



University of Glasgow | Experimental
Particle Physics

Physics at ATLAS From Top to Bottom: Jet Substructure in Boosted $t\bar{t}$ Events at 13 TeV and Development of Pixel Detector Modules for the Inner Tracker Upgrade

Adam Graham Rennie

School of Physics and Astronomy
College of Science and Engineering
University of Glasgow

Submitted in fulfilment of the requirements for the degree of Doctor of Philosophy

March 2023



Abstract

Operating at the energy frontier, between 2015 and 2018, CERN's Large Hadron Collider (LHC) collided protons at an unprecedented centre-of-mass energy of $\sqrt{s} = 13$ TeV, allowing for the production of massive particles in great abundance. Over this period, the ATLAS Experiment collected 139 fb^{-1} of data suitable for physics analysis; the largest pp dataset to date. Using this dataset, this thesis presents a measurement of the internal structure of jets arising from the production of the most massive of the Standard Model particles, the top quark, in $t\bar{t}$ pairs, in the case that these jets are produced with high momentum. These results are presented as differential cross-section measurements, unfolded to remove detector effects, and compared to state-of-the-art Monte Carlo simulations. It is found that some observables are modelled poorly by current predictions, and that these substructure observables are sensitive to the choice of parton shower modelling and modelling of the final state radiation. Also presented are testbeam studies which investigate the viability of various silicon pixel sensor designs to be used in the upcoming upgrade of the ATLAS tracking system. Here, it is found that all proposed designs meet the required specification for device efficiency after irradiation and are therefore suitable for use in the High-Luminosity LHC.

Declaration

The work presented in this thesis is my own, carried out as a member of the Experimental Particle Physics research group within the School of Physics and Astronomy, part of the University of Glasgow's College of Science and Engineering. It has not been submitted anywhere else for any other qualification.

This work was carried out within the ATLAS Collaboration, one of the experiments which uses CERN's Large Hadron Collider. This is a collaboration of thousands of people and as such, this work is built upon, and makes use of, tools and experience developed over decades within this collaboration. I contributed directly to the collaboration by participation in the ITk pixel planar sensor market survey. This involved data-taking at testbeam and contributions to the reconstruction and analysis of that data. I was also chosen to represent the collaboration in presenting "Highlights of the ATLAS top quark precision measurements" at the [XI International Conference on New Frontiers in Physics](#) and "ATLAS Virtual Visits: Where Worlds Collide" at [Lepton-Photon 2021](#).

Part I of this thesis provides an introduction to the field and the ATLAS Experiment and as such does not describe my own work.

Part II focuses on research carried out for a publication on jet substructure observables [1], for which I was the lead analyser, currently being submitted for publication in *Physical Review D*. The studies presented there are my own work with the exception of the unfolding procedure applied to produce the final results shown in Chapter 8, the implementation and validation of which is described in Section 6.2, and the χ^2 test described in Chapter 8. The technique described in Section 5.2, used to reduce the single top background, was subsequently applied in Ref. [2].

Part [III](#) covers research carried out into pixel detector modules for the upcoming upgrade to the ATLAS Detector. Chapter [9](#) provides an introduction to this upgrade and so is not my own work. Chapter [10](#) then describes the research itself. Each stage is my own work, from data-taking to reconstruction and analysis. The exception to this is the data taken during the October 2018 testbeam, for which I was not present. This work is the subject of an upcoming publication which is currently in internal review within the ATLAS ITk community.

Adam Rennie

Acknowledgements

The work presented in this thesis would not have been possible without the support of a great many people around me, from colleagues to friends and family. It is not possible to name each and every one here, but there are some who deserve to be recognised.

First and foremost, I must thank my supervisor, Craig Buttar. Sitting in (some of) his quantum mechanics lectures as an undergraduate, I could never have envisioned what was to come. His support and guidance over the last few years has been invaluable, and he has always been able to bring some insight to the problem at hand and keep an eye on the bigger picture. Add to this his genuinely caring approach to supervision and I am very grateful to have been his student. I'd also like to thank my second supervisor, Tony Doyle. Although not directly involved in my research, Tony has always been there to help when called upon and has provided many an interesting conversation over the course of the PhD.

To Federica Fabbri (AKA F. Federica), I owe an immeasurable debt for her endless patience and understanding over the last few years. Our exchange of physics know-how for English idioms has been one of the best deals I've ever struck. I simply cannot imagine a better person to have shown me the ropes. Over the course of my PhD, Federica has served in equal parts as patient teacher, role model, and friend. I wouldn't have been able to reach this point without her.

To my Canadian collaborators, Chris Garner and Pekka Sinervo, I must extend my thanks for several years of fruitful work and discussion. Dwayne Spiteri should also receive a shout-out for his assistance with the analysis after joining in its latter stages.

I have been immensely fortunate to have been surrounded by the other academics within Glasgow's Experimental Particle Physics research group for the duration of my PhD. For all of the insightful discussions and pointers along the way, I must thank Jay Howarth and Mark Owen. In this regard, Andy Buckley must receive a special

mention not only for discussions had but for *Practical Collider Physics* [3], which I only wish he'd published sooner. For unwavering support with all things outreach, I extend my gratitude to Aidan Robson. To the latecomers, Lucia Grillo and Mark Whitehead, I am grateful for their support and encouragement in the latter stages of this process with facing up to the post-PhD world.

My PhD experience would not have been the same were it not for the camaraderie with my immediate cohort at Glasgow — Bruno Borbely, Charlie Ruggles, David Friday, Laurynas Mincė, and Puwen Sun. People are often brought closer by shared trauma and this was certainly the case for the Glasgow Writers' Club; Bruno, David, and Laurynas. It was a privilege to have shared this experience with them, in particular.

Next up are the "GUMties" — Andy Blue, Dima Maneuski, Kenny Wraight, Liam Cunningham, Richard Bates, and Sneha Naik. As it pertains to my PhD, I must extend the deepest of admonishments to Dima and Kenny for involving me in the masochistic ritual they call "testbeam". Richard must also assume some responsibility for my decision to pursue a PhD in the first place.

On that topic, I must thank Emily Pender, Matt Sullivan, Ricardo González López, Tobias Fitschen, and Yanyan Gao for years of testbeam-related collaboration from data-taking at DESY to reconstruction, analysis, and paper-writing.

To the rest of the Glasgow group I am grateful for each of their unique contributions over the last few years; Dana Seman Bobulska, Elliot Watton, Ethan Simpson, Giuseppe Callea, Gordon Stewart, Harriet Watson, Jill Borland, Johnny Friend, John Nugent, Jonathan Jamieson, Jonathan Thompson, Mike Fenton, Murdo Traill, Neil Warrack, Paul Morrison, Sarah Boutle, and Will Madden.

To all the people, across several generations of the LTA, who made my two years in Geneva — COVID and all — as enjoyable as they were, I will always be thankful. These are Aaron O'Neill, Alex Götz, Conor Logie, Daniela Köck, David Monk, Dougie Smith, Eleanor Jones, Harry Cooke, James Iddon, Jordan Walker, Júlia Jakyess, Lakshan Ram, Luke Dyks, Martha Hilton, Maty Vozák, Melissa Yexley, Nikos Styl, Ryan Wood, Sara Mitchell, Shantam Taneja, Tobias Bisanz, and Victor Guillen.

Steph Hills has earned my gratitude for her organising of and assistance with a litany of public engagement activities, as well as her unparalleled cheerleading.

As co-founders of "Beans on Roast", Neil Moffat, Lluís Simon Argemi, and Lojius Lombigit will always have my appreciation for the establishment of a pillar of Glasgow

PPE which has now outstayed all of us, and for the many meals and social events enjoyed together.

I should thank both of my parents, who enabled me to pursue a physics degree and subsequently this PhD and without whom I would not have reached this point.

To a pair of my oldest friends, Jack and Kirsty, I am grateful as ever for their unwavering support during the many crises which unfolded over the course of this period of my life.

Finally, I must thank Savannah. Much of this PhD has been an experience shared with her and I am thankful to have had her by my side for the journey, and now as we each embark on our next adventure.

Contents

0. Introduction	1
I. Top Quark Physics and the ATLAS Experiment at the Large Hadron Collider	4
1. Theoretical Foundations	5
1.1. The Standard Model of particle physics	5
1.2. Top quark physics	11
1.3. Event simulation	16
2. The Large Hadron Collider and the ATLAS Experiment	21
2.1. The Large Hadron Collider	21
2.2. The ATLAS Experiment	23
2.2.1. Inner detector	26
2.2.2. Calorimeter system	31
2.2.3. Muon Spectrometer (MS)	34
2.2.4. Magnet system	37
2.2.5. Trigger system	37
2.2.6. Luminosity detectors	39
II. Measurement of Jet Substructure Observables in Boosted $t\bar{t}$ Events	42
3. Analysis Overview and Tools	43
3.1. Jets	43
3.1.1. Clustering	44
3.1.2. Particle Flow	52

3.1.3. Reclustering	56
3.1.4. Calibration	59
3.1.5. Substructure	62
3.1.6. Tagging	70
3.2. Analysis overview	73
4. Event Reconstruction	75
4.1. Data and Monte Carlo samples	75
4.1.1. Data samples	75
4.1.2. Monte Carlo samples	76
4.2. Detector-level objects	78
4.3. Particle-level objects	84
4.4. Study on different jet collections	85
5. Event Selection	93
5.1. Selection criteria	94
5.2. $m_{\ell b}$ cut	95
5.3. Mass window cut	103
5.4. Selection of hadronic top candidates	104
5.5. Background estimation	106
5.5.1. Monte Carlo backgrounds	106
5.5.2. Fake lepton estimate	107
5.6. Data/MC comparison	113
6. Substructure Observables and Unfolding	123
6.1. Selection of observables	123
6.2. Unfolding	134
6.2.1. Iterative Bayesian Unfolding	136
6.2.2. Optimisation	142
6.2.3. Validation	147
7. Uncertainties	154
7.1. Detector-related uncertainties	154
7.2. Modelling uncertainties	156
7.2.1. $t\bar{t}$ modelling	156
7.2.2. Background modelling	157
7.3. Summary of uncertainties	158

8. Results and Conclusions	165
 III. Development of Pixel Detector Modules for the ATLAS Inner Tracker	 177
9. The High-Luminosity LHC and the ATLAS Inner Tracker	178
9.1. Upgrade to High-Luminosity LHC	178
9.1.1. Motivation and physics prospects	179
9.1.2. Accelerator upgrades	181
9.2. The ATLAS Inner Tracker	184
9.3. The ITk pixel detector	188
 10. Analysis of Testbeam Data for ATLAS ITk Pixel Modules	 195
10.1. Testbeam setup	195
10.2. Devices under test (DUTs)	199
10.3. Testbeam campaigns	201
10.4. Reconstruction and analysis	201
10.4.1. EUTelescope	203
10.4.2. tbmon2	209
10.5. Results	211
10.5.1. October 2018	212
10.5.2. December 2018	214
10.6. Conclusions	219
 11. Conclusion	 221
 A. ATLAS Virtual Visits	 225
 Bibliography	 233
 List of figures	 256
 List of tables	 264

For my grandfather, Richard Graham Rennie, who never missed an opportunity to tell the story of the friend who failed his PhD.

“... we have to remember that what we observe is not nature in itself but nature exposed to our method of questioning.”

— Werner Heisenberg, 1901–1976

Chapter 0.

Introduction

“To be accepted as a paradigm, a theory must seem better than its competitors, but it need not, and in fact never does, explain all the facts with which it can be confronted.”

— Thomas S. Kuhn

My first awareness of CERN and the Large Hadron Collider (LHC) came as a secondary school pupil at the hands of the Daily Mail, who posed the question, “Are we all going to die next Wednesday?” [4,5]. Alas, the world went on and four years later came the discovery of the Higgs boson [6,7]. Mission accomplished for the LHC? Well, yes and no.

The final piece in the Standard Model (SM), the discovery of the Higgs boson represents the culmination of a century’s worth of research into the fundamental nature of the universe and the construction of the largest experiment ever conducted. The LHC was conceived with the explicit aim of conclusively demonstrating the existence or otherwise of the Higgs. After decades assembling and confirming the SM, the field of high-energy physics (HEP) now finds itself at a unique juncture. The SM is not a complete description of nature and cannot be the final theory. Despite this, it continues to be tremendously successful at describing observations made at experiments such as those at the LHC. For the first time since the advent of modern collider physics, there are no obvious candidates for the next piece in the puzzle.

Upon the discovery of the Higgs boson at the LHC, many in the field thought that the discovery of supersymmetry (SUSY) would follow soon thereafter. Despite over a

decade's worth of searching, this remains elusive and the SM remains frustratingly successful. So what now for the LHC and HEP at large?

The existence of the Higgs boson, and other particles discovered before it, is a prediction of the SM. Beyond the Standard Model (BSM) theories, such as SUSY, may make predictions which can be verified or falsified by experiments at the LHC, such as the ATLAS Experiment. Indeed, a part of ATLAS' physics program is in performing direct searches for such predictions. However, the search for new physics also encompasses methods for anomaly detection. Experiments such as ATLAS play a vital role in making precision measurements of SM predictions, both to further constrain current understanding and in the hope of finding instances of disagreement between observation and prediction.

Whether performing direct searches or making precision measurements and conducting anomaly detection, at least one thing is clear; wherever the new physics lies, it is well hidden. As such, sensitivity to this new physics will require the development of ever more powerful tools and analysis techniques. The LHC has now begun a third run of data-taking, but with no significant increase in energy as compared to Run 2, and given a lack of evidence for new physics in Runs 1 and 2, it seems unlikely that this will yield a major breakthrough. At the end of Run 3, however, a major upgrade to the accelerator will produce a significant increase in the luminosity delivered to the LHC experiments. Making the most of these new datasets will necessitate improvements to current modelling of many physics processes and the development of new analysis techniques, as well as the evolution of existing ones.

This thesis documents contributions to the development of some of the tools which will be vital in the search for the SM's successor in the decades to come. These are in the form of an analysis of jet substructure in highly boosted top quark pairs ($t\bar{t}$), and the development of pixel detector modules for the upgrade of ATLAS' tracking system for the high-luminosity phase of the LHC (HL-LHC).

As the most massive particle in the SM, the top quark sits in a unique position. Amongst many other interesting properties, its abundant production at the LHC makes its pair production a major background for many important processes and BSM searches, and it is expected to interact strongly with many postulated BSM particles. Going forward, it will therefore be important to be able to model this process well and to discriminate between it and others. Modelling is challenging as a result of the nature of the strong force. Quantum chromodynamics (QCD) dictates that quarks

such as the top are not observed in isolation by experiments such as ATLAS, but instead manifested as collimated sprays of hadrons called *jets*. Measurements of the *substructure* of these jets serve as a gateway for improved top-jet modelling and improved tools for discriminating between jets originating from top quarks and those arising from other processes.

To be able to make such measurements in the HL-LHC era, the ATLAS tracking sub-system must be replaced and upgraded in its entirety. The new system will be the Inner Tracker (ITk), made up of thousands of silicon tracking modules. The innermost and most precise of these are pixel detector modules. Also presented here are a set of measurements made at testbeam which aim to investigate the suitability of different sensor designs using prototypes of the devices which will be used in the ITk, after they have been subjected to the radiation damage they will experience over the lifetime of the HL-LHC.

The thesis is split into three parts. Part I begins with an introduction to the SM and the top quark in Chapter 1, followed by a description of the LHC and the ATLAS Detector in Chapter 2.

Part II describes the analysis of jet substructure observables in boosted $t\bar{t}$ events which makes up the bulk of this thesis. Chapter 3 provides a description of several analysis techniques which are important to this measurement, focused on the treatment of jets and jet substructure observables, as well as an overview of the strategy and structure of the analysis. Chapter 4 details the reconstruction of analysis objects from interaction with the detector, as well as studies performed in choosing which objects are best suited to the analysis. Chapter 5 deals with the event selection procedure applied to isolate the $t\bar{t}$ process, as well as the estimation of background processes. Chapter 6 details the selection procedure used to determine the substructure observables measured in the analysis, followed a description of the unfolding method used to remove the effects of interaction with the detector. Chapter 7 describes the various sources of uncertainty which are taken into account by the analysis. The results themselves are presented as unfolded distributions in Chapter 8.

Part III describes the work on the development of pixel detector modules for the ITk. Chapter 9 provides an introduction to the HL-LHC as well as descriptions of the ITk and the ITk pixel system. The testbeam analysis itself is then presented in its entirety in Chapter 10.

Finally, Chapter 11 presents a summary of the conclusions drawn from this thesis.

Part I.

Top Quark Physics and the ATLAS Experiment at the Large Hadron Collider

Chapter 1.

Theoretical Foundations

“If you wish to make an apple pie from scratch, you must first invent the universe.”

— Carl Sagan

1.1. The Standard Model of particle physics

The Standard Model of Particle Physics (SM) is a quantum field theory (QFT) which describes the elementary particles which are the universe’s building blocks, as well as the fundamental forces which govern their interactions. As a description of nature, it is the most successful physical theory humankind has produced thus far. In the SM picture, the universe is made up of a set of forces which act on and through a collection of particles, which are excitations of their associated quantum fields.

The forces are the electromagnetic (EM), and the weak and strong forces. These forces are each mediated by one of the gauge bosons, which are one of two classes of particles in the SM. Bosons are particles with integer spin, whilst the fermions have half-integer spin. The fermions can be divided again, into quarks and the leptons. Additionally, the fermions exist in three generations, with each successive generation essentially a more massive version of the previous.

The most familiar of the quarks are the first generation, the *up* and *down* quarks. Almost all ordinary matter which is familiar on an everyday, macroscopic scale, consists of these first generation quarks in bound states called hadrons. Sub-classes of

hadrons are mesons and baryons, which are bound states of two and three quarks, respectively. Protons and neutrons are baryons with composition uud and udd . Along with protons and neutrons, the electron, a first generation lepton, provides the other ingredient to form atoms.

Both the up and down quarks, as well as the leptons, have second and third generation partners. The second generation of quarks are the charm and strange quarks, and the third generation the top and bottom quarks. The second and third generations associated to the electron are the muon and tau lepton. Each of these also have an associated neutrino, a neutral lepton. These are the electron neutrino, muon neutrino, and tau neutrino (ν_e, ν_μ, ν_τ).

The forces acting on these twelve fermions are mediated by exchange of a collection of bosons. The strong force is mediated by the exchange of gluons, the EM force by the photon, and the weak force by the $W^\pm$ and Z bosons. Gluons and photons are massless, whilst the $W^\pm$ and Z bosons are some of the most massive particles in the SM. Above a certain energy scale, the EM and weak forces are unified into the *electroweak* (EW) force [8–10].

The SM is completed by a further, scalar, boson — the Higgs boson, discovered in 2012 by the ATLAS [6] and CMS [7] collaborations. The discovery of a Higgs-like particle was significant in its confirmation of the postulated Higgs field which is responsible for the Higgs mechanism. This was independently proposed by Robert Brout and Francois Englert [11], and Peter Higgs [12] as an explanation for the origin of mass within the SM.¹ EW symmetry breaking provides mass to the gauge bosons, with fermions gaining mass through Yukawa coupling to the Higgs field. A fermion's mass is proportional to the strength of its coupling to the Higgs field, meaning that a heavier particle couples more strongly to the field. The fermion masses, and hence their Yukawa couplings, are free parameters in the SM and therefore have required experimental measurement.

Together, all these pieces of the SM are shown in Figure 1.1.

The SM is formalised by the gauge symmetry

$$G_{\text{SM}} = SU(3)_C \otimes SU(2)_L \otimes U(1)_Y, \quad (1.1)$$

¹For this reason, it may be more appropriate to refer to the Brout-Englert-Higgs (BEH) field, but the more commonly adopted *Higgs field* is used here for brevity.




Figure 1.1: The Standard Model of particle physics. Each particle is shown with its mass, electrical charge, and spin. The shaded regions denote which forces are felt by which particles. Image from Ref. [13].

where $SU(3)_C$ represents the strong interaction governed by quantum chromodynamics (QCD), with the colour, C , the conserved current. $SU(2)_L \otimes U(1)_Y$ represents the electroweak interaction, governed by quantum electrodynamics (QED) and quantum flavour dynamics (QFD), with the conserved currents L and Y the weak isospin and weak hypercharge, respectively.

Which forces act on a given particle depend on its charge. All matter particles in the SM interact weakly. It is the weak interaction through the $W^\pm$ bosons which is responsible for the particle flavour-change which leads to radioactive decay. Only particles contained within a weak isospin doublet — such as (u, d) or (e, ν_e) — may change flavour. The weak interaction acts over relatively small distance scales because of the high mass of its bosons, which means that they have a short lifetime.

Particles carrying electric charge interact electromagnetically. In the SM, this means the $W^\pm$ bosons, the charged leptons, and the quarks. The photon is massless and as a result the EM force has an infinite range. In terms of the elementary charge, e , the up-type quarks (u, c, t) have an electric charge $+\frac{2}{3}e$, whilst the down-type quarks (d, s, b) have a charge of $-\frac{1}{3}e$. The electron, muon, and tau all have charge $-e$. The neutrinos have no electric charge.

Only particles with colour charge feel the strong interaction. The only particles which interact strongly are then the quarks and the gluons themselves. Gluons may interact with themselves through the strong force. The self-interaction gives rise to the property of confinement within QCD, which says that particles carrying colour charge cannot exist in isolation but must be confined to colour-singlet bound states; neutral hadrons. This principle of confinement in turn produces the phenomenon of *hadronisation*. This can be demonstrated by the example of two quarks produced at some energy. Here, each quark is bound to the other by a colour field contained in a tube between them which contains an energy which grows with increasing distance. As the quarks become further separated, more energy is stored within the field until passing the threshold required to spontaneously produce another pair of quarks from the vacuum. This process repeats until the energy of each pair of quarks is low enough that the production of new pairs can stop, with an ensemble of colourless hadrons the result. It is this collection of hadrons which are reconstructed as *jets* experimentally. This is discussed in detail in Section 3.1.

Each of the 17 particles shown in Figure 1.1 also has an *antimatter* particle — its *antiparticle* — which is a copy with the same quantum numbers but inverted charge.




Figure 1.2: Feynman diagrams for various SM interactions.

For the top quark, for example, there is the *antitop* quark, $\bar{t}$, which has electric charge $-\frac{2}{3}e$. These particles arise as the interpretation of the negative solutions of the Dirac equation. Each of the fermions has a distinct antiparticle. The neutral bosons are their own antiparticle, and the $W^\pm$ bosons are each other's antiparticle.

The interactions described by the SM can be represented in *Feynman diagrams*, which represent all time-orderings for a given process and may be related to the likelihood — encoded in a quantum-mechanical matrix element — of the process by Feynman rules. Examples of Feynman diagrams showing the effect of the SM forces can be seen in Figure 1.2. In Feynman diagrams, time runs from left to right, in that the left-hand side represents the initial state and the right-hand side the final state.

Interactions between bosons and fermions have a strength which is proportional to the coupling constant, g . The matrix element (ME) $\mathcal{M}$ for a given process includes a factor of g for each vertex in the diagram. The probability of an interaction is actually proportional to $\mathcal{M}\mathcal{M}^*$, so that each vertex contributes g^2 to the probability for the interaction. Usually, a dimensionless constant $\alpha \propto g^2$ is used instead of g . Each of the SM forces has its own constant which characterises the strength of the interaction. For the EM force, this is the fine-structure constant $\alpha = \frac{1}{137}$. The intrinsic strength of the weak force is determined by $\alpha_W \sim \frac{1}{30}$, and the strong interaction by $\alpha_S \sim 1$.²

² α_S is known as the strong coupling constant, which is something of a misnomer as it is not constant but has a value which is scale-dependent [14–19].

All of the components of the SM are brought together and encoded within the SM Lagrangian:

$$\mathcal{L}_{\text{SM}} = \underbrace{\frac{1}{4}F_{\mu\nu}F^{\mu\nu}}_{\mathcal{L}_{\text{Gauge bosons}}} + \underbrace{i\bar{\psi}\mathcal{D}\psi}_{\mathcal{L}_{\text{Fermions}}} + \underbrace{(\psi_i y_{ij} \psi_j \phi + \text{h.c.})}_{\mathcal{L}_{\text{Yukawa coupling}}} + \underbrace{(|\mathcal{D}_\mu \phi|^2 - V(\phi))}_{\mathcal{L}_{\text{Higgs}}}. \quad (1.2)$$

This is a heavily compacted form of the formula, which is split into terms encoding different aspects of the theory. These are:

1. $\mathcal{L}_{\text{Gauge bosons}}$: This term describes all of the gauge bosons in the SM, meaning the force carriers but not the Higgs boson. In this term, F is the field strength tensor which describes the strong force and the electroweak force. $F_{\mu\nu}F^{\mu\nu}$ is then the scalar product of this tensor. It contains kinetic terms for photons, and kinetic and interaction terms for the carriers of the weak and strong forces.
2. $\mathcal{L}_{\text{Fermions}}$: The interactions between the gauge bosons and the matter particles, the fermions, are described here. ψ represents the fermion fields. $\mathcal{D}$ is the covariant derivative describing the interactions of the gauge bosons without self-interactions. This term is sometimes shown with an addition of h.c., meaning Hermitian conjugate, but this is not necessary as $\mathcal{L}_{\text{Fermions}}$ is self-adjoint.
3. $\mathcal{L}_{\text{Yukawa coupling}}$: This term describes the Yukawa couplings of the fermions to the Higgs field. y_{ij} is the Yukawa matrix which prescribes the strength of these couplings for each particle, related to the mass of the particle. Unlike $\mathcal{L}_{\text{Fermions}}$, $\mathcal{L}_{\text{Yukawa}}$ is not self-adjoint, so the Hermitian conjugate, h.c., is added at the end, and describes the interactions between the antifermions and the Higgs field.
4. $\mathcal{L}_{\text{Higgs}}$: With $\mathcal{D}_\mu$ the covariant derivative of the gauge boson fields, $\mathcal{D}_\mu \phi$ describes the interaction of the massive gauge bosons with the Higgs field ϕ . $V(\phi)$ is the potential of the Higgs field, with its non-zero vacuum expectation value giving rise to spontaneous symmetry breaking.

Although it is an extraordinary theory, the SM is not a complete description of nature, falling short in a number of ways:

- **GRAVITY**: The SM does not describe gravity, and it is not clear how this would be incorporated into the theory. The omission of gravitational effects does not pose a problem for particle physics experiments as gravitational forces are orders of magnitude weaker than those described by the SM and are therefore negligible

by comparison. It is a problem for the theory, however, as it is a clear indication that the SM cannot represent a *final* theory of nature.

- **DARK MATTER:** Astronomical observations show that ordinary matter, of the kind described by the SM, accounts for only around 5% of the energy in the universe. Around 26% is comprised of matter which, if it interacts with the SM particles at all, does so only weakly. The SM does not provide an explanation for what this matter might be.
- **DARK ENERGY:** The remainder of the energy in the universe is made up of *dark energy*, which is driving the accelerated expansion of the universe. The SM also provides no explanation for this.
- **MATTER-ANTIMATTER ASYMMETRY:** When they meet, matter and antimatter particles annihilate. The SM predicts that matter and antimatter should exist in equal quantities. In such a universe, this should lead to widespread annihilation, which is inconsistent with the matter-dominated universe which is observed. This means that there is more matter than antimatter in the universe. The mechanism within the SM for producing such an asymmetry, however, is insufficient in accounting for the observed discrepancy.
- **NEUTRINO MASSES:** The SM does not explain how neutrinos gain mass. Observations [20–22] of neutrino oscillations have demonstrated that neutrinos do have mass, albeit orders of magnitude lower than the other particles in the SM.

All of this is to say that the SM must be a stepping stone towards some more complete theory. During its conception, it was hoped that the LHC collider would provide confirmation of the existence of the Higgs boson and associated field, as well as making discoveries of physics beyond the SM (BSM). Although the former hope has been realised, the latter remains out of reach for the time being.

1.2. Top quark physics

The top quark was discovered in 1995 by the D0 [23] and CDF [24] experiments at Fermilab, having been predicted in 1973 [25]. It is the most massive of all the SM particles and indeed it is for this reason that its discovery took so long; only the Tevatron and the LHC have produced collisions at sufficiently high energies.

At $172.69 \pm 0.30 \text{ GeV}$ [26], its mass lends the top quark many interesting properties. It decays very quickly, with a lifetime around $\tau = \mathcal{O}(10^{-25} \text{ s})$. Unique to the top quark, this is a shorter time than the characteristic time for hadronisation, $t = \mathcal{O}(10^{-23} \text{ s})$ [27]. This allows the top quark to be studied as a *bare* quark, reconstructed directly from its decay products. Because it is the only quark with such a mass, it is able to decay into a W boson and a down-type quark. As dictated by the $|V_{tb}|$ element of the CKM matrix [25, 26, 28], this down-type quark is almost exclusively a b -quark. The decay is then from a top quark into a W boson and a b -quark.

For the production of top-antitop pairs ($t\bar{t}$) then, there are two W bosons and two b -quarks produced. From there, $t\bar{t}$ events can be classified into three distinct channels: *all-hadronic*, *lepton+jets* (ℓ +jets), and *dilepton*. These channels are defined by how the two W -bosons decay. In the all-hadronic channel, both decay into a pair of quarks. In the ℓ +jets channel, one W decays into a pair of quarks and the other into a lepton and associated neutrino. Finally, if the event proceeds dileptonically, both W bosons decay into a lepton and associated neutrino.

The majority of $t\bar{t}$ events are all-hadronic or lepton+jets, 46% and 44% respectively, as shown in Figure 1.3. Although dileptonic events offer a “clean” signal, leading to good selection efficiency, there are far fewer such events because of the small branching ratio and the presence of two neutrinos leads to combinatorial challenges with the event reconstruction. The all-hadronic channel offers a far greater number of events to analyse but lacks the clean signature of the lepton with which to trigger. It is also difficult to separate the all-hadronic channel from other processes which also include several jets. In the ℓ +jets channel, the presence of the lepton allows for greater triggering efficiency with smaller backgrounds than the all-hadronic channel and more events than the dilepton channel.

At the LHC, the dominant $t\bar{t}$ production mode is gluon-gluon fusion, with the remainder through $q\bar{q}$ annihilation. These production modes are shown in Figure 1.4.

For any given process, there are an infinite number of possible Feynman diagrams which can be drawn. It is possible to add any number of arbitrarily soft emissions, or an arbitrary number of virtual gluon loops, as is shown in Figure 1.5 for $t\bar{t}$.

At a collider such as the LHC, the number of events observed for a given process is given by a combination of the luminosity delivered by the accelerator and the




Figure 1.3: Branching ratios for possible $t\bar{t}$ final states. Image from Ref. [29].




Figure 1.4: Leading-order $t\bar{t}$ production modes.




Figure 1.5: Feynman diagrams for $t\bar{t}$ production through $q\bar{q}$ annihilation.

Standard Model Production Cross Section Measurements

Status: February 2022




Figure 1.6: Summary of ATLAS cross-section measurements for several SM processes. The values here are corrected for the branching ratio of the process. Image from Ref. [30].

cross-section of the process. Staying with $t\bar{t}$ as the example, this can be expressed:

$$N_{t\bar{t}} = \sigma_{t\bar{t}} \int \mathcal{L}(t) dt, \quad (1.3)$$

where $\mathcal{L}$ is the instantaneous luminosity, which is a measure of the event rate, integrated over the time period data is taken over to give the integrated luminosity. $\sigma_{t\bar{t}}$ is the cross-section for $t\bar{t}$ production, which is a measure of the probability of the interaction, and $N_{t\bar{t}}$ is the number of events observed. Along with many other SM processes, a comparison between cross-section predictions and measurements made by ATLAS is shown in Figure 1.6.

The cross-section may be obtained by summing over all possible matrix elements. Because the vertices relevant to $t\bar{t}$ production are those of QCD, each vertex has a factor of α_s associated to it. At the scales present in LHC collisions, α_s is below unity, so a process becomes less and less probable for every additional QCD vertex with

an associated factor of α_s . So although Figures 1.5a and 1.5b have identical initial and final states, and thus represent the same process, the latter is less likely to occur because of the additional factor of α_s associated with the emission of an additional gluon.

In the language of perturbation theory, a cross-section can be calculated up to a given *order* through an expansion in α_s . The treatment of QCD as a perturbative theory is only valid at sufficient energy scales, and this is only possible because α_s tends to smaller values below unity at these energies, a behaviour known as asymptotic freedom. The “leading” order (LO) is for the lowest power of α_s , with each subsequent order — next-to-leading order (NLO), next-to-next-to-leading order (NNLO)... — adding a power of α_s and thus a lesser influence on the process. Calculations up to a given order in α_s are known as fixed-order calculations, and the most precise of these for $t\bar{t}$ production run up to NNLO [31–34]. To account for contributions from diagrams with EW vertices, EW corrections at NLO [35,36] are also factored in. Soft and collinear gluon emissions may also impact the cross-section and these are accounted for with next-to-next-to-leading logarithm (NNLL) re-summation [34,37–39].

In order to provide an explanation for many of the open questions which face the SM, many models of BSM introduce new particles and forces. In the attempt to observe this new physics, there are two main approaches which may be taken. Firstly, direct searches for these hypothetical particles may be performed. If a theory predicts a new particle associated with an extension or overhaul of the SM, the final state it is expected to produce may be measured in the same way as analyses focused on SM processes. The SM has been probed extensively at the LHC, up to higher energies than ever before, and so far without any sign of these new particles.

The second approach may be termed “anomaly detection”, and involves measuring SM processes in order to look for deviations between what is predicted by the SM and what is observed in the data, with the cause of these deviations being BSM physics.

Owing to its mass, the top quark is expected to play an important role in any new physics model. It is expected to couple strongly to many proposed BSM particles, as well as having the strongest coupling to the Higgs boson. Coupling to BSM particles could manifest as modifications to the top cross-section. As can be seen in Figure 1.6, top production at the LHC is also abundant, which makes it an important background for many other measurements of extreme regions of SM phase space and searches for BSM physics. Precise modelling of these processes is then of the utmost importance

for improved sensitivity in these channels. The Monte Carlo (MC) descriptions of top processes — discussed in the next section — may be improved by tuning derived from precise measurements which are made in data. This helps not only with identification of top processes in order to either accept and measure them as signal or reject them as background, but also with the development of analysis techniques. Given that there are BSM theories predicting particles which decay to $t\bar{t}$ pairs, which would be produced at very high energies, the resultant collimation of the decay products of the top quarks requires ever-improved techniques in these high-energy regions of phase space. With all this said, precision measurements in the top quark sector such as the analysis presented in this thesis are then of great importance to the ATLAS physics program.

1.3. Event simulation

The analyses performed at LHC experiments like ATLAS are dependent on simulations of events which allow for comparisons to be made between theoretical predictions and measurements made by the detector. This simulation of events is performed using Monte Carlo (MC) methods in a multi-staged approach shown in Figure 1.7.

The MC simulation of events [43] proceeds in several steps. First, the ME — introduced in Section 1.1 — must be calculated for the physics process being simulated. As before, here $t\bar{t}$ is used as an example. The ME then corresponds to the $pp \rightarrow t\bar{t}$ process. The ubiquitous soft and collinear radiation produced by partons in both the initial and final state are simulated as a parton shower (PS). The references made later in this thesis to NLO+PS predictions therefore refer to a matrix element calculation at NLO used in tandem with some PS scheme. Hadronisation which then proceeds from the partons must then be simulated by a non-perturbative method suitable for this lower energy regime.

Deposits in the detector for a given event are the result of far more than just the hard scatter of interest, however. Proton remnants are liable to interact, decay, and hadronise in processes referred to as the underlying event (UE) and multi-parton interaction (MPI). Under LHC conditions, for any proton bunch-crossing, there are likely to be tens of proton interactions, referred to as pileup (PU). These interactions and the additional radiation they contribute to an event can originate either from the




Figure 1.7: Shown are the numerous components which must be simulated for a $t\bar{t}$ event, as performed using PYTHIA [40]. From the incoming protons, shown in the pink circles either side of the diagram, the hard scatter process can be seen in the centre, producing a $t\bar{t}$ pair which subsequently decays. Also present are multi-parton interactions (MPI), the parton shower (PS), and the hadronisation process resulting in stable hadrons present in the final-state. Image from Ref. [41, 42].

same bunch-crossing as the hard-scatter of interest — in-time PU — or from adjacent bunch-crossings — out-of-time PU — and must also be simulated.

Owing to the fact that α_s takes a value which is scale-dependent, the modelling of the pp collisions taking place at the LHC is challenging as these events encompass a wide range of energy scales. The Factorisation Theorem [3, 44] states that the cross-section for a hard-scatter process can be factorised into two components, corresponding to low- and high-energy regions corresponding to non-perturbative and perturbative elements, respectively. A factorisation scale μ_f is chosen as the boundary between these perturbative and non-perturbative regimes, often around m_t^2 , the mass of the top quark. This then allows the cross-section for a hard scatter process, here $t\bar{t}$, to be expressed as

$$\sigma_{pp \rightarrow t\bar{t}} = \sum_{i,j}^{\text{parton flavours}} \int_0^1 \int_0^1 dx_1 dx_2 f_i(x_1, \mu_f) f_j(x_2, \mu_f) \hat{\sigma}_{ij \rightarrow t\bar{t}}, \quad (1.4)$$

where f_i and f_j are the parton distribution functions (PDF) with x_1 and x_2 the momentum fraction representing the momentum of the parton in the interaction as a proportion of the first and second proton's momentum, respectively. $\hat{\sigma}_{ij \rightarrow t\bar{t}}$ is the partonic cross-section of the $t\bar{t}$ process, evaluated perturbatively at the factorisation scale μ_f , as discussed in Section 1.2.

The distribution of energy within protons is described by PDFs, and is divided not only between the valence quarks — uud — but amongst the *sea* of quarks and gluons. A PDF $f_i(x, \mu_f)$ describes the probability that a parton $i \in \{q, \bar{q}, g\}$ with momentum fraction x is produced from the protons which are being collided, subject to the factorisation scale μ_f . The DGLAP equations [45–47] describe the energy dependence of the PDFs, which can be calculated perturbatively, but the momentum fractions cannot be calculated analytically. Precise measurements of the proton structure, such as those made by the ZEUS [48] and H1 [49] experiments at the HERA collider [50] are then used to fit parameterisations of these distributions. As well as μ_f , a renormalisation scale μ_r (usually $\mu_f = \mu_r$) is chosen for the renormalisation of the SM Lagrangian in order to account for both the soft and collinear emissions and high-energy loops which may enter at any order of the perturbative calculation of the matrix element.

PS algorithms are then used to model the evolution of the event following the fixed-order generation of the hard scattering process. These are used to provide an

approximation for the higher-order contributions from additional emissions and loops, whose inclusion in the perturbative calculation is not possible.

A parton, whether in the initial state radiation (ISR) consisting of the incoming partons, or the final state radiation (FSR), has some probability of splitting into two partons, each with some share of the original momentum. The PS calculates the branches and evolves the system until a scale around 1 GeV is reached using the Sudakov form factor [3],

$$\Delta(t_0, t) = \exp \left[- \int_{t_0}^t \frac{dt'}{t'} \int_{z_{\min}}^{z_{\max}} dz P(z) \right], \quad (1.5)$$

which gives the probability that a given parton evolves from virtuality $t_0 \rightarrow t$ ($t \equiv p^2$) and does not emit any resolvable partons, as determined by the momentum fraction limits $z_{\min}$ and $z_{\max}$. $P(z)$ is a DGLAP splitting function which dictates the probability of quark/gluon splittings.

The virtuality given here is the *ordering* variable, which prescribes the sequence of the splittings in the shower.³ The virtuality is only one possible choice, and different event generators use different ordering variables. The PYTHIA [40] and HERWIG [51] generators used in the analysis presented here use ordering variables which are momentum-based and angular, respectively. A *matching* process is used when interfacing the hard scatter parts of an event from the matrix element with the emissions from the PS to ensure that there is no double-counting of emissions. Different approaches are taken by different frameworks, with some applying an upper limit to the momentum of emissions from the PS and others subtracting the PS emissions with the highest energies from the ME.

With an ensemble of partons, the process of these partons forming colour singlet bound states — hadronisation — must be simulated. At this stage in the evolution of the event, the energies are much lower and α_s is too large to be treated perturbatively, so an alternative approach must be taken. There are two common methods used for modelling hadronisation — the Lund string model [52] and cluster hadronisation [53].

In the Lund string model, given that the potential between a quark and an antiquark — beyond very short distances — increases linearly, the force holding the pair together also increases constantly as they are pulled apart and the flux tube, or string, between

³ t is used as this makes the ordering variable in some way analogous to time.

the two becomes more concentrated. The string has an energy of around 1 GeV per fm such that there comes a point as the $q\bar{q}$ pair are separated that there is enough energy to produce a new $q\bar{q}$ pair, with the string being split in two and creating two new strings. These new $q\bar{q}$ will also be flying apart and so this process repeats, and so on. At some point, there is no longer enough energy to produce new pairs, and the process stops and hadrons are formed of groups of adjacent quarks and antiquarks. In addition to these mesons, baryons may be formed from groupings of quarks and *diquarks*. The Lund string model is the approach taken in PYTHIA.

HERWIG adopts the cluster model, which is based on *preconfinement*. The principle here is that the PS traces the colour flow and so the partons which result from the PS procedure can be clustered by their colour charge into colour singlets. Gluons are split into $q\bar{q}$ pairs to form clusters, and clusters decayed into stable hadrons.

Any interaction in the event not associated with the hard scatter process of interest is marked as UE. This definition is inclusive of MPI, dealt with by dedicated models, which is concerned with interactions between remnants that are not in colour-confined states and their subsequent hadronisation. As is also true with many of the parameters in the Lund string model, these models are often tuned to experimental data [54].

Pileup originates from interactions between other pairs of protons and is uncorrelated with the hard scatter process being simulated. Accordingly, it is simulated by overlaying the event with additional QCD hard scatter interactions simulated using the same method as the process of interest. These events then undergo a reweighting procedure such that the PU distributions match that in the data. This is done separately for each data-taking period.

The final stage is to simulate the detector response for each event. This is handled with an extremely detailed model of the ATLAS detector in GEANT4 [55], which provides a precise description of the expected signal in the detector given the stable particles produced by the event generation process. As such, the detector simulation is one of the most CPU-intensive operations in this whole chain. A slightly less precise model of the detector response has also been developed, called AFII [56], in order to reduce these computational demands, but this is not used in the context of the analysis presented in this thesis, and is better suited to high-level observables than the type of measurement undertaken in the chapters ahead.

More details on the specific MC predictions used in the analysis presented in this thesis are given in Section 4.1.

Chapter 2.

The Large Hadron Collider and the ATLAS Experiment

“My Lords, can my noble friend tell us what a large hadron collider is, and whether a smaller one might not do?”

— Lord Elton in the UK House of Lords, 18th July 1994

2.1. The Large Hadron Collider

The Large Hadron Collider (LHC) [57] is the world’s largest and highest energy particle accelerator, located at The European Organisation for Nuclear Research (CERN), just outside Geneva. Around 100 m underground, straddling the French-Swiss border, the LHC sits in a tunnel 27 km in circumference and collides protons with an energy of 6.5 TeV, giving a centre of mass energy of $\sqrt{s} = 13$ TeV.

The LHC chain comprises several smaller accelerators which feed into higher and higher energy accelerators before feeding the protons into the LHC to be accelerated up to 6.5 TeV. A bottle of hydrogen is the source for the whole process. This hydrogen is ionised and the resultant protons accelerated to 50 MeV by LINAC2, a linear accelerator. This feeds into LINAC4 which accelerates up to 160 MeV before the protons are fed into the Proton Synchrotron Booster, Proton Synchrotron and Super Proton Synchrotron to accelerate to 1.4 GeV, 26 GeV and 450 GeV, respectively. The Super Proton Synchrotron was itself once the most powerful particle accelerator in the world and is 6.9 km in




Figure 2.1: Schematic view of the accelerators and experiments which make up the CERN accelerator complex [60].

circumference. Only after this point are the protons injected into the LHC to accelerate each beam to 6.5 GeV. The entire CERN accelerator complex, including the LHC feeder chain, is shown in Figure 2.1.

The LHC's first run took place between 2010 and 2013, stepping from $\sqrt{s} = 7$ TeV to $\sqrt{s} = 8$ TeV, reaching 75% of the nominal instantaneous luminosity of $10^{34} \text{ cm}^{-2} \text{ s}^{-1}$. Run 2 raised the centre of mass energy to $\sqrt{s} = 13$ TeV and rose from the nominal luminosity to a peak at twice that, running from 2015 to 2018. Run 3 began in 2022 with $\sqrt{s} = 13.6$ TeV and again operating at $2 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$. Prior to the start of Run 3, the LHC was in shutdown to allow for maintenance as well as upgrades to the feeder chain, magnets and the first stage of upgrades to the detectors around the LHC ring in preparation for the eventual upgrade to High Luminosity LHC (HL-LHC) [58]. At this point, the LHC will reach five times its nominal luminosity with the aim of an order-of-magnitude increase to the integrated luminosity over the LHC's lifetime to 4000 fb^{-1} [59].




