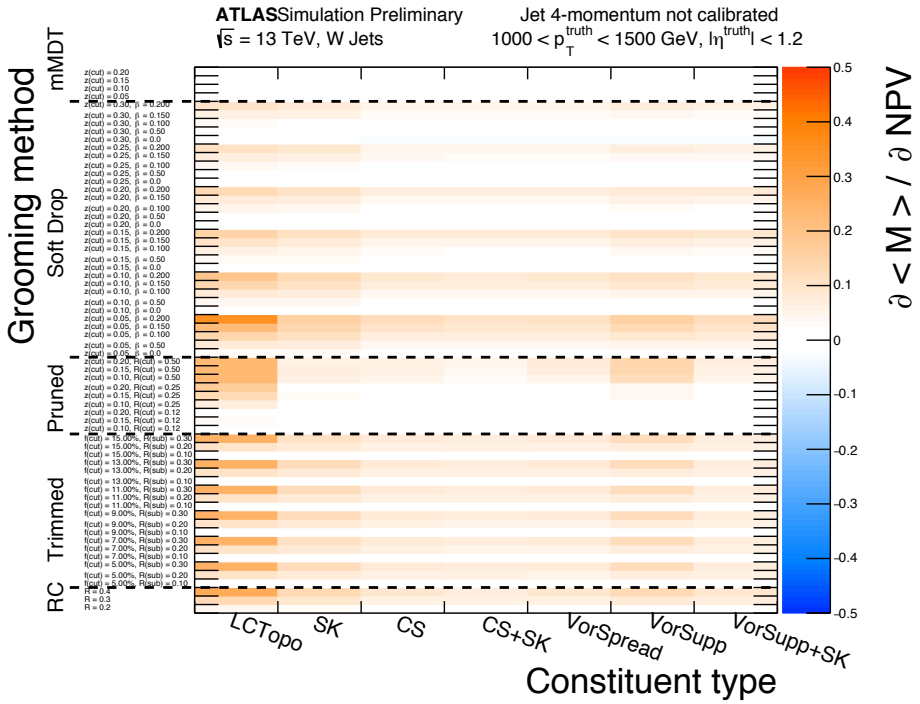
inputs are particularly pileup robust. Between these two inputs, the `CSSKLCTopo` is the most pileup robust in the lowest $p_{\text{T}}^{\text{true}}$ region of phase-space considered in these studies. Between the different jet groomers, the variations in pileup robustness are less strong. That said, perhaps the most obvious general trend is that grooming techniques which are more aggressive tend to provide jets with a smaller pileup dependence, as would be expected. This is illustrated, for example, in the case of Soft Drop when contrasting $\beta = 2.0$ and $\beta = 1.0$ grooming with the harsher $\beta = 0.0$ case.

7.4.1.3 D_2 Pileup Dependence

The D_2 variable is particularly powerful as it characterises the distribution of the hard energy between all constituents within the jet; as such, it has significant utility in assessing the '2-

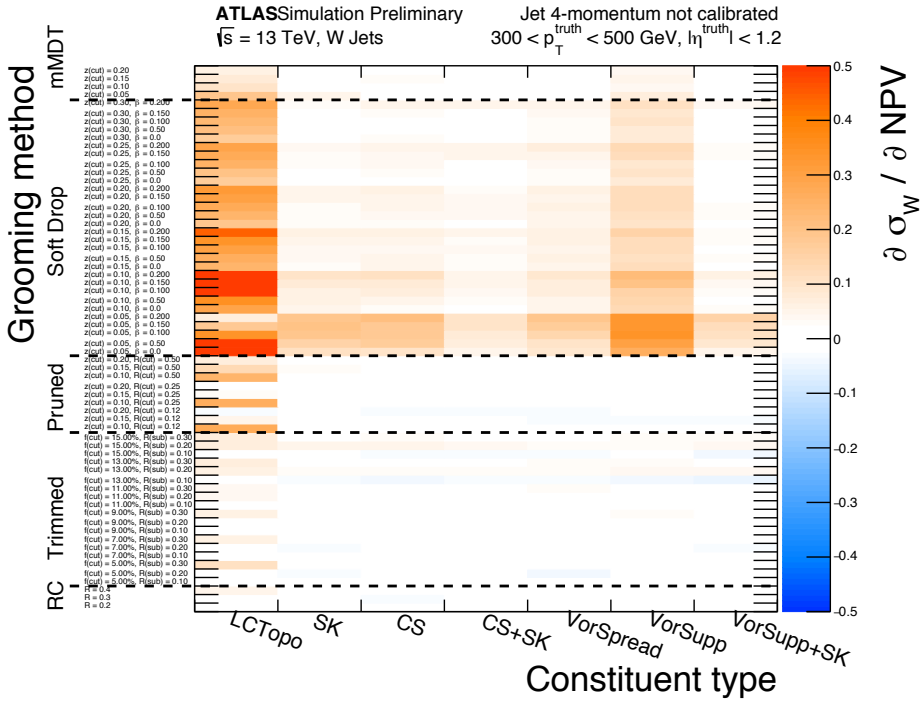


(a) Jet mass N_{PV}-dependence, $300 < p_{\mathrm{T}}^{\mathrm{true}} < 500$ GeV.

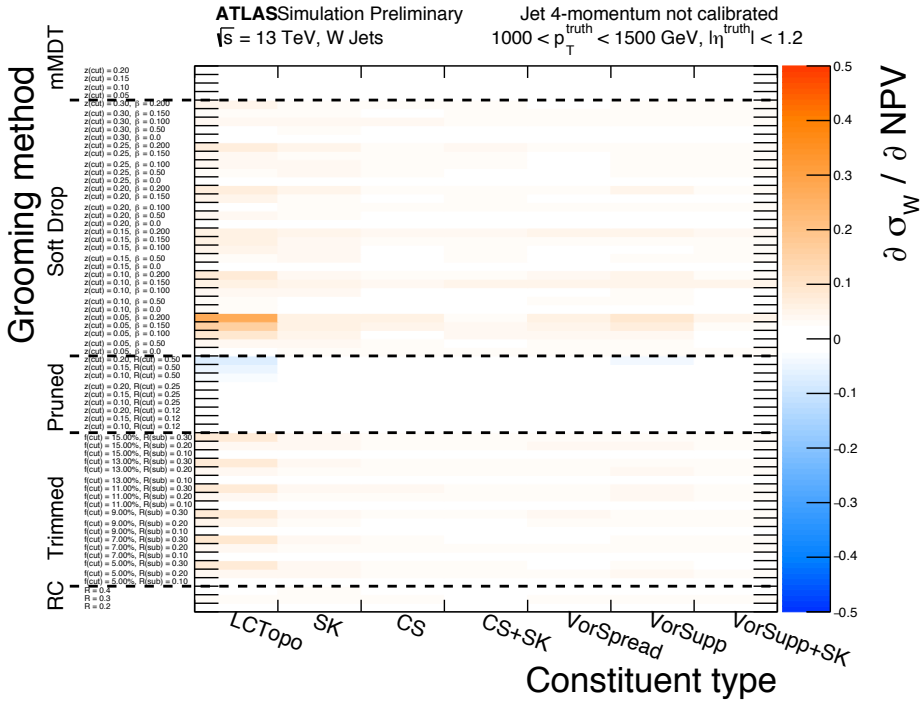


(b) Jet mass N_{PV}-dependence, $1000 < p_{\mathrm{T}}^{\mathrm{true}} < 1500$ GeV.

Figure 7.15: Heatmaps quantifying the slope of the linear fit to the average $W-$jet mass$-N_{\mathrm{PV}}$ dependence.



(a) Jet mass width N_{PV}-dependence, $300 < p_{\mathrm{T}}^{\mathrm{true}} < 500$ GeV.



(b) Jet mass width N_{PV}-dependence, $1000 < p_{\mathrm{T}}^{\mathrm{true}} < 1500$ GeV.

Figure 7.16: Heatmaps quantifying the slope of the linear fit to the $W-$jet mass with N_{PV} dependence.

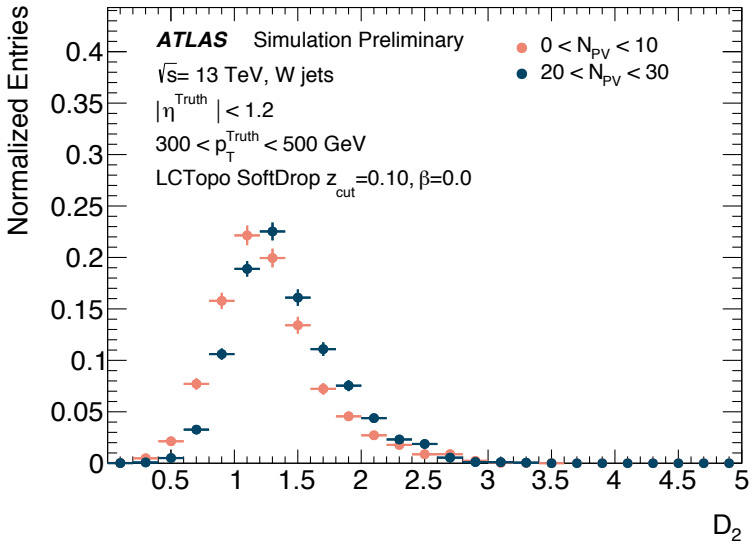


Figure 7.17: The D_2 distributions for `LCTopo` Soft Drop jets, in exclusive bins of N_{PV}. The change in this jet substructure variable with respect to the average N_{PV} value can be clearly seen.

pronginess' of the energy flow within the reconstructed jet. This variable is ideal for constructing W/Z boson taggers, but it is sensitive to pileup contributions which can affect the energy flow within the hadronic jet. Example pileup dependence is illustrated at low $p_{\mathrm{T}}^{\mathrm{true}}$ in Figure 7.17, and is most clearly observed when standard `LCTopo` inputs are used. Here, Soft Drop grooming is applied to the jets. As was concluded in the case of the jet mass pileup dependence, using constituent subtraction on the jet inputs improves the robustness of the substructure variable with respect to N_{PV}. An important application of D_2 is in exotics searches for New Physics (for example, see [124]), and therefore understanding its robustness against pileup for various different jet collections is an obvious metric to consider.

The approach considered in order to measure the D_2 pileup dependence is illustrated in the

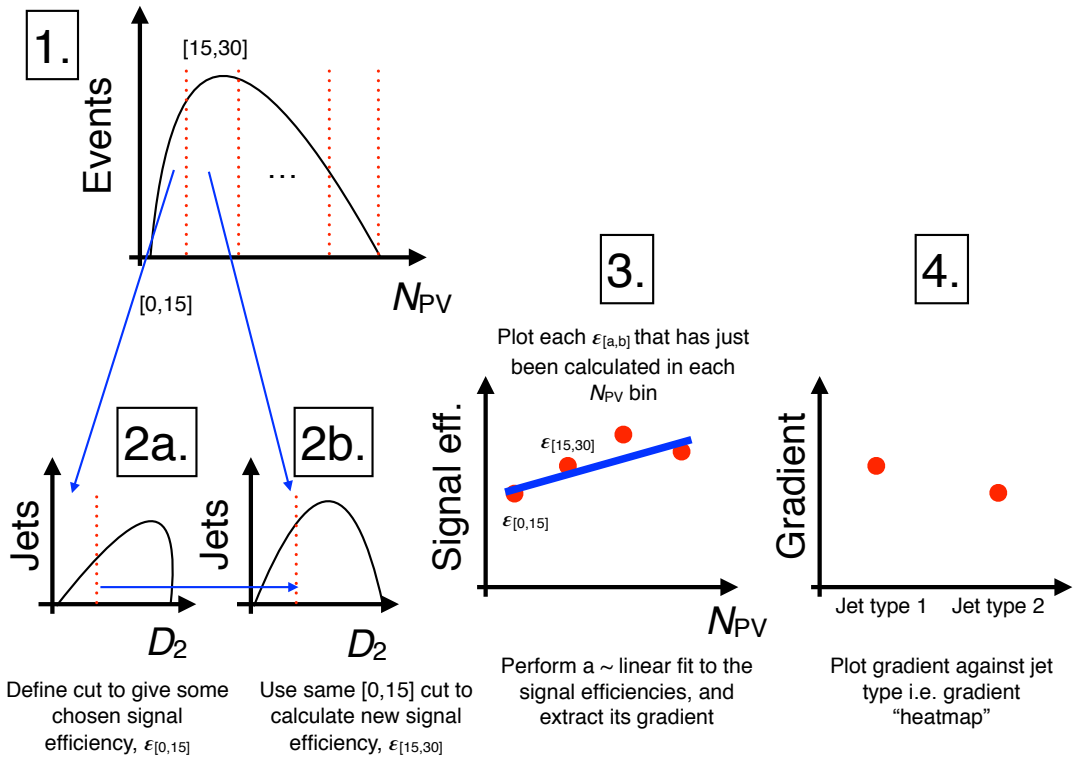
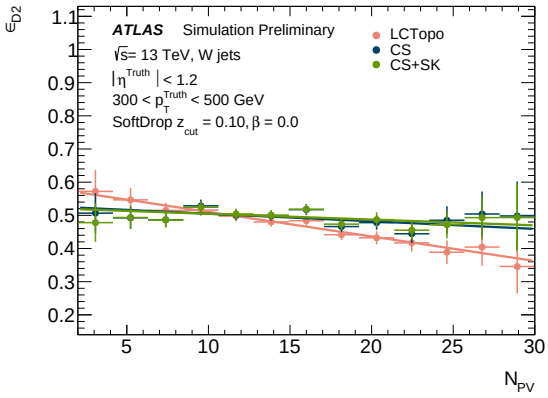
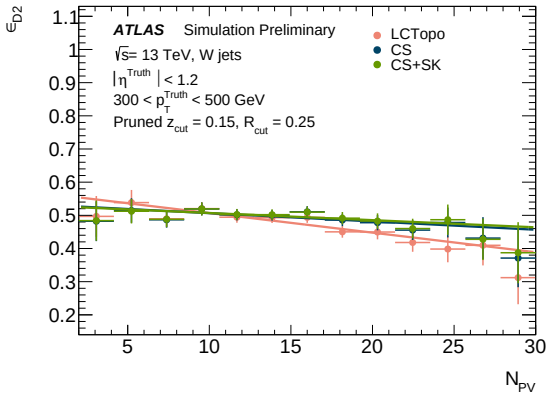


Figure 7.18: A schematic demonstrating how the pileup dependence on a jet substructure variable, such as D_2, can be assessed in a meaningful way. The distribution of all events is first split into bins of N_{PV}, and the substructure variable is studied in each bin.

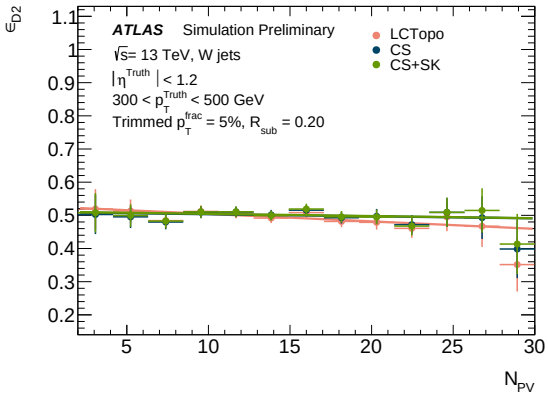
schematic in Figure 7.18. More generally, this sort of approach can be used to characterise the pileup dependence in the tails of numerous jet substructure variables. First, the N_{PV} distribution across all events is split into several bins, i.e. [0, 15], [15, 30], [30, 45], etc. Following this, the D_2 distribution is plotted for the events (such as $W-$jet signal events) in the lowest N_{PV} bin, and a cut on D_2 is defined that provides some desired signal efficiency. In these studies, a 50 % signal efficiency is required for jets with a mass above 50 GeV, using an initial N_{PV} bin in the range of $0 - 15$. The same cut is then used on the D_2 distribution described by the events in the next N_{PV} bin, and a new signal efficiency from this cut is calculated. This process is iterated through all N_{PV} bins, such that the signal efficiency can then be plotted with respect to the average N_{PV}. The gradient of the fit to this distribution then quantifies the D_2 pileup dependence on the jets considered in the sample of interest. This approach is extended to all jet definitions considered in the scan, and efficiencies for both the $W-$jets and QCD topologies have been studied. Example linear fits to N_{PV} vs. D_2 signal efficiency plots for $W-$jets are shown in Figure 7.19 for a few different jet definitions, where `LCTopo`, constituent subtraction, and constituent subtraction plus SoftKiller inputs are compared in each plot. The complete



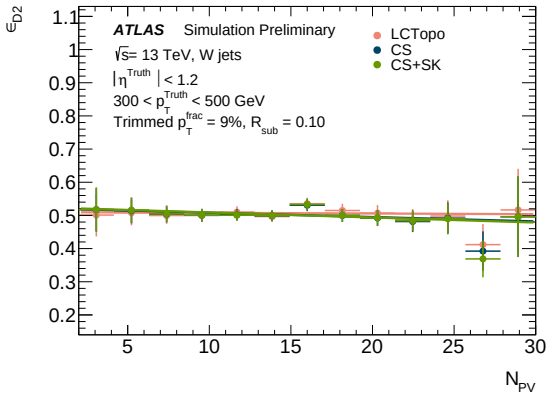
(a) Soft Drop, $z_{\text{cut}} = 0.10$, $\beta = 0.0$.



(b) Pruned, $z_{\text{cut}} = 0.15$, $R_{\text{cut}} = 0.25$.



(c) Standard trimming, $f_{\text{cut}} = 5.0$ %, $R_{\text{sub}} = 0.20$.

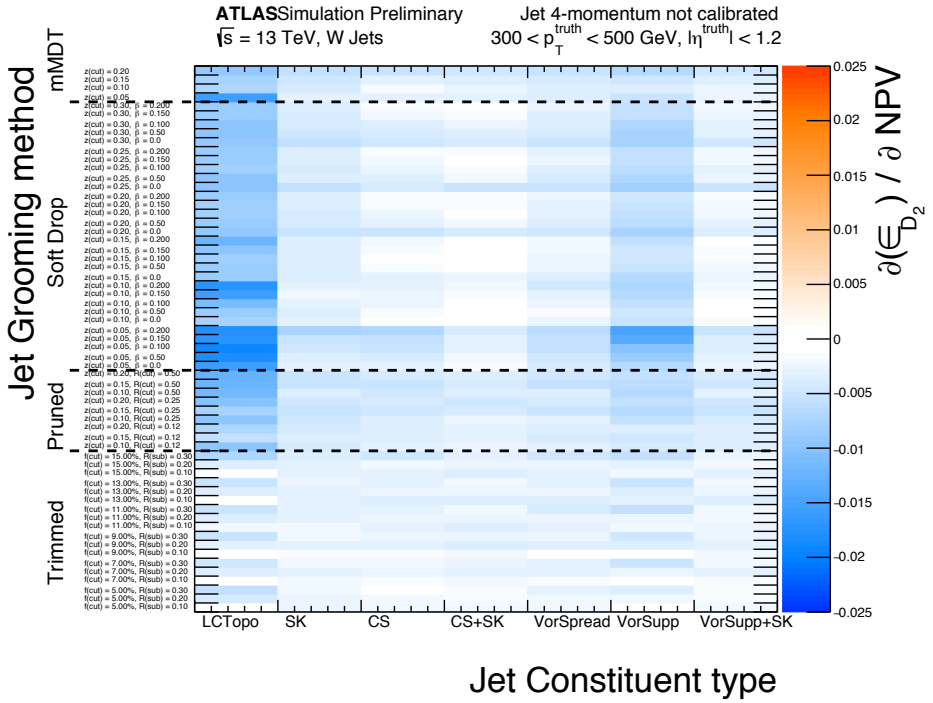


(d) Trimming, $f_{\text{cut}} = 9.0$ %, $R_{\text{sub}} = 0.10$.

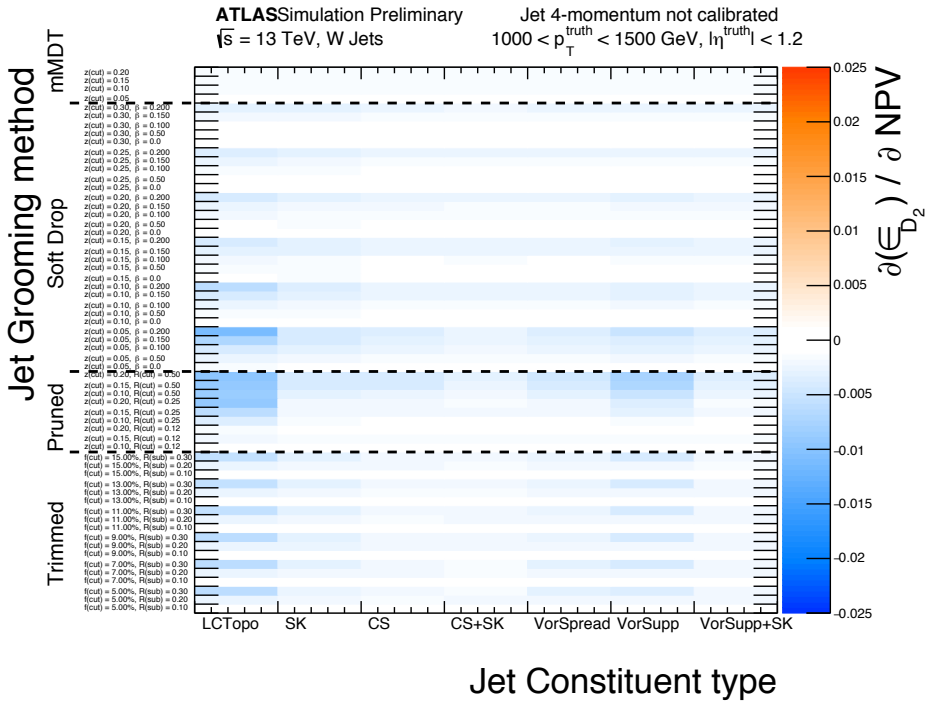
Figure 7.19: The pileup dependence on the $W-$jet D_2 efficiency, for different jet inputs and grooming algorithms, in the range $300 < p_{\text{T}}^{\text{true}} < 500$ GeV. The values of the different slopes from the linear fits are provided in Figure 7.20.

heatmaps, made from the linear fits for each jet collection, are summarised in Figures 7.20 and 7.21 for $W-$jets and QCD topologies respectively.

The conclusions drawn from studying the D_2 efficiency metric mirror those of the jet mass and jet mass resolution pileup dependence metric. Again the inputs based on `CSSKLCTopo` and `VorSuppSKLCTopo` are the most robust against pileup, with `CSSKLCTopo` outperforming the Voronoi-suppressed inputs slightly. As before, the trimmed algorithm is seen to be stable against pileup, even without using constituent subtraction or Voronoi-based techniques at the jet inputs level, and in particular smaller values of R_{sub} are more robust against pileup than larger values. For pruning, it is clear that larger values of R_{cut} and smaller values of z_{cut} reduce pileup robustness since they collectively reduce the amount of grooming applied to the initial jet. These effects are more apparent in the $W-$jets topology, which has the two-prong distribution

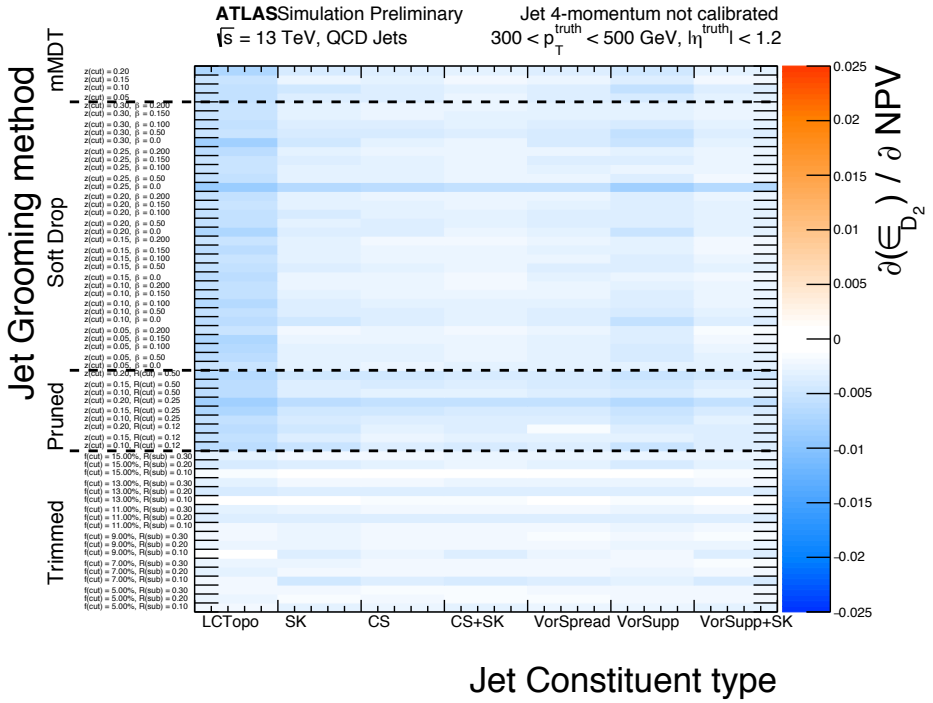


(a) $W-$jet D_2 efficiency N_{PV}-dependence, $300 < p_{\mathrm{T}}^{\mathrm{true}} < 500$ GeV.

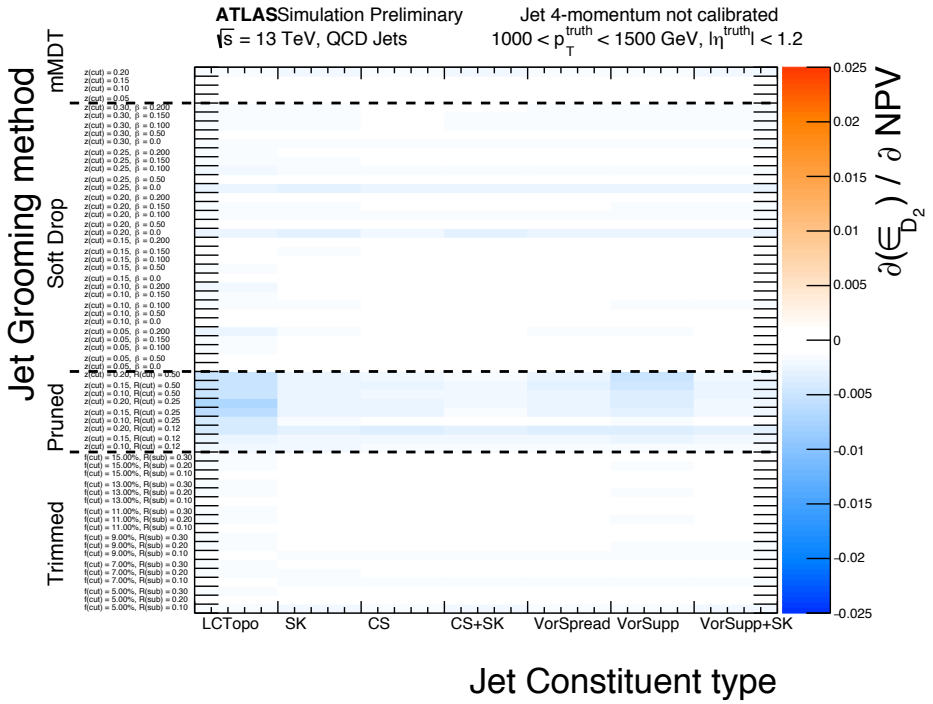


(b) $W-$jet D_2 efficiency N_{PV}-dependence, $1000 < p_{\mathrm{T}}^{\mathrm{true}} < 1500$ GeV.

Figure 7.20: Heatmaps quantifying the slope of the linear fit to the $W-$jet D_2 efficiency with N_{PV} dependence.



(a) QCD jet D_2 efficiency N_{PV}-dependence, $300 < p_{\mathrm{T}}^{\mathrm{true}} < 500$ GeV.



(b) QCD jet D_2 efficiency N_{PV}-dependence, $1000 < p_{\mathrm{T}}^{\mathrm{true}} < 1500$ GeV.

Figure 7.21: Heatmaps quantifying the slope of the linear fit to the QCD jet D_2 efficiency with N_{PV} dependence.

of energy flow which the D_2 variable is designed to identify.

7.4.1.4 Final Shortlist

The major conclusion from the initial set of studies in the inputs and grooming scan is that there is not a single obvious combination of jet input and groomer which is superior to all the others, when the results of all metrics are taken into account. However, there are clearly important trends which have been observed across these different metrics. Based on these general trends, a shortlist of promising new jet collections can be proposed. The finalist jet collections are summarised in Table 7.2, based on studies which consider a QCD jets background and a W-jets signal.

Constituent Type	Grooming Algorithm	Parametere Choice
`CSSKLCTopo`	Soft Drop	$z_{\text{cut}} = 0.1$, $\beta = 0$
`CSSKLCTopo`	Pruning	$z_{\text{cut}} = 0.15$, $R_{\text{cut}} = 0.25$
`CSSKLCTopo`	Trimming	$R_{\text{sub}} = 0.1$, $f_{\text{cut}} = 9$ %
`LCTopo`	Trimming	$R_{\text{sub}} = 0.2$, $f_{\text{cut}} = 5$ %
`EMTopo`	Reclustering	$R_{\text{small}} = 0.4$, $f_{\text{cut}} = 5$ %

Table 7.2: The shortlisted jet inputs and groomers for the initial scan.

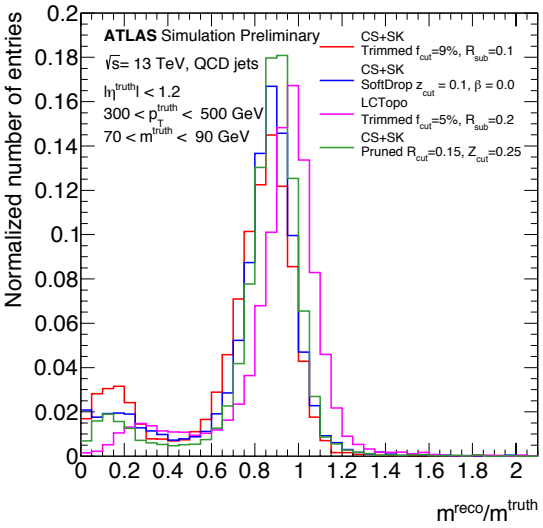
The `CSSKLCTopo` and `VorSuppSKLCTopo` inputs are the most pileup resilient, with `CSSK` slightly outperforming `VorSuppSK`. For this reason, the `CSSK` inputs are kept in the final shortlist. The trimmed `LCTopo` inputs which have been used as the standard large-radius jet definition in ATLAS also performed very well, and are kept for further study, together with a reclustered jet definition which uses $R = R_{\text{small}} = 0.4$ `EMTopo` jets as inputs. At high $p_{\text{T}}^{\text{true}}$, it was observed that the Soft Drop algorithm performs best at separating W signal jets from QCD based on selecting the jet mass alone. To ensure a fair comparison, it was decided that at least one example of each grooming algorithm should be selected, and so the best performing trimming and pruning configurations also make the shortlist. No mMDT grooming is included, but instead a Soft Drop configuration with $\beta = 0.0$ (which reduces to mMDT in that limit) is considered in the final shortlist. The second trimming configuration, which uses `CSSKLCTopo` inputs, was found to provide excellent pileup robustness, although the degree of grooming it applied to a given jet by is particularly aggressive. This aggressive trimming configuration further motivates the standard `AntiKt10LCTopoTrimmedPtFrac5SmallR20` being shortlisted so that the two algorithms may

be further compared.

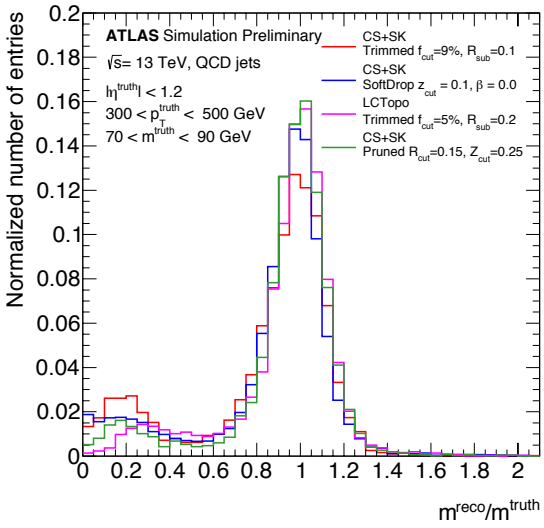
In order to study these finalist jet definitions in more detail, the JES and JMS corrections need to be derived and applied. The standard approach is adopted for each finalist jet collection: the Monte Carlo JES is first determined in bins of E^{true}, $p_{\text{T}}^{\text{true}}$, and η^{det} from a sample of QCD dijets, and applied to those same QCD dijets to correct the reconstruction-level jets to their matched truth-level jets in energy. Following this, those jets are then used as inputs to derive the JMS in bins of m^{true}, $p_{\text{T}}^{\text{true}}$, and η^{det}. The final JES + JMS calibration is then applied to both the QCD and $W-$jets samples considered in these studies. As previously discussed, it is expected that the closure in the jet mass response would be significantly improved after applying the JMS correction. This is illustrated in Figures 7.22 and 7.23 for regions of low- and high-$p_{\text{T}}^{\text{true}}$ respectively, selecting a mass around the W mass pole, for both the QCD and $W-$jets samples. Excellent closure in the mass response is recovered after the application of the JES + JMS calibration, derived for each shortlisted jet collection. The precise features in the jet mass response, such as the low response tails, vary between different groomers and inputs. This information is captured in the fractional jet mass resolution, calculated from the 68 % interquantile range of the jet mass response and the median response. The JMRs for the shortlisted jet definitions, as a function of $p_{\text{T}}^{\text{true}}$, are shown in Figure 7.24. The JMR behaviour is characteristic of the calorimeter jet mass: at higher transverse momentum the JMR degrades due to the Lorentz boost of the jet giving rise to an opening angle between the captured decay products which is comparable to the calorimeter granularity.

7.4.1.5 Tagging Performance Post-Calibration

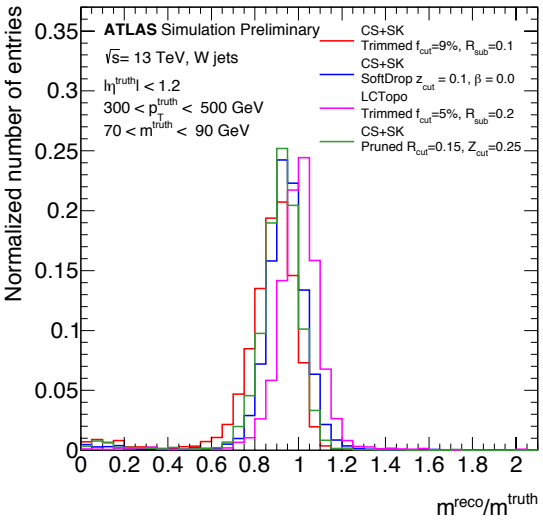
Following calibration, the shortlisted jet collections can now be tested as simple W boson taggers. Two such taggers are considered: one based on a simple cut placed on the W signal mass peak, and a second designed to mimic existing cut-based boson tagging at ATLAS [132], consisting of a cut on the signal boson mass peak and a requirement on the D_2 variable. Before summarising the performance on these taggers on the finalist jet collections, we must outline a systematic approach to studying tagging performance in particle physics. Assessing the tagging performance is essentially a counting experiment, where one computes the *tagging efficiency* and *background rejection* for a particular signal sample and background, based on a particular set of



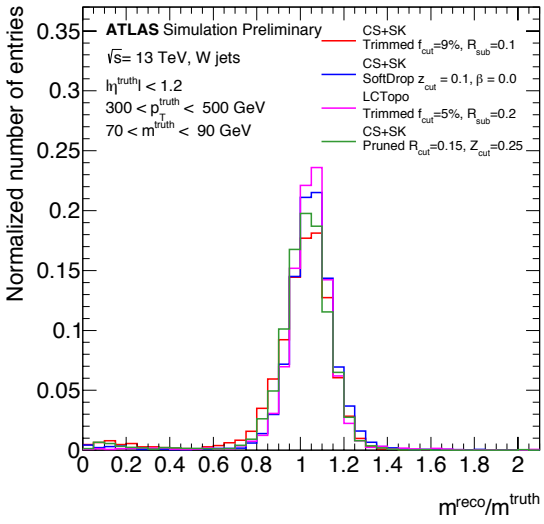
(a) JES-corrected QCD jet mass responses.



(b) JES+JMS-corrected QCD jet mass responses.

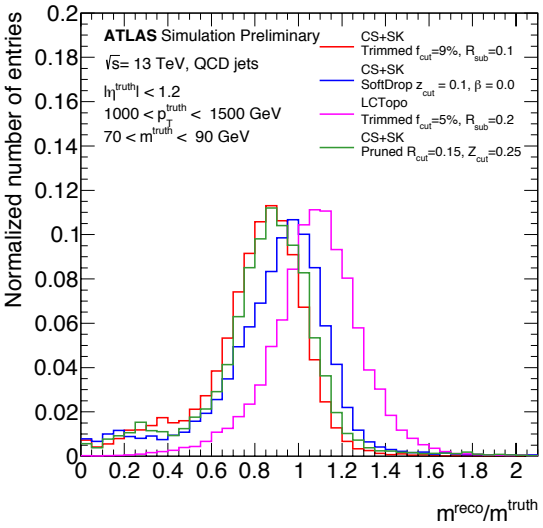


(c) JES-corrected $W-$jet mass responses.

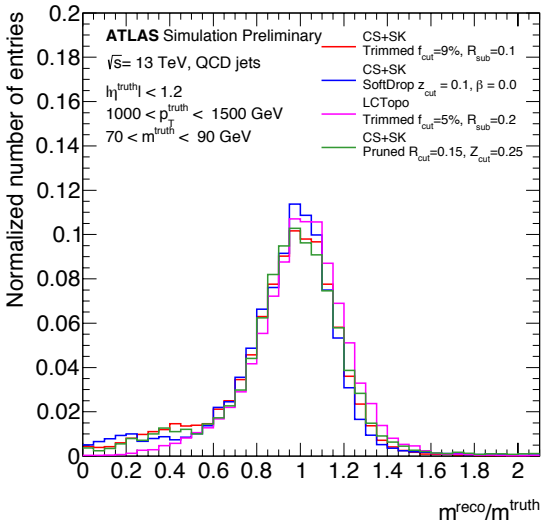


(d) JES+JMS-corrected $W-$jet mass responses.

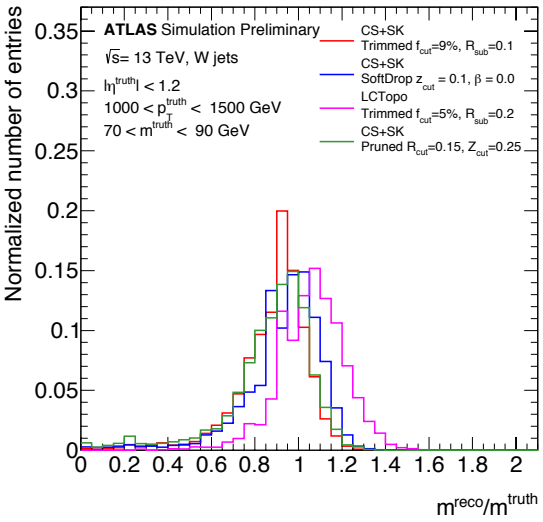
Figure 7.22: The JES and JES+JMS-corrected jet mass responses for the shortlisted jet collections, selecting events around the W mass and in the lower $p_{\mathrm{T}}^{\mathrm{true}}$ range $300 - 500$ GeV.



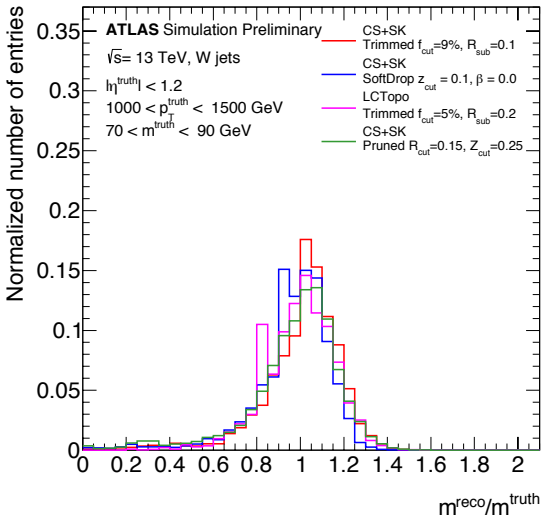
(a) JES-corrected QCD jet mass responses.



(b) JES+JMS-corrected QCD jet mass responses.

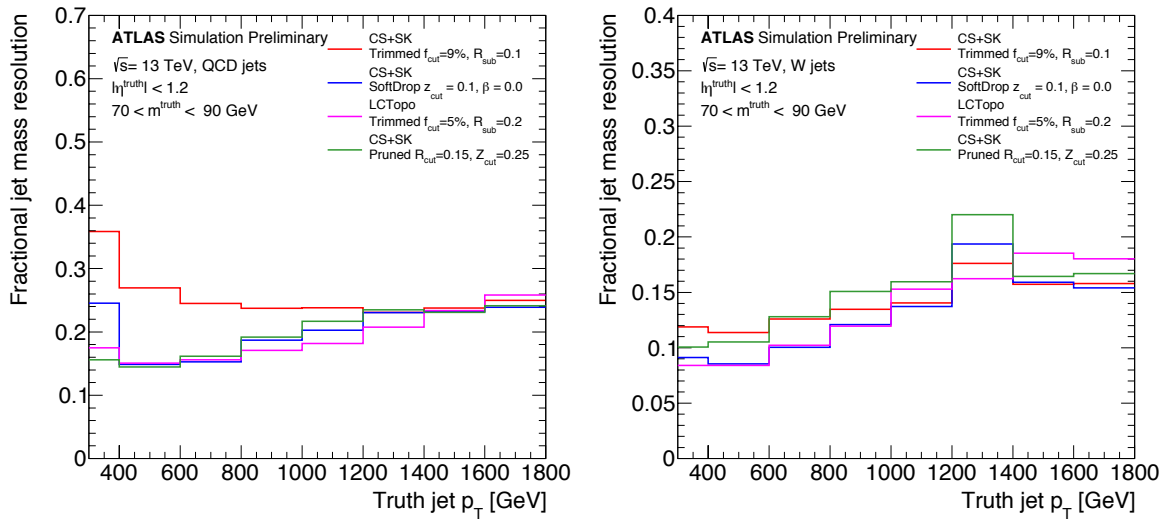


(c) JES-corrected $W-$jet mass responses.



(d) JES+JMS-corrected $W-$jet mass responses.

Figure 7.23: The JES and JES+JMS-corrected jet mass responses for the shortlisted jet collections, selecting events around the W mass and in the lower $p_{\mathrm{T}}^{\mathrm{true}}$ range $1000 - 1500$ GeV.



(a) QCD JMR, after JES+JMS corrections. **(b)** W JMR, after JES+JMS corrections.

Figure 7.24: Fractional jet mass resolutions for the shortlisted jet collections, after JES+JMS corrections have been applied, with respect to the truth-level jet transverse momentum. Both QCD and $W-$jet topologies are considered, and in both cases the JMR is quantified around the W mass pole.

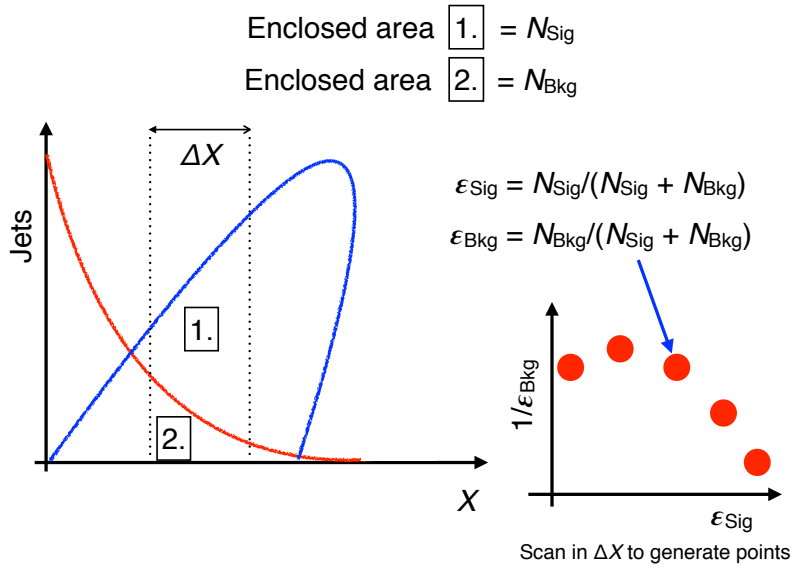


Figure 7.25: An illustration of how the tagging efficiency and background rejection factor are calculated from a tagger based on a generic cut (or set of cuts) on some general variable, X, using a window of size ΔX. The red line is the X-distribution for some undefined background, and the blue for some hypothetical signal. Varying the X window generates a set of background rejection and signal efficiency values. Generally X can be a single variable or a collection of variables, but here a single variable X is considered for clarity.

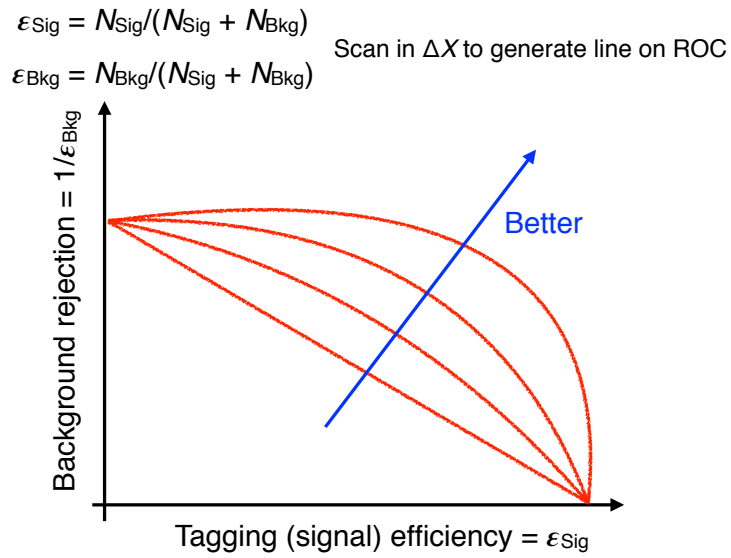


Figure 7.26: The typical form of a ROC curve, where each point is generated by considering the tagging efficiency and background rejection factor based on some variable X, with window ΔX.

cuts on a particular set of variables. In the cartoon provided in Figure 7.25 a variable X has been plotted for some signal (which has a large peak in some X-window) and a smoothly-falling background. For some interval in X, denoted ΔX, the total numbers of signal and background events are measured to be N_{Sig} and N_{Bkg} respectively. An algorithm which tags signal and rejects background will now have a tagging (or signal) efficiency,

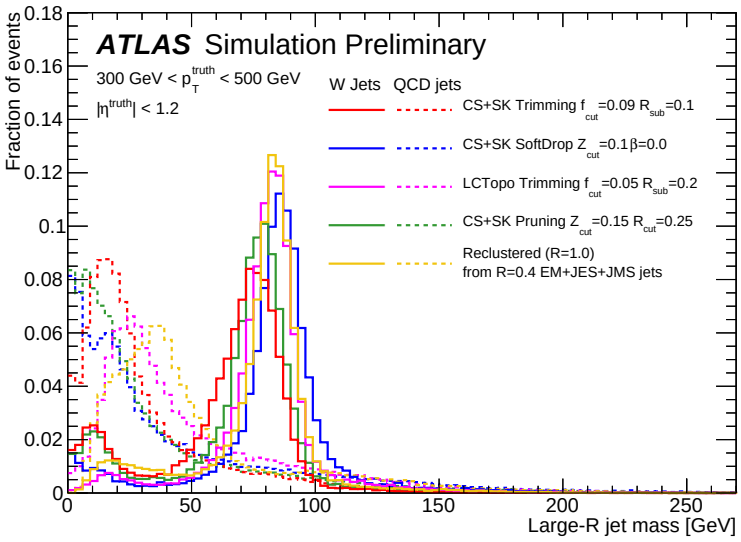
$$\epsilon_{\text{Sig}} = \frac{N_{\text{Sig}}}{N_{\text{Sig}} + N_{\text{Bkg}}}, \tag{7.21}$$

and background rejection factor,

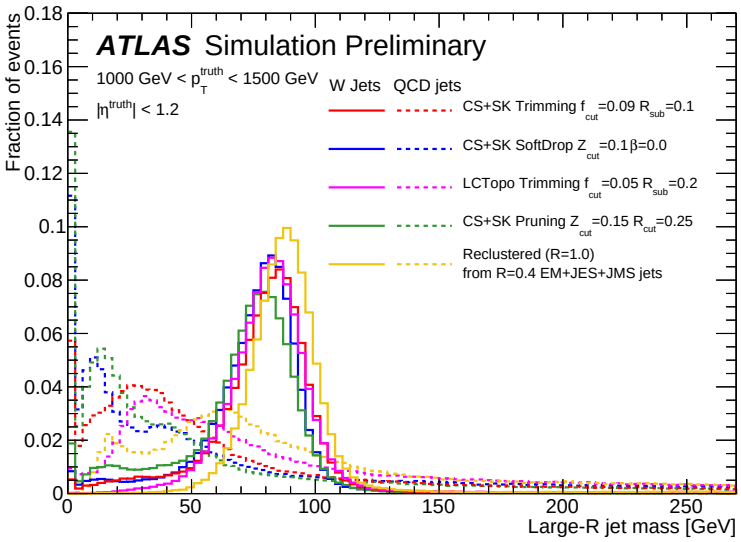
$$\frac{1}{\epsilon_{\text{Bkg}}} = \frac{N_{\text{Sig}} + N_{\text{Bkg}}}{N_{\text{Bkg}}}. \tag{7.22}$$

By scanning across the X distribution in windows of size ΔX, different tagging efficiencies and background rejections may be calculated. A scatter plot of these different tagging efficiencies against background rejections then forms the so-called *Receiver operating characteristic* (ROC) curve [133]. The typical form of a ROC curve is illustrated in Figure 7.26, where the best taggers have high tagging efficiencies and comparatively high background rejection factors.