Figure 2.2: The ATLAS Detector during construction. The barrel sections of the superconducting toroid magnet system can be clearly seen.

2.2. The ATLAS Experiment

ATLAS [61] (A Toroidal LHC ApparatuS), shown in Figure 2.2, is the largest general purpose particle detector ever constructed. It is housed at LHC Point 1, in a cavern approximately 100 m beneath the surface. At a length of 46 m, a height of 25 m, and a mass around 7000 t, it is the largest of all the LHC experiments.

Proton bunches are crossed in the centre of the detector, at the nominal interaction point (IP). ATLAS is designed to be a fully hermetic detector, covering as much of the volume surrounding the nominal IP as possible. This ensures that all particles emerging from a hard scatter event must traverse the detector and that the missing transverse energy in a given event may be calculated.

ATLAS employs a right-handed coordinate system in which the z -direction is the beam-direction, with x pointing towards the centre of the LHC and y in the vertical direction. ϕ and θ are the azimuthal and polar angles, respectively. θ is often snubbed in favour of the *rapidity*, defined for a particle with momentum in the beam-direction p_z and energy E by $y \equiv \frac{1}{2} \ln \frac{E+p_z}{E-p_z}$. Differences in rapidity are invariant under Lorentz boost, which makes it a preferable quantity to the angle θ . Knowledge of E and p_z can be difficult to come by at hadron colliders like the LHC, so in practice it is normally the *pseudorapidity* which is used. The pseudorapidity, $\eta \equiv -\ln \tan \frac{\theta}{2}$, approximates to the rapidity for highly relativistic particles. Two particles, i and j , can be said to be





Figure 2.3: Schematic view of the ATLAS detector, with constituent sub-detectors labelled [62].


separated by

$$\Delta R_{ij} = \sqrt{(\eta_i - \eta_j)^2 + (\phi_i - \phi_j)^2}.$$

The ATLAS Detector is made up of several sub-detector systems, each of which is specialised for a specific task. These sub-systems can be seen in Figure 2.3. Figure 2.4 shows the total integrated luminosity recorded by ATLAS over Run 2 as well as the pileup conditions over that period. How various kinds of particles interact with each of the sub-detector systems can be seen in Figure 2.5. The Inner Detector is the innermost system, encapsulating the interaction point in order to measure the tracks of charged particles. The calorimeters are arranged outside of the Inner Detector and aim to measure the energy of particles emanating from an interaction. The Muon Spectrometer aims to measure the tracks of muons as they traverse the rest of the detector. For measuring particle tracks, ATLAS utilises both a solenoid magnet system and a toroid magnet system in order to deflect the paths of charged particles. Readout and storage limitations require that a trigger system is in place to make quick decisions on which measurements are read out and stored to disk. Each of these sub-detector systems is described in this section.



(a) The total integrated luminosity over the whole Run 2 data-taking period [63].



(b) The mean number of interactions per bunch crossing for each of the calendar years entailed in Run 2, as well as for the whole run [63].

Figure 2.4: ATLAS operating status over the course of Run 2.




Figure 2.5: How known particles interact with the various sub-detectors within ATLAS [64].

2.2.1. Inner detector

The first of the sub-detectors in ATLAS is the Inner Detector (ID) [65]. In short, the ID is designed to make precise spatial measurements of particles as they pass through its many layers in order that tracks may be reconstructed from these points. This is done under a 2 T field generated by the solenoid magnet system in order that the momenta and charge of these particles may be inferred from their reconstructed tracks. Further particle identification is enabled by the resolution of both primary and secondary vertices, which also comes from evaluation of the reconstructed tracks.

The interactions of particles with matter can be characterised by the radiation length, X_0 , for electromagnetic interactions and by the nuclear interaction length, λ_I , for the development of hadronic showers by nuclear interaction. A radiation length prescribes the average distance over which the energy of an electron is reduced by a factor of $\frac{1}{e}$ by the Bremsstrahlung process as it traverses the material. For an energetic photon, it is also approximately $\frac{7}{9}$ the mean free path of the e^+e^- pair-production process in that material [66]. The nuclear interaction length describes how far on average a hadron will travel before hadronic interaction in the material. The nuclear interaction length is larger than the radiation length of a material.

For the ID, minimising the material budget in order to present as few radiation lengths or nuclear interaction lengths to the traversing particle as possible is an important consideration. This is in order to make a precise measurement of the trajectory of the particle whilst minimising the influence of that measurement on the particle. Minimising the material budget is a complex optimisation of the materials used for the active elements of the detector as well as their support structures and services. How these services are routed must also be optimised to present as few radiation lengths as possible.

The ID is comprised of three sub-systems. From the innermost and working out from the interaction point, these are the Pixel Detector, the Semiconductor Tracker (SCT), and the Transition Radiation Tracker (TRT). The radial extent of the ID can be seen in Figure 2.6a, with the arrangement of the sub-systems visible in Figure 2.6b.




Figure 2.6: The ATLAS Inner Detector.

Pixel Detector

The Pixel Detector [67, 69] is the closest part of the detector to the beamline. As such, it must be equipped with sufficiently high resolution to handle the extremely busy environment which exists so close to the interaction point, as well as being adequately robust against the harsh radiation conditions present there. It is made up of four layers of silicon pixel modules in the barrel section, and three layers in the end-caps. The innermost layer, the insertable *B*-layer (IBL) [70], was installed as an addition to the Pixel Detector during LHC Long Shutdown 1 (LS1) and differs in some respects from the rest of the pixel subsystem.

The original Pixel Detector covers a pseudorapidity region $|\eta| < 2.5$ and is made up of 1744 pixel modules capable of providing measurements with resolutions of $10\text{ }\mu\text{m}$ in the transverse plane and $110\text{ }\mu\text{m}$ in the longitudinal direction. These modules are *hybrid* pixel detectors, meaning that the silicon sensor and the electronic readout chip are manufactured separately and *bump-bonded* together. Each sensor is $250\text{ }\mu\text{m}$ thick and segmented into pixels of pitch $50 \times 400\text{ }\mu\text{m}^2$. Each module comprises a sensor with 16 FE-I3 [71] readout chips bump-bonded to it. The FE-I3 chip is segmented into an 18×160 matrix, with each cell connected to a pixel on the sensor. Altogether, this yields 47 232 pixels per module and a total of over 80 million readout channels.

For a given module, all 16 FE-I3 chips are combined by the Module Control Chip which handles triggering, timing, and readout. Readout speed depends on the layer of the Pixel Detector, with rates of 160 Mbit s^{-1} , 160 Mbit s^{-1} , and 40 Mbit s^{-1} for the *B*-Layer, Layer 1, and Layer 2 in the barrel respectively, and 80 Mbit s^{-1} for the end-cap disks.

During LS1, the IBL was installed to provide a fourth layer to the barrel section of the Pixel Detector. The IBL sits inside the *B*-Layer, 3.27 cm from the interaction point, and covers $|\eta| < 3.0$. The IBL furthers ATLAS' capability for vertexing and *b*-tagging through its proximity to the interaction point and greater granularity. It also provides an aspect of redundancy in that tracks in the ID may now comprise three hits in the Pixel Detector even with one layer in the barrel missing. This becomes increasingly advantageous as the effects of prolonged exposure to the harsh radiation environment degrades the performance of the *B*-layer.

The IBL is made up of 14 carbon staves tilted in ϕ by 14° , each of which has 32 FE-I4 [72] readout chips bump-bonded to silicon sensors. The FE-I4 chip has pixels of pitch $50 \times 250 \mu\text{m}^2$, segmented into a 336×80 array. Each staff also has an integrated CO_2 cooling pipe.

The IBL utilises two distinct sensor technologies; planar, as with the rest of the Pixel Detector modules, and 3D. In each staff, the central region holds twelve planar modules each with two FE-I4 chips. Each end of the staff then also has four single-chip 3D modules. This means that each staff then hosts a total of twenty IBL modules.

The 3D sensors have a thickness of $230 \mu\text{m}$ and utilise n^+ columns from the front side and p^+ columns from the back side. The planar sensors have an even lower thickness of $200 \mu\text{m}$.

Semiconductor Tracker (SCT)

The SCT is a silicon microstrip detector which extends out radially between 299 mm and 560 mm [73]. It is designed to ensure at least four measurements for a given charged particle traversing its layers.

The SCT comprises four barrel layers and nine end-cap disks on each side of the detector. There are a total of 4088 modules, with 2112 in the barrel region and 988 in




Figure 2.7: Schematics of an SCT module showing both offset layers as well as the base board providing mechanical and thermal structure.

each end-cap. The barrel region covers $|\eta| < 1.5$ and the end-caps extend coverage to $|\eta| < 2.5$.

In the barrel region, each module is made up of four silicon strip sensors arranged in back-to-back pairs which are offset by 40 mrad with respect to each other such that each layer provides two measurements. These sensors have a pitch of $80\ \mu\text{m}$ and a length of 6.4 cm, such that the two sensors placed end-to-end to form each of the pairs extend over 12 cm, with 768 active strips. This provides resolution of $17\ \mu\text{m}$ in $R - \phi$ and $580\ \mu\text{m}$ in z . Each strip is connected to a channel on an ABCD3TA front-end readout chip [74]. Each ABCD3TA has 128 channels so that there are 6 of these readout chips per sensor pair, or 12 per SCT module. Each side of the module is glued to a base board in the middle which forms the mechanical and thermal structure of the module allowing for sensor cooling and biasing. This can be seen in Figure 2.7.

In the end-cap region, sensors are trapezoidal with a mean pitch of $80\ \mu\text{m}$ and again arranged back-to-back and offset by 40 mrad. The inner disks have one sensor on each side, whilst the outer disks have two sensors on each side. All sensors, whether in the barrel or end-cap region, have a thickness of $285 \pm 15\ \mu\text{m}$.

For both the Pixel Detector and the SCT, operating temperature is around 10°C , with cooling provided by evaporating C_3F_8 at around $\sim -25^\circ\text{C}$, with the exception of the IBL. This is to ensure optimal performance over the lifetime of the detector components as they suffer radiation damage.

Transition Radiation Tracker (TRT)

The TRT is the outermost subsystem within the ID. In contrast with the Pixel Detector and the SCT, the TRT is a gaseous detector. It is made up of around 300 000 straw-tubes. These are drift tubes designed, as with the other components of the ID, to make measurements of particle trajectories and momenta. As in the other parts of the ID, there is a barrel section of the TRT as well as end-cap disks.

Each drift tube has a diameter of 4 mm [75]. The walls of these tubes are held at -1.5 kV, with a gold-plated tungsten wire of diameter $31\text{ }\mu\text{m}$ at ground potential running down the centre. Each tube is then filled with a gas mixture which is 70% Xe, 27% CO_2 , and 3% O_2 .

The barrel region of the TRT contains straws of length 144 cm which are arranged parallel to the beam, with radial coverage from 560 mm to 1080 mm [76, 77], out to $|\eta| < 1.0$. Straws are typically separated by 6.8 mm in $R - \phi$. There are three concentric rings totalling 73 layers of straws.

End-cap sections contain straws of length 37 cm which are arranged perpendicular to the beam, with radial coverage from 644 mm to 1004 mm [76, 77], $1.0 < |\eta| < 2.0$. Each end-cap comprises twenty wheels. The first twelve of these wheels contain eight layers of straws with 8 mm separation in z between each layer. The outer eight of these wheels are arranged identically but with 16 mm z -separation between layers. This gives 160 layers of straws in each end-cap.

The straws in the TRT behave as drift tubes. As a charged particle passes through a straw in the TRT, the gas mixture is ionised. The liberated electrons drift towards the wire in the centre of the straw, where they are collected and the resultant signal amplified. By relating the drift time of the electrons to the anode, a measurement of particle position is made.

The spaces between straws are filled with polypropylene fibres in the barrel section and polypropylene foils in the end-cap region. As a relativistic charged particle passes between the boundary between this material and the TRT straws, *transition radiation* is produced, in the form of soft X-rays. These photons can be absorbed by the Xe in the straws. The production of this transition radiation depends on the relativistic gamma factor, $\gamma = \frac{E}{m}$, of the particle in question. This means that particles which produce this transition radiation leave stronger signals in the TRT. Particle identification can then be carried out by considering the probability that the signal left by some particle in a

straw exceeds some threshold. Electrons, for example, produce transition radiation so would be expected to leave a stronger signal than charged pions. This allows for powerful discrimination between electrons and charged pions.

The TRT has only a resolution of $130\ \mu\text{m}$ in $R - \phi$ but tracks in the TRT contain many more hits - usually more than 30 - and are much longer than those in the Pixel Detector or SCT.

2.2.2. Calorimeter system

After the ID comes the system of calorimeters used by ATLAS to make measurements of particle energies. The calorimeters in ATLAS are sampling calorimeters; comprising alternating absorber and active layers. Incident particles will interact with the absorber layers, producing secondary particles which can be measured in the active layers in order to provide an energy measurement. In contrast to the ID, the absorber layers in the calorimeters aim to present many radiation lengths or nuclear interactions lengths to the traversing particles in order to minimise how much energy is allowed to escape the system. The calorimeter system extends out to $|\eta| = 4.9$.

The calorimeter system is itself divided into the electromagnetic calorimeters and the hadronic calorimeters. The electromagnetic calorimeter is primarily concerned with electrons and photons whilst the hadronic calorimeter primarily measures the energy of hadrons.

Electromagnetic calorimeters (ECAL)

A Liquid Argon (LAr) system is employed by the ECAL in ATLAS. The ECAL comprises barrel, end-cap and forward sections. Lead is used as the absorber and LAr as the active layer. An incident particle will interact in the lead absorber, initiating an EM shower. As this shower passes through the LAr, the Ar atoms are ionised. Drift of the electrons from the ionised Ar is initiated by an electric field and a signal is induced as these electrons come into contact with readout electrodes which run through the LAr. The layers themselves are arranged in an accordion geometry.

A pre-sampler estimates energy lost in $|\eta| < 1.8$. The barrel section covers $|\eta| < 1.475$, the end-cap sections cover $1.375 < |\eta| < 3.2$, and the EM layer of the forward calorimeter (FCAL) covers $|\eta| < 4.9$ but with a copper absorber.




Figure 2.8: Schematic view of the ATLAS detector with the LAr calorimeter sub-systems highlighted [62].

The LAr calorimeters are contained within a cryostat at 88 K in order to maintain the liquid state of the Argon. The layout of the LAr calorimeter system is shown in Figure 2.8.

Hadronic calorimeters (HCAL)

The HCAL operates by the same broad principle as the ECAL - by alternating layers of absorber and active material - but with some variation in which materials are used.

In the barrel region of the detector is the tile calorimeter, comprising a central barrel along with an extended barrel section on either side. In the tile calorimeter, steel plates are used as the absorber material whilst plastic scintillating tiles are used as the active material. Each of these sections is segmented into 64 modules.

Electron-hole pairs are created in the scintillating tiles by incident charged particles resulting from interaction in the absorber. When electrons return to the valence level, photons are emitted. These photons are absorbed and re-emitted by wavelength-shifting fibres. The re-emitted photons are then measured by photomultiplier tubes (PMTs), from which a signal can be read out.

Section	Sampling term, a	Noise term, b	Constant term, c
ECAL	10%	170 MeV	0.7%
HCAL	50%	/	3%
FCAL	100%	/	10%

Table 2.1: Design requirements for the terms in the energy resolution of the ATLAS calorimeters.

In addition to the tile calorimeter, the FCAL has hadronic calorimetry sections. In each end cap, there are two wheels which provide coverage over $1.5 < |\eta| < 3.2$. These use copper as the absorber and LAr as the active material. Beyond here, the outermost layers of the FCAL use tungsten as absorber and LAr as the active material, covering $|\eta| < 4.9$. The hadronic sections of the end caps are contained within the same cryostats as the EM sections.

The relative resolution of a calorimeter is given by

$$\frac{\sigma_E}{E} = \frac{a}{\sqrt{E}} \oplus \frac{b}{E} \oplus c, \quad (2.1)$$

where a is a *sampling* or *stochastic* term which reflects the development of the shower, b is a noise term arising from readout electronics and the effect of pileup, and c is a constant term which accounts for effects such as the dead material within the detector and imperfections in the construction of the calorimeter [3,78]. The $\oplus$ symbol denotes addition in quadrature.

Different sections of the ATLAS calorimeters were designed with different energy resolution requirements. These are given in Table 2.1, and were found to be satisfied by the detector [77]. At high energies, the constant term dominates the energy resolution of the calorimeter.

The noise term, b , is not included for all sub-systems in the design requirements as it depends heavily on operating conditions at the LHC and varies significantly depending on pileup and location within the detector, for example. This is in contrast to the sampling and constant terms which are largely constrained by the detector design. Additionally, the noise term is usually insignificant beyond low energies. It is measured throughout operation and may be mitigated as necessary through noise subtraction techniques in event reconstruction.




Figure 2.9: The ATLAS Muon Spectrometer.

2.2.3. Muon Spectrometer (MS)

With the exception of neutrinos, which are invisible to the detector, only muons usually traverse beyond the calorimeter system. The measurement of their tracks under ATLAS' magnetic field, and thus momenta, necessitates a bespoke sub-detector — the Muon Spectrometer, shown in Figure 2.9. The MS fulfils two key roles — momentum measurement and triggering. The MS can make measurements of charged particle momentum in $|\eta| < 2.7$ and perform triggering in $|\eta| < 2.4$. To achieve this mix of precision measurement and triggering, the MS uses four distinct technologies: monitored drift tubes (MDT), cathode strip chambers (CSC), resistive plate chambers (RPC), and thin gap chambers (TGC). The MDTs and CSCs are used for making precise measurements but lack the speed to meet the triggering criteria of the MS. This capability is provided by the RPCs and TGCs. As with the other sub-detectors, the MS comprises barrel and end-cap regions. Each of these sub-systems is shown in Figure 2.10.

The barrel section has three concentric layers of precision tracking chambers at radii of 5 m, 7.5 m, and 10 m. These layers are positioned on and between the coils of the barrel toroid magnet. The MS as a whole mirrors the symmetry in ϕ of the toroid magnet system in that it is divided into octants.




Figure 2.10: A schematic view of the ATLAS detector with the sub-systems of the muon spectrometer highlighted [80].

The end-cap chambers of the MS are arranged in wheels which sit either side of the end-cap toroid magnet. Chambers form wheels which are perpendicular to the beamline at 7.4 m, 10.8 m, 14 m, and 21.5 m along the z -axis.

Monitored Drift Tubes

The MDTs cover $|\eta| < 2.0$ in the innermost layer and $|\eta| < 2.7$ everywhere else. There are three barrel layers and four end-cap layers. The MDT chambers are made up of between three and eight drift tubes, providing a resolution of around $35\ \mu\text{m}$ per chamber, as well as a certain redundancy in the case of the failure of an individual drift tube. There are 1150 MDTs in the MS, each 30 mm in diameter and filled with a mixture of Ar and CO_2 . The anode is a gold-plated wire running through the centre of the tube, with the wall acting as the cathode. Electrons produced in the gas mixture near the cathode have a drift time around 700 ns. The MS is arranged such that a track should pass through at least three different chambers.

Cathode Strip Chambers

The CSCs cover $2.0 < |\eta| < 2.7$. They allow for $40\text{ }\mu\text{m}$ resolution in the bending plane (5 mm in the transverse plane) and are used in the innermost tracking layer as they can handle higher rates than the MDTs. The CSCs are multi-wire proportional chambers. These again utilise a mixture of Ar and CO_2 , with wires as anodes running in the z direction. Segmented strips perpendicular to these wires act as the cathodes, providing resolution of $40\text{ }\mu\text{m}$ in the bending plane and 5 mm in the transverse plane.

In order to maintain the resolution for these precision measurements, a sophisticated alignment system is required to monitor the position of the various components of these parts of the MS.

Resistive Plate Chambers

The RPCs make up the barrel ($|\eta| < 1.05$) section of the fast trigger chambers in the MS. Each RPC is formed of two overlapping and interlocking units, each of which has itself two detection layers. These layers are resistive plates separated by 2 mm , with the gap between plates filled with a gas mixture ionised by the passage of a muon. An electric field between the plates allows a signal to be read out from copper strips on the face of the plates. The RPCs have response time of $15\text{--}25\text{ ns}$ and timing resolution of 1.5 ns . In addition to this triggering function, the RPCs are used to improve the non-bending plane measurements which the MDTs cannot make.

Thin Gap Chambers

The TGCs form the end-cap ($1.05 < |\eta| < 2.4$) section of the fast trigger chambers in the MS. The TGCs have the same response time as the RPCs and a time resolution of 4 ns . Although this is not as high a time resolution, it is sufficient to allow the MS to perform bunch-crossing identification. Similarly to the RPCs, measurements in the TGCs can also be used to complement measurements in the MDTs. The TGCs are multi-wire proportional chambers with the distance between anode wires larger than the distance from anode to cathode. It is this that allows the TGCs to operate so quickly.




Figure 2.11: The ATLAS magnet system, comprising barrel and end-cap components to the superconducting toroid, and the central solenoid. Image from Ref. [81].

2.2.4. Magnet system

Momentum measurements made by the tracking systems in ATLAS are enabled by the bending of particle tracks under a magnetic field. In ATLAS, this magnetic field is provided by a system of four superconducting magnets and extends over $12\,000\text{ m}^3$. The magnet system can be seen in Figure 2.11.

The ID benefits from a 2 T magnetic field provided by a solenoid magnet aligned along the beam axis. This solenoid surrounds the ID, is 2.5 m in diameter, 5.8 m in length, and only 10 cm thick so as to minimise the material presented to impede particles prior to their interaction with the calorimeters. It is cooled with liquid helium to a temperature of 4.5 K, at which point it is superconducting.

The MS uses a magnetic field generated by a superconducting toroid magnet system. The toroid is made up of two end-caps - one on either side of the detector - and a central barrel region which comprises eight octants. This toroid system extends 26 m in length and is 22 m in diameter. It provides a 0.5 T field in the central region and a 1 T field in the end-cap regions. The barrel toroid is cooled to 4.6 K.

2.2.5. Trigger system

With a bunch crossing rate of 40 MHz, it is impossible for ATLAS to read out and store every signal which is induced in the detector. The trigger system is designed to perform fast reconstruction in order to determine which events could potentially




Figure 2.12: The ATLAS TDAQ system as it was configured for Run 2 [84].

be of physics interest and should therefore be read out and stored for offline analysis and which events should be discarded. The ATLAS trigger and data acquisition system (TDAQ) [82] is responsible for the online processing of events along with their selection and storage for offline analysis. Its structure is shown in Figure 2.12. The trigger is separated into two stages [83, 84] - the Level 1 (L1) trigger [85] and the High Level Trigger (HLT) [86].

The L1 trigger is a hardware-based trigger which takes coarse information from the calorimeters and muon system in order to very quickly cut the rate from 40 MHz to around 100 kHz within $2.5 \mu\text{s}$. It does this by considering event-level quantities such as the total energy in the calorimeters or by considering topological criteria such as angular separations. The L1 trigger can be further split into sub-systems.

The L1 calorimeter trigger (L1Calo) [87] takes signals from the calorimeters and runs them through a digitisation and pre-processing procedure. These signals are

passed to the Cluster Processor (CP), which identifies electron, photon, and τ -lepton candidates, and the Jet/Energy-sum Processor (JEP), which identifies jet candidates.

The L1 Muon trigger (L1Muon) [88] uses information from the fast components of the MS - the RPCs and the TGCs. In order to improve the proportion of selected particles emanating from the interaction point, there are coincidence requirements between layers of the TGCs and the tile calorimeter.

The Central Trigger Processor (CTP) decides on the L1 trigger using information from L1Calo and L1Muon, as well as a few other sub-detector systems. For events selected by the L1 trigger, the Front End (FE) electronics read out the data from the various sub-detector systems and pass it on to the Read-Out Drivers (RODs) and Read-Out System (ROS). From here, it is able to be passed to the HLT when necessary, along with Regions-Of-Interest (ROIs) also identified by the L1 trigger.

The HLT is a software-based trigger which makes initial rejections using fast trigger algorithms before making the final cuts with a more precise reconstruction process to reduce the rate to around 1 kHz within 200 ms. The reconstruction algorithms run on 40 000 Processing Units (PUs). These algorithms will typically take data fragments from ROIs then apply criteria based on some trigger condition to the reconstructed quantities to make a decision. An event accepted by the HLT is sent to CERN's Tier-0 [89] facility for storage and subsequent offline analysis.

2.2.6. Luminosity detectors

In relating the interaction rate of a process to its cross-section, the luminosity is a crucial parameter for ATLAS to measure. It is indispensable for cross-section measurements and represents a key systematic for many analyses. Real-time feedback from ATLAS to the LHC on the luminosity being delivered is also necessary for the optimisation of beam conditions. The ATLAS trigger system also uses the luminosity as an important input for calculating pre-scalings.

The luminosity delivered to ATLAS is measured by several systems. There are two dedicated luminosity detectors used by ATLAS - the Luminosity Cherenkov Integrating Detector (LUCID-2) and the Beam Conditions Monitor (BCM) - as well as complementary measurements provided by parts of the tracking and calorimetry systems [90].

Of these, LUCID-2 [91] is the primary provider of integrated luminosity measurements as well as online monitoring of the instantaneous luminosity. It is a Cherenkov detector designed to detect inelastic pp scattering, with stations 17 m either side of the interaction point at central ATLAS.

Each LUCID-2 detector is formed of sixteen photomultiplier tubes (PMTs) surrounding the beam pipe in groups of four and placed in the Target Absorber Secondaries (TAS) shielding [92]. There are an additional four PMTs are placed on top of the shielding, connected to quartz fibres which are placed between the four groups of PMTs closer to the beam pipe. All PMTs are shielded from stray magnetic fields. It is the quartz window of the PMTs themselves that act as the Cherenkov radiator which produce the photons which are amplified and measured by the PMT. The signal produced in the PMT is then proportional to the number of incident charged particles. In order to fulfil its goal, amongst other criteria, LUCID-2 must be sensitive to minimum-bias events, be able to resolve bunch-crossings given their 25 ns spacing, and be highly radiation hard.

Bunch crossings at the LHC result in multiple pp interactions - more than 30 on average during Run 2. LUCID-2 is able to measure the luminosity in two ways; first by hit counting, and second by total charge integration. The first method is able to relate the detection of an interaction to the luminosity through a Poisson distribution. The luminosity is related to pileup by

$$\mathcal{L} = \frac{\mu n_b f_r}{\sigma}, \quad (2.2)$$

where $\mathcal{L}$ is the instantaneous luminosity, μ is the mean number of interactions per bunch crossing — pileup — and n_b and f_r are the number of bunches and the LHC revolution frequency. It is assumed that the mean number of interactions follows a Poisson distribution, which states that the probability of n events occurring is given by

$$P(n; \mu) = \frac{\mu^n e^{-\mu}}{n!}, \quad (2.3)$$

where μ is the mean of the distribution, the pileup. Evaluating for the probability of zero events gives

$$p_0 = e^{-\mu}, \quad (2.4)$$

which can be solved to give an expression for the pileup,

$$\mu = -\ln(p_0). \quad (2.5)$$

The second method relies on the proportionality of the total charge per bunch crossing to the luminosity of that bunch crossing, which is independent of any statistical model.

Part II.

Measurement of Jet Substructure Observables in Boosted $t\bar{t}$ Events

Chapter 3.

Analysis Overview and Tools

The analysis which forms the main focus of this thesis makes use of several experimental methods, mainly concerned with the reconstruction and treatment of the objects called jets. This chapter begins by covering how jets are constructed from particles emanating from LHC collisions, how their origins may be identified, and how their structure may be measured. Following this, an introduction to the analysis of jet substructure observables in boosted $t\bar{t}$ events is given, with an overview of the analysis aims, strategy, and methods.

3.1. Jets

As dictated by confinement in QCD, quarks and gluons are not observed in the detector as isolated particles. Instead, what is observed in the detector are collimated sprays of colour-neutral particles which result from the showering and hadronisation of the original parton. These are reconstructed as jets [93, 94] and are ubiquitous at high energy colliders such as the LHC. Once reconstructed, jets are then used as the experimental proxies for the initiating parton in a given event.

At high Lorentz boost, the decay products of a massive particle, such as a top quark, become collimated. In this case, the jets corresponding to each of the decay products are close to one another and overlap. As shown in Figure 3.1, these jets may then be reconstructed as a single large jet corresponding to the original parton, in this case the top quark.




Figure 3.1: The lepton + jets decay channel for a boosted $t\bar{t}$ event. The hadronic decay products are collimated into a single large jet containing three distinct prongs.

For a jet to serve as a proxy for the initiating parton and thus offer a means to study the underlying dynamics, there must be as close a correspondence between a jet reconstructed at detector-level and one reconstructed for the same system at parton-level as possible. Coming up with a jet definition capable of this is not trivial and has been the subject of many decades of development. The parton-level represents an event frozen after the parton shower, before hadronisation, so includes all particles from the shower. The detector-level not only includes the effects of hadronisation but the interaction of the resultant particles with the detector such that it is defined in terms of tracks, calorimeter clusters, and other detector objects, of which more in Chapter 4. Additionally one may define the particle-level, which leads on from the parton-level and includes hadrons and their decay products, but prior to any interaction with the detector. The particle-level then consists of stable particles, defined as those with $c\tau > 10$ mm. The question of which particles belong to the jet is complicated by interference from pileup, underlying event, and other parts of the hard scatter.

3.1.1. Clustering

A jet definition comprises a clustering algorithm, which dictates how and which particles should be clustered into the jet, along with a set of associated parameters. In addition to the complications described above, there are a host of other characteristics

which should be exhibited by a good jet definition. The “Snowmass Accord” of 1990 laid out the criteria which should be met by such a jet definition [95]:

1. The jet definition should be easily implemented experimentally [96].
2. The jet definition should be simple to implement in theoretical calculations.
3. It should be defined at any order of perturbation theory.
4. It should not yield infinite cross sections at any order of perturbation theory.
5. It should yield a cross section which is relatively insensitive to hadronisation.

One of the most important parts of these requirements — linked to point 5 — is that the collection of jets in a given event produced by a jet definition should be infrared and collinear (IRC) safe. This means that the jet collection returned should be the same irrespective of the addition of some soft (IR) emission, or a collinear splitting. Having a jet definition which is IRC safe is key to the ability to make comparisons between experimental results and theory calculations. Soft emissions and collinear splittings are a necessary part of QCD events and their occurrence is difficult to predict, so a jet definition which is sensitive to such contributions will not be able to provide a clear mapping from the jet observed experimentally to that in the theory calculations.

Point 4 of the Snowmass conditions may also be violated by an IRC-unsafe algorithm. In the QCD calculations, infrared emissions and collinear splittings produce divergent tree-level matrix elements. Ordinarily, these are counter-balanced by divergent loop matrix elements which have the opposite sign. For an IRC unsafe jet algorithm, this may be disrupted such that there is no cancellation and infinite cross-section results. This possible behaviour of infrared and collinear unsafe algorithms is shown in Figures 3.2 and 3.3. If a jet algorithm is IRC unsafe and returns a different number of jets depending on infrared emission or collinear splitting, then the resulting collections for the event contribute to different cross-sections, which both become divergent.

Given these requirements for how a jet algorithm should behave, the question is then how a clustering algorithm should proceed in order to satisfy these requirements. The most obvious way to define a jet would be to simply define a cone with some radius and sum the momenta of the particles within that cone in order to calculate the kinematics of the jet and, by proxy, the initiating parton. Indeed, the first jet algorithm [98] was such a *cone algorithm*, developed for use in e^+e^- collisions, which




Figure 3.2: This diagram from Ref. [97] shows the effect of an infrared-unsafe algorithm. In a) and b), the result is the same. Two jets, the same two jets, are returned. In c), however, the soft radiation introduced has skewed the algorithm such that the two jets are merged into one and a different jet collection is returned.




Figure 3.3: This diagram from Ref. [97] shows the effect of a collinear-unsafe algorithm. a) and b), utilising a collinear-safe jet algorithm, return the same jet collection irrespective of the collinear splitting of the central parton. For c) and d), however, this splitting has caused the algorithm to return a modified jet collection. This algorithm is collinear-unsafe.

was used to classify events containing two jets. This description of the event was said to be satisfied if $1 - \epsilon$ of the total energy in the event was contained within two cones which each had a half opening angle of δ , with both ϵ and δ being a choice free to the experimenter. Jet algorithms developed from this basic cone idea for many years.

When confronted with a collection of particles from which jets are to be clustered, the immediate problem is where the cones representing the jets should be placed. Modern cone algorithms take an iterative approach, beginning from some seed particle i — usually the particle with greatest p_T . These are known as iterative cone (IC) algorithms. These proceed by first summing all particles j inside a cone of radius R , i.e. $\Delta R_{ij} < R$. The direction of the resulting cone is then used as the seed for the next iteration, with the process continuing until the direction of the resulting cone is stable, giving the final jet.

The development of these IC algorithms has itself been an iterative process. IC progressive removal (IC-PR) algorithms [99] arose from the need to deal with the case of overlapping cones, when particles are shared between two or more possible cones. IC-PR algorithms start with the hardest particle in the event as the seed, find the first jet corresponding to that seed, and then remove from the event all particles within that jet and repeat the process until there are no particles remaining. IC-PR clustering algorithms, as with many of the cone algorithms, are not collinear safe. Should the hardest particle in the event undergo a splitting, then it is possible that some other, softer particle in the event becomes the hardest and is used as the seed rather than the original hardest particle. This different seeding of the cone then returns a different collection of jets for that event. This scenario is what can be seen in Figure 3.3.

Two further approaches are split-merge and split-drop (IC-SM and IC-SD) algorithms [100]. First all stable cones from all seed particles are found, labelled as *protojets*. The hardest protojet is labelled a . The next-hardest protojet overlapping with a is labelled b . If there are no overlapping protojets, a is removed from the list of protojets and added to the final list of jets. Next, the total p_T shared between a and b is calculated. If the shared p_T is above some threshold (usually around 0.75) f then a and b are merged into one protojet. Otherwise, a and b are ‘split’, assigning particles to the protojet they are closest to. This is then repeated until all protojets are gone. This is the IC-SM algorithm. Under the IC-SD algorithm, the non-shared particles contained within the softer jet are simply dropped altogether. Both IC-SM and IC-SD are examples of clustering algorithms which are infrared unsafe, in that a new soft emission can provide a new seed which alters the collection of cones returned, which

subsequently is converted to the final set of jets output. This scenario is what is demonstrated in Figure 3.2, where the central emission of a gluon in c) acts as a seed for a cone containing both partons coming from the W decay. The cones corresponding to these partons are then merged with that seeded from the soft emission and one jet is returned in this region, rather than two.

The infrared unsafe behaviour of the IC-SM algorithm can be labelled IR_{2+1} . That is, the IR safety is manifest for neighbourhoods of two hard particles plus a soft one. The IC_{mp} -SM algorithm (a proposed fix known as the mid-point fix [101]) is IR_{3+1} . The IC-PR algorithm is Coll_{3+1} .

Throughout the running of the Tevatron experiments, many attempts were made to patch various cone algorithms to solve the IRC safety issues which plagued them, eventually leading to the Seedless Infrared-Safe Cone (SISCONE) algorithm [102].

The second class of jet clustering algorithms are the *sequential recombination* algorithms. These are generally simpler than the cone algorithms, and may also be used to represent the parton branching by providing a cluster sequence. This is a feature being exploited ever more in physics analyses using the substructure of jets at the LHC [103, 104] and in tagging applications, both of which are introduced below. Modern sequential recombination algorithms have now largely supplanted cone algorithms as the most widely used clustering technique.

Rather than collecting particles within a cone of some radius, sequential recombination algorithms use a distance measure to combine pairs of particles in order to construct a jet. The k_t family of sequential recombination algorithms are the favoured choice for modern experiments like ATLAS.

As with the cone algorithms, the sequential combination algorithms have their roots in e^+e^- experiments. The original form of the k_t algorithm [105] was a modification of the first sequential recombination algorithm, the JADE algorithm [106, 107]. In applying this algorithm at hadron colliders, a couple of further iterations produced a longitudinally invariant distance measure [108], and then a final modification which used the same distance measure except for the addition of a radius parameter R [109] analogous to the jet radius in the cone algorithms.

This final form of the k_t algorithm produced two distance measures, which can be generalised as

$$d_{ij} = \min(p_{T,i}^{2p}, p_{T,j}^{2p}) \frac{\Delta R_{ij}^2}{R^2}, \quad (3.1a)$$

$$d_{iB} = p_{T,i}^{2p}, \quad (3.1b)$$

with d_{ij} the distance measure, R the jet distance parameter, d_{iB} the *particle-beam* distance, and p a parameter. ΔR_{ij} is the distance in $\eta - \phi$ space between constituents i and j . This d_{iB} was introduced in earlier versions of the algorithm along with the concept of a *beam jet*, which particles may be recombined with rather than other protojets in the event. In the modern version of the algorithm presented in Equation 3.1, d_{iB} acts in a slightly different way and the concept of the particle beam is dropped.

Using these measures, the algorithm then proceeds in the following way:

1. Calculate all d_{ij} and d_{iB} in the event.
2. For each pair of particles, find the minimum of d_{ij} and d_{iB} .
3. If it is d_{ij} , recombine particles i and j into a new *protojet* and return to step 1.
4. Otherwise, i is taken to be a final state jet and removed from the list of protojets.
5. When no particles remain, stop.

The *inclusive* k_t algorithm is defined by taking $p = 1$ in Equation 3.1, and in fact this case came first, before the generalisation presented here. Here, for any two particles which are separated by less than R , d_{ij} will be smaller than d_{iB} and so those particles will be combined into a protojet. Otherwise, the particle in question will form its own jet. Of course, one consequence of this is that even a very soft particle which is sufficiently well separated from others will form its own jet. For this reason, as is done in the analysis presented in this thesis, a minimum p_T threshold should be applied to an event's final collection of jets.

The choice of R will determine the size of the jets produced by the clustering algorithm. In ATLAS, there are generally two classes of jets distinguished by this choice: small- R jets have $R = 0.4$, and large- R jets have $R = 1.0$. A larger R will generally produce a jet better able to capture fragmentation and hadronisation, but will be more susceptible to pileup and underlying event than a jet clustered with a

smaller R . Jets constructed with a lower value for R are also generally better able to resolve multiple hard partons.

Taking instead $p = 0$ in Equation 3.1 defines the Cambridge/Aachen (C/A) algorithm [110], so named for its roots in the Cambridge algorithm [111]. This is similar to the inclusive k_t algorithm but performs clustering independent of the p_T of the particles.

Taking $p = -1$ produces the anti- k_t algorithm [97]. Whereas the minimum function in the distance measure means that the k_t algorithm clusters soft particles first, the anti- k_t algorithm clusters hard particles first. The result of this behaviour is a collection of IRC safe jets centred on a hard core and approximately conical, as with the cone algorithms. The circular jets produced enable simplified triggering algorithms, as well as easing the process of energy correction for the jets to account for additional radiation in an event, and so are favoured by experimentalists as compared to other more jagged jet shapes. This can be seen in Figure 3.4, where a comparison is made to jets produced by SISCONE and the other variants of the k_t algorithm. These characteristics have made the anti- k_t clustering method the de facto choice for jet reconstruction at the LHC. Almost all ATLAS analyses use the anti- k_t algorithm, and the one presented in this thesis is no exception.

Where the anti- k_t algorithm falls short as compared to the inclusive k_t and C/A algorithms is in the utility of the clustering sequence of the jets produced. The clustering sequences produced by the inclusive k_t and C/A algorithms correspond to the p_T and angular-ordered parton showers, respectively, and so are commonly preferred when the task at hand is a study of the internal structure of jets; their substructure. Even with this application, however, it is common to first perform jet finding and reconstruction using the anti- k_t algorithm, before applying an alternative clustering method to the ensemble of particles contained within those jets. Indeed, that is done in many cases in the calculation of observables measured in the analysis presented here.

One aspect of the clustering which has not been discussed so far is what is meant by *particles*. At the parton- and particle-levels, these are as described at the beginning of this chapter, or possibly the protojets formed by recombination of particles. At the detector-level, there are a wealth of possible choices for the input. The most common choice for many years was to use clusters formed in the calorimeters as input to the clustering algorithms. It is also possible, however, to use tracks constructed in the detector's tracking system as input. More detail on these objects is given in Chapter 4.




Figure 3.4: This diagram from Ref. [94] provides a comparison between four jet algorithms — the three k_t sequential recombination algorithms, and the SIScone algorithm.

Another option is to use objects which are themselves constructed from combinations of objects from different parts of the detector. One such set of objects — used to reconstruct jets in the analysis presented in this thesis — are *particle flow* objects. A comparison of jet collections reconstructed from different objects in the context of observables measured in this analysis is presented in Section 4.4.

3.1.2. Particle Flow

The particle flow concept is neither new nor exclusive to ATLAS, but the description provided here will naturally focus on the approach used in ATLAS [112]. Whereas jets are traditionally clustered from either calorimeter clusters or from tracks, the particle flow approach aims to make use of both in order to benefit from the advantages of both detector systems. This provides a number of benefits.

For low energy particles, the tracker has superior momentum resolution to the energy resolution of the calorimeter. Beyond improving the resolution on some of these particles, this also extends the sensitivity of the detector to soft particles, for which those with a p_T below 400 MeV are below the noise thresholds which are used when seeding calorimeter clusters (more detail on this is given in Section 4.2). Such low- p_T particles may also be missed in jet clustering which uses only calorimeter clusters as the strong magnetic field within the detector can deflect their path such that they are outwith the cone reconstructed in the calorimeter by the time they reach those layers of the detector. When tracking information is also used, these particles may be clustered into the jet. In contrast, the calorimeter's superior energy resolution may be called upon for high energy particles. As well as the superior low- p_T momentum resolution provided by the tracker, a charged particle reconstructed in the tracker will have better angular resolution than it would in the calorimeter. The advantage of using particle flow objects over calorimeter clusters in terms of momentum and angular resolution can be seen in Figure 3.5.

One of the other main advantages of the use of particle flow objects is a robustness to pileup. This is possible because tracks allow for the identification of vertices and so particles which are associated to pileup vertices may be rejected and their impact on the resultant jets mitigated. This may be seen in Figure 3.6.

Although the inclusion of tracking information provides many advantages, the tracker cannot help either with neutral particles or particles falling outwith the accep-




Figure 3.5: Resolution comparison between jets clustered from calorimeter clusters and particle flow jets, simulated in MC for dijet events. Plots from Ref. [112].




Figure 3.6: Effects of pileup compared for particle flow jets and jets clustered from calorimeter clusters. Plots from Ref. [112].

tance of the ID, $|\eta| < 2.5$. Here, only calorimeter topoclusters (topological clusters, of which more in Chapter 4) can be used.

The challenge with this approach is in the avoidance of double-counting. If the energy of a particle is to be assigned by a track measurement, then the corresponding energy deposit in the calorimeter must be subtracted. To this end, the particle flow algorithm employs a cell-based subtraction system. The final output is a collection of objects — tracks, topoclusters modified through the subtraction scheme, and unmodified clusters. This collection of particle flow objects may then be used to cluster jets. The process for producing this set of objects is as follows.

To begin with, high-quality tracks are selected. This is done by application of the *tight* set of criteria within ATLAS, which require that each track has at least nine hits in the silicon layers of the tracking system and that there are no holes in the pixel layers. Tracks with $p_T > 40 \text{ GeV}$, or which are matched to electron or muon candidates, are vetoed.

The selected tracks are then matched to a single topocluster. Each track has its position extrapolated to the second layer of the ECAL and topoclusters are then ranked in increasing $\Delta R'$, defined as

$$\Delta R' = \sqrt{\left(\frac{\Delta\phi}{\sigma_\phi}\right)^2 + \left(\frac{\Delta\eta}{\sigma_\eta}\right)^2}, \quad (3.2)$$

where $\Delta\phi$ and $\Delta\eta$ represent the distance between the extrapolated track position and the barycentre of the topocluster, and σ_ϕ and σ_η are the angular widths of the topocluster. Topoclusters which may be matched to a track are required to have

$$\frac{E^{\text{clus}}}{p_{\text{trk}}} > 0.1, \quad (3.3)$$

with E^{clus} the energy of the cluster and p_{trk} the momentum of the track. Of the remaining topoclusters, the matched one is taken as that with the smallest $\Delta R'$. If no topocluster lies within a cone of radius $\Delta R' = 1.64$, then the particle is assumed not to have seeded a topocluster and the track is retained with no matched topocluster and no subsequent energy subtraction is carried out.