Now that we have developed a formal procedure for quantifying the performance of a tagger, let us return to the two tagging performance metrics applied to the shortlisted jet collections.



(a) $300 < p_{\mathrm{T}}^{\mathrm{true}} < 500$ GeV.

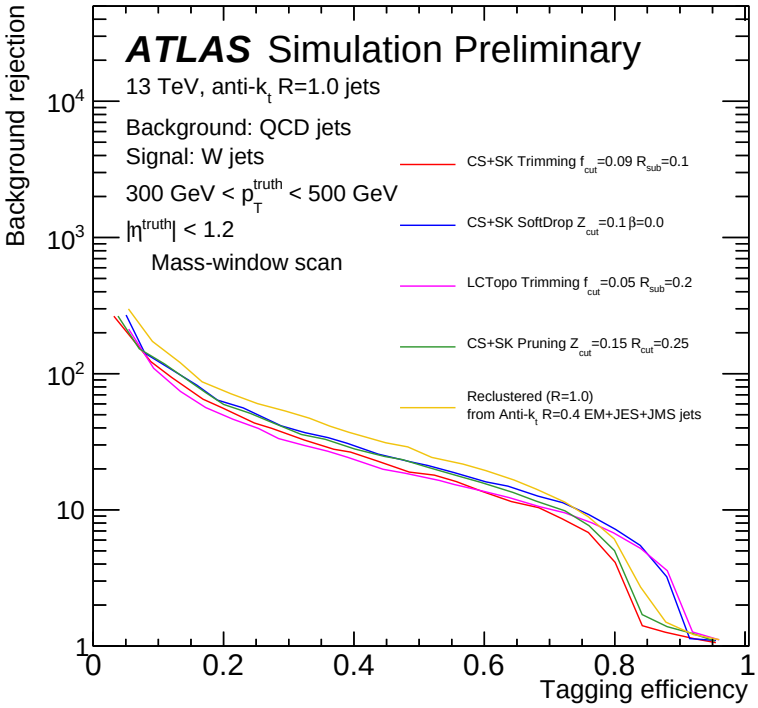


(b) $1000 < p_{\mathrm{T}}^{\mathrm{true}} < 1500$ GeV.

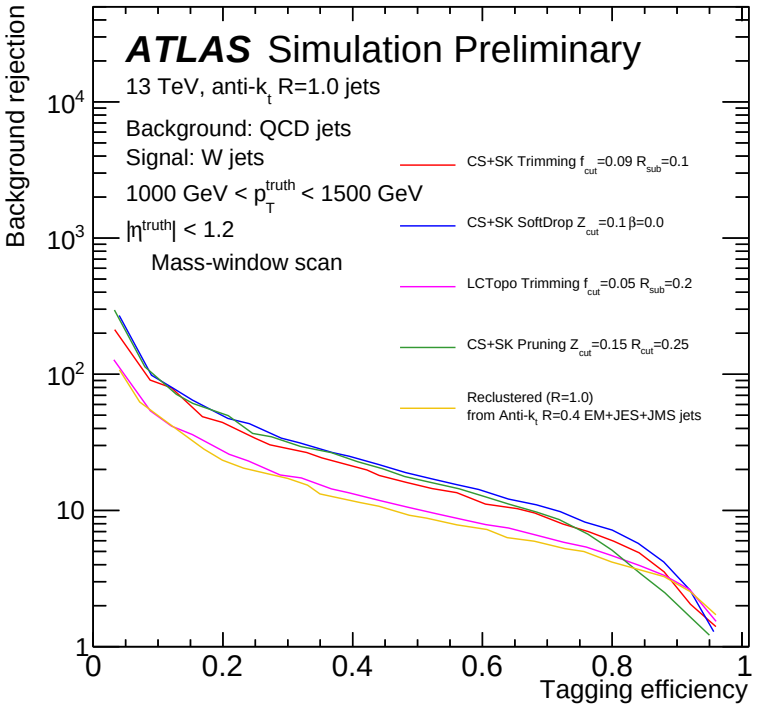
Figure 7.27: The jet mass distributions for $W-$jet signal and QCD background topologies, for different shortlisted jet definitions.

In the first such study, a mass-only tagger is created by cutting on the calorimeter jet mass for the signal W boson and background QCD jets. Such jet masses are shown, for two different $p_{\mathrm{T}}^{\mathrm{true}}$ bins, in Figure 7.27. The smallest possible mass window and central mass position are constructed around the W boson signal peak in order to provide some W boson signal tagging efficiency. For this fixed tagging efficiency, the background contribution within this mass window is used to obtain the background rejection. The signal tagging efficiency is then varied, and a new minimum W mass window consistent with this tagging efficiency is calculated. Then, as before, the background rejection is re-evaluated. By continuing to vary the tagging efficiency, and recalculate the new background rejection factor, the two example ROC curves in Figure 7.28 are obtained, for low and high jet transverse momentum bins. At lower transverse momenta, the performance between the different jet collections tends to be similar since the background QCD jet masses are typically pushed to lower values, and consequently have very similar mass tails falling within the signal peak for the different jet collections. The reclustered jet collection gives the optimal performance except at high tagging efficiency, where there is a decline. The trimmed and pruned jet collections also drop dramatically in this region, due to their particular grooming configurations which are harsh and therefore remove substantial amounts of signal, as well as background, jets. At higher $p_{\mathrm{T}}^{\mathrm{true}}$ the new `CSSKLCTopo` jet collections perform best, as expected since these jet collections were selected based on their superior performance under the background efficiency metric.

The second-tagger uses both mass and D_2 information. First, in a similar manner to the background efficiency metric calculation, the minimum mass window around the W signal is found such that 68 % of all signal jets fall into the mass range. With this fixed mass window, the D_2 distribution is then scanned (example distributions, after the 68 % mass window selection is applied, are shown in Figure 7.29), where the signal W jets have lower D_2 owing to their two-prong energy structure coming from the hadronic two-point decay of the W boson. Scanning in windows of D_2 then generates the 'mass + D_2' tagger ROC curves shown in Figure 7.30. The performance with this second tagger is dramatically different to the mass-only tagger: now the standard trimmed `LCTopo` jets perform best at lower and higher $p_{\mathrm{T}}^{\mathrm{true}}$, and this difference is particularly striking in the latter case. An understanding of this behaviour can in fact be gleaned from Figure 7.29. At higher transverse momenta in particular, it is clear that the new jet algorithms leave the $W-$jet D_2 relatively unchanged, but they do modify the QCD jets



(a) $300 < p_{\mathrm{T}}^{\mathrm{true}} < 500$ GeV.



(b) $1000 < p_{\mathrm{T}}^{\mathrm{true}} < 1500$ GeV.

Figure 7.28: The jet mass tagger ROC curves generated using $W-$jet signal and QCD background topologies, for different shortlisted jet definitions.

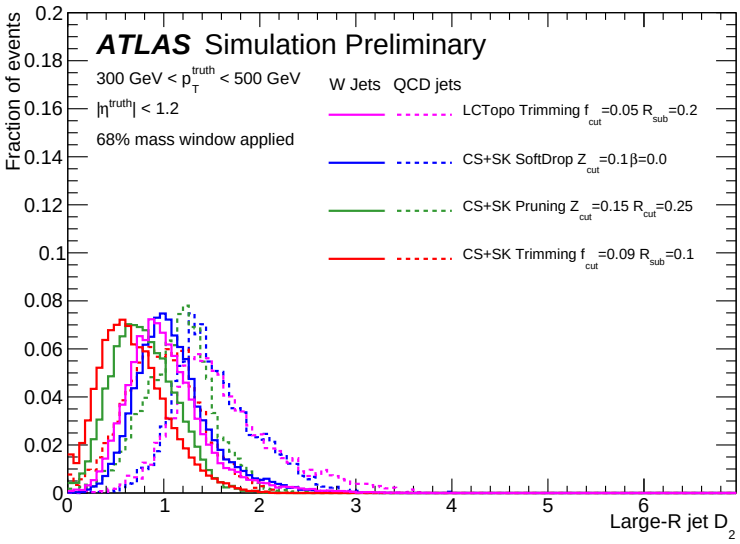
more significantly. For QCD jets these new algorithms have the effect of cutting away the radiation more aggressively than the case of standard trimming, resulting in far fewer clusters in the jet. This reduces the discrimination power between the QCD and $W-$jets, and hence we see the dramatic drop in tagging performance from these `CSSKLCTopo` plus groomer jet collections. A natural question is which component of the jet algorithm is removing the soft radiation in the jet: the constituent-subtraction applied to the jet inputs or the jet grooming itself? Studies which are not detailed here [123] revealed that in fact the choice of jet groomer plays the primary role in over- or under-removing radiation captured within the substructure of the reconstructed jet. This leads to a very important lesson: although aggressive grooming configurations in trimming, pruning, and Soft Drop give rise to optimal robustness against pileup, and maximize background rejection power, they also suffer from removing too much radiation from jets such that the important substructure information associated with the jet can be compromised. Therefore, for future studies, it was concluded that it is dangerous to rely on overly-aggressive grooming (e.g. $\beta = 0.0$ Soft Drop) which risks removing too many clusters and impairing the substructure discrimination utilised in conventional boson and top taggers. With this lesson in mind, we shall now summarise the most recent scan which followed the results outlined thus far. In this new scan, several important developments in Soft Drop grooming shall be investigated in the context of W boson and top-tagging.

7.4.2 The Updated Inputs and Grooming Scan

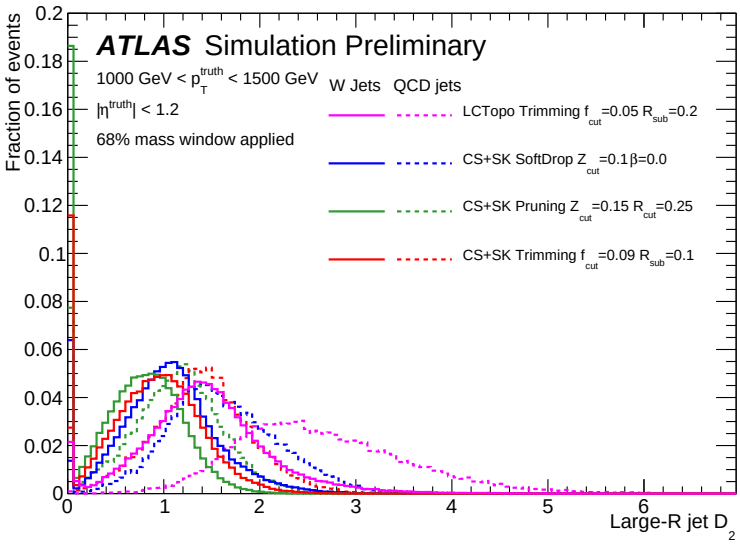
7.4.2.1 Tagging Performance

In the most recently published iteration of the jet inputs and grooming scan [131], the two latest developments in Soft Drop grooming (recursive and bottom-up Soft Drop [120]) have been included. Now only three jet inputs are considered, based on the performance from the previous inputs and grooming scan: `LCTopo` (unmodified), `CSSKLCTopo`, and `VorSuppSKLCTopo`. In recursive Soft Drop, $N = 3$, and $N = 5$ iterations are considered. A few of the most promising Soft Drop, pruned, and trimmed grooming parameter sets are also included in this new scan.

The background jets considered in this latest scan are again the QCD jets from Pythia 8 Monte Carlo. The signal jets are likewise $W-$jets, obtained by W-filtering the $W' \to WZ$ Monte Carlo, and matching reconstruction-level jets to a truth-level W boson within $\Delta R < 0.75$. For this signal and background, once again a 68 % signal fraction mass window is constructed

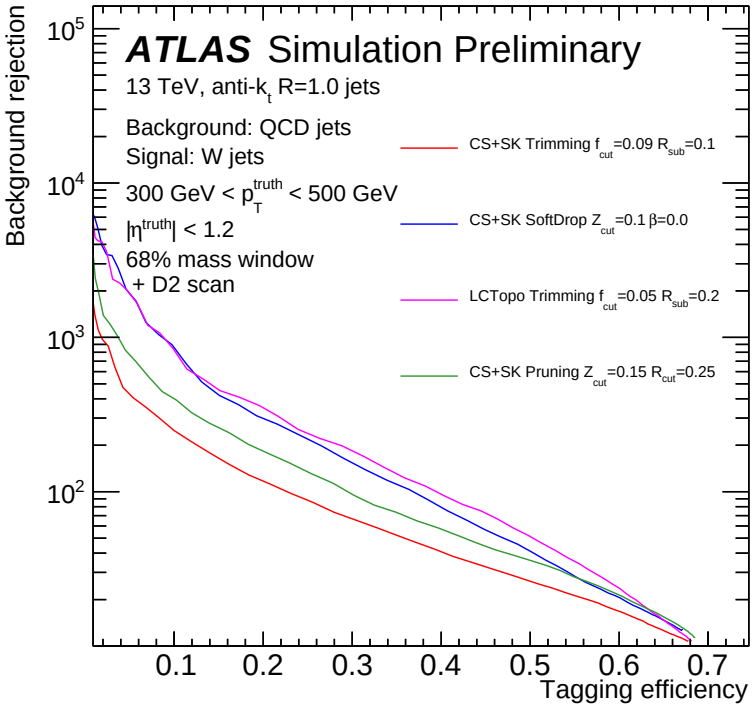


(a) $300 < p_{\mathrm{T}}^{\mathrm{true}} < 500$ GeV.

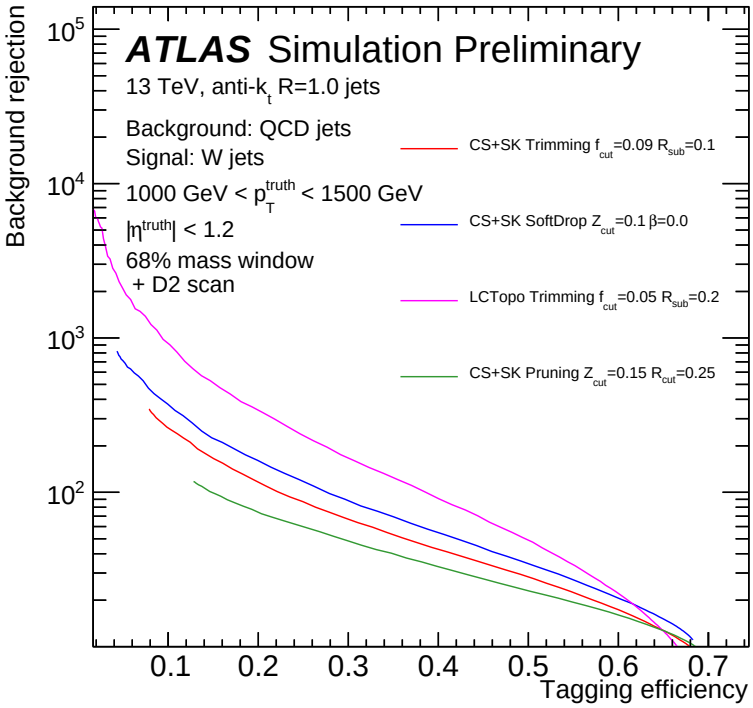


(b) $1000 < p_{\mathrm{T}}^{\mathrm{true}} < 1500$ GeV.

Figure 7.29: The jet D_2 distributions for $W-$jet signal and QCD background topologies, for different shortlisted jet definitions.

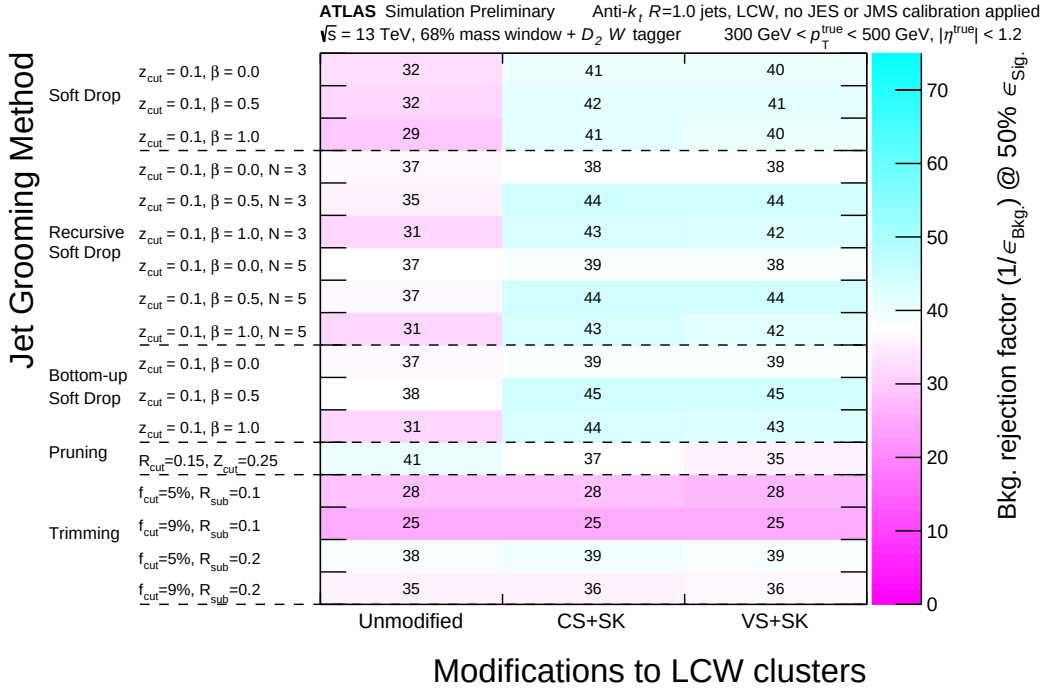


(a) $300 < p_{\mathrm{T}}^{\mathrm{true}} < 500$ GeV.

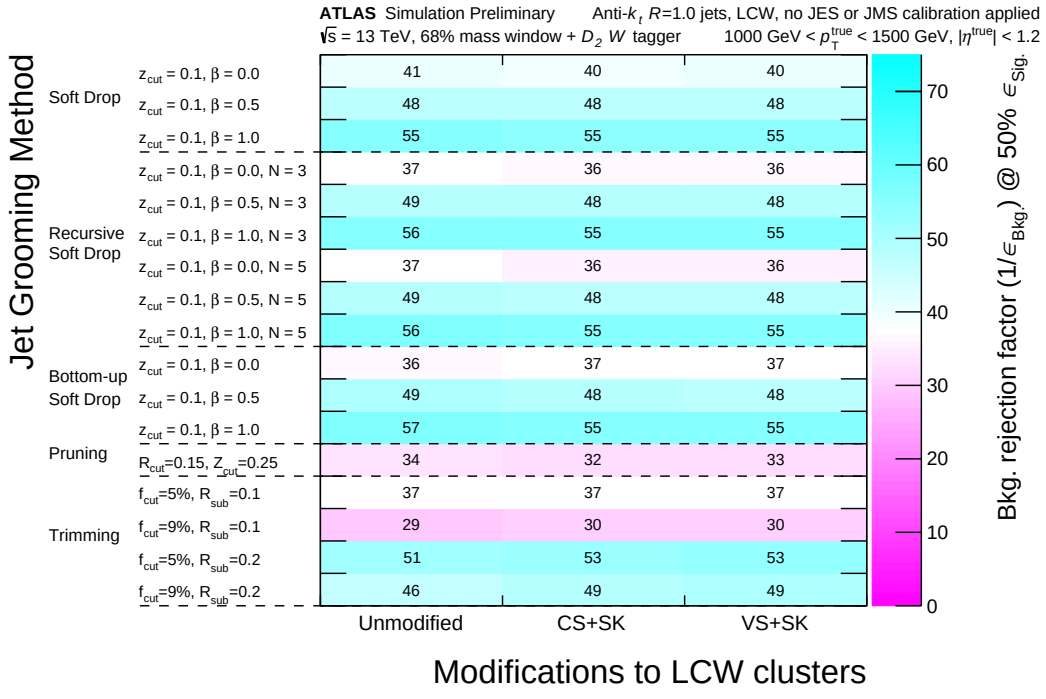


(b) $1000 < p_{\mathrm{T}}^{\mathrm{true}} < 1500$ GeV.

Figure 7.30: The jet 'mass + D_2' tagger ROC curves generated using $W-$jet signal and QCD background topologies, for different shortlisted jet definitions.



(a) QCD background rejections, $W-$jet signal, $300 < p_{\mathrm{T}}^{\mathrm{true}} < 500$ GeV.



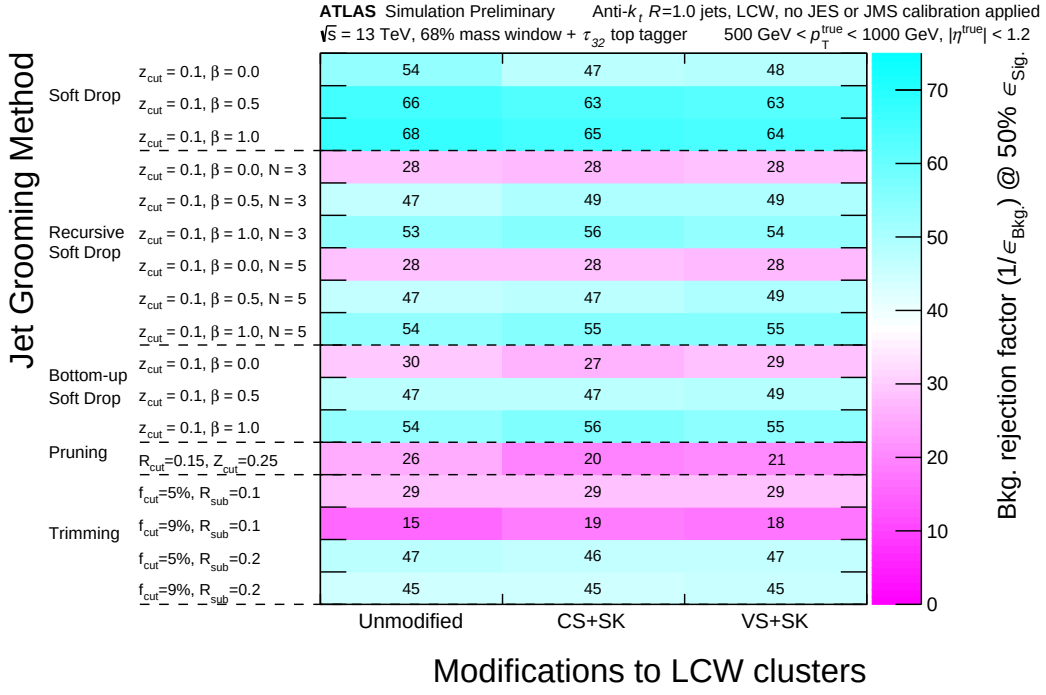
(b) QCD background rejections, $W-$jet signal, $1000 < p_{\mathrm{T}}^{\mathrm{true}} < 1500$ GeV.

Figure 7.31: Heatmaps for QCD background rejection factors, using a $W-$jet signal, for different new jet input and grooming collections.

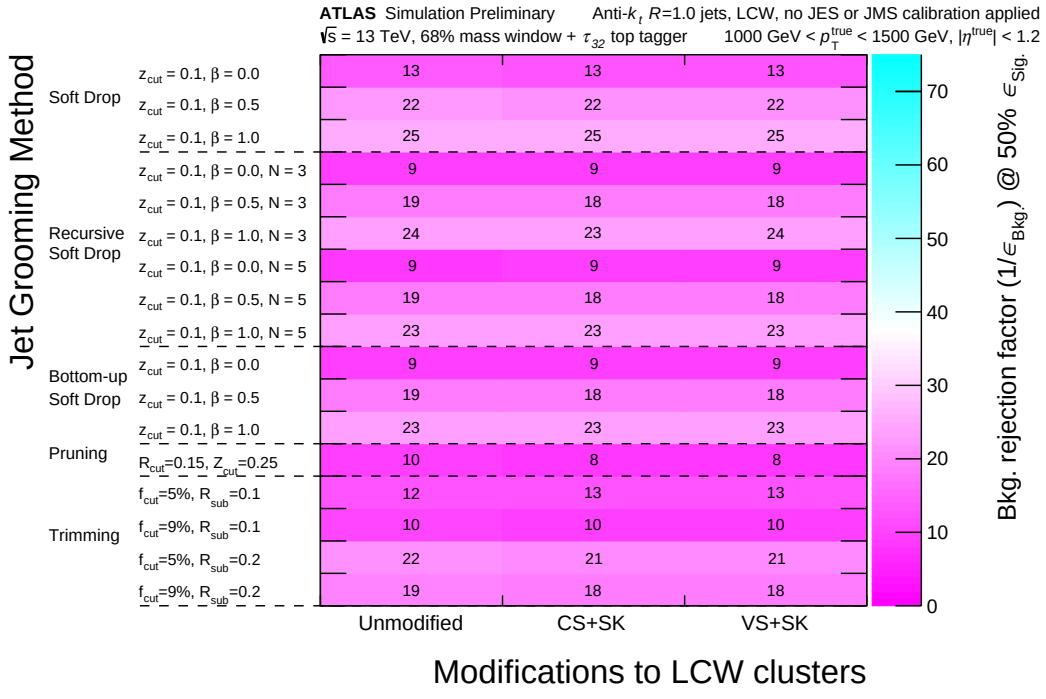
around the W boson mass peak, and the D_2 variable is scanned in order to generate the ROC performance curve for this 'mass + D_2' tagger. For different $p_{\mathrm{T}}^{\mathrm{true}}$ ranges the QCD background rejection factor has been calculated for a fixed signal tagging efficiency of 50 %, where the best performing jet collections will produce the largest background rejections. These results have been collected into the new heatmaps displayed in Figure 7.31 for high- and low-$p_{\mathrm{T}}^{\mathrm{true}}$ ranges. In the lower $300 - 500$ GeV $p_{\mathrm{T}}^{\mathrm{true}}$ range the pileup contributions to the jet masses are more significant, and therefore notable improvements in the background rejection are observed when switching to one of the input definitions based on constituent or Voronoi subtraction. Looser recursive Soft Drop or bottom-up Soft Drop groomers (such as $\beta = 0.5, 1.0$) provide marginal gains over 'vanilla' Soft Drop grooming. At higher transverse momentum, the improvements from different jet inputs are overshadowed by the groomer choice, and in fact it is clear that some grooming configuration of a Soft Drop approach, or trimming, offers very good background rejection, whilst pruning is clearly offering suboptimal performance.

A similar tagging study has been carried out for the top jets signal, obtained from the $Z' \to t\bar{t}$ described in previous chapters in this thesis. The top jets are required to have a minimum $p_{\mathrm{T}}^{\mathrm{true}}$ of 500 GeV so that the top quark decay products may be fully-contained within the $R = 1.0$ jet. The reconstructed jet must also be matched to a truth-level top quark within $\Delta R < 0.75$. As with the W tagger, first a 68 % top jet fraction mass peak is constructed using the smallest possible window around the top mass pole. Following this, the τ_{32} is calculated for the top jets and different values of τ_{32} are scanned, thus generating a new top-tagging ROC curve. Recall that lower τ_{32} values correspond to three-prong jet substructures (as would be expected from the dominant $t \to Wb$ decay.) This 'mass + τ_{32}' tagger is inspired by existing cut-based top-taggers already in use in ATLAS [132]. For a fixed 50 % top signal efficiency, another set of background rejection factor heatmaps have been generated for two transverse momentum bins, shown in Figure 7.32. It is interesting to note that, for the top-tagging, the largest gains in background rejection, at a moderate $p_{\mathrm{T}}^{\mathrm{true}}$ range of $500 - 1000$ GeV, come from the 'vanilla' Soft Drop grooming, whilst at much higher travserse momentum the background rejection is poorer across all combinations of inputs and groomers.

In addition to the background rejection, once again the pileup robustness of the jet mass has been studied for the $W-$jet signal, and is summarised by the heatmaps in Figure 7.33. Recall



(a) QCD background rejections, top jet signal, $500 < p_{\mathrm{T}}^{\mathrm{true}} < 1000$ GeV.



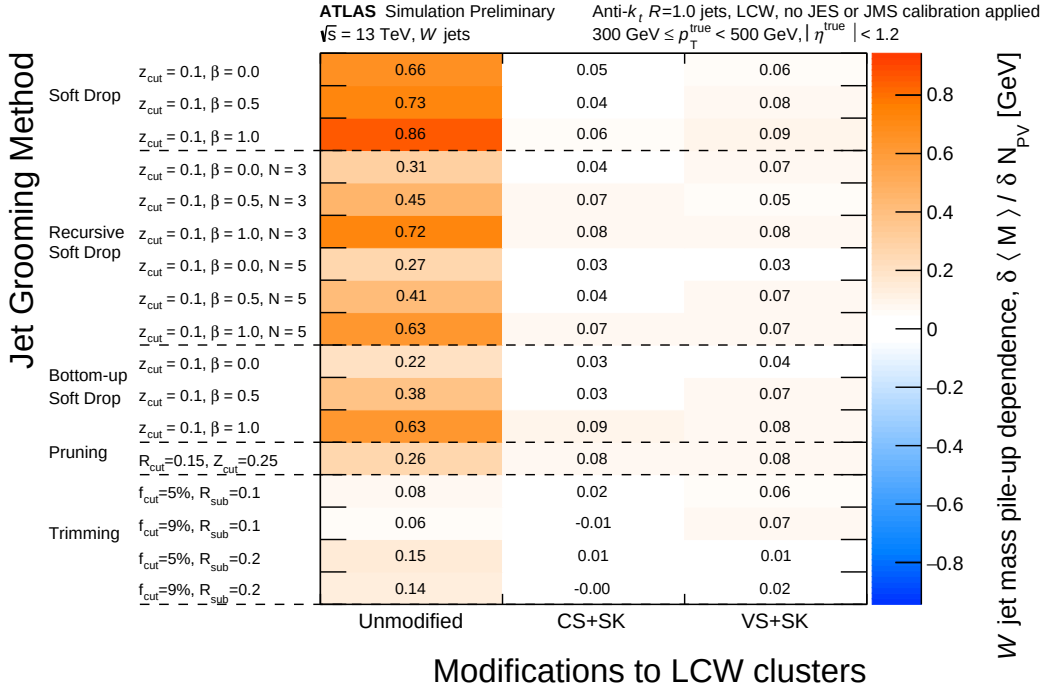
(b) QCD background rejections, top jet signal, $1000 < p_{\mathrm{T}}^{\mathrm{true}} < 1500$ GeV.

Figure 7.32: Heatmaps for QCD background rejection factors, using a top jet signal, for different new jet input and grooming collections.

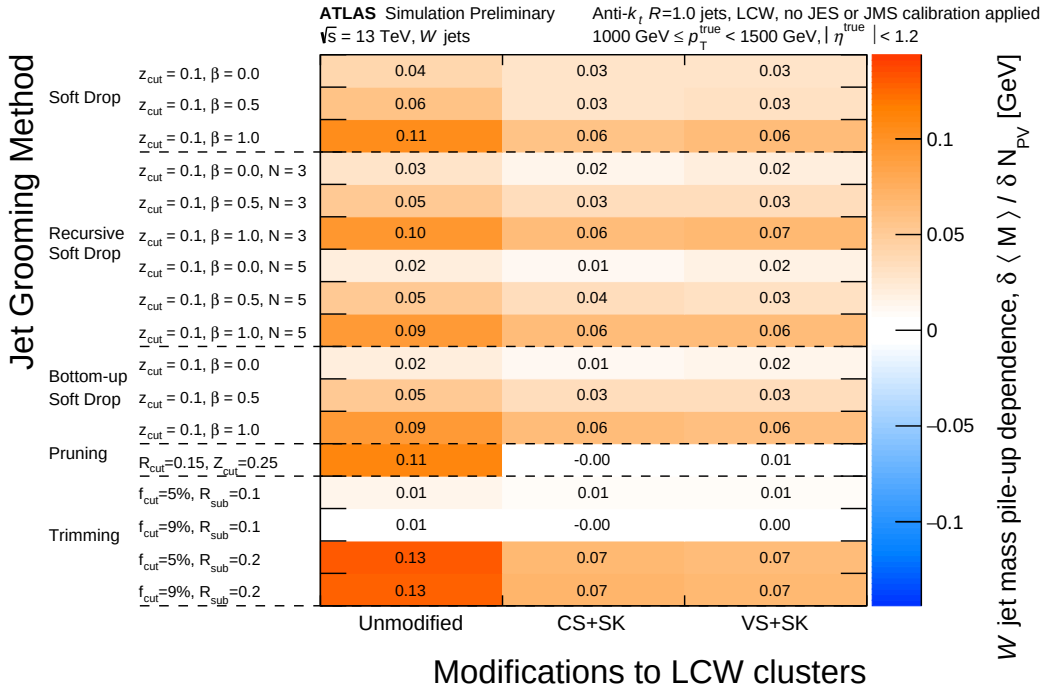
that the most robust jet collections have jet mass-N_{PV} gradients close to 0.0. Once again, for all jet groomers, the pileup robustness at $p_{\mathrm{T}}^{\mathrm{true}}$ 300 − 500 GeV is most clearly reduced by using constituent subtraction or Voronoi subtraction on the `LCTopo` inputs. At higher transverse momentum the differences between the various jets are markedly smaller, but there are a few promising jet definitions which make use of Soft Drop. Notably, excellent pileup robustness is observed in e.g. $f_{\mathrm{cut}} = 0.09$, $R_{\mathrm{sub}} = 0.1$ jet trimming, but it must be recalled that this aggressive grooming will remove substantial jet substructure information and lead to suboptimal tagging performance.

Based on this second scan, several particularly promising combinations of jet input and groomer have been calibrated at the JES and JMS, so that their tagging performance can be studied with biases between reconstruction- and truth-level Monte Carlo minimised. The ROC curves for the tagging performance of $W-$jets and top jets (which use the 'mass + D_2' and 'mass + τ_{32} taggers respectively) are displayed in Figures 7.34 and 7.35. Three different configurations of `CSSKLCTopo` $z_{\mathrm{cut}} = 0.1$ 'vanilla' Soft Drop ($\beta = 0.0$, 0.5, 1.0), and one of `CSSKLCTopo` with standard trimming, are considered. For comparison, the original unmodified `AntiKt10LCTopoTrimmedPtFrac5SmallR20` jets are included in solid black. The jet performance varies between signal topology and $p_{\mathrm{T}}^{\mathrm{true}}$ range, but general features in the tagging performance are apparent. In the lowest $p_{\mathrm{T}}^{\mathrm{true}}$ ranges (300 − 500 GeV for $W-$jets, and 500 − 1000 GeV for top jets) the $\beta = 1.0$ Soft Drop jets offer the best tagging performance based on the moderate grooming provided by $\beta = 1.0$ which successfully reduces wide-angle and soft contributions to the jet without aggressively grooming away the substructure information. At higher $p_{\mathrm{T}}^{\mathrm{true}}$ the gains from $\beta = 1.0$ Soft Drop are less apparent in the W tagger, where there is degeneracy in performance between Soft Drop and the two trimmed jet definitions. However, in the case of the top-tagger, the improvements from the $\beta = 1.0$ jet collection are still clear across all signal tagging efficiencies. Such observations make these `AntiKt10CSSKLCTopoSoftDropZCut10Beta100` jets potentially very promising in future boosted physics measurements and searches.

The studies described in this chapter have used the *calorimeter* jet mass definition. The more optimal combined mass performance is briefly discussed in Appendix B.

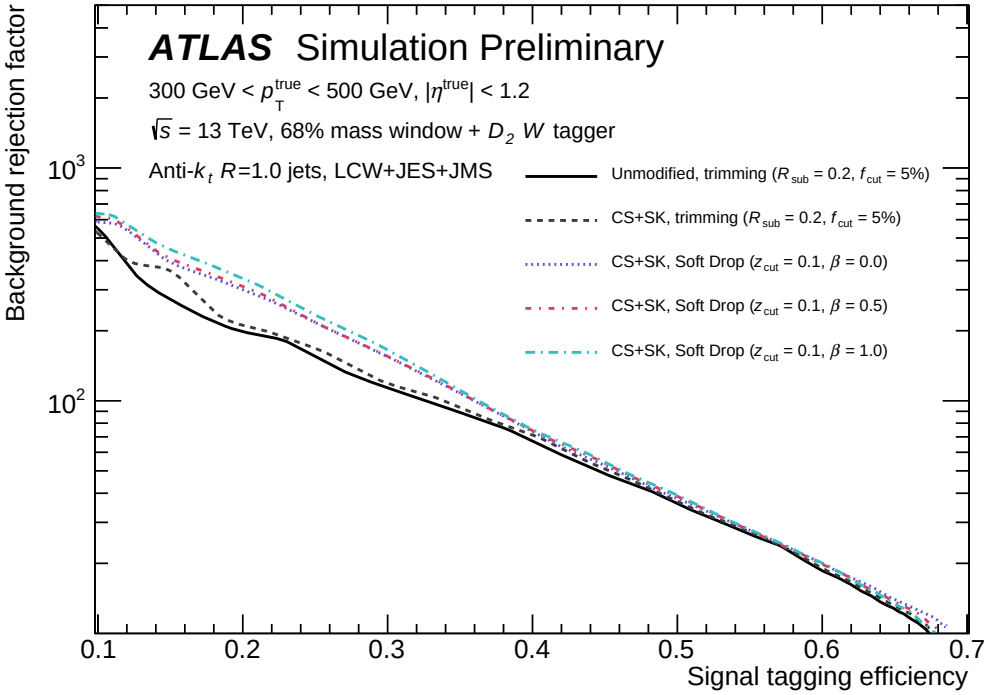


(a) Average $W-$jet mass pileup dependence, $300 < p_{\mathrm{T}}^{\mathrm{true}} < 500$ GeV.

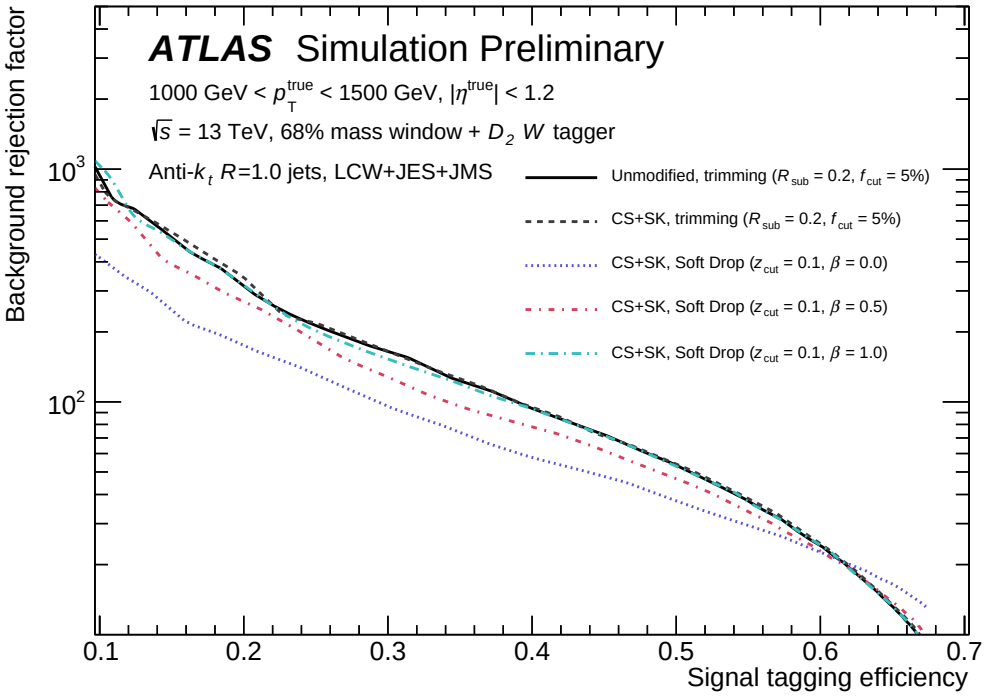


(b) Average $W-$jet mass pileup dependence, $1000 < p_{\mathrm{T}}^{\mathrm{true}} < 1500$ GeV.

Figure 7.33: The average $W-$jet mass pileup dependence, with respect to N_{PV}, for low and high-$p_{\mathrm{T}}^{\mathrm{true}}$ ranges.

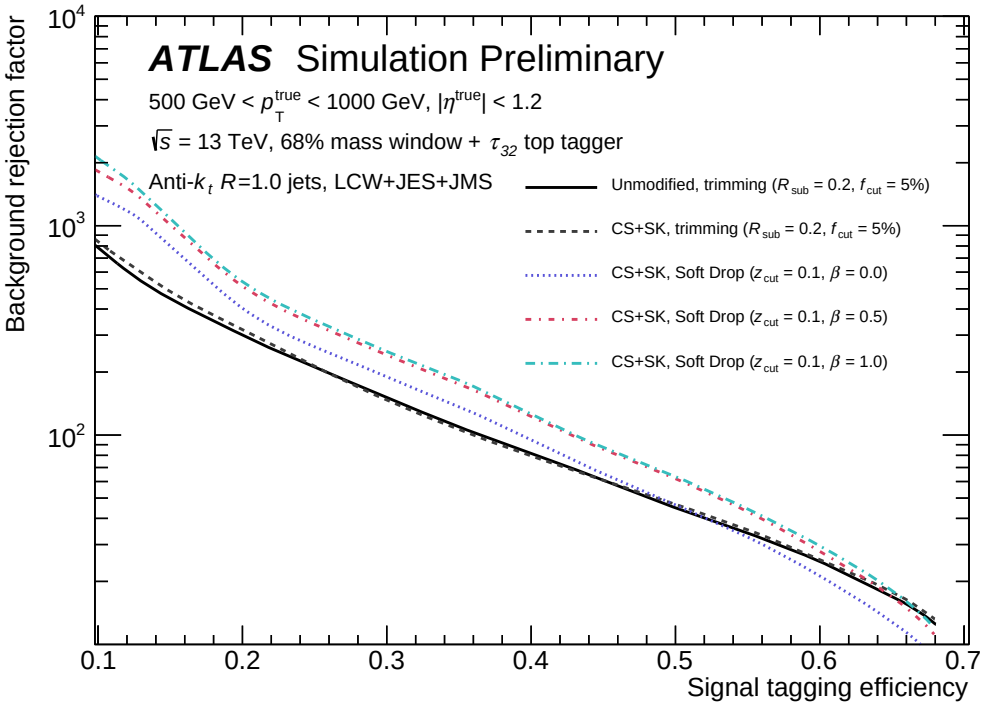


(a) Jet 'mass + D_2' tagger, $300 < p_{\mathrm{T}}^{\mathrm{true}} < 500$ GeV.

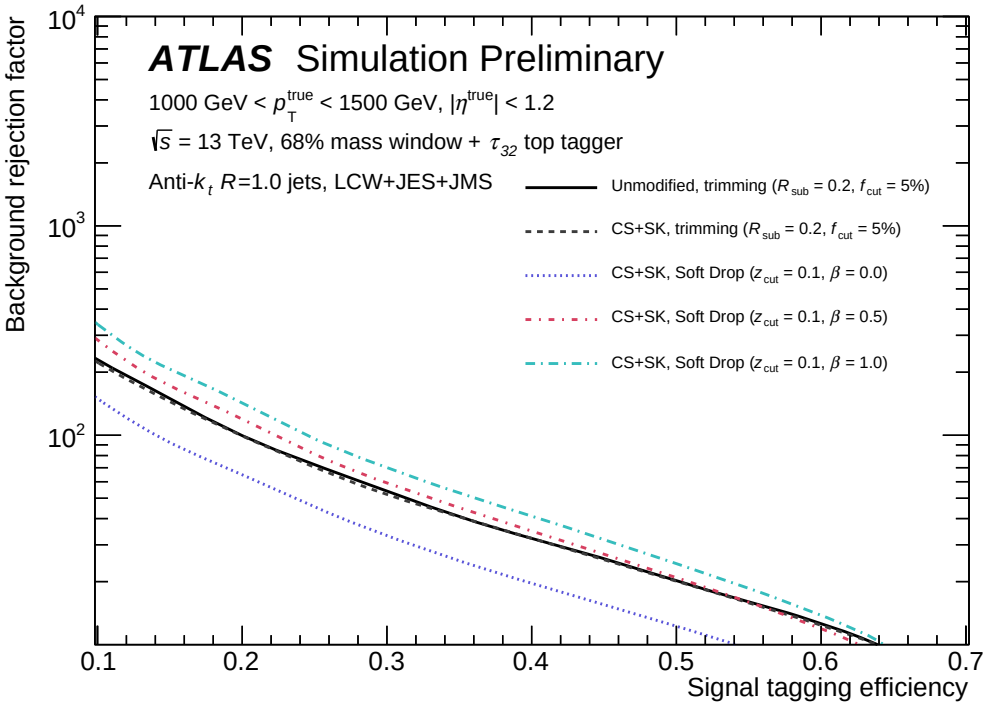


(b) Jet 'mass + D_2' tagger, $1000 < p_{\mathrm{T}}^{\mathrm{true}} < 1500$ GeV.

Figure 7.34: The jet 'mass + D_2' tagger ROC curves generated using $W-$jet signal and QCD background topologies, for different shortlisted jet definitions, based on the very latest inputs and grooming scan.



(a) Jet 'mass + τ_{32}' tagger, $500 < p_{\mathrm{T}}^{\mathrm{true}} < 1000$ GeV.



(b) Jet 'mass + τ_{32}' tagger, $1000 < p_{\mathrm{T}}^{\mathrm{true}} < 1500$ GeV.

Figure 7.35: The jet 'mass + τ_{32}' tagger ROC curves generated using top jet signal and QCD background topologies, for different shortlisted jet definitions, based on the very latest inputs and grooming scan.

7.4.3 The Future of Large-Radius Jets at ATLAS

In this chapter we have reviewed the standard $R = 1.0$ jet definition adopted in boosted object searches and measurements in ATLAS, interesting theoretical developments in pileup mitigation, jet grooming, and jet substructure, and finally how such breakthroughs can (and have) been used to construct new large-radius jet collections which have the potential to greatly improve pileup robustness and boosted object tagging in ATLAS analyses. It is becoming clear that the standard `AntiKt10LCTopoTrimmedPtFrac5SmallR20` jets, which have been the large-radius jet workhorse of ATLAS in recent years [90], can be improved upon using the latest developments in Soft Drop jet grooming and constituent subtraction of calorimeter energy deposits. The recent inputs and grooming studies presented here [123, 131] have brought these developments into the spotlight of the ATLAS jet physics community, but it is evident that we have merely scratched the surface. Future improvements, such as the use of particle flow and tracking information, have the potential to further revolutionise jet physics in colliders. Two particularly promising developments are the extension of particle flow jet inputs from $R = 0.4$ to $R = 1.0$ jets, the use of new tracking-based cluster variables (known as track-calo clusters [124]), and ultimately the combination of particle flow and track-calo cluster jet definitions. These advancements, together with the addition of constituent subtraction, SoftKiller, and Soft Drop grooming, will change the face of jet reconstruction of ATLAS as we enter Run 3 of the LHC. The work described in this Chapter has been central to future plans for large-radius jet reconstruction in ATLAS, and work is currently ongoing to finalise these plans so that new definitions with substantially improved tagging performance and pileup resilence can be deployed in good time for physics analyses for Run 3 of the LHC. Latest studies indicate that large-radius jet definitions based on particle flow inputs, together with constituent subtraction and SoftKiller, with additional Soft Drop (recursive or bottom-up) grooming applied will be very promising alternatives to the current trimmed definition.

Part III

The Hunt for New Matter in Multi-Jet Final States

If it disagrees with experiment, it's wrong. In that simple statement is the key to science.

– Richard P. Feynman,
The Character of Physical Law

8

Hunting for new phenomena in multi-jet final states

One of the strongest motivations for the construction of the LHC, and the ATLAS experiment, is to search for new physics and new matter at the high-energy frontier. The LHC has already proved to be enormously successful, culminating with the discovery of the Higgs boson [8, 9] by ATLAS and CMS towards the close of its first run. In parallel to these endeavours, it has long been hoped that the LHC will provide the necessary energy scale at which new, heavier, BSM particles might be discovered. In particular it is hoped that the discovery of supersymmetric particles might be achieved during the LHC run. The final part of this thesis attempts to synthesise the physics of Part I and Part II: we shall now focus on a search for new physics, motivated by the theory of supersymmetry, which relies on the principles of hadronic jet reconstruction to search for new forms of matter in final states where the hadronic jet activity is complex. We shall refer to this search as the 'multi-jets' analysis; such an analysis has been one of the benchmark searches for SUSY during Run 1 [134] and Run 2 [70, 71] of the LHC.

We shall focus on the most recent publication of this analysis [70], a search for SUSY using 36 fb^{-1} of ATLAS data, before briefly describing ongoing plans to extend this analysis to the full Run 2 ATLAS data-set, which has amassed 140 fb^{-1} of integrated luminosity. The full details of this 140 fb^{-1} analysis shall not be provided in this thesis, but this second piece of work is very mature, aiming for a publication very soon, and is being lead by the author. Some details on this ongoing work are provided in Chapter 10. The multi-jets analysis is

motivated through signatures where gluinos are pair-producted in the LHC proton-proton collisions, and can undergo long decay cascasdes, generating heavy-flavour quarks and bosons, and subsequently generating large multiplicities of hadronic jets, which may (or may not) have moderate $E_{\mathrm{T}}^{\mathrm{miss}}$ (see Section 4.4.6 for a definition). This feature puts the multi-jets analysis in at odds with many other searches for SUSY-motivated new physics, which typically rely on much cleaner final states with large amounts of $E_{\mathrm{T}}^{\mathrm{miss}}$ generated by the presence of an undetected, stable, lightest supersymmetric particle (LSP). A particularly attractive aspect of the multi-jets analysis is its sensitivity to a wider variety of new physics models; this shall be demonstrated in the next chapter, where the results of this search are used to constrain both R-parity-conserving and R-parity-violating SUSY signals.

8.1 Analysis Overview

The basic analysis strategy is to create *signal regions* from events which contain at least 8, 9, 10, or 11 `AntiKt4EMTopo` jets, zero leptons, and moderate $E_{\mathrm{T}}^{\mathrm{miss}}$. The jets are also required to have $p_{\mathrm{T}} > 50$ GeV in all signal regions. These final states are motivated by models of new physics where the pair production of gluinos can give rise to large jet multiplicities originating from the decay cascasde of supersymmetric particles; specific models which are used to motivate region definitions shall be described in Section 8.2. We also require that jets entering signal regions have a $E_{\mathrm{T}}^{\mathrm{miss}}$ significance (defined: $E_{\mathrm{T}}^{\mathrm{miss}}/\sqrt{H_{\mathrm{T}}}$) above 5.0 $\mathrm{GeV}^{1/2}$, where H_{T} comes from the sum of the jet p_{T}s in the event. This variable is important for reducing the contamination from standard model multi-jet events in the signal regions, and is indeed the key variable used in measuring the major multi-jets background arising from QCD SM processes. Further details of these kinematic variables are provided in Section 8.3.4.

Complex decay cascades provide copious numbers of heavy quarks, giving rise to jets initiated by botton quarks (referred to as b-jets) and top jets. To enhance sensitivty to such heavy flavour events, the analysis is further split into two streams: the *flavour stream* where `AntiKt4EMTopo` jets which pass b-tagging requirements [65] are selected; and the *fat-jet stream* which uses jet substructure information from large-radius jets to enhance the expected sensitivity to final states with moderately boosted top jets, as well as reducing the $t\bar{t}$ background contribution. Another major challenge of this analysis is the large SM background, originating from both leptonic activity and the production of SM hadronic QCD multi-jet events. Following the prediction of

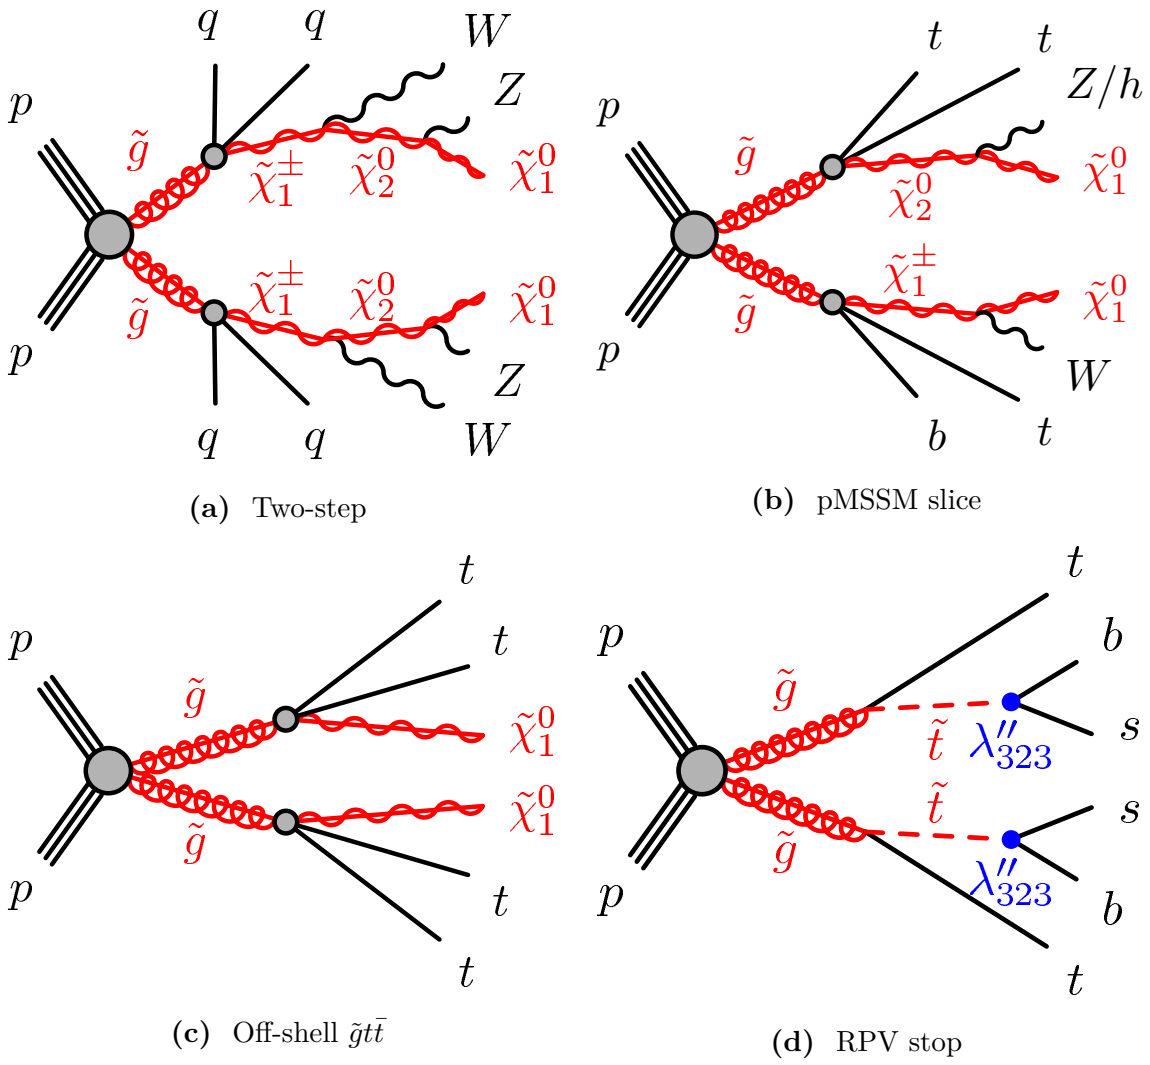
these background contributions, the data are compared with the SM expectations in order to determine if any statistically significant excess of events is observed, as would be expected from the presence of additional decay cascades arising from direct gluino production.

Although leptons are explicitly vetoed in the signal regions, it is still very possible that leptons which are either misidentified as jets, or which fall out of the acceptance of the ATLAS detector, enter these signal events. Such leptons originate dominantly from the semileptonic decay of top quarks and W bosons, and it is therefore necessary to estimate the contribtution to signal regions arising from these semileptonic top and 'W + jets' events. This is done using dedicated *control regions*: these are kinematically very close to the signal regions (i.e. they share the same jet multiplicities) but are rendered orthogonal to them by requiring the presence of an electron or muon. Monte Carlo estimates of the dominant top and W + jets backgrounds in these control regions are then scaled to the data measured there using a background-only fit [135], and the final normalisations extracted are then used to scale the leptonic backgrounds contaminating the signal regions. This control region approach shall be described in detail in Section 8.4.

The dominant background QCD cascade decays in the signal regions come from SM multi-jet processes, which can look near identical to cascade decays originating from heavy, central SUSY particle production. Monte Carlo simulations of 8, 9, 10, or 11-jet hadronic final states are not reliable predictions of such backgrounds, and would require immense computational power to generate a sufficiently large number of events. For these reasons, a fully data-driven approach must be adopted to estimate this background. This fully data-driven approach uses dedicated *template regions* to estimate the fully hadronic QCD jet background in the signal regions, and is referred to as the *template method* (originally described in [134]). The template method shall be outlined in Section 8.5.

8.2 Signal Models

The basic motivation behind simplified models has already been outlined in Section 3.3.2. In this analysis three simplified models are considered, along with a more complex phenomenological MSSM parameter space-inspired model [136] (the 'pMSSM slice'). The pMSSM slice, and the two-step, and $\tilde{g}t\bar{t}$ simplified models are used to aid the optimisation of kinematic variables and



(a) Two-step **(b)** pMSSM slice

(c) Off-shell $\tilde{g}t\bar{t}$ **(d)** RPV stop

Figure 8.1: The pMSSM slice and various simplified models considered in the multi-jets analysis.

jet substructure techniques, while the final model, referred to as the RPV stop model, is added in order to quantify the analysis sensitivity to R-parity-violating SUSY. The full statistical interpretations of these models shall be described in the next Chapter.

8.2.1 Two-Step Model

The two-step gluino decay model is one of the benchmark strongly-produced SUSY models, which has been in use since Run 1 of ATLAS [134]. In the first step the gluino undergos a decay, via a virtual squark, into a pair of first or second generation quarks and the lightest chargino, $\tilde{\chi}_1^{\pm}$. Subsequently, the chargino decays to a W boson and the next-to-lighest neutralino, $\tilde{\chi}_2^0$. In the final stage, this $\tilde{\chi}_2^0$ decays to a Z boson and the LSP, $\tilde{\chi}_1^0$ i.e. the lightest neutralino. Each step in the decay chain is assumed to have a 100 % branching fraction. The decay chain can be compactly summarised as:

$$pp \to \tilde{g}\tilde{g}, \quad \tilde{g} \to qq\tilde{\chi}_1^{\pm}, \tag{8.1}$$

$$\tilde{\chi}_1^{\pm} \to W^{\pm}\tilde{\chi}_2^0, \quad \tilde{\chi}_2^0 \to Z\tilde{\chi}_1^0. \tag{8.2}$$

This model is defined by mass parameters $m_{\tilde{g}}$ and $m_{\tilde{\chi}_1^0}$, where $m_{\tilde{\chi}_1^{\pm}} = \frac{1}{2}(m_{\tilde{g}} + m_{\tilde{\chi}_1^0})$ and $m_{\tilde{\chi}_2^0} = \frac{1}{2}(m_{\tilde{\chi}_1^{\pm}} + m_{\tilde{\chi}_1^0})$, and the pseudo-Feynman diagram is provided in Figure 8.1a. Such a model generates large jet multiplicities, where the jets typically originate from light quark production.

8.2.2 pMSSM Slice

The pMSSM slice is inspired by the results of a large Run 1 scan of the pMSSM parameter space [136], where it was found that the Run 1 multi-jets analysis is uniquely sensitive to a class of models where: 1) the gluino has a mass around 1.2 TeV; 2) there is an intermediate Higgsino-like triplet of SUSY particles; and 3) there exists a bino-like LSP with mass around 60 GeV. The decay spectrum for such a model [136] is shown in Figure 8.2. Based on this unique portion of phase-space, a new model is generated by varying the gluino and Higgsino mass parameters and requiring the LSP to be bino-like[1], with a corresponding 60 GeV mass. Additionally, squark and slepton masses are set to 5.0 TeV (so that they are decoupled from the gluino decay), and $\tan\beta$ is set to 10. These choices are motivated to ensure that the gluino decays are those which the multi-jets analysis is sensitive to, and that the correct Higgs mass is recovered. This is not a simplified model, and the decays of the intermediate sparticles do not have fixed branching fractions. This gives rise to a more varied set of final decays in the parameter space. Such branching fractions can result in heavy flavour b-jets and top jets being produced in the final state, as indicated by the pseudo-Feynman diagram for the slice in Figure 8.1b. The production of additional Z or Higgs bosons is also permitted by the decays considered in the parameter space this model spans.

[1] 100 % bino in fact.

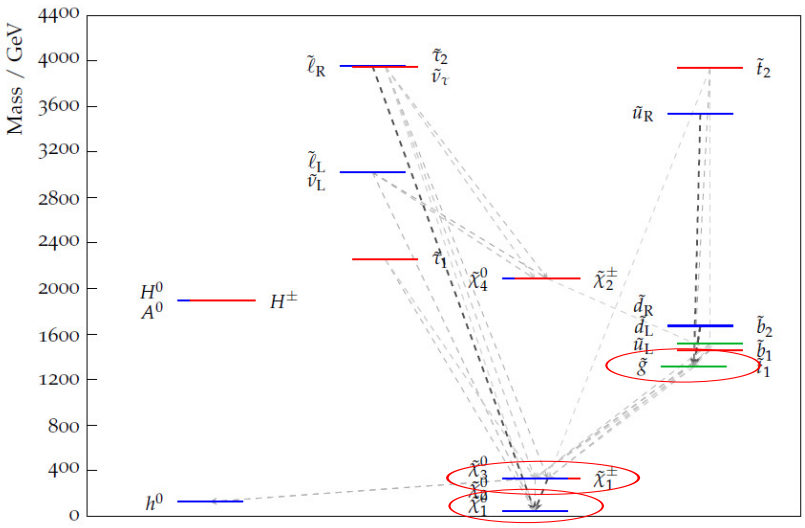


Figure 8.2: The typical mass spectrum for the original pMSSM parameter space for which the Run 1 analysis was found to be uniquely sensitive. This spectrum is generated using SoftSUSY [137] together with the PySLHA software [138].

8.2.3 $\tilde{g}t\bar{t}$ Model

An additional signal model which this analysis should have sensitivity to is the so-called $\tilde{g}t\bar{t}$ model, in which top quark-rich final states are produced by the direct decay of the gluino to $t\bar{t} + \tilde{\chi}_1^0$, via an off-mass-shell top squark. The pseudo-Feynman diagram is shown in Figure 8.1c. The masses of the gluino and the $\tilde{\chi}_1^0$ are varied to form a signal grid, where larger mass splittings will generate more boosted top jets. In this regime it is possible that the model sensitivity can be enhanced by using large-radius jets (trimmed or reclustered) and jet substructure techniques, such as tagging jets based on the jet mass and jet shape variables already described in Section 7.3. Theses studies shall be outlined in Section 8.3.6.

8.2.4 RPV $\tilde{t} \to \bar{s}\bar{b}$ Model

So far we have considered R-parity-conserving models, producing a stable LSP which could, for example, be interpreted as a WIMP dark matter candidate. Many strongly-produced gluino searches are not sensitive to RPV models, which produce final states without a stable LSP and therefore with lower $E_{\mathrm{T}}^{\mathrm{miss}}$. One of the most powerful features of the multi-jets search is that it is able retain sensivity to these RPV scenarios, because events with both low and moderate $E_{\mathrm{T}}^{\mathrm{miss}}$ can fall into its signal regions. Here we consider an RPV stop model which starts with a gluino pair decaying through $\tilde{g} \to \tilde{t} + t$. The last step of the decay

chain contains an RPV baryon-number-violating vertex where the stop LSP decays to heavy-flavour particles: $\tilde{t} \to \bar{s} + \bar{b}$. The signal grid is generated by scanning in $m_{\tilde{g}}$ and $m_{\tilde{t}}$, and the pseudo-Feynman diagram is given in Figure 8.1d.

8.3 Analysis Strategy and Physics Objects

ATLAS data for this analysis were collected from April − November, 2015, and April − November, 2016. The data are amassed from lists of ATLAS runs and luminosity blocks which pass a rigorous data-quality selection [139]. The final amount of data passing these selections, set by the ATLAS Data Quality coordination, totals 36.1 fb^{-1} of integrated luminosity to be used for the multi-jets search.

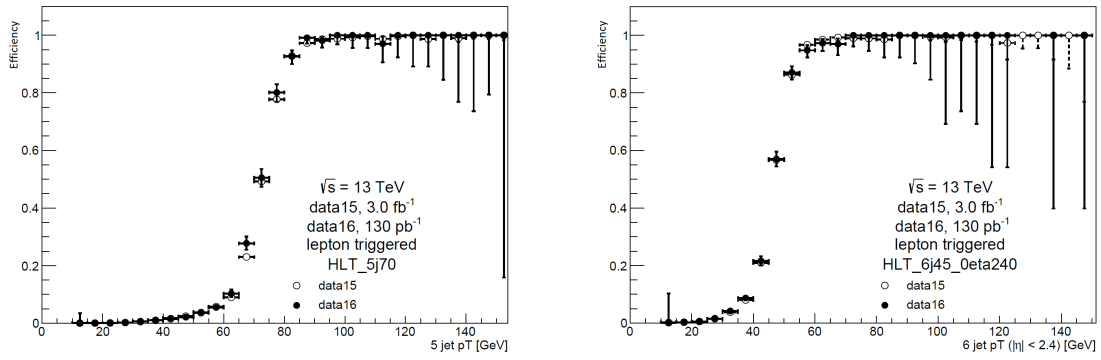
8.3.1 Trigger

The initial event selection is made from events passing the 5- and 6-jet high-level triggers, where these triggers take inputs from events which must pass the hardware-based Level-1 trigger, requiring at least 4 jets with $p_{\mathrm{T}} > 15$ GeV. An `OR` of two high-level triggers is considered: `HLT_5j70` and `HLT_6j45_0eta240`; these require either at least 5 or 6 jets, and are close to fully efficient at $p_{\mathrm{T}} > 70$ GeV or 45 GeV, respectively (and $|\eta| < 2.4$). The efficiencies of these triggers as a function of the jet p_{T} is shown in Figure 8.3, and it can be seen that the 5-jet trigger is fully efficient by 85 GeV, and the 6-jet trigger by 70 GeV. Unlike many SUSY searches, which trigger on $E_{\mathrm{T}}^{\mathrm{miss}}$, the multi-jets analysis only accepts events which cause these multi-jet triggers to fire. This allows us to achieve lower $E_{\mathrm{T}}^{\mathrm{miss}}$ thresholds than many other SUSY analyses. As already mentioned, this provides a greater sensitivity to RPV SUSY signatures or other BSM models which require little or no $E_{\mathrm{T}}^{\mathrm{miss}}$ in the final state.

8.3.2 Event Selection and Cleaning

Prior to the event selection which defines the analysis control and signal regions, the events must first pass a set of cleaning cuts and a preliminary selection.

The event cleaning cuts are applied in order to ensure that the events analysed are both free of detector problems and originate from the proton-proton collisions of the LHC. Only data runs and luminosity blocks which pass stringent data quality checks are selected as ‘good’



(a) `HLT_5j70` **(b)** `HLT_6j45_0eta240`

Figure 8.3: Trigger efficiency turn-on curves for the 5- and 6-jet high-level triggers used to fire on multi-jet events.

runs for users to analyse (the Good Runs Lists). Following this initial quality check, additional cleaning cuts are designed to veto events which have passed into areas of the detector which are no longer functioning, such as several inactive Tile calorimeter cells. Additional vetos are placed to remove muons originating from cosmic rays in the upper atmosphere, to select on good reconstructed primary vertices and to remove bad jets and leptons, as well as contributions from pileup in the event. Events with known recorded errors in the SCT, LAr and Tile calorimeters are also suitably vetoed. The standard event cleaning conditions are summarised in Table 8.1.

Name	Description
Data quality GRL	Events must be in time periods/luminosity blocks included in the 2015 or 2016 Good Runs Lists.
Good primary vertex	At least two tracks associated with the primary vertex.
Bad or cosmic muons	Events must not contains muons with $\sigma(q/p)/\lvert q/p\rvert > 0.2$ where this denotes the relative error on the charge-to-momentum ratio (*bad*); reject events with muon passing overlap removal with $\lvert z_0\rvert > 0.5$ mm and $\lvert d_0\rvert/\sigma(d_0) > 3$ (cosmic).
Bad jets	Reject 'fake' jets originating from non-collision sources, such as cosmic-ray showers and misreconstruction in the calorimeter. Cut on the pulse shapes of the calorimeter cell responses and the fraction of the jet energy in each calorimeter cell to distinguish these *good* jets [140].
Jet Vertex Tagger	Reject events when there is a jet with $60 < p_{\mathrm{T}} < 70$ GeV, $\lvert\eta\rvert < 2.4$, and JVT < 0.59 [141].
Bad Tile Cells	Reject events which are strongly affect by dead regions of the Tile calorimeter.
Additional LAr, Tile and SCT errors	Events with known errors reported in the LAr, Tile, or SCT are removed.

Table 8.1: A summary of the event cleaning cuts.

With the cleaning cuts now defined, the flow of data events from detector measurements to consideration in the multi-jets analysis region definitions (next Section) follows from the logical

AND of the following steps:

1. Events pass the Good Runs Lists selection.
2. Events pass the multi-jet trigger selection, a local OR of HLT_5j70 and HLT_6j45_0eta240.
3. Events pass event cleaning.

Following this initial set of selections and cleaning cuts, the physics objects which then pass overlap removal are used to select signal jets, leptons, and photons as described in Section 8.3.3.

8.3.3 Physics Objects Selection

Each event is reconstructed using the standard ATLAS reconstruction techniques described in Section 4.4, and the reconstruction of the $R = 0.4$ jets from EMTopo inputs is as described in Chapter 5. The physics object definitions used in this analysis are summarised in Table 8.2. Following the initial event selection and cleaning, there are two different sets of physics objects considered in the reconstruction: *baseline* and *signal*. The baseline objects are used to determine overlap removal decisions and to calculate the $E_{\mathrm{T}}^{\mathrm{miss}}$, whilst the signal objects are used in the final analysis event selections and have additional selection criteria. The basic overlap removal and object definitions are implemented through the use of a centralised ATLAS software package which applies the latest physics object definitions and recommendations as prescribed by the various combined performance groups.

An overlap removal algorithm is applied in order to prevent a particular reconstructed physics object being used twice in an event by removing it based on geometrical considerations. Muons and electrons are treated as independent objects, but photons and taus are treated as jets. The procedure adopted is applied sequentially, as follows:

1. If an electron and a muon share an inner detector track, the electron is removed.
2. If a photon has a ΔR(photon, electron) or ΔR(photon, muon) < 0.4, then it is removed.
3. Any jet which fails the loose b-tagging (80 % efficiency working point) selection is removed if:

   - ΔR(jet, electron) < 0.2; or

- the jet has no more than three tracks with $p_{\mathrm{T}} > 500$ MeV, or contains an inner detector track matched to a muon such that the transverse momentum of the jet is less than the muon and the muon track has more than 70 % of the sum of the transverse momenta of all tracks in the jet.

4. If ΔR(jet, electron (muon)) < 0.4, using the jets which survive requirement 3, then the electron (muon) is removed.

5. If ΔR(jet, photon) < 0.4, then the photon is removed.

Physics object	Selections
Baseline $R = 0.4$ jets	JVT > 0.59 [141] $p_{\mathrm{T}} > 20$ GeV $\lvert\eta\rvert < 2.8$
Baseline electrons	`LooseBLayer` ID [142] $p_{\mathrm{T}} > 10$ GeV $\lvert\eta\rvert < 2.47$
Baseline muons	`Medium` ID [64] $p_{\mathrm{T}} > 10$ GeV $\lvert\eta\rvert < 2.7$
Baseline photons	`Tight` ID [143] $p_{\mathrm{T}} > 25$ GeV $\lvert\eta\rvert < 2.37$
Signal $R = 0.4$ jets	Passes baseline overlap removal $p_{\mathrm{T}} > 50$ GeV or $p_{\mathrm{T}} > 80$ GeV $\lvert\eta\rvert < 2.0$
Signal electrons	Passes baseline overlap removal `Tight` ID [142] $p_{\mathrm{T}} > 25$ GeV `GradientLoose` isolation [144] $\frac{d_0}{\sigma_{d_0}} < 5$ $z_0 < 0.5$ mm
Signal muons	Passes baseline overlap removal $p_{\mathrm{T}} > 25$ GeV `GradientLoose` isolation [144] $\frac{d_0}{\sigma_{d_0}} < 3$ $z_0 < 0.5$ mm

Table 8.2: A summary of the basic physics objects used in the multi-jets analysis, together with their baseline and signal selections. Signal and baseline electrons, muons, and jets are used; additionally baseline photons are used in the overlap removal and $E_{\mathrm{T}}^{\mathrm{miss}}$ calculation.

In addition, large-radius $R = 1.0$ jets are also constructed by reclustering all $R = 0.4$ jets in the event, according to the procedure outlined in Section 5.3.2.2. b-jets are also considered, and

these are defined as signal jets which have also been identified with the MV2C10 algorithm [65], which uses a multi-variate discriminant to identify jets which contain a B-hadron, and is designed to give a b-jet selection efficiency of 70 % in simulated $t\bar{t}$ events.

8.3.4 Kinematic Variables

Following the physics object reconstruction, the signal objects are used in the calculation of important kinematic variables which help to define the final regions of the analysis. There are four variables which are of particular interest.

$\mathbf{H_T}$: the scalar sum of the p_T of all $R = 0.4$ jets in the event with $p_T > 40$ GeV and $|\eta| < 2.8$:[2]

$$H_T = \sum p_T^{jet}. \tag{8.3}$$

In the leptonic control regions (formally defined in Section 8.3.5), leptons which meet the same kinematic criteria as the jets are also used in the calculation of H_T.

$\mathbf{E_T^{miss}/\sqrt{H_T}}$: this is referred to in this analysis as the E_T^{miss} significance of the event, and is used to define the signal and control regions, and to estimate the data-driven QCD multi-jets background. The power of this variable rests on the observation that its shape is approximately invariant under any changes in the jet multiplicity, in events where the E_T^{miss} is dominated by jet mismeasurements (as would be the case for high jet multiplicity, multi-jet events[3]). This variable provides a metric for how much of the E_T^{miss} measured in an event is dominated by sources of real E_T^{miss} (such as neutrinos or neutralinos) and sources of fake E_T^{miss} (jet mismeasurement in the calorimeter): larger values of the E_T^{miss} significance are expected from events which are dominated by real soures of E_T^{miss}. This feature of the E_T^{miss} significance will be important for defining the analysis signal regions in Section 8.3.5. This definition of E_T^{miss} significance is not the true significance, which should be a dimensionless quantity, but we can start from this true significance in order to better understand the invarance of the $\frac{E_T^{miss}}{\sqrt{H_T}}$ shape. Assuming a Gaussian E_T^{miss} resolution, $\sigma(E_T^{miss})$, the true significance is a monotonic function of

$$\frac{(E_T^{miss})^2}{2\sigma(E_T^{miss})^2}. \tag{8.4}$$

[2]This selection is based on studies carried out in the Run 1 version of this analysis [134].
[3]See Appendix C for the full set of analysis $E_T^{miss}/\sqrt{H_T}$ plots.

We are interested in scenarios where the $E_{\mathrm{T}}^{\mathrm{miss}}$ is dominated by jet mismeasurement, and therefore the dominant component of the $\sigma(E_{\mathrm{T}}^{\mathrm{miss}})$ comes from the jet resolution:

$$\sigma(E_{\mathrm{T}}^{\mathrm{miss}})^2 \propto \sum \sigma(p_{\mathrm{T}}^{\mathrm{jet}})^2. \tag{8.5}$$

The jet resolution is obtained from the calorimeter energy resolution, and therefore can be approximated by $\sigma(p_{\mathrm{T}}^{\mathrm{jet}}) \propto \sqrt{p_{\mathrm{T}}^{\mathrm{jet}}}$ over a wide range of calorimeter energies. Recalling the definition of H_{T}, it now follows that

$$\sigma(E_{\mathrm{T}}^{\mathrm{miss}})^2 \propto H_{\mathrm{T}}. \tag{8.6}$$

Inserting this result into our original expression for the (Gaussian) true significance, we see that such a true significance may be approximated as a function of

$$\frac{(E_{\mathrm{T}}^{\mathrm{miss}})^2}{H_{\mathrm{T}}}, \tag{8.7}$$

which is equivalent to the definition $\frac{E_{\mathrm{T}}^{\mathrm{miss}}}{\sqrt{H_{\mathrm{T}}}}$ which we started out with.

$\mathbf{m_{\mathrm{T}}}$: this transverse mass variable is used when defining the leptonic control regions, and is calculated from the lepton p_{T} (p_{T}^{l}), the $E_{\mathrm{T}}^{\mathrm{miss}}$, and the ϕ-angle between the $E_{\mathrm{T}}^{\mathrm{miss}}$ and the lepton, $\Delta\phi(E_{\mathrm{T}}^{\mathrm{miss}}, p_{\mathrm{T}}^{l})$. It is defined as:

$$m_{\mathrm{T}} = \sqrt{2p_{\mathrm{T}}^{l}E_{\mathrm{T}}^{\mathrm{miss}}\Big(1 - \cos(\Delta\phi(E_{\mathrm{T}}^{\mathrm{miss}}, p_{\mathrm{T}}^{l}))\Big)}. \tag{8.8}$$

If a parent particle, P, undergos a decay into visible and invisible daughters, then the m_{T} distribution has a kinematic end-point at the rest mass of P. This makes m_{T} effective at reducing signal contamination in the control regions by requiring some upper value, $m_{\mathrm{T}} < x$ GeV. Since typical new physics particles will have larger m_{T}-values than SM particles.

$\mathbf{M_{J}^{\Sigma}}$: this event-level jet mass is the total mass of all $R = 1.0$ (reclustered) jets in the event with $p_{\mathrm{T}} > 100$ GeV and $|\eta| < 1.5$, and is defined as [145]:

$$M_{\mathrm{J}}^{\Sigma} = \sum_{j} m_{j}^{R\,=\,1.0}. \tag{8.9}$$

This variable is designed to be sensitive to the production of heavy flavour background, larger gluino masses, and to also work as a discriminating variable to reject $t\bar{t}$ backgrounds.

8.3.5 Region Definitions

The kinematic regions of interest are split into four categories: *signal*, *control*, *template*, and *validation* regions. The signal regions are designed to be rich in signal events, whilst having a much smaller contribution coming from Standard Model processes. The control regions are designed to be dominated by events coming from specific backgrounds, typically the largest Monte Carlo simulated backgrounds in the analysis, so that the Monte Carlo predictions for such backgrounds can be tested and appropriately normalised to match the data. Both the template and validation regions are designed to be dominated by the largest background in the analysis, that of Standard Model multi-jet production. Template regions are used to measure the SM multi-jet contribution in data, and therefore predict the size of this background in the signal regions, whilst the validation regions assess how well the multi-jet prediction closes with respect to the data, and also provides a measure of the uncertainties associated with the estimate of the multi-jet background from the template region. Let us briefly motivate the kinematic requirements in each of these regions.

Signal regions: Signal regions veto events which contain at least one electron or muon; such leptons may be produced with neutrinos, generating sources of real $E_{\mathrm{T}}^{\mathrm{miss}}$, which would otherwise contaminate the signal region. In order to be sensitivie to signals which do produce real $E_{\mathrm{T}}^{\mathrm{miss}}$, such as through the LSP production, a requirement of $E_{\mathrm{T}}^{\mathrm{miss}}/\sqrt{H_{\mathrm{T}}} > 5$ GeV$^{1/2}$ is imposed. This also serves to remove multi-jet events where the jets have been mismeasured in the calorimeter, generating fake $E_{\mathrm{T}}^{\mathrm{miss}}$. The long decay chains associated with the signal models produces final states with typically 7 or more hadronic jets, so we also require that there be at least 7, 8, 9, 10, or 11 jets in the final state. The signal regions are further subdivided into flavour and fat-jet streams. In the flavour stream, final states containing at least 0, 1, or 2 b-jets are considered in order to be sensitive to the b-jets which can be produced in these long decay cascades; in the fat-jet stream, the M_{J}^{Σ} variable is calculated from $R = 1.0$ jets (reclustered from the 7 or more $R = 0.4$ jets) in the event, and two regions are constructed for each jet multiplicity, requiring $M_{\mathrm{J}}^{\Sigma} > 340$ GeV or > 500 GeV. Such regions enhance the sensitivity to heavy flavour final states, such as top quark production, and the production of heavier gluinos. Jets in the flavour stream are required to have a threshold of $p_{\mathrm{T}} > 50$ GeV or 80 GeV, where the 5-jet and 6-jet triggers which fire on the events are close to fully-efficient. Jets in the fat-jet stream are required to pass a $p_{\mathrm{T}} > 50$ GeV threshold. For the $p_{\mathrm{T}} > 50$ GeV, jet multiplcities of at least 8, 9, 10, or 11

Selections	Flavour stream		Fat-jet stream
Lepton veto	No baseline e or μ remaining after overlap removal		
Jet $\lvert\eta\rvert$	< 2.0		
Trigger	HLT_6j45_0eta240	HLT_5j70	HLT_6j45_0eta240
Jet p_{T}	> 50 GeV	> 80 GeV	> 50 GeV
N_{jets}	$\geq \{8, 9, 10, 11\}$	$\geq \{7, 8, 9\}$	$\geq \{8, 9, 10\}$
$N_{b-\mathrm{jets}}$	$\geq \{0, 1, 2\}$		No selection
M_{J}^{Σ}	No selection		$\geq \{340, 500\}$ GeV
$E_{\mathrm{T}}^{\mathrm{miss}}/\sqrt{H_{\mathrm{T}}}$	> 5 GeV$^{1/2}$		

Table 8.3: A summary of the selections used to define the signal regions used in the multi-jets analysis.

are considered; for the 6-jet $p_{\mathrm{T}} > 80$ GeV regions, multiplicities of at least 7, 8, or 9 jets are considered. In the flavour stream, the combination of 4 $p_{\mathrm{T}} > 50$ GeV regions and 3 $p_{\mathrm{T}} > 80$ GeV regions, each with 0, 1, or 2 inclusive b-jets, totals to 21 signal regions. The inclusive 8, 9, and 10 $p_{\mathrm{T}} > 50$ GeV regions in the fat-jet stream, each with a 340 or 500 GeV requirements on M_{J}^{Σ}, adds an additional 6 signal regions, bringing the total up to 27. These selections are summarised in Table 8.3.

Control regions: Control regions are designed to be kinematically similar to signal regions, yet orthogonal in order to be rich in backgrounds and comparitively sparse in signal contamination. Orthogonality is achieved by requiring exactly one electron or muon in the event. To reduce statistical uncertainties, an inclusive N_{jets}-jet signal region is controlled by an inclusive ($N_{\mathrm{jets}} - 1$) jet control region. Similarly the requirement on $E_{\mathrm{T}}^{\mathrm{miss}}/\sqrt{H_{\mathrm{T}}}$ is also relaxed in order to increase the available statistics. There are two control regions for each signal region: the first with a b-jet veto applied, enhancing the contribution of W + jets backgrounds, and the second requiring at least a single b-jet in the event, enriching the region with $t\bar{t}$ background. The same signal region p_{T} thresholds also apply to the control regions. Leptons are also included as jets, for the purposes of selection, if they pass the same p_{T} and $\lvert\eta\rvert$ requirements as the jets in the event, in order to model the effects (on kinematic variables) of hadronic taus and jets being misidentified as leptons. Finally, $m_{\mathrm{T}} < 120$ GeV is also required in order to remove contributions from signal processes; this variable has a kinematic endpoint around m_W for leptonically-decaying Ws, but has no such bound when the lepton and $E_{\mathrm{T}}^{\mathrm{miss}}$ originate from different decays. The control regions are summarised in Table 8.4, and the different $E_{\mathrm{T}}^{\mathrm{miss}}/\sqrt{H_{\mathrm{T}}}$ thresholds are indicated in Table 8.5. In the control regions with higher jet multiplicities, this threshold is lowered in order

Lepton multiplicity	Exactly one signal e or μ remaining after overlap removal
Lepton p_{T}	> 20 GeV
m_{T}	< 120 GeV
Trigger	Same as signal regions
Jet p_{T}, $\lvert\eta\rvert$	Same as signal regions
N_{jets} (including lepton)	$N_{\mathrm{jets}}^{\mathrm{SR}} - 1$
$N_{b-\mathrm{jets}}$	$= 0$ (W + jets), ≥ 1 ($t\bar{t}$)
M_{J}^{Σ}	Same as signal regions
$E_{\mathrm{T}}^{\mathrm{miss}}/\sqrt{H_{\mathrm{T}}}$	Table 8.5

Table 8.4: A summary of the selections used to define the control regions used in the multi-jets analysis.

Signal stream	Control region N_{jets}		$E_{\mathrm{T}}^{\mathrm{miss}}/\sqrt{H_{\mathrm{T}}}$ selection
	Jet $p_{\mathrm{T}} > 50$ GeV	Jet $p_{\mathrm{T}} > 80$ GeV	
Flavour stream	7, 8	6	> 5 GeV$^{1/2}$
	9	7	> 4 GeV$^{1/2}$
	10	8	> 3 GeV$^{1/2}$
	$M_{\mathrm{J}}^{\Sigma} > 340$ GeV	$M_{\mathrm{J}}^{\Sigma} > 500$ GeV	
Fat-jet stream	7	-	> 5 GeV$^{1/2}$
	8	7	> 4 GeV$^{1/2}$
	9	8, 9	> 3 GeV$^{1/2}$

Table 8.5: A summary of the individual $E_{\mathrm{T}}^{\mathrm{miss}}/\sqrt{H_{\mathrm{T}}}$ selections for the control regions.

to increase the available statistics.

Template and validation regions: To extract a viable measurement of the shape of the multi-jet background a template region is defined. The signal regions associated with the p_{T} > 80 GeV threshold have corresponding template regions where exactly 5 jets and 0 leptons are required; the $p_{\mathrm{T}} > 50$ GeV signal regions have a corresponding 6 exclusive jet, 0 lepton template region. After extracting the multi-jet shape in these 5- and 6-jet template regions, the prediction (based on these template shapes) is tested in exclusive 6-jet and 7-jet validation regions respetively. We shall describe these regions, and the template background estimation procedure, in more detail in Section 8.5.

A shorthand notation is adopted for labelling the different regions in this analysis: first a `SR-`, `CR-`, `TR-`, or `VR-` is prepended for a signal, control, template, or validation region respectively. Following this, a piece is added of the form `<multiplicity><i/e><threshold>`, where `<multiplicity>` refers to the jet multiplicity, the `i` or `e` denote whether the jet

multiplicity selection is `inclusive` or `exclusive`, and the `<threshold>` corresponds to the jet p_{T} threshold of 50 or 80 GeV. The final component is stream specific: for the flavour stream, `<multiplicity><i/e>b` is appended, where the `<multiplicity>` corresponds to the b-jets in the event. For the fat-jet stream, `MJ<threshold>` is appended to correspond to the two M_{J}^{Σ} thresholds at 340 and 500 GeV. So for example, `SR-9ij50-0ib-MJ340` corresponds to a signal region with at least 9 jets, each with $p_{\mathrm{T}} > 50$ GeV, at least 0 b-jets in the event, and $M_{\mathrm{J}}^{\Sigma} > 340$ GeV. Or in the flavour stream, `SR-8ij80-2ib` corresponds to a signal region with at least 8 jets, with $p_{\mathrm{T}} > 80$ GeV, and at least two b-tagged jets in the event.

8.3.6 Jet Reclustering and Jet Substructure for Top-Tagging

Models such as the $\tilde{g}t\bar{t}$ produce final states with top quarks, where the mass-splitting between the LSP and the gluino can be large, resulting in boosted top production. Such boosted heavy flavour could be captured using large-radius jets, and the substructure information gleaned from these jets could be utilised to make a heavy-flavour tagger. In the multi-jets analysis large-radius jets have been investigated in order to tag heavy-flavour production. Two approaches have been considered:

- *Fat-jet stream*: here the $R = 0.4$ jets are reclustered into $R = 1.0$ jets, and the total mass of the reclustered jets in the event is used to increase the sensitivity to the production of heavy flavour, such as top jets. The M_{J}^{Σ} variable has already been described in Section 8.3.4.

- *Substructure top-tagging*: here `LCTopo` trimmed $R = 1.0$ jets are reconstructed in the multi-jet events, and the mass and τ_{32} substructure variables are used to construct a top-tagger and identify top quarks produced in $\tilde{g}t\bar{t}$-like scenarios.

As has already been stated, the final analysis uses the fat-jet stream approach in order to increase sensitivity to heavy flavour production, but this was only fully adopted after performing a direct comparison of the expected sensitivity to $\tilde{g}t\bar{t}$ signals obtained from the substructure top-tagger. For completeness, this sensitivity study shall now be briefly summarised.

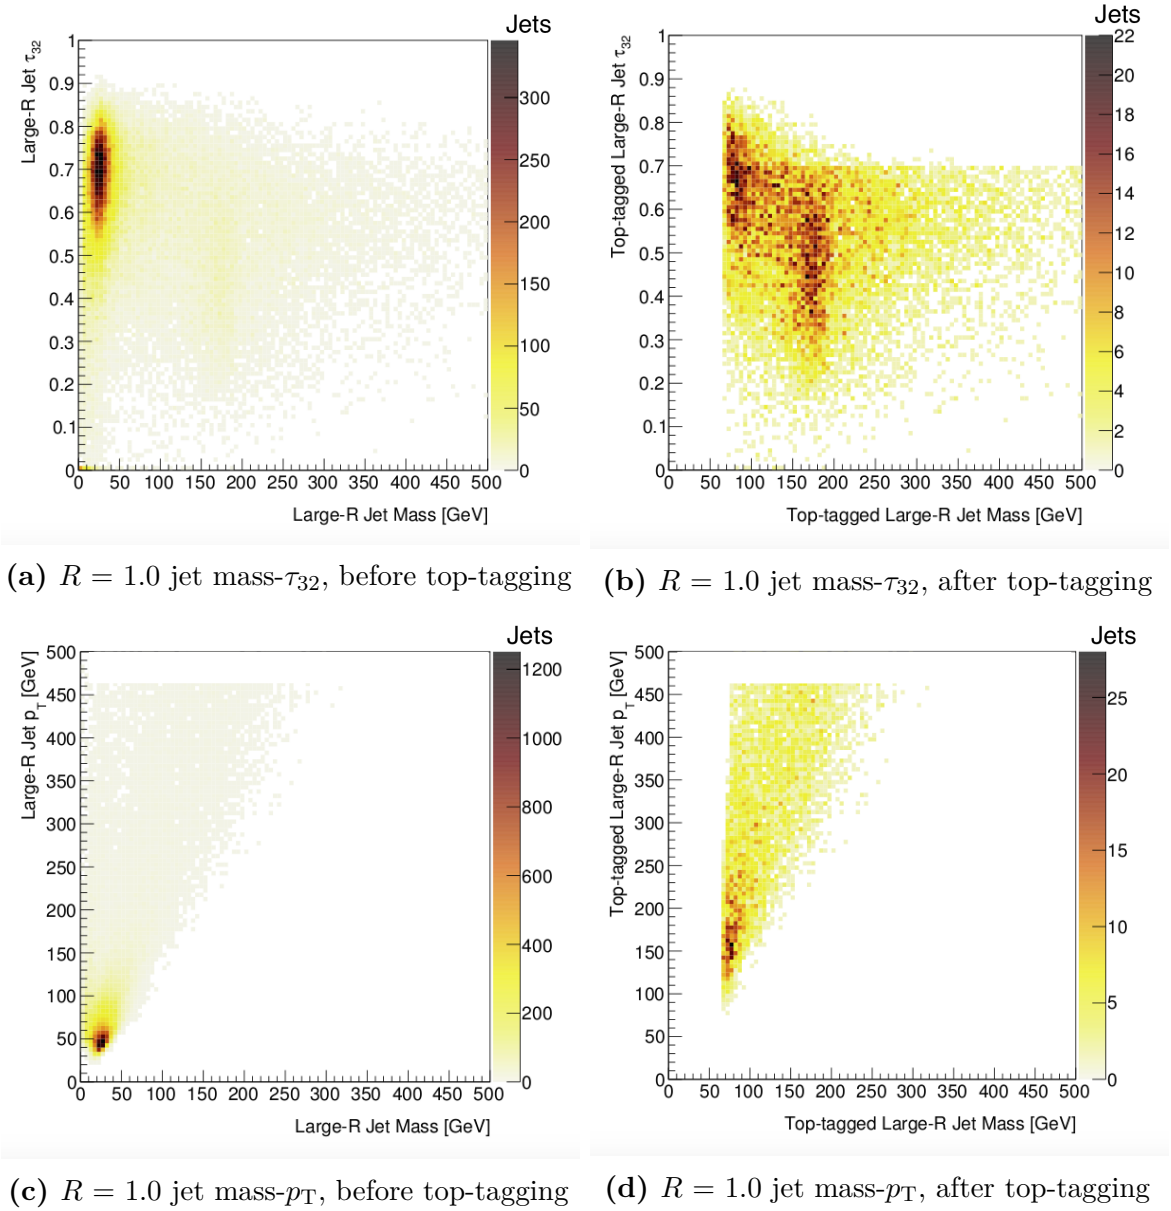
8.3.6.1 Top-Tagging in Multi-Jets

The approach of top-tagging described below is an expolaratory study carried out by the author, and, although it was not used in the final analysis, it was important in establishing the power of event-level variables in searches for new physics.

Top-tagging using jet-level substructure methods has been investigated substantially during Run 1 and Run 2 of ATLAS [132]. There are several different approaches to top-tagging within and outside of ATLAS, ranging from sophisticated machine-learning and Boosted Decision Tree taggers [146], to simpler taggers which select top-like candidates from a selection based on a few jet substructure variables. The latter are often referred to as cut-based taggers [132]. For the multi-jets analysis, it is necessary for the tagging efficiency to be high (70 % or higher) in order to preserve as much of the SUSY signal as possible, and consequently sophisticated machine-learning taggers are not a feasible option due to their low signal efficiency (40 – 50 % [132, 146]). This motivates the use of a simple, cut-based tagger which allows for the higher tagging efficiencies, at the cost of inferior background rejection.

For these studies a standard cut-based ATLAS top-tagger is adopted [132]. When given an input jet collection, this tagger selects each jet and uses the jet p_{T} in a lookup table which, for a given jet p_{T}, identifies the optimum cuts on the jet mass and jet τ_{32} (see Section 7.3 for the definition), which give the best background rejection for some fixed tagging efficiency provided by the user. An 80 % tagging efficiency working point is considered in these studies.

Figure 8.4 shows the correlation between jet mass and τ_{32}, and also the jet mass and p_{T}, before and after the standard top-tagger is applied to a sample $\tilde{g}t\bar{t}$ signal point, where `LCTopo` inputs have been reconstructed into `AntiKt10LCTopoTrimmedPtFrac5SmallR20` jets. The tagging has the expected effect of preserving jets which occupy the peak about the lower τ_{32}-value, whilst removing those which occupy the much higher τ_{32} peak. For larger values of p_{T} the tagger also correctly keeps jets which have an invariant mass around that of the top quark mass.



(a) $R = 1.0$ jet mass-τ_{32}, before top-tagging

(b) $R = 1.0$ jet mass-τ_{32}, after top-tagging

(c) $R = 1.0$ jet mass-p_{T}, before top-tagging

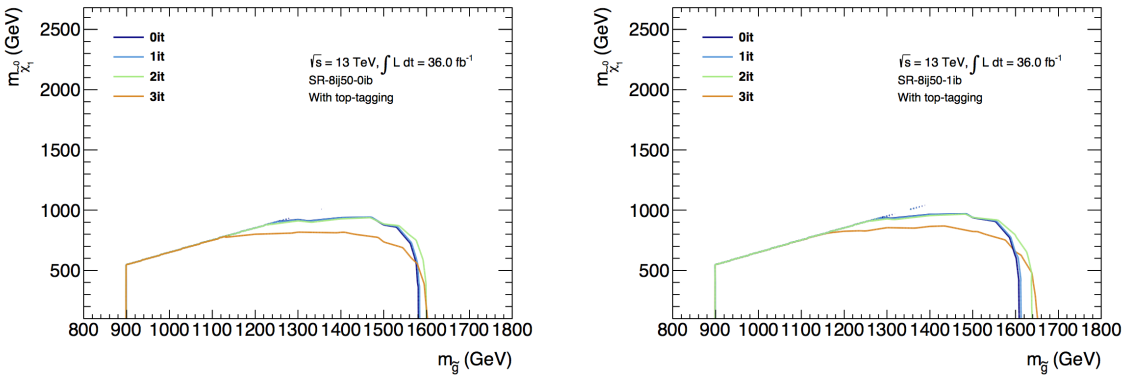
(d) $R = 1.0$ jet mass-p_{T}, after top-tagging

Figure 8.4: A summary of two-dimensional kinematic variable distributions for $R = 1.0$ jets before and after top-tagging. A $\tilde{g}t\bar{t}$ simulated signal point is considered. The large reduction in jets can be seen after applying the substructure cuts implemented in the top-tagger.

8.3.6.2 Expected Sensitivities from Top-Tagging and M_{J}^{Σ}

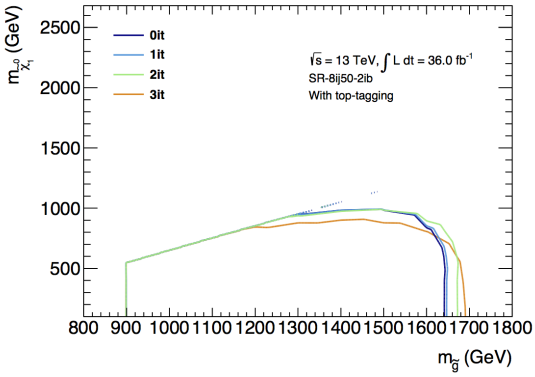
For the final sensitivity studies, the expected signal significance[4] is calculated for several different $\tilde{g}t\bar{t}$ signal mass points, assuming a dominating $t\bar{t}$ and QCD multi-jets background. New top-tagged signal region definitions have been constructed from the observation of reasonable discrimination between signal and background for final states with approximately 3 or more top-tagged jets.