The next step is to calculate the energy expected to be deposited by the particle responsible for each track by consideration of the matched topocluster. This is given

by

$$\langle E_{\text{dep}} \rangle = p^{\text{trk}} \left\langle \frac{E_{\text{ref}}^{\text{clus}}}{p_{\text{ref}}^{\text{trk}}} \right\rangle. \quad (3.4)$$

Again p^{trk} is the track momentum, and E^{clus} the energy of the cluster. The expectation value on the right-hand side has subscript *ref* as this value is determined in simulation by summing the energy contained in a cone of radius $R = 0.4$ around a track from a single particle, with no pileup, extrapolated to the second layer of the ECAL.

It may be that the particle under consideration deposited energy in two or more topoclusters. In this case, it can be necessary to add other topoclusters to the system to fully account for the particle's shower. This stage in the algorithm is known as split-shower recovery. A discriminant is used to determine whether or not to run this procedure. This is given by

$$S(E^{\text{clus}}) = \frac{E^{\text{clus}} - \langle E_{\text{dep}} \rangle}{\sigma(E_{\text{dep}})}, \quad (3.5)$$

with $\sigma(E_{\text{dep}})$ the spread of the expected energy deposited by the particle. Given the separation in the distribution of this discriminant for the cases where above 90% and above 70% of a particle's energy is deposited in the matched topocluster, the procedure is run in all cases for which $S(E^{\text{clus}}) < -1$. When this criterion is satisfied, any topoclusters within $\Delta R = 0.2$ of the track's position when extrapolated to the second ECAL layer are marked as matched to that track.

A cell-by-cell subtraction is then carried out to remove the particle's expected energy from the set of matched topoclusters. The simplest case is when the expected energy deposited by the track exceeds that in the matched topoclusters, in which case they are removed. Otherwise, the subtraction process begins from the layer of highest energy density (LHED), which is the layer where the gradient in energy density is at its maximum. Beginning from the track's position when extrapolated to the LHED, the most probable energy density profile is calculated for each layer and rings formed around that position in $\eta - \phi$ space for each layer. These rings are required to contain at least one cell and are concentric with equal spacing for each layer. Across all layers, rings are then ranked in descending energy density and the subtraction procedure commenced from the start of the ranked set, which is the innermost LHED ring by definition. A running sum of the subtracted energy is kept for comparison with $\langle E_{\text{dep}} \rangle$.

As long as the energy in the ring under consideration leaves this sum below $\langle E_{\text{dep}} \rangle$, it is removed. When the point is reached that the addition of the energy contained in a ring would push this sum over $\langle E_{\text{dep}} \rangle$, that ring has its energy scaled until the sum matches $\langle E_{\text{dep}} \rangle$, and the cell-subtraction procedure halted. The cell-subtraction procedure is shown schematically in Figure 3.7.

Should any energy remain in the matched topoclusters, it may be removed if adjudged to be consistent with shower fluctuations. The criterion here is that if this energy is within $1.5\sigma(E_{\text{dep}})$, then it is said to have arisen from a single particle and is removed.

The full particle flow algorithm just described is shown schematically in Figure 3.8. The resulting set of tracks, modified topoclusters, and unmodified topoclusters should then provide a set of objects capable of describing the event with improved resolution and without double-counting.

3.1.3. Reclustering

For a top quark produced with high momentum, its decay products may be reconstructed as a single jet of radius $R \approx \frac{p_T}{2m_t}$ [93]. For a large- R jet of radius $R = 1.0$, this means a minimum p_T of 350 GeV. For the analysis presented here, *reclustered* large- R jets (RC jets) are used [113]. These are clustered by using small- R particle flow jets as the inputs to the anti- k_t algorithm.

Because the particle flow jets from which the RC jets are constructed are already calibrated and trimmed, these steps are not required again for the resultant RC jet. This means that the systematic uncertainties associated with these large- R jets are simply inherited from the small- R particle flow jets and no additional prescription is needed. The two jet collections in the event then share the same set of uncertainties, and those associated with the RC jets are considerably smaller than for large- R jets clustered directly from topoclusters, as shown in Figure 3.9.

As with the approach taken for large- R jets in ATLAS clustered from topoclusters, contamination from pileup in RC jets is mitigated by removing from the jet any of the constituent small- R jets which make up less than 5% of the jet's total p_T [114, 115].

The use of RC jets also allows for more straightforward b -tagging, which is performed on small- R jets. An RC jet can be marked as b -tagged simply if one of the




Figure 3.7: Idealised view of the particle flow cell subtraction procedure. The extrapolated position of a track from a π^+ is shown by the blue star for two ECAL layers, with EMB2 representing the LHED. Calorimeter cells are red for the cluster belonging to the π^+ and green for those belonging to a π^0 . Starting from the innermost ring of the LHED in (a), rings are removed from those with highest energy density to those with lowest, as long as the removed energy remains less than that expected to be deposited. Once the procedure reaches (g), this final ring contains enough energy to take the total removed over the expected deposited energy, so is scaled until the energy removed matches that expected to have been deposited. Image from Ref. [112]




Figure 3.8: This diagram from Ref. [112] shows how the particle flow algorithm operates within ATLAS, outputting tracks, energy-subtracted topoclusters, and unmodified topoclusters.




Figure 3.9: Scale uncertainties as a function of a jet’s transverse momentum. The traditional large- R jets, clustered from calorimeter clusters, are shown in red. The blue represents large- R jets reclustered from small- R jets. Plots from Ref. [113].

constituent small- R jets has been b -tagged. For large- R jets clustered directly from topoclusters, this is not possible. Instead, a matching procedure is required to pair a large- R jet to a b -tagged small- R jet.

A complication presented by the use of RC jets in this analysis comes at the point of calculating substructure observables. Generally — and for the observables presented here, defined in Section 3.1.5 — substructure is calculated by passing the constituents from which the jet was clustered to the appropriate algorithm for that observable. Because the RC jets used in this analysis are clustered from objects which have themselves already been clustered, they generally comprise only a small number of constituents. In many cases, it is not possible to calculate substructure observables from such a small number of constituents, and even if it were, this would provide a far less granular set of inputs to the calculation than would the original jet constituents.

In order to calculate the substructure of the RC jets, a new ensemble of particles is created by merging each set of constituents associated to the small- R jets used to cluster the RC jet. This set of constituents is then clustered using the anti- k_t algorithm to produce a new jet object with a complete set of constituents and cluster history which may be passed to the substructure calculation. There is evidence that reclustered large- R jets provide substructure-based discrimination at least as powerful as traditional large- R jets clustered directly from constituents [114].




Figure 3.10: This diagram from Ref. [116] demonstrates the various stages in the JES calibration within ATLAS.

3.1.4. Calibration

Various aspects of a jet must be corrected before use in the measurement, in order to account for the effects of interaction with the detector, pileup, and possible offsets. A calibration procedure is applied to bring the jets' energy scale at detector-level closer to that at particle-level and to improve agreement between the data and simulation.

Jet Energy Scale (JES) The JES calibration procedure [116] is shown in Figure 3.10, with each of the steps having some effect on the four-momentum of the jet, with the aim of matching the energy scale of the jets to that of the jets simulated at particle-level. The first step in the calibration is designed to account for pileup. Pileup can arise either in-time, the consequence of multiple pp collisions in a bunch-crossing, or out-of-time, where particles may interfere from neighbouring bunch-crossings. Pileup is distributed approximately uniformly throughout an event and has the effect of increasing the energy of a jet as the pileup particles are clustered into it. Owing to the uniformity of the pileup, the correction is implemented through a p_T density-based scaling. The p_T density, ρ , is given as the ratio of the p_T of a jet to its area A in $\eta - \phi$ space. The median p_T density is estimated using jets with $|\eta| < 2.0$. The ratio of the jet's p_T minus the median p_T density to the original jet p_T is then applied as a correction factor to the jet's p_T . In this way, the greater ρ , the more the jet's p_T is scaled down.

Because this first pileup correction is derived from the central regions of the calorimeter, where occupancy is generally low, a second correction is applied to

account for residual pileup dependencies owing to the differing sensitivity in forward regions or high-occupancy regions as would be expected inside a hard jet. These dependencies of the pileup effects are found to be linear in the number of primary vertices in an event, N_{PV} , and the number of interactions in a bunch crossing, μ . Including the ρ -based correction, the total correction to the jet p_T is given by

$$p_T^{\text{corr}} = p_T^{\text{reco}} - \rho \times A - \alpha \times (N_{PV} - 1) - \beta \times \mu, \quad (3.6)$$

where α and β are coefficients derived from linear fits in p_T and η , and p_T^{reco} is the p_T of the jet as it is initially reconstructed in the detector. As compared to jets clustered using only topological clusters, the dependence of particle flow jets on in-time pileup — seen in the dependence on N_{PV} — is reduced significantly for $|\eta| < 2.5$ owing to the use of tracks in the reconstruction.

After pileup corrections, calibration based on dijet events in MC is applied to match the energy scale of the reconstructed jets to that at truth-level. This calibration targets effects on both the jet's energy and its direction. These may arise from the non-compensating nature of the calorimeters, meaning differing responses to the EM and non-EM components of hadronic showers, or from energy that is lost to *dead material* in the detector. That is, parts of the detector which are not sensitive to incident particles. Biases to the jet's direction may arise in jets reconstructed in combinations of calorimeter technologies which may have differing response and granularity. The calibration is applied in two stages, using pileup corrected jets which are matched to truth jets within $\Delta R = 0.3$, with all jets isolated. The first stage applies a calibration factor which is defined as the inverse of the average jet energy response, $\mathcal{R}$, which is extracted as the mean of a Gaussian fit of the ratio of the reconstructed energy of the jet and the energy of the matched jet at truth-level. This distribution is binned in the energy of the truth-level jet and η_{det} . The second stage in the calibration applies another factor designed to account for effects on the reconstructed jet η . This is derived from $\eta^{\text{reco}} - \eta^{\text{truth}}$, parameterised in E^{truth} and η_{det} . The calibration up to this stage is known as MCJES.

The next step is the *global sequential calibration* (GSC). Following the MCJES calibration, the jet response may still vary depending on its composition, how it was initiated, or its location in the detector. The GSC corrections aim to improve the jet resolution — given by the width of a Gaussian fit of the jet p_T response distribution — without changing the response itself. The same matching requirements as for the MCJES

calibration are applied. Corrections are applied for each of six observables by the inverse of the jet response, parameterised in p_T^{true} and $|\eta_{\text{det}}|$, and applied sequentially. These observables are:

1. f_{charged} — the fraction of the jet's p_T coming from constituents which are charged particles.
2. f_{Tile0} — the fraction of the jet's energy which was recorded in the first layer of the Tile calorimeter.
3. f_{LAr3} — the fraction of the jet's energy which was recorded in the third layer of the LAr calorimeter.
4. n_{trk} — the number of tracks associated to the jet.
5. w_{trk} — the track width; average distance in $\eta - \phi$ space between the axis of the jet and the tracks associated to the jet, weighted by p_T .
6. n_{segments} — punchthrough correction; the number of muon track segments associated to the jet.

Finally, a correction derived from data — *in-situ* — is used to correct for residual differences between the data and MC. These differences may arise from either simulation of the detector or modelling of physics processes. It may be that detector materials, geometry, response, etc. differ slightly between the simulated and true detector. Similarly, simulation of the hard scatter at bunch crossing, the underlying event, formation of the jets, pileup, and interactions with the detector material cannot be perfect. The *in-situ* calibration is based on the double ratio

$$c = \frac{\mathcal{R}_{\text{in-situ}}^{\text{data}}}{\mathcal{R}_{\text{in-situ}}^{\text{MC}}}, \quad (3.7)$$

where $\mathcal{R}_{\text{in-situ}}^{\text{data}}$ and $\mathcal{R}_{\text{in-situ}}^{\text{MC}}$ are the jet response in the data and MC. This is calculated using a reference object which has been well calibrated in order to balance the p_T of a jet, and is defined as the average ratio of the p_T of the jet to the p_T of the reference object being used. These response terms are calculated in bins of the reference object p_T before c is transformed into a function of the jet p_T and applied as the calibration to the jet. In fact several different processes are used in deriving the corrections in order that the calibration is as generally applicable as possible. In this way, the calibration can be applied to jets up into the multi-TeV range in p_T and across the $|\eta|$ spectrum.

Jet Energy Resolution (JER) In addition to the JES, knowledge of the JER is required. The JER accounts for the finite precision of the detector and is important for precision measurements of SM processes such as top pair production, as well as in the reconstruction of missing momentum, E_T^{miss} . The relative JER as a function of the jet p_T can be parameterised as [116]:

$$\frac{\sigma(p_T)}{p_T} = \frac{N}{p_T} \oplus \frac{S}{\sqrt{p_T}} \oplus C, \quad (3.8)$$

with N , S , and C noise, stochastic, and constant terms, respectively. The noise term accounts for electronic noise from the readout front-ends, as well as from pileup. Its $\frac{1}{p_T}$ scaling means that it is most significant at lower values of the jet p_T . The stochastic term accounts for statistical fluctuations in energy deposition, and is most significant up to a few hundred GeV, owing to the $\frac{1}{\sqrt{p_T}}$ scaling. Finally, the constant term covers several effects which all apply as some constant fraction of the jet's p_T . Some of these are similar to those covered in the GSC stage of the JES calibration, such as energy lost to dead material and differing response in different regions of the calorimeters. The constant term is most significant for jets with the highest p_T , above 400 GeV.

Measuring the JER requires that jets' momentum must be measured with high precision. Dijet events are used with the same *in-situ* method as for the JER calibration, along with an additional step for the estimation of the noise term in Equation 3.8. At the end of the calibration procedure, an additional smearing is applied to better match data and simulation. This smearing is applied in regions of the jet p_T for which the JER in the simulation is better than in the data, and is done such that the average resolution is matched between the two. In regions where the resolution in the data is better than in the simulation, no smearing is applied.

3.1.5. Substructure

Thus far, much has been said on what jets are, and how they may be reconstructed and treated experimentally. An emerging sub-field within jet physics which has come to the fore since the beginning of the LHC-era is that of jet *substructure* [93,94,117–121], which aims to characterise the radiation within a jet, not just its kinematics as a complete object.

Through parton showering and hadronisation, a given jet will form a signature related to its origin. A jet from a gluon splitting, $g \rightarrow gg$, for example, is expected to have higher multiplicity and greater extent in $\eta - \phi$ than a quark jet with $q \rightarrow qg$. This is because the former splitting has probability proportional to colour factor $C_A = 3$, whereas the colour factor for the latter splitting is $C_F = 4/3$. This information may be used to distinguish between the two types of jet. This same principle — using the angular and momentum distributions of particles within a jet in order to discriminate against different processes — finds a variety of applications at the LHC. Observables based on these characteristics of jets are used in substructure *taggers*, discussed in the next section, in order to identify which process produced a given final state, where several processes may produce the same signature. Several such observables are considered and measured in the analysis presented here.

The study of jet substructure is not restricted to use in analysis tools such as taggers, however. The use of substructure techniques allows for precision measurement of the SM, with the ability to contrast SM predictions with data over a wide range of energy scales, as well as serving as a probe of the strong force [104, 122]. Although mismodelling of substructure observables may not exactly prescribe the underlying cause of the discrepancy, differential measurements of the disagreement can be used to tune simulation parameters such that the MC predictions better describe the data [123].

Part of the desire for IRC-safe observables — as described in Section 3.1.1 — is that they have calculable cross-sections in perturbative QCD. It has also been shown that there are a broader set of observables which are not IRC-safe but can still be computed in perturbative QCD. Such observables are said to be *Sudakov-safe* [124, 125].

While there has been discussion on how a jet should be clustered and what its constituent *particles* should be, this must be considered anew in the context of substructure observables. It is not uncommon for the algorithms used to construct these observables to take the ensemble of particles used in clustering the jet and perform the clustering again with a different algorithm. Whilst anti- k_t jets are almost universally favoured in event reconstruction, it is common for substructure algorithms to then cluster those jets again using either the k_t or C/A algorithm, which cluster soft particles first and by angular ordering, respectively. The clustering sequence for these algorithms is then related to the QCD branching.

In the analysis presented here, tracks which have been *ghost-associated* to jets are used to construct the substructure observables. Here, track four-vectors are assigned

negligible magnitude and included in the jet reconstruction process. They are then considered to be associated to the jet they are clustered into, but have no impact on the kinematics of that jet. As compared to calorimeter clusters, tracks have improved angular resolution which is propagated through to the observables themselves. This is also useful in providing access into the collinear region which is sensitive to non-perturbative QCD effects.

Although they formed no part of the programme laid out in the proposals for the LHC experiments, substructure techniques are now an integral part of the physics capabilities of experiments like ATLAS, allowing for discrimination between processes, tests of theoretical calculations, heightened sensitivity to BSM physics, and the improvement of the simulations used to model these processes at the LHC.

Although for the analysis presented here, the substructure is being calculated using tracks, observables can be defined in general, regardless of the constituents being used. Defined below are several families of substructure observables which are considered in Chapter 6 for measurement.

Generalised angularities [126] At a hadron collider, the generalised angularities can be defined as

$$\lambda_{\beta}^{\kappa} = \sum_{i \in J} z_i^{\kappa} \left(\frac{\Delta R(i, \hat{n})}{R} \right)^{\beta}. \quad (3.9)$$

For the constituents of a jet J , z is the fraction of the jet's p_T carried by the i^{th} constituent. $\Delta R(i, \hat{n})$ is the distance in $\eta - \phi$ space between the i^{th} constituent and the jet axis, $\hat{n}$. R is the jet radius. κ and β are parameters chosen to weight the p_T and angular terms, respectively. Configurations with $\kappa = 1$ are infrared and collinear (IRC) safe, whilst configurations with $\kappa \neq 1$ are IRC unsafe but Sudakov safe.

The variants investigated here are:

1. λ_0^0 : the particle multiplicity.
2. λ_0^2 : the jet p_T dispersion.
3. $\lambda_{0.5}^1$: the Les Houches Angularity (LHA).
4. λ_1^1 : the jet width.

5. λ_2^1 : the jet thrust.

The jet p_T dispersion is highly correlated with the particle multiplicity and thus a scaled p_T dispersion is defined:

$$\lambda_0^{2*} = p_T^{d,*} = \sqrt{\left(p_T^d - \frac{1}{N}\right) \frac{N}{N-1}} \quad (3.10)$$

Here, p_T^d is the standard p_T dispersion and N is the particle multiplicity. In this way, $\lambda_0^{2*} \rightarrow 0$ for equally distributed constituents and $\lambda_0^{2*} \rightarrow 1$ when most of the momentum is carried by a single particle.

Energy Correlation Functions (ECF) [127, 128] The general formula for calculating the ECFs at a hadron collider is

$$\text{ECF}(N, \beta) = \sum_{i_1 < i_2 < \dots < i_N \in J} \left(\prod_{a=1}^N p_{T,i_a} \right) \left(\prod_{b=1}^{N-1} \prod_{c=b+1}^N \Delta R(i_b, i_c) \right)^\beta \quad (3.11)$$

with i always denoting a constituent of the jet, and N the number of prongs. The sum runs over all constituents in the jet. ΔR is the distance in $\eta - \phi$ space between pairs of constituents with combinations determined by the product terms. ECFs are calculated here for both $\beta = 1$ and $\beta = 0$, which weight the angular terms. The ECFs are only IRC safe for $\beta > 0$. It is expected that for a jet of N prongs, $\text{ECF}(N+1) \ll \text{ECF}(N)$.

A simple example can be used to intuit this behaviour from the formulation. Consider a two-pronged jet — perhaps originating from a hadronically decaying W boson — with only three constituents. If there are two constituents corresponding to the W decay products, these will be relatively hard. That then leaves one other constituent which will likely be soft in comparison and unordered with respect to the two-pronged structure of the jet. If one evaluates Equation 3.11 in this case, then for ECF2, one arrives at Equation 3.14b, and for ECF3, Equation 3.14c. ECF3 is sent to a small value by the soft p_T term in the product. ECF2, on the other hand, has two terms that are sent to small values by the soft p_T term, but one which is the product of the two hard terms. ECF2, therefore, assumes a greater value. This is of course an oversimplification, but the logic extends to more realistic cases with more constituents.

Dimensionless ratios of the ECFs can also be used to probe N -pronged substructure, as is also done below for the N -subjettiness observables. The C_N observable is given

by

$$C_{\mathcal{N}}^{(\beta)} = \frac{\text{ECF}(\mathcal{N} + 1, \beta) \text{ECF}(\mathcal{N} - 1, \beta)}{\text{ECF}(\mathcal{N}, \beta)^2}. \quad (3.12)$$

For a jet of N subjets, the leading order substructure comprises N hard prongs. Small values of $C_{\mathcal{N}}$ suggest that the higher-order radiation is soft or collinear with respect to the leading-order structure. $C_{\mathcal{N}}$ will tend to higher values if there is higher-order radiation which is not strongly ordered with respect to the leading order structure, suggesting that the jet has greater than N subjets. The dimensionless observable $C_{\mathcal{N}}$ is not invariant under Lorentz boost, and scales as $\gamma^{-\beta}$ for Lorentz boost γ [127]. That is to say that greater boost pushes C_N to smaller values.

One advantage that the ECFs and associated observables hold over some others is the insensitivity to recoil effects which results from the fact that there is no explicit reclustering of the jet, unlike, for example, the N -subjettiness observables.

Another observable useful for identifying two-pronged structure, which is invariant under Lorentz boost, $D_2^{(\beta)}$, is given by

$$D_2^{(\beta)} = \frac{\text{ECF}(3, \beta) \text{ECF}(1, \beta)^3}{\text{ECF}(2, \beta)^3}. \quad (3.13)$$

It can be seen that the observables which probe N -pronged substructure require $\text{ECF}(N + 1)$ for it to be possible to calculate them. As such, the ECFs calculated here are given by:

$$\text{ECF1} = \sum_{i \in J} p_{T_i} \quad (3.14a)$$

$$\text{ECF2}(\beta) = \sum_{i < j \in J} p_{T_i} p_{T_j} (\Delta R_{ij})^\beta \quad (3.14b)$$

$$\text{ECF3}(\beta) = \sum_{i < j < k \in J} p_{T_i} p_{T_j} p_{T_k} (\Delta R_{ij} \Delta R_{ik} \Delta R_{jk})^\beta \quad (3.14c)$$

$$\text{ECF4}(\beta) = \sum_{i < j < k < l \in J} p_{T_i} p_{T_j} p_{T_k} p_{T_l} (\Delta R_{ij} \Delta R_{ik} \Delta R_{il} \Delta R_{jk} \Delta R_{jl} \Delta R_{kl})^\beta. \quad (3.14d)$$

The ECFs themselves, as described in equation 3.14, are highly correlated with the particle multiplicity of the jet. To combat this, normalised versions of the ECFs are

defined as:

$$\text{ECF2}^{\text{NORM}}(\beta) = \frac{\text{ECF2}(\beta)}{\text{ECF1}^2} \quad (3.15a)$$

$$\text{ECF3}^{\text{NORM}}(\beta) = \frac{\text{ECF3}(\beta)}{\text{ECF1}^3} \quad (3.15b)$$

$$\text{ECF4}^{\text{NORM}}(\beta) = \frac{\text{ECF4}(\beta)}{\text{ECF1}^4}. \quad (3.15c)$$

From this point on, the ECFs referred to will be these normalised versions.

***N*-subjettiness ratios [129]** A jet's *N*-subjettiness provides a measure of the degree to which that jet is compatible with comprising *N* or fewer subjets. It is given by:

$$\tau_N = \frac{1}{d_0} \sum_k p_{t,k} \min \{ \Delta R_{1,k}, \Delta R_{2,k}, \dots, \Delta R_{N,k} \}, \quad (3.16a)$$

$$d_0 = \sum_k p_{t,k} R_0. \quad (3.16b)$$

The constituents of the jet are reclustered using the exclusive k_t algorithm [109], which clusters soft constituents first, up to the harder, larger-angle proto-jets last. The clustering continues until *N* subjets are returned. The calculation of τ_N then runs over the constituents of the jet, with $p_{t,k}$ the p_T of the k^{th} constituent. Here $\Delta R_{N,k}$ is the distance in $\eta - \phi$ space between the axis of the N^{th} subjet and the k^{th} jet constituent, while R_0 is the jet radius parameter.

Here, the winner-takes-all (WTA) recombination scheme is used, whereby the subjet axes are defined by the hardest particle contained therein. The *N*-subjettiness

observables used in the studies presented here are then given by:

$$\tau_0(\beta) = \sum_{i \in J} p_{T_i} R^\beta \quad (3.17a)$$

$$\tau_1(\beta) = \frac{1}{\tau_0(\beta)} \sum_{i \in J} p_{T_i} \Delta R_{a_1, i}^\beta \quad (3.17b)$$

$$\tau_2(\beta) = \frac{1}{\tau_0(\beta)} \sum_{i \in J} p_{T_i} \min(\Delta R_{a_1, i}^\beta, \Delta R_{a_2, i}^\beta) \quad (3.17c)$$

$$\tau_3(\beta) = \frac{1}{\tau_0(\beta)} \sum_{i \in J} p_{T_i} \min(\Delta R_{a_1, i}^\beta, \Delta R_{a_2, i}^\beta, \Delta R_{a_3, i}^\beta) \quad (3.17d)$$

$$\tau_4(\beta) = \frac{1}{\tau_0(\beta)} \sum_{i \in J} p_{T_i} \min(\Delta R_{a_1, i}^\beta, \Delta R_{a_2, i}^\beta, \Delta R_{a_3, i}^\beta, \Delta R_{a_4, i}^\beta) \quad (3.17e)$$

Here, the subjet axis is a_n , ΔR the angular distance and β a parameter to weight the angular separation. In the studies presented here, β is set to unity.

The N -subjettiness observables operate in a somewhat similar manner to the ECFs described above. For N -pronged substructure, τ_N tends towards smaller values. If one considers a jet originating from the decay of a boosted top quark, it should have three-pronged structure. Considering τ_3 , as expanded in Equation 3.17d, this should tend towards lower values as most constituents will be around one of the three subjet axes corresponding to the prongs of the jet and therefore send the angular term in the product to a small value. If one were to evaluate τ_2 , as expanded in Equation 3.17c, for this same jet, one would expect it to tend to higher values, the reason being that the two reconstructed subjet axes do not correspond to the three hard cores that exist in the jet. This means that whilst perhaps there are constituents in two of the cores which are close to the reconstructed subjet axes, there will be many constituents in the third core which are far from both reconstructed subjet axes. The angular terms for these constituents will then send the sum corresponding to τ_2 to higher values.

It was stated that the N -subjettiness encodes how consistent a jet is with being formed of N or fewer subjets. The case of characterising a three-pronged jet with τ_2 was considered above. Let us now consider characterising a two-pronged jet with τ_3 . In this case, it is likely that τ_3 will still tend towards smaller values. This is because, of the three reclustered subjet axes, it is likely that two will correspond to the real hard cores of the two-pronged jet. This means that much of the radiation within the jet will be close to those axes and thus the angular terms will take small values sending the observable to a low value. The same would be true for a one-pronged jet. It was

shown in the previous example that τ_N goes to higher values for jets with more than N prongs, so from these two cases, it is clear that τ_N does indeed quantify consistency with N or fewer subjets.

With this in mind, it is often instructive to consider the ratios of the N -subjettiness observables, in order to focus on precisely N subjets. A jet originating from a hadronically decaying boosted top quark will have three-pronged substructure. It is then expected that τ_3 will tend towards lower values, whilst τ_2 will tend towards higher values. The former suggests consistency with three or fewer subjets, whilst the latter suggests inconsistency with two or fewer subjets, leading to the conclusion that the jet is consistent with exactly three subjets. Taking the ratio between τ_3 and τ_2 captures this in a single observable. For a jet with three-pronged substructure, this ratio will take small values, with the numerator being small and the denominator large.

The ratios $\tau_{21} \equiv \frac{\tau_2}{\tau_1}$, $\tau_{32} \equiv \frac{\tau_3}{\tau_2}$, and $\tau_{43} \equiv \frac{\tau_4}{\tau_3}$ are used in the analysis presented here. These are useful for identifying two, three, and four-pronged jets, respectively. The ratios assume values between zero and one.

k_t splitting scales The splitting scale observables [130] are calculated by reclustering the jet with the exclusive k_t algorithm. The observables are given by:

$$\sqrt{d_{ij}} = \min(p_{Ti}, p_{Tj}) \times \Delta R_{ij}. \quad (3.18)$$

The two variants studied here, d_{12} and d_{23} , give the splitting scale of the final and penultimate merging, respectively. p_{Ti} is the p_T of the i^{th} protojet and ΔR the distance in $\eta - \phi$ space. For a contained top-jet then, $\sqrt{d_{12}}$ is expected to have a value around $m_t/2$ and d_{23} a value around $m_W/2$.

Q_W This is the minimum pair-wise invariant mass which can be calculated from the three protojets at the end of the clustering procedure using the exclusive k_t algorithm.

Jet Eccentricity By considering the eigenvalues, v , of the energy-weighted covariance matrix M of the distances in $\eta - \phi$ space between jet constituents and the jet axis, one may define the eccentricity:

$$\epsilon = 1 - \frac{v_{\min}}{v_{\max}} \quad (3.19)$$

The aforementioned covariance matrix is given by:

$$M = \sum_i E_i \begin{pmatrix} (\Delta\eta_{i,\hat{n}_r})^2 & \Delta\eta_{i,\hat{n}_r}\Delta\phi_{i,\hat{n}_r} \\ \Delta\eta_{i,\hat{n}_r}\Delta\phi_{i,\hat{n}_r} & (\Delta\phi_{i,\hat{n}_r})^2 \end{pmatrix} \quad (3.20)$$

Here the subscript $\hat{n}$ represents the jet axis, while the sum is over the jet constituents indexed with the subscript i . The eccentricity tends to zero for a perfectly circular jet, and to one for an elliptical jet.

3.1.6. Tagging

Jet tagging is a method used to identify jets which have different origins. The ability to perform this sort of discrimination is key to isolating certain final states and consequently is of the utmost importance to a huge number of ATLAS analyses, from precision measurements of SM processes to searches for rare ones and BSM signatures [131]. In the context of the analysis presented here, tagging is pertinent in two regards. Firstly, a b -tagging algorithm is needed to isolate the events used in the measurement, and secondly, the results of the measurement may be of use in top tagging algorithms [132] used to identify the process being measured here.

b -tagging algorithms are able to exploit characteristics of jets arising from b -quarks, shown in Figure 3.11, to distinguish them from those arising from light quarks and gluons. The hadronisation of a b -quark produces, amongst the rest of the jet, a b -hadron which contains the b -quark. These hadrons have a long lifetime which allows them to travel macroscopic distances, on the order of a few mm before decaying. This produces a secondary vertex at the decay of the b -hadron. The high mass of the b -quark means that the decay products may have relatively large p_T and thus be produced at large angle with respect to the original b -jet direction. When extrapolated inwards, the tracks corresponding to this secondary vertex will be displaced from the primary vertex by a distance known as the *impact parameter*, d_0 . A secondary consequence of the high mass of the b -quark is a high track multiplicity at the secondary vertex. The experimental signature of a b -quark, then, is a jet containing a secondary vertex with tracks with high- d_0 .

The approach to b -tagging in ATLAS follows a two-staged strategy. First, several low-level algorithms are used to construct features of the jet using the tracks and






Figure 3.11: The typical characteristics and geometry of a b -jet. Diagram from Ref. [133]


reconstructing vertices from those tracks. There are two algorithms based on the impact parameters — IP2D and IP3D [134] — which are part of this ensemble of low-level algorithms. Also used is a tagger based on a recurrent neural network, using tracks as input, called RNNIP [134]. An algorithm called SV1 [135] is used to tag the secondary vertex in a jet, before the JetFitter multi-vertexing algorithm [136] is used to reconstruct the full b -hadron decay chain. In the second stage, the DL1r algorithm [137, 138] takes the results of these low-level algorithms and the jet kinematics — a total of 31 variables — as input to a deep neural network and then outputs a discriminant which can be cut on. The performance of the complete algorithm is then expressed in terms of its ability both to identify b -jets and to reject c -jets and light quark jets. The DL1r algorithm offers several “working points” which represent a cut on the output discriminant which provides a certain efficiency. The 77% working point, for example, applies a cut on the discriminant value which correctly identifies a true b -jet 77% of the time. A comparison of the DL1r performance with previous b -tagging algorithms can be seen in Figure 3.12, along with the dependence of the b -tagging efficiency and c -jet rejection rate on the jet p_T .



(a) For a given b -tagging efficiency, from top to bottom are the background rejection, light-flavour rejection, and c -jet rejection. The latter two are given as a ratio to the MV2c10 algorithm.



(b) b -tagging efficiency at 77% working point as a function of the jet p_T , with the ratio relative to the MV2c10 algorithm.



(c) c -jet rejection at 77% working point as a function of the jet p_T , with the ratio relative to the MV2c10 algorithm.

Figure 3.12: Performance of the DL1r b -tagging algorithm as compared to two predecessors — DL1 and MV2c10. Plots from Ref. [138].

The top tagger used in ATLAS [139] is based on a deep neural network which takes as input high-level jet substructure observables of the kind measured in the analysis presented in this thesis. It aims to identify the distinctive three-pronged structure expected in a large- R jet containing the decay products of a top quark. The algorithm outputs a discriminant, as in the b -tagging, which may be cut on to identify events containing a top quark. As a result, if these substructure observables which are used in the taggers are themselves poorly modelled, this can lead to inaccurate estimates of both tagging efficiencies and misidentification rates for those taggers. This has knock-on consequences for physics analyses seeking to use these taggers to isolate a particular process. This may then limit the sensitivity of future measurements and BSM searches, for which substructure tagging is expected to be of increasing importance. This is because heavy BSM particles which decay into lighter SM particles which would be produced at high boost, give collimated jets with rich substructure.

There has also been evidence in recent studies that tagging performance can be improved by using machine learning techniques directly on jet constituents, as opposed to high-level observables constructed from those constituents [140, 141]. This would seem to suggest that there is room for improvement on both the modelling of the dynamics of these processes and the observables themselves, which may not be capturing all of the available information for discrimination.

3.2. Analysis overview

The analysis presented here is fundamentally a differential cross-section analysis, performed using 139 fb^{-1} of data collected by the ATLAS detector from pp collisions at $\sqrt{s} = 13 \text{ TeV}$; the full Run 2 dataset. Measuring the cross-section of the $t\bar{t}$ process means counting the number of events recorded, subject to some selection criteria, so that that tally may be compared to a number of predictions based on the underlying physics. Such cross-section measurements are frequently used to probe not only the inclusive cross-section given pp collisions at a particular collision energy, but how the cross-section is dependent on the $t\bar{t}$ decay channel or on the kinematics or other properties of the system. A measurement may aim, for example, to measure one of the three decay channels — all-hadronic, ℓ +jets, and dileptonic. These channels are defined by the decay of the W bosons which result from the decay of each top quark into a W boson and a b -quark. The W may decay either hadronically into a

pair of quarks or leptonically into a lepton-neutrino pair. The all-hadronic, ℓ +jets, and dilepton channels are given by the case where both, one, or none of the W -bosons decay hadronically. Once a channel has been selected, the cross-section may be measured differentially, meaning as a function of some other property such as the p_T of the large- R jet corresponding to the top quark. Double-differential measurements may then further divide such distributions into separate regions. An example of this could be the mass of that same large- R jet. Cross-section measurements such as these have formed a key part of the top physics programme for both ATLAS [2, 142–145] and CMS [146–151]. They represent a vital test of theoretical calculations as well as a means to tune and improve the Monte Carlo used to simulate these processes.

The analysis presented here measured both differential and double-differential cross-sections in both the all-hadronic and ℓ +jets $t\bar{t}$ channels, specifically when the top quarks are produced with sufficient boost that their decay products overlap and are collimated. Only the ℓ +jets channel is presented in this thesis. The cross-section is measured as a function of several jet substructure observables and then double-differentially as a function of both the p_T and mass of the large- R jet corresponding to the top quark. Measurements of jet substructure have been made in several contexts by both ATLAS [103, 152–154] and CMS [122, 155]. The investigation of the top quark decay, showering, and hadronisation effects provides a test of theoretical calculations and may improve sensitivity to the effects of BSM physics through improvements to modelling, and performance of future analysis tools, as well as possible modification to the substructure itself.

Following the event reconstruction and selection in order to isolate the ℓ +jets boosted $t\bar{t}$ process, the substructure of the hadronic jet is calculated here for several observables, before the distributions are *unfolded* to remove the effects of interaction with the detector. The final results are given with a number of uncertainties considered which are related to the detector and reconstruction of objects and events, as well as the modelling of these processes in MC. The unfolded data are compared to the predictions of several NLO+PS models.

Chapter 4.

Event Reconstruction

The events analysed in this thesis make use of several reconstructed physics objects. This chapter describes the data and Monte Carlo simulated samples from which these objects are reconstructed, followed by definitions of the objects themselves. Objects are described at detector-level and particle-level. The detector-level refers to signals left in the detector by incident particles. Particle-level is defined by stable particles ($c\tau > 10$ mm) without the effects of interaction with the detector. As such, detector-level objects are reconstructed in both data and simulation samples, whereas particle-level objects are reconstructed only in simulated samples.

High-level objects — representing particles and their decay products — are constructed from low-level objects at detector-level. These low-level objects are topological clusters in the calorimeter systems and tracks within the ID and the muon system.

Following object definitions, a study is presented which compares different definitions for large- R jets and their associated substructure.

4.1. Data and Monte Carlo samples

4.1.1. Data samples

The data used for this analysis were collected between 2015 and 2018, comprising the full ATLAS Run 2 dataset. This corresponds to an integrated luminosity of 139 fb^{-1} collected during $\sqrt{s} = 13 \text{ TeV}$ pp collisions at the LHC. In fact, more than 139 fb^{-1} was collected, but this subset is that which was marked as “good for physics” as per

ATLAS' data quality requirements [156]. The uncertainty on the luminosity for this dataset is 1.7% [157]. This luminosity has recently been updated to 140 fb^{-1} , with a lower uncertainty, but this change was found to have a negligible impact on the analysis presented here, so the original luminosity was retained.

4.1.2. Monte Carlo samples

Events simulated in Monte Carlo are interfaced with a simulation of the ATLAS Detector implemented in GEANT4 [55] in order to reproduce the effects of interaction with the detector. The effect of pileup is simulated by overlaying the event containing the hard scatter with events which are generated in PYTHIA8 [158]. This is done using the A3 tune [159], interfaced with the NNPDF2.31o set of PDFs [160]. For each of the three data-taking periods comprising Run 2 — 2015/16, 2017, and 2018 — the pileup profile in simulated samples is reweighted in order to match the average number of collisions per bunch-crossing in the data.

$t\bar{t}$ prediction Several samples are produced for the $t\bar{t}$ process; a nominal sample and several alternative samples which are compared to the nominal sample in order to probe the modelling of the process and calculate the consequent uncertainties.

The nominal $t\bar{t}$ prediction — POWHEG+PYTHIA8 — is produced using matrix elements generated at NLO using POWHEGBOX v2 [161–164] interfaced with PDFs from the NNPDF3.0NLO set [165]. Matching of NLOME to PS is governed by the h_{damp} parameter in POWHEG, which is set to $1.5 \times m_t$, the mass of the top quark, which in turn is set to 172.5 GeV. h_{damp} is a tunable parameter which regulates the p_T of the hardest emission in the ISR against which the $t\bar{t}$ system recoils. In this way, it is used to control the cutoff scale for the ME calculation, below which the PS takes over. Its main impact is therefore expected to be on the p_T distribution for the $t\bar{t}$ system. The scale for renormalisation and factorisation is set using a function of the form $\sqrt{m_t^2 + p_T^2}$. The generated events are then passed to PYTHIA8.230 for simulation of the parton shower and hadronisation, with the A14 tune [166] and using the NNPDF2.31o set of PDFs [160]. The final stage is the simulation of heavy flavour quarks, b and c . This is handled separately by EVTGEN v1.6.0 [167]. The resultant sample is then normalised to the TOP++2.0 [168–174] cross-section at NNLO which includes resummation of NNLL soft gluon terms. This results in a cross-section of $\sigma(t\bar{t})_{\text{NNLO+NNLL}} = 832 \pm 51 \text{ pb}$.

Several other samples are generated which are variations on the nominal $t\bar{t}$ prediction. These are identical to the sample described above except for some specific variation.

A sample is generated to probe the effect of the h_{damp} parameter by setting $h_{\text{damp}} = 3 \times m_t$.

Another set of samples [175] is used to allow estimation of uncertainties due to the choice of scale in the hard scatter and the simulation of the parton shower. The nominal sample is reweighted to produce several different samples without a complete generation of fresh samples for each variation. These samples correspond to half and double the nominal value of both renormalisation scale μ_R and factorisation scale, μ_F . For the scale of the showering, Var3c variations for the A14 tune are applied, and for FSR, μ_R^{FSR} also has its nominal value doubled and halved.

In order to assess the modelling of parton shower and hadronisation, another sample is generated which uses the same POWHEG generator setup as described above for the nominal sample but a different model for the parton shower and hadronisation. This is HERWIG 7.13 [53, 176], with PDFs from the MMHT2014L0 set [177] and the H7UE parameters [51, 176]. Heavy flavour hadron decays are simulated in the same way as for the nominal POWHEG+PYTHIA8 prediction.

Another sample is generated to assess the modelling of the hard scatter. This sample replaces POWHEG with MADGRAPH_AMC@NLOV2.6.0 [178] for the hard scattering calculation. This is interfaced with the same PDF set as for the nominal sample, and uses the same choice for renormalisation and factorisation scales. To preserve spin correlations, top quarks are made to decay using MADSPIN [179, 180] at LO. The functional form $\mu_q = H_T/2$ is used for the shower starting scale, with H_T the scalar sum of the p_T of all outgoing partons. These events are interfaced with Pythia8.230 as for the nominal sample, also using the A14 tune and the same set of PDFs, with heavy flavour hadrons also decaying in the same way. The nominal prediction has matrix element corrections turned on. These are designed to avoid the double-counting of radiation in matrix element and showering. In this alternate sample, these corrections are not present. For this reason, another variation of the nominal prediction is generated which is identical to the POWHEG+PYTHIA8 sample but with the matrix element corrections turned off. This is so that a comparison between the nominal sample and alternate hard scattering sample is not sensitive to differences arising from these matrix element corrections.

Finally, in order to derive an uncertainty arising from the top mass, two further variations of the nominal prediction are generated. These set the top mass to 169 GeV and 176 GeV.

Background modelling The tW process, which is an important background in this analysis, is simulated using the same generator and showering setup as for the nominal $t\bar{t}$ prediction, using the five-flavour scheme in POWHEG and the diagram removal (DR) scheme to deal with interference with $t\bar{t}$ production. Another sample is generated in the same way but using the diagram subtraction (DS) scheme instead, and this is used to assess the uncertainty associated with this choice. The s -channel and t -channel single top processes are simulated in the same way, but with no need for DR/DS. In the case of the t -channel, the four-flavour scheme is employed in POWHEG.

V +jets, with $V = W/Z$, is modelled using SHERPA v2.2.1 [181]. This generates LO and NLO matrix elements up to four and two partons, respectively, with calculations from the COMIX [182] and OPENLOOPS 1 [183–185] libraries. Parton showering is handled by SHERPA [186] with the MEPS@NLO prescription [187–190]. The NNPDF3.0nnlo PDF set is used.

$t\bar{t}V$, with $V = W/Z$, is modelled at NLO using the MADGRAPH_AMC@NLO v2.3.3 [178] generator with the NNPDF3.0nnlo PDF set. These events are then interfaced with PYTHIA8.210, using the same tune and PDF set as for the nominal $t\bar{t}$ prediction. Heavy flavour hadron decays are simulated with EVTGEN v1.2.0. $t\bar{t}H$ is modelled using the same setup as for the nominal $t\bar{t}$ prediction.

The diboson sample is produced using SHERPA v2.2.0 to generate matrix elements at LO and NLO for up to three and one additional partons, respectively. Parton showering based on Catani-Seymour factorisation [182, 186] is handled with SHERPA using the MEPS@NLO prescription and the NNPDF2.0nnlo PDF set. Virtual QCD corrections come from the OPENLOOPS 1 library.

4.2. Detector-level objects

Calorimeter clusters *Topological clusters* (topoclusters) are constructed in the calorimeter system as a measurement of the energy deposited in the detector by incident particles. A topocluster is built from a seed cell which is selected for having an energy

above a certain threshold. This is determined by the cell significance [191]:

$$\zeta_{\text{cell}}^{\text{EM}} = \frac{E_{\text{cell}}^{\text{EM}}}{\sigma_{\text{noise,cell}}^{\text{EM}}}, \quad (4.1)$$

where $E_{\text{cell}}^{\text{EM}}$ is the energy recorded in that cell and $\sigma_{\text{noise,cell}}^{\text{EM}}$ the expected noise in that cell from the electronics and effects of pileup, as measured at the EM scale. Local cell reweighting (LCW) is sometimes used to correct the calorimeter response to that expected for hadrons in a separate calibration, but the measurements presented here are performed exclusively with clusters calibrated to the EM scale.

The cell significance is required to be above four for a seed cell. Any cells which neighbour the seed and have significance above two are then merged with the seed to form a *proto-cluster*. An iterative process continues to merge neighbouring cells with significance above two until none remain. Finally, all remaining neighbour cells which have registered positive signals are added to the proto-cluster.

Proto-clusters which contain two or more local energy maxima — cells with $E_{\text{cell}}^{\text{EM}} > 500 \text{ GeV}$ and with neighbours with energy below that — are also subject to a splitting stage. The clustering process is repeated for each pair of maxima, with cells shared between the two having their energy divided up according to the energy of each maximum and the cell's location.

Tracks The reconstruction of primary tracks in the ID proceeds in several stages. First, clusters are formed by grouping connected hits in a given pixel or strip sensor, where a hit is a signal above the charge threshold. Clusters in pixel layers correspond to a *space-point*, whilst in strip layers, clusters on both sides of the layer are combined to form a space-point. A space-point provides a 3D measurement. Space-points can then be used to build tracks. First, track seeds are constructed from triplets of space-points. An iterative algorithm is used [192], performing a loose search first for track candidates built from the seeds by incorporating space-points from the rest of the silicon layers. This search establishes many combinatorial options before employing an ambiguity solver method to select the final set of tracks. This solver orders track candidates with a *track score* and gives preference to those with a higher score. Clusters which have been matched to multiple track candidates are then tackled by the ambiguity solver, requiring shared clusters to meet certain criteria or be removed from that

track candidate and its score recalculated. The ambiguity solver will also reject track candidates which have

- $p_T < 400 \text{ MeV}$
- fewer than seven clusters in the SCT and pixel detector
- outwith $|\eta| < 2.5$
- more than two holes in the entire track
- more than one or two shared clusters per layer for pixel and SCT sub-detectors, respectively
- more than one or two holes in the track in pixel and SCT sub-detectors, respectively
- $|d_0^{\text{BL}}| > 2.0 \text{ mm}$ or $|z_0^{\text{BL}} \sin \theta| > 3.0 \text{ mm}$, with d_0^{BL} the transverse impact parameter relative to the beamline, z_0^{BL} the longitudinal difference along the beamline between where d_0^{BL} is measured and the primary vertex, and θ the track's polar angle.

For the analysis presented in this thesis, tracks are required to pass the ATLAS TightPrimary working point, which defines an additional set of criteria which must be satisfied. This requires that each track candidate has $p_T > 500 \text{ MeV}$ and $|\eta| < 2.5$. It also stipulates a minimum number of hits in the silicon layers of the ID — 9 and 11 for $|\eta| < 1.65$ and $|\eta| > 1.65$, respectively. Additionally, no holes are allowed in the pixel detector and a hit must exist in the IBL or B-layer of the pixel detector if expected by the track trajectory.

In the analysis presented in this thesis, tracks are used to construct jet substructure observables. This is done by ghost-association [193] of the tracks to jets — tracks are clustered into a jet along with the usual constituents as described below, but with zero magnitude so that they have no impact on the kinematics of the resultant jet. Tracks are then considered to be associated to the jets into which they have been clustered. In this analysis, tracks are ghost-associated to the small- R EMPflow jets described below. As tracks are used to calculate the substructure of large- R reclustered jets (described below), it is the tracks associated to the small- R EMPflow jets from which those large- R reclustered jets are reclustered which are then passed to the algorithm used to construct the substructure observable for that large- R reclustered jet.

Interactions between multiple particles and particle decays produce divergent tracks. The points which these tracks emanate from are called *vertices*. ID tracks are used to reconstruct these vertices. The point in an event with the greatest total track p_T is labelled the primary vertex, with all others labelled as secondary vertices. The primary vertex represents the location of the hard scattering in the event, whilst secondary vertices can arise from a number of processes and are often used to identify particle decays.

Tracks corresponding to muons are constructed from hits in the muon spectrometer matched to ID tracks. For the reconstruction of MS tracks, track segments are first identified in individual stations of MDTs using a Hough transform [194]. Track candidates combine segments from the different stations, with three-dimensional track candidates then incorporating information in the non-bending plane from the RPCs and TGCs. A χ^2 fit of the hits in each track candidate is performed, accounting for alignment and interaction with detector material, with outliers then removed and hits on this trajectory added to the track. The fit is then repeated.

Tracks must contain a minimum of two segments unless in the transition region between barrel and endcap regions, in which case one segment of high quality is enough. In the case where two tracks have a shared segment, both are kept if they also contain a segment in the outermost layer which they do not share. Otherwise, an overlap removal procedure assigns segments to the track with best χ^2 fit. The full procedure is described in Ref. [195].

Electrons The reconstruction of electrons combines information from the inner detector (ID) with calorimetry information [196,197]. Tracks from the ID are matched to deposits in the electromagnetic calorimeter (ECAL). These tracks must be associated to the primary vertex. This association is done by applying criteria on the transverse impact parameter, d_0 , and the longitudinal impact parameter, z_0 , as measured relative to the beamline. For a given electron, it is required that tracks have $|d_0/\sigma(d_0)| < 5$ and $|\Delta z_0| < 0.5$ mm. The ECAL clusters are required to fall within the central acceptance, $|\eta| < 2.47$, and have $E_T > 27$ GeV. Clusters must also be outwith the transition region between the barrel and end-cap sections of the ECAL, $1.37 < |\eta| < 1.52$.

Electrons are distinguished from other objects with similar signatures using a likelihood-based method requiring that electrons satisfy the ATLAS *tight* working

point. This corresponds to a background rejection factor of 3.5 and selection efficiency of 80%.

Electron candidates also have an isolation requirement in order to ensure they are sufficiently separated from other activity. The ATLAS *tight* requirement sets criteria on the amount of energy allowed surrounding the electron candidate, for both tracks and calorimeter deposits. The sum of the transverse energy in a cone of radius $R = 0.2$ around the candidate is not allowed to exceed 6% of its p_T . The same threshold is required for the sum of track p_T in a cone of variable size up to $R = 0.2$.

Muons The reconstruction of muons combines tracks in the ID with those found in the muon spectrometer [198]. Tracks from the ID must be associated to the primary vertex. This is done by requiring that ID tracks have $|d_0/\sigma(d_0)| < 3$ and $|\Delta z_0| < 0.5\text{mm}$. As with electron reconstruction, combined tracks are required to have $p_T > 27\text{ GeV}$ and $|\eta| < 2.5$.

Muons must satisfy the ATLAS *medium* identification working point. The criteria set by this requirement are related to the quality of the tracks constructed in the MS and their association to ID tracks. MS tracks must have a minimum of three hits, with at least two layers of the MDTs. For tracks with $|\eta| < 0.1$, only one MDT layer is required and a hole in the MDTs is permitted.

Again, similarly to electron reconstruction, muons are subject to isolation requirements as set by the ATLAS *tight* requirements. The sum of the transverse energy in a cone of radius $R = 0.2$ around the candidate is not allowed to exceed 15% of its p_T for topoclusters in the calorimeters. In the ID, the the sum of track p_T in a cone of variable size up to $R = 0.3$ is not allowed to exceed 4%.

E_T^{miss} In an ideal collision in ATLAS, the transverse momenta for all objects reconstructed in an event should sum to zero. The missing transverse energy in an event, E_T^{miss} , may originate from limits on the measurements made by the detector, or from particles which do not interact with the detector material. In the context of the analysis presented in this thesis, the E_T^{miss} is used as a proxy for neutrinos, which are the only SM particles which are not measured by the detector.

E_T^{miss} is the magnitude of the negative vector sum of the p_T of all selected and calibrated objects in an event, as well as a soft term which is calculated from the p_T of

ID tracks associated to the primary vertex which are not also associated to any of the reconstructed objects.

Small- R jets These are reconstructed using particle flow objects [112]. The particle flow algorithm — discussed in Chapter 3 — makes use of both calorimeter and tracking information from the detector to provide optimal resolution for charged and neutral particles at a range of energy scales. The anti- k_t jet clustering algorithm [97] is used to cluster these particle flow objects with a radius parameter of $R = 0.4$. In order to suppress jets originating from pileup, the recommended tight cut on the ATLAS Jet Vertex Tagger (JVT) discriminant is applied, requiring jets with $p_T < 60$ GeV to have $\text{JVT} > 0.5$. Jets are required to have $p_T > 25$ GeV and $|\eta| < 2.5$ such that they fall within the acceptance of the inner detector. The jets are calibrated using the standard approach in order to correct the energy of the jet such that it, on average, matches the particle-level energy [199].

b -tagging The DL1r multivariate algorithm [137] is used to identify small- R jets which originate from a b -quark and thus contain a b -hadron. This algorithm takes as input information on secondary vertices, jet kinematics, impact parameters, and outputs from other tagging algorithms in order to calculate a discriminant which can be used to select jets. A cut is applied to this discriminant corresponding to the 77% efficiency working point of the DL1r algorithm, as measured in simulated $t\bar{t}$ events.

Large- R reclustered jets RC jets are constructed using the small- R jets previously described as input to the anti- k_t jet clustering algorithm, using a radius parameter $R = 1.0$. The calibrations and uncertainties for these RC jets are propagated from the small- R jets used as input. Most pileup mitigation is thus inherited by the RC jets from the techniques applied to the small- R jets. In addition to this, a trimming technique is applied to the RC jets which removes all small- R jets comprising less than 5% of the p_T of the RC jet. This is with the aim of removing small- R jets originating from pileup.

RC jets are discarded if they comprise only one constituent. This is because of a lack of mass calibration for the small- R jets. These small- R jets have an energy calibration and so in the case of an RC jet with at least two constituents, the mass of that jet is calculated using the calibrated kinematics of the input small- R jets. With only one such constituent, this is not possible. The number of events containing RC jets

with only one constituent is very small. By the end of the event selection procedure, only 1.41% of events in the signal MC contained an RC jet constructed from only one small- R jet. This small statistical loss is more than made up for by the number of additional events passing selection as a result of using the RC jets. This is discussed in Section 4.4. Large- R RC jets such as these offer lower uncertainties and improved mass resolution by comparison to conventional Large- R jets reconstructed from clusters in the calorimeter [113], as described in Section 3.1.3.

Overlap An overlap removal procedure is applied to mitigate any double counting of objects which are independently reconstructed. An electron is rejected if it shares a track with a muon or if it has a track overlapping with another electron. For jets which are close to an electron, that is with $\Delta R(j, e) < 0.2$, the closest jet fulfilling this criterion is removed. If there are then other jets close to an electron, that is with $\Delta R(j, e) < 0.4$, that electron is removed. Jets which are either close to a muon, that is with $\Delta R(j, \mu) < 0.2$, or share a track with that muon are removed if they comprise fewer than three tracks. Any muon which is then close to a jet, that is with $\Delta R(j, \mu) < 0.4$, is removed.

4.3. Particle-level objects

Leptons are defined such that they do not come from a hadron decay, either directly or through the decay of a τ . In this way, leptons must come from electroweak decay without the need for additional matching to a W boson. Leptons are then *dressed*, meaning that their four-momenta are summed with any stable photons within $\Delta R = 0.1$. It is then required that these dressed leptons have $p_T > 27 \text{ GeV}$ and $|\eta| < 2.5$.

At particle-level, all neutrinos should come from electroweak decay as determined above, with the E_T^{miss} then calculated from the four-momentum sum of the selected neutrinos.

Small- R jets at particle-level are built by clustering all stable particles with $|\eta| < 5.0$, except for the selected leptons and photons they may have been dressed with. The anti- k_t jet clustering algorithm is used with a jet radius parameter of $R = 0.4$.

At particle-level, b -tagging is done by ghost-matching [200] b -hadrons to jets. These b -hadrons are required to have $p_T > 5 \text{ GeV}$ and any jet with at least one ghost-matched b -hadron is marked as being b -tagged.

Reclustered large- R jets are constructed at particle-level using the same method as at detector-level. The RC jets use the anti- k_t jet clustering algorithm with a radius parameter of $R = 1.0$ and using the particle-level small- R jets as input. The same trimming is also applied, removing any small- R jet with less than 5% of the p_T of the RC jet.

As when constructing the jet substructure using tracks at detector-level, at particle-level only the charged components of a jet are used as input to the substructure calculation. It has been shown that the correlation between the substructure calculated from the whole jet and that using only the charged components is very strong.

4.4. Study on different jet collections


The choice to use tracks — that is, only the charged component of the jet — to study the substructure, was motivated by two factors:

- The improved resolution provided by tracks, as compared to calorimeter clusters.
- The propagation of track information allows for a bottom-up uncertainty on the substructure observables, which tends to be smaller than those associated with the calorimeter clusters.

The resolution improvement is illustrated in Figures 4.1a and 4.1b, which compare large- R jet substructure using calorimeter and track information, respectively. The move to using only track information to determine substructure resulted in reduced migration for the substructure observables of interest in the analysis.


In choosing to use reclustered large- R jets in the ℓ +jets channel as opposed to the $R = 1.0$ LCTopo large- R jets constructed from calorimeter clusters, a study was conducted to investigate the consequences of this choice. The move to reclustered jets was motivated by several factors:

- Reduced migration in the substructure observables being measured.
- Improved resolution on the mass of the hadronic top candidate jet.



(a) LCTopo large- R jet collection with substructure calculated from topological clusters.

(b) LCTopo large- R jet collection with substructure calculated from tracks.



(c) Reclustered large- R jet collection with substructure calculated from tracks.

Figure 4.1: Migration matrices for τ_{32} .





Figure 4.2: Detector-level τ_{32} . There are approximately 30 000 more events in the data which pass the event selection when using the reclustered jet collection.

- Reduced sensitivity to pileup.
- Smaller uncertainties as a result of using only a single jet collection.


The reclustered jets use $R = 0.4$ particle flow jets as input. The substructure of the resultant jet is then calculated using tracks which are ghost-matched to the sub-jets. As illustrated in Figures 4.1b and 4.1c, the move to using reclustered jets rather than large- R jets further reduces the migrations for the substructure observables of interest in the analysis.

Moving to the reclustered jet collection also resulted in a significant increase in the number of events passing the event selection, as shown in Figure 4.2. It was hypothesised that this could be explained for the most part by two effects; the difference in jet mass resolution, and the difference in jet p_T .

If the reclustered jets tend to be more energetic then more jets may pass the cut at 350 GeV (this is a part of the event selection, discussed in Chapter 5) than in the case where the LCTopo large- R jet collection is being used. It was found that the reclustered jets are indeed more energetic on average than the large- R jets. This can be seen in Figure 4.3, which shows the p_T difference between the leading large- R jet in each event and its ΔR -matched reclustered jet, as a proportion of the p_T of the large- R jet. This effect is also observed at particle-level, as shown in Figure 4.4.




(a) LCTopo large- R jet collection used for event selection.




(b) Reclustered large- R jet collection used for event selection.

Figure 4.3: Detector-level distribution of the difference in the p_T between the leading large- R jet and the ΔR -matched reclustered large- R jet. The difference is expressed as a proportion of the p_T of the LCTopo large- R jet.



(a) LCTopo large- R jet collection used for event selection.



(b) Reclustered large- R jet collection used for event selection.

Figure 4.4: Detector and particle-level distribution of the difference in the p_T of the leading large- R jet and the ΔR -matched reclustered large- R jet. The difference is expressed as a proportion of the p_T of the LCTopo large- R jet and utilises only the signal Monte Carlo simulation. The distributions are normalised such that the total area is equal to unity.




Figure 4.5: Detector-level mass distribution for the hadronic top candidate jet, having had the mass window cut applied during the event selection procedure.

As is evidenced in Figure 4.5, the leading large- R jet mass distribution in the case of the reclustered large- R jets exhibits a sharper peak than the LCTopo jet collection, which falls off more gently, particularly at lower values. This would seem to increase the probability, in the case of the LCTopo large- R jet collection, that a jet falls below the 110 GeV cut-off enforced by the mass window during the event selection.

In to order evaluate these effects, the event selection was run twice. The first time, the standard selection criteria were applied, with p_T and mass window cuts — and any other cuts on the large- R jets — applied on the reclustered large- R jets. The leading reclustered jet in the event is then ΔR -matched to an LCTopo large- R jet. Distributions are then plotted of the p_T of the LCTopo jets in question, provided that they have p_T below 350 GeV. In other words, a p_T distribution is plotted of the LCTopo jets which would have been eliminated by the 350 GeV cut in the event selection. The second time, the procedure is repeated but with the selections inverted. That is, the LCTopo jets are used for the selection procedure and the p_T distribution of the matched reclustered jets plotted, provided that they failed the p_T cut. The event numbers in these distributions, Figure 4.6, can be used to compare to the observed discrepancy.

A distribution is also then plotted in both selection scenarios of the difference in mass between the LCTopo jet and its matched reclustered counterpart, as a proportion of the mass of the LCTopo jet, provided that the non-selected jet falls outwith the





Figure 4.6: Detector-level p_T distribution for the hadronic top candidate jet. Each distribution shows the jets which failed the 350 GeV p_T cut when that cut was passed by the ΔR -matched jet from the other jet collection.

mass window. To avoid double-counting with the previous distributions of p_T , the non-selected jet is required to pass the p_T cut. This can be seen in Figure 4.7.

As can be seen in these distributions, the p_T and mass window cuts applied in the event selection *do* impact the LCTopo jets more than the reclustered jets. As the reclustered jets are, on average, more energetic than the LCTopo jets, there are an order of magnitude more LCTopo jets failing the p_T cut when the reclustered jets are used for selection than in the inverse scenario. As can be seen particularly in Figure 4.7b, more LCTopo jets are removed by the mass window cut at lower masses where the fall from the peak is more gradual, although the difference between the two jet collections here is not so dramatic as for the p_T cut.

These findings are summarised in Table 4.1. For each selection — whether applied on the LCTopo or reclustered large- R jets — the number of ΔR -matched non-selected jets which do not pass either the p_T cut or the mass window cut is shown. The differences are then expressed in terms of the number of events passing for the LCTopo jets as compared to the reclustered jets.



(a) LCTopo jet collection used for the event selection.

(b) RC jet collection used for the event selection.

Figure 4.7: Detector-level mass difference between LCTopo large- R jet and ΔR -matched reclustered jet counterpart when the jet from the collection which is not used for event selection falls outwith the mass window. Both jets are required to have passed the p_T cut. The difference is expressed as a proportion of the LCTopo jet mass.

Selection	p_T cut	Mass window cut	Delta
Large- R	1896	4135	+6031
RC	19951	9342	-29293
Overall	-18055	-5207	-23262

Table 4.1: For both selections applied to the LCTopo Large- R jets and the reclustered large- R jets, the number of events in the non-selected jet collection which fail the cut in question. The combination of these two cuts leads to significantly more LCTopo jets failing than for the reclustered jets.




Figure 4.8: Mean value of the τ_{32} distribution as a function of pileup for both the topocluster large- R jet collection and reclustered large- R jet collection. Calculated from signal MC in the ℓ +jets channel. The topocluster large- R jets exhibit a greater dependence on pileup, with τ_{32} values rising (less “top-like”) as pileup increases.

By taking into account the effect of these two cuts on the different jet collections, slightly more than 80% of the observed discrepancy is accounted for. It was therefore concluded that the use of the reclustered large- R jets does not constitute a problem for the analysis.

Another consideration for the analysis is the effect of pileup on the substructure observables. This will change depending on whether these observables are calculated from the topocluster large- R jets or the reclustered large- R jets. Generally, the reclustered large- R jets are more resistant to pileup than the topocluster large- R jets in their kinematics. Figure 4.8 shows that this is also the case for jet substructure observables in boosted $t\bar{t}$.

Chapter 5.

Event Selection

Introduced in this chapter are the event selection procedures for the ℓ +jets channel of the analysis, shown in Figure 5.1. The starting points for the event selection were based on the standard approach within the ATLAS top cross-section sub-group, as employed in previous ℓ +jets [143] $t\bar{t}$ cross-section measurements. From here, alterations were made as per the specific requirements of the substructure analysis and the availability of more modern techniques.

The event selection criteria for each channel are described, followed by studies conducted by the author which motivated choices in the ℓ +jets channel. The criteria described here apply at the detector-level, as applied to both data and Monte Carlo simulation.




Figure 5.1: The ℓ +jets $t\bar{t}$ decay topology. This is the final state which is targeted by the event selection procedure.

Sample	Event Yield
Data	83069
$t\bar{t}$ (ℓ +jets)	89881
Single top	2213
Fake leptons	1503
W +jets	1490
$t\bar{t}V$ ($t\bar{t}W$, $t\bar{t}Z$, $t\bar{t}H$)	913
Other	433
Data/ $\sum$ predictions	0.86

Table 5.1: Event yields after the detector-level selection process, for both data and Monte Carlo simulation samples.

Following the application of the event selection procedure, the resultant event yields are shown in Table 5.1.

5.1. Selection criteria

One of the lepton triggers [83, 88, 201] is required to be passed for each event, with the event containing one lepton with $p_T > 27 \text{ GeV}$ corresponding to the passed trigger. The transverse mass of the leptonically decaying W boson, defined as

$$m_T^W = \sqrt{2p_T^\ell E_T^{\text{miss}} \left(1 - \cos \Delta\phi \left(p_T^\ell, E_T^{\text{miss}} \right) \right)}, \quad (5.1)$$

is used along with the E_T^{miss} to reduce the contribution of the multi-jet background and fake leptons. The E_T^{miss} for an event must exceed 20 GeV and the sum of the E_T^{miss} and m_T^W must be greater than 60 GeV .

In order to select the boosted topology, the leading large- R reclustered jet in the event is required to have $p_T > 350 \text{ GeV}$. This jet is identified as the hadronic top jet. In the case that there is more than one large- R reclustered jet above this threshold, the jet which has mass closest to the mass of the top quark, $m_t = 172.5 \text{ GeV}$, is chosen as the hadronic top jet. The motivation for this choice came from a study which is presented in Section 5.4. In order to reduce the contribution from jets originating

from hadronically decaying boosted W bosons, a mass window cut is applied to the hadronic top candidate jet. This requires $122.5 \text{ GeV} < m_{\text{jet}} < 222.5 \text{ GeV}$. The boundaries of this window are discussed in a study presented in Section 5.3.

Each event must contain at least one b -tagged small- R particle flow (EMPFlow) jet. Owing to the leptonic decay of one top quark, there must be a b -tagged small- R EMPFlow jet which is close to the lepton, $\Delta R(\ell, b\text{-jet}) \leq 1.5$, with the closest b -tagged small- R EMPFlow jet defined as the leptonic b -jet.

The final requirement in the event selection is that the invariant mass of the system formed by the lepton and the leptonic b -jet, $m_{\ell b}$, is below 120 GeV . This is motivated by several factors investigated in a study presented in Section 5.2. The $m_{\ell b}$ cut reduces contributions from fake leptons, as well as from the tW -process single top background which has a large associated modelling uncertainty. The combination of the $m_{\ell b}$ cut and mass window on the measured jet also allow for the removal of any requirement to have a b -tagged small- R jet within the hadronic top candidate jet. This can both induce a bias on the substructure of the jet and negatively impact the signal-to-background ratio at high- p_T .

These requirements are summarised in Table 5.2.

5.2. $m_{\ell b}$ cut

One of the aims in the development of the event selection procedure in the ℓ +jets channel was that any potential bias on the substructure of the jet to be measured be minimised. The standard approach in $t\bar{t}$ selection in previous analyses has been to require a b -tagged jet inside the large- R jet which is identified as the hadronic top candidate jet — $\Delta R(\text{top} - \text{jet}, b\text{-jet}) \leq 1.0$. As the b -tagging algorithm is substructure-dependent, there is a risk of biasing the jet by this criterion. Another drawback of the b -tagging approach is the drop in the b -tagging efficiency at high- p_T , which becomes more pronounced in the boosted region. These considerations and criteria led to the development of the cut on the invariant mass of the lepton- b -jet, $m_{\ell b}$, in the event. The $m_{\ell b}$ is calculated by four-vector addition of the lepton and closest b -jet. The effect of this cut, as well as any comparison to the b -tagging approach, is investigated by applying the rest of the event selection criteria first so as to isolate the impact of each approach.

Objects	Detector-level Selection	
Leptons	Electrons $p_T > 27 \text{ GeV}$ $ \eta < 1.37$ $1.52 < \eta < 2.47$ $ d_0/\sigma d_0 < 5$ $ \Delta z_0 < 0.5 \text{ mm}$	Muons $p_T > 27 \text{ GeV}$ $ \eta < 2.47$ $ d_0/\sigma d_0 < 3$ $ \Delta z_0 < 0.5 \text{ mm}$
E_T^{miss}, m_T^W	$E_T^{\text{miss}} > 20 \text{ GeV}$ $E_T^{\text{miss}} + m_T^W > 60 \text{ GeV}$	
Large- R reclustered jets	$122.5 \text{ GeV} < m_{\text{jet}} < 222.5 \text{ GeV}$ $p_T > 350 \text{ GeV}$ $\Delta R(\ell, \text{hadronic top}) \geq 1.0$ $N_{\text{constituents}} > 1$	
EMPFLOW small- R jets	$N_{\text{jets}} > 3$ $p_T > 25 \text{ GeV}$ $ \eta < 2.5$ $\text{JVT} > 0.5 \text{ (if } p_T < 60 \text{ GeV)}$	
b -tagging	DL1r at 77% WP At least one b -tagged small- R jet At least one $\Delta R(b\text{-jet}, \ell) \leq 1.5$	
$m_{\ell b}$	$m_{\ell b} < 120 \text{ GeV}$	

Table 5.2: Summary of the event selection procedure in the lepton + jets channel.




Figure 5.2: Leading large-jet mass prior to applying any final cut. Two peaks can be observed, at the W and t masses.

After all cuts except for the $m_{\ell b}$ cut, there is still a significant contribution from boosted W bosons in the events. This can be seen in figure 5.2.

The possible bias induced by the application of the b -tagging algorithm is shown in Figure 5.3. Here, the event selection procedure is applied as normal up to the point that the b -tagging or $m_{\ell b}$ requirement would be made. Events are selected for which the hadronic top candidate jet at detector-level can be matched to a small- R b -jet at particle-level. This selects “real” b -jets inside the hadronic top candidate jet. There is no b -tagging requirement on this hadronic top candidate jet at detector-level. Finally, on those selected hadronic top candidate jets, the requirement at detector level to have $\Delta R(\text{top} - \text{jet}, b\text{-jet}) \leq 1.0$ is applied. These two distributions — with and without detector-level b -tagging requirement — can then have their shape compared to isolate the effect of the b -tagging algorithm.

It can be seen that the b -tagging algorithm has a slight preference towards lower values of τ_{32} . This could produce results which are artificially more “top-like”.

The solution is to propose an alternative cut which does not suffer these same drawbacks. The proposed alternative is a cut on the mass of the lepton- b system, $m_{\ell b}$, on the leptonic side of the event. Figure 5.4 shows the distribution of $m_{\ell b}$ prior to




Figure 5.3: Distribution of the τ_{32} substructure observable at detector-level, for the hadronic top candidate jet. Plotted are the distributions for which **no b -tagging requirement** is applied at detector-level and the subset of those jets which **do contain a b -tagged small- R jet** at detector-level. The area under both is normalised to unity to emphasise shape differences.




Figure 5.4: $m_{\ell b}$ distribution at detector-level prior to the final stage of selection.




Figure 5.5: Distribution of the τ_{32} substructure observable at detector-level. Histograms are shown prior to the application of the event selection, up to the application of the $m_{\ell b}$ or b -tagging requirement, and for different combinations of final requirements. The area under all distributions is normalised to unity to emphasise shape differences.

applying any final cut. That is, after the application of the event selection but without any b -tagging requirement on the hadronic top candidate jet or a cut on the $m_{\ell b}$.

One final benefit of the $m_{\ell b}$ cut is a reduction of the number of fakes in the event. For an event in which a fake lepton is present, there is no reason for there to be a b -jet sufficiently close to this fake to calculate the $m_{\ell b}$.

It was decided to apply an upper limit on the $m_{\ell b}$ of 120 GeV, in tandem with a mass window cut on the leading large jet, $110 \text{ GeV} < m_{\text{Large Jet}} < 230 \text{ GeV}$, in order to further reduce the contribution of the W . The application of the upper limit on the $m_{\ell b}$ at 120 GeV was a choice made to reproduce the 77% signal efficiency of the MV2c10 b -tagging algorithm, which was being used at a 77% working point. This choice was made for the purposes of comparison between the two approaches. Henceforth, references to the $m_{\ell b}$ cut include the mass window on the hadronic top candidate jet. Figure 5.5 compares the shape of the τ_{32} observable as affected by the b -tagging cut and the $m_{\ell b}$ cut.

Before settling on this approach, a comparison was made between the $m_{\ell b}$ cut and the conventional cut requiring a b -tag inside the large jet. Figure 5.6 compares the signal efficiency of both approaches as a function of the p_T of the hadronic top candidate jet. Figure 5.7 shows the background fraction for each approach as a function




Figure 5.6: The proportion of signal events passing selection, as a function of p_T , for the m_{lb} cut and the b -tagging cut.




Figure 5.7: The proportion of accepted events which are background, as a function of p_T , for the m_{lb} cut, the b -tagging cut, and when no final cut is applied.

of the p_T of the hadronic top candidate jet, and Figure 5.8 the ROC curve for the m_{lb} cut. This comparison study was done using the anti- k_t large- R LCTopo jets and MV2c10 b -tagging algorithm.




Figure 5.8: The ROC curve, showing background rejection against signal efficiency, for the $m_{\ell b}$ cut. Each point is an upper limit on the $m_{\ell b}$, between 100–140 GeV, for 1 GeV increments. The point at which the 77% signal efficiency is met is an upper limit of 120 GeV.

Another instructive comparison is made by considering the double ratio,

$$\text{Double Ratio} = \frac{\left(\frac{\text{Signal}}{\text{Total}}\right)_{\text{Passing Cut}}}{\left(\frac{\text{Signal}}{\text{Total}}\right)_{\text{Before Cut}}}, \quad (5.2)$$

which gives an indication of the signal-to-noise behaviour of the two approaches. It is constructed from the ratio of the signal fraction after the cut is applied to that before its application. In this way, a value above one indicates that the proportion of events which are signal is improved by the cut and a value below one indicates the opposite. This is shown in Figure 5.9.

From these comparisons, it can be seen that the $m_{\ell b}$ cut outperforms the previous approach in terms of raw efficiency but also by a significant margin at high- p_T . In the highest bin, there is a $\sim 17\%$ improvement given by the $m_{\ell b}$ cut, whereas the conventional approach is actually degrading the signal-to-noise ratio.

It is, of course, important to see that this new cut is not inadvertently introducing some new correlation, essentially a correlation between the substructure of the jet and the $m_{\ell b}$. It can be seen from figure 5.10 that this is not the case.




Figure 5.9: Double ratio computed for both b -tagging and $m_{\ell b}$ cuts as a function of the large jet p_T .




Figure 5.10: 2D histogram showing τ_{32} vs. $m_{\ell b}$. It can be seen that there is no correlation between the two observables.




Figure 5.11: The fraction of events coming from each Monte Carlo sample at detector-level.

As well as being interested in whether or not the background is reduced by this new cut, it is also important to know how the *composition* of the background is affected. Some backgrounds, notably the single top process, carry larger uncertainties than others. It is then desirable to reduce these as much as is possible.

Figure 5.11 makes clear that the $m_{\ell b}$ cut removes proportionally more background than the conventional approach requiring a b -tag inside the large jet. The $m_{\ell b}$ cut also removes a higher proportion of single top contribution.

5.3. Mass window cut

Following comparison of the shape of the distribution of the τ_{32} observable with that in the all-hadronic channel, which initially had a different mass window being applied to hadronic top candidate jets, a study was conducted to investigate the effect of the width of the window.

The window employed in the ℓ +jets channel required $110 \text{ GeV} < m_{\text{jet}} < 230 \text{ GeV}$, which was slightly more relaxed than that used in the all-hadronic channel — $122.5 \text{ GeV} < m_{\text{jet}} < 222.5 \text{ GeV}$. Figure 5.12 shows normalised distributions in the ℓ +jets channel for both windows.

It can be seen that the selection using the slightly tighter mass window produced distributions shifted slightly towards lower values of τ_{32} as compared to the more relaxed mass window. These lower values are more consistent with the expected signal




Figure 5.12: τ_{32} distributions in the ℓ +jets channel for two different mass window requirements on the hadronic top candidate jet. The area under each distribution is normalised to unity to emphasise shape differences.


shape. This would suggest that the tighter mass window excludes a higher proportion of the spurious contributions coming from hadronically decaying boosted W bosons.

As a result, the $122.5 \text{ GeV} < m_{\text{jet}} < 222.5 \text{ GeV}$ window was adopted as the standard in both ℓ +jets and all-hadronic channels.

5.4. Selection of hadronic top candidates

As introduced in Section 5.1, the process for identifying the hadronic top candidate in the ℓ +jets channel was initially to select the leading large- R reclustered jet in the event. This essentially encodes the assumption that, because the hadronic top jet is expected to have high p_T , the leading jet in the event will be the hadronic top jet. Although rare, it is possible that there is another high- p_T large- R reclustered jet in an event which does not originate from the hadronic top decay. If this jet was the leading jet, it would be identified as the hadronic top candidate and the substructure calculation in that event would use the wrong jet.

To investigate and quantify this effect, events were selected for which these definitions diverge. That is, events for which the leading large- R reclustered jet is *not* the jet with mass closest to the top mass, $m_t = 172.5 \text{ GeV}$. A relatively small fraction of events — around 1% — fall into this category. For these events, distributions are plotted of the substructure observable τ_{32} for each definition. These distributions are shown in Figure 5.13.



(a) Distribution of the substructure observable τ_{32} for the hadronic top candidate jet when this is defined as the large- R reclustered jet with mass closest to the top mass.

(b) Distribution of the substructure observable τ_{32} for the hadronic top candidate jet when this is defined as the leading large- R reclustered jet in the event.

Figure 5.13: Events for which the two possible methods for choosing the hadronic top candidate jet return different jets.

As defined in Section 3.1.5, τ_{32} is an observable which tends to lower values when a jet is well-described by a three-prong configuration, as would be expected for a hadronic top quark jet. The distribution in Figure 5.13a, with the hadronic top candidate being chosen as the large- R reclustered jet with mass closest to the top mass, has shape consistent with the expected signal. The distribution in Figure 5.13b, with the hadronic top candidate being chosen as the leading large- R reclustered jet in the event, however, has shape more consistent with the expected background.

For this reason, it was concluded that the former definition returned more reliably the large- R reclustered jet originating from the hadronically decaying top and this approach was subsequently adopted in the event selection procedure.




Figure 5.14: Leading-order single top production.

5.5. Background estimation

5.5.1. Monte Carlo backgrounds

With the exception of the fake lepton background — described below — all of the background processes present in this analysis are estimated using Monte Carlo simulation samples, described in Section 4.1.

The leading background in the analysis comes from single top production, shown in Figure 5.14. Whilst all three production processes are included in the estimate, the dominant mode is tW production, shown in Figure 5.14c. This is because the final state can be very similar to the ℓ +jets $t\bar{t}$ signature. Following the decay of the top quark into a W boson and a b -quark, if one W decays hadronically and the other leptonically, then the final state is identical to the signal with the exception of one additional b -jet in the $t\bar{t}$ process. This difference is what is exploited by the $m_{\ell b}$ cut to reduce this background. The selection criteria require that the lepton in the event be far away from the hadronic top candidate jet. In the case of a tW event, this is likely to mean that it is the W originating from the top which decays hadronically, and the other W in the event which decays leptonically. The ability to form an $m_{\ell b}$ system is contingent on the presence of a b -jet nearby to the lepton in the event. The ℓ +jets $t\bar{t}$ process will satisfy this requirement, but the tW process will not as the leptonic W decay does not have an associated b -jet.

After single top, the most important backgrounds come from the W +jets process and events containing a fake lepton. Other backgrounds considered are the Z +jets process, $t\bar{t}V$, where $V = Z/W/H$, and the diboson processes VV , with $V = Z/W$, which is grouped into the “other” category. The fraction of selected events which originate from these background processes is 6%.

5.5.2. Fake lepton estimate

Fakes, in the context of the analysis presented here, are events which pass the event selection procedure but do not contain an isolated lepton which comes from the decay of a boson. Leptons which do not originate from boson decay may be non-prompt, arising from hadronic decay, or resulting from a jet which is mis-reconstructed as an electron. These background processes are typically not particularly well simulated due to their dependence on detector performance and because many of these fakes arise from the multijet process. The nominal method adopted for estimating this background is then a data-driven one, the *matrix method* [202].

The matrix method proceeds by reweighting events selected using the same criteria as in the nominal selection procedure, but with a *loose* set of criteria on lepton definition. With respect to the reconstruction criteria described in Chapter 4, this loose selection removes the isolation requirements and relaxes the ID requirements. The event weights are calculated by combination of the lepton efficiencies for both real and fake leptons with the tight criteria. These efficiencies are extracted centrally by the ATLAS Top Reconstruction group from control regions enriched in either real or fake leptons.

Because these control regions are in a phase-space which differs from that of this analysis, with the efficiencies then parameterised by the p_T and $|\eta|$ of the lepton, a second, bespoke data-driven method for estimating the fake lepton background was also investigated.

Of the selection procedure applied in Chapter 5, it is the cuts on E_T^{miss} , $E_T^{\text{miss}} + m_T^W$, and $m_{\ell b}$ which reduce contributions from the fake lepton background. Control regions which are enriched in fakes are defined by inverting these cuts, requiring $E_T^{\text{miss}} < 20 \text{ GeV}$, $E_T^{\text{miss}} + m_T^W < 60 \text{ GeV}$, and $m_{\ell b} > 170 \text{ GeV}$.

A Monte Carlo template comprising boosted all-hadronic $t\bar{t}$ and W/Z +jets events is used as the basis for this alternative estimate of the fake lepton background. The strategy is to fit this template to the data in the fake-enriched control region, then apply the result in the signal region in order to produce an estimate of the fake lepton background.

First of all, the normalisation of the $t\bar{t}$ Monte Carlo sample in the signal region is extracted from a fit to the data, in order to correct for the over-prediction of the process in the simulation. This is done by first subtracting the sum of the background processes in the simulation from the data in order to leave the signal process. The

Distribution	CR	Fakes norm.	$t\bar{t}$ norm.	Fakes in SR
E_T^{miss}	$< 20\text{GeV}$	17.57 ± 2.84	0.894 ± 0.003	930.99
$E_T^{\text{miss}} + m_T^W$	$< 60\text{GeV}$	18.41 ± 1.90	0.908 ± 0.002	866.69
$m_{\ell b}$	$> 170\text{GeV}$	16.97 ± 4.85	0.865 ± 0.003	838.77

Table 5.3: Results of the fitting and scaling process for the fake template method.

resulting distribution is then fitted in the signal region using a function of the form

$$f(x) = p_0 + p_1 e^{-p_2 x} + p_3 x + p_4 (z - p_5 x), \quad (5.3)$$

where z represents the histogram corresponding to the signal process in the Monte Carlo, and x each of the distributions listed in Table 5.3. In this way, the parameter p_4 is used to scale the $t\bar{t}$ Monte Carlo to match the data and is extracted as the normalisation for that sample.

The $t\bar{t}$ signal is then scaled using the extracted normalisation factor, and then subtracted from the data. The data has then had both the sum of background processes and the signal processes subtracted. Without an estimate for the fake lepton background already included in the background processes, this is what should be left over in the data after this subtraction, particularly in the fake-enriched control regions.


It is this fake-enriched region of the reduced data which is then used to fit the normalisation for the fake template sample. This fit is performed in the control region using a function of the form

$$f(x) = p_0 z, \quad (5.4)$$

with z representing the histogram corresponding to the fake template sample. The parameter p_0 is then extracted and used to scale the sample beyond the control region and into the signal region.

Table 5.3 shows the results of the fitting procedure, with scalings for the $t\bar{t}$ signal process and fake template shown, as well as the resulting event yields in the signal regions.

Figures 5.15 and 5.16 show the components of the estimate and resulting sample for the template fit method, as compared to the matrix method, for E_T^{miss} and $E_T^{\text{miss}} + m_T^W$



(a) Control and signal regions for the template fit method. The data and signal Monte Carlo are shown in solid black and green, respectively. The fit to the data with backgrounds subtracted, used for the $t\bar{t}$ normalisation, is shown in red. The dashed black line shows the reduced data, with scaled signal and background processes subtracted. In magenta is shown the final scaled fake template sample. It is possible to see by the reduced data that there are a greater number of fake events in the control region.

(b) In blue is the matrix method, the nominal estimate for the fake lepton background. In magenta is the scaled fake template sample, and in black the reduced data, with the sum of scaled signal and background processes subtracted.


Figure 5.15: Detector-level E_T^{miss} , event counts divided by bin width.

distributions. Although scaled to the data, the template method is limited by the low number of events in the sample, and the resulting scaled distribution does not have a regular shape.

The final sample is obtained by scaling by the mean of the normalisation factors obtained from the three distributions. Figure 5.17 shows both the template fit and matrix methods for the τ_{32} and D_2 substructure observables. The shape derived from the matrix method is more in line with expected shape for background processes in these observables, but the template fit method is again limited by the underlying statistics before scaling.

Figure 5.18 shows the detector-level distribution for the τ_{32} substructure observable using each method for the estimate of the fake lepton background. The template fit method results in slightly under half the number of events as compared to the matrix method, with a slightly different shape to the distribution.

Owing to the limitations of the data-driven template fit method for estimating the fake lepton background, the matrix method was chosen as the nominal estimate for



(a) Control and signal regions for the template fit method. The data and signal Monte Carlo are shown in solid black and green, respectively. The fit to the data with backgrounds subtracted, used for the $t\bar{t}$ normalisation, is shown in red. The dashed black line shows the reduced data, with scaled signal and background processes subtracted. In magenta is shown the final scaled fake template sample. It is possible to see by the reduced data that there are a greater number of fake events in the control region.

(b) In blue is the matrix method, the nominal estimate for the fake lepton background. In magenta is the scaled fake template sample, and in black the reduced data, with the sum of scaled signal and background processes subtracted.

Figure 5.16: Detector-level $E_T^{miss} + m_T^W$, event counts divided by bin width.




Figure 5.17: Detector-level distributions for the fake lepton estimate in two of the measured substructure observables. In blue is the nominal estimate, with the fake template method shown in red. The number of events is divided by the bin width.




Figure 5.18: Detector-level distribution of the τ_{32} substructure observable, for different estimates of the fake lepton background. The number of events is divided by the bin width.

this background. An uncertainty is then applied which is the difference between the nominal estimate and the alternative method, the template fit. As is apparent from the comparison between methods shown in Figure 5.18, this uncertainty is relatively insignificant to the analyses owing to the small impact of the fake lepton background on the overall distributions.

Figure 5.19 shows the detector-level $|\eta|$ distribution of leptons in the ℓ +jets channel. It can be seen that in the region $1.0 < |\eta| < 1.5$, there is an increase in the contribution of the fake lepton background.

Figure 5.20 shows the event weights as a function of the lepton $|\eta|$ for the nominal fakes sample, with the number of events on the z-axis. With the sample divided into leptons identified as electrons and those identified as muons, it can be seen that there is a group of events containing a fake muon which have high event weights in the region identified as having a high contribution from the fake background, $1.0 < |\eta| < 1.5$.

Figure 5.21 shows the same $|\eta|$ distribution as Figure 5.19, but for only the nominal fake lepton sample, separated into electron and muon channels. It can be seen that the increase in the number of events in the region covering $1.0 < |\eta| < 1.5$ is coming from the muons, as was suggested by the cluster of high event weights. When then cutting out events with high weights (above 1.5), the striking peak in the muon distribution is






Figure 5.19: Detector-level distribution of lepton $|\eta|$ in the ℓ +jets channel. An increase in the contribution from the fake lepton background estimation can be seen in $1.0 < |\eta| < 1.5$.



(a) Electrons



(b) Muons

Figure 5.20: Event weights as a function of lepton $|\eta|$ for events in the nominal fakes sample.




Figure 5.21: For both electrons and muons, distribution of events in the nominal fakes sample as a function of $|\eta|$.

removed. It is then concluded that the source of the peak in this region is indeed the high weights in this region for the muons.

Figure 5.22 shows the $|\eta|$ distribution for the nominal fake lepton sample, split into electron and muon channels. Here, an uncertainty band is added which is calculated as the difference in that bin to the fake lepton sample derived from the alternative method, as is done in the calculation of the systematics for the ℓ +jets channel. From this band it is concluded that the size of the uncertainty applied in the analysis is sufficiently large as to absorb the effect of the high event weights for the muons in the $1.0 < |\eta| < 1.5$ region for the nominal fake lepton sample. It should also be noted that the scale of this background relative to other contributions in Figure 5.19 is exaggerated by the logarithmic y -axis.