[4]The contours correspond to values where the signal Z significance ($\frac{X-\mu}{\sigma^2}$) is equal to 1.64, where it is assumed that the data yield matches the simulated SM prediction.

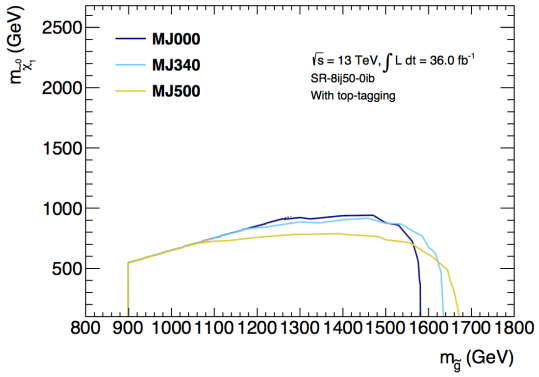


(a) `8ij50-0ib` + top-tagging

(b) `8ij50-1ib` + top-tagging



(c) `8ij50-2ib` + top-tagging



(d) `8ij50-0ib-MJ340`

Figure 8.5: The expected significance for `8ij50-0/1/2ib-0/1/2/3it` for the $\tilde{g}t\bar{t}$ model and the expected significance for the `8ij50-0ib-MJ340` region.

The impact of the top-tagging on the sensitivity to the gluino and LSP mass is quantified by first defining additional regions where there are top-tagged large-R jets. Four regions are considered: `0it`, `1it`, `2it`, and `3it` for 0, 1, 2, 3 inclusive top-tagged large-R jets respectively. A final summary of the *expected* signal sensitivity from top-tagging is presented in Figures 8.5 for a region where at least 8 small-radius jets with $p_{\mathrm{T}} > 50$ GeV, and zero leptons, is required. Similar behaviour is observed in regions with other large jet multiplicities. These plots should be compared with the expected sensitivity to the same $\tilde{g}t\bar{t}$ signal grid point achieved by applying jet reclustering and the M_{J}^{Σ} cuts described in Section 8.3.4 . This is also shown in Figure 8.5, again with the usual requirement of zero leptons and at least 8 small-radius jets in the final state.

Comparing the different sensitivities obtained from top-tagging demonstrates that a very similar reach is obtained from the separate top-tagging and M_{J}^{Σ} selections, where the latter approach provides a greater sensitivity to larger gluino masses. This is an interesting and

important lesson for searches targeting strongly-produced SUSY scenarios. Techniques such as top-tagging, which use substructure variables applied to individual jets, perform best for searches which target final states where the jets are highly boosted and the final state structure is simple, typically consisting of 1 or 2 jets, and perhaps a lepton. Such final states are typical among ATLAS exotics searches, where conventional top-tagging is utilised frequently. This is not the case for the models of SUSY considered in the multi-jets analysis; here the final states contain many jets, and therefore, for the same collision energy, the average jet momentum is lower and the jet consequently less boosted. For these less boosted jets a better variable to use is *not* a jet-level one, but an event-level variable. For example, the M_{J}^{Σ}, which is the event-level jet mass, is seen to be more sensitive to the (moderately) boosted $\tilde{g}t\bar{t}$ gluinos than the jet-by-jet N-subjettiness and mass variables of the top-tagger. The M_{J}^{Σ} variable encodes information about the mass of the entire event, and this acts as an effective proxy for heavy flavour production, which is typically massive, but not necessarily strongly boosted. Such event-level information therefore gives M_{J}^{Σ}, and similar variables, an advantage over conventional jet-level substructure information.

8.4 Leptonic Background Estimation

Standard Model processes which produce multi-jet events with one or more leptons are catagorised as *leptonic* backgrounds. Although a lepton veto is applied to the signal regions, it is still possible for leptonic backgrounds to contaminate them. The lepton veto only applies to electrons and muons, hence hadronically decaying tau leptons (originating from top or W bosons) contribute. Hadronically decaying taus are treated as jets, and the real $E_{\mathrm{T}}^{\mathrm{miss}}$ from the associated neutrino will also enter the signal regions. Additionally, there are cases where the electons or muons may fall out of the detector acceptance, or be misreconstructed in the detector, generating more $E_{\mathrm{T}}^{\mathrm{miss}}$. Finally, it is also possible for $E_{\mathrm{T}}^{\mathrm{miss}}$ originating from the $Z \rightarrow \nu\nu$ decay to enter our signal regions. With these effects established, it is important to estimate the contamination of leptonic backgrounds in the signal region; this is achieved through Monte Carlo simulation of such backgrounds. The two dominating Monte Carlo backgrounds are leptonically and semileptonically decaying $t\bar{t}$, and leptonically-decaying W bosons which are produced in association with jets (referred to as the 'W + jets' background). Smaller backgrounds include the production of leptonically-decaying Z bosons and associated jets ('Z + jets'), single top production, diboson production of WW, WZ, or ZZ, and the production of $t\bar{t}$ in association with heavier particles − this final background covers the three-top, four-top,

$t\bar{t}+W$, $t\bar{t}+Z$, $t\bar{t}+H$, and $t\bar{t}+WW$ topologies, and is collectively referred to as '$t\bar{t}X$'. These additional backgrounds are comparatively very small and are estimated in the signal regions through Monte Carlo simulation alone.

SM Background	Matrix Element	Parton Shower
$t\bar{t}$	Powheg	Pythia 6.428
W/Z + jets	Sherpa 2.2.1	Sherpa 2.2.1
Single top quarks	Powheg	Pythia 6.428
Diboson	Sherpa 2.1.1	Sherpa 2.1.1
$t\bar{t}X$	MadGraph 5 v.2.2.2	Pythia 8.186
$t\bar{t}H$	MadGraph 5_aMC@NLO v.2.2.3	Pythia 8.186

Table 8.6: The Monte Carlo matrix element and parton shower details for the various simulated standard model backgrounds.

SM Background	PDF	Underlying Event Tune
$t\bar{t}$	CT10 [114] (matrix element), CTEQ6L1 [114] (showering)	Perugia 2012 [147]
W/Z + jets	NNPDF30NNLO [114]	From Sherpa authors
Single top quarks	CT10 (matrix element), CTEQ6L1 (showering)	Perugia 2012
Diboson	CT10	From Sherpa authors
$t\bar{t}X$	NNPDF23LO [114]	A14 [115]
$t\bar{t}H$	NNPDF30NLO [114] (matrix element), NNPDF23LO (showering)	A14

Table 8.7: The Monte Carlo PDF and underlying event tune details for the various simulated standard model backgrounds.

These leptonic backgrounds, and their associated Monte Carlo generators, are summarised in Tables 8.6 and 8.7. For all background simulations, the ATLAS detector response is modelled using the full simulation of GEANT 4 [112] (a fast simulation [148], based on a simpler parameterisation of the ATLAS electromagnteic and hadronic calorimeter, is used to model the signal). Leptonic $t\bar{t}$ and single top are simulated using Powheg [149] interfaced to Pythia 6.428 [150] for partonic showering. The production of W/Z + jets, and dibosons, is modelled using Sherpa 2.2.1 and Sherpa 2.1 [95] respectively. With one exception, the $t\bar{t}X$ backgrounds are simulated using MadGraph 5 v.2.2.2 [116] interfaced with Pythia 8.186 [113, 115] for the parton shower; $t\bar{t}H$ is instead simulated using MadGraph 5_aMC@NLO v.2.2.3 [116], again interfaced with Pythia 8.186. In additional, the Monte Carlo events are reweighted in order to account for:

- Differences between the pileup distributions as measured in ATLAS data and the expected pileup profile simulated in the Monte Carlo.

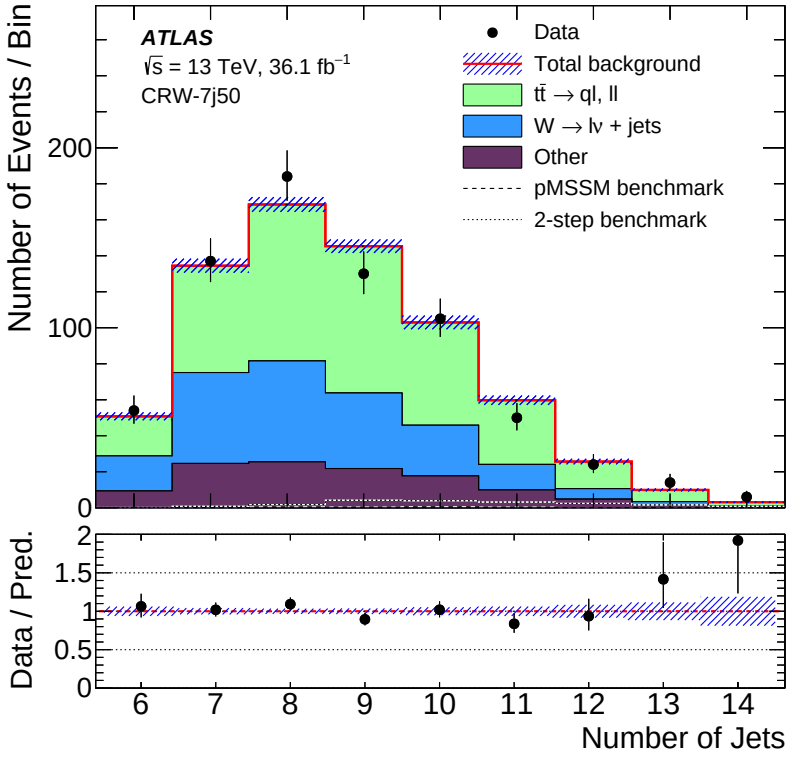
- Differences between the efficiencies of b-jets, electrons, and muons measured in data to those simulated in Monte Carlo [65].

The contributions from the two largest leptonic backgrounds are first estimated using Monte Carlo events in the dedicated $t\bar{t}$ and W + jets control regions defined in Section 8.3.5 and Table 8.4. A global fit of the simulated backgrounds to data is carried out in these control regions (see Section 9.1.1 for details) in order to obtain normalisation factors which are applied to the leptonic backgrounds in the signal regions to improve the background prediction there. The control regions must be rich in the two largest leptonic backgrounds in order for the normalisations from the fit to be derived with a good statistical precision. Each signal region is controlled by two control regions according to the nomenclature outlined in Section 8.3.5. For example, the `SR-8ij50-0ib` region is controlled by the `CR1l-7ij50-0eb` and `CR1l-7ij50-1ib` for W + jets and $t\bar{t}$ backgrounds respectively. The distributions of the number of jets in these two control regions is shown in Figure 8.6 after the application of the normalisations from the fit to data. The two largest simulated backgrounds can be clearly seen (blue corresponds to W + jets, green to $t\bar{t}$), and the uncertainties shown are statistical. The agreement between data and Monte Carlo, post-fit, is very good and provides confidence in our approach of extracting background estimates in such kinematic regions of phase-space. It should be noted that the final normalisations which are derived incorporate the effects of theoretical, statistical, and systematic uncertainties. The treatment of these uncertainties is described in Section 8.6.

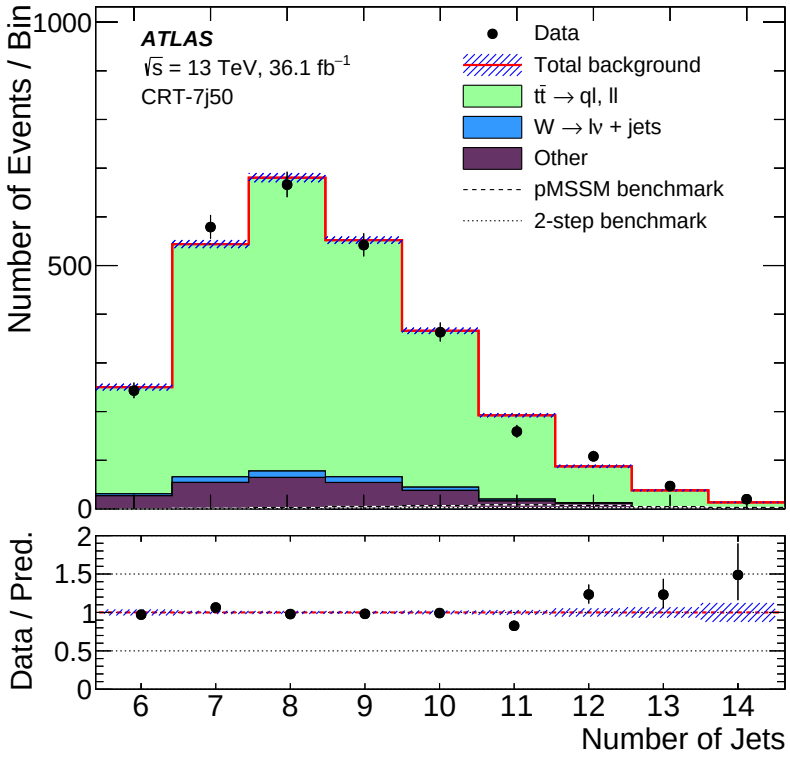
8.5 QCD Multi-Jet Background Estimation

The dominant background in this analysis comes from Standard Model QCD multi-jets events. It would be expected that such final states would have little or no $E_{\mathrm{T}}^{\mathrm{miss}}$, however detector resolution effects can lead to mismeasurement of the jets, generating sources of fake $E_{\mathrm{T}}^{\mathrm{miss}}$. Due to a lack of reliable Monte Carlo for simulating such high jet multiplicity final states, a full data-driven approach must be adopted, which we shall refer to as the *template method*.

The template method is an 'ABCD' approach, where the $E_{\mathrm{T}}^{\mathrm{miss}}/\sqrt{H_{\mathrm{T}}}$ distribution is used to extract the background estimate. Let us first define a template region which corresponds to a $E_{\mathrm{T}}^{\mathrm{miss}}/\sqrt{H_{\mathrm{T}}}$ distribution which is fully measured in data. Two template regions are considered in order to estimate the multi-jets background in the $p_{\mathrm{T}} > 50$ GeV and $p_{\mathrm{T}} > 80$ GeV signal



(a) CR1l-7ij50-0eb



(b) CR1l-7ij50-1ib

Figure 8.6: The jet multiplicity ditributions using the selection for the W + jets and $t\bar{t}$ control regions. Normalisations from the background-only fit have also been applied.

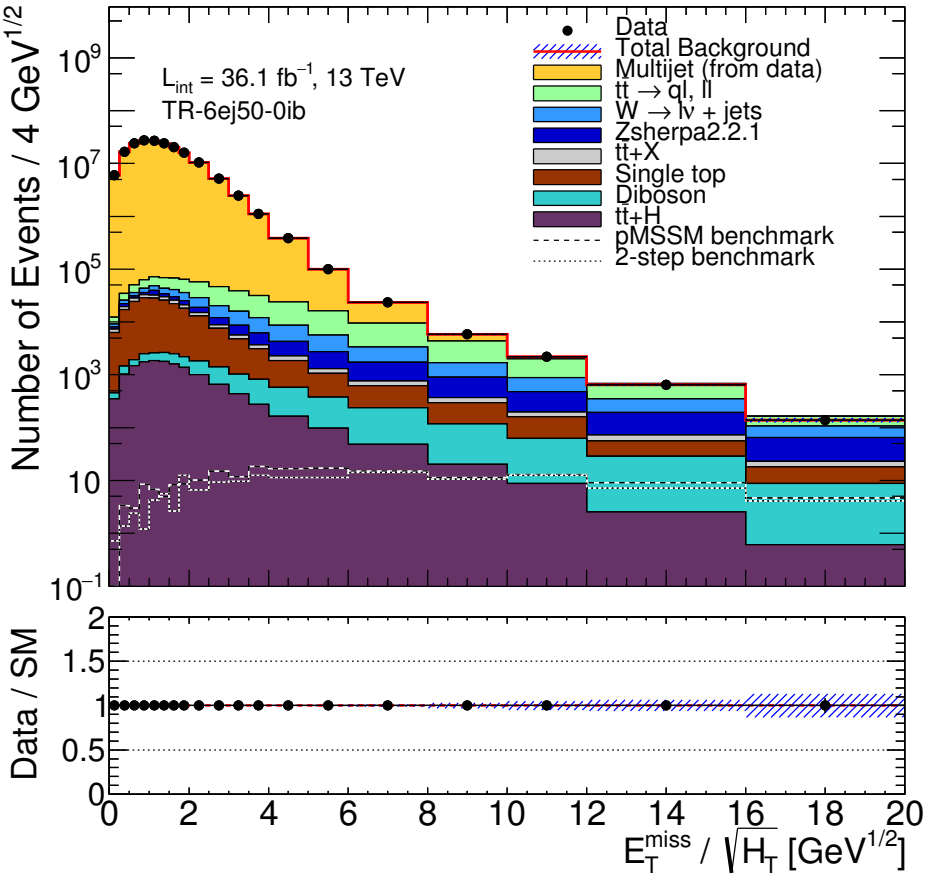


Figure 8.7: $E_{\mathrm{T}}^{\mathrm{miss}}/\sqrt{H_{\mathrm{T}}}$ distribution for events in the region of exactly 6 jets, each with a jet transverse momentum of 50 GeV or greater, and 0 inclusive b-tagged jets. The lower panel shows the ratio of the observed data to the total standard model background, including the statistical uncertainty in the hatched blue. This ratio is, by construction, unity.

regions separately: these are denoted `TR-6ej50-*` and `TR-5ej80-*`, where `*` is a placeholder for additional kinematic cuts which depend on the particular signal region of interest. An example template region is shown in Figure 8.7 and corresponds to `TR-6ej50-0ib` i.e. exactly 6 jets, with $p_{\mathrm{T}} > 50$ GeV, and 0 inclusive b-tagged jets in the event. The various Monte Carlo backgrounds are overlaid, together with the data (dark circles), and the large orange contribution comes from the QCD multi-jets background is also fully measured in data here. The two largest Monte Carlo backgrounds, semileptonic $t\bar{t}$ and W + jets, are normalised according to a background-only fit. The peak at lower $E_{\mathrm{T}}^{\mathrm{miss}}/\sqrt{H_{\mathrm{T}}}$ is dominated by the fake $E_{\mathrm{T}}^{\mathrm{miss}}$ contributions from jet mismeasurement in the calorimeter and leads to the large population of QCD multi-jets events.

At the heart of the template method is the approximate invariance of the shape of the $E_{\mathrm{T}}^{\mathrm{miss}}/\sqrt{H_{\mathrm{T}}}$ distribution as the jet multiplicity changes, which was justified in Section 8.3.4.

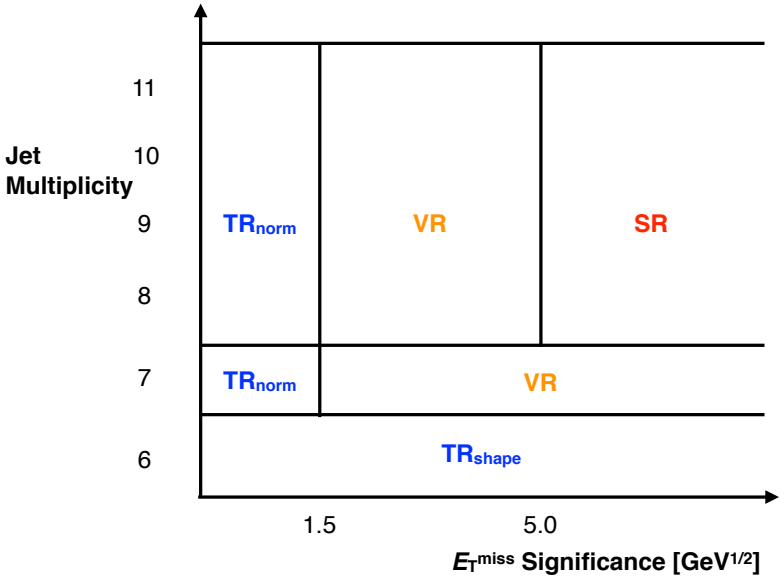
In the template region, the shape and event yield is measured at $E_{\mathrm{T}}^{\mathrm{miss}}/\sqrt{H_{\mathrm{T}}} < 1.5$ GeV$^{1/2}$, where the multi-jets background dominates, and at $E_{\mathrm{T}}^{\mathrm{miss}}/\sqrt{H_{\mathrm{T}}} > 5.0$ GeV$^{1/2}$, which is the same $E_{\mathrm{T}}^{\mathrm{miss}}/\sqrt{H_{\mathrm{T}}}$ cut used to define the signal regions. The $E_{\mathrm{T}}^{\mathrm{miss}}/\sqrt{H_{\mathrm{T}}}$ distribution is then considered for one of the signal regions, which is blinded at $E_{\mathrm{T}}^{\mathrm{miss}}/\sqrt{H_{\mathrm{T}}} > 5.0$ GeV$^{1/2}$. The shape and yield of events at $E_{\mathrm{T}}^{\mathrm{miss}}/\sqrt{H_{\mathrm{T}}} < 1.5$ GeV$^{1/2}$ is then extracted in that region in order to obtain a normalisation to the shape measured in the original template region. Denoting the yields in the < 1.5 GeV$^{1/2}$ and > 5.0 GeV$^{1/2}$ by N_A and N_C for the template region, and N_B and N_D for the signal regions, the invariance of the significance shape implies that $\frac{N_C}{N_A} = \frac{N_D}{N_B}$ holds, and therefore the multi-jet background estimate in the signal region is determined from:

$$N_D = N_B \frac{N_C}{N_A}. \tag{8.10}$$

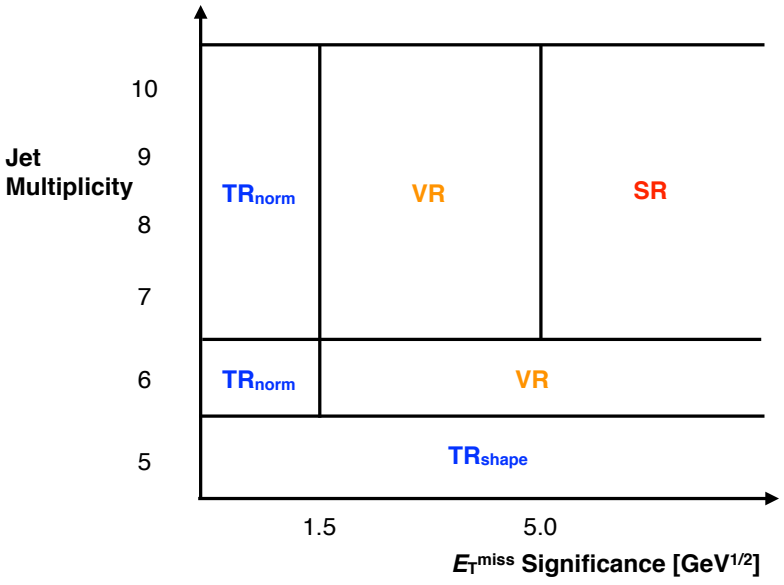
Before using this technique to estimate the multi-jet background in a signal region, it is first necessary to check that the prediction matches the data well, and this is achieved by testing the template prediction in *validation regions*. Predictions extracted from the `TR-6ej50-*` and `TR-5ej80-*` templates are validated in the `VR-7ej50-*` and `VR-6ej80-*` $E_{\mathrm{T}}^{\mathrm{miss}}/\sqrt{H_{\mathrm{T}}} \geq 1.5$ GeV$^{1/2}$ regions respectively, and also in the 1.5 GeV$^{1/2} \leq E_{\mathrm{T}}^{\mathrm{miss}}/\sqrt{H_{\mathrm{T}}} < 5.0$ GeV$^{1/2}$ region of the signal region under consideration. As a concrete example, the template method is illustrated for the `SR-8ij50-0ib` and `SR-7ij80-0ib` regions in the schematic in Figure 8.8, showing the relationship between the template, validation, and signal regions in $E_{\mathrm{T}}^{\mathrm{miss}}/\sqrt{H_{\mathrm{T}}}$ − jet multiplicity space.

For the case of `SR-8ij50-0ib`, the event yield is measured in the `TR-6ej50-0ib` $E_{\mathrm{T}}^{\mathrm{miss}}/\sqrt{H_{\mathrm{T}}}$ < 1.5 GeV$^{1/2}$ to be $N_{6j}^{\mathrm{data}}(1.5)$ and in the $E_{\mathrm{T}}^{\mathrm{miss}}/\sqrt{H_{\mathrm{T}}} > 5.0$ GeV$^{1/2}$ to be $N_{6j}^{\mathrm{data}}(5.0)$. This is directly measured from data and will therefore include all backgrounds; the total Monte Carlo background estimates in these regions are denoted $N_{6j}^{\mathrm{MC}}(1.5)$ and $N_{6j}^{\mathrm{MC}}(5.0)$. The background estimate is now tested in the `VR-7ej50-0ib` region, using a normalisation derived from the events with $E_{\mathrm{T}}^{\mathrm{miss}}/\sqrt{H_{\mathrm{T}}} < 1.5$ GeV$^{1/2}$, which shall be similarly denoted $N_{7j}^{\mathrm{data}}(1.5)$. Again the yield from the Monte Carlo-only background is $N_{7j}^{\mathrm{MC}}(1.5)$. Putting this together, the estimate of the multi-jet background in the validation region is:

$$N_{7j}^{\mathrm{multi-jet}}(5.0) = \frac{N_{7j}^{\mathrm{data}}(1.5)}{N_{6j}^{\mathrm{data}}(1.5)} \times \Big[N_{6j}^{\mathrm{data}}(5.0) - N_{6j}^{\mathrm{MC}}(5.0) \Big]. \tag{8.11}$$

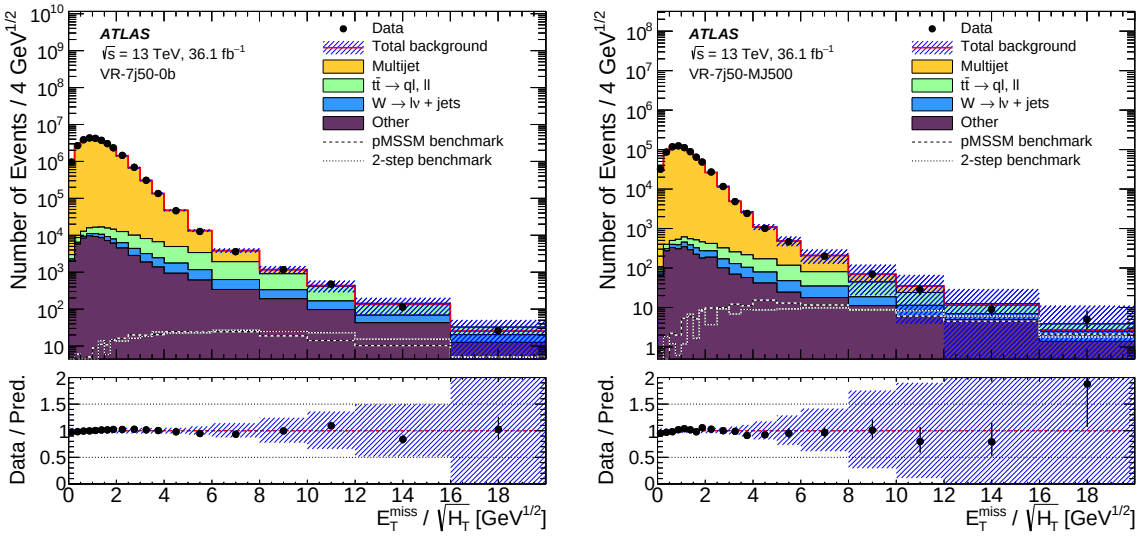


(a) Template method, $p_{\mathrm{T}} > 50$ GeV regions.



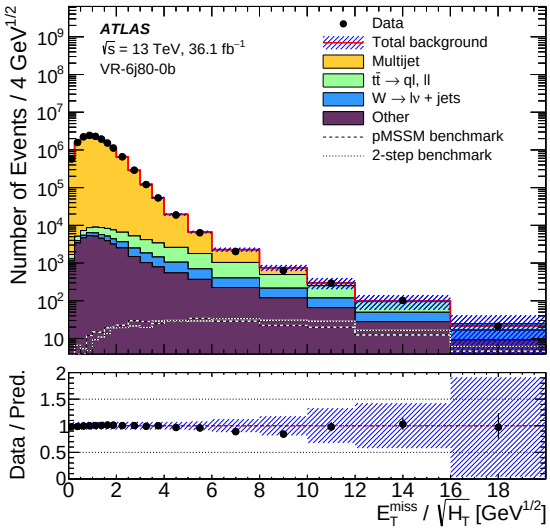
(b) Template method, $p_{\mathrm{T}} > 80$ GeV regions.

Figure 8.8: A schematic illustrating the template method for the two different p_{T} thresholds in multi-jets. The region where the background shape is extracted in data is denoted by 'shape', and the region where the normalisation is extracted, for rescaling the shape, is denoted by 'norm'.

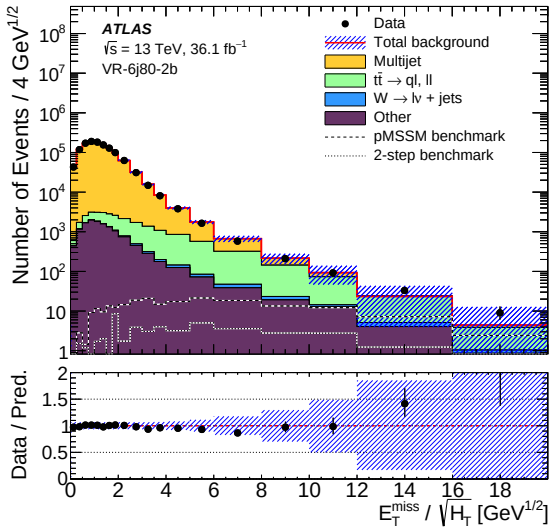


(a) VR-7ej50-0ib

(b) VR-7ej50-MJ500



(c) VR-6ej80-0ib



(d) VR-6ej80-2ib

Figure 8.9: Example exclusive 6- and 7-jet validation regions, used to assess the closure of the QCD background estimates.

A similar procedure follows for the multi-jet estimate in SR-8ij50-0ib:

$$N_{8j}^{\text{multi-jet}}(5.0) = \frac{N_{8j}^{\text{data}}(1.5)}{N_{6j}^{\text{data}}(1.5)} \times \Big[N_{6j}^{\text{data}}(5.0) - N_{6j}^{\text{MC}}(5.0)\Big]. \tag{8.12}$$

The $E_{\text{T}}^{\text{miss}}/\sqrt{H_{\text{T}}}$ distributions for a small selection of different validation regions are shown in Figure 8.9. The total background prediction, including the multi-jet from the template, is observed to close well with respect to the data, and is entirely covered within the statistical and systematic uncertainties (to be evaluated shortly).

Before finishing with the template method, one final detail needs to be addressed. It has been measured that the H_{T} in the event is correlated with the jet multiplicity [134]: events with a larger jet multiplicity should, on average, posses a larger measured H_{T}. Consequently there will always be *some* $E_{\mathrm{T}}^{\mathrm{miss}}/\sqrt{H_{\mathrm{T}}}$-dependence on jet multiplicity, despite the initial assumption of the invariance of its shape between different regions, However, it is possible to reduce this dependence by binning the template region in H_{T}, where lower bin boundaries of 0, 600, 900, and 1200 GeV have been found to be sufficient to remove the dependence of the template region shape on H_{T}. Another important factor affecting the shape of the template is the relative composition of heavy-flavour (such as b-quark initiated) jets in the final state. This effect is accounted for by assessing the flavour-dependence of the template method as a systematic uncertainly, and shall be discussed in detail in Section 8.6.3.

8.6 Treatment of Uncertainties

In addition to the statistical uncertainties there are also several sources of systematic uncertainty which need to be considered. The systematic uncertainties are divided in three classes: detector, theory, and the template systematics. The detector systematics are applied to the Monte Carlo backgrounds and arise from the detector simulation and physics object reconstruction effects on such backgrounds. The theory systematics are also applied to the Monte Carlo backgrounds and arise from differences between Monte Carlo generators, partonic showering, and radiation effects when simulating different backgrounds. Finally, the template method itself comes with its own set of systematic uncertainies, which are applied to the background prediction from this same method. These include uncertainties due to the non-closure of the template prediction with respect to data, the flavour-dependence of the template, and the choice of H_{T}-binning in the template calculation. These systematics enter into the fitting procedure (to be described in Sections 9.1.1 and 9.2.1) as Gaussian constraints with widths calculated from the variation of the nominal yield with the systematic of interest.

8.6.1 Detector Systematics

The detector systematics describe the uncertainties, applied to Monte Carlo only, which incorporate the effects of detector simulation and physics object reconstruction. Dominating detector systematics come from: uncertainties on b-tagging in the flavour stream signal regions, accounting for up to 4 % in some regions; uncertainties due to the measurement of the jet

energy scale and jet energy resolution, varying from $6 - 12$ % in certain signal regions; and an uncertainty in the simulation of the soft term of the $E_{\mathrm{T}}^{\mathrm{miss}}$, which has an effect of up to 8 % in the fat-jet stream. Additional uncertainties, such as those on the detector luminosity of the recorded data-set, are small (< 1 %). Briefly, the full list of detector systematics corresponds to the following:

Jet energy scale: The jet energy scale calibration, described in Chapter 5, corrects the energy of a jet, measured in both data and reconstruction-level Monte Carlo, to the energy of the original parton. Uncertainties due to differences in Monte Carlo modelling of the calibration, the physical modelling of different processes, and the jet flavour, trigger inefficiences, and pileup effects on calibrated data collectively give rise to 73 nuisance parameters. The correlations between these parameters can be reduced to produce a final set of 4 nuisance parameters [151], which are then varied independently in the final fit.

Jet energy resolution: The spread of the reconstructed jet energies, as measured in the detector, around the true values is characterised by the jet energy resolution. A single nuisance paramter is assigned for this uncertainty, which is evaluated by smearing the jet energies with a Gaussian function, where the width of the Gaussian is characterised by the difference between the energy resolution measured in data and measured in Monte Carlo.

$b-$**tagging**: b-tagging scale factors account for differences in the b-tagging efficiencies between data and Monte Carlo, and these differences come with associated uncertainties [152]. These scale factor uncertainties are propagated through the final fit. The procedure of b-tagging in ATLAS is briefly described in Section 4.4.4.

$\mathbf{E}_{\mathrm{T}}^{\mathrm{miss}}$ **soft term**: The soft component of the $E_{\mathrm{T}}^{\mathrm{miss}}$ has uncertainties associated with its scale and resolution. Up and down variations on the $E_{\mathrm{T}}^{\mathrm{miss}}$ scale, and the smearing of the soft term in order to account for the differences between data and Monte Carlo, are used to assess the magnitude of this uncertainty.

Electron and photon energy scale, **and resolution**: Two nuisance parameters are applied to cover these uncertainties: the first covers the scale uncertainty on the electron and photon

response, and the second covers the total resolution uncertainty. Effects of correlations between uncertainties are neglected.

Electron selection efficiency: Uncertainties arising from electron reconstruction, identification, and isolation selections are accounted for by varying the electron scale factors.

Muon energy scale and resolution: A single set of up and down variables covers the uncertainty on the absolute muon energy scale, which again quantifies the difference between the energy as measured in data and Monte Carlo. The tracks in the inner detector and muon spectrometer are also independently varied according to the resolutions of those detectors; the effects of this on the combined track measurement is calculated in order to account for the resolution uncertainties.

Muon selection efficiency: Additional muon scale factors are varied to account for the uncertainty on their measurements and selections, and the statistical and systematic components of this uncertainty are treated separately.

Pileup reweighting: The pileup reweighting scale factor, used to correct the Monte Carlo pileup distribution to its measurement in data, is varied upwards and downwards in order to account for the uncertainties on the overlaid pileup distribution in the Monte Carlo.

8.6.2 Theory Systematics

The theory systematics assess the uncertainties coming from the matrix element generation and the subsequent showering and decay of the initial particles produced in the matrix element, as well as additional uncertainties coming from radiative emissions of the decaying particles. It was found that the largest theoretical uncertainty comes from the simulation of QCD radiation from the top quarks in $t\bar{t}$ events, amounting to an uncertainty of up to 25 % of the total $t\bar{t}$ background prediction. The uncertainties due to partonic showering and additional radiative effects were found to be smaller (or comparable) in size. The calculation of the theory systematics is split into three major Monte Carlo backgrounds: $t\bar{t}$, W + jets (the largest backgrounds), and the remaining minor backgrounds for which theory systematics were not explictly calculated but

rather global, conservative estimates are applied.

Leptonic and semileptonic $t\bar{t}$: Uncertainties from the generation of the matrix element are assessed by comparing the predictions from two different generators; the nominal sample is produced using POWHEG for the matrix element generation, and PYTHIA 6.428 for the partonic showering. A sample using MADGRAPH 5_aMC@NLO in place of POWHEG provides the variation for the matrix element generation uncertainty. Additionally, alternative samples using PYTHIA 8.186 and HERWIG++ [97] in place of PYTHIA 6.428 provide the variation for the partonic showering uncertainty. Finally, the uncertainty arising from radiative emissions (i.e. gluon emissions above and below the nominal expected rates) are investigated using two variants of the nominal Monte Carlo settings which increase/decrease the rate of QCD radiation overlaid on the base physics process; these are referred to as RADHI and RADLO respectively.

W + jets: The W + jets backgrounds are generated using SHERPA 2.2.1 and come with their own set of Monte Carlo weights which correspond to different variations of the Monte Carlo renormalisation, factorisation, resummation, and matching scales [153]. The renormalisation and factorisation scales, μ_R and μ_F respectively, set the energy scales for truncating the calculation of the matrix element coupling constant and parton distribution functions [154], such that divergences are factored out. These scales are varied according to the following sets of configurations:

$$\{\mu_R, \mu_F\} = \{0.5, 0.5\}, \{0.5, 1.0\}, \{1.0, 0.5\}, \{1.0, 1.0\}, \{1.0, 2.0\}, \{2.0, 1.0\}, \{2.0, 2.0\}, \quad (8.13)$$

where $\{x, y\}$ refers to a varying the μ_R and μ_F scales by factors of x and y of the nominal scales. The resummation scale corresponds to the scale at which the calculation of the parton shower is truncated, and the matching scale sets the energy at which a given emission is treated as part of the matrix element or as part of the partonic shower. In order to determine the variations which generate the final Monte Carlo weights, these scales are varied between $1/\sqrt{2}$ and $\sqrt{2}$ of the magnitude of the nominal scale.

Minor backgrounds: The remaining backgrounds predict comparatively small numbers of events in the signal regions, as such they do not require a detailed calculation of the-

ortical uncertainties. Rather, global and conservative uncertainties[5] on each of these minor backgrounds are applied:

- Z + jets: 40 %
- diboson: 50 %
- single top: 30 %
- $t\bar{t}X$: 30 %

8.6.3 Template Systematics

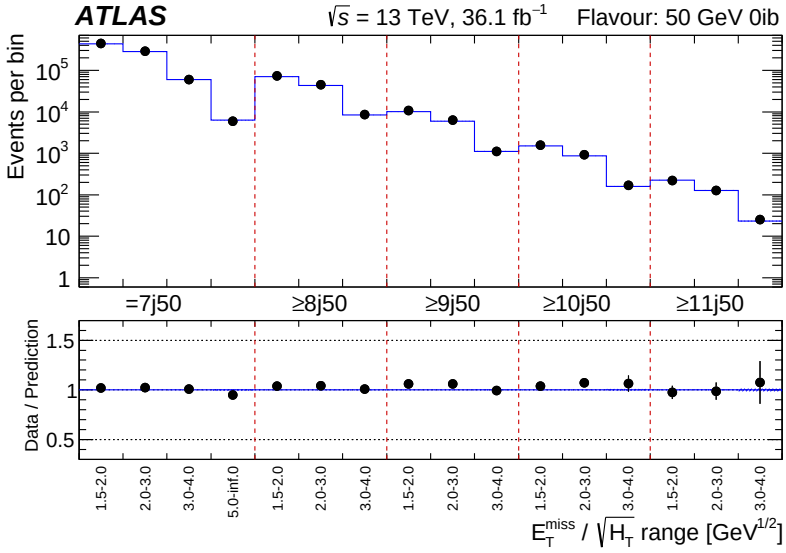
The three uncertainties contributing to the total template systematic originate from the non-closure of the template prediction with respect to data, the flavour-dependence of the template arising from the semileptonic decays of B hadrons, and the effects of H_{T} binning on the template.

8.6.3.1 Template non-closure uncertainty

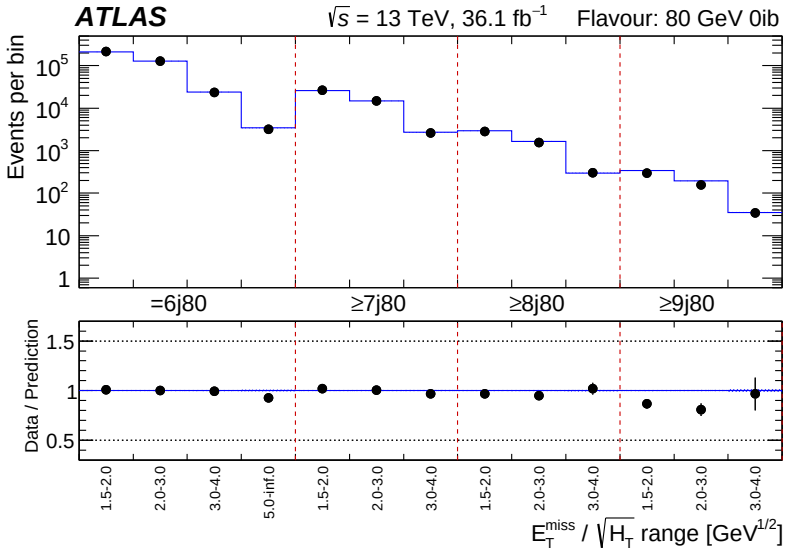
The largest systematic uncertainty on the template prediction arises from its non-closure: the degree to which the prediction deviates from data in regions where little signal is expected. To quantify this non-closure, the differences between the prediction and data are assessed in a series of validation regions. The closure of the prediction is checked against the number of observed events in regions with $E_{\mathrm{T}}^{\mathrm{miss}}/\sqrt{H_{\mathrm{T}}} \in (1.5, 2.0)$, $(2.0, 2.5)$, and $(2.5, 3.5)$ $\mathrm{GeV}^{1/2}$, for ≥ 8, ≥ 9, and ≥ 10 jet regions; it is also checked in the $E_{\mathrm{T}}^{\mathrm{miss}}/\sqrt{H_{\mathrm{T}}} > 5.0$ $\mathrm{GeV}^{1/2}$ exclusive 7-jet validation region. The closure is defined as the difference between the total background prediction and the observed number of events in data in each region. For a given signal region the symmetrical systematic uncertainty applied is the maximal deviation in any of the closure regions of the same jet multiplicity or lower, for the same M_{J}^{Σ} and/or b-jet requirement. The non-closure is therefore evaluated separately for each of the different b-jet multiplicities and jet p_{T} thresholds in the flavour stream, and separately for the two different M_{J}^{Σ} thresholds in the fat-jet stream.

Example degrees of non-closure of the total background estimation is provided in Figure 8.10 for the `0ib` regions in the $p_{\mathrm{T}} > 50$ GeV and $p_{\mathrm{T}} > 80$ GeV threshold flavour stream, and the $M_{\mathrm{J}}^{\Sigma} > 340$ GeV in the fat-jet stream. Across the different signal regions, the non-closure uncertainty varies from 7.40 % to 29.70 %, and is generally larger in regions with high inclusive jet multiplicity.

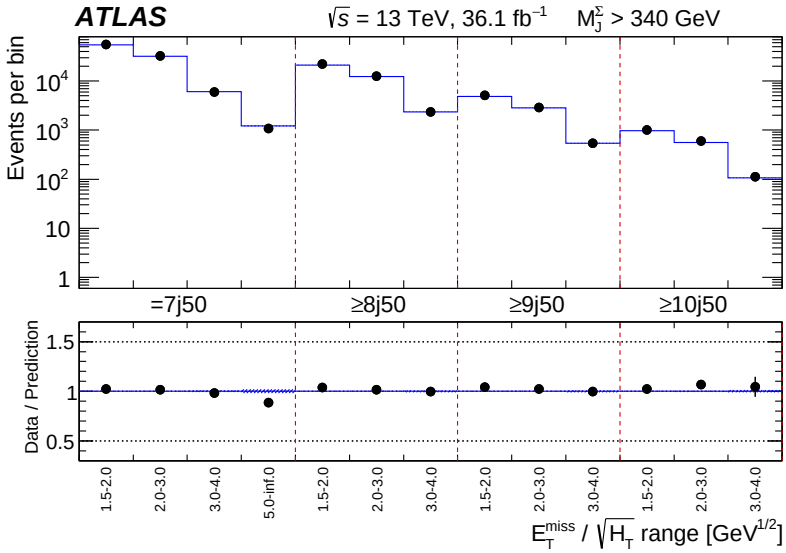
[5] These uncertainty estimates are based on Run 1 studies [134].



(a) Non-closure, $p_{\mathrm{T}} > 50$ GeV, `0ib`.



(b) Non-closure, $p_{\mathrm{T}} > 80$ GeV, `0ib`.



(c) Non-closure, $p_{\mathrm{T}} > 50$ GeV, `0ib-MJ340`.

Figure 8.10: The non-closures in the background estimation for several different signal regions in the flavour and fat-jet streams.

8.6.3.2 Template flavour uncertainty

The usual approach to building the template region requires no selection on the number of b-jets. However, events with b-jets are more likely to have real $E_{\mathrm{T}}^{\mathrm{miss}}$ due to the semileptonic decays of B-hadrons, which would have the effect of broadening the $E_{\mathrm{T}}^{\mathrm{miss}}/\sqrt{H_{\mathrm{T}}}$ distribution. As this broadening is not accounted for in the template estimation, it is necessary to derive an uncertainty which covers its impact on the $E_{\mathrm{T}}^{\mathrm{miss}}/\sqrt{H_{\mathrm{T}}}$ distribution.

In order to evaluate the size of this effect on the template prediction two separate template regions are constructed (let us take the example of the template prediction for the $p_{\mathrm{T}} > 50$ GeV regions): the first from a selection requiring exactly 6 jets and no b-jets (`6ej50-0eb`), and the second from a selection requiring exactly 6 jets and at least 1 b-jet (`6ej50-1ib`). A similar procedure is applied to the 5-jet template used when considering the $p_{\mathrm{T}} > 80$ GeV multi-jet prediction. These two contributions are added to give a 'flavour-split' yield, F_s, the total yield from the template prediction and the leptonic backgrounds. This F_s estimate is then added in a linear combination with the same yield expected from the original (`0ib`) 'flavour-blind' template, F_b:

$$\mathcal{T} = fF_b + (1-f)F_s, \tag{8.14}$$

where f is constrained to lie between 0 and 1. In fact, because there are three b-jet multiplicities considered in this analysis, it is necessary to construct three different types of flavour-blind and flavour-split templates:

- `0ib` flavour-blind, `0eb` + `1ib` flavour-split (as above).
- `1ib` flavour-blind, `1eb` + `2ib` flavour-split.
- `2ib` flavour-blind, `2eb` + `3ib` flavour-split.

Independently for each b-jet multiplicity, this background estimate $\mathcal{T}$ calculation is performed in the same validation regions used to assess the template non-closure uncertainty. In each of these regions, the optimal f value is determined through a χ^2 minimisation with respect to the E observed data:

$$\chi^2 = \frac{(E - \mathcal{T})^2}{\sigma^2}, \tag{8.15}$$

where σ is the sum in quadrature of the statistical uncertainties in the data and total background prediction. The uncertainty on f is determined by varying f until the agreement in the validation regions has deteriorated sufficiently to create an increase in the χ^2 of the data − prediction comparison of 1. Equivalently, Δf comes from the change in f necessary to modify each minimum χ^2-value by the amount $\Delta\ \chi^2 = 1$. Once the optimum f value has been determined in the validation region corresponding to a given signal region, the systematic uncertainty in the $E_{\mathrm{T}}^{\mathrm{miss}}/\sqrt{H_{\mathrm{T}}} > 5.0$ GeV$^{1/2}$ signal region is evaluated as the difference between the mixed template and the flavour-blind template, i.e. $\mathcal{T}$(optimum f) − $\mathcal{T}(f = 1)$. This systematic is generally subdominant compared to the non-closure systematic, and is typically largest in the higher jet multiplicity, higher b-jet multiplicity signal regions: the per signal region uncertainty varies from < 1 % to 19.10 %.

8.6.3.3 Template H_{T} binning uncertainty

As already stated, the template estimate is binned in H_{T} in order to factorise out low-H_{T} effects which distort the $E_{\mathrm{T}}^{\mathrm{miss}}$ resolution, other than jet mismeasurement. The choice of H_{T} binning is based on the quality of the closure of the template, and a good choice is found to be H_{T} = (0, 600), (600, 900), (900, 1200), (1200, ∞) GeV. An alternative binning of (0, 600), (600, 800), (800, 1000), (1000, 1200), (1200, 1400), (1400, ∞) GeV is considered, and the difference between the two predictions is taken as a (small) systematic uncertainty. This difference is evaluated across all valdiation and signal regions and is found to amount to typically no more than a couple of %, smaller than the non-closure and flavour-dependence systematics.

8.6.4 Systematics Summary

Before moving to the unblinded results in the next Chapter, we close by presenting the summary of the various systematics considered across the different analysis signal regions, shown in Figure 8.11. The uncertainties are split into the statistical, theory, experimental (detector), and template components. The theoretical uncertainties tend to dominate the total uncertainty in each signal region, however there are also substantial contributions from the experimental and template uncertainties, particularly at higher jet multiplicities. It should be noted that the total

uncertainty in each region is in fact calculated through the background-only fit, to be described in Section 9.1.1, and correlations between the various systematics are accounted for in doing so.

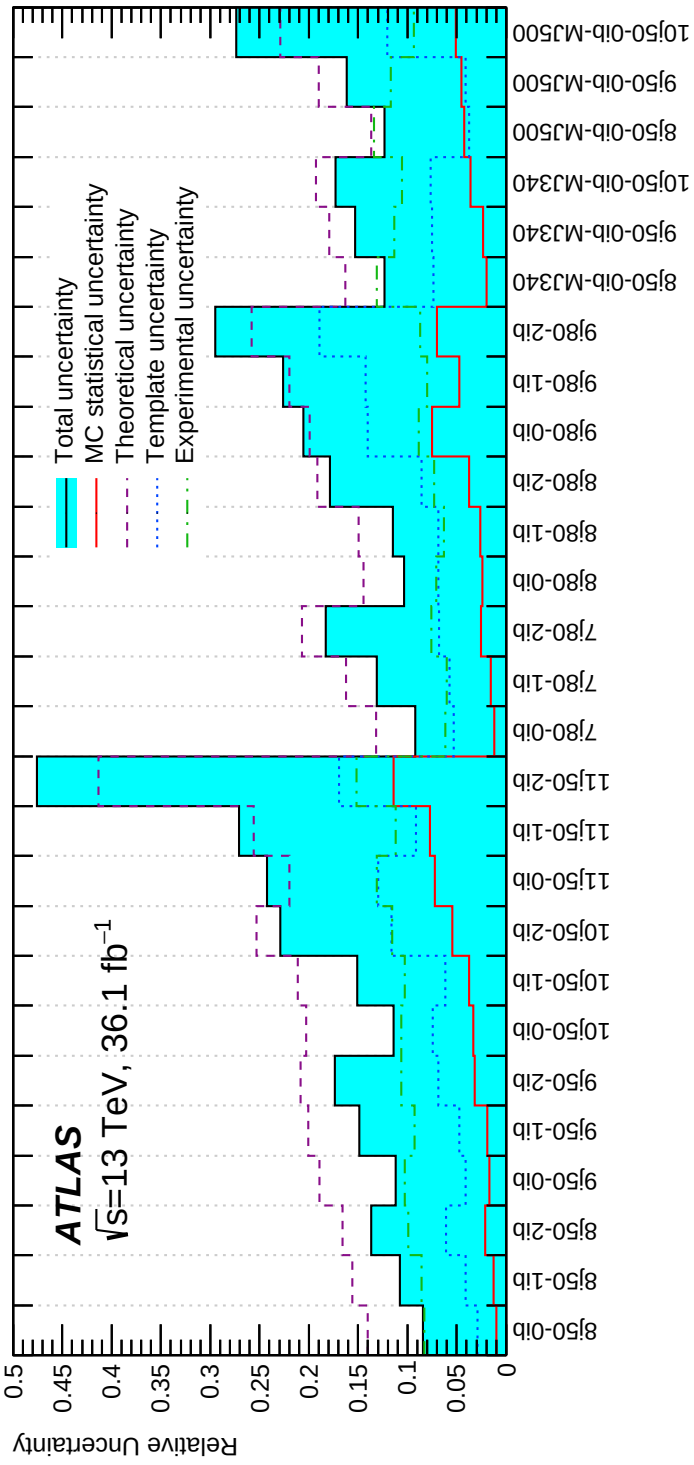


Figure 8.11: A summary of the various systematic (and statistical) uncertainties across the 27 analysis signal regions.

9

Constraining new phenomena in multi-jet final states

We now turn to the unblinded results of the multi-jets analysis, using the 2015 and 2016 ATLAS data-set. In addition, interpretations of various simplified models, and the pMSSM slice, are also discussed and the latest constraints on the masses of gluinos, neutralinos, and stops generated by these different signals are presented. We shall also determine the expected and observed sensitivities to several other R-parity-violating models of SUSY, by making use of the analysis sensivity to both moderate- and low-$E_{\mathrm{T}}^{\mathrm{miss}}$ final states.

9.1 Results with 36 fb^{-1} of ATLAS Data

Before unblinding the signal regions and comparing the observed data with the total background prediction, the individual components of the total background, together with their uncertainties, are fitted in order to improve this final prediction. This is achieved using a background-only fit, which takes information from the $t\bar{t}$ and W + jets control regions and extracts normalisation factors, $\mu_{t\bar{t}}$ and $\mu_{W+\mathrm{jets}}$, which fit the backgrounds to data. A similar normalisation for the data-driven multi-jets estimate, $\mu_{\mathrm{multi-jets}}$ is also extracted using the template validation regions.

9.1.1 Background-Only Fit

The background-only fit is applied using the statistical techniques in the HistFactory software package [155], which is in turn steered using the ATLAS HistFitter package [135]. The analysis is first split into different *channels*, which correspond to different regions of phase-space

defined by specific selection criteria, namely a given signal region and the two associated control regions. Each channel is simply an event count of the different backgrounds (referred to as *samples*) and data. Since we are considering simple event counts, often in channels with few events, an appropriate probability density function for each channel, c, is the Poissonian:

$$\text{Pois}(n_c|\nu_c) = \frac{\nu_c^{n_c}}{n_c!}e^{-\nu_c}, \tag{9.1}$$

where ν_c and n_c are, respectively, the expected and observed number of events in channel c. The expected number of events is determined from the predicted backgrounds b i.e. the Monte Carlo predictions for $t\bar{t}$ and W + jets, and is given by

$$\nu_c = \sum_b \lambda_c \mu_b \eta_{b,c}(\alpha) S_{b,c}. \tag{9.2}$$

In the above, λ_c denotes the true value of the detector luminosity and μ_b are the various normalisations for the different backgrounds b (which we wish to obtain through this background-only fit). The $\eta_{b,c}(\alpha)$ represent the relative values of ν_c when a parameter is changed by some fraction α of its standard deviation. In this way, varying a parameter by a $\pm\sigma$ uncertainty corresponds to $\alpha = \pm 1$, and therefore $\eta_{b,c}(+1)$ and $\eta_{b,c}(-1)$ denote the 1σ up and down systematic variations respectively. This variable therefore directly incorporates systematic uncertainties into the fit. By definition, $\eta_{b,c}(0) = 1$, i.e. corresponds to the nominal value of ν_c. The $S_{b,c}$ corresponds to the (unscaled) expected number of events in channel c from the background b (per unit integrated luminosity): it is the product of the background b's total cross-section, its acceptance for falling into channel c, and the efficiency with which the detector reconstruction-level physics objects fall into channel c, given their truth-level properties.

In addition to this Poissonian, we also need to account for the measured detector luminosity, L_0, and its associated uncertainty, Δ_L. These values are determined from beam-separation scans (similar to what is described in [156]) and, for the total ATLAS 2015 and 2016 data-set, are measured to be 36.21 $\pm$ 0.77 fb^{-1}. The true luminosity, λ, is constrained by a Gaussian probability distribution, $\mathcal{G}(L_0|\lambda, \Delta_L)$, centred on the measured luminosity. This feature will also need to be added to our final likelihood for the background-only fit.

In this fit the values of α are allowed to float. Allowing for this can result in 'profiling' of the systematic uncertainties, whereby the uncertainties can be shifted to values with large $|\alpha|$ in order to give a better fit to the data in channel c. Since the calibrations and variations provided by the various performance groups are typically trusted at values not too far above $|\alpha| = 1$, it is important to ensure that profiling of the uncertainties is constrained to α-values close to what is recommended. This is achieved through the application of Gaussian constraints on α for each systematic variation p: $\mathcal{P}_p(a_p, \alpha_p)$, where the a_p are chosen such that $\mathcal{P}$ is centred on 0 with width 1.

Putting all of this together, we now arrive at the full expression for the likelihood function of the background-only fit:

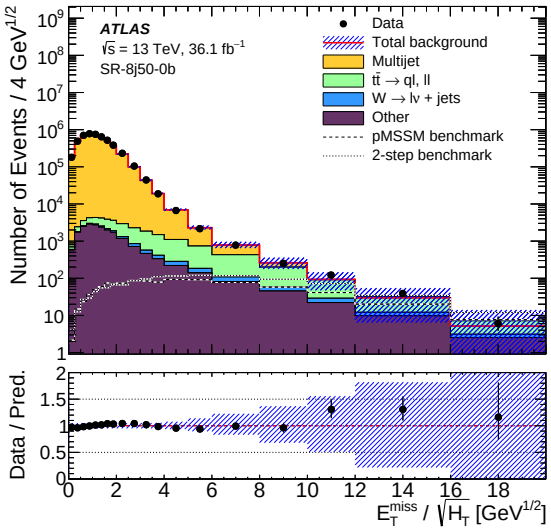
$$\mathcal{L}(\mu, \alpha|\mathbf{n}, \mathbf{a}) = \prod_{\mathbf{c}} \text{Pois}(\mathbf{n_c}|\nu_\mathbf{c})\mathcal{G}(\mathbf{L_0}|\lambda, \mathbf{\Delta_L}) \prod_{\mathbf{p}} \mathcal{P}_\mathbf{p}(\mathbf{a_p}, \alpha_\mathbf{p}), \tag{9.3}$$

where μ and $\mathbf{n}$ are vectors of all scale factors μ_b and observed events n_c in each channel c. By maximising this likelihood, $\mu_{W+\text{jets}}$ and $\mu_{t\bar{t}}$ can be extracted for each of the signal regions.

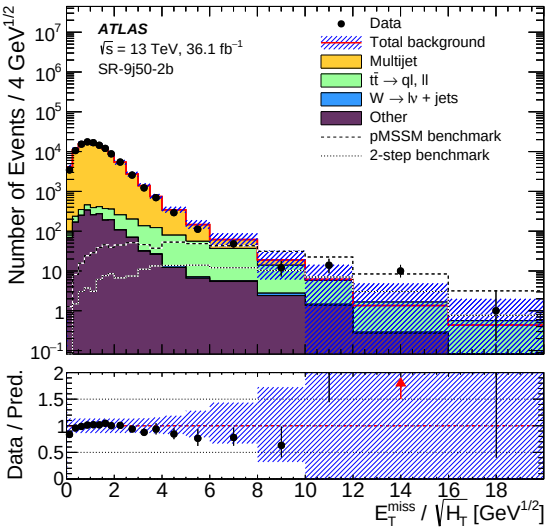
9.1.2 Unblinded Results

The leptonic background estimations, multi-jet template validation, and background-only fit evaluation of the μ normalisations and systematic uncertainties takes place while the analsysis is still blinded. That is to say, data points are not added to the $E_\text{T}^\text{miss}/\sqrt{H_\text{T}} > 5.0$ GeV$^{1/2}$ signal regions. Following the background-only fit, the normalisations are applied to the $t\bar{t}$ and W + jets backgrounds (and systematics) and the analysis is unblinded in the signal regions. Example unblinded $E_\text{T}^\text{miss}/\sqrt{H_\text{T}}$ distributions in the `SR-8ij50-0ib`, `SR-9ij50-2ib` and `SR-10ij50-0ib-MJ340` signal regions are shown in Figure 9.1[1]. The dominant multi-jet background prediction from the template method, as well as the Monte Carlo leptonic background, are overlaid with the data. Two example signal model mass points, for the two-step and pMSSM models, are also shown; these have larger event yields at higher values of $E_\text{T}^\text{miss}/\sqrt{H_\text{T}}$, as would be expected. Comparing the SM background with the data, it is clear that there is no statistically significant excess above the SM expectation in these signal regions (see Table 9.1).

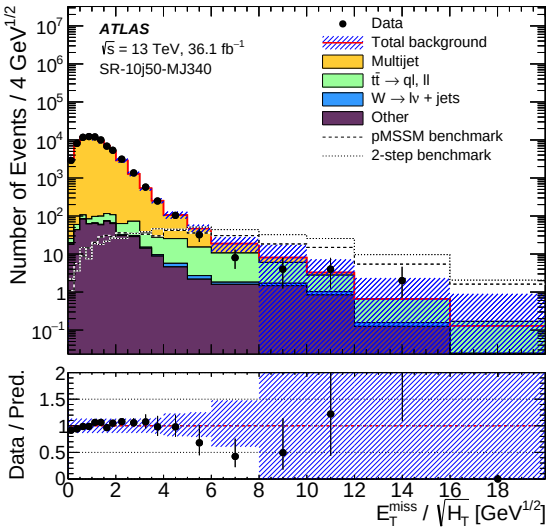
[1]For completeness, the $E_\text{T}^\text{miss}/\sqrt{H_\text{T}}$ distributions for the remaining signal regions are provided in Appendix C.



(a) SR-8ij50-0ib



(b) SR-9ij50-2ib



(c) SR-10ij50-0ib-MJ500

Figure 9.1: Example $E_{\mathrm{T}}^{\mathrm{miss}}/\sqrt{H_{\mathrm{T}}}$ distributions for three (unblinded) signal regions, where no statistically significant excesses in data (see Table 9.1), with respect to the injected two-step and pMSSM signals, are reported. These two example 'benchmark' signals have parameters $(m_{\tilde{g}}, m_{\tilde{\chi}_1^0}) = (1400, 200)$ GeV and $(m_{\tilde{g}}, m_{\tilde{\chi}_1^\pm}) = (1400, 200)$ GeV respectvely.

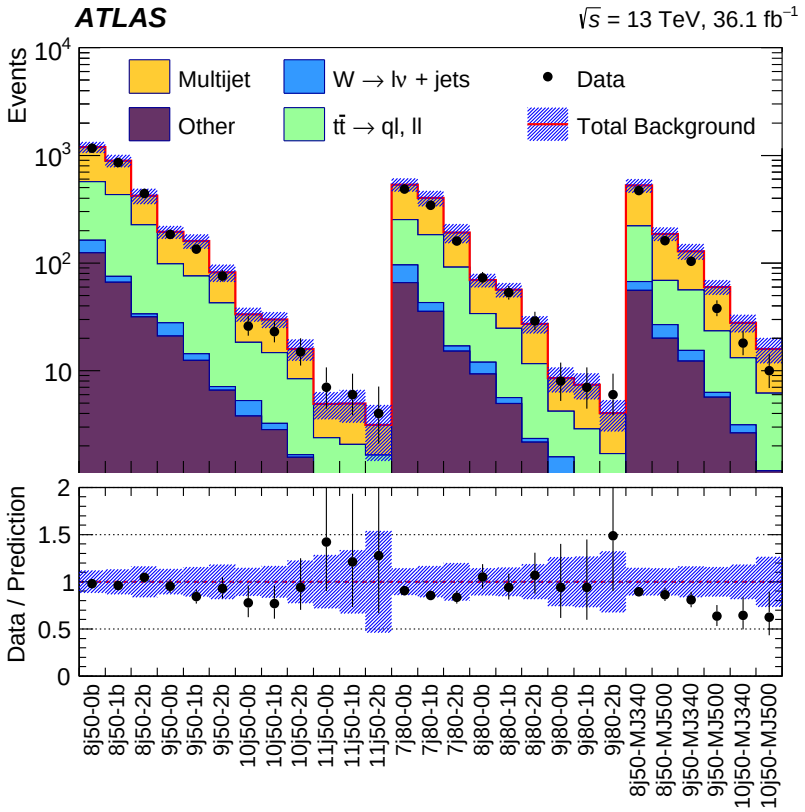


Figure 9.2: The event yields in each of the 27 signal yields, together with the expected SM backgrounds. No statistically significant excesses are reported.

The event yields in each of the 27 signal regions, following background normalisation, are summarised in Figure 9.2. The number of standard deviations from the SM expectation for the observed events in each region is indicated in the ratio on the lower pad. No statistically significant excess is observed in any of the signal regions, with the largest deviation being a 1.8σ deficit in `SR-9ij50-0ib-MJ500`. The total number of expected and observed events, per signal region, are listed in Table 9.1. With no statistically significant excesses observed, we may now set both model dependent- and model-independent limits on new physics processes, including the models described in Section 8.2. This shall form the focus of Section 9.2.

9.2 Interpretations with 36 fb^{-1} of ATLAS Data

9.2.1 Limit Setting

9.2.1.1 Model-independent Fits

In the model-independent fit the data in the signal region and the background prediction from the background-only fit are used to estimate the minimum number of physics events that *any*

Signal Regions	Multi-jet Events	Leptonic Events	Total SM Events	Observed Events
SR-8ij50-0ib	622 ± 42	570 ± 140	1190 ± 140	1169
SR-8ij50-1ib	460 ± 50	430 ± 110	890 ± 140	856
SR-8ij50-2ib	196 ± 39	226 ± 57	422 ± 81	442
SR-9ij50-0ib	96 ± 11	98 ± 24	194 ± 28	185
SR-9ij50-1ib	84 ± 15	76 ± 20	160 ± 31	135
SR-9ij50-2ib	39 ± 12	42.5 ± 9.5	82 ± 19	76
SR-10ij50-0ib	15.1 ± 3.0	18.3 ± 3.9	33.5 ± 5.1	26
SR-10ij50-1ib	15.3 ± 3.7	14.7 ± 3.3	30.0 ± 5.9	23
SR-10ij50-2ib	7.6 ± 3.1	8.4 ± 1.8	16.0 ± 4.2	15
SR-11ij50-0ib	2.54 ± 0.76	2.4 ± 1.2	4.9 ± 1.2	7
SR-11ij50-1ib	2.88 ± 0.84	2.1 ± 1.4	5.0 ± 1.3	6
SR-11ij50-2ib	1.49 ± 0.72	1.6 ± 1.5	3.1 ± 1.5	4
SR-7ij80-0ib	282 ± 32	253 ± 69	535 ± 74	486
SR-7ij80-1ib	219 ± 28	183 ± 60	402 ± 74	343
SR-7ij80-2ib	100 ± 17	91 ± 34	191 ± 44	160
SR-8ij80-0ib	35.7 ± 4.6	33.8 ± 8.3	70 ± 10	73
SR-8ij80-1ib	31.6 ± 5.7	24.8 ± 6.4	56 ± 10	53
SR-8ij80-2ib	15.5 ± 3.8	11.6 ± 3.3	27.1 ± 6.0	29
SR-9ij80-0ib	4.3 ± 1.3	4.2 ± 1.8	8.5 ± 2.0	8
SR-9ij80-1ib	4.5 ± 1.3	2.9 ± 1.5	7.4 ± 1.8	7
SR-9ij80-2ib	2.34 ± 0.95	1.69 ± 0.89	4.0 ± 1.2	6
SR-8ij50-0ib-MJ340	306 ± 54	220 ± 55	526 ± 72	471
SR-8ij50-0ib-MJ500	118 ± 18	69 ± 20	187 ± 54	161
SR-9ij50-0ib-MJ340	73 ± 15	56 ± 15	129 ± 23	104
SR-9ij50-0ib-MJ500	36.5 ± 6.3	23.3 ± 7.0	60 ± 10	38
SR-10ij50-0ib-MJ340	14.6 ± 3.8	13.2 ± 3.5	27.9 ± 5.7	18
SR-10ij50-0ib-MJ500	9.8 ± 2.6	6.2 ± 3.3	16.0 ± 4.7	10

Table 9.1: The expected SM backgrounds, shown separately for the multi-jet and leptonic contributions, calculated from the template method, Monte Carlo simulation, and the background-only fit. The observed 2015 + 2016 data yields are also shown.

new physics process should predict in order to be excluded by that signal region; equivalently this corresponds to an upper limit on the cross-section for BSM processes, given the luminosity of the data. The fit considers one signal region, and all associated control regions, at a time; all possible values of the signal strength normalisation factor, μ_{sig} are scanned, until the p-value of the signal plus background hypothesis is smaller than 0.05, or equivalently excluded at the 95 % confidence level [157].

The model-independent fit begins by modifying the expected number of events ν_c by adding a single signal event in the signal regions, with a normalisation given by the signal strength parameter, μ_{sig} (hereafter simply denoted μ). No assumptions are made about the nature of this signal, and the effects of signal contamination in the control regions and systematic

uncertainties are neglected. μ is the only parameter of interest in the fit; all other parameters (such as background normalisations) are absorbed into a list of nuisance parameters, θ, such that the likelihood may be written as a funtion of μ and θ. To compute the upper limits on the BSM events, the profile log-likelihood ratio [157] is employed as the test-statistic, which is derived from the relative likelihood ratio:

$$\lambda(\mu) = \frac{\mathcal{L}(\mu, \hat{\hat{\theta}}(\mu))}{\mathcal{L}(\hat{\mu}, \hat{\theta})}, \tag{9.4}$$

where $\hat{\mu}$ and $\hat{\theta}$ are the μ- and θ-values which maximise $\mathcal{L}$, and $\hat{\hat{\theta}}(\mu)$ is a curve with points corresponding to the maximum value of $\mathcal{L}$ for a given μ. The $\lambda(\mu)$-value is close to 1.0 when the agreement between the data and the hypothesis value of μ is good.

The profile log-likelihood ratio test statistic is defined:

$$q_\mu = \begin{cases} -2\ln\lambda(\mu) & \mu \geq \hat{\mu}, \\ 0 & \mu < \hat{\mu} \end{cases} \tag{9.5}$$

where in this scheme data with $\mu < \hat{\mu}$ are not considered part of the final rejection region on the statistical tests. To make progress, it is necessary to obtain a p-value, which can be calculated from the probability density function of this test statistic, $f(q_\mu|\mu)$. For large numbers of observed events, an estimate of this function is obtained using the *asymptotic approximation* [158]. However, if the number of observed events is low then it is necessary to obtain an $f(q_\mu|\mu)$ using Monte Carlo technique in which many pseudo-data-sets are sampled from the original $\mathcal{L}$ of interest. Once the probability density has been correctly evaluated, the p-value for the signal plus background hypothesis is computed from the integral

$$p_{s+b} = \int_{q_{\mu,\mathrm{obs}}}^{\infty} f(q_\mu|\mu) dq_\mu. \tag{9.6}$$

μ is scanned until a $p_{s+b} < 0.05$, at which an exclusion at the 95 % confidence level on the upper limit of the number of signal events in the signal region is attained, here denoted S^{95}_{obs}. This can also be interpreted as an upper limit on the visible BSM cross-section: $\langle\epsilon\sigma\rangle^{95}_{\mathrm{obs}} = S^{95}_{\mathrm{obs}}/\mathcal{L}_{\mathrm{int}}$, where ϵ corresponds to the fraction of the total cross-section to which the signal region is sensitive. It is also possible to calculate the expected upper limit, S^{95}_{exp}, using the test statistic for the background-only hypothesis:

$$q_0 = \begin{cases} -2\ln\lambda(0) & \hat{\mu} \geq 0, \\ 0 & \hat{\mu} < 0 \end{cases} \tag{9.7}$$

where negative values are not considered in the final rejection region. The background-only p-values, p_b, is defined to be $1-p_0$, where

$$p_0 = \int_{q_{0,\mathrm{obs}}}^{\infty} f(q_0|0)dq_0. \tag{9.8}$$

All of the above assumes that a signal can only lead to a positive enhancement in the number of event counts (the convention adopted in ATLAS) and thus the presence of new physics is associated with an excess of events in a given region. The expected and observed upper limits on BSM physics, together with the upper limit on the visible cross-section and the 95 % confidence level for the background-only hypothesis, CL_b, is summarised in Table 9.2 for each signal region in the analysis. Knowledge of these upper limits is useful for theorists who may wish to perform reinterpretations of the analysis results.

In the above treatment we have only associated *excesses* in data with the presence of new physics, and have not accounted for the possibility of a downward fluctuation in data. In order to ensure that such downward fluctuations don't result in new physics signal being incorrectly excluded at the 95 % confidence level, ATLAS uses the conservative CL_s prescription [159], in which the p-value from the signal-plus-background hypothesis is normalised to the p-value for the background-only hypothesis:

$$\mathrm{CL}_s = \frac{p_{s+b}}{1-p_b}. \tag{9.9}$$

Whilst CL_s is not a true confidence level, it is always higher than CL_{s+b}, and a signal strength μ_{sig} is considered excluded if it has $\mathrm{CL}_s < 0.05$. This CL_s prescription is used to calculate both the model-independent results in Table 9.2, and the model-dependent limits (discussed in the next section).

9.2.1.2 Model-dependent Fits

In the event that no statistically significant excesses are reported, limits may be placed on sets of specific models using the same likelihood function and test statistic considered in the model-independent fit. This model-dependent fit is used to reject the hypothesis from certain

Signal Regions	$\langle\epsilon\sigma\rangle^{95}_{\text{obs}}$ [fb]	S^{95}_{obs}	S^{95}_{exp}	Discovery p_0
SR-8ij50-0ib	7.2	260	270^{+90}_{-70}	0.50
SR-8ij50-1ib	6.4	230	150^{+80}_{-60}	0.50
SR-8ij50-2ib	4.6	170	160^{+50}_{-40}	0.40
SR-9ij50-0ib	1.5	53	58^{+20}_{-15}	0.50
SR-9ij50-1ib	1.2	44	55^{+18}_{-14}	0.50
SR-9ij50-2ib	1.0	35	38^{+12}_{-9}	0.50
SR-10ij50-0ib	0.30	11	15^{+6}_{-4}	0.50
SR-10ij50-1ib	0.31	11	15^{+6}_{-4}	0.50
SR-10ij50-2ib	0.31	11	12^{+5}_{-3}	0.50
SR-11ij50-0ib	0.23	8.5	$6.3^{+3.0}_{-1.5}$	0.21
SR-11ij50-1ib	0.21	7.4	$6.5^{+2.6}_{-1.7}$	0.34
SR-11ij50-2ib	0.19	6.9	$6.0^{+2.2}_{-1.3}$	0.35
SR-7ij80-0ib	3.1	110	130^{+40}_{-30}	0.50
SR-7ij80-1ib	2.7	100	120^{+40}_{-30}	0.50
SR-7ij80-2ib	1.7	60	72^{+22}_{-17}	0.50
SR-8ij80-0ib	0.80	29	27^{+10}_{-7}	0.40
SR-8ij80-1ib	0.62	22	24^{+9}_{-7}	0.50
SR-8ij80-2ib	0.49	18	16^{+6}_{-5}	0.41
SR-9ij80-0ib	0.22	7.6	$7.9^{+3.4}_{-2.0}$	0.50
SR-9ij80-1ib	0.21	7.5	$7.5^{+2.8}_{-2.1}$	0.50
SR-9ij80-2ib	0.22	8.0	$5.9^{+2.6}_{-1.4}$	0.20
SR-8ij50-0ib-MJ340	2.9	100	130^{+40}_{-30}	0.50
SR-8ij50-0ib-MJ500	1.0	36	48^{+17}_{-13}	0.50
SR-9ij50-0ib-MJ340	0.87	32	42^{+14}_{-11}	0.50
SR-9ij50-0ib-MJ500	0.32	12	20^{+8}_{-6}	0.50
SR-10ij50-0ib-MJ340	0.25	9.1	14^{+6}_{-4}	0.50
SR-10ij50-0ib-MJ500	0.22	7.9	11^{+4}_{-3}	0.50

Table 9.2: Summary of the results of the model-independent fit: from left to right, the 95 % CL limits on the visible cross-section, the expected number of signal events, and the number of signal events given the expected number of background events (including $\pm 1\sigma$ excursions on the expected value), and finally the discovery p-value, p_0.

signal models, namely those outlined in Section 8.2. The best expected signal region is included in the fit, together with the control regions and their associated background normalisation, and μ_{sig} is allowed to float. The CL_s is calculated for $\mu_{\text{sig}} = 1$ at different signal grid points (e.g. a grid built from the mass of the gluino and the neutralino), creating an exclusion contour in the parameter-space of the signal grids, usually the particle mass points. A given signal point is excluded at the 95 % confidence level if its $\text{CL}_s < 0.05$. A 95 % exclusion contour is drawn from the interpolation between points at, or around, $\text{CL}_s = 0.05$. Having reported no statistically significant excess, we shall first review the limits placed on the models in Section 8.2, before discussing how this analysis can be used to investigate the sensitivity to

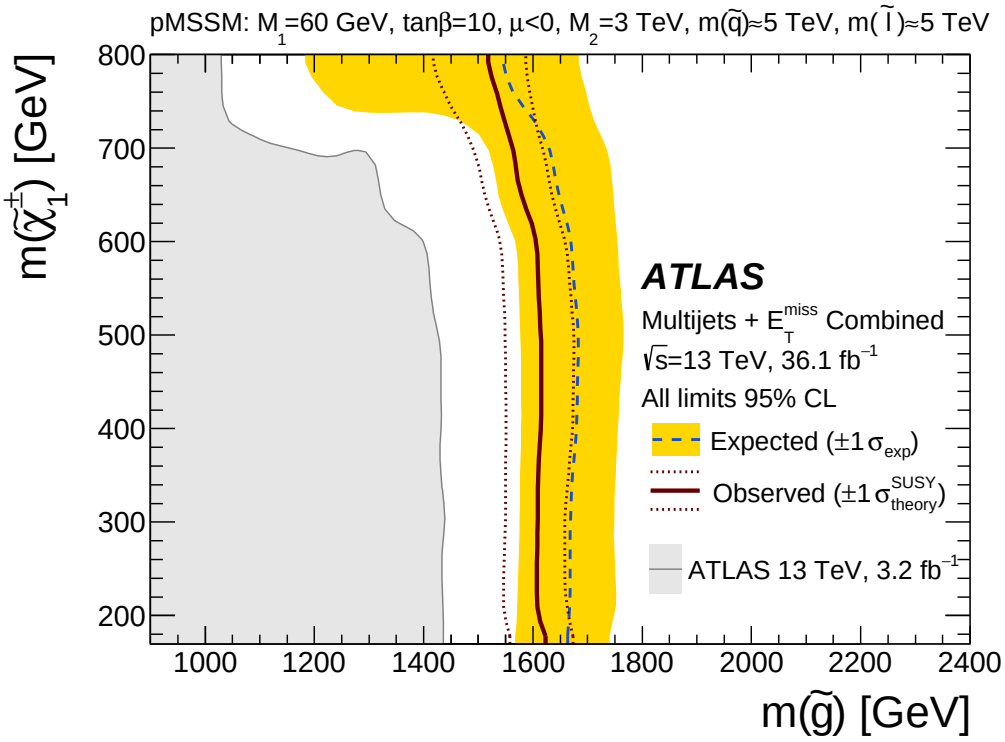


Figure 9.3: Expected and observed limits on the pMSSM slice model.

new regions of the R-parity-violating SUSY parameter space.

9.2.2 Constraining Gluinos, Neutralinos, and Stops

In the final sensitivity calculation the 27 signal regions are combined by considering, for each signal grid point, the particular region which gives the best expected sensitivity calculated from applying the signal hypothesis to the expected results from the Monte Carlo backgrounds and multi-jet template estimation. The combined regions are used to calculate the final observed limits, which test the signal hypothesis against the actual observed event yields as measured in data. Systematic uncertainties are also included in the final limits.

In the pMSSM slice the gluino masses are excluded at the 95 % CL up to 1800 GeV, for chargino masses of up to 600 GeV. For the R-parity-conserving simplified models, the gluino masses can also be excluded up to about 1800 GeV in the case of the two-step models, whilst in the $\tilde{g}t\bar{t}$ simplified model, the more boosted gluino mass points can be excluded upto about 1540 GeV, whilst the compressed neutralino LSP masses can be excluded up to about 520 GeV. These exclusion limits are shown in Figures 9.3 and 9.4 for the pMSSM slice and the two-step and $\tilde{g}t\bar{t}$ signals respectively. These limits should be interpreted with some caution because they

apply to specific models within the SUSY parameter space. However they do provide indications that multi-jets is able to constrain the gluino mass up to values encroaching on 2.0 TeV, which can be disfavoured by certain naturalness arguments.[2] For completeness, a comparison of the multi-jets sensitivity to the two-step model with other ATLAS SUSY analyses is provided in Figure 9.5, which again indicates that gluinos are being excluded up to around 2.0 TeV at the 95 % confidence level.

As has been briefly discussed, the stop model considered in this analysis contains an R-parity-violating coupling for the direct decay of the stop quark to purely heavy-flavour SM particles ($\bar{b}\bar{s}$ for the simplified model considered here). In this RPV scenario, gluino masses can be excluded up to about 1400 GeV and the exclusion on the stop mass reaches about 900 GeV. This 900 GeV mass could be interpreted by some as being in tension with the expected masses for the production of stop quarks in natural SUSY scenarios, and is competitive with the limits from other models considered in direct searches for stop production [160]. The limits are shown in Figure 9.6.

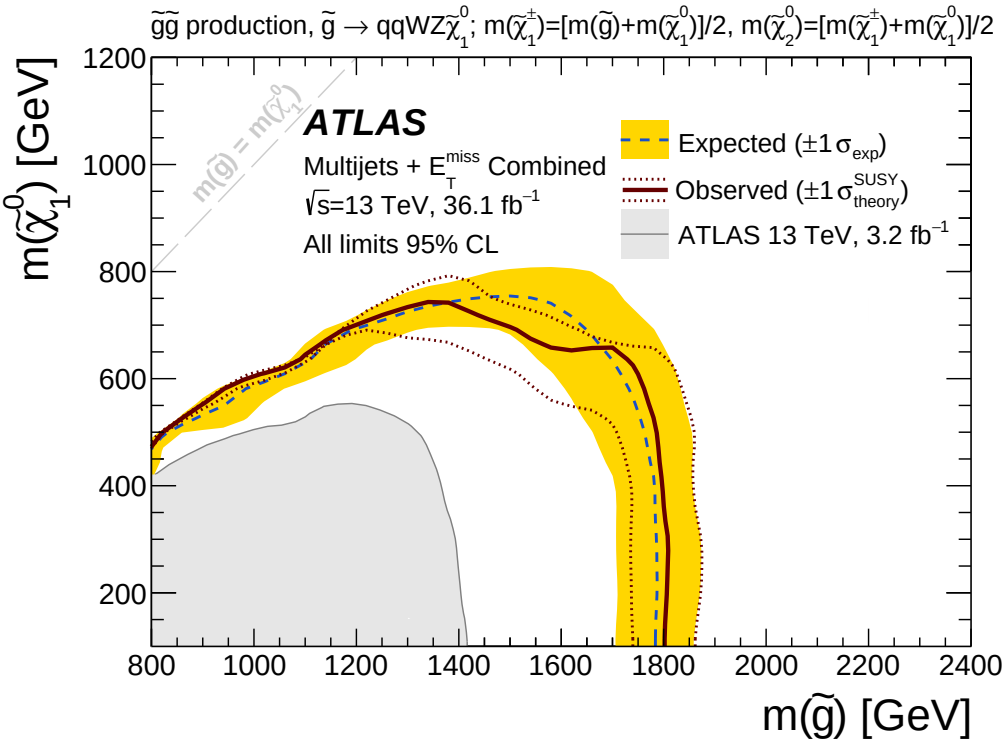
In principle it is possible to extend the exclusion power of this analysis to other simplified models, and particularly to different classes of RPV simplified models where the neutralino can be long-lived, or decay promptly. The analysis regions are sensitive to final states with little and moderate $E_{\mathrm{T}}^{\mathrm{miss}}$, and therefore to a variety of scenarios which depend on the strength of the RPV coupling. The analysis sensitivity to the R-parity-violating phase-space, beyond setting limits on sparticle masses, is the subject of [161] and the next section.

9.2.3 Constraining R-Parity Violating SUSY

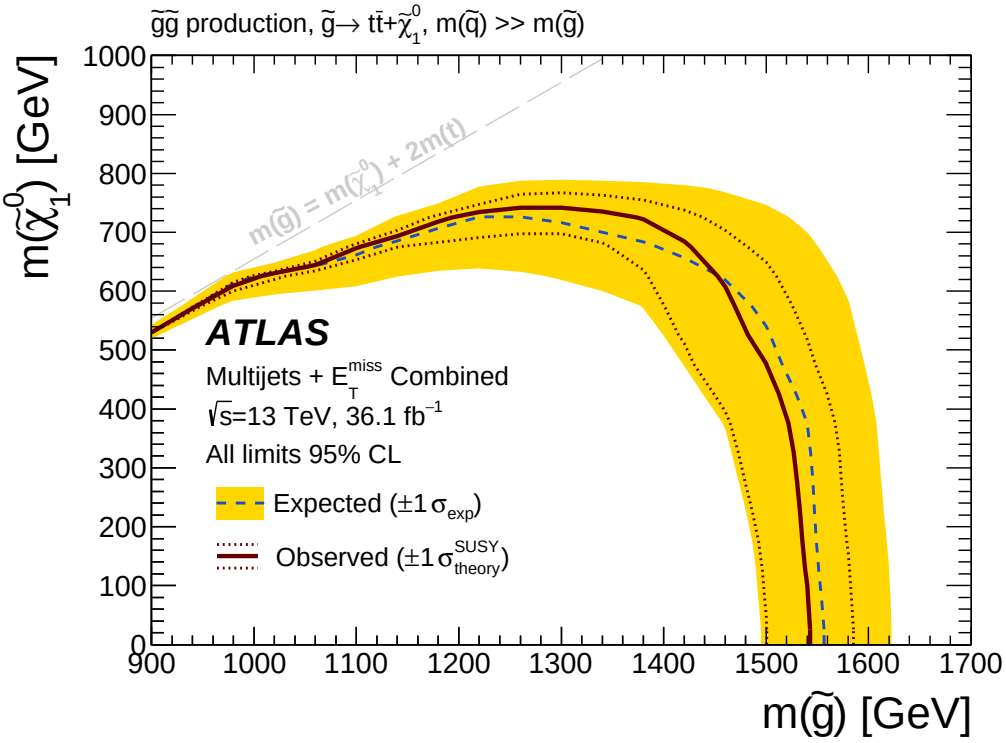
9.2.3.1 The 'RPC-meets-RPV' Spectrum

Since the multi-jets analysis is sensitive to a variety of SUSY scenarios, both R-parity-conserving and violating, it is interesting to quantify how this sensitivity changes with respect to the size of the R-parity-violating coupling where, in the limit of zero R-parity couplings RPC SUSY is recovered. There is also a clear relationship between the strength of the coupling and the distance the sparticle involved in the decay propagates within ATLAS: smaller couplings will be

[2] A word of caution here: it is important to recall that natural SUSY is most strongly constrained by limits on Higgsino production, as briefly discussed in Chapter 3.



(a) Two-step



(b) $\tilde{g}t\bar{t}$

Figure 9.4: Expected and observed limits on the two-step and $\tilde{g}t\bar{t}$ RPC SUSY simplified models.

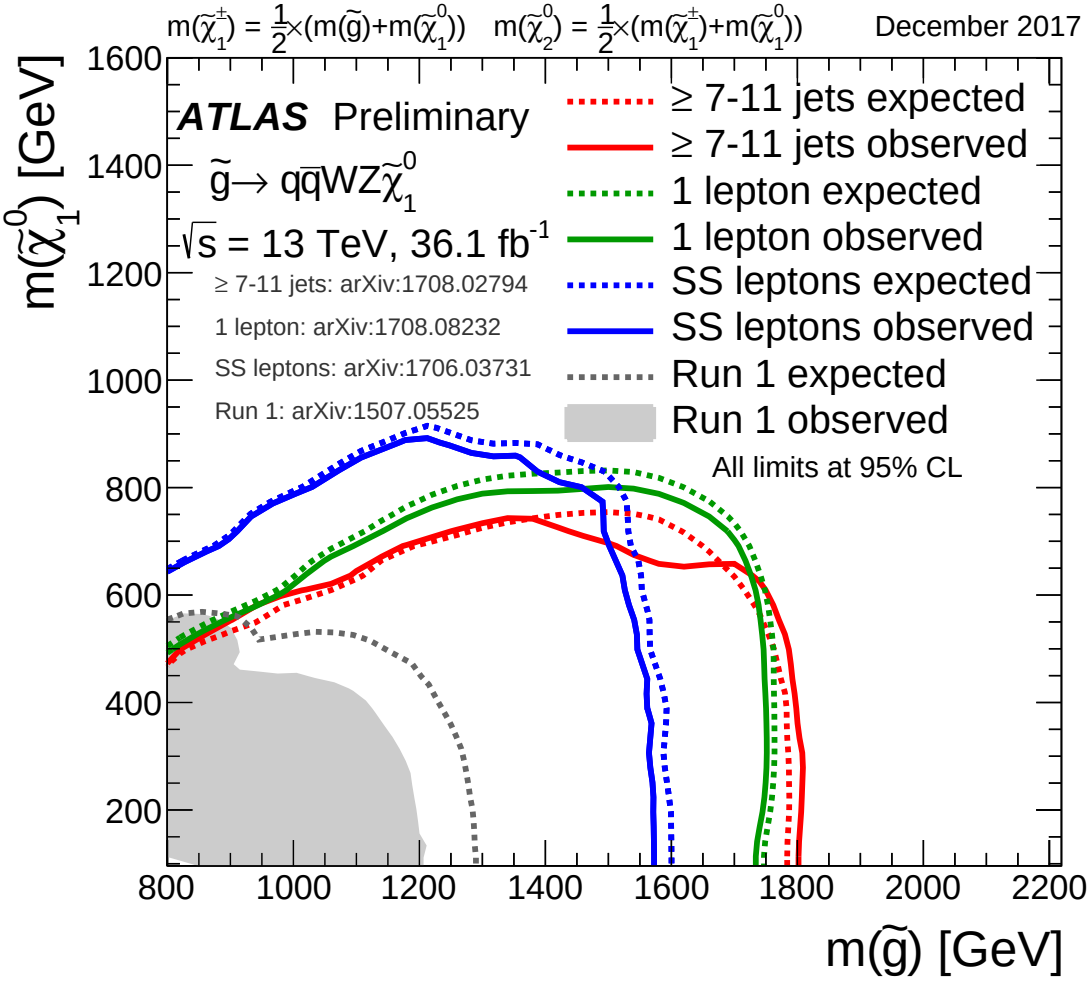


Figure 9.5: A summary of the 95 % CL exclusions placed on the two-step model by a variety of SUSY analyses. The observed multi-jets limit is shown by the solid red line.

associated with more long-lived sparticles which propagate some distance through the ATLAS tracker, and even into the calorimeter; larger couplings correspond to more prompt RPV decays which occur close to the interaction point. This relationship is illustrated schematically in Figure 9.7 together with versions of the $\tilde{g}q\bar{q}$ and $\tilde{g}t\bar{t}$ models where the baryon-number-violating coupling, λ'', is allowed to float. In those models the neutralino is taken as the RPV SUSY particle which decays according to λ''. The relationship between the distance, L, (in centimetres) the neutralino propagates and the coupling at its decay vertex is [162]:

$$L = \frac{0.9\beta\gamma}{\lambda''^2}\left(\frac{m_{\tilde{q}}}{100\text{ GeV}}\right)^4\left(\frac{1\text{ GeV}}{m_{\tilde{\chi}_1^0}}\right)^5, \tag{9.10}$$

where the LSP neutralino χ_1^0 is assumed to be dominated by its bino-component, and have a mass of 200 GeV, and the squark masses are set to 3.0 TeV and 5.0 TeV in the $\tilde{g}q\bar{q}$ and $\tilde{g}t\bar{t}$

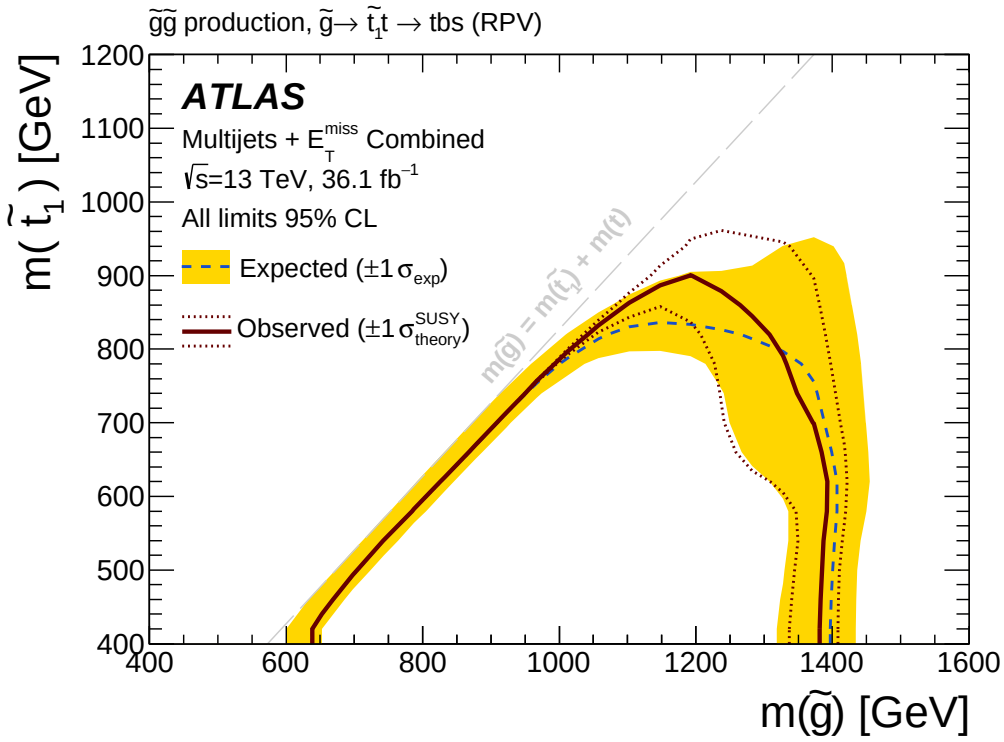


Figure 9.6: Expected and observed limits on the RPV stop simplified model.

models respectively. The $\tilde{g}q\bar{q}$ model has a non-zero λ''_{112} coupling of the neutralino to lighter generation quarks ($q = \{u, d, s, c\}$), and the $\tilde{g}t\bar{t}$ model has a non-zero λ''_{323} coupling of the neutralino to heavy-flavour, $\chi_1^0 \to tbs$. The models with larger couplings, giving rise to prompt RPV decays, will produce less $E_{\mathrm{T}}^{\mathrm{miss}}$ in the final states, and strongly produced SUSY analyses which rely on an $E_{\mathrm{T}}^{\mathrm{miss}}$ trigger will lose sensitivity once the average $E_{\mathrm{T}}^{\mathrm{miss}}$ associated with the neutralino falls below the trigger threshold. Analyses such as the multi-jets search, which have no $E_{\mathrm{T}}^{\mathrm{miss}}$ trigger threshold, should not be affected as severely and hence retain some sensitivity to these lower $E_{\mathrm{T}}^{\mathrm{miss}}$, prompt decays.