5.6. Data/MC comparison

As seen in Table 5.1, the event yield given by the sum of $t\bar{t}$ prediction and all backgrounds is 14% higher than that in the data. This effect is not new and has been observed in boosted $t\bar{t}$ in several previous analyses [2, 142, 143, 149, 203–205]. This overprediction in the calculation of the $t\bar{t}$ cross-section can be attributed to missing higher-order contributions in the NLO calculation as a function of the p_T of the top. To correct for this, the $t\bar{t}$ prediction is normalised to the data by a simple scaling factor.




Figure 5.22: Number of events in the nominal fakes sample as a function of $|\eta|$, with uncertainty calculated as the difference to the fakes sample derived from the alternative method.

This factor is calculated as

$$\text{SF} = \frac{N_{\text{data}} - N_{\text{backgrounds}}}{N_{\text{signal}}}, \quad (5.5)$$

with N the number of events in that sample, in the case of $N_{\text{backgrounds}}$ the sum of the background processes. This results in a scale factor of 0.85 ± 0.03 , which is applied to the $t\bar{t}$ prediction.

Figures 5.23–5.30 provide a comparison between the Monte Carlo predictions and the data for a range of distributions focused on the kinematics of various parts of the system. These are not the distributions which are being measured in the analysis but are intended to provide an indication of how well the system is being described in simulation. In general, there is good agreement between the Monte Carlo and the data, with the data lying within the detector and background uncertainties.




Figure 5.23: Number of events as a function of detector-level final state control observables in the ℓ +jets channel. Both the nominal MC predictions for signal and background and the data are shown, with the total prediction being normalised to the number of events present in the data. The lower pad shows the ratio between the data and total prediction. The hatched band shows the sum of statistical and detector uncertainties.




Figure 5.24: Number of events as a function of detector-level final state control observables in the ℓ +jets channel. Both the nominal MC predictions for signal and background and the data are shown, with the total prediction being normalised to the number of events present in the data. The lower pad shows the ratio between the data and total prediction. The hatched band shows the sum of statistical and detector uncertainties.




Figure 5.25: Number of events as a function of detector-level final state control observables in the ℓ +jets channel. Both the nominal MC predictions for signal and background and the data are shown, with the total prediction being normalised to the number of events present in the data. The lower pad shows the ratio between the data and total prediction. The hatched band shows the sum of statistical and detector uncertainties.




Figure 5.26: Number of events as a function of detector-level final state control observables in the ℓ +jets channel. Both the nominal MC predictions for signal and background and the data are shown, with the total prediction being normalised to the number of events present in the data. The lower pad shows the ratio between the data and total prediction. The hatched band shows the sum of statistical and detector uncertainties.




Figure 5.27: Number of events as a function of detector-level final state control observables in the ℓ +jets channel. Both the nominal MC predictions for signal and background and the data are shown, with the total prediction being normalised to the number of events present in the data. The lower pad shows the ratio between the data and total prediction. The hatched band shows the sum of statistical and detector uncertainties.




Figure 5.28: Number of events as a function of detector-level final state control observables in the ℓ +jets channel. Both the nominal MC predictions for signal and background and the data are shown, with the total prediction being normalised to the number of events present in the data. The lower pad shows the ratio between the data and total prediction. The hatched band shows the sum of statistical and detector uncertainties.




Figure 5.29: Number of events as a function of detector-level final state control observables in the ℓ +jets channel. Both the nominal MC predictions for signal and background and the data are shown, with the total prediction being normalised to the number of events present in the data. The lower pad shows the ratio between the data and total prediction. The hatched band shows the sum of statistical and detector uncertainties.


(a) Lepton p_T .

Figure 5.30: Number of events as a function of detector-level final state control observables in the ℓ +jets channel. Both the nominal MC predictions for signal and background and the data are shown, with the total prediction being normalised to the number of events present in the data. The lower pad shows the ratio between the data and total prediction. The hatched band shows the sum of statistical and detector uncertainties.

Chapter 6.

Substructure Observables and Unfolding

“ Briefly summarised, what I did can be described as simply an act of desperation.”

— Max Planck

Detailed in this chapter are the studies undertaken in order to select the final set of substructure observables to be measured in the analysis. The criteria by which the observables included in the analysis results were chosen is discussed, along with studies conducted to investigate these criteria. Finally, the unfolding procedure for the analysis is outlined.

6.1. Selection of observables

As detailed in Chapter 4, substructure observables are defined by associating tracks to the hadronic top candidate jet at detector-level, or the charged particles at particle-level. It is the small- R jets from which the large- R jets are reclustered that have the tracks associated to them.

Table 6.1 summarises the families of observables defined in Section 3.1.5. From these, a range were chosen and from this range, a subset was selected for measurement. Some observables were included in the study in more than one form. In particular,

Observable	Motivation	Reference(s)
N -Subjettiness Ratios	Tagging, sensitive to FSR/ISR variations, sensitive to generator variation	[93, 122, 154]
Energy-Correlation Functions	Tagging, sensitive to FSR variations, sensitive to generator variation	[93, 122, 127, 154]
ECF observables	Tagging, sensitive to FSR variations, sensitive to generator variation	[93, 122, 127, 154]
k_t Splitting Functions	Sensitive to FSR/ISR variations, sensitive to generator variation	[154, 206–208]
Eccentricity	Sensitive to FSR variations	[122, 209]
Generalised Angularities	Sensitive to FSR/ISR variations, sensitive to generator variation	[122, 126]
Q_W	Sensitive to FSR/ISR variations, sensitive to generator variation	[139, 207, 210]

Table 6.1: Categories of substructure observables which were considered for inclusion in the measurement, along with motivations for each. Also provided are references to their definitions and previous measurement.

multiple versions of the ECF observables were considered, using different values for the angular weighting parameter, β .

In order to make the final selection, observables were evaluated with reference to the following criteria:

- Whether the observable is employed in ATLAS tagging algorithms.
- How well the observable is described by the nominal Monte Carlo prediction.
- How sensitive the observable is to different Monte Carlo predictions.
- How sensitive the observable is to systematic variations of the nominal Monte Carlo prediction.


- The correlations between candidate observables. In the case that observables are highly correlated with each other, preference is given to those with lower uncertainties and superior resolution.

Figures 6.1–6.7 provide comparisons between the data and nominal Monte Carlo prediction for the eight observables selected for measurement, as well as comparisons between the nominal Monte Carlo prediction and alternative predictions. The comparisons between the data and the nominal prediction provide an indication for how well described the observables are by that prediction, as well as exhibiting the uncertainties and resolution for each. The particle-level comparisons to the nominal prediction show which observables are sensitive to different predictions as well as different tunings of the nominal prediction. This is quantified by a χ^2 calculation relative to the nominal prediction, using only the statistical uncertainty. Figure 6.8 summarises the χ^2 per degrees of freedom obtained from the particle-level study for all observables considered.

The comparisons between the data and the nominal Monte Carlo prediction across all observables show that shapes are generally quite well described, with some exceptions. Observables designed to probe two-prong structure within a jet — and therefore employed in many such taggers — such as D_2 , the related C_2 , and τ_{21} , are not well described. This is especially at the lower end of the distribution, the region considered to be more two-prong like. This is also true of τ_{32} , which is designed to probe three-prong structure. The p_T -dispersion, $p_T^{d,*}$, is also poorly described.

There are some observables which experience large detector uncertainties in some regions, such as the jet constituent multiplicity. There are others which have poor resolution and as a result are limited to a small number of bins in their distribution, such as $\text{ECF4}(\beta = 0)$.


The particle-level comparisons to the nominal prediction for each observable show that, in general, the observables are most sensitive to the use of different predictions. For many observables, the greatest χ^2 value is seen between the nominal prediction and Powheg+Herwig7. Many also exhibit significant sensitivity to the different FSR tunes of the nominal prediction. Most observables show less sensitivity to the different ISR tunes, though the LHA, ECF2 , d_{12} , and τ_{21} are those which exhibit the most. The $\beta = 0$ variants of the ECFs and their derived observables, as well as several of the N -subjettiness observables show the greatest sensitivity to the different FSR tunes.



(a) Distributions at particle-level for the $t\bar{t}$ signal Monte Carlo prediction. The nominal prediction, Powheg+Pythia8, is shown in black. Variations of the FSR and ISR are shown in red and blue, respectively, and an alternative prediction made using Powheg+Herwig7 is shown in green. For each prediction, the χ^2 per degree of freedom is calculated relative to the nominal prediction. The ratio of each prediction to the nominal prediction is shown below the distribution. The area under each prediction is normalised to unity to emphasise shape differences.

(b) Distributions at detector-level for signal and background Monte Carlo predictions, and for the data. The hatched band represents the combination of statistical and detector-related uncertainties on the Monte Carlo predictions. The sum of the predictions is normalised to the data. The ratio of the data to the Monte Carlo prediction is shown in the region below the distribution.


Figure 6.1: Number of events as a function of the τ_{32} substructure observable as measured for the hadronic top candidate jet in the ℓ +jets channel.



(a) Distributions at particle-level for the $t\bar{t}$ signal Monte Carlo prediction. The nominal prediction, Powheg+Pythia8, is shown in black. Variations of the FSR and ISR are shown in red and blue, respectively, and an alternative prediction made using Powheg+Herwig7 is shown in green. For each prediction, the χ^2 per degree of freedom is calculated relative to the nominal prediction. The ratio of each prediction to the nominal prediction is shown below the distribution. The area under each prediction is normalised to unity to emphasise shape differences.

(b) Distributions at detector-level for signal and background Monte Carlo predictions, and for the data. The hatched band represents the combination of statistical and detector-related uncertainties on the Monte Carlo predictions. The sum of the predictions is normalised to the data. The ratio of the data to the Monte Carlo prediction is shown in the region below the distribution.


Figure 6.2: Number of events as a function of the τ_{21} substructure observable as measured for the hadronic top candidate jet in the ℓ +jets channel.



(a) Distributions at particle-level for the $t\bar{t}$ signal Monte Carlo prediction. The nominal prediction, Powheg+Pythia8, is shown in black. Variations of the FSR and ISR are shown in red and blue, respectively, and an alternative prediction made using Powheg+Herwig7 is shown in green. For each prediction, the χ^2 per degree of freedom is calculated relative to the nominal prediction. The ratio of each prediction to the nominal prediction is shown below the distribution. The area under each prediction is normalised to unity to emphasise shape differences.

(b) Distributions at detector-level for signal and background Monte Carlo predictions, and for the data. The hatched band represents the combination of statistical and detector-related uncertainties on the Monte Carlo predictions. The sum of the predictions is normalised to the data. The ratio of the data to the Monte Carlo prediction is shown in the region below the distribution.


Figure 6.3: Number of events as a function of the ECF2 substructure observable as measured for the hadronic top candidate jet in the ℓ +jets channel.



(a) Distributions at particle-level for the $t\bar{t}$ signal Monte Carlo prediction. The nominal prediction, Powheg+Pythia8, is shown in black. Variations of the FSR and ISR are shown in red and blue, respectively, and an alternative prediction made using Powheg+Herwig7 is shown in green. For each prediction, the χ^2 per degree of freedom is calculated relative to the nominal prediction. The ratio of each prediction to the nominal prediction is shown below the distribution. The area under each prediction is normalised to unity to emphasise shape differences.

(b) Distributions at detector-level for signal and background Monte Carlo predictions, and for the data. The hatched band represents the combination of statistical and detector-related uncertainties on the Monte Carlo predictions. The sum of the predictions is normalised to the data. The ratio of the data to the Monte Carlo prediction is shown in the region below the distribution.


Figure 6.4: Number of events as a function of the D_2 substructure observable as measured for the hadronic top candidate jet in the ℓ +jets channel.



(a) Distributions at particle-level for the $t\bar{t}$ signal Monte Carlo prediction. The nominal prediction, Powheg+Pythia8, is shown in black. Variations of the FSR and ISR are shown in red and blue, respectively, and an alternative prediction made using Powheg+Herwig7 is shown in green. For each prediction, the χ^2 per degree of freedom is calculated relative to the nominal prediction. The ratio of each prediction to the nominal prediction is shown below the distribution. The area under each prediction is normalised to unity to emphasise shape differences.

(b) Distributions at detector-level for signal and background Monte Carlo predictions, and for the data. The hatched band represents the combination of statistical and detector-related uncertainties on the Monte Carlo predictions. The sum of the predictions is normalised to the data. The ratio of the data to the Monte Carlo prediction is shown in the region below the distribution.


Figure 6.5: Number of events as a function of the C_3 substructure observable as measured for the hadronic top candidate jet in the ℓ +jets channel.



(a) Distributions at particle-level for the $t\bar{t}$ signal Monte Carlo prediction. The nominal prediction, Powheg+Pythia8, is shown in black. Variations of the FSR and ISR are shown in red and blue, respectively, and an alternative prediction made using Powheg+Herwig7 is shown in green. For each prediction, the χ^2 per degree of freedom is calculated relative to the nominal prediction. The ratio of each prediction to the nominal prediction is shown below the distribution. The area under each prediction is normalised to unity to emphasise shape differences.

(b) Distributions at detector-level for signal and background Monte Carlo predictions, and for the data. The hatched band represents the combination of statistical and detector-related uncertainties on the Monte Carlo predictions. The sum of the predictions is normalised to the data. The ratio of the data to the Monte Carlo prediction is shown in the region below the distribution.

Figure 6.6: Number of events as a function of the $p_T^{d,*}$ substructure observable as measured for the hadronic top candidate jet in the ℓ +jets channel.



(a) Distributions at particle-level for the $t\bar{t}$ signal Monte Carlo prediction. The nominal prediction, Powheg+Pythia8, is shown in black. Variations of the FSR and ISR are shown in red and blue, respectively, and an alternative prediction made using Powheg+Herwig7 is shown in green. For each prediction, the χ^2 per degree of freedom is calculated relative to the nominal prediction. The ratio of each prediction to the nominal prediction is shown below the distribution. The area under each prediction is normalised to unity to emphasise shape differences.

(b) Distributions at detector-level for signal and background Monte Carlo predictions, and for the data. The hatched band represents the combination of statistical and detector-related uncertainties on the Monte Carlo predictions. The sum of the predictions is normalised to the data. The ratio of the data to the Monte Carlo prediction is shown in the region below the distribution.

Figure 6.7: Number of events as a function of the LHA substructure observable as measured for the hadronic top candidate jet in the ℓ +jets channel.




Figure 6.8: χ^2 per degree of freedom for each considered substructure, relative to the nominal Monte Carlo prediction at particle-level. Comparisons are made to the alternative Monte Carlo prediction — Powheg+Herwig7 — and to variations of the FSR and ISR within the nominal prediction. The χ^2 calculation considers only statistical uncertainties.

Because the $\beta = 0$ variants of the ECFs and their derived observables are not IRC safe [127], as well as their large χ^2 , often large uncertainties, and an observed anti-correlation with $p_T^{d,*}$, it was decided that β should be set to unity for the measurement.

τ_{32} and D_2 are some of the most-used observables in tagging algorithms. Because of this, they were chosen to be measured double-differentially, in windows of the hadronic top candidate jet p_T and mass.

Table 6.2 summarises the observables selected for measurement and the reasons for their selection.

6.2. Unfolding

Theoretical predictions are, of course, made independently of the effects of interaction with any particular detector. Measurements made by experiments such as ATLAS are therefore limited by the difficulty of using those results for comparisons both with those of other experiments and with new theories. There are then two options to improve this situation.

First, each new theory could be passed through the detector simulation as happens when analyses are first performed. There are several issues with this approach. It does not solve the issue with comparing results measured at different experiments. The detector simulation procedure is also very computationally expensive and much of its inner workings are available only internally within the collaboration, so that a theorist with a theory to test against a given measurement would have to request that ATLAS spend CPU hours to produce the comparison. It's possible that a theory may require comparison against multiple measurements, which may have been made with differing detector configurations and thus require multiple independent simulations. All of this is at best cumbersome, and more realistically, not possible.

The second approach is to take the inverse approach and attempt to remove the effects of interaction with the detector and thus return the data itself to particle-level. The resultant detector-independent measurements would then be valid for comparison both against new and improved predictions and across experiments. This process of removing detector effects is called *unfolding*. This comes from the fact that this is the inverse of the process of *folding* the effects of detector interaction into the true

Observable	Reason for selection
τ_{32}	• Poorly described by nominal Monte Carlo Prediction • Sensitive to variations in FSR • Used by top-tagging algorithms • Low correlation with other observables
τ_{21}	• Used by top-tagging algorithms • Moderate correlation with other observables
τ_3	• Poorly described by nominal Monte Carlo prediction • Sensitive to variations in FSR • Highly correlated with τ_4 but with superior resolution
ECF2 ($\beta = 1$)	• High resolution • Low uncertainties • Low correlation with other observables
D_2 ($\beta = 1$)	• Used in two-prong and top tagging algorithms • Sensitive to variations in FSR • Highly correlated with τ_4 but with superior resolution
C_3 ($\beta = 1$)	• Poorly described by nominal Monte Carlo Prediction • Sensitive to variations in FSR • Low correlation with other observables
$p_T^{d,*}$	• Poorly described by nominal Monte Carlo Prediction • Low correlation with other observables
LHA	• Poorly described by nominal Monte Carlo Prediction • Low correlation with other observables • Sensitive to variations in ISR

Table 6.2: Substructure observables chosen for measurement in the analysis, with the reasons for their selection.

underlying distributions. This forward-folding process includes both systematic and stochastic effects, with many of these having some correction applied already in the process of producing the reconstructed objects and resultant observables. This process can be seen by [211]:

$$\int_{\Omega} K(d, t) f(t) dt + b(d) = g(d), \quad (6.1)$$

with $K(d, t)$ the *kernel* representing the detector effects, $f(t)$ the truth level distribution, $b(d)$ the background distribution and $g(d)$ the detector-level distribution. The set of detector effects are then convoluted with the underlying distribution to produce what's observed when reconstructed. Unfolding represents an inversion of this process.

The unfolding process addresses those residual detector biases which are not so well understood or easily corrected for. Because the forward-folding process is stochastic, meaning that a given particle-level distribution does not have a straightforward one-to-one mapping to a unique detector-level distribution, the unfolding process is inherently probabilistic and is categorised as an *ill-posed* problem in that the many-to-one nature of the mapping encoded in the folding function means its inversion is not straightforward.

6.2.1. Iterative Bayesian Unfolding

The most common approach taken in tackling this problem in ATLAS — and the one taken in the analysis presented in this thesis — is a regularised unfolding method put forward by D'Agostini [212], which has come to be known as iterative Bayesian unfolding (IBU), implemented via the RooUnfold software package [213].

The folding process for a given observable is encoded in its *response matrix*. The response matrix $\mathcal{R}_{ij}$ accounts for limited detector resolution and represents the probability for an event which is produced in bin i at particle-level to be reconstructed in bin j at detector-level. That is to say that the smearing induced by the detector resolution may cause the value of a given observable to *migrate* from one bin at particle-level to another at detector-level, with the degree of migration visible by the diagonality of the matrix. Here, the *migration matrix*, $\mathcal{M}_{ij}$, is used in place of the response matrix, where the former is simply a row-normalised version of the latter, such that each bin represents the fraction of events in a given particle-level bin which are reconstructed




Figure 6.9: Row-normalised migration matrices for measured substructure observables in the ℓ +jets channel.




Figure 6.10: Row-normalised migration matrices for measured substructure observables in the ℓ +jets channel.

in a given detector-level bin. Figures 6.9–6.12 show the migration matrices for the substructure observables selected for measurement.

Given that $\mathcal{M}_{ij}$ encodes the folding process, from particle-level to detector-level, the unfolding process is governed by the inverse, $\mathcal{M}_{ij}^{-1}$. The inversion of this matrix is




Figure 6.11: Row-normalised migration matrices for measured substructure observables in the ℓ +jets channel.




Figure 6.12: Row-normalised migration matrices for measured substructure observables in the ℓ +jets channel.

not straightforward and can result in large fluctuations in the unfolded distribution. The *Bayesian* part of IBU refers to the method used to approximate the matrix inversion in a regularised manner through the use of Bayes' theorem, given by

$$P(A|B) = \frac{P(B|A) \cdot P(A)}{P(B)}, \quad (6.2)$$

which is expressed in terms of conditional probabilities. $P(A|B)$, the posterior probability of A given B , is the probability of observing A , given that B is true. Similarly $P(B|A)$ gives the probability of observing B , given that A is true. $P(A)$ and $P(B)$ are the prior and marginal probabilities, respectively, which give the probability that A or B are true without any further knowledge, so without condition.

Using the *law of total probability*, which states that for a given sample space including both A and B , partitioned into $\{A_j\}$, then $P(B)$ can be expressed as

$$P(B) = \sum_j P(B|A_j) \cdot P(A_j). \quad (6.3)$$

Expressing this in the context of the unfolding discussion here gives

$$P(T_i|R_j) = \frac{P(R_j|T_i) \cdot P(T_i)}{\sum_{k=1}^{n_T} P(R_j|T_k) \cdot P(T_k)}, \quad (6.4)$$

with $P(T_i|R_j)$ the probability of observing an event in truth bin T_i given that an event is observed to be reconstructed in bin j at detector-level. Similarly, $P(R_j|T_i)$ is the probability of observing an event to be reconstructed in bin j at detector-level given that it was generated in bin i at particle-level. $P(T_i)$ is the prior probability of an event in bin i at particle-level. It can be seen then that $P(R_j|T_i)$ represents the forward-folding process and can be replaced by $\mathcal{M}_{ij}$, and $P(T_i|R_j)$ is the unfolding mechanism which the method aims to provide, so can be replaced with $\mathcal{M}_{ij}^{-1}$:

$$\mathcal{M}_{ij}^{-1} = \frac{\mathcal{M}_{ij} \cdot P(T_i)}{\sum_{k=1}^{n_T} \mathcal{M}_{ik} \cdot P(T_k)}. \quad (6.5)$$

In addition to the unfolding procedure governed by $\mathcal{M}_{ij}^{-1}$, there are certain corrections which must be applied to address some of the shortcomings of the detector




Figure 6.13: Efficiency and acceptance corrections for substructure observables.

reconstruction. The first of these is the efficiency correction,


$$f_j^{\text{eff}} = \frac{N_j^{\text{detector} \wedge \text{particle}}}{N_j^{\text{particle}}}, \quad (6.6)$$

which accounts for the limited reconstruction efficiency of the detector by consideration of events which pass the selection criteria at particle-level but not at detector-level. $N_j^{\text{detector} \wedge \text{particle}}$ is the number of events in bin j which pass both detector *and* particle-level, and N_j^{particle} the number of events passing the particle-level selection.

The acceptance correction, given by

$$f_{\text{acc}}^i = \frac{N_i^{\text{detector} \wedge \text{particle}}}{N_i^{\text{detector}}}, \quad (6.7)$$

parameterises the effect of events which originate outside of the fiducial region being reconstructed inside it, and is given by the ratio of events passing both detector- and particle-level selections to the number passing the detector-level selections. The distributions of efficiency and acceptance corrections are shown in Figures 6.13 and 6.14 for the τ_{32} and D_2 substructure observables' 1D and 2D distributions, respectively.



(a) τ_{32} split into regions of the p_T of the hadronic top jet.

(b) D_2 split into regions of the mass of the hadronic top jet.

Figure 6.14: Efficiency and acceptance corrections for double-differential measurement. The x -axis represents the concatenation of bins from each region, as per Equation 6.11.

$\mathcal{M}_{ij}^{-1}$, f_j^{eff} , and f_{acc}^i are all evaluated using the nominal $t\bar{t}$ MC prediction. Bringing together the unfolding procedure and corrections, an expression is constructed which gives the number of events in bin j of the distribution resulting from the unfolding of the detector-level distribution to particle-level:

$$N_j^{\text{unfolded}} = \frac{1}{f_j^{\text{eff}}} \sum_i \mathcal{M}_{ij}^{-1} f_{\text{acc}}^i (N_i^{\text{data}} - N_i^{\text{bkg}}). \quad (6.8)$$

Here N_i^{data} and N_i^{bkg} are the number of events in bin i of the distribution in the data and the sum of the background MC samples, respectively.

Given that the cross-section is related to the number of events by

$$N = \sigma \mathcal{L}, \quad (6.9)$$

where $\mathcal{L} = \int L(t) dt$ is the integrated luminosity, the expression for the unfolded differential cross-section is then

$$\frac{d\sigma_j}{dX_j} = \frac{N_j^{\text{unfolded}}}{\mathcal{L} \cdot \Delta X_j}, \quad (6.10)$$

for an observable X with bin-width ΔX_j in bin j .

Having constructed a description of the unfolded differential cross-section, now the *iterative* aspect of IBU comes into play. In Equation 6.5, which is contained within Equation 6.10, the prior is taken in the first instance to be particle-level Monte Carlo produced by the nominal prediction. This induces a bias towards the MC simulation of the SM, which is of course known to be an imperfect description of nature. This bias is reduced by performing the unfolding procedure iteratively, where each iteration performed after the first uses the output of Equation 6.8 for the previous iteration as the prior.

Equation 6.10 may also be used to provide unfolded double-differential distributions. In order to pass these distributions through the unfolding, these 2D histograms are transformed into 1D histograms. This is done by constructing, for an observable X , a new observable X' which concatenates the full distribution in X for each bin of the second observable, Y , as:

$$X' = X + i_Y \times \Delta X, \quad (6.11)$$

with i_Y the bin index for observable Y and ΔX the range of the observable X . In the context of the analysis presented here, X — the internal variables — are the substructure observables and Y — the external variables — are either the jet mass or p_T .

6.2.2. Optimisation

There are ultimately two parameters which must be optimised in the IBU procedure described here. These are the binning of the histograms to be unfolded and the number of iterations used to produce the final distributions.

Binning Because IBU acts on histograms, it is a bin-dependent procedure, meaning that the choice of binning for each of the unfolded observables has an impact on the performance of the unfolding.

The aim of the optimisation process for the binning of the histograms is to reduce biases which may be induced by the unfolding procedure whilst simultaneously minimising the statistical uncertainty across the resultant distributions. In order to

attain ideal performance and stability of the unfolding, two criteria for the binning optimisation were used: firstly that each particle-level bin must have 40% of events reconstructed at detector-level appear in the same bin; secondly that the statistical uncertainty in each bin on the diagonal of the migration matrix should not exceed 5%. These requirements represent the first step of the optimisation procedure, with manual fine-adjustment being applied according to fluctuations in modelling uncertainties and the results of the pull and stress tests described in Section 6.2.3. The study presented in Section 4.4 on the use of tracks to define the substructure was partially motivated by the improved resolution at this 40% diagonality requirement.

Number of iterations The number of iterations used in the unfolding procedure also has an impact on the outcome. In short, fewer iterations produce stronger regularisation manifested in smaller statistical fluctuations, but this comes at the cost of bias towards the prior used for the first iteration.

The optimisation of the number of iterations used, N_{iter} , is performed by minimisation of the average correlation factor, ρ_{avg} , for each observable [214]:

$$\rho_j^x = \sqrt{1 - \left(V_{jj}^x \cdot (V^{-1})_{jj}^x \right)^{-1}}, \quad (6.12)$$

$$\rho_{\text{avg}}^x = \frac{1}{N_{\text{bins}}^x} \sum_{j=1}^{N_{\text{bins}}^x} \rho_j^x. \quad (6.13)$$

Here, j is the bin number for observable x , N_{bins}^x the number of bins for that observable's distribution, and V^x the statistical covariance matrix. This matrix is evaluated using 1000 *pseudo-experiments*. These are constructed by smearing each bin in the detector-level MC distribution using a Poisson distribution. The bin content in that bin is used to set the mean of the Poisson distribution and the width is set to the statistical uncertainty in that bin, with the nominal corrections then used to unfold each of the pseudo-experiments.

As the unfolding procedure is iterated, the optimum number of iterations for a given observable is chosen to be that which minimises ρ_{avg}^x , indicating the smallest migration effects. For the observables measured in this analysis, the results of this optimisation are shown in Table 6.3, whilst the evolution of ρ_{avg}^x through the iteration

Observable	Number of Iterations (N_{iter})
τ_{32}	7
τ_3	4
τ_{21}	7
ECF2	7
D_2	6
C_3	10
p_T -dispersion	8
Les Houches Angularity	7

Table 6.3: Optimised number of iterations of the unfolding mechanism for each of the substructure observables measured in this analysis, as determined by the minimisation of the average correlation factor. The final choice for the number of iterations also accounts for the growth of statistical uncertainties and the χ^2 between iterations.





Figure 6.15: Evolution of the average correlation factor ρ_{avg}^x through the iteration of the unfolding procedure.

can be seen in Figures 6.15 and 6.16 for τ_{32} and D_2 in single- and double-differential distributions, respectively.

Often, the approach in ATLAS has been to use four iterations for IBU, without bespoke optimisation per analysis and observable as performed here. Similarly, the requirement on the fraction of events on the diagonal of the migration matrices has frequently been chosen to be 60%, rather than the 40% used here. The fact that the optimal number of iterations by the minimisation of ρ_{avg}^x tends to be higher than this “standard” number is related to the diagonality criteria.



(a) τ_{32} split into regions of the p_T of the hadronic top jet.

(b) D_2 split into regions of the mass of the hadronic top jet.

Figure 6.16: Evolution of the average correlation factor ρ_{avg}^x through the iteration of the unfolding procedure for double-differential distribution of substructure observables.

In order further to assess the behaviour of the unfolding mechanism as a function of the number of iterations used, both the stability and the impact on statistical uncertainty were investigated.

A χ^2 comparison is used to check how the unfolding changes between iterations. For each iteration, the χ^2 between that iteration and the previous one is calculated, with the expectation that over time this falls as the bias from the prior is reduced. The results of this test are shown in Figures 6.17 and 6.18 for τ_{32} and D_2 in single- and double-differential distributions, respectively. As can be seen, the χ^2 does indeed reduce as the number of unfolding iterations is increased. More specifically, the rate of change of this difference is seen to decrease after a few iterations, with the steepest drop to begin with. This is especially evident given the logarithmic scale of these plots. This would suggest that the choices for the number of iterations in the analysis are valid, given that the reduction of the χ^2 is in general very small after around five iterations.

Although the bias in the unfolding is reduced with iteration, it is expected that statistical uncertainties grow with the number of iterations, as fluctuations become amplified. For this reason, that statistical uncertainty in each bin of each distribution for each observable for each iteration is also investigated. For a reasonably central bin in the distribution, this behaviour is shown in Figure 6.19 for D_2 and the Les Houches Angularity. The expected increase with N_{iter} is observed, but the rate at which this increases does not taper off in the same way as seen for the χ^2 between iterations, with lower numbers of iterations favoured. The number of iterations resulting from




Figure 6.17: Evolution of the χ^2 between iterations of the unfolding procedure for distributions of substructure observables.




Figure 6.18: Evolution of the χ^2 between iterations of the unfolding procedure for double-differential distributions of substructure observables.




Figure 6.19: Evolution of the statistical uncertainty in a given bin as a function of the number of iterations of the unfolding procedure.

the optimisation of the average correlation factor, reported in Table 6.3, for some observables rose to over 7, a relatively high number. On consideration of the growth of statistical uncertainties, the number of iterations used for each observable was set to 6. This provides a balance between reducing the bias induced by the regularisation whilst constraining statistical fluctuations.

6.2.3. Validation

Following the optimisation of the parameters used in the unfolding procedure, a series of tests are implemented in order to check that the process is not overly sensitive to statistical fluctuations and that biases induced by the mechanism are under control.

Closure test The most basic test of the unfolding is the *closure* test, designed to verify that the unfolding procedure can indeed recover an underlying particle-level distribution for unseen events.

The nominal MC prediction, Powheg+Pythia8, is used for this test. This sample is split into two halves denoted half0 and half1. With half0 used as pseudo-data to be unfolded, half1 is used to derive the corrections used by the IBU mechanism. The particle-level distributions for both half0 and half1 can then be compared to the unfolded pseudo-data. If the unfolding is behaving as expected, the particle-level MC should be within the statistical uncertainties for the unfolded pseudo-data.

For each particle-level half of the sample, the closure or otherwise with the pseudo-data may be determined by a χ^2 test, with

$$\chi_{\text{closure}}^2 = V^T C^{-1} V, \quad (6.14)$$

where V contains the residuals between the unfolded pseudo-data and the particle-level MC and C is a covariance matrix which is constructed from the sum of covariance matrices for both half0 and half1.

Figure 6.20 shows the closure test for the τ_{32} and D_2 substructure observables, with reasonable closure for both. Table 6.4 provides the χ^2 values and corresponding p -values for each observable.




Figure 6.20: Relative differential cross-section for substructure observables at particle-level. half0 and half1 represent the nominal particle-level MC prediction split into two equal halves. U_{half0} corresponds to sample half0 after it has been unfolded using corrections derived from half1.

Observable	χ^2/ndf	$p\text{-value}$
τ_{32}	17/13	0.20
τ_3	8.3/12	0.76
τ_{21}	18.25/15	0.25
ECF2	22.6/20	0.31
D_2	19/17	0.32
C_3	21.6/14	0.09
$p_T\text{-dispersion}$	19.0/13	0.12
Les Houches Angularity	18.5/18	0.42

Table 6.4: Results of the closure test for each measured observable. The provided χ^2 and p -values quantify agreement between the unfolded pseudo-data and the particle-level half0, with half1 having been used to derive the corrections.

Pull test The next statistical test is the pull test, designed to assess biases in the reconstructed bin values accounting for statistical fluctuation. The result of the pull test can then be used as an indication for the validity of the estimation of the statistical uncertainty.

As with the iteration optimisation using ρ_{avg}^x , 1000 pseudo-experiments are generated, where for each bin at detector-level a Poisson distribution is used to smear that distribution. For each observable, all of these pseudo-experiments are passed through the unfolding procedure using the nominal corrections. The unfolded distribution for each pseudo-experiment can then be compared to the particle-level distribution used to generate the distributions originally. A pull variable,

$$p_j^i = \frac{XS_j^i - T^i}{\sigma_i}, \quad (6.15)$$

is defined, with p_j^i the pull for pseudo-experiment j in bin i . XS_j^i is the unfolded cross-section in bin i for pseudo-experiment j , and T^i the particle-level cross-section in bin i . σ_i is the error estimated by the unfolding procedure. Calculating this pull across these pseudo-experiments produces a distribution for each bin of the observable distribution. This pull distribution is then fitted with a Gaussian curve, and the mean and standard deviation of this Gaussian extracted for each bin as μ_{pull} and σ_{pull} , respectively. If the estimate of the uncertainty from the unfolding, σ_i , is reasonable, then one expects μ_{pull} to be around zero and σ_{pull} around unity.

Figures 6.21 and 6.22 show the distributions of μ_{pull} and σ_{pull} for τ_{32} and D_2 in single- and double-differential distributions, respectively.

Stress test The final validation test is the Data/MC *stress* test, which aims to quantify the bias induced by the unfolding technique. These may arise because of the dependence of the unfolding machinery itself on MC simulation, which could bias the unfolding towards the existing prediction and hence limit sensitivity to the data. The shape of the final unfolded distributions should not depend, for example, on the shape of the distributions from which the corrections were defined. Modelling uncertainties related to the choice of simulation are covered in Chapter 7, but there are also biases arising from the regularisation in the unfolding which is used to limit the effect of statistical fluctuations.




Figure 6.21: Unfolding pull test for substructure observables, showing the mean and standard deviation of the Gaussian fit to the distribution of the pull variable across 1000 pseudo-experiments. The error bands represent the uncertainties from the fit on the pull distribution.




Figure 6.22: Unfolding pull test for substructure observables, showing the mean and standard deviation of the Gaussian fit to the distribution of the pull variable across 1000 pseudo-experiments. The error bands represent the uncertainties from the fit on the pull distribution.

The regularisation imposed by the unfolding mechanism depends on the agreement between detector-level and particle-level distributions, as these are used to derive the corrections. The regularisation is then expected to be stronger when there is a greater discrepancy between the data and the MC. The strategy with the stress test is to use the unfolding procedure with nominal corrections to unfold distributions with a shape which has been altered.

This is done by re-weighting the particle-level and detector-level with a continuous *stress function*, unfolding the re-weighted detector-level distribution, and comparing the result to the re-weighted particle-level distribution. In the case that the nominal configuration of the unfolding mechanism induces a bias towards the particle-level distribution from which the corrections were derived, some non-closure would be expected. For the fairness of comparison between detector- and particle-level distributions, only events which passed selection for both are considered.

For each observable, three stress functions are used. All observables are stressed using functions which are based on the data/MC agreement for the τ_{32} observable and the top jet p_T , as these both exhibit poor agreement. A third function is then defined for each observable based on the agreement for that observable.

For the stress functions derived from the substructure observables, a second-order polynomial is fit to the distribution of $\frac{\text{data}}{\text{MC}} - 1$ for that observable, which has the form

$$a_0 + a_1x + a_2x^2. \quad (6.16)$$

The parameters of the fit are then varied both up and down by twice the error on that parameter, as given by the fit. Of these variations, the one which results in the larger stress on the observable is used as the stress function for the test.

The stress function derived from the p_T of the top candidate jet is constructed by again fitting $\frac{\text{data}}{\text{MC}} - 1$ for the p_T distribution, this time with a linear function, and then increasing the gradient of that fit by a factor of two.

The stress functions derived from τ_{32} and the top p_T are shown in Figure 6.23.

The results of the stress test are shown in Figures 6.24 and 6.25 for the τ_{32} and D_2 observables, for single- and double-differential distributions, respectively.

The results of the stressing of observables are evaluated by taking the ratio of the unfolded detector-level re-weighted distribution to the particle-level re-weighted




Figure 6.23: Stress functions applied in the stress test for all measured observables.




Figure 6.24: Particle-level distributions for substructure observables. The lines represent a ratio to the re-weighted particle-level distribution. The dashed lines are for the nominal particle-level prediction, whilst the solid lines are for the re-weighted detector-level MC which has been unfolded. The ratio for the unfolded distribution should be compatible with unity in order for the test to have achieved closure.




Figure 6.25: Particle-level distributions for substructure observables, further divided into regions of the top jet kinematics. The lines represent a ratio to the re-weighted particle-level distribution. The dashed lines are for the nominal particle-level prediction, whilst the solid lines are for the re-weighted detector-level MC which has been unfolded. The ratio for the unfolded distribution should be compatible with unity in order for the test to have achieved closure.

distribution. If the unfolding is sufficiently unbiased, this ratio should sit at unity. For the stress test to be passed, the ratio shown in the solid lines in Figures 6.24 and 6.25 must be compatible with unity to within the displayed error band, which gives the statistical uncertainty on the detector-level re-weighted distribution. It can be seen that this bias is covered by the statistical uncertainty for the observables shown, and this is also true for all measured observables.

Chapter 7.

Uncertainties

“... There’s no problem so bad that you can’t make it worse.”

— Chris Hadfield, *An Astronaut’s Guide to Life on Earth* [215]

The results presented in this thesis are subject to several sources of uncertainty. These may result from the limitations imposed by the physical characteristics of the detector and the reconstruction of physics objects from the detector, but also from the theoretical assumptions which are woven into the analysis through modelling choices. This chapter details these sources of uncertainty.

The evaluation of the impact of each uncertainty on the particle-level results is performed as a part of the unfolding procedure. Each source of systematic uncertainty is varied one at a time, with a new pseudo-dataset produced for each of these variations. This modified sample is then subject to the full analysis flow, all the way through the unfolding back to the particle-level. The resulting uncertainty for a given variation is then taken as the difference between the resultant distribution and the nominal prediction. When this has been done for every variation, the uncertainties are added in quadrature to produce an overall uncertainty.

7.1. Detector-related uncertainties

Leptons The performance of lepton reconstruction and identification varies between the data and the MC. Scale factors are used to account for the difference between the two. The associated uncertainties are estimated by varying these scale factors.

A similar method of scale factor variation is used to account for uncertainties in the lepton trigger efficiencies. This is done as prescribed by the ATLAS e/γ and muon performance groups

There are also uncertainties arising from the variation of the lepton momentum scale and resolution, which affects the lepton momentum. This is to account again for the difference between data and MC.

E_T^{miss} Uncertainties in the E_T^{miss} modify both the magnitude and direction of its four-vector. The uncertainties in E_T^{miss} are propagated from the energy scales and resolutions of leptons and jets. Scale and resolution uncertainties also exist which are associated with the soft term. These are independent from those coming from the leptons and jets.

Small- R jets The Category Reduction scheme supplied by the ATLAS Jet- E_T^{miss} group is used to evaluate the uncertainty on the jet energy scale (JES) for the $R = 0.4$ jets. This comprises 30 nuisance parameters, each with associated up/down variations and p_T and η dependencies, which parameterise the JES uncertainties.

The full-JER scheme provided by the ATLAS Jet- E_T^{miss} group is used to evaluate the uncertainty on the jet energy resolution (JER). The JER is parameterised by stochastic, noise, and constant terms. If the JER in MC is smaller than the JER in data, then it is matched to the average data resolution by smearing. The MC jets are then smeared using the result of the previous smearing and a 1σ variation. If the data JER is below the MC, the difference between the two is taken as an uncertainty. This scheme provides 13 variations in total, all modifying the energy of the jet.

The jet vertex tagger (JVT) has a different efficiency in data to MC, with scale factors derived to reflect this. The uncertainty associated with the JVT cut is evaluated by varying this scale factor.

b -tagging performance varies between simulation and data and thus must be calibrated. These calibrations have associated uncertainties which are propagated by the variation associated to the scale factors derived for this purpose. There are 9 variations for the b -jet calibration and 4 for each of the c -jet and light-jet calibrations, with a further two used to account for the extrapolation to high- p_T .

Reclustered large- R jets The systematic uncertainties associated with the reclustered large- R jets are simply propagated through from the constituent small- R jets, without the need for separate treatment.

Tracks Because tracks are used to define the substructure measurement, there are multiple sources of systematic uncertainties which are evaluated for them [216]. These are the impact of the ATLAS TightPrimary criteria, discussed in Section 4.2, as well as the impact of fake leptons, and the resolution on the impact parameters and the alignment process which may also affect the impact parameters. A final uncertainty is included to account for the challenges posed by the high track density present in high- p_T jets.

7.2. Modelling uncertainties

Uncertainties resulting from the modelling of events in Monte Carlo are evaluated using the alternative samples described in Section 4.1. The uncertainties are calculated by taking the difference between nominal and alternative samples following the unfolding procedure.

7.2.1. $t\bar{t}$ modelling

The uncertainties due to scale choice in the hard scatter and parton shower are evaluated using the POWHEG+PYTHIA8 samples reweighted to half and double the renormalisation scale, μ_R , and factorisation scale, μ_F . The uncertainty on the tuning choice for PYTHIA8 is evaluated using the samples with variations of Var3c and A14. The uncertainty on the FSR is evaluated using the samples with half and double the nominal μ_R^{FSR} value. Uncertainties on ISR resulting from choice of the h_{damp} parameter are evaluated with the POWHEG+PYTHIA8 sample with modified h_{damp} . Uncertainties on the choice of generator and parton shower models are evaluated using the MADGRAPH_AMC@NLOV2.6.0 and HERWIG7 samples, respectively. The uncertainty due to the top mass is evaluated using the $m_t' = m_t \pm 3.5 \text{ GeV}$ samples.

The uncertainty resulting from the choice of PDF is evaluated using the eigenvectors of the PDF4LHC30 PDF set. This estimate differs slightly in that the comparison is made

not to the NNPDF2.31o PDF set used in the nominal prediction but to the 0th eigenvector of the PDF4LHC30 prescription.

The stress tests implemented in the unfolding revealed a small non-closure at the extremities of some distributions, resulting from the poor Monte Carlo modelling of the top p_T . For this reason, a non-closure uncertainty is also included in the results.

7.2.2. Background modelling

Uncertainties due to the modelling of background processes are evaluated by comparison to samples where the modelling of that process is varied.

As the leading background in the analysis, several variations are considered for the single top simulation. Uncertainties on renormalisation and factorisation scales, the generator tune, FSR, and parton showering are all evaluated in the same way as described above for $t\bar{t}$ signal sample. Additionally, an uncertainty on the method for removing interference with $t\bar{t}$ is calculated by taking the difference between diagram removal (DR) and diagram subtraction (DS) approaches. This DR/DS uncertainty only applies to the tW production channel. There is also an uncertainty on the NNLO cross-section to which the sample is normalised as a result of the choice of PDFs and α_S . This uncertainty comes to 3.9%, 4.2%, and 5.4% for s -channel, t -channel, and tW respectively.

For the W +jets background, uncertainty is evaluated by comparison to samples with μ_R and μ_F doubled or halved with respect to their nominal value.

For the $t\bar{t}V$ processes, uncertainties on the cross-section calculations come to 13.3%, 12.0%, and 9.9% for $t\bar{t}Z$, $t\bar{t}W$, and $t\bar{t}H$ processes, respectively. A conservative choice is made to simply apply the 13.3% uncertainty to all three processes.

For both the Z +jets and diboson processes, a conservative 50% uncertainty [217] is applied. This is justified as these backgrounds are small for this analysis.

The uncertainty on events containing a fake lepton is evaluated by comparison of the nominal matrix method estimate with the alternative data-driven sample described in Section 5.5.2. This uncertainty is below 1% across most distributions, reaching 5% at the extremities.




Figure 7.1: Distribution of uncertainties as a function of the substructure of the large- R reclustered jet which is the hadronic top candidate, unfolded to particle-level. Statistical uncertainties are shown in an orange band, while the yellow band shows the sum of statistical and systematic uncertainties. Sources of systematic uncertainties are further separated out by the various coloured lines shown.

7.3. Summary of uncertainties

Figures 7.1–7.4 show the distribution of uncertainties on the normalised distribution of each of the eight substructure observables being measured. The leading sources of uncertainty come from the modelling uncertainties in FSR and parton showering. Track uncertainties also become significant for some of the observables. Uncertainties defined in this chapter which have minimal effect are grouped together under *other*. These include uncertainties for JVT, pileup, luminosity, flavour-tagging, leptons, E_T^{miss} , and PDF uncertainties.

In general, because the results are presented as normalised distributions, uncertainties arising from the JES and b -tagging are limited in their impact. This is because they generally impact on the event rate and so this effect does not manifest in the shape of the substructure observables. For the double-differential distributions however, where the cross-section is given as a function of both τ_{32} or D_2 and the mass of the hadronic top candidate jet or its p_T , the JES does play a significant role in the high mass region. The uncertainties for these distributions are shown in Figures 7.5–7.8.




Figure 7.2: Distribution of uncertainties as a function of the substructure of the large- R reclustered jet which is the hadronic top candidate, unfolded to particle-level. Statistical uncertainties are shown in an orange band, while the yellow band shows the sum of statistical and systematic uncertainties. Sources of systematic uncertainties are further separated out by the various coloured lines shown.




Figure 7.3: Distribution of uncertainties as a function of the substructure of the large- R reclustered jet which is the hadronic top candidate, unfolded to particle-level. Statistical uncertainties are shown in an orange band, while the yellow band shows the sum of statistical and systematic uncertainties. Sources of systematic uncertainties are further separated out by the various coloured lines shown.




Figure 7.4: Distribution of uncertainties as a function of the substructure of the large- R reclustered jet which is the hadronic top candidate, unfolded to particle-level. Statistical uncertainties are shown in an orange band, while the yellow band shows the sum of statistical and systematic uncertainties. Sources of systematic uncertainties are further separated out by the various coloured lines shown.




Figure 7.5: Uncertainties on the distribution of the $t\bar{t} \ell + \text{jets}$ production cross-section as a function of the τ_{32} substructure observable for the large- R reclustered jet which is the hadronic top jet candidate, split into three regions of the mass of the jet. Statistical uncertainties are shown in an orange band, while the yellow band shows the sum of statistical and systematic uncertainties.




Figure 7.6: Uncertainties on the distribution of the $t\bar{t} \ell + \text{jets}$ production cross-section as a function of the τ_{32} substructure observable for the large- R reclustered jet which is the hadronic top jet candidate, split into five regions of the p_T of the jet. Statistical uncertainties are shown in an orange band, while the yellow band shows the sum of statistical and systematic uncertainties.




Figure 7.7: Uncertainties on the distribution of the $t\bar{t} \ell + \text{jets}$ production cross-section as a function of the D_2 substructure observable for the large- R reclustered jet which is the hadronic top jet candidate, split into three regions of the mass of the jet. Statistical uncertainties are shown in an orange band, while the yellow band shows the sum of statistical and systematic uncertainties.




Figure 7.8: Uncertainties on the distribution of the $t\bar{t} \ell + \text{jets}$ production cross-section as a function of the D_2 substructure observable for the large- R reclustered jet which is the hadronic top jet candidate, split into five regions of the p_T of the jet. Statistical uncertainties are shown in an orange band, while the yellow band shows the sum of statistical and systematic uncertainties.

Chapter 8.

Results and Conclusions

This chapter introduces the normalised differential cross-sections as a function of the substructure observables selected in Chapter 6. The distributions have been unfolded to particle-level as described in Section 6.2 and are compared to several alternative predictions described in Section 4.1.2. These are shown in Figures 8.1–8.4.




Figure 8.1: Distribution of the $t\bar{t} \ell$ +jets production cross-section as a function of the substructure of the large- R reclustered jet which is the hadronic top candidate, unfolded to particle-level. Distributions are shown for several NLO predictions and the data. The two lower pads show the ratio between the MC predictions and the data. Of these two pads, the upper one shows the nominal and alternative models, whilst the lower one shows variations of the nominal prediction. Statistical uncertainties are shown in an orange band, while the yellow band shows the sum of statistical and systematic uncertainties.




Figure 8.2: Distribution of the $t\bar{t}\ell$ +jets production cross-section as a function of the substructure of the large- R reclustered jet which is the hadronic top candidate, unfolded to particle-level. Distributions are shown for several NLO predictions and the data. The two lower pads show the ratio between the MC predictions and the data. Of these two pads, the upper one shows the nominal and alternative models, whilst the lower one shows variations of the nominal prediction. Statistical uncertainties are shown in an orange band, while the yellow band shows the sum of statistical and systematic uncertainties.




Figure 8.3: Distribution of the $t\bar{t}\ell$ +jets production cross-section as a function of the substructure of the large- R reclustered jet which is the hadronic top candidate, unfolded to particle-level. Distributions are shown for several NLO predictions and the data. The two lower pads show the ratio between the MC predictions and the data. Of these two pads, the upper one shows the nominal and alternative models, whilst the lower one shows variations of the nominal prediction. Statistical uncertainties are shown in an orange band, while the yellow band shows the sum of statistical and systematic uncertainties.




Figure 8.4: Distribution of the $t\bar{t}\ell$ +jets production cross-section as a function of the substructure of the large- R reclustered jet which is the hadronic top candidate, unfolded to particle-level. Distributions are shown for several NLO predictions and the data. The two lower pads show the ratio between the MC predictions and the data. Of these two pads, the upper one shows the nominal and alternative models, whilst the lower one shows variations of the nominal prediction. Statistical uncertainties are shown in an orange band, while the yellow band shows the sum of statistical and systematic uncertainties.

The τ_{32} distribution, shown in Figure 8.1a, rises from the second bin around 0.2 to a peak around 0.4, before falling off gradually towards 0.8 and then falling sharply in the highest bin for the least top-like structure. In general, the Monte Carlo overestimates at low values of τ_{32} and underestimates at higher values. In other words, the simulation is more top-like than the data.

The τ_{21} distribution, shown in Figure 8.1b, rises from zero up to around 0.5 before falling off again at higher values. It exhibits generally good agreement with the data, with slight overestimation by the Monte Carlo at lower values.

τ_3 , shown in Figure 8.2a, rises from 0 to a peak around 0.035, before falling off again at the upper end of the distribution. The Monte Carlo provides an overestimate of the data at low values and an underestimate towards the higher end of the distribution. This is reflected in the distribution of τ_{32} .

ECF2, shown in Figure 8.2b, peaks at the centre of the distribution, around 0.2, with tails falling off to either side to the extremities of the spectrum. There is good agreement across the distribution between the Monte Carlo prediction and the data, within the combined uncertainties.

The C_3 distribution, shown in Figure 8.3a, rises to a peak around 0.25 before falling to near zero at the higher end of the distribution around 0.45. As was seen with τ_{32} , the Monte Carlo overestimates at the lower end of the spectrum and provides an underestimate at the upper end. In the same way as for τ_{32} , this represents a more top-like structure in the simulation than in the data.

The Les Houches Angularity, shown in Figure 8.3b, peaks around 0.6, falling to almost zero at 0.2 and more sharply at the upper end of the spectrum to near zero around 0.8. Across the full distribution, there is good agreement between the Monte Carlo and the data within the combined uncertainties.

D_2 , shown in Figure 8.4a, peaks around 1.2. As for ECF2, the Monte Carlo and the data are in good agreement across the distribution.

Finally, $p_T^{d,*}$ is shown in Figure 8.4b. The distribution rises sharply from just under 0.2 to a peak just below 0.3, before falling off more gradually to near zero around 0.7. The spectrum is poorly described by the Monte Carlo, with underestimation at the lower end of the distribution and overestimation at the upper end. This corresponds to the Monte Carlo predicting jets which have their p_T more concentrated amongst constituents than is the case in the data.

By comparison across the measured observables, some general observations can be made.

The nominal POWHEG+PYTHIA8 prediction is in reasonably good agreement for the majority of the observables measured in this analysis. The exceptions to this are τ_3 , τ_{32} , and $p_T^{d,*}$. The POWHEG+PYTHIA8 FSR-up variation shows poorer agreement with the data, whilst the FSR-down variation generally improves the agreement, which suggests the MC is overestimating the FSR. The model which succeeds the best at describing the data is MADGRAPH_AMC@NLO+PYTHIA8. As compared to the nominal prediction, the POWHEG+ HERWIG7 sample also fits the data better.

Whilst some of the predictions do a better job than others, none model τ_{32} or τ_3 very well. C_3 is poorly modelled by the nominal prediction, though agreement with the data is better for MADGRAPH_AMC@NLO+PYTHIA8, POWHEG+ HERWIG7, and POWHEG+PYTHIA8 (FSR down) samples. This would seem to suggest that the NLO predictions are not able to accurately model the three-pronged structure which results in jets from top quark decays, particularly the nominal prediction.

The two-prong probes — τ_{21} and D_2 — and ECF2 agree well with the data across the board, and are also sensitive to the FSR variation. Considering the previous observation on the poor modelling of the three-pronged observables, this would seem to suggest that the poor modelling takes effect before FSR effects are included.

As for the other three-prong discriminator, C_3 , this is also not modelled particularly well. This could arise for the same reasons as the poor modelling of τ_3 and τ_{32} , but may also be affected by the p_T of the jets, as C_3 scales under Lorentz boost.

As $p_T^{d,*}$ is a metric for the momentum distribution within a jet, it is sensitive to the choice of hadronisation model. In this regard, HERWIG7 seems to do a better job than PYTHIA8 of describing the distribution. By contrast, the Les Houches Angularity — an observable more sensitive to the ISR — is described well by each of the NLO+PS predictions considered.

Figures 8.5 and 8.6 show normalised double-differential cross-sections. Here, the cross-section is probed as a function of both the substructure observables and the kinematics of the hadronic top candidate jet. The τ_{32} and D_2 distributions are displayed in regions of both the jet p_T and mass. As τ_{32} and D_2 rank highly for use in three- and two-prong taggers which aim to discriminate between jets with those




Figure 8.5: Distribution of the $t\bar{t} \ell$ +jets production cross-section as a function of the τ_{32} substructure observable for the large- R reclustered jet which is the hadronic top jet candidate, unfolded to particle level. Shown is a comparison between the data and the nominal prediction, with results divided into regions of the jet kinematics.

structures and others with simpler substructure, the correlation with the jet mass and p_T are of interest.

Figures 8.7–8.10 show the ratio in these regions of the predictions to the data across the regions of jet mass and p_T , for both τ_{32} and D_2 .

Observations of the various models for each observable are made with reference not only to the distributions shown here but to a χ^2 test used to quantify the agreement or otherwise between the unfolded spectra and the various NLO+PS predictions, with associated p -value also extracted.

The χ^2 is calculated as

$$\chi^2 = V^T C^{-1} V, \quad (8.1)$$

where V is the vector of residuals between the unfolded data and predictions and C the covariance matrix which includes statistical uncertainties as well as the bin-to-bin correlations which are induced by the regularisation of the unfolding procedure and the systematic uncertainties. The results from this χ^2 calculation and the extracted p -values are shown in Tables 8.1 and 8.2. Respectively, these tables display the χ^2 results




Figure 8.6: Distribution of the $t\bar{t} \ell + \text{jets}$ production cross-section as a function of the D_2 substructure observable for the large- R reclustered jet which is the hadronic top jet candidate, unfolded to particle level. Shown is a comparison between the data and the nominal prediction, with results divided into regions of the jet kinematics.




Figure 8.7: Distribution of the $t\bar{t} \ell + \text{jets}$ production cross-section as a function of the τ_{32} substructure observable for the large- R reclustered jet which is the hadronic top jet candidate, unfolded to particle level and split into three regions of the mass of the jet. Shown is the ratio between NLO predictions and the data. Statistical uncertainties are shown in an orange band, while the yellow band shows the sum of statistical and systematic uncertainties.




Figure 8.8: Distribution of the $t\bar{t}\ell$ +jets production cross-section as a function of the τ_{32} substructure observable for the large- R reclustered jet which is the hadronic top jet candidate, unfolded to particle level and split into five regions of the p_T of the jet. Shown is the ratio between NLO predictions and the data. Statistical uncertainties are shown in an orange band, while the yellow band shows the sum of statistical and systematic uncertainties.




Figure 8.9: Distribution of the $t\bar{t}\ell$ +jets production cross-section as a function of the D_2 substructure observable for the large- R reclustered jet which is the hadronic top jet candidate, unfolded to particle level and split into three regions of the mass of the jet. Shown is the ratio between NLO predictions and the data. Statistical uncertainties are shown in an orange band, while the yellow band shows the sum of statistical and systematic uncertainties.




Figure 8.10: Distribution of the $t\bar{t}\ell$ +jets production cross-section as a function of the D_2 substructure observable for the large- R reclustered jet which is the hadronic top jet candidate, unfolded to particle level and split into five regions of the p_T of the jet. Shown is the ratio between NLO predictions and the data. Statistical uncertainties are shown in an orange band, while the yellow band shows the sum of statistical and systematic uncertainties.

for the nominal, alternative predictions, and the variations on the nominal prediction. PWG+PY8 refers to the POWHEG+PYTHIA8 sample, PWG+H7 to the POWHEG+HERWIG7 sample, and MC@NLO+PY8 to the MADGRAPH_AMC@NLO+PYTHIA8 sample.

When splitting the τ_{32} distribution into regions of the jet mass, shown in Figure 8.7, the behaviour is similar as for the inclusive τ_{32} spectrum. Across the three mass regions, there is generally good agreement between the predictions and the data in the central regions of τ_{32} , with poor agreement for low and high values of τ_{32} . Generally, the POWHEG+HERWIG7 sample outperforms the nominal POWHEG+PYTHIA8 sample, with the exception of the low mass region. Overall, the FSR-down variation of the nominal prediction gives the most consistent agreement across the three mass regions.

In the case that the τ_{32} distributions are viewed in regions of the jet p_T — shown in Figure 8.8 — similar trends are exhibited. For high and low values of τ_{32} , modelling is generally poor. Again, there doesn't seem to be any great distinction between any of the p_T regions and the inclusive spectrum. As before, POWHEG+HERWIG7 describes the data better than the nominal prediction for the majority p_T spectrum. The FSR-

Observable	PWG+PY8		PWG+H7		MC@NLO+PY8	
	χ^2/NDF	$p\text{-value}$	χ^2/NDF	$p\text{-value}$	χ^2/NDF	$p\text{-value}$
τ_{32}	53/12	<0.01	19/12	0.08	15/12	0.24
τ_{21}	14/14	0.44	8/14	0.91	16/14	0.31
τ_3	37/11	<0.01	41/11	<0.01	13/11	0.31
ECF2	24/18	0.16	13/18	0.78	15/18	0.68
C_3	11/13	0.61	6/13	0.96	3/13	1.00
LHA	14/17	0.69	10/17	0.92	21/17	0.22
D_2	19/16	0.28	17/16	0.40	21/16	0.20
$p_T^{d,*}$	27/12	<0.01	12/12	0.48	11/12	0.52
τ_{32} vs. m	151/42	<0.01	75/42	<0.01	57/42	0.06
τ_{32} vs. p_T	148/50	<0.01	101/50	<0.01	54/50	0.31
D_2 vs. m	58/42	0.05	62/42	0.03	59/42	0.04
D_2 vs. p_T	70/56	0.10	65/56	0.20	70/56	0.11

Table 8.1: χ^2 and p -value for each unfolded distribution, relative to the data, for the nominal prediction and the alternative predictions.

Observable	PWG+PY8		PWG+PY8 (FSR-UP)		PWG+PY8 (FSR-DOWN)	
	χ^2/NDF	$p\text{-value}$	χ^2/NDF	$p\text{-value}$	χ^2/NDF	$p\text{-value}$
τ_{32}	53/12	<0.01	164/12	<0.01	39/12	<0.01
τ_{21}	14/14	0.44	41/14	<0.01	7/14	0.92
τ_3	37/11	<0.01	132/11	<0.01	23/11	0.02
ECF2	24/18	0.16	29/18	0.04	24/18	0.15
C_3	11/13	0.61	36/13	<0.01	4/13	0.99
LHA	14/17	0.69	13/17	0.76	19/17	0.35
D_2	19/16	0.28	33/16	<0.01	18/16	0.35
$p_T^{d,*}$	27/12	<0.01	55/12	<0.01	25/12	0.01
τ_{32} vs. m	151/42	<0.01	411/42	<0.01	75/42	<0.01
τ_{32} vs. p_T	148/50	<0.01	357/50	<0.01	111/50	<0.01
D_2 vs. m	58/42	0.05	114/42	<0.01	45/42	0.33
D_2 vs. p_T	70/56	0.10	105/56	<0.01	92/56	<0.01

Table 8.2: χ^2 and p -value for each unfolded distribution, relative to the data, for the nominal prediction and the FSR-up and FSR-down variations of the nominal prediction.

down variation of the nominal prediction greatly improves the agreement with the data, though this trend is less pronounced for high values of τ_{32} .

Overall, it does not seem that the poor modelling of τ_{32} is particularly better or worse either at high or low jet p_T or high or low jet mass.

As for the single-differential results, D_2 is generally better described by the NLO predictions than τ_{32} . When viewed as a function of the jet mass — as in Figure 8.9 — most of the simulated samples describe the data well in the three mass regions. The exception to this is in the low mass region, where the lower end of the D_2 spectrum is less well described. Again, the FSR-down variation of the nominal prediction improves the agreement with the data. Unlike for the τ_{32} distributions, it is not clear that the POWHEG+ HERWIG7 sample performs any better than the nominal POWHEG+PYTHIA8 prediction.

A similar pattern emerges for D_2 when viewed as a function of the jet p_T , as shown in Figure 8.10. Once more, the modelling is superior to that for τ_{32} , but with some regions of the distributions which do not agree particularly well with the data. The greatest disagreements tend to be seen at the lower end of the D_2 spectrum, an effect which is greatest for the low- p_T region. Again, POWHEG+ HERWIG7 performs better than the nominal POWHEG+PYTHIA8 prediction.

So, although D_2 is generally well described, this success of the NLO predictions begins to unravel for low jet mass and high jet p_T .

A clear picture does emerge, however, to show that the HERWIG7 parton showering model provides better agreement with the data than the nominal POWHEG+PYTHIA8 prediction. Additionally, the FSR-up variation of this nominal prediction produces poor agreement with the data, whilst the FSR-down variation improves the description of the observation.

Overall, the relatively poor description of many of these substructure observables serves to underline the need for improved modelling of substructure in boosted $t\bar{t}$ events in the future.

Given the sensitivity of some of these observables to the variation of FSR, it may be possible going forward to perform a fit to extract the value of α_s which best replicates the observed data. This is one example of how such a measurement could be used to produce Monte Carlo tunes which better describe the data for processes such as these.

Part III.

Development of Pixel Detector Modules for the ATLAS Inner Tracker

Chapter 9.

The High-Luminosity LHC and the ATLAS Inner Tracker

Following the conclusion of Run 3, the LHC will undergo a major upgrade during Long Shutdown 3 (LS3) from 2026–2029. Run 4 will then mark the beginning of the High-Luminosity LHC (HL-LHC) era. The conditions which will be present at the HL-LHC necessitate major upgrades of the LHC experiments. ATLAS will completely replace its inner tracking system — the Inner Detector (ID) — with a new tracker, the Inner Tracker (ITk). This chapter details the upgrade of the LHC as well as the ATLAS ITk upgrade. Particular attention is paid to the ITk pixel detector sub-system, in order to motivate studies presented in Chapter 10.

9.1. Upgrade to High-Luminosity LHC

The HL-LHC is designed to deliver an integrated luminosity of around 4000 fb^{-1} over its lifetime [59], facilitating an order of magnitude increase in the amount of physics data recorded by the experiments around the ring. In order to do this, the instantaneous luminosity will reach $7.5 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$ — 7.5 times the nominal LHC luminosity. As well as a significant increase in the luminosity, the collision energy is also due to be increased to $\sqrt{s} = 14 \text{ TeV}$.

The upgrades to the accelerator complex which are required to meet these specifications are due to be undertaken during LS3, building upon upgrades implemented ahead of Run 3, with HL-LHC operation from the start of Run 4. The HL-LHC is then expected to operate until around 2040. This can be seen in Figure 9.1.




Figure 9.1: Schedule for the HL-LHC upgrade [218].

A consequence of the increase in luminosity is a significant increase in pileup — the number of simultaneous interactions per bunch crossing. At the HL-LHC, it is expected that pileup will reach $\mu = 200$, an approximately four times increase on the levels observed during Run 2 and expected in Run 3. This presents significant challenges for the LHC experiments and their ability to produce physics results from the delivered collisions.

9.1.1. Motivation and physics prospects

Measurements of the interactions and self-interaction of the Higgs boson are a crucial component of the LHC physics program. These are key direct probes of the SM. It is expected that the HL-LHC will allow ATLAS to measure the Higgs production cross-sections to a precision between 2.4% and 7.7% [219], constrain the couplings to around 2% to 7% and achieve a significance on the Higgs self-coupling λ_{HHH} of 3.0σ when combining the main decay channels of the HH production [220]. It is possible that this could rise to 4.0σ when combined with CMS measurements [221].

As for other SM processes, vector boson scatterings like WW , WZ , and ZZ are good candidates for precision measurement at the HL-LHC [222]. These processes

Standard Model Production Cross Section Measurements

Status: February 2022




Figure 9.2: Cross-sections for many SM processes accessible to ATLAS [30]. Measurements from all data-taking periods are included as well as theoretical predictions. Measurements are corrected for branching fraction.

are generally identified by the $VVjj$ event signature. As can be seen in Figure 9.2, this is a signature with an extremely small cross-section which has naturally limited opportunities for measurement up to now.

Searches for BSM physics are also of huge importance to the LHC experiments and cover a broad range of analyses from searches for supersymmetry and dark matter to searches for flavour-changing neutral currents (FCNC) in top decays. Many of these types of processes have very small cross-sections, leaving measurements dominated by statistical uncertainties and searches in need of more data. These statistical limitations will be mitigated by the order-of-magnitude increase in luminosity. Analyses of some of these processes will also benefit from enhanced cross-sections resulting from the slight increase to the centre-of-mass energy. To take the example of FCNC, ATLAS has performed searches across a variety of decay and production modes in recent years, in tqg [223], tqH [224], tqZ [225], and $tq\gamma$ [226], without finding evidence for new

physics. It is expected, however, that the HL-LHC could allow for as much as an order of magnitude increase in sensitivity in these kinds of measurements [227]. Sensitivity to chargino and neutralino production — processes with very small cross-sections — is expected to benefit hugely from the increased luminosity on offer at HL-LHC, with sensitivity near the TeV mass scale [228]. Similarly, searches for dark matter focusing on large E_T^{miss} are enhanced by both the increased luminosity and centre-of-mass energy.

In top physics, there have been recent analyses across many sub-groups in ATLAS which would benefit from the opportunities provided at the HL-LHC. In properties, for example, the measurement of charge asymmetry in $t\bar{t}\gamma$ [229] — a process with a small cross-section — is statistically limited, whilst in the cross-section group, the recent measurement of single top production in the s -channel [230] was also statistically limited and could not claim significance to the level of observation.

Apart from the increases in luminosity and centre-of-mass energy, detector upgrades necessitated by the move to high-luminosity — detailed in Section 9.2 — may in some cases bring their own opportunities. The improved forward lepton acceptance in ATLAS, for example, is expected to allow for improvements to measurements of the electroweak mixing angle $\sin \theta_W$ [231] as well as making it possible to measure the mass of the W boson to a precision below 10 MeV [232]. This is of heightened interest following the surprising measurement published by CDF in 2022 [233], which was to comparable precision.

It is clear that the upgrade to HL-LHC opens up an even broader and deeper physics program than that of the current manifestation of the LHC, from precision measurement to searches for vanishingly rare processes. Establishing that the HL-LHC environment is desirable is, of course, only the first step. A vast series of upgrades to the LHC machine is required to realise this set of conditions.

9.1.2. Accelerator upgrades

At the LHC, the beam energy is ultimately determined by the strength of the magnetic field generated by the 1232 superconducting dipole magnets spaced around the ring. These magnets deflect the trajectory of particles as they move through the ring, with an 8 T magnetic field providing the force which “bends” their path. There is one

beam-pipe for each of the two beams, which travel in opposite directions around the LHC¹. Each dipole is cooled by liquid helium held at 1.9 K.

In addition to the dipole magnets at the LHC, there are a host of other magnets used to shape and control the beam. Most notable of these are the quadrupole magnets used for focusing the beam. The field generated by the quadrupoles is designed such that its zero-point is at the centre of the beam. The polarity is established such that particles in the horizontal plane are directed towards the centre of the beam, should they be off-axis, whilst particles in the vertical plane are deflected away from the centre of the beam — focusing and de-focusing, respectively. The gradient of the field increases outwards from the zero-point such that the strongest focusing or de-focusing happens furthest from the centre of the beam.

The luminosity provided by the LHC is essentially determined by the number of proton bunches making up the the beam, the number of protons in each bunch, and the beam cross-section at the interaction point (IP). At the LHC, there are 2808 proton bunches, each containing 1.20×10^{11} protons and spaced by 25 ns, and with a transverse size around 40 μm . The four main experiments on the LHC are each positioned in a long straight section (LSS) of the tunnel, 280 m long. This allows for a series of magnets to focus and direct the beams together such that the bunch-crossing at the IP happens with a very small cross-section. Most important are the inner-triplets — focusing quadrupoles on either side of the IP. These are shown schematically in Figure 9.3. The beam-size at any given point around the ring is related to the beta function at that point. At the IP, where the beams cross, the beta function is labelled β^* . The luminosity which can be delivered to the experiments is in large part determined by β^* .

With reference to these factors, the HL-LHC will aim to reduce β^* at the ATLAS IP. There will also be an upgrade to the injector system in order to raise the bunch population without increasing the size of the beam.

Increasing the bunch population will require a larger crossing angle between the bunches in order to avoid collisions between the protons either side of the IP as the beams move closer together. This increase in the crossing angle also has the effect of lowering the luminosity. So although β^* is lowered by upgrades to the injector scheme and magnets in the LSS either side of the IP, much of this gain is then mitigated by

¹CERN maintains that any resemblance to the design of the Glasgow Subway is purely coincidental.




Figure 9.3: Schematic of the HL-LHC magnets directly before the ATLAS interaction point [59]. Those marked *Q* are quadrupoles whilst those marked *D* are separation and recombination dipoles. Q1–Q3 are the inner triplet quadrupoles. The beam is shown in grey.

the necessary increase in the crossing angle. This can be seen in the lower portion of Figure 9.4a.

The degradation of the instantaneous luminosity which results from the higher bunch crossing angle can be resolved by the use of novel radio-frequency crab cavities². The crab cavities tilt the bunches in the transverse plane as they approach the IP so as to increase the effective cross-section they present to each other and thus increase the number of collisions. They then reverse the process as the bunches pass through to the other side of the IP and continue along the LHC ring. Figure 9.4b shows the geometric reduction factor — the reduction of the instantaneous luminosity due to the crossing angle — as a function of β^* . The advantage of the crab cavities is then made clear — in this configuration, the HL-LHC is able to have a significantly lower β^* than in Run 2, whilst having a smaller geometric reduction.

²So-called either because of their “pinching effect” on the bunches, or because the bunches are then seen to march sideways towards each other.




Figure 9.4: Effect of the RF crab cavities on β^* at the HL-LHC.

9.2. The ATLAS Inner Tracker

The huge increase in luminosity at the HL-LHC will present the ATLAS detector with a huge increase in the number of particles passing through the detector. These effects are particularly acute closest to the IP. The main consequences of the new environment for the detector are:

1. Detector elements will be subject to far greater levels of radiation damage than previously experienced, which will degrade the performance of the sensing elements and electronics.
2. The large numbers of interactions and resulting high occupancy in the detector will exceed the limits of the current electronics in their readout rate and triggering capabilities.
3. Occupancy levels resulting from the high levels of pileup will degrade the ability of the detector to resolve objects. In other words, the granularity of the detector will be insufficient to cope with the extremely busy environment generated. This can be seen in Figure 9.5, which depicts an event display for a $t\bar{t}$ event in the high-pileup environment which will be present at the HL-LHC. The vast number




Figure 9.5: Event display for a simulated $t\bar{t}$ event, with pileup $\mu = 200$. The tracks coloured cyan emanate from the primary vertex, whilst those in magenta come from secondary vertices, i.e. pileup [235].

of vertices underlines the challenge of separating out tracks originating from the primary vertex and of differentiating different tracks from one another.

Of course by the end of Run 3, the ID will also have had to withstand radiation damage from around a decade of operation and much would need to be replaced even without the new performance requirements imposed by the HL-LHC.

The upshot of all of these factors is that the ID will be replaced in its entirety by the ITk. In contrast with the ID, the ITk will be an all-silicon tracking detector — that is, the ITk will not seek to replicate the TRT. The ITk will also extend tracking coverage out to $|\eta| < 4.0$. The layout for the ITk is shown in Figure 9.6. The layers closest to the beam are made of pixel modules whilst those slightly further out are strip modules. The ITk will allow for a minimum of 9 and 13 hits per track in the barrel and forward regions, respectively.

The strip detector [237] will be made of four barrel layers and six end-cap disks. The inner two layers in the barrel will use Short Strip (SS) modules whilst the outer two will use Long Strip (LS) modules. The six end-cap disks are made of 32 petal structures each built from modules arranged in six rings. Because these petals are wedge-shaped, they are built from six different sensor designs in order to form the six rings. Modules




Figure 9.6: The planned layout for the ATLAS ITk [236], showing one quadrant of the detector. The x -axis of the diagram runs along the z -axis in the ATLAS coordinate system, i.e. the beam direction, whilst the y -axis represents the radial distance from the IP. Pixel layers are shown in red and strip layers in blue.

are then attached to support structures in both barrel and end-cap regions which house cooling and services. Glued back-to-back, modules in the barrel are rotated by 26 mrad on each side to give a stereo angle of 52 mrad providing two-dimensional resolution in space. The same is done in the end-cap regions, with a stereo angle of 40 mrad. The strip detector will cover the region in $|\eta| < 2.7$ and extend out to 1 m radially.

Although the exact configuration varies by type, each module is made up of a few basic units. At least one *hybrid* and one *powerboard* are glued onto a 320 μm thick n^+ -in-p silicon sensor. The hybrid is a PCB housing the front-end (FE) readout chips — the ATLAS Binary Chip (ABC) — and the Hybrid Controller Chip (HCC). The HCC amalgamates the data from the ABCs and sends the event to the End-of-Substructure³ (EoS) card on the support structure. The powerboard performs DC-DC conversion for the FE as well as providing information on voltages, currents, and temperatures. Depending on position in the detector, strips have length 1.5–6.0 cm and pitch 70–81 μm .

Altogether, the strip detector will be built from around 18 000 modules covering an area around 164 m^2 , with almost 5×10^7 channels. This represents an order-

³And you thought you'd seen the end of the substructure in this thesis.




Figure 9.7: Graphic of the ATLAS ITk [236]. The beamline is represented in blue, running through the centre of the image and surrounded in the first instance by the pixel detector and the strip detector outside that.




Figure 9.8: Number of hits per track expected across both pixel and strip detectors in the ITk, as a function of track pseudorapidity [236].

of-magnitude increase in the number of channels as compared to the ID strip detector, the SCT. Over its lifetime, the expected fluence for the ITk strip detector is $1.6 \times 10^{15} n_{\text{eq}} \text{ cm}^{-2}$, where n_{eq} is the 1 MeV neutron-equivalent. Testbeam measurements have shown that ITk strip modules will still be able to meet their performance goals by the end of this 4000 fb^{-1} period, mainly the track resolution and hit efficiency.

Both pixel and strip structures are shown in Figure 9.7, whilst Figure 9.8 shows the expected number of hits in the tracker as a function of η .

9.3. The ITk pixel detector

Before discussing the specifics of the pixel detector to be installed in the ITk, it is worth touching on the basic operation principles of pixel detectors in general.

All pixel detectors presented here are *hybrid* pixel detectors, meaning that they consist of a semiconductor sensor, in which a signal is generated, and separate electronics which are then joined to that sensor to amplify and read out the signal.

In a semiconductor like silicon, there is a gap of 1.12 eV between the valence band and the conduction band, corresponding to the highest filled energy level and the next available energy level. When a charged particle passes through such a material, it releases electron-hole pairs as an electron is promoted to the conduction band and a hole, or lack of electron, is left behind. These electrons and holes act as charge carriers within the silicon. In silicon there is an intrinsic charge carrier density which arises from thermally generated electron-hole pairs. In order to use the electron-hole pairs generated by a charged particle to make a measurement of that particle, this intrinsic carrier density must be overcome as the free charge carriers will swamp and recombine with the generated pairs.

This is achieved by doping. A semiconductor is labelled n-type after the addition of donor ions from Group 5 of the periodic table, which have energy levels at the lower end of the conduction band and so provide electrons. With the addition of acceptor ions from Group 3, which provide energy levels at the top of the valence band and so absorb electrons to supply holes, a semiconductor is labelled p-type. So in an n-type semiconductor there is an abundance of free electrons and in a p-type semiconductor, an abundance of free holes.

When an n-type semiconductor is brought into contact with a p-type semiconductor, a *pn-junction* is formed. Here, majority carriers diffuse across the boundary of the junction. Electrons from the n-type region cross into the p-type region where they combine with holes, and holes from the p-type region cross into the n-type region, where they combine with electrons. The donor and acceptor ions are fixed in place in the crystal and cannot move but now attain positive and negative charges respectively. In the n-type material, this produces a positively charged region near the junction and in the p-type material, a negatively charged region near the junction. This has the effect of generating an electric field which repels the majority carriers in each type of material away from the boundary, so presenting a barrier to the carriers. These regions

PIXEL MODULES	Pitch	Sensor thickness	Channels	Modules	Area
ID	$50 \times 400 \mu\text{m}^2$	$250 \mu\text{m}$	9×10^7	2024	1.9 m^2
ITk	$50 \times 50 \mu\text{m}^2$	$150 \mu\text{m}$	5×10^9	9164	12.8 m^2


Table 9.1: A comparison between characteristics of the pixel subsystems in the ATLAS ID and ATLAS ITk. Caveats for the ID come from the IBL which has sensor thicknesses of $230 \mu\text{m}$ and $200 \mu\text{m}$ for 3D and planar sensors, respectively, with pitch $50 \times 250 \mu\text{m}^2$. For the ITk, it should be noted that the L1 planar sensors are $100 \mu\text{m}$ thick and that the barrel 3D sensors have pixel pitch of $25 \times 100 \mu\text{m}^2$.

at the junction lose their free charge carriers and so this is known as the *depletion zone*. By applying a *reverse bias* — applying a voltage across the junction — the width of this depletion zone can be extended.


In a hybrid pixel detector, these pn-junctions are manufactured by placing implants of one doping type onto a bulk (in this case silicon) of the opposite kind. This forms the sensor. The sensor is said to be fully depleted when the depletion zone extends throughout the full thickness of the silicon. The bias voltage at which the sensor becomes fully depleted is called the *depletion voltage*. As charged particles pass through this depletion zone and generate electron-hole pairs, the electrons and holes drift to opposite contacts, where the charge is collected, amplified, and read out. Thermally generated electron-hole pairs will always be present — this is called *leakage current* — but in a well-behaved pixel detector, this will be far below the charge resulting from an incident particle.

Each pixel implant on the sensor is connected to a channel on the readout chip in a process known as bump-bonding, in which a solder bump joins the two together and allows the FE to collect the signal from the pixel. Each pixel channel has a *threshold*, which determines what is considered a hit. A signal in that pixel which is higher than that threshold is considered a hit, whilst anything below is not. The aim therefore is to set this threshold above the thermal noise but not so high that the signal induced by a real particle passing through the detector is below this cut-off. Figure 9.9 shows charge collection in a single pixel cell and the assembly of a full pixel module.

The pixel detector in the ITk is to be 33 mm from the beam at closest approach and will extend radially to 40 cm. The beam-pipe itself extends to 25 mm. The pixel detector will cover the region in $|\eta| < 4.0$. Table 9.1 provides a comparison of some key figures between the current ID and the ITk pixel detector.



(a) A single pixel in a hybrid pixel detector. Charge is liberated as an ionising particle traverses the depletion zone. By means of the applied electric field, this charge is collected and then read out by the electronics.



(b) Exploded view of a general hybrid pixel detector. The FE chip is bump-bonded to the sensor, with a read out channel for each pixel.

Figure 9.9: Schematics showing the general makeup of a hybrid pixel detector [238].

The high fluences expected at the HL-LHC have several effects on the pixel modules which make up the subsystem, stemming from the damage to the silicon crystal lattice. Higher bias voltages are required to deplete the sensor. The leakage current or dark current — the current drawn in the absence of any incident particles — increases, which can potentially introduce a thermal runaway effect as a result of a feedback loop between the operating temperature of the device and the leakage current. The module's charge collection capabilities are reduced, which can have a direct effect on the hit efficiency and increase the fake rate. This happens when there are hits originating from electrical noise rather than real interactions between incident particles and the sensor. The expected fluences across the ITk are shown in Figure 9.10, with the pixel detector expected to experience up to $9 \times 10^{15} n_{\text{eq}} \text{ cm}^{-2}$.

Figures 9.5 and 9.8 illustrate the very high particle and thus track densities which are expected at the HL-LHC. This presents a challenge for the tracking algorithms to the detriment of their performance in that it is much more difficult to resolve tracks from one another in such a busy environment. The solution to this is to increase the granularity of the detector by shrinking the pixels on each device, with respect to the ID. This comes at a cost, however, in that the number of channels which must then




Figure 9.10: Expected fluence across the ATLAS detector over the lifetime of the HL-LHC [239]. It is worth noting that the values are given per fb^{-1} of integrated luminosity.

be read out is increased by more than a factor of five in the ITk pixel detector, so that readout rates and bandwidths must be able to cope with this.

As well as altering the technologies of the pixel detector system to maintain performance in the new environment generated at the HL-LHC, the pixel detector in the ITk will provide greater coverage in the forward regions of the detector — out to $|\eta| < 4.0$ — than in the ID. This is motivated by the benefits this extended coverage brings to the measurement of certain physics processes, as well as improved rejection of pileup.

Of course, each of these requirements and associated upgrades is constrained by the detector itself; by power requirements, space for cooling, readout, powering services, and by the material budget, which must be minimised. The material budget has a direct impact on the performance of the tracker. All material a particle has to traverse impacts its trajectory. The more material presented to the particles emanating from the IP, the lower the measurement efficiency and higher the fake rate. For the ITk pixel detector, the material budget is kept low by using thin sensors, thin FE readouts, serial powering of modules, and adjustment to the layout of certain layers. New for the ITk is the use of inclined pixel planes for the outer three layers in the first section of end-cap disks, visible in red in Figure 9.6 and the centre of Figure 9.7. This tilt in the modules ensures that particles cross those layers at an incident angle closer to perpendicular and therefore reduce the amount of material they must travel through.

The spatial resolution — and by extension, the resolution on reconstructed tracks — of a pixel detector is determined in large part by the size of the pixels. This is vital for the detector's ability to resolve different tracks. Having high spatial resolution close to the beamline is especially important for the vertex reconstruction required for b -tagging. The granularity of the ITk pixel detector is greatly improved with respect to the ID as a result of having smaller pixel pitch. With the exception of the innermost layer, pixels have pitch $50 \times 50 \mu\text{m}^2$. Rectangular $25 \times 100 \mu\text{m}^2$ pixels are employed closest to the beamline in order to maximise resolution on the transverse impact parameter, d_0 .

Visible in Figure 9.11, the innermost layer of the ITk pixel detector will also make use of 3D sensors, as was done for portions of the ID IBL. In these sensors, electrodes are inserted into the depth of the silicon rather than on the surface as in the case of the planar sensors. This has several effects. The charge liberated in the bulk drifts horizontally and has a much shorter distance to travel through the silicon to the electrode than in the planar sensor. In the 3D sensor the distance travelled is determined by the spacing of the electrodes whereas in the planar sensor it is the thickness of the silicon. This means that 3D modules can be operated at lower bias voltages. Critically, 3D sensors have heightened radiation hardness as compared to their planar counterparts as the effects of charge-trapping are minimal over that short drift distance. Despite this, the two innermost layers are designed to be completely replaced after 2000 fb^{-1} . All sensors are silicon, with the planar sensors $150 \mu\text{m}$ thick and the 3D sensors in the innermost layer $100 \mu\text{m}$ thick. Both 3D and planar sensors use n-type implants in a p-type substrate produced using a single-sided process.

The readout chip to be used in the ITk pixel detector will be called the ITkPix [239], which will have dimensions of 400×384 pixels. This is based on the RD53 readout chip developed by the RD53 Collaboration [241] — a joint effort between the ATLAS and CMS collaborations.

Figure 9.12 shows the structure of the basic unit of the ITk pixel detector — the module. For all but the L0 3D modules, pixel modules are *quads*. This means that they are built from four ITkPix FEs which are flip-chip bump-bonded to the sensor. This constitutes a *bare* module. A flexible PCB is then glued to the surface of the sensor and wire-bonded to the FEs to create a complete, assembled module. The L0 3D modules are *triplets*. These are three single bare modules — three sensors bump-bonded to three FEs — with a single triplet flex glued on top.




Figure 9.11: The two different types of sensors employed by the ITk pixel detector [240]. 3D sensors are favoured in the innermost layers because of their robustness against radiation damage.




Figure 9.12: Exploded schematic of an ITk pixel quad module [239]. Using the flip-chip process, four ITkPix FE chips are bump-bonded to a single sensor. A flexible PCB is then glued on top and wire-bonded to the FE chips.

With all this put together, the pixel detector is required to have track reconstruction efficiencies above 99% and 85% for muons and electrons, respectively, and to maintain a fake rate below 10^{-5} . Pixel occupancy is to be below 1% and performance maintained with up to 15% of channels out of action [236].

Testing that devices meet their design specifications is a challenge undertaken by performing a menagerie of quality control and systems tests. A large part of assessing the performance of the finished modules is taking real data at testbeam. Contributions to this characterisation effort made by the author are presented and summarised next, in Chapter 10.

Chapter 10.

Analysis of Testbeam Data for ATLAS ITk Pixel Modules

Ahead of the installation of the ATLAS ITk, the performance of the pixel modules to be used (described in Chapter 9) must be validated. This chapter describes research into the performance of prototypes of these new pixel modules carried out at testbeam. Results presented here make use of data taken over the course of several testbeam campaigns using various different devices. Described here are measurements made of devices developed by the UK ITk community, with many of the methods developed during work on the ATLAS-wide planar sensor market survey (MS). The effects of different biasing structures, device irradiation, and different operating parameters such as the bias voltage and threshold are all investigated.

10.1. Testbeam setup

Testbeams for the purpose of testing ITk pixel modules have taken place at both the SPS facility at CERN and Deutsches Elektronen-Synchrotron (DESY) in Hamburg. As mentioned in Chapter 2, the SPS acts as a feeder accelerator for the LHC. It also provides a variety of testbeam options by steering the primary 450 GeV proton beam onto an array of targets to provide secondary beams of electrons, muons, and hadrons, as shown in Figure 10.1. Results presented here were produced with an 120 GeV pion beam.

The synchrotron at DESY is used to generate an electron beam of 1–6 GeV in order to probe the performance of the devices [242]. Results presented here used a 4 GeV




Figure 10.1: The testbeam facility at the CERN SPS [243].




Figure 10.2: The testbeam facility at DESY [242].

electron beam. The configuration of the DESY testbeam hall can be seen in Figure 10.2. Electrons incident on a carbon fibre primary target generate Bremsstrahlung photons. These photons undergo e^+e^- pair production when incident on a converter target. A dipole magnet then separates out the electrons which are sent through a collimator and into the testbeam hall.

In the testbeam hall is the *telescope*, comprising six planes. Up to two devices under test (DUTs) are placed in the centre, such that there are three telescope planes on either side. As the beam passes through the telescope, it leaves hits in the telescope planes, as well as each DUT. The hits in the telescope planes are reconstructed into tracks. By comparison of the position of these tracks and the presence of an associated hit in the




Figure 10.3: Simulation of the telescope configuration at DESY generated in Allpix² [244]. The beam enters from the right-hand-side of the frame. The DUTs are represented centrally inside a polystyrene cool box shown in blue.




Figure 10.4: Photograph of the telescope configuration at DESY. The DUTs are contained in a cooling box in the centre of the frame, behind the black piping.