9.2.3.2 Results

The $\tilde{g}q\bar{q}$ and $\tilde{g}t\bar{t}$ are generated by varying the gluino mass and the λ'' coupling. The observed sensitivity from the multi-jets analysis, calculated from the usual model-dependent fit, is shown by the solid red curve in Figures 9.8 and 9.9 , where the expected sensitivity is also given by the dotted red curve. For the first time in ATLAS, limits have now been set across the fully theoretically motivated λ''_{112} and λ''_{323} spectra, for different gluino masses. Several RPC- and RPV-focused SUSY searches are considered. In Figure 9.8 the multi-jets sensitivity to $\tilde{g}q\bar{q}$

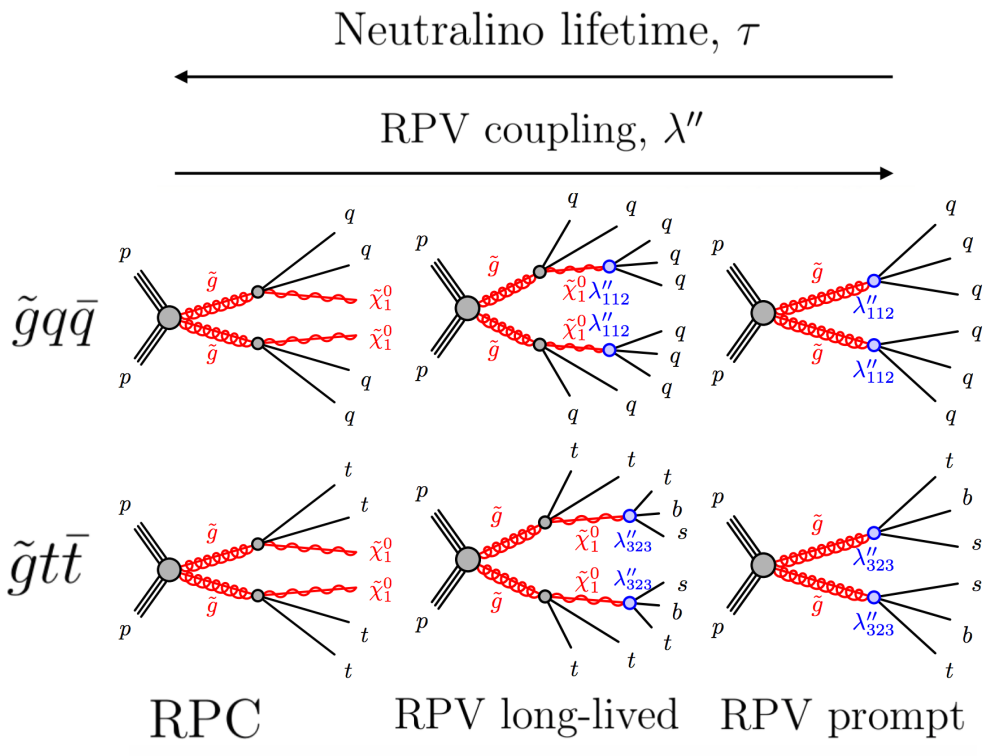


Figure 9.7: The RPC-RPV continuum from RPC, long-lived, and prompt decaying neutralinos in ATLAS, for the two simplified models described.

is compared with another SUSY search which is optimised to search for gluino production, leading to final states with 2 to 6 jets [163]. This analysis is generally more sensitive than the multi-jets search at lower λ''_{112} where the neutralino is more long-lived and the final states typically have a jet multiplicity lower than 7 (which multi-jets is optimised for). In Figure 9.9 the multi-jets sensitivity to the $\tilde{g}t\bar{t}$ model is compared with several other SUSY analyses. The analyses indicated by the green [164] and black [165] curves are optimised to search for RPV SUSY scenarios, and are therefore the most sensitive, for a given gluino mass, at larger λ''_{323} where the RPV neutralino decays more promptly. An interesting comparison is with the multi-b analysis [166], which has a limit indicated by the blue curve. This analysis searches for RPC gluino production with final states with 3 or more b-tagged jets, not dissimilar to certain signal region strategies in the multi-jets analysis. This analysis is more sensitive than the multi-jets search at lower λ''_{323} because its signal regions are explicitly optimised to exclude RPC $\tilde{g}t\bar{t}$ signals with heavier gluinos, unlike the multi-jets analysis where a more model-independent view is taken in its optimsiation. Interestingly, the multi-b sensitivity falls sharply at larger λ''_{323}, whilst the multi-jets sensitivity remains reasonably constant. Here we are seeing the effect of the $E_{\mathrm{T}}^{\mathrm{miss}}$ trigger which is used in the multi-b analysis, where there is a loss in sensitivity

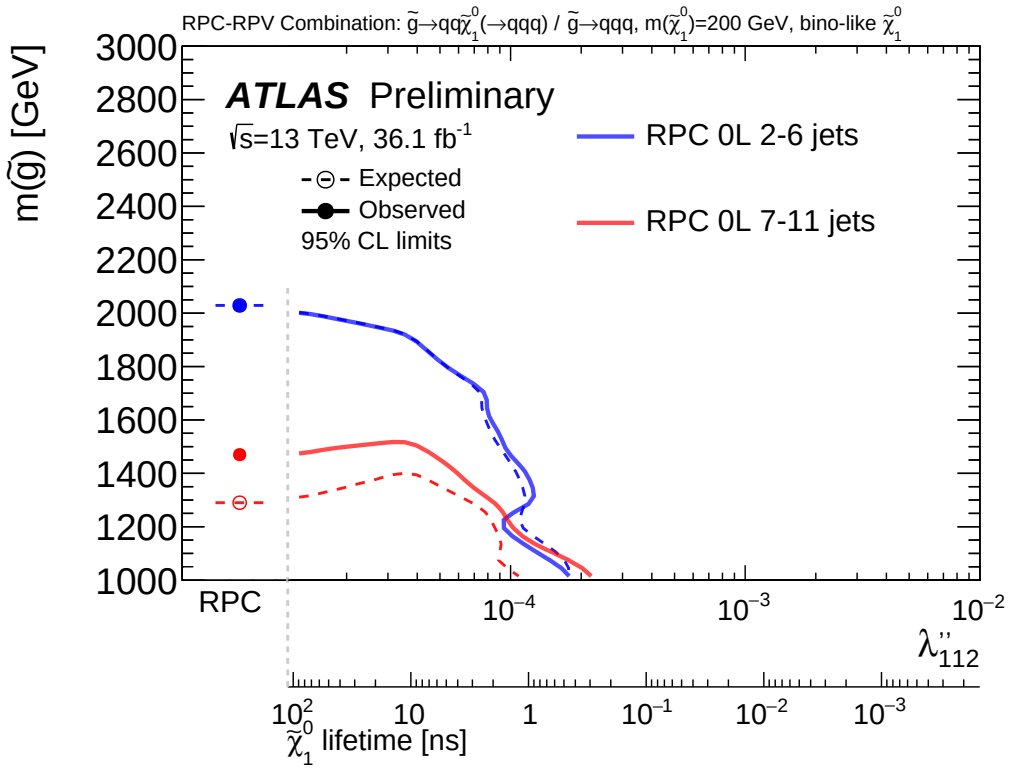


Figure 9.8: The expected and observed limits set by a variety of SUSY searches [70, 163], including multi-jets, in the 'RPC-meets-RPV' spectrum of the $\tilde{g}q\bar{q}$ simplified model.

to the prompt RPV neutralino decay which generates lower $E_{\mathrm{T}}^{\mathrm{miss}}$ final states. The multi-jets analysis instead relies on a series of jet triggers, and so maintains a sensitivity to the high jet multiplicity, lower $E_{\mathrm{T}}^{\mathrm{miss}}$ prompt RPV processes at larger λ''_{323}.

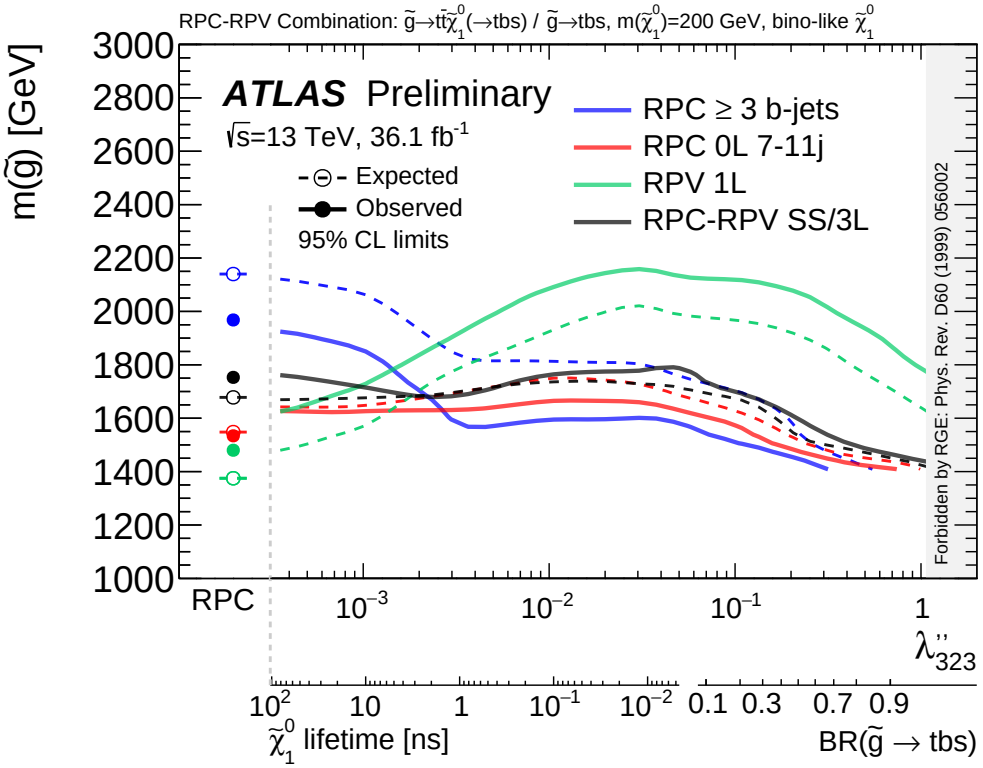


Figure 9.9: The expected and observed limits set by a variety of SUSY searches [70, 164–166], including multi-jets, in the 'RPC-meets-RPV' spectrum of the $\tilde{g}t\bar{t}$ simplified model.

10
Multi-Jets with Full Run 2 Data

To close Part III of this thesis, we shall briefly review the status of the latest iteration of the multi-jets analysis, which is using the full Run 2 data-set and is aiming for publication in 2019. This latest quest for new matter is making use of several innovative developments in jet physics, including optimisations of the reclustering algorithm, the technology of particle flow jet reconstruction, and the ongoing development of taggers which can classify quark- and gluon-initiated jets in complex multi-jet final states. These exciting endeavours are now briefly outlined.

10.1 Jet Reclustering Optimisation

The jet reclustering algorithm, used to construct the M_{J}^{Σ} observable, reclusters jets into new jets of some larger radius [96]. In the 36 fb^{-1} iteration of the multi-jets analysis, the reclustering approach adopted is to combine the $R = 0.4$ jets into new ones of size $R = 1.0$, without applying jet grooming of any kind, which would remove additional soft contributions to the jet. It is believed that these $R = 1.0$ jets should be sufficient to capture sufficiently boosted heavy objects, however the reclustering procedure has never been optimised.

For the full Run 2 analysis, the jet reclustering definition has been optimised by studying reclustered jets of differing reclustered radii (0.6, 0.8, 1.0, 1.2, 1.4, 1.5, and 1.6) and studying the *expected* sensitivity to different SUSY models in the inclusive 8-jet signal region, scaled to an integrated luminosity of 120 fb^{-1} [1]. In addition to changing the reclustered jet radius, the

[1] Comparable to the projected full Run 2 integrated luminosity at the time of these studies

event-level jet mass has also been modified in order to find, for a given reclustered jet definition, the jet mass cut which provides the largest gain in signal sensitivity. As well as fixed-radius reclustering, variable-radius [167] (VR) reclustering is also investigated. In the VR approach, individual jets are reclustered into larger jets of some radius R which explicitly depends on the jet p_{T} according to $R = \rho/p_{\mathrm{T}}$, where $\rho = 2m_X$ defines the mass scale of the reclustered jet. Different mass scales are considered: m_X = (90, 125, 175, 300, 500, 1000) GeV; and two maximum reclustered radii set: $R_{\mathrm{max}} = 1.8$ and 2.5. The effect of the maximum radius is in fact negligible. The effects of varying jet trimming [89] parameters on the reclustered jets are also measured, and also found to be negligible. The studies summarised here therefore use untrimmed jets.

An important conclusion of this scan is that sensitivity to compressed[2] regions of SUSY phase-space can be enhanced by using larger reclustered jet radii, so that a decaying object with a larger opening angle (between its respective decay products) can be fully captured within the reclustered jet cone. Conversely, smaller jet radii can be more sensitive to boosted regions of phase-space where the average opening angle between decay products will shrink enough to allow the parent particle to be fully contained within the smaller reclusterd jet. However, it is also possible to instead keep the jet radius fixed at $R = 1.0$ and vary the M_{J}^{Σ} cut on the reclustered jet. It is observed that standard $R = 1.0$ reclustered jets, with M_{J}^{Σ} cuts at a340 and 500 GeV, provide as much sensitivity to typical SUSY models as the most promising other, more exotic reclustered jet definitions considered in the scan. This point is illustrated in Figures 10.1 (which shows the performance of different jet radii for $M_{\mathrm{J}}^{\Sigma} = 500$ GeV is first shown), 10.2 (showing the $R = 1.0$ jet definition covering a nice range of sensitivities to the simplified model for different M_{J}^{Σ}) and finally 10.3 (which shows different VR jets with $M_{\mathrm{J}}^{\Sigma} > 500$ GeV imposed). These studies are carried out using $\tilde{g}t\bar{t}$ signal model points.

Studies using the two-step, pMSSM, and RPV stop signal models show similar expected sensitivities between various reclustered jet definitions. Scanning through the different jet collections, it was observed that the $R = 1.0$ untrimmed jets provided a good sensitivity to a variety of signal models, which could be easily tuned by varying the M_{J}^{Σ} parameter. Again, consistent with results from previous iterations, M_{J}^{Σ} of 340 and 500 GeVare found to offer the

[2]Where the mass-splitting between the gluino and LSP is small.

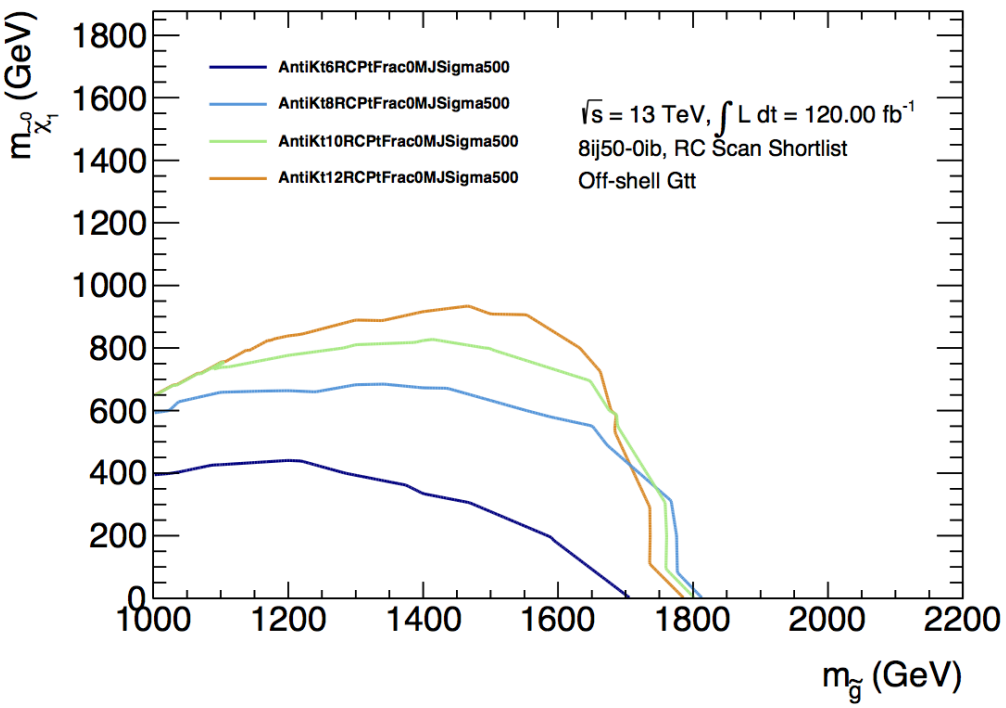


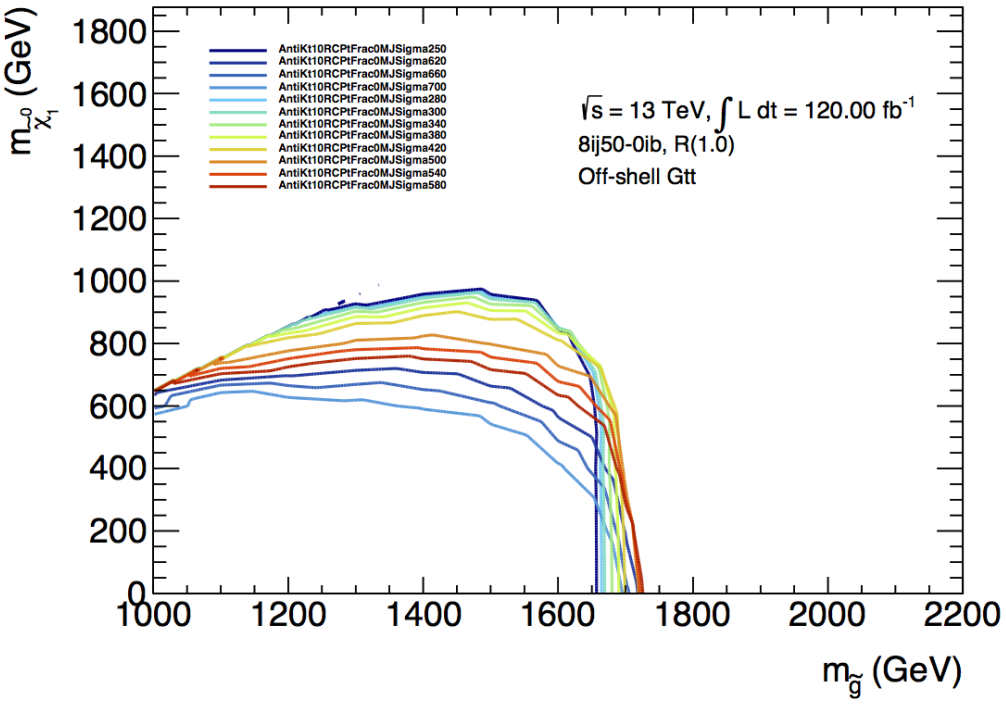
Figure 10.1: Different jet radii for $M_{\mathrm{J}}^{\Sigma} > 500$ GeV. For example, `AntiKt12RCPtFrac0MJSigma500` refers to jets with reclustered radius $R = 1.2$, no trimming applied, and $M_{\mathrm{J}}^{\Sigma} > 500$ GeV imposed.

best sensitivity to the more compressed and boosted regions of phase-space respectively.

For the case of variable-radius reclustering it is observed that no single combination of a VR reclustered jet and some M_{J}^{Σ} cut provides the best possible signal sensitivity across all of the available phase-space for the various SUSY signal models considered in the multi-jets analysis. Therefore, putting all of these observations together, it is concluded that the original $R = 1.0$ jets with M_{J}^{Σ} cuts at 340 and 500 GeV offer the best signal sensitivity, and so these are adopted as the reclustered jet definition for the full Run 2 iteration of the multi-jets analysis.

10.2 Object-Based $E_{\mathrm{T}}^{\mathrm{miss}}$ Significance

The event-level variable considered in the template method is the $E_{\mathrm{T}}^{\mathrm{miss}}$ significance; recall that this variable quantifies the relative likelihood that a measurement of $E_{\mathrm{T}}^{\mathrm{miss}}$ in the event originates from a real source (such as a neutrino or some other undetected new physics particle) or a fake source (e.g. jet resolution mismeasurement). The definition considered thus far is in fact a 'macroscopic' approximation to the true $E_{\mathrm{T}}^{\mathrm{miss}}$ significance, which would in reality depend



(a) Scan in M_{J}^{Σ} for $R = 1.0$ jets.



(b) Different jet definitions considered.

Figure 10.2: Different M_{J}^{Σ} for $R = 1.0$ jets. For example, `AntiKt10RCPtFrac0MJSigma580` refers to jets with reclustered radius $R = 1.0$, no trimming, and $M_{\mathrm{J}}^{\Sigma} > 580$ GeV imposed.

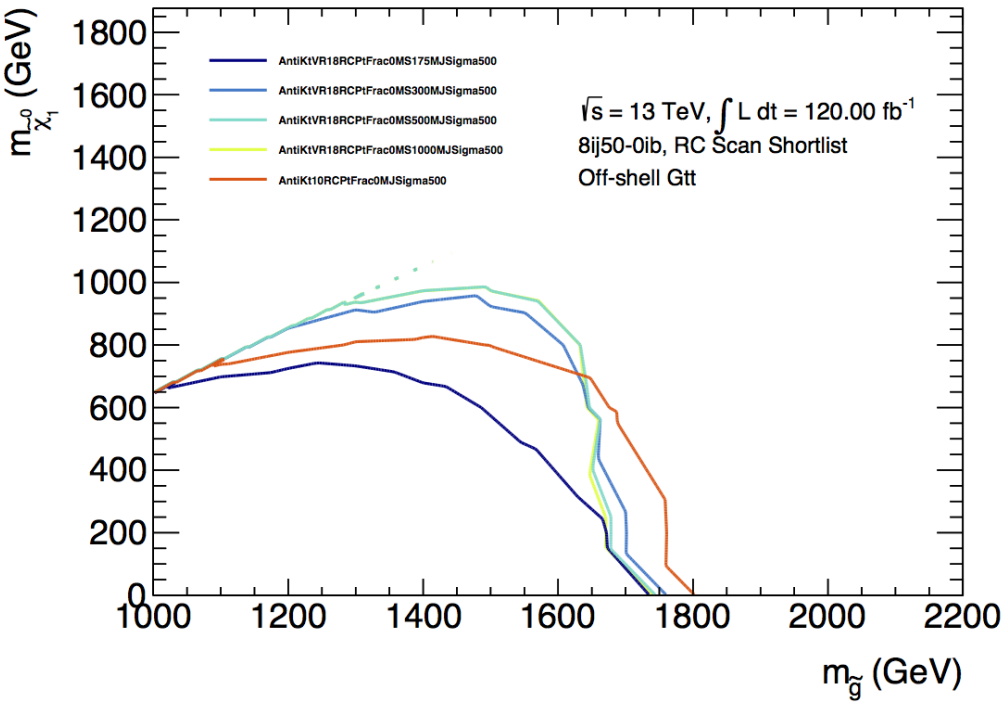


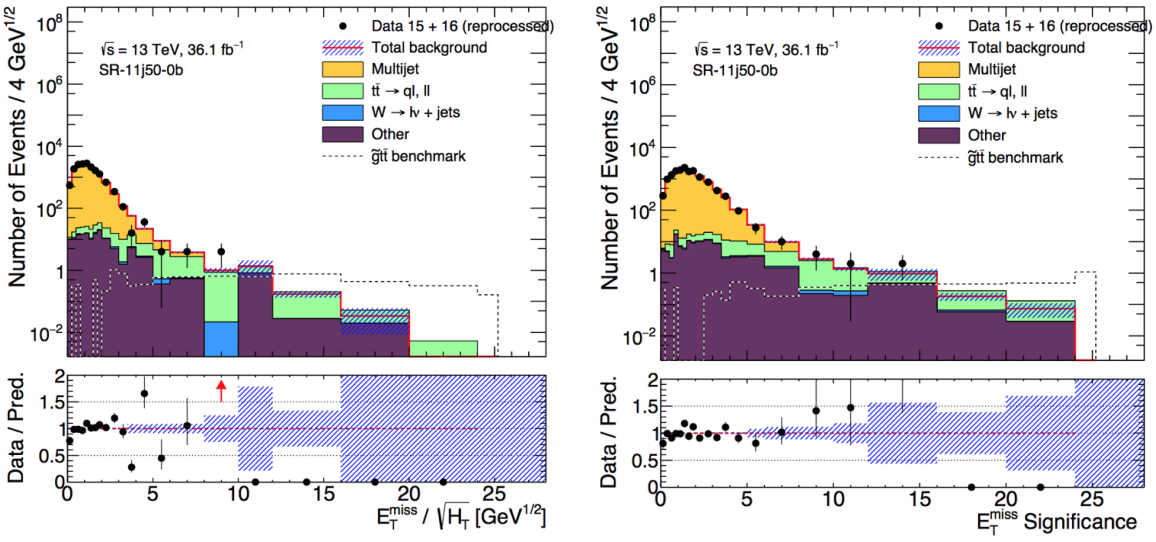
Figure 10.3: Different VR reclustered jet definitions for $M_{\mathrm{J}}^{\Sigma} > 500$ GeV. For example, `AntiKtVR18RCPtFrac0MS300MJSigma500` refers to jets which are produced from VR reclustering with $R_{\mathrm{max}} = 1.8$, $m_X = 300$ GeV, and $M_{\mathrm{J}}^{\Sigma} > 500$ GeV is applied.

on the resolutions from all physics objects in the event and not just the jets. A more general, object-based $E_{\mathrm{T}}^{\mathrm{miss}}$ significance [168] can be calculated from the hypothesis test that the real $\mathbf{E}_{\mathrm{T}}^{\mathrm{miss}}$ carried by invisible particle is zero, against the hypothesis that it is different from zero.

This significance quantifies the likelihood that, given there is a measurement of $\mathbf{E}_{\mathrm{T}}^{\mathrm{miss}}$, it is from a real source rather than a mismeasurement or unwanted detector effect. Under the assumption that the detector physics objects which enter into this calculation have transverse momenta with probability distributions which are Gaussian and described by a covariance matrix V_i, the significance can be reduced to a χ^2 expression:

$$\mathcal{S}^2 = (\mathbf{E}_{\mathrm{T}}^{\mathrm{miss,real}})^{\mathbf{T}}(\mathbf{\Sigma_i V_i})^{-\mathbf{1}}(\mathbf{E}_{\mathrm{T}}^{\mathrm{miss,real}}). \tag{10.1}$$

The covariance matrix encodes the resolutions of the different hard physics objects which make up the $\mathbf{E}_{\mathrm{T}}^{\mathrm{miss}}$ calculation. Five objects enter into this calculation: electrons, photons, hadronically-decaying taus, small-radius jets, and muons. In addition the track-soft term (TST) is also added: recall that this TST is calculated from the transverse momenta of reconstructed inner detector tracks which are not associated with any hard object in the detector, but which



(a) Using $E_{\mathrm{T}}^{\mathrm{miss}}/\sqrt{H_{\mathrm{T}}}$ **(b)** Using object-based $E_{\mathrm{T}}^{\mathrm{miss}}$ significance

Figure 10.4: A comparison of the object-based and macroscopic $E_{\mathrm{T}}^{\mathrm{miss}}$ significances for an 11-jet signal region in the 2015+2016 multi-jets analysis.

are associated to the primary vertex of the event and therefore do not originate from pp collisions associated with pileup. In addition the covariance matrix also contains the correlations between the components of the particular object resolution which are longitudinal and transverse to the direction of the total $\mathbf{E}_{\mathrm{T}}^{\mathrm{miss}}$ in the event.

The object-based definition of the $E_{\mathrm{T}}^{\mathrm{miss}}$ significance is being adopted for the full Run 2 multi-jets analysis. The new definition is more correct than the 'macroscopic' definition, and preliminary studies indicate that the closure of the background prediction with respect to data is much improved in more challenging regions of phase-space when the object-based significance is adopted. This is illustrated in Figure 10.4, where the 36 fb^{-1} [3] `SR-11ij50-0ib` region is compared for both the $E_{\mathrm{T}}^{\mathrm{miss}}/\sqrt{H_{\mathrm{T}}}$ significance and the object-based significance.

10.3 Particle Flow Jets

One particularly attractive prospect for the full Run 2 analysis is the use of `AntiKt4EMPFlow` jets as the small-radius jet definition (see Section 6.1.6 for more details). In using these jets the analysis will benefit from the improved jet energy resolution at lower p_{T} from the additional tracking information, as well as improvements to the $E_{\mathrm{T}}^{\mathrm{miss}}$ resolution. Studies also indicate that

[3] The full Run 2 analysis is blinded at the time of writing.

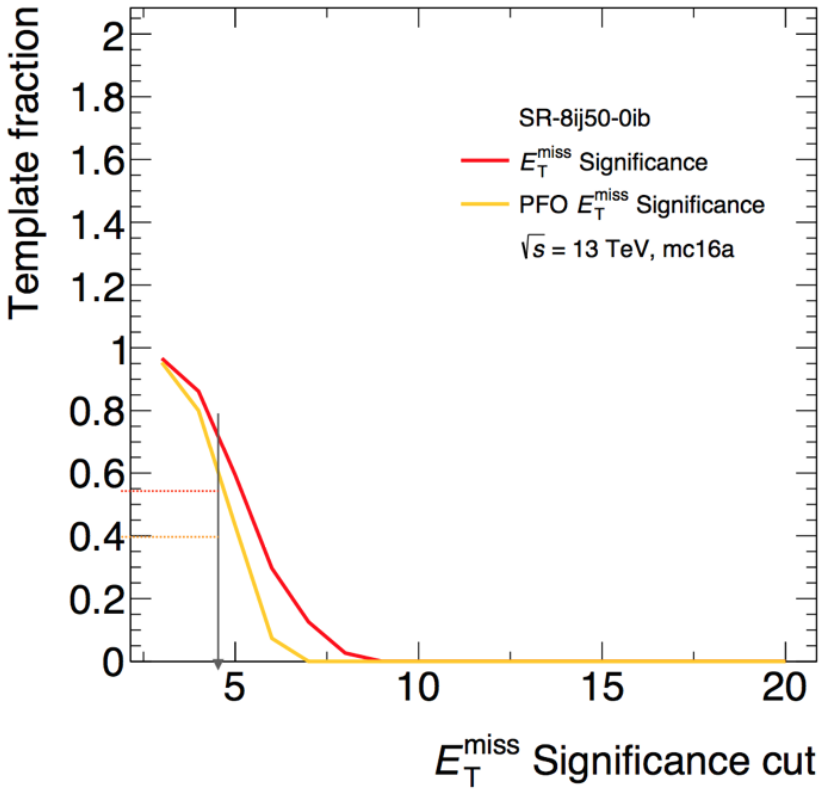


Figure 10.5: The template fraction of the total background prediction, for different requirements on the object-based $E_{\mathrm{T}}^{\mathrm{miss}}$ significance given in the x-axis, for both `EMTopo` and `EMPFlow` object definitions.

the relative proportion of multi-jet background expected in the signal regions will also be reduced using the new particle flow jets. This is shown in Figure 10.5 where the template fraction of the total background is calculated for different lower bounds on the object-based $E_{\mathrm{T}}^{\mathrm{miss}}$ in the `SR-8ij50-0ib` signal region definition (note that the multi-jet background is estimated from the usual template method and the other backgrounds are simulated in Monteestimatesdata in the 8-jet signal region is explicitly used of course). It can be seen that for a signal region based on $E_{\mathrm{T}}^{\mathrm{miss}}$ significance above 4.0 or 5.0 (now unitless, as this is the object-based definition), there is a 30 % reduction in the multi-jet background in the signal region when the PFO significance is used, which calculates the object-based $E_{\mathrm{T}}^{\mathrm{miss}}$ significance using the particle flow information.

10.4 Quark-gluon Tagging

An important conclusion of the studies in Section 8.3.6 is that the largest improvements in sensitivity estimates to SUSY models considered in the multi-jets analysis comes from constructing event-level variables, such as M_{J}^{Σ}, and not jet-by-jet variables such as those used in conventional top-tagging. One of the major challenges when analysing complex multi-jet

final states is differentiating the quark-initiated jets from the gluon-initiated jets, where the latter come from higher-order showering effects of the initial hard-scatter event. It would be advantageous to develop a rigorous procedure for identifying such gluon-initiated jets so that they can be removed in the final selection, whilst preserving the jets which are produced from the original hard-scatter event. Several quark-gluon taggers are available for such a task. A simple tagger measures the number of tracks associated to a jet, n_{track} [168]. Gluon-initiated jets tend to undergo more radiative emissions, and therefore such jets are wider with respect to their axis, and have more tracks associated to them. Therefore by selecting jets which have fewer associated tracks, it is possible to remove some gluon-initiated jets from the event. Cutting on n_{track} works well at the jet level, but at the event-level such a variable suffers from the same problems as the substructure variables considered in top-tagging. In the multi-jets analysis, the task is therefore to construct an event-level quark-gluon discriminating variable based on the initial n_{track} information from the individual jets in the event.

For the n_{track} ATLAS tagger [168], a jet is selected and its p_{T} is measured. A given track multiplicity value, $n_{\text{track},X}$, is considered for each jet, and the value depends on the jet p_{T}. This $n_{\text{track},X}$ is the typical upper limit on the track-multiplicity for a quark-initiated jet in the p_{T}-range of the jet of interest. The jet is then classified as either quark or gluon-initiated according to this simple prescription: if $n_{\text{track}} > n_{\text{track},X}$ the jet is tagged as gluon-initiated, otherwise it is considered to be quark-initiated.

This classification can be extended to the event-level in a simple way. For each quark-tagged jet, a counter Q can be incremented by 1. Then, summing over all of the jets in the event, the event-level Q_{J}^{Σ} is defined:

$$Q_{\text{J}}^{\Sigma} = \sum_i Q_i, \tag{10.2}$$

where the sum is over all i quark-tagged jets in the event. This definition can be generalised by adding appropriate weights in the following way:

$$Q_{\text{J}}^{\Sigma} = \sum_i (w_i^q Q_i + w_i^g G_i). \tag{10.3}$$

Here the Q_i and G_i start from 0, and the former is incremented by 1 for a quark and the latter by -1 for a gluon. The w_i are defined in order to add larger weightings for jets with either few or many associated tracks. The weights themselves can be optimised. For example, if a jet has a measured number of associated tracks n_{track}, then one possible definition for these weights is:

$$w^{q/g} = 1.0 + \frac{|n_{\text{track}} - n_{\text{track},X}|}{\sqrt{n_{\text{track},X}}}. \tag{10.4}$$

The use of such event-level quark-gluon tagging variables, and variants thereof, remains under investigation.

10.5 Summary

In this chapter we have briefly described that latest developments in the multi-jets analysis with the full Run 2 ATLAS data-set. The author is serving as one of the leaders of this analysis, and is spearheading the incorporation of some of the latest jet and $E_{\text{T}}^{\text{miss}}$ reconstruction techniques in order to enhance the sensitivity to new physics at the LHC. This analysis is now in a very mature stage, with most of the details surrounding the signal region definitions, systematic uncertainties, and statistical fitting procedure finalised and will further our understanding of the fundamental building blocks of reality using the very latest LHC physics data in an upcoming publication.

11 Conclusions

The LHC has now completed its second full run of data-taking, amassing approximately 140 fb^{-1} of integrated luminosity suitable for physics analysis. This is an extraordinary, record-breaking achievement in experimental particle physics. At the time of writing, many ATLAS analyses which use this full Run 2 data-set are being prepared, including the latest iteration of the multi-jets search, which the author is leading. These analyses will also benefit from the latest technology in jet reconstruction, an area which the author has contributed to actively and with great enthusiasm throughout Run 2 of ATLAS, and which was outlined in Part II of this thesis.

The first major component of original work in this thesis has focused on research and development in hadronic jet physics, including 1) improvements which allow ATLAS experimentalists to better measure jet kinematic variables, such as the jet mass, and assess their uncertainties through the introduction of new jet mass definitions, resolutions, and approaches to extracting resolution systematics; 2) the development of calibrations at the jet energy and jet mass scales which are pivotal to all ATLAS analyses utilising jet reconstruction; 3) the introduction, and subsequent performance analysis of, new inputs to jet reconstruction and new approaches to pileup mitigation in large-radius jets, achieved via a combination of constituent-level pileup subtraction techniques and new jet groomers beyond the standard approach of jet trimming (such as Soft Drop and variants thereof). Putting all of this together, the work in Part II of this thesis has led to entirely new large-radius jet definitions, and jet mass calculations, within the ATLAS collaboration which will enhance the precision of new particle physics measurements, and

increase our sensitivity to new physics scenarios in dedicated searches, both in supersymmetry and in other exotic contexts.

In the final part of this thesis, the techniques of jet reconstruction have been applied in a benchmark search for new physics, motivated through models where gluinos are directly produced and subsequently decay through complex cascades which generate final states with many hadronic jets and $E_{\mathrm{T}}^{\mathrm{miss}}$. With 36 fb^{-1} of ATLAS data, the quest for new matter in these multi-jet final states has not produced any statistically significant deviation from the predictions of the Standard Model, but this has allowed the author to constrain a variety of models of new physics, including the masses of gluinos in R-parity-conserving supersymmetric models to up to 1.8 TeV. Further to this, R-parity-violating supersymmetric scenarios, where the lightest supersymmetric particle also decays, have been uniquely constrained by capitalising on the low $E_{\mathrm{T}}^{\mathrm{miss}}$ final-states permitted in multi-jets signal regions. In this context, the mass of stop particles have been constrained to around 1.4 TeV and, for the first time at the LHC, new constraints on the spectrum of the R-parity-violating decay vertex have been established, indicating that limits on both prompt and long-lived RPV SUSY scenarios can be examined up to moderate gluino masses around (and beyond) 1.5 TeV. Whilst these constraints could begin to be interpreted as being in tension with the expected supersymmetric spectra arising in natural models of SUSY[1], it is worth bearing in mind that it is the job of an experimentalist to make detailed, thorough, and trustworthy measurements of our observable universe so that current physics hypotheses can be assessed, and new and exciting theories of physics can emerge from the conclusions, ready for a new set of stringent and dedicated measurements to be made. In this spirit, experimentalists at the LHC are now readying themselves for the third run of data-taking in the early 2020s, where we shall continue to test the Standard Model of particle physics and resume our quest for the production of new (and unexpected) forms of matter at ever higher luminosities and collision energies.

[1]It is important to recall that there is also a degree of model-dependence in the limits discussed in this thesis.

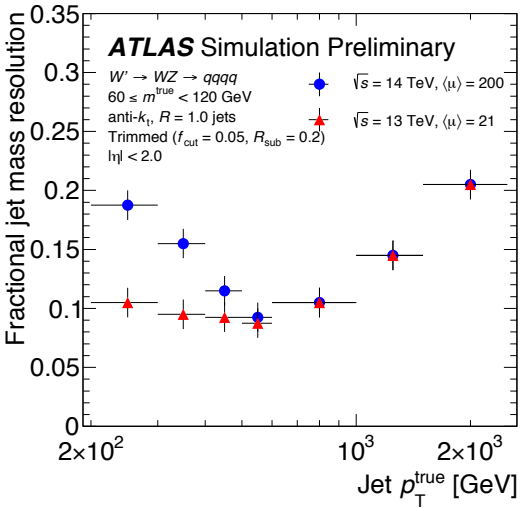
Appendices

A

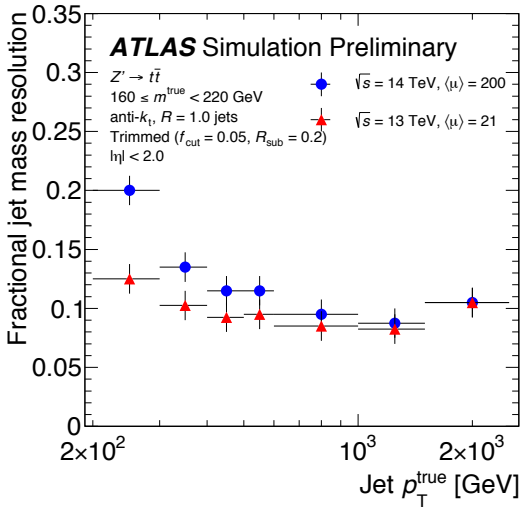
Appendix: Jet Mass Resolution at the HL-LHC

One important motivation for understanding the pileup robustness of different jet definitions, and the impact of different jet input and grooming procedures on jet reconstruction, is quantifying how the jet performance evolves in environments where the pileup contribution is significant. An important use-case is for the High-Luminosity LHC (HL-LHC) [169], which is expected to come online in the late 2020s. The average number of interactions per bunch-crossing in the HL-ATLAS experiment is projected to be around 200, significantly larger than the $\mathcal{O}(50)$ recorded during Run 2. Such dramatic increases in pileup will only deteriorate the resolution of jet properties measured in the calorimeter, if for example the current $R = 1.0$ ATLAS jet definition (`AntiKt10LCTopoTrimmedPtFrac5SmallR20`) remains unmodified. This deterioration is illustrated in Figure A.1, where the jet mass resolution is shown with respect to the jet $p_{\mathrm{T}}^{\mathrm{true}}$, for simulations of the same QCD, W/Z-jet, and top jet topologies considered in Chapters 6 and 7. The calorimeter mass definition is used, and the simulated jets have been calibrated at the jet energy and jet mass scales in both the Run 2 ($\sqrt{s} = 13$ TeV) and HL-LHC ($\sqrt{s} = 14$ TeV) cases. The large degradation in the resolution at higher transverse momenta, due to the finite granularity of the calorimeter, is evident (and expected). The red shows the expected JMR using simulated Run 2 2015 + 2016 data conditions, where the average number of interactions per bunch-crossing is 21; in the blue, HL-LHC conditions are simulated with an average of 200 interactions per bunch-crossing. A detector simulation of the upgraded ATLAS inner tracker detector (ITk) [170] is also used for the HL-LHC JMR[1]. In all

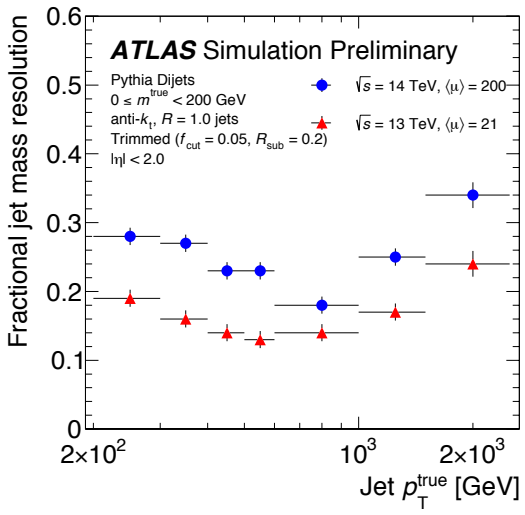
[1]Technical point: the ITk Step-1.9 detector geometry is used in the HL-LHC reconstruction simulation, consistent with what was being adopted for other ITk detector and physics studies at the time.



(a) $W' \to WZ \to q\bar{q}q\bar{q}$



(b) $Z' \to t\bar{t}$



(c) QCD jets

Figure A.1: The fractional calorimeter jet mass resolution as a function of the truth jet transverse momentum. The resolution is determined from jets produced via beyond Standard Model $W' \to WZ \to q\bar{q}q\bar{q}$ and $Z' \to t\bar{t}$ signals, and the QCD jet process, using `AntiKt10LCTopoTrimmedPtFrac5SmallR20` reconstruction. The JMR is shown separately for typical ATLAS Run 2 conditions and for HL-LHC conditions with an upgraded ATLAS detector based on the ITk Step-1.9 detector geometry (red and blue respectively). The jets are reconstructed from LCW-calibrated calorimeter cell clusters, and separate Monte Carlo-based jet energy scale and jet mass scale calibrations have been derived for the Run 2 jets and HL-LHC jets, with good closure in the mass response observed after calibration.

three topologies, the JMR performance at low $p_{\mathrm{T}}^{\mathrm{true}}$ is significantly worse in the high pileup environment. This is a dramatic illustration of the limitations in current jet reconstruction at ATLAS, and demonstrates the need for more effective pileup mitigation at the level of the jet constituents if we are to utilise jet substructure techniques at the HL-LHC with improved or comparable systematic uncertainties to what we currently have.

B

Appendix: Soft Drop Combined Mass

Following the latest inputs and grooming scan, the jet definition using $z_{\text{cut}} = 0.10$, $\beta = 1.0$ Soft Drop on `CSSKLCTopo` inputs has been studied in the combined jet mass definition. The combined jet mass weights are again constructed from a sample of Pythia 8 Monte Carlo QCD dijets, and used to construct the combined mass from the separate calorimeter and track-assisted masses. The resulting jet mass resolutions are shown in Figure B.1, where both the combined and calorimeter masses for the standard trimmed and Soft Drop jets have been overlaid. The Soft Drop jet mass resolution will not necessarily be smaller than the trimmed; a more interesting comparison is the pileup dependence of the combined mass. It was observed in Chapter 6 that the combined mass performance at lower jet $p_{\text{T}}^{\text{true}}$ could be non-optimal, in part due to the presence of pileup in the calculation of the weights.

In Figures B.2 and B.3 the Soft Drop and trimmed combined masses are shown. Three different cases are considered: first an inclusive case where there is no μ-cut applied to the QCD dijets sample which is used to derive the combined mass weights; then the low pileup case, where only dijet events with average $\mu < 25$ are used to calculate the combined mass weights; and finally the high pileup case, where only dijet events with average $\mu > 35$ are considered in the combined mass weight calculation. It can be seen that, in the case of trimmed jets, the combined mass resolution at lower transverse momenta significantly worsens when higher pileup events are used to construct the weights, as anticipated from previous observations. This is not the case for the Soft Drop combined mass, where the combined mass resolution is superior

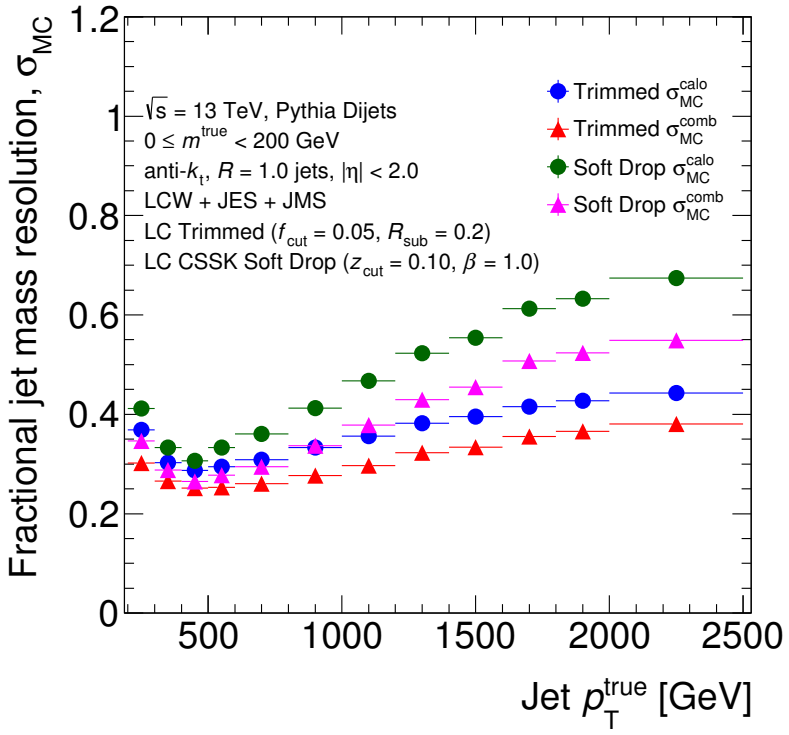
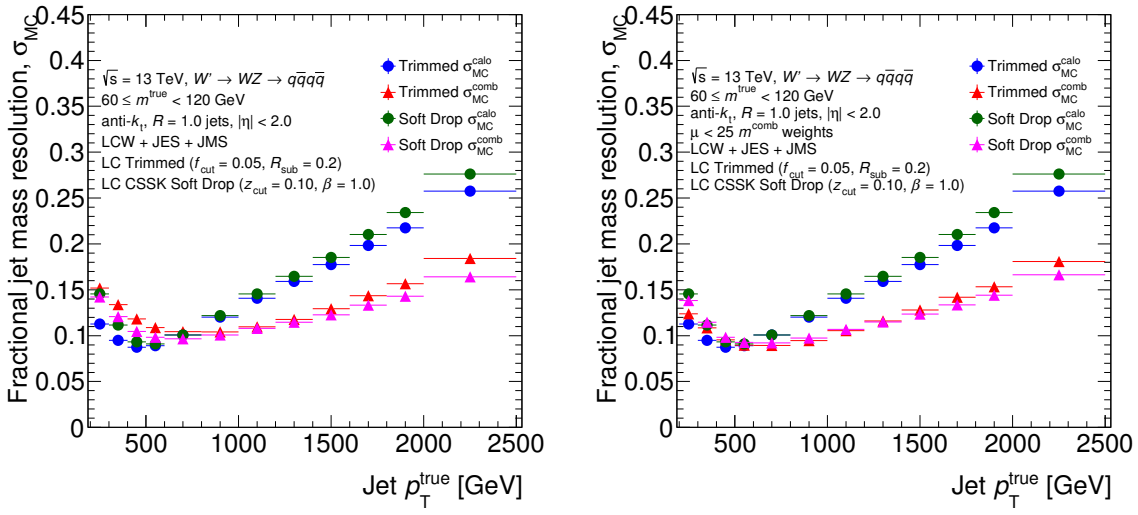


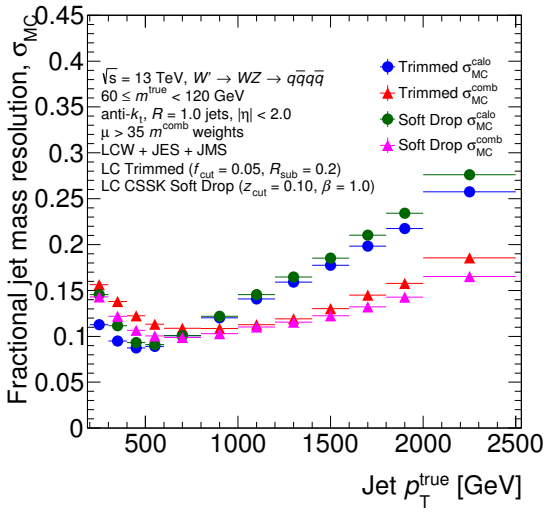
Figure B.1: The QCD dijets mass resolution for four different cases: standard trimmed jets calorimeter mass and combined mass; Soft Drop `CSSKLCTopo` jets calorimeter mass and combined mass.

to the calorimeter mass resolution in all three pileup scenarios, indicating that the inclusion of constituent subtraction in the `LCTopo` inputs directly improves the pileup robustness of the combined mass weights, and hence the mass resolution itself.



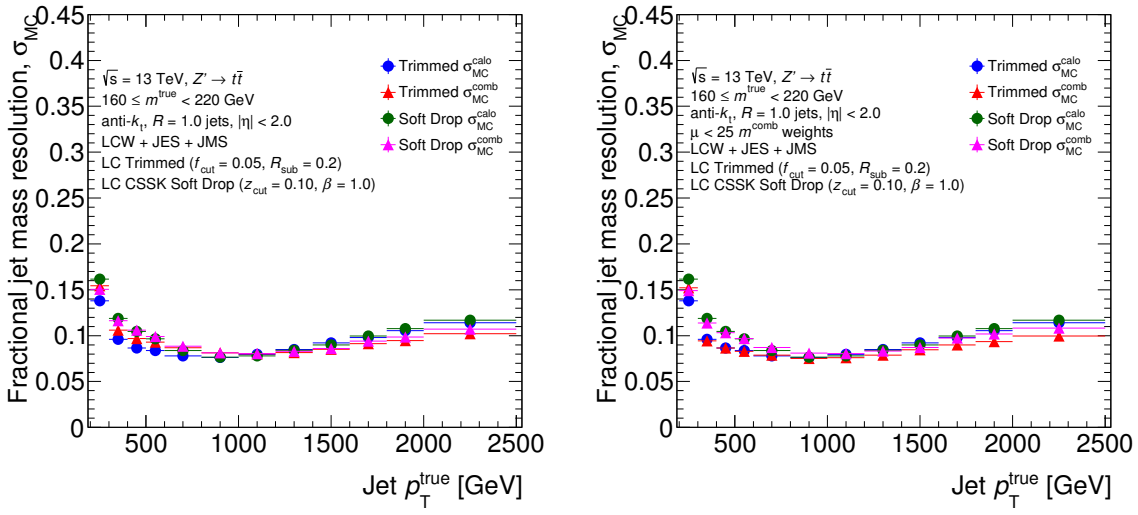
(a) Inclusive $\langle\mu\rangle$.