DUT, the *efficiency* of the DUT can be measured. A simulation of this setup can be seen in Figure 10.3.

The beam telescope itself is an EUDET-like [245] setup comprising six MIMOSA-26 [246] pixel detectors, and can be seen in Figure 10.4. Owing to their $18.4 \times 18.4 \mu\text{m}^2$ pixel pitch, these planes have extremely high spatial resolution, which is beneficial to constructing tracks within the telescope that allow hits on a DUT to be determined at a resolution beyond that of the DUT itself. The $50 \mu\text{m}$ thickness of these sensors also means small contributions to multiple scattering.

The trigger logic unit (TLU) takes signals from extremely fast scintillator planes placed at either end of the telescope. In the case of coincident signals in all scintillator planes, a trigger is distributed to all telescope planes and DUTs. Because the telescope planes have an integration time of $115 \mu\text{s}$, which is orders of magnitude longer than the DUT integration time of 25 ns , there is a risk of out-of-time contributions in the reconstructed tracks entering the device efficiency calculation. That is to say that for a given trigger, hits may be read out for a given event which originated from previous triggers. Given that the efficiency is generally defined as

$$\epsilon = \frac{N_{\text{DUT}}}{N_{\text{T}}}, \quad (10.1)$$

with N_{DUT} the number of reconstructed tracks with a hit in the DUT and N_{T} the total number of reconstructed tracks, contributions from out-of-time tracks will result in a lower device efficiency. This is solved by the addition of an additional plane — an FEI4 module with the same integration time as the DUTs — to serve as a timing reference. All tracks to be used for analysis are then required to be associated to a signal in the reference plane. In conjunction with the TLU, the EUDAQ framework then takes the individual streams from each device and merges them into a single output.

The key requisite for the timing reference plane is that it have comparable integration time to that of the DUTs. For the measurements presented in this thesis, an additional FEI4 plane was used as the reference plane, but in the case where two DUTs are being measured, one DUT may be used as the reference plane when measuring the other.

The telescope used was of the same design and layout for both the SPS and the DESY testbeams. The most prominent difference between the two campaigns is the




Figure 10.5: A quad module based on the RD53A. Here, four RD53A readout chips are bump-bonded to a silicon sensor. A flexible PCB is then glued on top of the sensor. Here, the sensor and readout chips can be seen protruding past the edge of the flex, with wire bonds connecting the contacts on the flex to the RD53A chips.

greater effect of multiple scattering at DESY. This is because of the lower energy of the beam and is a contribution to the uncertainty on the track resolution in the telescope.

10.2. Devices under test (DUTs)

The devices being tested using the telescope setup at DESY are hybrid pixel modules built from a passive high-resistivity n -in- p silicon sensor which has a flexible PCB glued on and is then bump-bonded to a front-end readout chip. An example can be seen in Figure 10.5. Presented here are studies on modules utilising different sensor designs, summarised in Table 10.1. All have pixel pitch of $50 \times 50 \mu\text{m}^2$, a thickness of $150 \mu\text{m}$, and have been irradiated to a fluence of $3.4 \times 10^{15} n_{\text{eq}} \text{cm}^{-2}$. This was done using 25 MeV protons extracted from the Karlsruhe Kompakt Zyklotron [247] at the Karlsruher Institut für Technologie (KIT).

Each used a different punchthrough bias (PTB) structure. This biasing structure serves an important function in the production of pixel modules. Given the complexity




Figure 10.6: Possible biasing structures for a pixel detector sensor [240].

DUT ID	Pitch (μm)	Thickness (μm)	PTB
DUT-A	50×50	150	none
DUT-B	50×50	150	zigzag

Table 10.1: Silicon sensor designs for the hybrid pixel modules included in the testbeam campaigns.

of the processes involved, the various components of an assembled module should be tested prior to assembly. The sensor and readout chip, for example, should be electrically tested before they are bump-bonded together in order to maintain high yields in the final collection of modules. Ordinarily, measuring the IV curve of the sensor is usually not possible until after bump-bonding. By incorporating a biasing structure, contact can be made with the pixels on the sensor in order to measure IV curves. Although very useful in production, these structures can also cause the degradation of the sensor efficiency. In spite of the consequences for the efficiency of charge collection, the benefits in production of using PTB structures mean that their removal is highly undesirable. Possible biasing configurations are shown in Figure 10.6. Specifically, the *zigzag* structure shown in Figure 10.6c is expected to mitigate some of the efficiency problems by moving the bias rail closer to the pixel implant. If the bias rail is further away from the pixel implant, it is possible that the depletion region created by the bias does not extend fully, resulting in lowered efficiency. Of the devices presented here, one has no PTB structure and the other the zigzag PTB structure, as per Table 10.1.

The modules presented here utilise the RD53A [248] front-end readout chip which has been developed as the first prototype of the eventual ITk readout chips to be produced en masse. The RD53A has dimensions of 400×192 pixels, and an area of $20.1 \times 11.6 \text{ mm}^2$. It has a thickness of $150 \mu\text{m}$ and was designed to have radiation tolerance up to 500 Mrad. The RD53A was developed with three possible modes — *linear*, *differential*, and *synchronous* — such that it is segmented into three sections corresponding to these modes. This means that a given set of testbeam measurements will utilise only a third of the 400 rows of pixels on the sensor. The devices presented here used either the linear or differential front-ends. ATLAS plans to use the differential front-end, while CMS has opted for the linear mode.

10.3. Testbeam campaigns

The results presented here were produced from data taken over two testbeam campaigns. The first was at the SPS in October 2018, and the second at DESY in December 2018.

10.4. Reconstruction and analysis

Experience working on the ATLAS ITk planar sensor MS led to the overhaul of the framework in use for the reconstruction and analysis of testbeam data within the UK ITk community. An overview of this new framework is shown in figure 10.7.

The raw data is converted and reconstructed by the EUTelescope software, in tandem with the HTCondor batch processing system. Here, parameters such as the cluster size, cluster multiplicity, number of tracks, and track residuals are calculated. Following the reconstruction, the analysis of the reconstructed data is carried out — again on HTCondor — with the tbmon2 software. Here, global and in-pixel efficiencies, amongst other characteristics, may be extracted for the device in question. For both the reconstruction and analysis, a single run is first treated step-by-step locally in order to determine an optimal set of parameters — discussed in Section 10.4.1 — for the rest of the batch.

Following the batch processing, the output is checked for any runs which may have crashed at any point. If a run crashed during reconstruction, the source of the




Figure 10.7: The analysis flow from start to finish in the UK testbeam framework. Through a chain completely contained within a single repository, raw data is taken as input and output as plots of the results.

crash is ascertained and rectified and the run repeated through reconstruction and analysis. If the run was successfully reconstructed but crashed during analysis, then the source of the crash is ascertained and rectified and the analysis step only repeated. This procedure is followed until such a point as all runs have been successfully reconstructed and analysed.

The processed data is then gathered into a JSON file which stores the various pieces of information and associates them to their run, device, batch, bias, threshold, etc. This JSON file then contains the sum total of the analysis efforts and is passed as input to a custom Streamlit web-app which is able to interpret the JSON and display the results in a flexible, user-friendly manner.

10.4.1. EUTelescope

The reconstruction of the testbeam data is carried out using the EUTelescope [249,250] software package. This proceeds in five steps:

1. **CONVERTER:** The first step is converting the raw data into an LCIO common format. Additionally, *noisy pixels* are flagged at this stage. These are individual pixels which register above a user-defined *firing frequency*. This is simply the rate at which a particular pixel registers a “hit”. Cutting on this rate can help to mask certain problematic pixels at the level of the reconstruction software. Here, a firing frequency of 0.1% was adopted for the DUTs and 0.5% for the telescope planes. Some pixels may have already been masked at the level of the hardware in the tuning process during data acquisition. In this case, these pixels cannot be unmasked in the software as they were turned off when the data was taken.
2. **CLUSTERING:** Pixels which fired in the event are grouped into clusters. Each cluster has its centre defined as the charge-weighted centre of the pixels from which it is comprised,

$$\bar{x} = \frac{1}{Q} \sum_{i=0}^N x_i q_i, \quad (10.2)$$

with Q the total charge of the cluster, x_i the index of the i^{th} pixel in the cluster, and q_i the charge of that pixel. These positions are calculated in x and y directions independently, before the results are saved in a new LCIO collection. Positions




Figure 10.8: The cluster size in an example batch of data-taking. It can be seen that the vast majority of clusters are made up of only a few pixels.

at this point are defined in the local frame of reference for that plane. It is at this point that any clusters which contain a noisy pixel are removed from the dataset for the purposes of the next steps in the reconstruction. It is possible at this stage to view the effect of different choices of the firing frequency cut by looking at the number of noisy pixels present for various choices of this parameter, as well as the firing frequency of those offending pixels. As the firing frequency used for this cut is raised, the fraction of noisy pixels on a module decreases. At this point, it is possible to investigate the cluster sizes, both in terms of an overall mean and in x and y directions. An example of this can be seen in Figure 10.8. It is also possible to view the cluster signal and event cluster multiplicity.

3. HITMAKER: The key function of this stage is to transform the coordinate system in use from a local system in terms of pixel positions to a global system in terms of (x, y) positions. It is possible to investigate the correlations between hits in the various planes, as shown in Figure 10.9. These are important to determine that the planes are in the correct orientation. This step also performs a pre-alignment, ahead of the full alignment in the next step, by minimising the difference in positions of hits between the planes.




Figure 10.9: Correlations between a DUT and telescope plane for an example batch of data-taking. In this case, the sensor has pixel pitch of $25 \times 100 \mu\text{m}^2$, which can be seen in the difference in granularity between the two coordinates.

4. **ALIGNMENT:** Plane positions are adjusted using track candidates in this global coordinate system. This is done by minimising the residuals in the telescope planes. These residuals are defined as the difference between the track position and the hit position. These residuals should ideally be gaussian and centred on zero. To perform this alignment of the planes, the *general broken lines* [251] method is used. This helps to mitigate effects from multiple scattering. The user can decide both how many iterations of the alignment procedure to execute and how many track candidates should be used.
5. **FIT:** The final step is to use this updated geometry to then fit the tracks to the DUTs. During this stage, tracks are formed from hits in the aligned telescope planes and compared to hits in the DUT planes. The quality of this fit is assessed by a χ^2 calculation which can be seen in Figure 10.10. The resolution of the DUTs should then be given by the width of the residuals, which can be seen in Figure 10.11. In principle, for binary hits, this should be $\sigma = \frac{\text{pixel pitch}}{\sqrt{12}}$. In practice, this is not achieved due to contributions from multiple scattering, which are more prominent for the data taken at the DESY testbeams. The separate contributions from the intrinsic resolution of the device and the effects of multiple scattering explain why the observed residuals are in fact a convolution between a Gaussian distribution and a uniform distribution.

There are several input files which guide the reconstruction. *Steering* files for each step specify processors which are to be used. A *configuration* file specifies paths for input and output files, as well as setting parameters such as the number of events to




Figure 10.10: χ^2 per degree of freedom distribution for the GBL fit used to build the final collection of tracks in the telescope.




Figure 10.11: Distribution of residuals for an example DUT and batch of testbeam data. The residuals represent the difference in position between the cluster on the DUT and the tracks as built from the telescope. The RMS of the distribution gives the expected resolution in that dimension, which is a convolution of the sensor's spatial resolution and the effects of multiple scattering. In this case, the sensor has pixel pitch $25 \times 100 \mu\text{m}^2$, which can be seen in the difference in the width of the distributions in the two coordinates.

run over or how many tracks to use for alignment. It's also here that cuts such as the firing frequency may be set. The GEAR framework is used for the geometry of the setup, with *gear files* specifying the devices in use and the physical configuration of the setup. The *runlist* contains a list of the *runs* of data taken during the testbeam. This associates parameters to the corresponding run.

At each of these stages in the reconstruction process, there are several outputs. The most notable are the histograms, which contain the distributions which can be investigated at each step, and the *tbtrack* files which are output from the final fitting step. These are passed into the *tbmon2* software for analysis.

Previously, reconstruction of testbeam data was done locally, step-by-step. This was a time-consuming, labour-intensive effort. With the data separated into runs, roughly equal periods of data-taking, each run can take as long as 36 hours to reconstruct, depending on how many iterations of alignment are performed, how many tracks are used, as well as how many events are in the run.

Additionally, the gear files specifying the configuration of the telescope, had to be altered according to the run being analysed. This sort of arrangement also led to, at best, redundancy, with each analyser creating their own gear files and config files. At worst, this could cause divergence in the method being applied to runs by different analysers.

Further redundancy arose with each analyser often investigating the reconstructed data in their own ad-hoc manner, with many often writing their own analysis scripts to achieve the same results as others.

One of the author's primary contributions to the MS work was the establishment of a central framework to address these issues. Here, configuration files, gear files, analysis scripts etc. are hosted on a central Git repository. In this way, the user's installation of EU Telescope is controlled from within this new framework. All gear files, configs, and scripts are here such that they may be shared and standardised across the whole reconstruction effort.

During testbeam, runs are grouped into *batches*. These corresponded to physical configurations of the telescope. As such, gear and config files are created for each batch. This step may even be taken whilst data are being taken such that, as soon as a batch is completed, the reconstruction can begin.

A batch may consist of over 100 runs, for longer batches. If each run takes up to 36 hours to reconstruct, the time taken to perform this reconstruction becomes a bottleneck in the evaluation of the performance of the DUTs and, by extension, the development of these devices. A major advantage to this new software framework has been the development of a method to perform the reconstruction using the HTCondor [252] job scheduling software. Previous attempts had been made to parallelise testbeam reconstruction but none fulfilled the necessary criteria and thus had not been adopted.

Under this new system, an entire batch may be submitted to the HTCondor system at CERN. With the HTCondor software, *jobs* are matched to an available CPU to run that job. By using CERN's HTCondor system, the job scheduler has access to the shared file-system, allowing it to access all data necessary without that data having to be copied to another location first. It may then output back to that same shared file-system such that all members of the team can access the entirety of the input and output as necessary.

In a given batch, each run is sent as a job to the HTCondor system, configured such that each job runs each step in the reconstruction consecutively. In this way, all ~ 100 runs in a batch can be reconstructed simultaneously with minimal input from the user required.

For a given testbeam, there may be $\mathcal{O}(10)$ batches. It is clear to see that this way of performing the reconstruction represents a major step forward and has allowed the analysis of the prototype modules to proceed in a timely manner.

Following the success of this MS framework, a similar project was launched within the UK pixel testbeam community. This built upon many of the ideas implemented in the MS repository and extended them. In the same way, common configuration files, gear files, etc. are generated. As compared to the MS framework, however, the UK framework represents a more collaborative workflow. In the language of the Git version control software, each user is encouraged to first create their own *branch* with which to work. Each branch is a copy of the main project, which maintains its history. The user may then modify and extend the project without impacting the central *master* branch from which others work. When a new feature, configuration, bug-fix, etc. is finished on a branch, these branches may then be periodically reviewed and merged into the master branch of the framework in order that these modifications

are regularly propagated through to each user, provided they stay up-to-date with the master branch.

Another advantage to this new way of performing the reconstruction and analysis of the testbeam data is that a common installation of EUTelescope and tbmon2 can be used by these frameworks as necessary. Although an independent project, the UK framework operates on the same principles as the project developed for the MS; centrally hosted methods and configuration files etc., which are then propagated to EUTelescope and/or tbmon2 when the data are being processed. As such, a user may have a single installation of EUTelescope and tbmon2, each of which is controlled by the testbeam project. In this way, the installations of EUTelescope and tbmon2 are never directly modified. In the case where a single user is working on multiple projects — say the MS and separate UK analyses — they may then modify each of those projects and run their reconstruction and analysis without worrying about interfering with the other. This is analogous to having an operating system with multiple users. Each user may have their own files, settings, etc., without fear of interfering with the other user, but ultimately it is the same operating system carrying out the desired processes, albeit guided by the user. In this analogy, the installations of EUTelescope and tbmon2 are like the operating system, with the UK repository acting as the user.

10.4.2. tbmon2

The analysis of the reconstructed testbeam data is performed using the tbmon2 software. This takes the tbtrack files which are output from the reconstruction in order to calculate many of the metrics used to assess the performance of the sensors.

The most important of these quantities is the global efficiency of the device, given by equation [10.1](#)

Another important quantity is the *in-pixel* efficiency, which gives the efficiency across the area of an individual pixel. This can be seen in Figure [10.12](#), which shows an in-pixel efficiency map. This can be used to identify factors affecting the efficiency of the sensors. Biasing structures, for example, may sometimes be seen as localised areas of low efficiency between pixels.

There has been some attention devoted to defining a *fiducial region* for the measurement of sensor efficiencies which account the effect of the PTB structure. Results




Figure 10.12: Example of an in-pixel efficiency map. In this case, the entirety of the sensor is divided into 4×4 pixel (pitch $50 \times 50 \mu\text{m}^2$) blocks, with the map formed from the average of these blocks. It is then possible to see how the efficiency varies across individual pixels. The location of the punchthrough bias dots and biasing rails is also shown, as well as the position of the pixels.

are then to be quoted both globally and within the fiducial region for devices which incorporate these features.

10.5. Results

Of the devices listed in Table 10.1, DUT-A was measured at both October 2018 and December 2018 testbeam campaigns, with both linear and differential front-ends tested. DUT-B — the device with zigzag PTB — was measured at the December 2018 testbeam, with linear front-end.

The analysis aimed to evaluate the efficiency of the sensor in a variety of ways. Results are presented as:

1. Global efficiency values giving the inclusive efficiency across all pixels on the sensor. As mentioned above, a fiducial efficiency is also provided for devices incorporating a PTB structure. For this fiducial region, the areas containing the punchthrough dots are removed.
2. In-pixel efficiency maps showing the efficiency at a sub-pixel level. In order to accumulate the necessary number of events, the map actually represents an averaging over the surface of the sensor. The sensor is divided into 4×4 pixel blocks, with each position on the final in-pixel efficiency map being an average of the values across all of these blocks.

The parameters which vary in these results, whether between different devices or operational parameters, are:

1. V_{bias} in general increasing the bias voltage will improve the signal-to-noise ratio of the device and thus one would expect the efficiency to improve in tandem. Irradiated devices may be required to operate at higher V_{bias} in order to compensate for effects such as induced charge trapping within the bulk.
2. **Threshold:** the threshold which is set for a given device must fall in something of a *Goldilocks zone*. It must be set high enough to clear the noise floor and thus avoid fake hits. Equally, it must set low enough that the charge produced in the sensor by an incident particle passes that level and is registered as a hit. Again, devices which have been irradiated will likely require a higher threshold than those which have not, as the noise levels will rise in line with radiation damage.
3. **Punchthrough bias structure:** devices with PTB structures are expected to exhibit lower efficiency as a result of charge collection at the biasing rail. The magnitude

of this efficiency drop, as well as the extent to which it is mitigated by different PTB structures and how localised these effects are, are all unclear.

Results are evaluated with reference to the criteria set for the MS:

- 97% for 150 μm sensor irradiated to $2 \times 10^{15} n_{\text{eq}} \text{cm}^{-2}$ at 400 V_{bias}
- 97% for 150 μm sensor irradiated to $5 \times 10^{15} n_{\text{eq}} \text{cm}^{-2}$ at 600 V_{bias}

Note that results shown in this section do not have error bars on the plots as they are small enough to be invisible and inconsequential.

10.5.1. October 2018

Figure 10.13 shows a threshold scan for the device without PTB structure and at $V_{\text{bias}} = 600 \text{ V}$. At lower threshold, the efficiency for both front-ends is at $> 99\%$, comfortably above the required level. As the threshold is increased, this efficiency drops below the required level. This is likely because at that level, most of the noise has already been cut away, but signal is beginning to drop below the threshold, impacting the efficiency.

Although both front-ends exhibit this drop in efficiency as the threshold is increased from around 1200 e^- to around 1600 e^- , it appears more pronounced in the differential front-end. For the linear front-end, when the threshold is increased further to 2100 e^- , the efficiency drops again.

Figures 10.14 and 10.15 show the in-pixel efficiency maps for these same data points, for the differential and linear front-ends, respectively. The efficiency drops in a reasonably uniform manner across the sensor in the case of the differential front-end. In the case of the linear front-end, the efficiency clearly drops significantly in the corner regions of the pixels as a result of charge-sharing, driving the drop of the overall efficiency for the device despite the other areas retaining a high efficiency.




Figure 10.13: Device efficiency as a function of threshold, in number of electrons. Shown here are linear and differential front-ends, at $V_{\text{bias}} = 600$ V, for the module with no PTB structure.




Figure 10.14: In-pixel efficiency map representing the average efficiency of a 4×4 pixel array scanned across the full sensor. Shown here is the module with no PTB structure and using the differential front-end, at $V_{\text{bias}} = 600$ V, with varied threshold.




Figure 10.15: In-pixel efficiency map representing the average efficiency of a 4×4 pixel array scanned across the full sensor. Shown here is the module with no PTB structure and using the differential front-end, at $V_{\text{bias}} = 600 \text{ V}$, with varied threshold.

10.5.2. December 2018

Figure 10.16 shows a threshold scan for both the device using zigzag PTB and the device using no PTB, both using the linear front-end, with each line representing a different V_{bias} . The behaviour exhibited is broadly similar — as the threshold is raised, the efficiency rises before reaching a plateau, reflecting the point at which most noise has been removed but before any significant cutting of signal. This behaviour is seen most clearly for the device with zigzag PTB. The effect of a higher V_{bias} appears to be to shift the point of this plateau to a lower threshold, as would be expected. In other words, a greater V_{bias} means that the device reaches its efficiency plateau at a lower threshold. Although for some of the other entries, there are too few points to observe this same behaviour, it seems reasonable to assume it would hold for the other configurations as well. In the case of the scan for the device with zigzag PTB at $V_{\text{bias}} = 400 \text{ V}$, given the rise in efficiency in the step from $\text{THL} = 1027 e^-$ to $\text{THL} = 1200 e^-$ and that it reaches almost 98% efficiency at the latter threshold, in light of the expected lower efficiency level of the plateau for that device given the PTB, it seems extremely unlikely that it would see a further gain on increasing the threshold again. For the device without PTB, it would seem that at the points in the threshold scan have already reached this plateau, which is at a higher efficiency than for that with the PTB structure. Again, this is to be expected. One would expect that points at a lower threshold would reveal similar behaviour for this device. As with the threshold scan in Figure 10.13 for the October 18 testbeam, one would expect in all cases that the efficiency of the devices would drop if the threshold continued to be raised. For all devices and V_{bias} points in this threshold scan for the December 2018 testbeam, the 97% required efficiency is satisfied.




Figure 10.16: Device efficiency as a function of threshold. Shown are both the device with zigzag PTB, and the device with no PTB structure, each at $V_{\text{bias}} = 400 \text{ V}$, $V_{\text{bias}} = 500 \text{ V}$, and $V_{\text{bias}} = 600 \text{ V}$. All measurements are with the linear front-end.




Figure 10.17: In-pixel efficiency map representing the average efficiency of a 4×4 pixel array scanned across the full sensor. Shown here is the module with zigzag PTB structure and using the linear front-end, at $V_{\text{bias}} = 600 \text{ V}$, with different thresholds.

Figure 10.17 shows the in-pixel efficiency maps for the threshold scan in yellow in Figure 10.16, corresponding to the device with zigzag PTB structure at $V_{\text{bias}} = 600 \text{ V}$ and using the linear front-end. At $\text{THL} = 900 \text{ e}^-$, it is clear that the device has not reached the plateau, is still under significant noise influence, and has an efficiency which is relatively low across the sensor. In the subsequent two threshold points, once the efficiency plateau has been reached, the regions of low efficiency around the punchthrough dots are clear to see. Despite this, as the rest of the sensor achieves a very high efficiency, the device comfortably crosses the 97% required.

Figure 10.18 shows a bias scan for both the device using zigzag PTB and the device using no PTB, both using the linear front-end. Both exhibit the same behaviour in that

their efficiency increases with the V_{bias} before reaching a plateau. The device with the zigzag PTB structure reaches this plateau slightly earlier and that plateau is located at a lower efficiency. This is likely due to the contributions from the regions of low efficiency around the punchthrough dots. Despite this, both the device with zigzag PTB and that without reach the 97% efficiency required, rising to 98.4% and 99.7% respectively. This means that both devices reach the required efficiency below the required V_{bias} .

Although both devices are at constant threshold during the bias scan, they are not at the same threshold as each other. The device with no PTB is at $\text{THL} = 1027 \text{ e}^-$, whilst the one with PTB is at $\text{THL} = 1200 \text{ e}^-$. When making a comparison between the two to assess the impact of the PTB, this must be taken into account. Whilst not a quantification, it is likely that the gap in efficiency between the two is actually slightly larger than it appears in Figure 10.18. This is because a threshold of $\text{THL} = 1200 \text{ e}^-$ is likely to yield a slightly better efficiency than a threshold of $\text{THL} = 1027 \text{ e}^-$, on consideration of the threshold scans shown in Figure 10.16. This effect, whilst it must not be forgotten, is also likely to be near-negligible, however. Although there is a significant increase in efficiency observed for the device with the PTB structure when stepping from $\text{THL} = 1000 \text{ e}^-$ to around $\text{THL} = 1200 \text{ e}^-$, this device is already at a threshold of $\text{THL} = 1200 \text{ e}^-$ in Figure 10.18. Given that the device with no PTB exhibits minimal gain in efficiency in the step from around $\text{THL} = 1000 \text{ e}^-$ to around $\text{THL} = 1200 \text{ e}^-$, it can be assumed that even if it were at $\text{THL} = 1200 \text{ e}^-$ rather than $\text{THL} = 1027 \text{ e}^-$, its efficiency would be relatively unaffected and therefore of minimal consequence to the above comparison.

Figures 10.19 and 10.20 show the in-pixel efficiency maps corresponding to the bias scan in Figure 10.18, for devices with and without PTB structure respectively, at $\text{THL} = 1200 \text{ e}^-$ and $\text{THL} = 1027 \text{ e}^-$. For the device with PTB, this structure is immediately apparent at low V_{bias} as localised areas of low efficiency, even against a background of relatively low efficiency across the rest of the sensor. As the V_{bias} is increased and the efficiency across the rest of the sensor rises and reaches the plateau, the area around the punchthrough dots persists as localised low efficiency. In spite of these regions, the device comfortably passes the 97% efficiency required by $V_{\text{bias}} = 400 \text{ V}$, the benchmark V_{bias} . For the device without PTB, it is at low V_{bias} that the structure is most apparent. Although the efficiency across the sensor is relatively low, the pixel corner regions appear as regions of particularly low efficiency. When the V_{bias} is raised and the device reaches the efficiency plateau, these localised low




Figure 10.18: Device efficiency as a function of V_{bias} . Shown are both the device with no PTB at $\text{THL} = 1027 \text{ e}^-$ and with zigzag PTB at $\text{THL} = 1200 \text{ e}^-$.

efficiency regions retreat and a high efficiency is achieved uniformly across the sensor. The device comfortably passes the 97% required, reaching near 100% efficiency.




Figure 10.19: In-pixel efficiency map representing the average efficiency of a 4×4 pixel array scanned across the full sensor. Shown here is the module with zigzag PTB structure and using the linear front-end, at $\text{THL} = 1200 \text{ e}^-$, with varied V_{bias} .




Figure 10.20: In-pixel efficiency map representing the average efficiency of a 4×4 pixel array scanned across the full sensor. Shown here is the module with no PTB structure and using the linear front-end, at $\text{THL} = 1027 \text{ e}^-$, with varied V_{bias} .

10.6. Conclusions

Results have been presented for prototype pixel detector modules being developed for use in the ATLAS ITk. These are hybrid detectors formed from a $150 \mu\text{m}$ sensor with different punchthrough bias structures bump-bonded to an RD53A readout chip using both linear and differential readout modes. Each module was irradiated to $3.4 \times 10^{15} n_{\text{eq}} \text{ cm}^{-2}$.

Over two testbeam campaigns at the SPS and DESY in October 2018 and December 2018 respectively, measurements were taken using an EUDET-like telescope setup. The data taken at these two testbeam campaigns was reconstructed using EUTelescope and analysed with tbmon2. A new software framework was developed to standardise and automate this process.

A series of bias and threshold scans were performed on the devices and in-pixel efficiency maps constructed, with each device required to surpass a 97% efficiency at the required V_{bias} .

For both bias scans and threshold scans, the efficiency of both devices is seen to increase as the bias or threshold is increased until it levels out at a plateau. If the threshold is then increased further, it can be seen to begin to cut signal and thus degrade the efficiency. The device without any PTB structure is seen to reach its efficiency plateau for lower bias or threshold than for the device with zigzag PTB. The efficiency plateau for the device with no PTB is also at greater efficiency than for the device with zigzag PTB.

In-pixel efficiency maps reveal how structures at a sub-pixel level on the sensor can affect the efficiency. For the device with PTB, low efficiency regions around the punchthrough dots are consistently observed and ultimately limit the efficiency of the device. For the device with no PTB, the corner regions of pixels are seen to have a lower efficiency than the rest of the pixel, due to the effects of charge sharing. Unlike the regions around punchthrough dots, however, these localised dips in efficiency are overcome by increased threshold or V_{bias} , leaving a sensor with a uniformly high efficiency.

Whether no PTB or zigzag PTB, both devices achieved the required 97% efficiency.

Chapter 11.

Conclusion

“...while we can never have sufficiently good arguments in the empirical sciences for claiming that we have actually reached the truth, we can have strong and reasonably good arguments for claiming that we have made progress towards the truth.”

— Karl Popper

This thesis has presented the measurement of jet substructure observables as both differential and double-differential cross-sections in the ℓ +jets channel of boosted $t\bar{t}$, using 139 fb^{-1} of $\sqrt{s} = 13\text{ TeV}$ pp collision data collected with the ATLAS detector at the LHC. Testbeam measurements of prototype pixel detector modules as part of the R&D process towards the upgrade of the ATLAS inner tracking system were also presented.

Substructure observables were measured in $t\bar{t}$ events selected to contain a top-quark which decayed hadronically and was reconstructed as a large- R reclustered jet with high transverse momentum. Backgrounds were reduced through a series of selection criteria, including a cut on the mass of the lepton- b -jet system from the leptonically decaying top quark which reduces interference with the single top tW process and reduces bias on the measured jet by replacing a b -tagging requirement on that large- R jet.

Eight substructure observables were selected for measurement, and tracks associated to the reconstructed large- R jets were used for the calculation of these observables, providing improved resolution over using calorimeter clusters. Double-differential dis-

tributions were constructed by further dividing the substructure spectra into kinematic regions of the mass of the large- R jet and of its p_T .

The relative cross-sections for each of these distributions were extracted at particle-level after the application of an unfolding procedure used to remove the effects of interaction with the detector. They were then compared to several NLO+PS predictions, including some with variations of the final state radiation (FSR).

Agreement with the data varies for different substructure observables and different NLO+PS predictions. The nominal POWHEG+PYTHIA8 prediction generally describes the ECF2, D_2 , and τ_{21} observables but predicts more top-like distributions for τ_{32} , τ_3 , and C_3 than are observed in the data, and does not model $p_T^{d,*}$ well.

The POWHEG+ HERWIG7 model, which uses a different parton shower and hadronisation model, generally provides better agreement with the data, with the exception of τ_3 , and the double-differential measurements. Similarly, the model with alternative matrix element generation, MADGRAPH_AMC@NLO+PYTHIA8, provides generally better agreement with the data than the nominal sample. The nominal prediction is mostly in better agreement with the data when using the FSR-down variation, and poorer agreement with the FSR-up variation. The double-differential distributions exhibit poor modelling across all predictions.

These results underline the need for improvements to the modelling of jet substructure in the high-momentum regime for $t\bar{t}$ production at the LHC.

When presenting the physics motivations for building ATLAS and CMS at the LHC, there was no mention of jet substructure. The kinds of precision techniques now using jet substructure were not thought to be possible in the dense, complex environment present at the LHC. These techniques are being used in ever more complex environments to probe ever more subtle effects in ever more precise ways. A few such recent examples are ALICE's observation of the dead cone effect in QCD [104], which may be possible in $t\bar{t}$ at ATLAS using the Run 3 dataset [253], the exploration of the Lund jet plane [103], which again is being investigated in $t\bar{t}$, the advances in top-tagging using both high-level substructure observables and low-level constituents [140, 141] in tandem with machine learning techniques, and the optimisation of α_s in simulation using jet substructure measurements [122, 123].

Figure 11.1 shows the trend in publications on jet substructure for each year from 2000 onwards. If there exists a correlation between trends in publication and interest




Figure 11.1: Number of publications on [INSPIRE-HEP](#) involving the phrase “jet substructure” for each year of the 21st century up to the end of 2022.

within the field, then it is clear that interest in jet substructure measurements and applications of jet substructure techniques has only been rising since the advent of the LHC era, with no sign of slowing down any time soon. The continuing exploitation of such methods will allow ATLAS to improve the precision of measurements and sensitivity of searches over the years to come.

In the testbeam measurements presented here, prototype devices were measured over testbeam campaigns conducted at CERN and DESY, towards the development of pixel detector modules for the upgraded ATLAS tracking system for the HL-LHC, the ITk. These devices were irradiated to the level expected over their operation lifetime. Devices were used which had a punchthrough biasing structure used for quality control during the production phase, and compared to devices with no such structure.

Data was recorded using an EUDET-like beam telescope, and reconstructed and analysed using a new software framework using the HTCondor batch computing system with EU Telescope and tbmon2, in a multi-staged approach.

The performance of the pixel detector modules was assessed by the global device efficiency, as well as using in-pixel efficiency maps in order to visualise localised effects across the sensor. Results were presented both as a function of the device threshold and V_{bias} , and required to meet criteria set out during the ATLAS ITk planar pixel market survey such that the efficiency was above 97%.

At the required threshold and V_{bias} settings, all devices are seen to meet the efficiency criteria. Localised dips in the efficiency are observed near to the punchthrough dots on the sensor, but these are not expressed in a significant way.

The results of these testbeam measurements were fed into the broader set of considerations being made when deciding on the final configuration of the ITk pixel detector modules and were able to demonstrate that neither the radiation hardness of the devices nor the use of punchthrough biasing structures during production presents a problem for their use at the HL-LHC. Furthermore, the additional techniques developed for the reconstruction and analysis of these data have been adopted by the wider community and will continue to reap benefits for future measurements.

Appendix A.

ATLAS Virtual Visits

“What these approaches have in common is the conviction that science from the point of view of only the dominant social groups is at best partial and at worst delusional.”

— Geoffrey Gorham, *Philosophy of Science* [254]

Detailed in this appendix is the work of the author on helping to develop the ATLAS Virtual Visit program, and coordinating that program amongst UK institutes. Although not a part of the ATLAS physics program, such public engagement activities are of great importance and the author is indebted to science communicators for fostering his interest in science in the first place. For this work, the author was awarded the University of Glasgow’s *School of Physics and Astronomy Public Engagement Prize 2022*.

Virtual visits to the ATLAS experiment present a way for those who are otherwise unable to visit CERN in person to see the detector and have their curiosities sated by physicists working on the experiment.

It is not always possible, for example, for school groups to afford to travel to CERN. Many groups may not consider this sort of trip when planning allocation of resources to an excursion. There are, of course, also practical challenges for larger groups wishing to visit. Only six persons are allowed to descend to see the ATLAS detector at one time, for example. Whilst not a like-for-like replacement, virtual visits are a fantastic alternative which allow such groups to be reached and engage with ATLAS, the science being conducted at CERN, and the scientists conducting it.

In the past, these virtual visits have been conducted from the ATLAS Visitor Centre (AVC). The AVC was equipped with cameras and microphones and located adjacent to the ATLAS control room, separated by only a glass division. Visits were not to the detector itself but entailed perhaps a presentation, mention of the control room, and some time for questions and discussion. Visits to the detector itself were carried out in smaller groups in person. For around two years, the AVC has been undergoing renovation and has therefore been unavailable for visits.

In addition to this, the COVID-19 pandemic stopped the possibility of in-person visits in March 2020 and these have not been possible to date. The training courses associated with access to the ATLAS cavern and for guide purposes had only begun late in January of 2020 and were in high demand, with a long waiting list. There was therefore only a short window in which it was possible to complete these courses and gain the requisite permissions. On top of this, any access to CERN at all, even for those working there, was highly restricted from March 2020 until late in the Summer that year.

For various reasons, many of those physicists working on ATLAS who may ordinarily have wished to take available outreach opportunities at CERN found themselves displaced from Geneva and thus unable to participate in any such activities, should they be possible.

All of this conspired against opportunities to showcase the ATLAS detector. Although the author nominally took over the role of UK coordinator for ATLAS virtual visits at the end of March 2020, there was very little activity involving visits to the detector until later in the year for the above reasons.

Beginning at the end of the Summer, it was decided that a strategy for the resumption of these virtual visits should be devised. Due to the closure of the AVC, it was decided that virtual visits should take place from the cavern itself, in front of the detector. It seemed natural then, that their format should take on something more akin to that of a traditional in-person visit to the detector, with associated exposition, rather than the more general overview of the collaboration which the virtual visits had often offered.

The cavern presents a few hurdles to overcome which had not been present. Unlike the AVC, the cavern is not equipped with audiovisual equipment. On this front there are two main issues.

First, the detector is big. Although the cavern it occupies is also big, the detector almost completely fills it. This makes encapsulating the entirety — or anything approaching the entirety — of the detector in a single shot virtually impossible. Even when just at one end of the detector, a guide must present almost with their back against the wall. A setup that can make use of a wide-angle lens, in order that as much of the detector be in shot as possible, and allows for some mobility, to allow for a few different angles showcasing different parts of the detector, was therefore deemed desirable.

Secondly, the cavern is loud. The many cooling systems and perennial throng of busy technicians, engineers, etc. working on the detector mean that the cavern has a rather high level of constant background noise, amidst which even the most enthusiastic voices of science communication can be lost. A headset or microphone for the presenter would solve this in large part, in having the microphone close to their mouth. A second criterion would be for this to be wireless so that the presenter is not restricted in their movement either.

With consideration of these criteria and after some testing of various options and configurations, a wireless lapel microphone system was purchased along with a compact tripod with smartphone fixing. As in figure [A.1](#), the presenter may speak through the lapel microphone whilst listening to other participants with a pair of bluetooth earphones whilst interfacing a smartphone with a videoconferencing software such as Zoom in order to conduct the visit in a flexible way. Using a smartphone equipped with a wide-angle lens, here the Samsung Galaxy S10+, both the presenter and detector may be in shot together whilst the presenter is also able to move around different parts of the cavern.

Although highly customisable depending on the audience, event, and presenter, the format of these visits until now has been broadly similar. Following a brief introduction between the ATLAS personnel and participants, an overview talk is given on particle physics as a field, CERN as an organisation, the Large Hadron Collider, and the ATLAS Experiment. This may be delivered by the presenter from the cavern or by another ATLAS physicist connected to the call from their office or home. Following this, the *visit* section takes place with a live view of the detector and an explanation of the various subsystems, ideally with some linking explanation to the broader physics aims discussed in the beginning. After this, some time is left at the end for a Q&A session.




Figure A.1: A virtual visit to the ATLAS detector in progress, March 2021

For this final part, and for the overview talk at the beginning, a shortlist has been drawn up of ATLAS physicists at UK institutes who have expressed an interest in being involved in this particular outreach activity. Ahead of a virtual visit then, a call may go out on the associated mailing list looking for some of these people to volunteer to join a panel for the Q&A session. Many of those on this list are early career researchers who may otherwise be at CERN on a long term attachment (LTA).

Not only does this allow these displaced physicists to participate in and contribute to these efforts, but the participants are able to engage with a broader cross-section of the collaboration and hear from members with vastly different backgrounds, research interests, and paths to their current position.

From revamping the project in late 2020, until departing CERN at the end of the summer of 2021, the author conducted virtual visits to the benefit of over 4000 people if one also includes viewers of live streams as well as participants in the video call. These visits included, but were not limited to:

- UK Government Department for International Trade and the Department for Business, Energy, and Industrial Strategy. Guests included Dr Mike Short CBE, the

Chief Scientific Adviser for the DIT, and Eleanor Baha MBE, Deputy Director in the Department for International Trade's office in the British Embassy in Geneva.

- The *Frontiers Project*, an online school for science teachers. This was streamed in parallel to Facebook Live, recorded [here](#).¹ Comment from the organisers: "The streaming was followed by 3700 persons! It must be a record high! In addition about 100 teachers followed the sessions and created 9 new educational scenaria for all student ages."
- Several university groups and societies.
- Several schools across all nations of the UK.
- Cardiff Science Festival.
- An event with the University of Glasgow for LGBT History Month.

The specifics of the methods mentioned here, along with some video resources which may be used for virtual visits, are documented in detail [here](#)², with the idea that anyone wishing to become involved with ATLAS virtual visits may use it for reference.

Figure [A.2](#) shows a poster designed and presented on behalf of the ATLAS Collaboration which makes the case for the ATLAS Virtual Visit program as a fantastic tool for public engagement, particularly for reaching through to groups who may not ordinarily be exposed to CERN. This poster was presented at [Lepton-Photon 2021](#) and [Institute of Physics HEPP & APP 2022](#), as well as the [ATLAS UK 2022](#) meeting.

¹<https://www.facebook.com/frontierseu/videos/339226680553930>

²<https://cernbox.cern.ch/index.php/s/SaR6ew46JfT3fML>

ATLAS VIRTUAL VISITS

...where worlds collide!

Scan the QR code on the left for more information on ATLAS virtual visits and to book one of your own. Scan the QR code on the right for a video giving a taste of what an ATLAS Virtual Visit can look like!

Since the programme's inception, ATLAS Virtual Visits have been given for thousands of people. 2021 saw 155 visits to over 30 different countries! The ATLAS Virtual Visits programme is unique in its ability to reach groups unable to travel to CERN. For much of the pandemic, that's been everyone interested in ATLAS!

atlas.cern/discover/visit/virtual-visit

Find out more and book now!

Virtual Visits can be tailored to any audience; school students, science festivals, government departments, university societies, etc. They are a fantastic avenue for groups without the means to visit Geneva to engage with ATLAS physics and physicists.

https://youtu.be/uR_tntr768

Watch a mock virtual visit!

Virtual Visitors can tour ATLAS from anywhere in the world. Multiple groups can also participate from multiple remote locations simultaneously. Flexibility is one of the many advantages of the ATLAS Virtual Visits programme!

Virtual Visits can be booked at a time convenient to the visitors and their guide. Regular Open Visits are also held which are open for registration for any interested member of the public!

Guides can give their tour with a smartphone and videoconferencing software. They then have the flexibility to travel around the detector demonstrating its various components. There is also a bespoke audiovisual setup in the brand new ATLAS Visitor Centre!

Guides represent a wide variety of nationalities, backgrounds, institutions, and physics interests. A typical Virtual Visit might begin with an overview of particle physics, CERN, the LHC, and ATLAS, before the live demonstration of the detector and finally a Q&A.

1. 2. 3. 4. 5. 6. 7. 8.

ATLAS EXPERIMENT
Adam Rennie on behalf of the ATLAS Collaboration

Lepton-Photon 2021

University of Glasgow
adam.ennie@cern.ch

Figure A.2: Poster on the ATLAS Virtual Visit program [255].

Colophon

This thesis was written in L^AT_EX using the “hepthesis” class [[256](#)].

Bibliography

- [1] ATLAS Collaboration, *Measurement of jet substructure in boosted $t\bar{t}$ events with the ATLAS Detector using 140 fb^{-1} of $\sqrt{s} = 13\text{ TeV}$ pp collisions*, [ATLAS CONF Note](#) (2023) . Cited on p. [iii](#).
- [2] ATLAS Collaboration, *Measurements of differential cross-sections in top-quark pair events with a high transverse momentum top quark and limits on beyond the Standard Model contributions to top-quark pair production with the ATLAS detector at $\sqrt{s} = 13\text{ TeV}$* , [JHEP 06 \(2022\) 063 \[2202.12134\]](#). Cited on pp. [iii](#), [74](#), and [113](#).
- [3] A. Buckley, C. White and M. White, *Practical Collider Physics*. IOP Publishing, Dec., 2021, [10.1088/978-0-7503-2444-1](#). Cited on pp. [vi](#), [18](#), [19](#), and [33](#).
- [4] M. Hanlon, <https://www.dailymail.co.uk/sciencetech/article-1052354>, *Are we all going to die next Wednesday?*, 2008. Cited on p. [1](#).
- [5] F. Macrae, <https://www.dailymail.co.uk/sciencetech/article-1051070>, *Landmark experiment to unlock secrets of Big Bang could cause end of world, say scientists in court bid to halt it*, Sept., 2008. Cited on p. [1](#).
- [6] ATLAS Collaboration, *Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC*, [Phys. Lett. B 716 \(2012\) 1 \[1207.7214\]](#). Cited on pp. [1](#) and [6](#).
- [7] CMS Collaboration, *Observation of a new boson at a mass of 125 GeV with the CMS experiment at the LHC*, [Phys. Lett. B 716 \(2012\) 30 \[1207.7235\]](#). Cited on pp. [1](#) and [6](#).
- [8] S. Weinberg, *A Model of Leptons*, [Physical Review Letters 19 \(1967\) 1264](#). Cited on p. [6](#).
- [9] A. Salam and J. C. Ward, *Weak and electromagnetic interactions*, [Il Nuovo Cimento \(1955-1965\) 11 \(1959\) 568](#). Cited on p. [6](#).