(b) $\langle\mu\rangle < 25$.



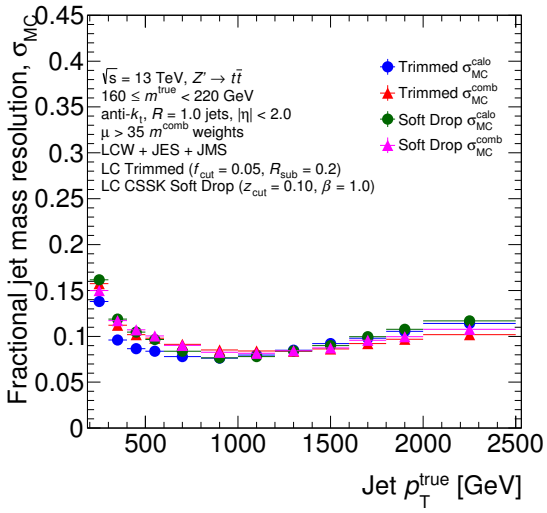
(c) $\langle\mu\rangle > 35$.

Figure B.2: The W/Z-jet mass resolution for different ranges in $\langle\mu\rangle$ and four different cases: standard trimmed jets calorimeter mass and combined mass; Soft Drop `CSSKLCTopo` jets calorimeter mass and combined mass.



(a) Inclusive $\langle\mu\rangle$.

(b) $\langle\mu\rangle < 25$.



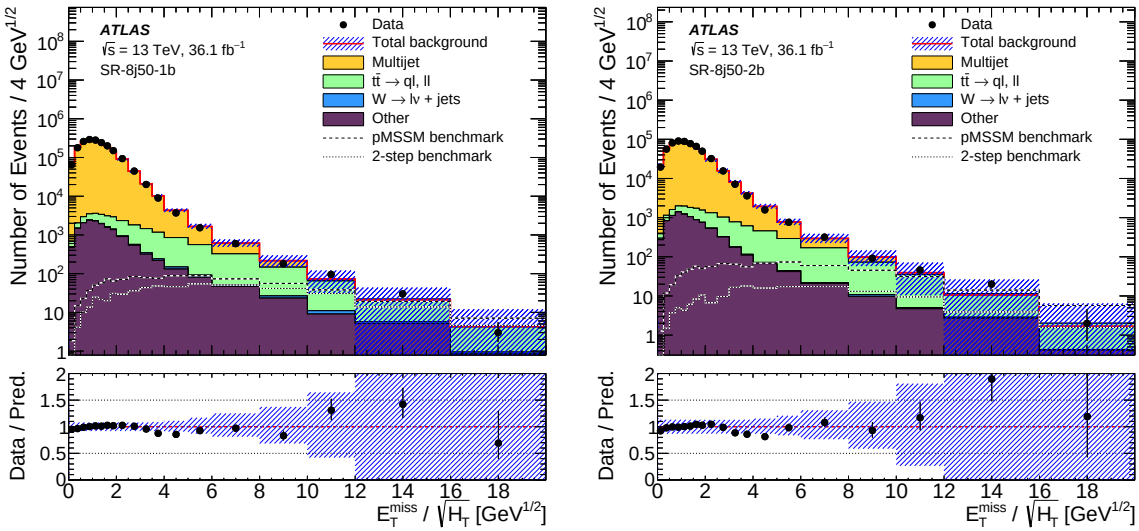
(c) $\langle\mu\rangle > 35$.

Figure B.3: The top jet mass resolution for different ranges in $\langle\mu\rangle$ and four different cases: standard trimmed jets calorimeter mass and combined mass; Soft Drop `CSSKLCTopo` jets calorimeter mass and combined mass.

C

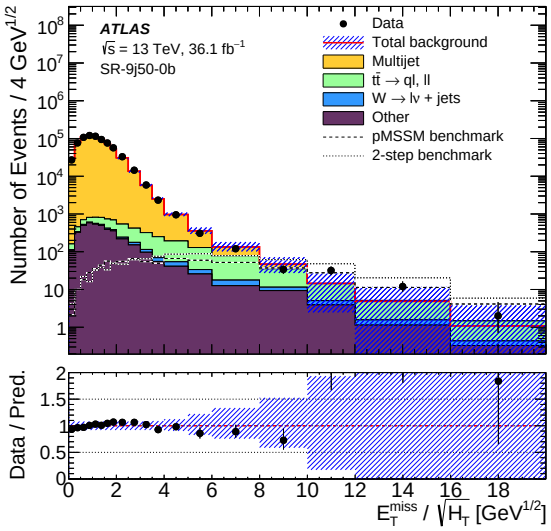
Appendix: $E_{\mathrm{T}}^{\mathrm{miss}}/\sqrt{H_{\mathrm{T}}}$ Distributions in all Multi-Jets Signal Regions

This appendix summarises the remaining set of $E_{\mathrm{T}}^{\mathrm{miss}}/\sqrt{H_{\mathrm{T}}}$ distributions in each of the signal regions in the multi-jets search which have not already been described in the main body of this thesis. Collectively these plots were included as material in [70]. The x-axes have variable bin widths (up to 4 GeV$^{1/2}$) and the y-axes indicate the numbers of events per 4 GeV$^{1/2}$. The plots show the major Monte Carlo backgrounds (semileptonic $t\bar{t}$ and W + jets), the data-driven template QCD prediction, and the other combined backgrounds, together with 2015 + 2016 ATLAS data. Two 'benchmark' signal points are also added, for the two-step and pMSSM models, with parameters $(m_{\tilde{g}}, m_{\tilde{\chi}_1^0}) = (1400, 200)$ GeV and $(m_{\tilde{g}}, m_{\tilde{\chi}_1^\pm}) = (1400, 200)$ GeV respectively.

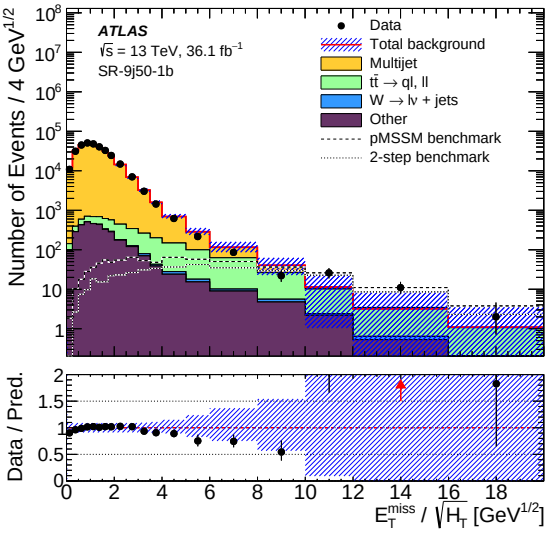


(a) SR-8ij50-1ib

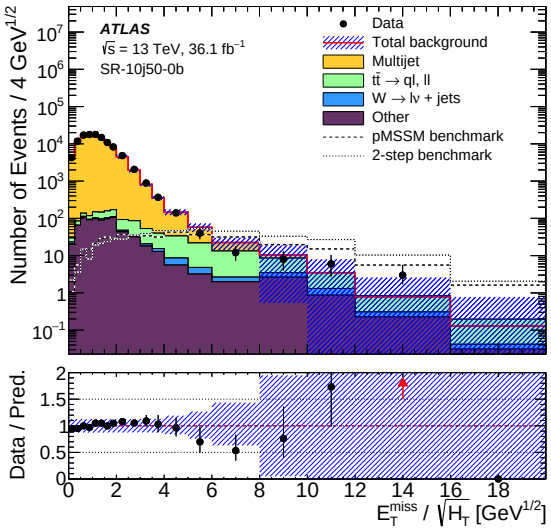
(b) SR-8ij50-2ib



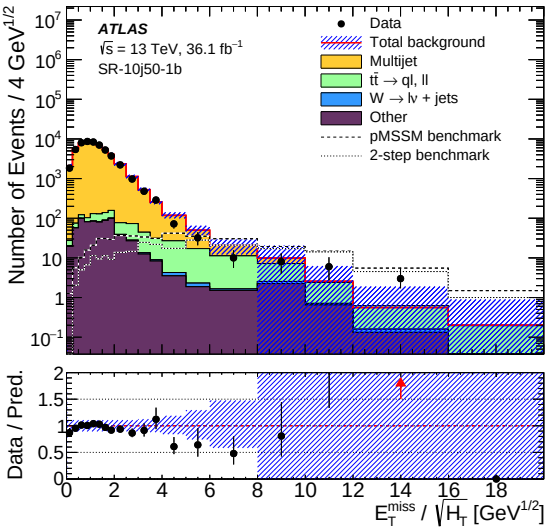
(c) SR-9ij50-0ib



(d) SR-9ij50-1ib

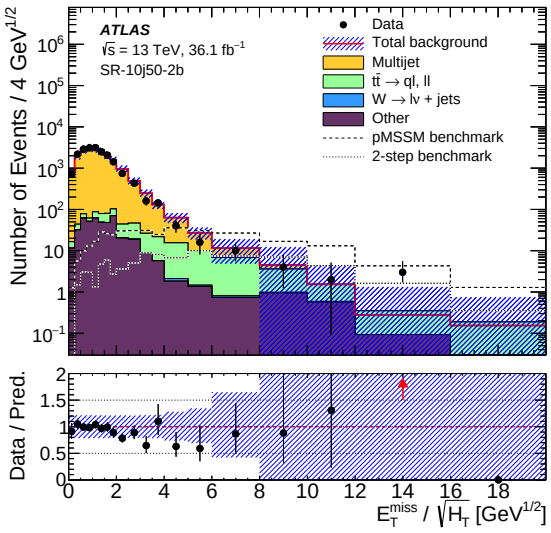


(e) SR-10ij50-0ib

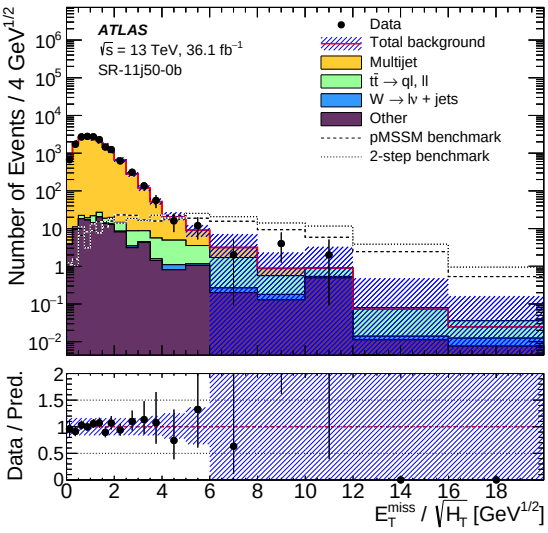


(f) SR-10ij50-1ib

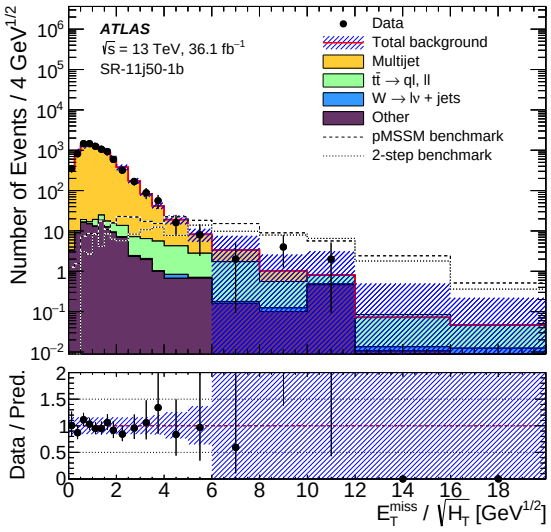
Figure C.1: $E_{\mathrm{T}}^{\mathrm{miss}}/\sqrt{H_{\mathrm{T}}}$ distributions in various signal regions, where each region is indicated in a subfigure.



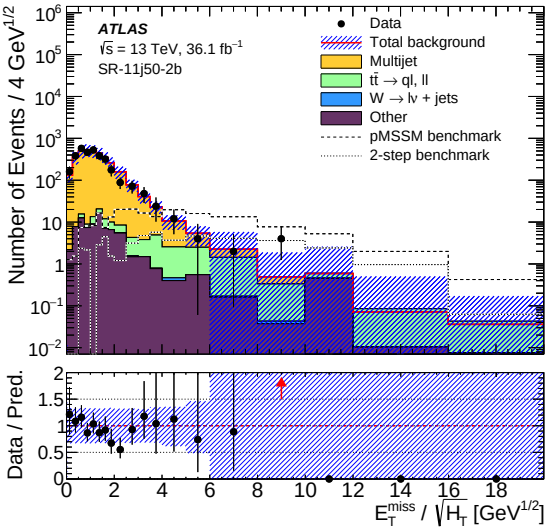
(a) SR-10ij50-2ib



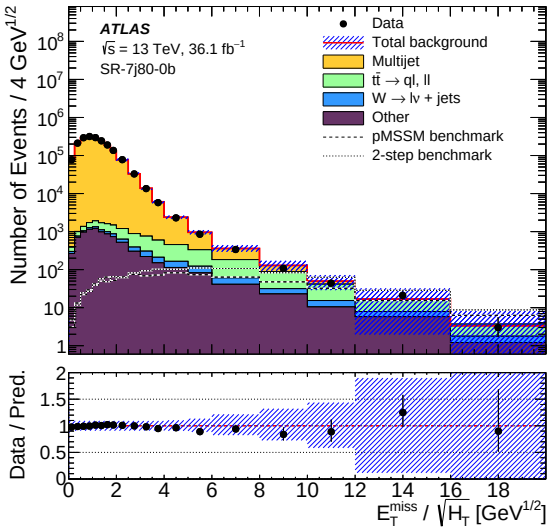
(b) SR-11ij50-0ib



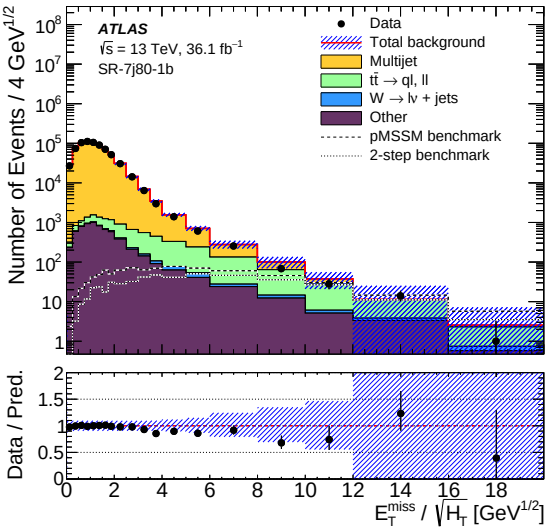
(c) SR-11ij50-1ib



(d) SR-11ij50-2ib

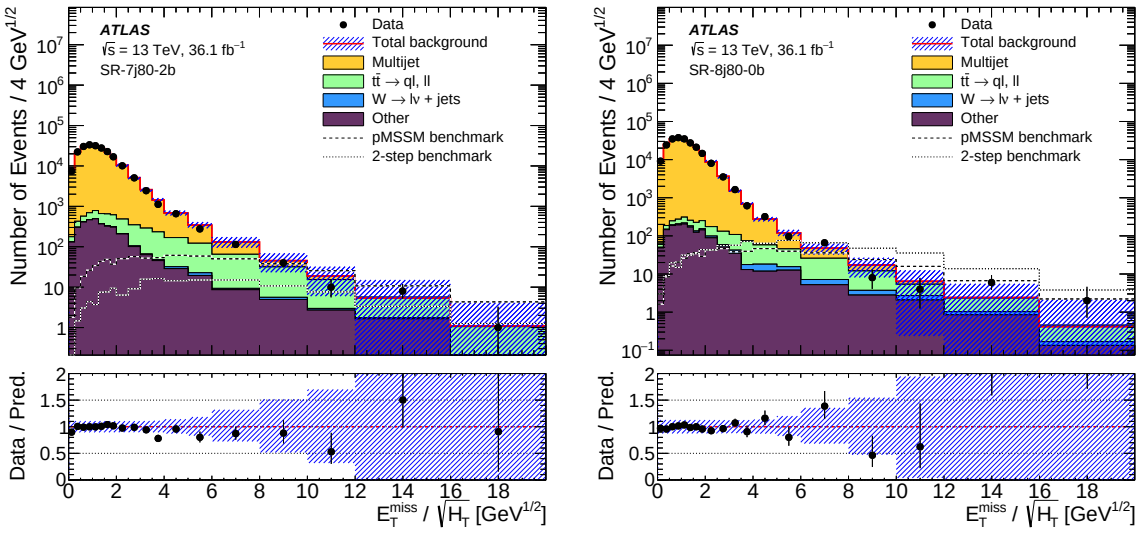


(e) SR-7ij80-0ib



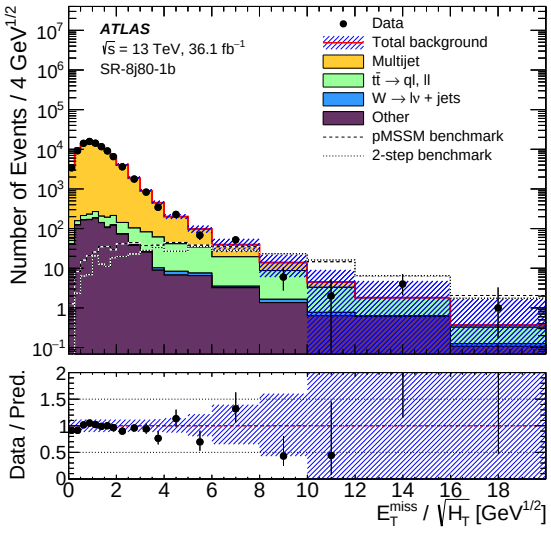
(f) SR-7ij80-1ib

Figure C.2: $E_{\mathrm{T}}^{\mathrm{miss}}/\sqrt{H_{\mathrm{T}}}$ distributions in various signal regions, where each region is indicated in a subfigure.

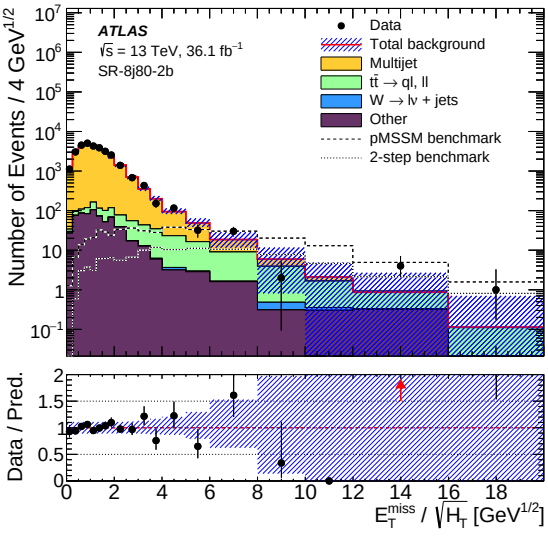


(a) SR-7ij80-2ib

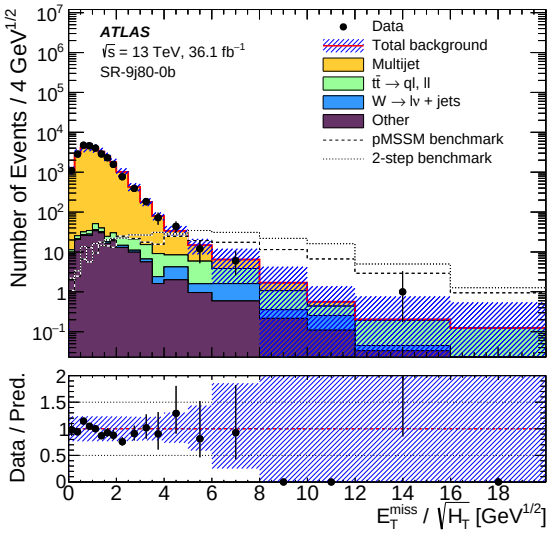
(b) SR-8ij80-0ib



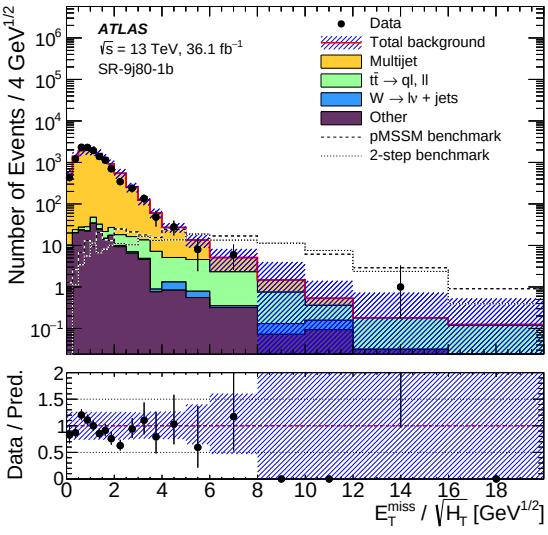
(c) SR-8ij80-1ib



(d) SR-8ij80-2ib

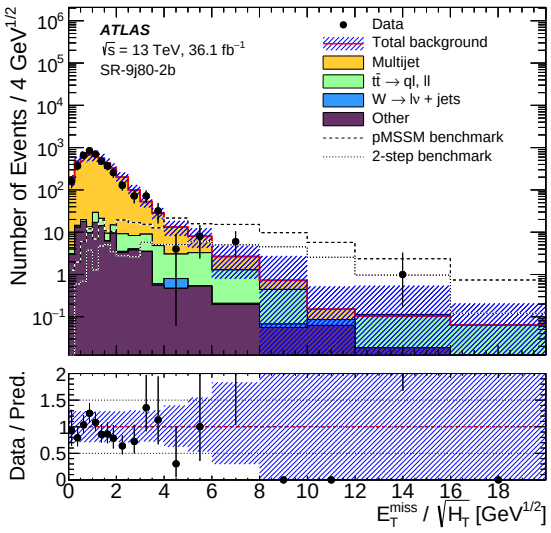


(e) SR-9ij80-0ib

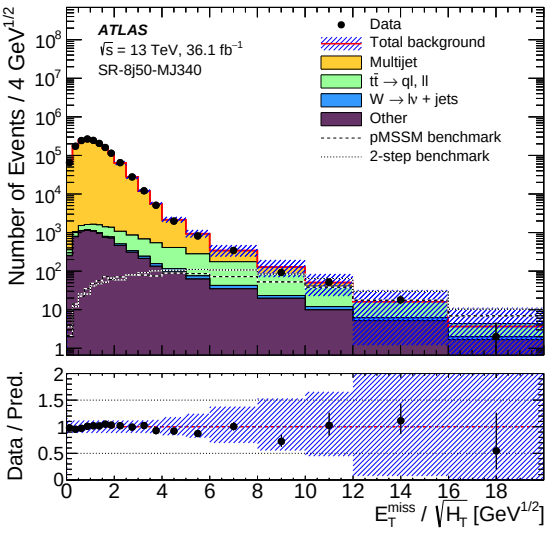


(f) SR-9ij80-1ib

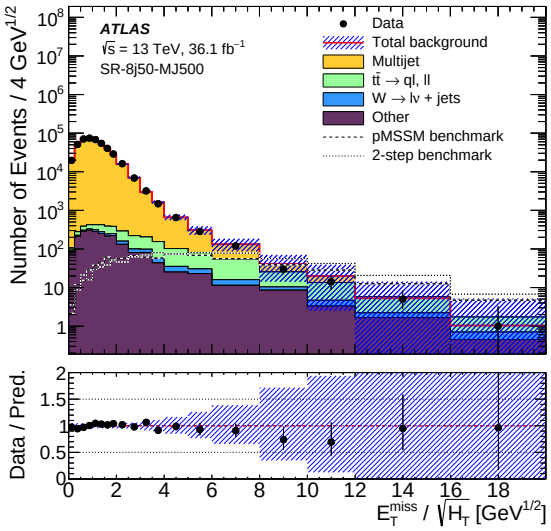
Figure C.3: $E_{\mathrm{T}}^{\mathrm{miss}}/\sqrt{H_{\mathrm{T}}}$ distributions in various signal regions, where each region is indicated in a subfigure.



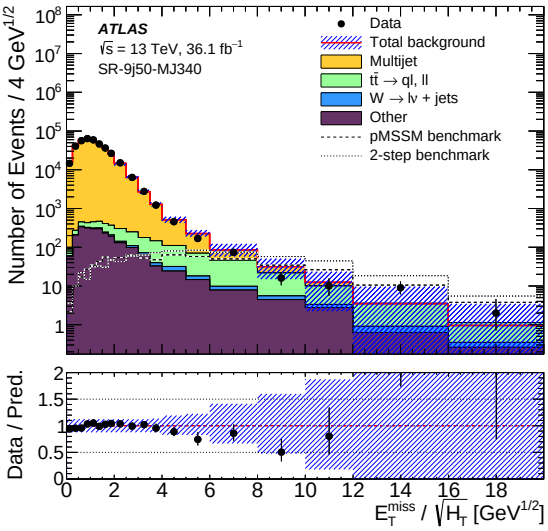
(a) SR-9ij80-2ib



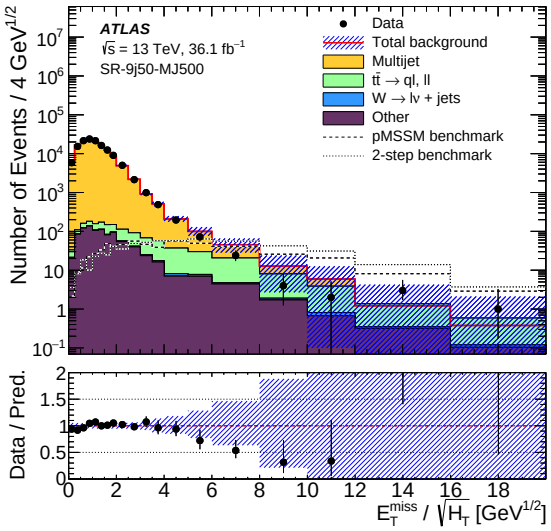
(b) SR-8ij50-0ib-MJ340



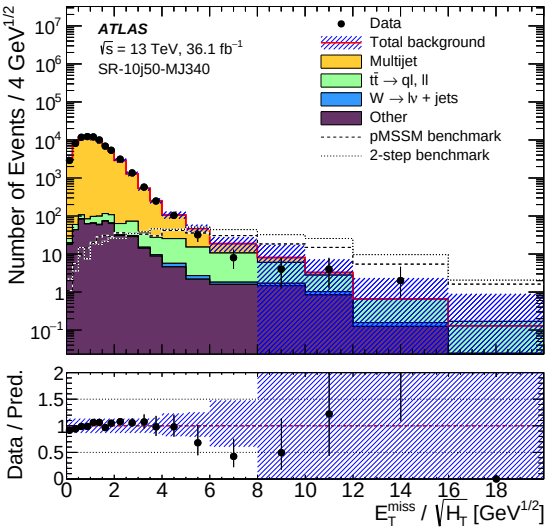
(c) SR-8ij50-0ib-MJ500



(d) SR-9ij50-0ib-MJ340



(e) SR-9ij50-0ib-MJ500



(f) SR-10ij50-0ib-MJ340

Figure C.4: $E_{\mathrm{T}}^{\mathrm{miss}}/\sqrt{H_{\mathrm{T}}}$ distributions in various signal regions, where each region is indicated in a subfigure.

Glossary

Large Hadron Collider (LHC): The world's largest and most powerful particle accelerator, situated at CERN in Geneva.

Beyond the Standard Model (BSM): New physics theories which extend the Standard Model in some way.

Supersymmetry (SUSY): A class of BSM theory that predicts the existence of a symmetry between fermion and boson states, generating superpartner particle states.

R**-parity**: A multiplicative quantum number that is $+1$ for Standard Model particles and -1 for SUSY particles.

R**-parity conserving (RPC)**: A particulary class of SUSY model in which the R-parity quantum number is conserved at each interaction vertex.

R**-parity violating (RPV)**: A particulary class of SUSY model in which the R-parity quantum number may be violated at a given interaction vertex.

Sparticle: A superpartner of a Standard Model particle, hence a member of the SUSY sector in a SUSY theory.

Lightest supersymmetric particle (LSP): The lightest sparticle in a given SUSY model. In RPC models this will be stable and in cases where it is only weakly interacting it could be a dark matter candidate.

Gluino: Superpartner of the gluon.

Higgsino: Superpartner of a Higgs boson.

Wino: Superpartner of a W boson.

Bino: Superpartner of a B boson.

Local cluster weighting (LCW): A calibration applied to topoclusters correcting e.g. non-compensation in the hadronic response.

Jet mass (energy) response: the ratio of the reconstructed jet mass (energy) to the truth-level jet mass (energy) when considering a Monte Carlo calibration of the mass (energy).

Gaussian smearing: A procedure for modifying the simulated jet resolution to account for additional detector effects, using a multiplicative Gaussian function of the resolution.

Forward-folding: A procedure for modifying the simulate jet resolution to account for additional detector effects, using a linear transformation which includes a term derived directly from studies in data.

Closure: The degree of closeness of the calibrated jet mass (or energy) response to unity, where the deviation from unity quantifies the non-closure.

Grooming: an algorithm which removes soft and wide-angled contributions from a jet, typically to remove pileup contributions. Examples include trimming, pruning, and Soft Drop.

Missing transverse momentum, $E_{\mathrm{T}}^{\mathrm{miss}}$: an event-level variable, which is the magnitude of the negative vector sum of all visible momenta in the event.

List of Figures

1.1 The connections between theoretical developments, jet physics, and experimental searches for new physics described within this thesis. 3

2.1 Feynman diagram corresponding to the interaction vertex obtained from the $-ie\bar{\psi}\gamma^{\mu}A_{\mu}\psi$- term in the QED Lagrangian. 10

2.2 The leading order loop corrections for the quark and gluon self-energies, where the effects of the vacuum polarisation are included for the gluon loop. These are important 1-loop corrections to the quark and gluon propagators that contribute to the running of the strong coupling. 12

2.3 A summary of different measurements of the running of the strong coupling with energy scale Q. Results are shown for deep inelastic scattering experiments, electron-positron annihilation experiments, as well as pp and $p\bar{p}$ collider experiments. From [6]. 13

2.4 The weak interaction vertex for electron and electron-antineutrino transitions. The factor of $(1 - \gamma^5)$ arises from the need to project out the left-handed components of the particle wavefunctions, since the weak interaction does not conserve parity. 15

2.5 The 'Mexican hat' Higgs potential in the complex plane, after spontaneous symmetry breaking where the vacuum expectation value of the Higgs field is generated. The circle traces out the minimum of the potential, with radius related to the magnitude of the Higgs vacuum expectation value. In this complex plane there are both radial and circular degrees of freedom, again generated by spontaneous symmetry breaking. 18

3.1 The two alternative neutrino mass hierarchies between the mass eigenstates of the neutrinos, ν_1, ν_2, and ν_3. For each mass eigenstate, the relative contributions from the weak eigenstates (ν_e, ν_μ, ν_τ) are also provided as a function of the (currently unknown) CP-violating phase, $\delta_{\mathtt{CP}}$. The mass splittings, as determined by solar and atmospheric neutrino experiments, are also indicated. From [24]. . 27

3.2 The three complementary approaches to dark matter searches. LHC experiments such as ATLAS focus on the production of dark matter from collisions of known particles. From [26]. 28

3.3 The running of the SM couplings (black, dotted lines) and running of the same couplings in a BSM extension based on supersymmetry (solid, coloured lines). In this case, the unification is achieved via the minimal supersymmteric standard model (MSSM), a supersymmetric model which we shall shortly review. The masses of new SUSY particles are varied as a common threshold between 750 GeV and 2.5 TeV, and the $\alpha_S(m_Z)$ between 0.117 and 0.120. From [22]. 30

3.4 Feynman diagrams for the first-order loop corrections to the Higgs mass, coming from a fermion loop (left) and a scalar loop (right) respectively. 31

3.5 Leading-order gluino production diagrams in proton-proton collisions, where the t-channel diagram on the right is suppressed if squark masses are much greater than gluino masses. 39

3.6 Production cross-sections for several different SUSY processes, from [42], where the gluino production cross-section assumes all squarks are decoupled. For a given particle mass, the relatively large gluino production cross-section is evident. These cross-sections were calculated using Reference [43]. 40

4.1 A cut-away schematic of the ATLAS detector at the Large Hadron Collider, reproduced from Reference [45], and showing the various detector subsystems. The forward detector subsystems, which are not considered in any analysis in this thesis, are not shown. 43

4.2 The LHC accelerator complex at CERN, including the various pre-accelerators. Shown are the linear accelerator (LINAC 2), the Proton Synchrotron Booster, the Proton Synchrotron (PS), Super Proton Synchrotron (SPS), and the Large Hadron Collider (LHC). The four main experiments on the LHC ring are also shown, although their relative positions on the ring are not to scale. Image taken from Reference [50]. 45

4.3 The luminosity weighted distribution of the mean number of interactions per bunch-crossing (*pileup*), split by year throughout Run 2 of ATLAS. Taken from [51]. 46

4.4 The cumulative luminosity vs. day in year delivered to ATLAS during stable beams for high-energy pp processes, shown for the various years through Run 1 and Run 2 data-taking. Taken from [51]. 46

4.5 A cut-away view of the ATLAS inner detector system used in Run 2, including the IBL and pixel subsystems, the SCT, and the TRT. The IBL subsystem is closest to the beam and was added for Run 2 data-taking. The end-caps are shown in the upper panel, and the barrel region in the lower panel. 49

4.6 A cut-away schematic of the electromagnetic and hadronic calorimeter systems in ATLAS, including the barrel and end-cap regions. Image taken from [45]. . . 51

4.7 A cut-away schematic of the muon spectrometer and its various detector subsystems. Image taken from [45]. 53

5.1 A toy model of simple two-two scattering, where the additional processes of multi-partonic interactions (underlying event) and showering of the partons produced from the initial scatter give rise to complex final states consisting of multiple hadrons (commonly pions), leptons from hadronic decay, and radiation. Reconstructing and reinterpreting such processes measured using calorimetry is one of the major challenges in contemporary hadron-hadron collider experiments. 64

5.2 Feynman diagrams for the s-channel $q\bar{q} \to q'\bar{q}'$ process, showing the hard-scatter and ISR gluon emission. 65

5.3 A candidate top event observed during the 2011 ATLAS data run, where a leptonically-decaying top candidate is inferred from lepton identification in the electromagnetic calorimeter, and a hadronically-decaying top candidate is inferred using information from both the electromagnetic and hadronic calorimeters. The yellow dots correspond to energy deposits in the calorimeters, and the coloured lines correspond to tracks reconstructed in the inner detectors. Jet reconstruction is a complex process which can utilise information from both tracking and energy deposition. From [78]. 66

5.4 An illustration of some of the different Monte Carlo inputs which can be given to a jet reconstruction algorithm: partons, leading-order and beyond, can form parton jets; jets can also be formed from the partonic shower, and finally the full chain of hadrons at truth-level. This latter example gives rise to truth-level jets. From [80]. 67

5.5 The topoclustering process, from initial seeding and growth, to the final topoclusters, in spherical polar coordinates based on the ϕ and θ position of the initial energy deposits in the calorimeter. This clustering is based on dijet events in ATLAS 2010 data-taking conditions simulated with PYTHIA 6 Monte Carlo. The thick black lines indicate the topocluster boundaries. From [82]. 70

5.6 An illustration of collinear safe and unsafe jet reconstruction algorithms. From [68]. 71

5.7 An illustration of infrared safe and unsafe jet reconstruction algorithms. From [68]. 72

5.8 The three different classes of sequential recombination algorithms, used to construct $R = 1.0$ jets from inputs defined in the ϕ and rapidity (y) frames for some simulation of events. The final jets are signified by the coloured objects in the $y - \phi$ plane, and the inputs have heights proportional to their transverse momenta. From [68]. 73

5.9 The number of input clusters, N, against the computational time, t, in seconds, for various different jet reconstruction algorithms. The sequential recombination algorithms are the most computationally efficient for a given number of input clusters, often by several orders of magnitude. From [68]. 74

5.10 A cartoon summarising the basic arrangement of clusters in the anti-k_t algorithm adopted by ATLAS. Clusters are grouped in $\eta - \phi$ space based their separation from other clusters and from the beam itself. This distance measure depends on the geometrical separation of clusters in $\eta - \phi$, the inverse square of the cluster transverse momenta, and the desired radius of the final jets. 75

5.11 A summary of the complete jet energy scale calibration chain for `AntiKt4EMTopo` jets in ATLAS. The initial origin correction is followed by the Monte Carlo JES calibration. After this a correction to the flavour-dependence of the jet is applied using Monte Carlo, before a final *in situ* calibration is determined using events measured in ATLAS data. From [91]. 77

5.12 The N_{PV} and μ dependencies before and after applying the full jet area and residual corrections. Applying the full residual correction, where the correction is derived in bins of the jet η^{det}, almost entirely removes the pileup dependence. These corrections are derived for `AntiKt4EMTopo` jets using Monte Carlo dijet events produced in Pythia, at reconstruction and truth-level. From [91]. 80

5.13 The average jet energy response and residual jet η values, before any η-JES correction has been applied, as function of the detector η. The JES correction is obtained from the inverse of the energy response. From [91]. 82

5.14 The five observables used to derive the GSC correction for `AntiKt4EMTopo` jets, using dijets events simulated in Pythia, at both the truth and reconstruction-levels. These corrections are applied sequentially, with the Tile0 correction applied first and the muon segments correction apply at the end. The responses shown are before any GSC-correction is applied. From [91]. 85

5.15 The η-intercalibration correction as a function of the detector η, for `AntiKt4EMTopo` jets using the 2015 ATLAS data-set. The smallest corrections are in the central region, which is dominated by statistical uncertainties. The systematic uncertainty components, such as the difference between the Powheg + Pythia 8 (used as the nominal sample) and Sherpa dijets Monte Carlo, are discussed in Section 5.3.3. The final correction (thick black line) is obtained from a smoothed fit to the intercalibration points (open circles). These corrections are obtained in different intervals of $p_{\mathrm{T}}^{\mathrm{avg}}$: $40 - 55$ GeV in this case. From [91]. 87

5.16 The geometry of an event considered in the direct balance method. The balance is calculated from the transverse momentum ratio of the leading jet J_1 and some well-defined reference Z boson or photon, at some angle $\Delta\phi$ to J_1. From [91]. . . 88

5.17 The MPF and direct balance methods, using a Z and γ reference respectively, as a function of the jet transverse momentum, measured in the central region of the detector, $|\eta^{\text{det}}| < 0.8$. The balances are calculated using events in 2015 data (black), Monte Carlo POWHEG + PYTHIA (red), and Monte Carlo SHERPA (blue). The lower pad ratio gives the size of the correction to apply to data, where the POWHEG + PYTHIA Monte Carlo is the nominal sample used for the correction and the differences between this and the SHERPA sample provides a handle on the generator uncertainty. From [91]. 89

5.18 The final combined *in situ* correction for `AntiKt4EMTopo` jets measured in 2015 ATLAS data, as a function of the jet transverse momentum. The relative size of the correction is within a few percent of unity. The major uncertainties are discussed briefly in Section 5.3.3. From [91]. 90

5.19 A cartoon illustrating the trimming a circular jet. From [90]. 90

5.20 The `AntiKt10LCTopoTrimmedPtFrac5SmallR20` POWHEG + PYTHIA dijet Monte Carlo jet energy response as a function of η^{reco} for different bins of E^{true} given by different colours and patterns of points. The closure for jets with energies above 300 GeV is well below even the 1 % level. 92

5.21 The fractional JES uncertainties for `AntiKt4EMTopo` jets as a function of the jet transverse momentum and η. The total uncertainty comes from all of the separate contributions of the Monte Carlo JES and *in situ* calibration uncertainties. From [91]. 93

6.1 A summary of the complete jet mass scale calibration chain for `AntiKt4EMTopo` and `AntiKt10LCTopoTrimmedPtFrac5SmallR20` jets in ATLAS. Currently only the larger-radius jets are calibrated at the *in situ* level. Both the calorimeter and track-assisted jet mass definitions need to be separately calibrated. 99

6.2 Example calorimeter and track-assisted jet mass responses, fitted with a Gaussian (red line), for `AntiKt4EMTopo` jets. These jets have been calibrated up to the jet energy scale and are shown for a particular m^{true}, $p_{\text{T}}^{\text{true}}$, and $|\eta^{\text{det}}|$ bin, with the ranges given in each plot. The mean and width are provided, and they depend on the $p_{\text{T}}^{\text{true}}$. 100

6.3 Example calorimeter and track-assisted jet mass responses, fitted with a Gaussian (red line), for `AntiKt10LCTopoTrimmedPtFrac5SmallR20` jets. These jets have been calibrated up to the jet energy scale and are shown for a particular m^{true}, $p_{\text{T}}^{\text{true}}$, $|\eta^{\text{det}}|$ bin, with the ranges given in each plot. The mean and width are provided, and they depend on the $p_{\text{T}}^{\text{true}}$. 101

6.4 The uncalibrated (dashed line) and calibrated (solid line) reconstruction-level jet masses for the calorimeter, track-assisted, and track-only jet mass definitions. The calibrations are obtained from QCD dijets Monte Carlo, and applied to a simulated sample enriched in W/Z jets. `AntiKt10LCTopoTrimmedPtFrac5SmallR20` jets are considered. From [98]. 103

6.5 The average calorimeter and track-assisted jet mass response from QCD dijets as a function of the truth-level transverse momentum, before and after JMS calibration. Three different truth-level jet mass bins are shown, and the closure following calibration is within 2 % in all cases. `AntiKt10LCTopoTrimmedPtFrac5SmallR20` jets are considered. From [98]. 104

6.6 The final combined *in situ* correction for `AntiKt10LCTopoTrimmedPtFrac5SmallR20` jets measured in 2015+2016 ATLAS data, as a function of the jet transverse momentum. The relative size of the correction is within a few percent of unity. From [104]. 105

6.7 The fractional jet mass resolution for the `AntiKt10LCTopoTrimmedPtFrac5SmallR20` jets used to determine the weights for the combined mass. 108

6.8 The three different jet mass definitions for `AntiKt10LCTopoTrimmedPtFrac5SmallR20` jets for a simulated sample of W/Z jets. The jet mass resolution is provided as a function of the jet transverse momentum, and the resolution improvement from the combined jet mass is clearly seen. 109

6.9 A flow chart summarising the particle flow object reconstruction algorithm. Tracks are first matched to clusters. The charged track component's energy of then subtracted to avoid double-counting the reconstructed energy. Finally the shower calorimeter shower splitting is assessed and additional calorimeter clusters are added in appropriately. From [105]. 111

6.10 The MC-based calorimeter JMR for m^{true} of large-radius jets as a function of jet $p_{\text{T}}^{\text{true}}$, for four jet topologies: W/Z jets, top jets, $H \to b\bar{b}$ jets, and QCD jets. Representative ranges of jet m^{true} or $m^{\text{true}}/p_{\text{T}}^{\text{true}}$ are shown. 115

6.11 The MC-based JMR for m^{calo} and m^{comb} of large-radius jets as a function of truth jet transverse momentum $p_{\text{T}}^{\text{true}}$ in jet topology-dependent m^{true} intervals. Results for four jet topologies are shown: W/Z jets, top jets, $H \to b\bar{b}$ jets, and QCD jets. 116

6.12 The MC-based JMR for m^{calo} of small-radius jets as a function of truth jet transverse momentum $p_{\text{T}}^{\text{true}}$ in jet topology-dependent m^{true} intervals. Results for four jet topologies are shown: W/Z jets, top jets, $H \to b\bar{b}$ jets, and QCD jets. For each topology three distributions are shown for jets reconstructed from EM topoclusters, particle flow objects (PFO) and LCW topoclusters respectively. The black lines indicate the $p_{\text{T}}^{\text{true}}$ scale with $2m^{\text{true}}/(R = 0.4)$, corresponding to the boundaries where the particle of interest can be reasonably contained within the jet. 117

6.13 The jet mass response for large-radius jets. The calorimeter jet mass definition is used. A Gaussian fit has been overlaid for illustrative purposes. The $p_{\text{T}}^{\text{calo}}$ bins have been chosen for demonstrative purposes as examples of distributions which are or are not reasonably Gaussian. 120

6.14 The non-Gaussianity, $(\sigma_{\text{MC}}^{\text{calo,smeared}}/\sigma_{\text{MC}}^{\text{calo}}) - 1.2$, for large-radius jets, as a function of truth jet transverse momentum. The calorimeter jet mass definition is used. A 20.0 % Gaussian smearing has been used. The red line gives the mean value, averaged across the points. The green lines provide the standard deviation across the points. 121

6.15 The non-Gaussianity, $(\sigma_{\mathrm{MC}}^{\mathrm{comb,smeared}}/\sigma_{\mathrm{MC}}^{\mathrm{comb}}) - 1.2$, for large-radius jets, as a function of truth jet transverse momentum. The combined jet mass definition is used. A 20.0 % Gaussian smearing has been used. The red line gives the mean value, averaged across the points. The green lines provide the standard deviation across the points. 122

6.16 The non-Gaussianity, $(\sigma_{\mathrm{MC}}^{\mathrm{calo,smeared}}/\sigma_{\mathrm{MC}}^{\mathrm{calo}}) - 1.2$, for small-radius EM-topocluster-based jets, as a function of truth jet transverse momentum. The calorimeter jet mass definition is used. A 20.0 % Gaussian smearing has been used. The red line gives the mean value, averaged across the points. The green lines provide the standard deviation across the points. 123

6.17 Summary of the large-radius jet *in situ* jet mass response and relative jet mass resolution, parametrised in several bins of the jet transverse momentum. These values were extracted by selecting $t\bar{t}$ events with final states containing a lepton and jets. The first two bins are to be used in W-enriched samples, and the last two bins in top-enriched samples. From [99]. 125

6.18 The ratio $\sigma_{\mathrm{FF}}/\sigma_{\mathrm{MC,corr}}$ for W and top jets using the κ- factors derived in previous studies, and a variable Gaussian smearing, with m^{true} windows as specified in the plots. The jets are smeared in a p_{T}-dependent fashion. The error bars represent the uncertainty due to the finite sample size. 128

6.19 The Δ ratios for W and top quark jets. Mass and p_{T} windows, as specified on the plots, reflect the mass and p_{T} where forward-folding κ parameters have been derived. The jets are smeared in a p_{T}-dependent fashion. The up variations are calculated by considering a smearing of $x_\kappa + \delta x_\kappa$. The down variations are calculated by considering a smearing of $x_\kappa - \delta x_\kappa$. The error bars represent the uncertainty due to the finite sample size. The large difference in the final bin of Δ_- is due to a large $\delta\kappa$ and an undersmeared $x_\kappa - \delta x_\kappa$ (see Table 6.3). 129

7.1 A sample simulated event display with an average of 20 interactions per pp bunch-crossing. The z-axis describes the relative sign of the transverse momentum of each Voronoi cell, relative to its square root. The same event is shown before and after Voronoi subtraction, where in the latter case several negative cells are generated through the subtraction process. From [125]. 134

7.2 A sample simulated event display with an average of 20 interactions per pp bunch-crossing. The z-axis describes the relative sign of the transverse momentum of each Voronoi cell, relative to its square root. The same event is shown after Voronoi subtraction, followed by either an additional Voronoi suppression or Voronoi spreading. From [125]. 135
7.3 A simulated Pythia 8 QCD dijets event-display using pileup conditions typical for the HL-LHC. The effects of Voronoi subtraction and suppression, followed by SoftKiller, are shown. From [125]. 138
7.4 A cartoon of the pruning algorithm, including wide-angle and soft conditions. From [90]. 140
7.5 A cartoon of the 'vanilla' variant of Soft Drop, based on those in [90]. 142
7.6 A cartoon of the recursive variant of Soft Drop. From [120]. 143
7.7 The τ_{32} distribution for a truth-level sample of top jets and QCD jets. The top jets have a smaller average τ_{32} as expected from their strongly three-prong energy structure. From [129]. 147
7.8 The jet mass distributions for $W-$jet and QCD jet topologies, based on applying Soft Drop grooming to `LCTopo` inputs, for both high and low ranges of $p_{\mathrm{T}}^{\mathrm{true}}$. The z_{cut} is fixed at 0.1, and β-values of 0.0, 1.0, and 2.0 are considered. 151
7.9 The jet mass distributions for $W-$jet and QCD jet topologies, based on applying Soft Drop grooming to `LCTopo` inputs, for both high and low ranges of $p_{\mathrm{T}}^{\mathrm{true}}$. The β is fixed at 0.0, and z_{cut}-values of 0.05, 0.1, 0.15, and 0.2 are considered. 152
7.10 The jet mass distribution for `AntiKt10LCTopoTrimmedPtFrac5SmallR20` jets, with the QCD background in red and the $W-$jet signal in blue. A Gaussian fit (black) is applied to the W signal, and the dotted lines indicate the mass window containing 68 % of the desired signal. 153
7.11 Heatmaps showing the efficiency for QCD jets to fall into the smallest $W-$jet mass window containing 68 % of the signal, both at low and higher $p_{\mathrm{T}}^{\mathrm{true}}$. 154
7.12 The jet mass of a reconstructed W boson, in two exclusive bins of N_{PV}, for Soft Drop jets made from unmodified `LCTopo` inputs. 155
7.13 The pileup dependence on the average $W-$jet mass, for different jet inputs and grooming algorithms, in the range $300 < p_{\mathrm{T}}^{\mathrm{true}} < 500$ GeV. The values of the different slopes from the linear fits are provided in Figure 7.15 157

7.14 The pileup dependence on the $W-$jet mass width, for different jet inputs and grooming algorithms, in the range $300 < p_{\mathrm{T}}^{\mathrm{true}} < 500$ GeV. The values of the different slopes from the linear fits are provided in Figure 7.16. 158

7.15 Heatmaps quantifying the slope of the linear fit to the average $W-$jet mass$-N_{\mathrm{PV}}$ dependence. 159

7.16 Heatmaps quantifying the slope of the linear fit to the $W-$jet mass with N_{PV} dependence. 160

7.17 The D_2 distributions for `LCTopo` Soft Drop jets, in exclusive bins of N_{PV}. The change in this jet substructure variable with respect to the average N_{PV} value can be clearly seen. 161

7.18 A schematic demonstrating how the pileup dependence on a jet substructure variable, such as D_2, can be assessed in a meaningful way. The distribution of all events is first split into bins of N_{PV}, and the substructure variable is studied in each bin. 162

7.19 The pileup dependence on the $W-$jet D_2 efficiency, for different jet inputs and grooming algorithms, in the range $300 < p_{\mathrm{T}}^{\mathrm{true}} < 500$ GeV. The values of the different slopes from the linear fits are provided in Figure 7.20. 163

7.20 Heatmaps quantifying the slope of the linear fit to the $W-$jet D_2 efficiency with N_{PV} dependence. 164

7.21 Heatmaps quantifying the slope of the linear fit to the QCD jet D_2 efficiency with N_{PV} dependence. 165

7.22 The JES and JES+JMS-corrected jet mass responses for the shortlisted jet collections, selecting events around the W mass and in the lower $p_{\mathrm{T}}^{\mathrm{true}}$ range 300 − 500 GeV. 167

7.23 The JES and JES+JMS-corrected jet mass responses for the shortlisted jet collections, selecting events around the W mass and in the lower $p_{\mathrm{T}}^{\mathrm{true}}$ range 1000 − 1500 GeV. 168

7.24 Fractional jet mass resolutions for the shortlisted jet collections, after JES+JMS corrections have been applied, with respect to the truth-level jet transverse momentum. Both QCD and $W-$jet topologies are considered, and in both cases the JMR is quantified around the W mass pole. 169

7.25 An illustration of how the tagging efficiency and background rejection factor are calculated from a tagger based on a generic cut (or set of cuts) on some general variable, X, using a window of size ΔX. The red line is the X-distribution for some undefined background, and the blue for some hypothetical signal. Varying the X window generates a set of background rejection and signal efficiency values. Generally X can be a single variable or a collection of variables, but here a single variable X is considered for clarity. 169

7.26 The typical form of a ROC curve, where each point is generated by considering the tagging efficiency and background rejection factor based on some variable X, with window ΔX. 170

7.27 The jet mass distributions for $W-$jet signal and QCD background topologies, for different shortlisted jet definitions. 171

7.28 The jet mass tagger ROC curves generated using $W-$jet signal and QCD background topologies, for different shortlisted jet definitions. 173

7.29 The jet D_2 distributions for $W-$jet signal and QCD background topologies, for different shortlisted jet definitions. 175

7.30 The jet 'mass + D_2' tagger ROC curves generated using $W-$jet signal and QCD background topologies, for different shortlisted jet definitions. 176

7.31 Heatmaps for QCD background rejection factors, using a $W-$jet signal, for different new jet input and grooming collections. 177

7.32 Heatmaps for QCD background rejection factors, using a top jet signal, for different new jet input and grooming collections. 179

7.33 The average $W-$jet mass pileup dependence, with respect to N_{PV}, for low and high-$p_{\mathrm{T}}^{\mathrm{true}}$ ranges. 181

7.34 The jet 'mass + D_2' tagger ROC curves generated using $W-$jet signal and QCD background topologies, for different shortlisted jet definitions, based on the very latest inputs and grooming scan. 182

7.35 The jet 'mass + τ_{32}' tagger ROC curves generated using top jet signal and QCD background topologies, for different shortlisted jet definitions, based on the very latest inputs and grooming scan. 183

8.1 The pMSSM slice and various simplified models considered in the multi-jets analysis. 189

8.2 The typical mass spectrum for the original pMSSM parameter space for which the Run 1 analysis was found to be uniquely sensitive. This spectrum is generated using SoftSUSY [137] together with the PySLHA software [138]. 191

8.3 Trigger efficiency turn-on curves for the 5- and 6-jet high-level triggers used to fire on multi-jet events. 193

8.4 A summary of two-dimensional kinematic variable distributions for $R = 1.0$ jets before and after top-tagging. A $\tilde{g}t\bar{t}$ simulated signal point is considered. The large reduction in jets can be seen after applying the substructure cuts implemented in the top-tagger. 203

8.5 The expected significance for `8ij50-0/1/2ib-0/1/2/3it` for the $\tilde{g}t\bar{t}$ model and the expected significance for the `8ij50-0ib-MJ340` region. 204

8.6 The jet multiplicity ditributions using the selection for the W + jets and $t\bar{t}$ control regions. Normalisations from the background-only fit have also been applied. 208

8.7 $E_{\mathrm{T}}^{\mathrm{miss}}/\sqrt{H_{\mathrm{T}}}$ distribution for events in the region of exactly 6 jets, each with a jet transverse momentum of 50 GeV or greater, and 0 inclusive b-tagged jets. The lower panel shows the ratio of the observed data to the total standard model background, including the statistical uncertainty in the hatched blue. This ratio is, by construction, unity. 209

8.8 A schematic illustrating the template method for the two different p_{T} thresholds in multi-jets. The region where the background shape is extracted in data is denoted by 'shape', and the region where the normalisation is extracted, for rescaling the shape, is denoted by 'norm'. 211

8.9 Example exclusive 6- and 7-jet validation regions, used to assess the closure of the QCD background estimates. 212

8.10 The non-closures in the background estimation for several different signal regions in the flavour and fat-jet streams. 219

8.11 A summary of the various systematic (and statistical) uncertainties across the 27 analysis signal regions. 222

9.1 Example $E_{\mathrm{T}}^{\mathrm{miss}}/\sqrt{H_{\mathrm{T}}}$ distributions for three (unblinded) signal regions, where no statistically significant excesses in data (see Table 9.1), with respect to the injected two-step and pMSSM signals, are reported. These two example 'benchmark' signals have parameters ($m_{\tilde{g}}$, $m_{\tilde{\chi}_1^0}$) = (1400, 200) GeV and ($m_{\tilde{g}}$, $m_{\tilde{\chi}_1^\pm}$) = (1400, 200) GeV respectvely. 226

9.2 The event yields in each of the 27 signal yields, together with the expected SM backgrounds. No statistically significant excesses are reported. 227

9.3 Expected and observed limits on the pMSSM slice model. 232

9.4 Expected and observed limits on the two-step and $\tilde{g}t\bar{t}$ RPC SUSY simplified models. 234

9.5 A summary of the 95 % CL exclusions placed on the two-step model by a variety of SUSY analyses. The observed multi-jets limit is shown by the solid red line. . 235

9.6 Expected and observed limits on the RPV stop simplified model. 236

9.7 The RPC-RPV continuum from RPC, long-lived, and prompt decaying neutralinos in ATLAS, for the two simplified models described. 237

9.8 The expected and observed limits set by a variety of SUSY searches [70, 163], including multi-jets, in the 'RPC-meets-RPV' spectrum of the $\tilde{g}q\bar{q}$ simplified model. 238

9.9 The expected and observed limits set by a variety of SUSY searches [70, 164–166], including multi-jets, in the 'RPC-meets-RPV' spectrum of the $\tilde{g}t\bar{t}$ simplified model. 239

10.1 Different jet radii for $M_{\mathrm{J}}^{\Sigma} > 500$ GeV. For example, `AntiKt12RCPtFrac0MJSigma500` refers to jets with reclustered radius $R = 1.2$, no trimming applied, and $M_{\mathrm{J}}^{\Sigma} >$ 500 GeV imposed. 242

10.2 Different M_{J}^{Σ} for $R = 1.0$ jets. For example, `AntiKt10RCPtFrac0MJSigma580` refers to jets with reclustered radius $R = 1.0$, no trimming, and $M_{\mathrm{J}}^{\Sigma} > 580$ GeV imposed. 243

10.3 Different VR reclustered jet definitions for $M_{\mathrm{J}}^{\Sigma} > 500$ GeV. For example, `AntiKtVR18RCPtFrac0MS30` refers to jets which are produced from VR reclustering with $R_{\mathrm{max}} = 1.8$, $m_X =$ 300 GeV, and $M_{\mathrm{J}}^{\Sigma} > 500$ GeV is applied. 244

10.4 A comparison of the object-based and macroscopic $E_{\mathrm{T}}^{\mathrm{miss}}$ significances for an 11-jet signal region in the 2015+2016 multi-jets analysis. 245

10.5 The template fraction of the total background prediction, for different requirements on the object-based $E_{\mathrm{T}}^{\mathrm{miss}}$ significance given in the x-axis, for both `EMTopo` and `EMPFlow` object definitions. 246

A.1 The fractional calorimeter jet mass resolution as a function of the truth jet transverse momentum. The resolution is determined from jets produced via beyond Standard Model $W' \to WZ \to q\bar{q}q\bar{q}$ and $Z' \to t\bar{t}$ signals, and the QCD jet process, using `AntiKt10LCTopoTrimmedPtFrac5SmallR20` reconstruction. The JMR is shown separately for typical ATLAS Run 2 conditions and for HL-LHC conditions with an upgraded ATLAS detector based on the ITk Step-1.9 detector geometry (red and blue respectively). The jets are reconstructed from LCW-calibrated calorimeter cell clusters, and separate Monte Carlo-based jet energy scale and jet mass scale calibrations have been derived for the Run 2 jets and HL-LHC jets, with good closure in the mass response observed after calibration. 253

B.1 The QCD dijets mass resolution for four different cases: standard trimmed jets calorimeter mass and combined mass; Soft Drop `CSSKLCTopo` jets calorimeter mass and combined mass. 255

B.2 The W/Z-jet mass resolution for different ranges in $\langle\mu\rangle$ and four different cases: standard trimmed jets calorimeter mass and combined mass; Soft Drop `CSSKLCTopo` jets calorimeter mass and combined mass. 256

B.3 The top jet mass resolution for different ranges in $\langle\mu\rangle$ and four different cases: standard trimmed jets calorimeter mass and combined mass; Soft Drop `CSSKLCTopo` jets calorimeter mass and combined mass. 257

C.1 $E_{\mathrm{T}}^{\mathrm{miss}}/\sqrt{H_{\mathrm{T}}}$ distributions in various signal regions, where each region is indicated in a subfigure. 259

C.2 $E_{\mathrm{T}}^{\mathrm{miss}}/\sqrt{H_{\mathrm{T}}}$ distributions in various signal regions, where each region is indicated in a subfigure. 260

C.3 $E_{\mathrm{T}}^{\mathrm{miss}}/\sqrt{H_{\mathrm{T}}}$ distributions in various signal regions, where each region is indicated in a subfigure. 261

C.4 $E_{\mathrm{T}}^{\mathrm{miss}}/\sqrt{H_{\mathrm{T}}}$ distributions in various signal regions, where each region is indicated in a subfigure. 262

List of Tables

2.1 The SM fermions, each possessing a spin of $+\frac{1}{2}$. The top six are the quarks; the bottom six the leptons. There are also corresponding anti-fermions of opposite electric charge. Masses are taken from the 2018 Review of Particle Physics [6]. . 6

2.2 The SM bosons, together with the forces they mediate. Masses are taken from the Review of Particle Physics [6]. 7

3.1 The particle content of the (minimal) supersymmetric Standard Model (MSSM), prior to electroweak symmetry breaking. 36

6.1 The basic jet selections, truth-matching, and isolation requirements for the jets considered in the Monte Carlo jet mass calibration. 99

6.2 Overview of the simulated samples used in these jet mass resolution studies. . . 113

6.3 The different κ-values considered in the forward-folding procedure, extracted from Figure 6.17. The top two rows are the κ values derived from W jets, and the bottom two rows are derived from top jets. The corresponding $p_{\mathrm{T}}^{\mathrm{true}}$-values are given. The equivalent x_κ-values for a Gaussian smearing, assuming a Gaussian detector response, are also given. In the final row, the $\kappa - \delta\kappa$ of 0.62 would produce a non-real $x_\kappa - \delta x_\kappa$, and so this lower bound is set to 0 (no smearing). 126

6.4 The relationship between the different JMR definitions. The nominal JMRs are presented in the first column. The second and third column represent the nominal JMRs which have been modified by additional systematic variations, coming from δx_κ in the Gaussian smearing approach and $\delta\kappa$ in the forward-folding approach. Systematic variations are not named (N. A. = Not Applicable) for the nominal MC-based JMR. 130

7.1 A different jet groomers, and grooming parameteters, considered in the initial jet inputs and grooming scan. 149

7.2 The shortlisted jet inputs and groomers for the initial scan. 163

8.1 A summary of the event cleaning cuts. 193
8.2 A summary of the basic physics objects used in the multi-jets analysis, together with their baseline and signal selections. Signal and baseline electrons, muons, and jets are used; additionally baseline photons are used in the overlap removal and $E_{\mathrm{T}}^{\mathrm{miss}}$ calculation. 195
8.3 Signal regions 199
8.4 Control regions 200
8.5 Control Region $E_{\mathrm{T}}^{\mathrm{miss}}/\sqrt{H_{\mathrm{T}}}$ selections 200
8.6 The Monte Carlo matrix element and parton shower details for the various simulated standard model backgrounds. 206
8.7 The Monte Carlo PDF and underlying event tune details for the various simulated standard model backgrounds. 206

9.1 The expected SM backgrounds, shown separately for the multi-jet and leptonic contributions, calculated from the template method, Monte Carlo simulation, and the background-only fit. The observed 2015 + 2016 data yields are also shown. 228
9.2 Summary of the results of the model-independent fit: from left to right, the 95 % CL limits on the visible cross-section, the expected number of signal events, and the number of signal events given the expected number of background events (including $\pm 1\sigma$ excursions on the expected value), and finally the discovery p-value, p_0. 231

References

[1] C. P. Burgess. "Introduction to Effective Field Theory". In: *Ann.Rev.Nucl.Part.Sci* 57 (2007), pp. 329–362. arXiv: 0701053 [hep-th].

[2] F. Englert and R. Brout. "Broken Symmetry and the Mass of Gauge Vector Mesons". In: *Phys. Rev. Lett.* 13 (1964), pp. 321–323.

[3] P. Higgs. "Broken Symmetries and the Masses of Gauge Bosons". In: *Phys. Rev. Lett.* 13 (1964), pp. 508–509.

[4] C. R. Hagen G. S. Guralnik and T. W. B. Kibble. "Global Conservation Laws and Massless Particles". In: *Phys. Rev. Lett.* 13 (1964), pp. 585–587.

[5] J. J. Thomson. "XL. Cathode Rays". In: *Philsophical Magazine Series 5* 44 (1897), p. 293.

[6] Particle Data Group. "Review of Particle Physics". In: *Phys. Rev. D* 98 (2018), p. 1898.

[7] SNO Collaboration. "Measurement of the rate of $\nu_e + d \rightarrow p + p + e^-$ interactions produced by ^8B solar neutrinos at the Sudbury Neutrino Observatory". In: *Phys. Rev. Lett.* 87 (2001), p. 6. arXiv: 0106015 [hep-ex].

[8] ATLAS Collaboration. "Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC". In: *Phys. Lett. B.* 716 (2012), pp. 1–29. arXiv: 1207.7214 [hep-ex].

[9] CMS Collaboration. "Observation of a new boson at a mass of 125 GeV with the CMS experiment at the LHC". In: *Phys. Lett. B* 716 (2012), p. 30. arXiv: 1207.7235 [hep-ex].

[10] Murray Gell-Mann. "Symmetries of Baryons and Mesons". In: *Phys. Rev.* 125 (1962), p. 18.

[11] ALICE Collaboration Foka P. "Overview of result from ALICE at the CERN LHC". In: *J. Phys. Conf. Ser.* 455 (2013), p. 15.

[12] et al. Goeckler M. "A Determination of the Lambda Parameter from Full Lattice QCD". In: *Phys. Rev. D* 73 (2005), p. 25. arXiv: 0502212 [hep-ph].

[13] C. S. et al. Wu. "Experimental Test of Parity Conservation in Beta Decay". In: *Phys. Rev.* 105 (1957), pp. 1413–1415.

[14] ALEPH Collaboration et al. "Precision Electroweak Measurements on the Z Resonance". In: *Phys. Rept.* 427 (2006), pp. 257–454. arXiv: 0509008 [hep-ex].

[15] P. A. M. Dirac. "The quantum theory of the emission and absorption of radiation". In: *Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences* 114 (1927), pp. 243–265. eprint: rspa.1927.0039.

[16] Yu. A. Golfand and E. P. Likhtman. "Extension of the Algebra of Poincare Group Generators and Violation of p Invariance". In: *JETP Lett.* 13 (1971). [Pisma Zh. Eksp. Teor. Fiz. **13** (1971) 452], p. 323.

[17] D. V. Volkov and V. P. Akulov. "Is the Neutrino a Goldstone Particle?" In: *Phys. Lett. B* 46 (1973), p. 109.

[18] J. Wess and B. Zumino. "Supergauge Transformations in Four-Dimensions". In: *Nucl. Phys. B* 70 (1974), p. 39.

[19] Abdus Salam and J. A. Strathdee. "Supersymmetry and Nonabelian Gauges". In: *Phys. Lett. B* 51 (1974), p. 353.

[20] J. Wess and B. Zumino. "Supergauge Invariant Extension of Quantum Electrodynamics". In: *Nucl. Phys. B* 78 (1974), p. 1.

[21] S. Ferrara and B. Zumino. "Supergauge Invariant Yang-Mills Theories". In: *Nucl. Phys. B* 79 (1974), p. 413.

[22] Stephen P. Martin. "A Supersymmetry primer". In: *Adv. Ser. Direct. High Energy Phys.* 18 (1998), p. 1. arXiv: `hep-ph/9709356`.

[23] Pierre Fayet. "Supersymmetry and Weak, Electromagnetic and Strong Interactions". In: *Phys. Lett. B* 64 (1976), p. 159.

[24] P. Vogel X. Qian. "Neutrino Mass Hierarchy". In: *PPNP* 83 (2015), pp. 1–30. arXiv: `hep-ex/1505.01891`.

[25] M. N. Rebelo E. Kh. Akhmedov G. C. Branco. "Seesaw mechanism and structure of neutrino mass matrix". In: *Phys. Lett. B* 478 (2000), pp. 215–223. arXiv: `hep-ph/9911364`.

[26] T. M. Undagoitia and L. Rauch. "Dark matter direct-detection experiments". In: *J. Phys. G* 43 (2016), p. 78. arXiv: `physics.ins-det/1509.08767`.

[27] W. K. Jr. Ford V. Rubin N. Thonnard. "Rotational Properties of 21 Sc Galaxies With a Large Range of Luminosities and Radii from NGC 4605 (R=4kpc) to UGC 2885 (R=122kpc)". In: *Astrophysical Journal, Part 1* 238 (1980), pp. 471–487.

[28] M. et al Markevtich. "Direct constraints on the dark matter self-interaction cross-section from the merging galaxy cluster 1E0657-56". In: *Astrophysics. J.* 606 (2004), pp. 819–824. arXiv: `astro-ph/0309303`.

[29] R. Silva W. J. C. de Silva H. S. Gimenes. "Extended Λ CDM model". In: *Astroparticle Physics* 105 (2019), pp. 37–43. arXiv: `07797 [hep-co]`.

[30] John McDonald. "Generation of WIMP Miracle-like Densities of Baryons and Dark Matter". In: (2012). arXiv: `1201.3124 [hep-ph]`.

[31] Andrei D Sakharov. "Violation of CP invariance, C asymmetry, and baryon asymmetry of the universe". In: *JETP Lett.* 5 (1967), pp. 24–27.

[32] G. 't Hooft. "Naturalness, Chiral Symmetry, and Spontaneous Chiral Symmetry Breaking". In: *Recent Developments in Gauge Theories. Ed. by G.'t Hooft et al.* Boston, MA: Springer US (1980), pp. 135–157.

[33] J. Mandula S. Coleman. "All Possible Symmetries of the S Matrix". In: *Phys. Rev.* 159 (1967), pp. 1251–1256.

[34] O. Schlotterer S. Krippendorf F. Quevedo. "Cambridge Lectures on Supersymmetry and Extra Dimensions". In: (2010). arXiv: `1011.1491 [hep-ph]`.

[35] Y. Nir K. Blum Y. Grossman and G. Perez. "Combining K^0-$\bar{K}^0$ Mixing and D^0-$\bar{D}^0$ Mixing to Constrain the Flavor Structure of New Physics". In: *Phys. Rev. Lett.* 102 (2009), pp. 1–4.

[36] Pierre Fayet. "Spontaneously Broken Supersymmetric Theories of Weak, Electromagnetic and Strong Interactions". In: *Phys. Lett. B* 69 (1977), p. 489.

[37] Riccardo Barbieri, S. Ferrara, and Carlos A. Savoy. "Gauge Models with Spontaneously Broken Local Supersymmetry". In: *Phys. Lett. B* 119 (1982), p. 343.

[38] Michael Dine and Willy Fischler. "A Phenomenological Model of Particle Physics Based on Supersymmetry". In: *Phys. Lett. B* 110 (1982), p. 227.

[39] Lisa Randall and Raman Sundrum. "Out of this world supersymmetry breaking". In: *Nucl. Phys. B* 557 (1999), p. 79. arXiv: `hep-th/9810155`.

[40] J. M. Alberto Casas et al. "What is a Natural SUSY scenario?" In: *JHEP* 06 (2015), p. 41. arXiv: hep-ph/1407.6966.

[41] T. Kuwahara B. Bajc J. Hisano and Y. Omura. "Threshold Corrections to Dimension-six Proton Decay Operators in Non-minimal SUSY $SU(5)$ GUTs". In: *Nucl. Phys. B* 910 (2016), pp. 1–22. arXiv: 1603.03568.

[42] William James Fawcett and Alan Barr. "Supersymmetry searches at the LHC and their interpretations". Presented 29 Sep 2017. Feb. 2018. URL: http://cds.cern.ch/record/2305551.

[43] Christoph Borschensky et al. "Squark and gluino production cross sections in pp collisions at $\sqrt{s} = 13$, 14, 33 and 100 TeV". In: *Eur. Phys. J.* C74.12 (2014), p. 3174. arXiv: 1407.5066 [hep-ph].

[44] D. et al. Alves. "Simplified Models for LHC New Physics Searches". In: *J. Phys. G: Nucl. Part. Phys.* 39 (2012), p. 30. arXiv: hep-ph/1105.2838.

[45] G. Aad et al. "The ATLAS Experiment at the CERN Large Hadron Collider". In: *JINST* 3 (2008), S08003.

[46] Lyndon Evans and Philip Bryant. "LHC Machine". In: *JINST* 3 (2008), S08001.

[47] S. Chatrchyan et al. "The CMS Experiment at the CERN LHC". In: *JINST* 3 (2008), S08004.

[48] A. Augusto Alves Jr. et al. "The LHCb Detector at the LHC". In: *JINST* 3 (2008), S08005.

[49] ATLAS Collaboration. *Luminosity determination in pp collisions at $\sqrt{s} = 13$ TeV using the ATLAS detector at the LHC*. Tech. rep. ATLAS-CONF-2019-021. Geneva: CERN, Jan. 2019. URL: http://cdsweb.cern.ch/record/2677054.

[50] David Hall and Christopher Hays. "" Presented 29 August, 2014. Jan. 2015. URL: http://cds.cern.ch/record/1339842.

[51] ATLAS Collaboration. "Luminosity Public Results Run 2". In: (2018). URL: https://twiki.cern.ch/twiki/bin/view/AtlasPublic/LuminosityPublicResultsRun2.

[52] Karolos Potamianos. "The upgraded Pixel detector and the commissioning of the Inner Detector tracking of the ATLAS experiment for Run-2 at the Large Hadron Collider". In: *PoS* EPS-HEP2015 (2015), p. 261. arXiv: 1608.07850 [physics.ins-det].

[53] N. Ilic. "Performance of the ATLAS Liquid Argon Calorimeter after three years of LHC operation and plans for a future upgrade". In: *JINST* 9 (2014), p. C03049.

[54] M. Tylmad. "Pulse shapes for signal reconstruction in the ATLAS Tile Calorimeter". In: *2009 16th IEEE-NPSS Real Time Conference*. May 2009, pp. 543–547.

[55] E. Diehl. "Calibration and Performance of the ATLAS Muon Spectrometer". In: *Particles and fields. Proceedings, Meeting of the Division of the American Physical Society, DPF 2011, Providence, USA, August 9-13, 2011.* (2011), p. 7. arXiv: 1109.6933 [physics.ins-det].

[56] ATLAS Collaboration. "Measurement of b-hadron pair production with the ATLAS detector in proton-proton collisions at $\sqrt{s} = 8$ TeV". In: *JHEP* 11 (2017), p. 49. arXiv: 1705.03374 [hep-ex].

[57] ATLAS Collaboration. "Measurement of $D^{\pm}$, $D^{\pm}$ and $D_s^{\pm}$ meson production cross sections in pp collisions at $\sqrt{s} = 7$ TeV with the ATLAS detector". In: *Nucl. Phys. B* 907 (2016), pp. 717–763. arXiv: 1512.02913 [hep-ex].

[58] Morad Aaboud et al. "Performance of the ATLAS Trigger System in 2015". In: *Eur. Phys. J.* C77.5 (2017), p. 317. arXiv: 1611.09661 [hep-ex].

[59] R. Achenbach et al. "The ATLAS level-1 calorimeter trigger". In: *JINST* 3 (2008), P03001.

[60] Eduard Simioni. "The Topological Processor for the future ATLAS Level-1 Trigger: from design to commissioning". In: (2014), p. 5. arXiv: 1406.4316 [hep-ex].

[61] ATLAS Collaboration. "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.* 121.081801 (2018), p. 20. arXiv: 1804.03496 [hep-ex].

[62] *Vertex Reconstruction Performance of the ATLAS Detector at* $\sqrt{s} = 13$ TeV. Tech. rep. ATL-PHYS-PUB-2015-026. Geneva: CERN, July 2015. URL: http://cds.cern.ch/record/2037717.

[63] ATLAS Collaboration. "Topological cell clustering in the ATLAS calorimeters and its performance in LHC Run 1". In: *Eur. Phys. J.C.* (2016), p. 88. arXiv: 1603.02934 [hep-ph].

[64] 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.* C76.5 (2016), p. 292. arXiv: 1603.05598 [hep-ex].

[65] *Optimisation of the ATLAS b-tagging performance for the 2016 LHC Run.* Tech. rep. ATL-PHYS-PUB-2016-012. Geneva: CERN, June 2016. URL: http://cds.cern.ch/record/2160731.

[66] ATLAS Collaboration. "Identification and energy calibration of hadronically decaying tau leptons with the ATLAS experiment in pp collisions at $\sqrt{s} = 8$ TeV". In: *Eur. Phys. J.* C75.7 (2015), p. 303. arXiv: 1412.7086 [hep-ex].

[67] *Expected performance of missing transverse momentum reconstruction for the ATLAS detector at* $\sqrt{s} = 13$ *TeV.* Tech. rep. ATL-PHYS-PUB-2015-023. Geneva: CERN, July 2015. URL: http://cds.cern.ch/record/2037700.

[68] Gavin P. Salam. "Towards Jetography". In: *Eur.Phys.J.* C67 (2010), p. 95. arXiv: 0906.1833 [hep-ph].

[69] Georges Aad et al. "Jet energy measurement with the ATLAS detector in proton-proton collisions at $\sqrt{s} = 7$ TeV". In: *Eur. Phys. J.* C73.3 (2013), p. 2304. arXiv: 1112.6426 [hep-ex].

[70] ATLAS Collaboration. "Search for new phenomena with large jet multiplicities and missing transverse momentum using large-radius jets and flavour-tagging at ATLAS in 13 TeV pp collisions". In: *JHEP* 12 (2017), p. 034. arXiv: 1708.02794 [hep-ex].

[71] ATLAS Collaboration. "Search for new phenomena in final states with large jet multiplicities and missing transverse momentum with ATLAS using $\sqrt{s} = 13$ TeV proton-proton collisions". In: *Phys. Lett.* B757 (2016), pp. 334–355. arXiv: 1602.06194 [hep-ex].

[72] Richard P. Feynman. "Very High-Energy Collisions of Hadrons". In: *Phys. Rev. Lett.* 23 (1969), p. 3.

[73] F. Close. "The quark parton model". In: *Rep. Prog. Phys.* 42 (1979), p. 51.

[74] S. Mandelstam. "Determination of the Pion-Nucleon Scattering Amplitude from Dispersion Relations and Unitarity. General Theory". In: *Phys. Rev.* 112 (1958), p. 17.

[75] P. A. M. Dirac. "The quantum theory of the electron". In: *Proc. Royal. Soc. A.* (1928), p. 15.

[76] Juan Rojo Jun Gao Lucian Harland-Lang. "The Structure of the Proton in the LHC Precision Era". In: *Phys. Reports* 742 (2018), pp. 1–121. arXiv: 1709.04922 [hep-ph].

[77] Zoltan Nagy and Davison E. Soper. "What is a parton shower?" In: *Phys. Rev. D.* 98 (2018), p. 64. arXiv: 1705.08093 [hep-ph].

[78] Georges Aad et al. "A search for $t\bar{t}$ resonances in lepton + jets events with highly boosted top quarks collected in pp collisions at $\sqrt{s} = 7$ TeV with the ATLAS detector". In: *JHEP* 041 (2012), p. 25. arXiv: 1207.2409 [hep-ex].

[79] M. A. Dobbs et al. "Les Houches Guidebook to Monte Carlo Generators for Hadron Collider Physics". In: *Les Houches* (2003), p. 49. arXiv: 0403045 [hep-ph].

[80] Gavin P. Salam. "CTEQ MCNet School Lectures on Jets". In: ().

[81] E. W. N. Glover and David A. Kosower. "Recombination Methods for Jets in $p\bar{p}$ Collisions". In: (1996), p. 11. arXiv: 9510420 [hep-ph].

[82] Georges Aad et al. "Topological cell clustering in the ATLAS calorimeters and its performance in LHC Run 1". In: *Eur. Phys. J.* C77 (2017), p. 490. arXiv: 1603.02934 [hep-ex].

[83] J. Huth et al. "Towards a Standardization of Jet Definition". In: *Proc. Research Directions for the Decade, Snowmass 1990, ed. E. L. Berger (World Scientific, Singapore)* ().

[84] Gavin P. Salam and Gregory Soyez. "A practical Seedless Infrared-Safe Cone jet algorithm". In: *JHEP* 05 (2007), p. 042. arXiv: 0704.0292 [hep-ph].

[85] Matteo Cacciari, Gavin P. Salam, and Gregory Soyez. "The Anti-k(t) jet clustering algorithm". In: *JHEP* 04 (2008), p. 063. arXiv: 0802.1189 [hep-ph].

[86] Dokshitzer, Yu.L., Leder, G. D., Moretti, S., Webber, B. R. ""Better Jet Clustering Algorithms"". In: *JHEP* (1997), p. 32. arXiv: 9707323 [hep-ph].

[87] S. D. Ellis and Davision E. Soper. "Successive Combination Jet Algorithm For Hadron Collisions". In: *Phys.Rev.D* 48 (1993), p. 18. arXiv: 9305266 [hep-ph].

[88] Matteo Cacciari, Gavin P. Salam, and Gregory Soyez. "FastJet user manual". In: (2011), p. 69. arXiv: 1111.6097 [hep-ph].

[89] L-T. Wang D. Krohn J. Thaler. "Jet Trimming". In: *JHEP* (2010), p. 20. arXiv: 0912.1342 [hep-ph].

[90] 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* 076.09 (2013), p. 72. arXiv: 1306.4945 [hep-ex].

[91] M. Aaboud et al. "Jet energy scale measurements and their systematic uncertainties in proton-proton collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector". In: *Phys. Rev.* D96.7 (2017), p. 072002. arXiv: 1703.09665 [hep-ex].

[92] Matteo Cacciari, Gavin P. Salam, and Gregory Soyez. "The Catchment Area of Jets". In: *JHEP* 0804 (2008), p. 42. arXiv: 0802.1188 [hep-ph].

[93] Matteo Cacciari and Gavin P. Salam. "Pileup subtraction using jet areas". In: *Phys.Lett.B* 659 (2008), p. 12. arXiv: 0707.1378 [hep-ph].

[94] Torbjorn Sjöstrand, Stephen Mrenna, and Peter Z. Skands. "A Brief Introduction to PYTHIA 8.1". In: *Comput. Phys. Commun.* 178 (2008), pp. 852–867. arXiv: 0710.3820 [hep-ph].

[95] T. Gleisberg et al. "Event generation with SHERPA 1.1". In: *JHEP* 02 (2009), p. 007. arXiv: 0811.4622 [hep-ph].

[96] B. Nachman et al. "Jets from Jets: Re-clustering as a tool for large radius jet reconstruction and grooming at the LHC". In: *JHEP* (2015), p. 20. arXiv: 1407.2922 [hep-ph].

[97] M. Bahr et al. "Herwig++ Physics and Manual". In: *Eur. Phys. J.* C58 (2008), pp. 639–707. arXiv: 0803.0883 [hep-ph].

[98] ATLAS Collaboration. *Jet mass reconstruction with the ATLAS Detector in early Run 2.* Tech. rep. ATLAS-CONF-2016-035. Geneva: CERN, July 2016. URL: https://cds.cern.ch/record/2200211.

[99] ATLAS Collaboration. *In-situ measurements of large-radius jet reconstruction performance.* Tech. rep. ATLAS-CONF-2017-063. Geneva: CERN, July 2017. URL: https://cds.cern.ch/record/2275655.

[100] ATLAS Collaboration. *ATLAS Combined Mass resolution and W mass reconstruction performance*. Tech. rep. JETM-2017-002. Geneva: CERN, Feb. 2017. URL: https://atlas.web.cern.ch/Atlas/GROUPS/PHYSICS/PLOTS/JETM-2017-002/.

[101] ATLAS Collaboration. *Jet mass resolution in ATLAS using Run 2 Monte Carlo simulation*. Tech. rep. ATL-PUB-2018-015. Geneva: CERN, July 2018. URL: https://cds.cern.ch/record/2631339.

[102] ATLAS Collaboration. *Measurement of large radius jet mass reconstruction performance at* $\sqrt{s}$ *= 8 TeV using the ATLAS detector*. Tech. rep. ATLAS-CONF-2016-008. Geneva: CERN, Mar. 2016. URL: http://cds.cern.ch/record/2139642.

[103] Jerome Friedman Trevor Hastie Robert Tibshirani. "The Elements of Statistical Learning". In: *Springer* (2001), p. 737.

[104] ATLAS Collaboration. "In situ calibration of large-R jet energy and mass in 13 TeV proton-proton collisions with the ATLAS detector". In: *EPJC* (2018), p. 069. arXiv: 1708.02794 [hep-ex].

[105] ATLAS Collaboration. "Jet reconstruction and performance using particle flow with the ATLAS Detector". In: *Eur. Phys. J.C.* (2017), p. 67. arXiv: 1703.10485 [hep-ph].

[106] D. Adams et al. "Towards an Understanding of the Correlations in Jet Substructure". In: *Eur. Phys. J.* C75.9 (2015), p. 409. arXiv: 1504.00679 [hep-ph].

[107] ATLAS Collaboration. "Jet mass and substructure of inclusive jets in $\sqrt{s} = 7$ TeV pp collisions with the ATLAS experiment". In: *JHEP* (2012), p. 54. arXiv: 1203.4606 [hep-ph].

[108] M. Perelstein. "Little Higgs Models and Their Phenomenology". In: *Prog.Part.Nucl.Phys.* 58 (2007), pp. 247–291. arXiv: 0512128 [hep-ph].

[109] P. Langacker. "The Physics of Heavy Z' Gauge Bosons". In: *Rev.Mod.Phys.* 81 (2009), pp. 1199–1228. arXiv: 0801.1345 [hep-ph].

[110] Sundrum. R. Randall L. "A Large Mass Hierarchy from a Small Extra Dimension". In: *Phys. Rev. Lett.* 83 (1999), pp. 3370–3373. arXiv: 9905221 [hep-th].

[111] ATLAS Collaboration. "The ATLAS Simulation Infrastructure". In: *Eur. Phys. J.* C70 (2010), pp. 823–874. arXiv: 1005.4568 [physics.ins-det].

[112] S. Agostinelli et al. "GEANT4: A Simulation toolkit". In: *Nucl. Instrum. Meth.* A506 (2003), pp. 250–303.

[113] Torbjorn Sjostrand, Stephen Mrenna, and Peter Z. Skands. "A Brief Introduction to PYTHIA 8.1". In: *Comput. Phys. Commun.* 178 (2008), pp. 852–867. arXiv: 0710.3820 [hep-ph].

[114] Richard D. Ball et al. "Parton distributions for the LHC Run II". In: *JHEP* 1504 (2015), p. 040. arXiv: 1410.8849 [hep-ph].

[115] *ATLAS Run 1 Pythia8 tunes*. Tech. rep. ATL-PHYS-PUB-2014-021. Geneva: CERN, Nov. 2014. URL: http://cds.cern.ch/record/1966419.

[116] 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. arXiv: 1405.0301 [hep-ph].

[117] ATLAS Collaboration. "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. arXiv: 1306.4945 [hep-ex].

[118] Mrinal Dasgupta et al. "Towards an understanding of jet substructure". In: *JHEP* 09 (2013), p. 029. arXiv: 1307.0007 [hep-ph].

[119] Andrew J. Larkoski et al. "Soft Drop". In: *JHEP* 05 (2014), p. 146. arXiv: 1402.2657 [hep-ph].

[120] F. Dreyer et al. "Recursive Soft Drop". In: *J. High Energ. Phys.* 06 (2018), p. 093. eprint: 1804.03657 (hep-ph).

[121] Peter Berta et al. "Particle-level pileup subtraction for jets and jet shapes". In: *JHEP* 06 (2014), p. 092. arXiv: 1403.3108 [hep-ex].

[122] Matteo Cacciari, Gavin P. Salam, and Gregory Soyez. "SoftKiller, a particle-level pileup removal method". In: *Eur. Phys. J.* C75.2 (2015), p. 59. arXiv: 1407.0408 [hep-ph].

[123] ATLAS Collaboration. "Impact of Alternative Inputs and Grooming Methods on Large-R Jet Reconstruction in ATLAS". In: (2017). URL: https://cds.cern.ch/record/2297485.

[124] ATLAS Collaboration. *Search for diboson resonances in hadronic final states in 79.8 $fb^1 of pp collisions at \sqrt{s}$ = 13 TeV with the ATLAS detector*. Tech. rep. ATLAS-CONF-2018-016. Geneva: CERN, May 2018. URL: http://cdsweb.cern.ch/record/2621302.

[125] ATLAS Collaboration. *Constituent-level pile-up mitigation techniques in ATLAS*. Tech. rep. ATLAS-CONF-2017-065. Geneva: CERN, July 2017. URL: https://cds.cern.ch/record/2281055.

[126] Stephen D. Ellis, Christopher K. Vermilion, and Jonathan R. Walsh. "Recombination Algorithms and Jet Substructure: Pruning as a Tool for Heavy Particle Searches". In: *Phys. Rev.* D81 (2010), p. 094023. arXiv: 0912.0033 [hep-ph].

[127] J. Thaler C. Frye A. Larkoski and K. Zhou. "Casimir Meets Poisson: Improved Quark/Gluon Discrimination with Counting Observables". In: *JHEP* 09 (2017), p. 48. arXiv: 1704.06266 [hep-ph].

[128] Jonathan M. Butterworth et al. "Jet substructure as a new Higgs search channel at the LHC". In: *Phys. Rev. Lett.* 100 (2008), p. 242001. arXiv: 0802.2470 [hep-ph].

[129] Jesse Thaler and Ken Van Tilburg. "Identifying Boosted Objects with N-subjettiness". In: *JHEP* 03 (2011), p. 015. arXiv: 1011.2268 [hep-ph].

[130] Andrew J. Larkoski, Gavin P. Salam, and Jesse Thaler. "Energy Correlation Functions for Jet Substructure". In: *JHEP* 06 (2013), p. 108. arXiv: 1305.0007 [hep-ph].

[131] ATLAS Collaboration. "Comparative performance of large-R jet reconstruction with constituent-level pile-up mitigation and various grooming algorithms". In: (2018). URL: https://atlas.web.cern.ch/Atlas/GROUPS/PHYSICS/PLOTS/JETM-2018-003/.

[132] ATLAS Collaboration. "Performance of top-quark and W-boson tagging with ATLAS in Run 2 of the LHC". In: (2018). arXiv: 1808.07858 [hep-ex].

[133] T. Fawcett. "An introduction to ROC analysis". In: *Pattern Recognition Letters* 27 (2006), pp. 861–874.

[134] ATLAS Collaboration. "Search for new phenomena in final states with large jet multiplicities and missing transverse momentum at $\sqrt{s}$ = 8.0 TeV proton-proton collisions using the ATLAS experiment". In: *JHEP* 130 (2013), p. 49. arXiv: 1308.1841 [hep-ex].

[135] M. Baak et al. "HistFitter software framework for statistical data analysis". In: *Eur. Phys. J.* C75 (2015), p. 153. arXiv: 1410.1280 [hep-ex].

[136] ATLAS Collaboration. "Summary of the ATLAS experiment's sensitivity to supersymmetry after LHC Run 1 — interpreted in the phenomenological MSSM". In: *JHEP* 10 (2015), p. 134. arXiv: 1508.06608 [hep-ex].

[137] B. Allenach. "SOFTSUSY: a program for calculating supersymmetric spectra". In: *Comput. Phys. Commun.* 143 (2002), pp. 305–331. arXiv: 0104145 [hep-ph].

[138] Andy Buckley. "PySLHA: a Pythonic interface to SUSY Les Houches Accord data". In: (). arXiv: 1305.4194 [hep-ph].

[139] A Corso-Radu et al. "Data Quality Monitoring Framework for the ATLAS experiment: Performance achieved with colliding beams at the LHC". In: *J. Phys. Conf. Ser.* 331 (2011), p. 6.

[140] *Data-Quality Requirements and Event Cleaning for Jets and Missing Transverse Energy Reconstruction with the ATLAS Detector in Proton-Proton Collisions at a Center-of-Mass Energy of $\sqrt{s}$ = 7 TeV*. Tech. rep. ATLAS-CONF-2010-038. Geneva: CERN, July 2010. URL: https://cds.cern.ch/record/1277678.

[141] *Tagging and suppression of pileup jets with the ATLAS detector*. Tech. rep. ATLAS-CONF-2014-018. Geneva: CERN, May 2014. URL: https://cds.cern.ch/record/1700870.

[142] *Electron efficiency measurements with the ATLAS detector using the 2015 LHC proton-proton collision data*. Tech. rep. ATLAS-CONF-2016-024. Geneva: CERN, June 2016. URL: https://cds.cern.ch/record/2157687.

[143] Morad Aaboud et al. "Measurement of the photon identification efficiencies with the ATLAS detector using LHC Run-1 data". In: *Eur. Phys. J.* C76.12 (2016), p. 666. arXiv: 1606.01813 [hep-ex].

[144] Jean-Baptiste De Vivie De Regie et al. *ATLAS electron, photon and muon isolation in Run 2*. Tech. rep. ATL-COM-PHYS-2017-290. This note contains the Moriond 2017 recommendations. It will be updated when new recommendations become available. Geneva: CERN, Mar. 2017. URL: https://cds.cern.ch/record/2256658.

[145] Sonia El Hedri et al. "Learning How to Count: A High Multiplicity Search for the LHC". In: *JHEP* 136 (2013), p. 36. arXiv: 1302.1870 [hep-ph].

[146] ATLAS Collaboration. "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: (2017). URL: http://cdsweb.cern.ch/record/2259646.

[147] Peter Zeiler Skands. "Tuning Monte Carlo Generators: The Perugia Tunes". In: *Phys. Rev. D* 82 (2010), p. 074018. arXiv: 1005.3457 [hep-ph].

[148] (for the ATLAS Collaboration) Lukas W. "Fast Simulation for ATLAS: Atlfast-II and ISF". In: *J. Phys.: Conf. Ser.* 396 (2012), p. 9.

[149] 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. arXiv: 1002.2581 [hep-ph].

[150] Torbjorn Sjostrand, Stephen Mrenna, and Peter Z. Skands. "PYTHIA 6.4 Physics and Manual". In: *JHEP* 05 (2006), p. 026. arXiv: hep-ph/0603175 [hep-ph].

[151] ATLAS Collaboration. *Jet Uncertainties Twiki*. https://twiki.cern.ch/twiki/bin/view/AtlasProtected/JetUncertainties2015ICHEP2016.

[152] ATLAS Collaboration. *Flavour Tagging Twiki*. https://twiki.cern.ch/twiki/bin/view/AtlasProtected/BTagCalib2015.

[153] M. Schönherr. "Parton shower matching and merging". In: *Proceedings of Corfu Summer Institute 2017 'School and Workshops on Elementary Particle Physics and Gravity* (2018), p. 10. arXiv: 1804.10643 [hep-ph].

[154] Hung-Liang Lai et al. "New parton distributions for collider physics". In: *Phys. Rev. D* 82 (2010), p. 074024. arXiv: 1007.2241 [hep-ph].

[155] *HistFactory: A tool for creating statistical models for use with RooFit and RooStats*. Tech. rep. CERN-OPEN-2012-016. Geneva: CERN, 2012. URL: https://cds.cern.ch/record/1456844?/.

[156] ATLAS Collaboration. "Improved luminosity determination in pp collisions at $\sqrt{s} = 7$ TeV using the ATLAS detector at the LHC". In: *Eur. Phys. J. C.* 73 (2013), p. 40. arXiv: 1302.4393 [hep-ex].

[157] Eilam Gross. "LHC Statistics for Pedestrians". In: (2008). URL: https://cds.cern.ch/record/1099994.

[158] Glen Cowan et al. "Asymptotic formulae for likelihood-based tests of new physics". In: *Eur. Phys. J.* C71 (2011). [Erratum: Eur. Phys. J.C73,2501(2013)], p. 1554. arXiv: 1007.1727 [physics.data-an].

[159] A L Read. "Presentation of search results: the CL s technique". In: *Journal of Physics G: Nuclear and Particle Physics* 28.10 (2002), p. 2693. URL: http://stacks.iop.org/0954-3899/28/i=10/a=313.

[160] ATLAS Collaboration. "Search for top-squark pair production in final states with one lepton, jets, and missing transverse momentum using 36 fb^1 of $\sqrt{s} = 13$ TeV pp collision data with the ATLAS detector". In: *JHEP* 108 (2018), p. 79. arXiv: 1708.02794 [hep-ex].

[161] ATLAS Collaboration. *Reinterpretation of searches for supersymmetry in models with variable R-parity-violating coupling strength and long-lived R-hadrons*. Tech. rep. ATLAS-CONF-2018-003. Geneva: CERN, Mar. 2018. URL: http://cdsweb.cern.ch/record/2308391.

[162] Dreiner, H., and Ross, G. G. "R-Parity Violation at Hadron Colliders". In: *Nucl. Phys. B* 367 (1991), p. 597.

[163] ATLAS Collaboration. "Search for squarks and gluinos in final states with jets and missing transverse momentum using 36 fb^1 of $\sqrt{s} = 13$ TeV pp collision data with the ATLAS detector". In: *Phys. Rev. D.* 97 (2018), p. 48. arXiv: 1712.02332 [hep-ex].

[164] ATLAS Collaboration. "Search for new phenomena in a lepton plus high jet multiplicity final state with the ATLAS experiment using $\sqrt{s} = 13$ Tev proton-proton collision data". In: *JHEP* 88 (2017), p. 51. arXiv: 1704.08493 [hep-ex].

[165] ATLAS Collaboration. "Search for supersymmetry in final states with two same-sign or three leptons and jets using 36 fb^1 of $\sqrt{s} = 13$ TeV pp collision data with the ATLAS detector". In: *JHEP* 084 (2017), p. 43. arXiv: 1706.03731 [hep-ex].

[166] ATLAS Collaboration. "Search for supersymmetry in final states with missing transverse momentum and multiple b-jets in proton-proton collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector". In: *JHEP* 107 (2018), p. 55. arXiv: 1711.01901 [hep-ex].

[167] L-T. Wang D. Krohn J. Thaler. "Jets with Variable R". In: *JHEP* (2009), p. 20. arXiv: 0903.0392 [hep-ph].

[168] ATLAS Collaboration. *Object-based missing transverse momentum significance in the ATLAS detector*. Tech. rep. ATLAS-CONF-2018-038. Geneva: CERN, July 2018. URL: https://cds.cern.ch/record/2630948.

[169] Apollinari, G. (ed.) *et al*. "High-Luminosity Large Hadron Collider (HL-LHC): Technical Design Report V.0.1". In: *CERN Yellow Reports: Monographs* (2017), p. 599. eprint: 1306.4945.

[170] ATLAS Collaboration. "Technical Design Report for the ATLAS Inner Tracker Strip Detector". In: CERN-LHCC-2017-005 (Apr. 2017). URL: https://cds.cern.ch/record/2257755.