- [10] G. t. Hooft, *Renormalizable Lagrangians for massive Yang-Mills fields*, *Nuclear Physics B* **35** (1971) 167. Cited on p. 6.
- [11] F. Englert and R. Brout, *Broken Symmetry and the Mass of Gauge Vector Mesons*, *Physical Review Letters* **13** (1964) 321. Cited on p. 6.
- [12] P. W. Higgs, *Broken Symmetries and the Masses of Gauge Bosons*, *Physical Review Letters* **13** (1964) 508. Cited on p. 6.
- [13] https://upload.wikimedia.org/wikipedia/commons/0/00/Standard_Model_of_Elementary_Particles.svg, *Standard Model of Elementary Particles*, Sept., 2019. Cited on p. 7.
- [14] CMS Collaboration, *Determination of the strong coupling constant $\alpha_s(m_Z)$ from measurements of inclusive $W^\pm$ and Z boson production cross sections in proton–proton collisions at $\sqrt{s} = 7$ and 8 TeV*, *JHEP* **06** (2020) 018 [1912.04387]. Cited on p. 9.
- [15] CMS Collaboration, *Measurement of $t\bar{t}$ normalised multi-differential cross sections in pp collisions at $\sqrt{s} = 13$ TeV, and simultaneous determination of the strong coupling strength, top quark pole mass, and parton distribution functions*, *Eur. Phys. J. C* **80** (2020) 658 [1904.05237]. Cited on p. 9.
- [16] CMS Collaboration, *Measurement of the ratio of the inclusive 3-jet cross section to the inclusive 2-jet cross section in pp collisions at $\sqrt{s} = 7$ TeV and first determination of the strong coupling constant in the TeV range*, *Eur. Phys. J. C* **73** (2013) 2604 [1304.7498]. Cited on p. 9.
- [17] ATLAS Collaboration, *Determination of the strong coupling constant α_s from transverse energy–energy correlations in multijet events at $\sqrt{s} = 8$ TeV using the ATLAS detector*, *Eur. Phys. J. C* **77** (2017) 872 [1707.02562]. Cited on p. 9.
- [18] ATLAS Collaboration, *Measurement of dijet azimuthal decorrelations in pp collisions at $\sqrt{s} = 8$ TeV with the ATLAS detector and determination of the strong coupling*, *Phys. Rev. D* **98** (2018) 092004 [1805.04691]. Cited on p. 9.
- [19] ATLAS Collaboration, *Measurement of transverse energy–energy correlations in multi-jet events in pp collisions at $\sqrt{s} = 7$ TeV using the ATLAS detector and determination of the strong coupling constant $\alpha_s(m_Z)$* , *Phys. Lett. B* **750** (2015) 427 [1508.01579]. Cited on p. 9.

- [20] The Super-Kamiokande Collaboration, *Evidence for oscillation of atmospheric neutrinos*, *Physical Review Letters* **81** (1998) 1562. Cited on p. 11.
- [21] 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*, *Physical Review Letters* **87** (2001) 071301. Cited on p. 11.
- [22] T2K Collaboration, K. Abe, N. Abgrall, H. Aihara, T. Akiri, J. B. Albert et al., *Evidence of Electron Neutrino Appearance in a Muon Neutrino Beam*, *Physical Review D* **88** (2013) 032002. Cited on p. 11.
- [23] S. Abachi, B. Abbott, M. Abolins, B. S. Acharya, I. Adam, D. L. Adams et al., *Observation of the Top Quark*, *Physical Review Letters* **74** (1995) 2632. Cited on p. 11.
- [24] F. Abe, H. Akimoto, A. Akopian, M. G. Albrow, S. R. Amendolia, D. Amidei et al., *Observation of Top Quark Production in $p\bar{p}$ Collisions with the Collider Detector at Fermilab*, *Physical Review Letters* **74** (1995) 2626. Cited on p. 11.
- [25] M. Kobayashi and T. Maskawa, *CP Violation in the Renormalizable Theory of Weak Interaction*, *Prog. Theor. Phys.* **49** (1973) 652. Cited on pp. 11 and 12.
- [26] Particle Data Group, *Review of Particle Physics*, *Progress of Theoretical and Experimental Physics* **2022** (2022) 083C01. Cited on p. 12.
- [27] B. Martin and G. Shaw, *Particle Physics*. Wiley, 4 ed., 2017. Cited on p. 12.
- [28] N. Cabibbo, *Unitary Symmetry and Leptonic Decays*, *Physical Review Letters* **10** (1963) 531. Cited on p. 12.
- [29] F. N. A. Laboratory, “Useful diagrams of top signals and backgrounds.” https://www-d0.fnal.gov/Run2Physics/top/top_public_web_pages/top_feynman_diagrams.html. Cited on p. 13.
- [30] ATLAS collaboration, *Standard Model Summary Plots February 2022*, tech. rep., CERN, Geneva, 2022. Cited on pp. 14 and 180.
- [31] P. Bärnreuther, M. Czakon and A. Mitov, *Percent-Level-Precision Physics at the Tevatron: Next-to-Next-to-Leading Order QCD Corrections to $q\bar{q} \rightarrow t\bar{t} + X$* , *Physical Review Letters* **109** (2012) 132001. Cited on p. 15.

- [32] M. Czakon and A. Mitov, *NNLO corrections to top-pair production at hadron colliders: the all-fermionic scattering channels*, *Journal of High Energy Physics* **2012** (2012) 54. Cited on p. 15.
- [33] M. Czakon and A. Mitov, *NNLO corrections to top pair production at hadron colliders: the quark-gluon reaction*, *Journal of High Energy Physics* **2013** (2013) 80. Cited on p. 15.
- [34] M. Czakon, P. Fiedler and A. Mitov, *Total Top-Quark Pair-Production Cross Section at Hadron Colliders Through $\mathcal{O}(\alpha_s^4)$* , *Physical Review Letters* **110** (2013) 252004. Cited on p. 15.
- [35] J. Kühn, A. Scharf and P. Uwer, *Weak interaction effects in top-quark pair production at hadron colliders*, *The European Physical Journal C* **51** (2007) 37. Cited on p. 15.
- [36] M. Czakon, D. Heymes, A. Mitov, D. Pagani, I. Tsinikos and M. Zaro, *Top-pair production at the LHC through NNLO QCD and NLO EW*, *Journal of High Energy Physics* **2017** (2017) 186. Cited on p. 15.
- [37] M. Beneke, P. Falgari, S. Klein and C. Schwinn, *Hadronic top-quark pair production with NNLL threshold resummation*, *Nuclear Physics B* **855** (2012) 695. Cited on p. 15.
- [38] M. Cacciari, M. Czakon, M. Mangano, A. Mitov and P. Nason, *Top-pair production at hadron colliders with next-to-next-to-leading logarithmic soft-gluon resummation*, *Physics Letters B* **710** (2012) 612. Cited on p. 15.
- [39] A. Ferroglia, B. D. Pecjak and L. L. Yang, *Soft-gluon resummation for boosted top-quark production at hadron colliders*, *Physical Review D* **86** (2012) 034010. Cited on p. 15.
- [40] T. Sjöstrand, S. Ask, J. R. Christiansen, R. Corke, N. Desai, P. Ilten et al., *An introduction to PYTHIA 8.2*, *Comput. Phys. Commun.* **191** (2015) 159 [1410.3012]. Cited on pp. 17 and 19.
- [41] C. Bierlich, S. Chakraborty, N. Desai, L. Gellersen, I. Helenius, P. Ilten et al., <http://arxiv.org/abs/2203.11601>, *A comprehensive guide to the physics and usage of PYTHIA 8.3*, Mar., 2022. Cited on p. 17.
- [42] T. Sjöstrand, <http://arxiv.org/abs/2208.05284>, *Soft QCD theory*, Aug., 2022. Cited on p. 17.

- [43] A. Buckley, J. Butterworth, S. Gieseke, D. Grellscheid, S. Hoche, H. Hoeth et al., *General-purpose event generators for LHC physics*, *Physics Reports* **504** (2011) 145. Cited on p. 16.
- [44] R. K. Ellis, W. J. Stirling and B. R. Webber, *QCD and collider physics*, Cambridge monographs on particle physics, nuclear physics, and cosmology. Cambridge University Press, Cambridge, 2003, [10.1017/CBO9780511628788](https://doi.org/10.1017/CBO9780511628788). Cited on p. 18.
- [45] V. N. Gribov and L. N. Lipatov, *Deep inelastic $e p$ scattering in perturbation theory*, *Sov. J. Nucl. Phys.* **15** (1972) 438. Cited on p. 18.
- [46] Y. L. Dokshitzer, *Calculation of the Structure Functions for Deep Inelastic Scattering and $e^+ e^-$ Annihilation by Perturbation Theory in Quantum Chromodynamics.*, *Sov. Phys. JETP* **46** (1977) 641. Cited on p. 18.
- [47] G. Altarelli and G. Parisi, *Asymptotic freedom in parton language*, *Nuclear Physics B* **126** (1977) 298. Cited on p. 18.
- [48] ZEUS Collaboration, “The Zeus Detector: Technical Proposal.” <https://cds.cern.ch/record/1478623>, 1986. Cited on p. 18.
- [49] H1 collaboration, *The H1 detector at HERA*, *Nucl. Instrum. Meth. A* **386** (1997) 310. Cited on p. 18.
- [50] “HERA - A Proposal for a Large Electron Proton Colliding Beam Facility at DESY.” <https://cds.cern.ch/record/1480597>, 1981. Cited on p. 18.
- [51] J. Bellm et al., *Herwig 7.1 Release Note*, [1705.06919](https://arxiv.org/abs/1705.06919). Cited on pp. 19 and 77.
- [52] B. Andersson, G. Gustafson, G. Ingelman and T. Sjöstrand, *Parton fragmentation and string dynamics*, *Phys. Rept.* **97** (1983) 31. Cited on p. 19.
- [53] M. Bähr et al., *Herwig++ physics and manual*, *Eur. Phys. J. C* **58** (2008) 639 [0803.0883]. Cited on pp. 19 and 77.
- [54] ATLAS Collaboration, *Measurement of charged-particle distributions sensitive to the underlying event in $\sqrt{s} = 13$ TeV proton–proton collisions with the ATLAS detector at the LHC*, *JHEP* **03** (2017) 157 [1701.05390]. Cited on p. 20.
- [55] GEANT4 Collaboration, S. Agostinelli et al., *GEANT4 – a simulation toolkit*, *Nucl. Instrum. Meth. A* **506** (2003) 250. Cited on pp. 20 and 76.

- [56] ATLAS Collaboration, “The simulation principle and performance of the ATLAS fast calorimeter simulation FastCaloSim.” ATL-PHYS-PUB-2010-013, 2010. Cited on p. 20.
- [57] L. Evans and P. Bryant, *LHC Machine*, *Journal of Instrumentation* **3** (2008) S08001. Cited on p. 21.
- [58] CERN, *CERN Yellow Reports: Monographs, Vol 4 (2017): High-Luminosity Large Hadron Collider (HL-LHC) Technical Design Report V. 0.1*, . Cited on p. 22.
- [59] I. B. Alonso, O. Brüning, P. Fessia, M. Lamont, L. Rossi, L. Taviani et al., *High-Luminosity Large Hadron Collider (HL-LHC)*, . Cited on pp. 22, 178, 183, and 184.
- [60] E. Mobs, <https://cds.cern.ch/record/2684277>, *The CERN accelerator complex*, 2019. Cited on p. 22.
- [61] The ATLAS Collaboration, G. Aad, E. Abat, J. Abdallah, A. A. Abdelalim, A. Abdesselam et al., *The ATLAS Experiment at the CERN Large Hadron Collider*, *Journal of Instrumentation* **3** (2008) S08003. Cited on p. 23.
- [62] J. Pequeno, <https://cds.cern.ch/record/1095928?ln=en>, *Computer generated image of the ATLAS Liquid Argon*, Mar., 2008. Cited on pp. 24 and 32.
- [63] ATLAS Collaboration, <https://twiki.cern.ch/twiki/bin/view/AtlasPublic/LuminosityPublicResultsRun2>, *Luminosity Summary Plots*, 2020. Cited on p. 25.
- [64] J. Pequeno and P. Schaffner, <https://cds.cern.ch/record/1505342>, *How ATLAS detects particles: diagram of particle paths in the detector*, Jan., 2013. Cited on p. 25.
- [65] ATLAS Collaboration, *ATLAS Inner Detector: Technical Design Report Volume 1*, tech. rep., CERN, Geneva, Apr., 1997. Cited on p. 26.
- [66] M. Thomson, *Modern Particle Physics*. Cambridge University Press, 2013. Cited on p. 26.
- [67] B. Mandelli, *The Pixel Detector of the ATLAS Experiment for the Run 2 at the Large Hadron Collider*, *Nuclear and Particle Physics Proceedings* **273-275** (2016) 1166. Cited on p. 27.

- [68] J. Pequenaó, <https://cds.cern.ch/record/1095929?ln=en>, *Computer generated image of the ATLAS Muons subsystem*, Mar., 2008. Cited on p. 27.
- [69] ATLAS Collaboration, *ATLAS Insertable B-Layer: Technical Design Report*, tech. rep., CERN, Geneva, Sept., 2010. Cited on p. 27.
- [70] A. La Rosa, *The ATLAS Insertable B-Layer: from construction to operation*, *arXiv:1610.01994 [physics]* (2016) . Cited on p. 27.
- [71] I. Perić, L. Blanquart, G. Comes, P. Denes, K. Einsweiler, P. Fischer et al., *The FEI3 readout chip for the ATLAS pixel detector*, *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* **565** (2006) 178. Cited on p. 27.
- [72] M. Garcia-Sciveres et al., *The FE-I4 pixel readout integrated circuit*, *Nucl. Instrum. Meth. A* **636** (2011) S155. Cited on p. 28.
- [73] ATLAS Collaboration, *Operation and performance of the ATLAS semiconductor tracker*, *JINST* **9** (2014) P08009 [1404.7473]. Cited on p. 28.
- [74] F. Campabadal, C. Fleta, M. Key, M. Lozano, C. Martinez, G. Pellegrini et al., *Design and performance of the ABCD3TA ASIC for readout of silicon strip detectors in the ATLAS semiconductor tracker*, *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* **552** (2005) 292. Cited on p. 29.
- [75] A. Vogel, *ATLAS Transition Radiation Tracker (TRT): Straw tube gaseous detectors at high rates*, *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* **732** (2013) 277. Cited on p. 30.
- [76] ATLAS Collaboration, *Performance of the ATLAS Transition Radiation Tracker in Run 1 of the LHC: tracker properties*, *JINST* **12** (2017) P05002 [1702.06473]. Cited on p. 30.
- [77] ATLAS Collaboration, *The ATLAS Experiment at the CERN Large Hadron Collider*, *JINST* **3** (2008) S08003. Cited on pp. 30 and 33.
- [78] N. Ilic, *Performance of the ATLAS Liquid Argon Calorimeter after three years of LHC operation and plans for a future upgrade*, *Journal of Instrumentation* **9** (2014) C03049. Cited on p. 33.

- [79] ATLAS Collaboration, *Commissioning of the ATLAS Muon Spectrometer with cosmic rays*, *Eur. Phys. J. C* **70** (2010) 875 [[1006.4384](#)]. Cited on p. 34.
- [80] J. Pequenaio, <https://cds.cern.ch/record/1095926>, *Computer generated image of the ATLAS inner detector*, Mar., 2008. Cited on p. 35.
- [81] ATLAS Collaboration, <https://cds.cern.ch/record/338080>, *ATLAS Magnet System: Magnet Project Technical Design Report, Volume 1*, 1997. Cited on p. 37.
- [82] ATLAS Collaboration, <https://cds.cern.ch/record/1602235>, *ATLAS TDAQ System Phase-I Upgrade: Technical Design Report*, 2013. Cited on p. 38.
- [83] ATLAS Collaboration, *Performance of the ATLAS trigger system in 2015*, *Eur. Phys. J. C* **77** (2017) 317 [[1611.09661](#)]. Cited on pp. 38 and 94.
- [84] ATLAS Collaboration, *Operation of the ATLAS trigger system in Run 2*, *JINST* **15** (2020) P10004 [[2007.12539](#)]. Cited on p. 38.
- [85] ATLAS Collaboration, <https://cds.cern.ch/record/381429>, *ATLAS Level-1 Trigger: Technical Design Report*, 1998. Cited on p. 38.
- [86] ATLAS Collaboration, <https://cds.cern.ch/record/616089>, *ATLAS High-Level Trigger, Data Acquisition and Controls: Technical Design Report*, 2003. Cited on p. 38.
- [87] R. Achenbach, P. Adragna, V. Andrei, P. Apostologlou, B. Åsman, C. Ay et al., *The ATLAS Level-1 Calorimeter Trigger*, *Journal of Instrumentation* **3** (2008) P03001. Cited on p. 38.
- [88] ATLAS Collaboration, *Performance of the ATLAS muon triggers in Run 2*, *JINST* **15** (2020) P09015 [[2004.13447](#)]. Cited on pp. 39 and 94.
- [89] M. Elsing, L. Goossens, A. Nairz and G. Negri, *The ATLAS Tier-0: Overview and operational experience*, *Journal of Physics: Conference Series* **219** (2010) 072011. Cited on p. 39.
- [90] F. Lasagni Manghi, LUCID collaboration, “LUCID: The ATLAS Luminosity Detector.” *ATL-FWD-PROC-2018-008*, 2019. [10.22323/1.340.0280](#). Cited on p. 39.

- [91] G. Avoni, M. Bruschi, G. Cabras, D. Caforio, N. Dehghanian, A. Floderus et al., *The new LUCID-2 detector for luminosity measurement and monitoring in ATLAS*, *Journal of Instrumentation* **13** (2018) P07017. Cited on p. 40.
- [92] M. Bruschi, *The new ATLAS/LUCID detector*, in 2015 4th International Conference on Advancements in Nuclear Instrumentation Measurement Methods and their Applications (ANIMMA), (Lisbon, Portugal), pp. 1–6, IEEE, Apr., 2015, DOI. Cited on p. 40.
- [93] R. Kogler et al., *Jet Substructure at the Large Hadron Collider: Experimental Review*, *Rev. Mod. Phys.* **91** (2019) 045003 [1803.06991]. Cited on pp. 43, 56, 62, and 124.
- [94] G. P. Salam, *Towards Jetography*, *The European Physical Journal C* **67** (2010) 637. Cited on pp. 43, 51, and 62.
- [95] J. E. Huth et al., *Toward a standardization of jet definitions*, in 1990 DPF Summer Study on High-energy Physics: Research Directions for the Decade (Snowmass 90), pp. 0134–136, 12, 1990. Cited on p. 45.
- [96] S. D. Ellis, Z. Kunszt and D. E. Soper, *One-jet inclusive cross section at order α_s^3 . Gluons only*, *Physical Review D* **40** (1989) 2188. Cited on p. 45.
- [97] M. Cacciari, G. P. Salam and G. Soyez, *The anti- k_t jet clustering algorithm*, *Journal of High Energy Physics* **2008** (2008) 063. Cited on pp. 46, 50, and 83.
- [98] G. Sterman and S. Weinberg, *Jets from Quantum Chromodynamics*, *Physical Review Letters* **39** (1977) 1436. Cited on p. 45.
- [99] G. Arnison, A. Astbury, B. Aubert, C. Bacci, G. Bauer, A. Bézaguet et al., *Hadronic jet production at the CERN proton-antiproton collider*, *Physics Letters B* **132** (1983) 214. Cited on p. 47.
- [100] G. C. Blazey, J. R. Dittmann, S. D. Ellis, V. D. Elvira, K. Frame, S. Grinstein et al., <http://arxiv.org/abs/hep-ex/0005012>, *Run II Jet Physics: Proceedings of the Run II QCD and Weak Boson Physics Workshop*, May, 2000. 10.48550/arXiv.hep-ex/0005012. Cited on p. 47.
- [101] M. H. Seymour, *Jet Shapes in Hadron Collisions: Higher Orders, Resummation and Hadronization*, *Nuclear Physics B* **513** (1998) 269. Cited on p. 48.
- [102] G. P. Salam and G. Soyez, *A practical Seedless Infrared-Safe Cone jet algorithm*, *Journal of High Energy Physics* **2007** (2007) 086. Cited on p. 48.

- [103] ATLAS Collaboration, *Measurement of the Lund Jet Plane Using Charged Particles in 13 TeV Proton–Proton Collisions with the ATLAS Detector*, *Phys. Rev. Lett.* **124** (2020) 222002 [2004.03540]. Cited on pp. 48, 74, and 222.
- [104] ALICE Collaboration, S. Acharya, D. Adamova, A. Adler, J. Adolfsson, G. Aglieri Rinella et al., *Direct observation of the dead-cone effect in quantum chromodynamics*, *Nature* **605** (2022) 440. Cited on pp. 48, 63, and 222.
- [105] S. Catani, Y. L. Dokshitzer, M. Olsson, G. Turnock and B. R. Webber, *New clustering algorithm for multijet cross sections in $e+e$ annihilation*, *Physics Letters B* **269** (1991) 432. Cited on p. 48.
- [106] W. Bartel, L. Becker, R. Felst, D. Haidt, G. Knies, H. Krehbiel et al., *Experimental studies on multijet production in $e+e$ annihilation at PETRA energies*, *Zeitschrift für Physik C Particles and Fields* **33** (1986) 23. Cited on p. 48.
- [107] S. Bethke, J. Allison, K. Ambrus, R. J. Barlow, W. Bartel, C. K. Bowdery et al., *Experimental investigation of the energy dependence of the strong coupling strength*, *Physics Letters B* **213** (1988) 235. Cited on p. 48.
- [108] S. Catani, Y. Dokshitzer, M. Seymour and B. Webber, *Longitudinally-invariant k -clustering algorithms for hadron-hadron collisions*, *Nuclear Physics B* **406** (1993) 187. Cited on p. 48.
- [109] S. D. Ellis and D. E. Soper, *Successive Combination Jet Algorithm For Hadron Collisions*, *Physical Review D* **48** (1993) 3160. Cited on pp. 48 and 67.
- [110] M. Wobisch and T. Wengler, *Hadronization Corrections to Jet Cross Sections in Deep-Inelastic Scattering*, *arXiv:hep-ph/9907280* (1999) . Cited on p. 50.
- [111] Y. L. Dokshitzer, G. D. Leder, S. Moretti and B. R. Webber, *Better Jet Clustering Algorithms*, *Journal of High Energy Physics* **1997** (1997) 001. Cited on p. 50.
- [112] ATLAS Collaboration, *Jet reconstruction and performance using particle flow with the ATLAS Detector*, *Eur. Phys. J. C* **77** (2017) 466 [1703.10485]. Cited on pp. 52, 53, 57, and 83.
- [113] ATLAS Collaboration, “Jet reclustering and close-by effects in ATLAS Run 2.” ATLAS-CONF-2017-062, 2017. Cited on pp. 56, 58, and 84.

- [114] B. Nachman, P. Nef, A. Schwartzman, M. Swiatlowski and C. Wanotayaroj, *Jets from Jets: Re-clustering as a tool for large radius jet reconstruction and grooming at the LHC*, *Journal of High Energy Physics* **2015** (2015) 75. Cited on pp. 56 and 58.
- [115] D. Krohn, J. Thaler and L.-T. Wang, *Jet Trimming*, *Journal of High Energy Physics* **2010** (2010) . Cited on p. 56.
- [116] ATLAS Collaboration, *Jet energy scale and resolution measured in proton–proton collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector*, *Eur. Phys. J. C* **81** (2020) 689 [2007.02645]. Cited on pp. 59 and 62.
- [117] R. Kogler, *Advances in Jet Substructure at the LHC: Algorithms, Measurements and Searches for New Physical Phenomena*, vol. 284 of *Springer Tracts in Modern Physics*. Springer International Publishing, Cham, 2021, 10.1007/978-3-030-72858-8. Cited on p. 62.
- [118] J. M. Butterworth, A. R. Davison, M. Rubin and G. P. Salam, *Jet substructure as a new Higgs search channel at the LHC*, *Physical Review Letters* **100** (2008) 242001. Cited on p. 62.
- [119] A. J. Larkoski, I. Moult and B. Nachman, *Jet substructure at the large hadron collider: A review of recent advances in theory and machine learning*, *Physics Reports* **841** (2020) 1. Cited on p. 62.
- [120] B. Nachman, S. Rappoccio, N. Tran, J. Bonilla, G. Chachamis, B. M. Dillon et al., *Jets and Jet Substructure at Future Colliders*, *arXiv:2203.07462 [hep-ex, physics:hep-ph]* (2022) . Cited on p. 62.
- [121] S. Marzani, G. Soyez and M. Spannowsky, *Looking inside jets: an introduction to jet substructure and boosted-object phenomenology*, *arXiv:1901.10342 [hep-ex, physics:hep-ph]* (2019) . Cited on p. 62.
- [122] CMS Collaboration, *Measurement of jet substructure observables in $t\bar{t}$ events from proton–proton collisions at $\sqrt{s} = 13$ TeV*, *Phys. Rev. D* **98** (2018) 092014 [1808.07340]. Cited on pp. 63, 74, 124, and 222.
- [123] D. Kar and P. Sarmah, *Effect of new jet substructure measurements on Pythia8 tunes*, *arXiv:2101.11395 [hep-ex, physics:hep-ph]* (2021) . Cited on pp. 63 and 222.
- [124] A. J. Larkoski and J. Thaler, *Unsafe but Calculable: Ratios of Angularities in Perturbative QCD*, *Journal of High Energy Physics* **2013** (2013) 137. Cited on p. 63.

- [125] A. J. Larkoski, S. Marzani and J. Thaler, *Sudakov Safety in Perturbative QCD*, *Physical Review D* **91** (2015) 111501. Cited on p. 63.
- [126] A. J. Larkoski, J. Thaler and W. J. Waalewijn, *Gaining (mutual) information about quark/gluon discrimination*, *Journal of High Energy Physics* **2014** (2014) 129. Cited on pp. 64 and 124.
- [127] A. J. Larkoski, G. P. Salam and J. Thaler, *Energy correlation functions for jet substructure*, *Journal of High Energy Physics* **2013** (2013) 108. Cited on pp. 65, 66, 124, and 134.
- [128] A. J. Larkoski, I. Moult and D. Neill, *Analytic boosted boson discrimination*, *Journal of High Energy Physics* **2016** (2016) 117. Cited on p. 65.
- [129] J. Thaler and K. Van Tilburg, *Identifying Boosted Objects with N-subjettiness*, *Journal of High Energy Physics* **2011** (2011) 15. Cited on p. 67.
- [130] A. Collaboration, *Identification of high transverse momentum top quarks in pp collisions at $\sqrt{s} = 8$ TeV with the ATLAS detector*, *Journal of High Energy Physics* **2016** (2016) 93. Cited on p. 69.
- [131] T. Faucett, S.-C. Hsu and D. Whiteson, *Learning to identify semi-visible jets*, *JHEP* **12** (2022) 132 [2208.10062]. Cited on p. 70.
- [132] A. J. Larkoski, I. Moult and D. Neill, *Building a Better Boosted Top Tagger*, *Physical Review D* **91** (2015) 034035. Cited on p. 70.
- [133] N. Bartosik, “b-tagging diagram.” https://upload.wikimedia.org/wikipedia/commons/b/ba/B-tagging_diagram.png. Cited on p. 71.
- [134] ATLAS Collaboration, “Identification of Jets Containing b -Hadrons with Recurrent Neural Networks at the ATLAS Experiment.” ATL-PHYS-PUB-2017-003, 2017. Cited on p. 71.
- [135] ATLAS Collaboration, “Secondary vertex finding for jet flavour identification with the ATLAS detector.” ATL-PHYS-PUB-2017-011, 2017. Cited on p. 71.
- [136] ATLAS Collaboration, “Topological b -hadron decay reconstruction and identification of b -jets with the JetFitter package in the ATLAS experiment at the LHC.” ATL-PHYS-PUB-2018-025, 2018. Cited on p. 71.

- [137] ATLAS Collaboration, *ATLAS b -jet identification performance and efficiency measurement with $t\bar{t}$ events in pp collisions at $\sqrt{s} = 13$ TeV*, *Eur. Phys. J. C* **79** (2019) 970 [1907.05120]. Cited on pp. 71 and 83.
- [138] ATLAS Collaboration, *ATLAS flavour-tagging algorithms for the LHC Run 2 pp collision dataset*, 2211.16345. Cited on pp. 71 and 72.
- [139] ATLAS Collaboration, *Performance of top-quark and W-boson tagging with ATLAS in Run 2 of the LHC*, *Eur. Phys. J. C* **79** (2019) 375 [1808.07858]. Cited on pp. 73 and 124.
- [140] ATLAS Collaboration, “Constituent-Based Top-Quark Tagging with the ATLAS Detector.” ATL-PHYS-PUB-2022-039, 2022. Cited on pp. 73 and 222.
- [141] Y. Lu, A. Romero, M. J. Fenton, D. Whiteson and P. Baldi, *Resolving Extreme Jet Substructure*, *arXiv:2202.00723 [hep-ex, physics:hep-ph]* (2022) . Cited on pp. 73 and 222.
- [142] ATLAS Collaboration, *Differential $t\bar{t}$ cross-section measurements using boosted top quarks in the all-hadronic final state with 139 fb^{-1} of ATLAS data*, *JHEP* **04** (2023) 080 [2205.02817]. Cited on pp. 74 and 113.
- [143] ATLAS Collaboration, *Measurements of top-quark pair differential and double-differential cross-sections in the ℓ +jets channel with pp collisions at $\sqrt{s} = 13$ TeV using the ATLAS detector*, *Eur. Phys. J. C* **79** (2019) 1028 [1908.07305]. Cited on pp. 74, 93, and 113.
- [144] ATLAS Collaboration, *Measurements of top-quark pair single- and double-differential cross-sections in the all-hadronic channel in pp collisions at $\sqrt{s} = 13$ TeV using the ATLAS detector*, *JHEP* **01** (2021) 033 [2006.09274]. Cited on p. 74.
- [145] ATLAS Collaboration, *Measurement of the $t\bar{t}$ production cross-section in the lepton+jets channel at $\sqrt{s} = 13$ TeV with the ATLAS experiment*, *Phys. Lett. B* **810** (2020) 135797 [2006.13076]. Cited on p. 74.
- [146] CMS Collaboration, *Measurement of the $t\bar{t}$ production cross section, the top quark mass, and the strong coupling constant using dilepton events in pp collisions at $\sqrt{s} = 13$ TeV*, *Eur. Phys. J. C* **79** (2019) 368 [1812.10505]. Cited on p. 74.

- [147] CMS Collaboration, *Measurement of differential cross sections for the production of top quark pairs and of additional jets in lepton+jets events from pp collisions at $\sqrt{s} = 13$ TeV*, *Phys. Rev. D* **97** (2018) 112003 [[1803.08856](#)]. Cited on p. 74.
- [148] CMS Collaboration, *Measurements of $t\bar{t}$ differential cross sections in proton–proton collisions at $\sqrt{s} = 13$ TeV using events containing two leptons*, *JHEP* **02** (2019) 149 [[1811.06625](#)]. Cited on p. 74.
- [149] CMS Collaboration, *Measurement of differential $t\bar{t}$ production cross sections using top quarks at large transverse momenta in pp collisions at $\sqrt{s} = 13$ TeV*, *Phys. Rev. D* **103** (2021) 052008 [[2008.07860](#)]. Cited on pp. 74 and 113.
- [150] CMS Collaboration, *Measurement of differential $t\bar{t}$ production cross sections in the full kinematic range using lepton+jets events from proton–proton collisions at $\sqrt{s} = 13$ TeV*, *Phys. Rev. D* **104** (2021) 092013 [[2108.02803](#)]. Cited on p. 74.
- [151] CMS Collaboration, *Measurement of the inclusive and differential $t\bar{t}\gamma$ cross sections in the dilepton channel and effective field theory interpretation in proton–proton collisions at $\sqrt{s} = 13$ TeV*, *JHEP* **05** (2022) 091 [[2201.07301](#)]. Cited on p. 74.
- [152] ATLAS Collaboration, *Jet mass and substructure of inclusive jets in $\sqrt{s} = 7$ TeV pp collisions with the ATLAS experiment*, *JHEP* **05** (2012) 128 [[1203.4606](#)]. Cited on p. 74.
- [153] ATLAS Collaboration, *Measurement of the cross-section of high transverse momentum vector bosons reconstructed as single jets and studies of jet substructure in pp collisions at $\sqrt{s} = 7$ TeV with the ATLAS detector*, *New J. Phys.* **16** (2014) 113013 [[1407.0800](#)]. Cited on p. 74.
- [154] ATLAS Collaboration, *Measurement of jet-substructure observables in top quark, W boson and light jet production in proton–proton collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector*, *JHEP* **08** (2019) 033 [[1903.02942](#)]. Cited on pp. 74 and 124.
- [155] CMS Collaboration, *Study of quark and gluon jet substructure in Z+jet and dijet events from pp collisions*, *JHEP* **01** (2021) 188 [[2109.03340](#)]. Cited on p. 74.
- [156] ATLAS Collaboration, *ATLAS data quality operations and performance for 2015–2018 data-taking*, *JINST* **15** (2020) P04003 [[1911.04632](#)]. Cited on p. 76.

- [157] ATLAS Collaboration, “Luminosity determination in pp collisions at $\sqrt{s} = 13$ TeV using the ATLAS detector at the LHC.” ATLAS-CONF-2019-021, 2019. Cited on p. 76.
- [158] T. Sjöstrand, S. Mrenna and P. Skands, *A brief introduction to PYTHIA 8.1*, *Comput. Phys. Commun.* **178** (2008) 852 [[0710.3820](#)]. Cited on p. 76.
- [159] ATLAS Collaboration, “The Pythia 8 A3 tune description of ATLAS minimum bias and inelastic measurements incorporating the Donnachie–Landshoff diffractive model.” ATL-PHYS-PUB-2016-017, 2016. Cited on p. 76.
- [160] R. D. Ball et al., *Parton distributions with LHC data*, *Nucl. Phys. B* **867** (2013) 244 [[1207.1303](#)]. Cited on p. 76.
- [161] S. Frixione, G. Ridolfi and P. Nason, *A positive-weight next-to-leading-order Monte Carlo for heavy flavour hadroproduction*, *JHEP* **09** (2007) 126 [[0707.3088](#)]. Cited on p. 76.
- [162] P. Nason, *A new method for combining NLO QCD with shower Monte Carlo algorithms*, *JHEP* **11** (2004) 040 [[hep-ph/0409146](#)]. Cited on p. 76.
- [163] S. Frixione, P. Nason and C. Oleari, *Matching NLO QCD computations with parton shower simulations: the POWHEG method*, *JHEP* **11** (2007) 070 [[0709.2092](#)]. Cited on p. 76.
- [164] S. Alioli, P. Nason, C. Oleari and E. Re, *A general framework for implementing NLO calculations in shower Monte Carlo programs: the POWHEG BOX*, *JHEP* **06** (2010) 043 [[1002.2581](#)]. Cited on p. 76.
- [165] NNPDF collaboration, *Parton distributions for the LHC run II*, *JHEP* **04** (2015) 040 [[1410.8849](#)]. Cited on p. 76.
- [166] ATLAS Collaboration, “ATLAS Pythia 8 tunes to 7 TeV data.” ATL-PHYS-PUB-2014-021, 2014. Cited on p. 76.
- [167] D. J. Lange, *The EvtGen particle decay simulation package*, *Nucl. Instrum. Meth. A* **462** (2001) 152. Cited on p. 76.
- [168] M. Beneke, P. Falgari, S. Klein and C. Schwinn, *Hadronic top-quark pair production with NNLL threshold resummation*, *Nucl. Phys. B* **855** (2012) 695 [[1109.1536](#)]. Cited on p. 76.

- [169] M. Cacciari, M. Czakon, M. Mangano, A. Mitov and P. Nason, *Top-pair production at hadron colliders with next-to-next-to-leading logarithmic soft-gluon resummation*, *Phys. Lett. B* **710** (2012) 612 [[1111.5869](#)]. Cited on p. 76.
- [170] P. Bärnreuther, M. Czakon and A. Mitov, *Percent-Level-Precision Physics at the Tevatron: Next-to-Next-to-Leading Order QCD Corrections to $q\bar{q} \rightarrow t\bar{t} + X$* , *Phys. Rev. Lett.* **109** (2012) 132001 [[1204.5201](#)]. Cited on p. 76.
- [171] M. Czakon and A. Mitov, *NNLO corrections to top-pair production at hadron colliders: the all-fermionic scattering channels*, *JHEP* **12** (2012) 054 [[1207.0236](#)]. Cited on p. 76.
- [172] M. Czakon and A. Mitov, *NNLO corrections to top pair production at hadron colliders: the quark-gluon reaction*, *JHEP* **01** (2013) 080 [[1210.6832](#)]. Cited on p. 76.
- [173] M. Czakon, P. Fiedler and A. Mitov, *Total Top-Quark Pair-Production Cross Section at Hadron Colliders Through $O(\alpha_s^4)$* , *Phys. Rev. Lett.* **110** (2013) 252004 [[1303.6254](#)]. Cited on p. 76.
- [174] M. Czakon and A. Mitov, *Top++: A program for the calculation of the top-pair cross-section at hadron colliders*, *Comput. Phys. Commun.* **185** (2014) 2930 [[1112.5675](#)]. Cited on p. 76.
- [175] ATLAS Collaboration, “Studies on top-quark Monte Carlo modelling with Sherpa and MG5_aMC@NLO.” ATL-PHYS-PUB-2017-007, 2017. Cited on p. 77.
- [176] J. Bellm et al., *Herwig 7.0/Herwig++ 3.0 release note*, *Eur. Phys. J. C* **76** (2016) 196 [[1512.01178](#)]. Cited on p. 77.
- [177] L. A. Harland-Lang, A. D. Martin, P. Motylinski and R. S. Thorne, *Parton distributions in the LHC era: MMHT 2014 PDFs*, *Eur. Phys. J. C* **75** (2015) 204 [[1412.3989](#)]. Cited on p. 77.
- [178] J. Alwall, R. Frederix, S. Frixione, V. Hirschi, F. Maltoni, O. Mattelaer et al., *The automated computation of tree-level and next-to-leading order differential cross sections, and their matching to parton shower simulations*, *JHEP* **07** (2014) 079 [[1405.0301](#)]. Cited on pp. 77 and 78.
- [179] S. Frixione, E. Laenen, P. Motylinski and B. R. Webber, *Angular correlations of lepton pairs from vector boson and top quark decays in Monte Carlo simulations*, *JHEP* **04** (2007) 081 [[hep-ph/0702198](#)]. Cited on p. 77.

- [180] P. Artoisenet, R. Frederix, O. Mattelaer and R. Rietkerk, *Automatic spin-entangled decays of heavy resonances in Monte Carlo simulations*, *JHEP* **03** (2013) 015 [[1212.3460](#)]. Cited on p. 77.
- [181] E. Bothmann et al., *Event generation with Sherpa 2.2*, *SciPost Phys.* **7** (2019) 034 [[1905.09127](#)]. Cited on p. 78.
- [182] T. Gleisberg and S. Höche, *Comix, a new matrix element generator*, *JHEP* **12** (2008) 039 [[0808.3674](#)]. Cited on p. 78.
- [183] F. Buccioni, J.-N. Lang, J. M. Lindert, P. Maierhöfer, S. Pozzorini, H. Zhang et al., *OpenLoops 2*, *Eur. Phys. J. C* **79** (2019) 866 [[1907.13071](#)]. Cited on p. 78.
- [184] F. Cascioli, P. Maierhöfer and S. Pozzorini, *Scattering Amplitudes with Open Loops*, *Phys. Rev. Lett.* **108** (2012) 111601 [[1111.5206](#)]. Cited on p. 78.
- [185] A. Denner, S. Dittmaier and L. Hofer, *COLLIER: A fortran-based complex one-loop library in extended regularizations*, *Comput. Phys. Commun.* **212** (2017) 220 [[1604.06792](#)]. Cited on p. 78.
- [186] S. Schumann and F. Krauss, *A parton shower algorithm based on Catani–Seymour dipole factorisation*, *JHEP* **03** (2008) 038 [[0709.1027](#)]. Cited on p. 78.
- [187] S. Höche, F. Krauss, M. Schönherr and F. Siegert, *A critical appraisal of NLO+PS matching methods*, *JHEP* **09** (2012) 049 [[1111.1220](#)]. Cited on p. 78.
- [188] S. Höche, F. Krauss, M. Schönherr and F. Siegert, *QCD matrix elements + parton showers. The NLO case*, *JHEP* **04** (2013) 027 [[1207.5030](#)]. Cited on p. 78.
- [189] S. Catani, F. Krauss, B. R. Webber and R. Kuhn, *QCD Matrix Elements + Parton Showers*, *JHEP* **11** (2001) 063 [[hep-ph/0109231](#)]. Cited on p. 78.
- [190] S. Höche, F. Krauss, S. Schumann and F. Siegert, *QCD matrix elements and truncated showers*, *JHEP* **05** (2009) 053 [[0903.1219](#)]. Cited on p. 78.
- [191] ATLAS Collaboration, *Topological cell clustering in the ATLAS calorimeters and its performance in LHC Run 1*, *Eur. Phys. J. C* **77** (2017) 490 [[1603.02934](#)]. Cited on p. 79.
- [192] ATLAS Collaboration, *Performance of the ATLAS track reconstruction algorithms in dense environments in LHC Run 2*, *Eur. Phys. J. C* **77** (2017) 673 [[1704.07983](#)]. Cited on p. 79.

- [193] M. Cacciari and G. P. Salam, *Pileup subtraction using jet areas*, *Physics Letters B* **659** (2008) 119. Cited on p. 80.
- [194] J. Illingworth and J. Kittler, *A Survey of the Hough Transform*, *Computer Vision, Graphics, and Image Processing* **44** (1988) . Cited on p. 81.
- [195] ATLAS Collaboration, *Muon reconstruction and identification efficiency in ATLAS using the full Run 2 pp collision data set at $\sqrt{s} = 13$ TeV*, *Eur. Phys. J. C* **81** (2021) 578 [2012.00578]. Cited on p. 81.
- [196] ATLAS Collaboration, *Electron efficiency measurements with the ATLAS detector using 2012 LHC proton–proton collision data*, *Eur. Phys. J. C* **77** (2017) 195 [1612.01456]. Cited on p. 81.
- [197] ATLAS Collaboration, *Electron and photon performance measurements with the ATLAS detector using the 2015–2017 LHC proton–proton collision data*, *JINST* **14** (2019) P12006 [1908.00005]. Cited on p. 81.
- [198] ATLAS Collaboration, *Muon reconstruction performance of the ATLAS detector in proton–proton collision data at $\sqrt{s} = 13$ TeV*, *Eur. Phys. J. C* **76** (2016) 292 [1603.05598]. Cited on p. 82.
- [199] ATLAS Collaboration, *Jet energy scale measurements and their systematic uncertainties in proton–proton collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector*, *Phys. Rev. D* **96** (2017) 072002 [1703.09665]. Cited on p. 83.
- [200] M. Cacciari, G. P. Salam and G. Soyez, *The Catchment Area of Jets*, *Journal of High Energy Physics* **2008** (2008) 005. Cited on p. 85.
- [201] ATLAS Collaboration, *Performance of electron and photon triggers in ATLAS during LHC Run 2*, *Eur. Phys. J. C* **80** (2020) 47 [1909.00761]. Cited on p. 94.
- [202] ATLAS Collaboration, “Estimation of non-prompt and fake lepton backgrounds in final states with top quarks produced in proton–proton collisions at $\sqrt{s} = 8$ TeV with the ATLAS Detector.” ATLAS-CONF-2014-058, 2014. Cited on p. 107.
- [203] ATLAS Collaboration, *Measurements of $t\bar{t}$ differential cross-sections of highly boosted top quarks decaying to all-hadronic final states in pp collisions at $\sqrt{s} = 13$ TeV using the ATLAS detector*, *Phys. Rev. D* **98** (2018) 012003 [1801.02052]. Cited on p. 113.

- [204] ATLAS Collaboration, *Measurements of top-quark pair differential cross-sections in the lepton+jets channel in pp collisions at $\sqrt{s} = 13$ TeV using the ATLAS detector*, *JHEP* **11** (2017) 191 [[1708.00727](#)]. Cited on p. 113.
- [205] ATLAS Collaboration, *Measurement of the differential cross-section of highly boosted top quarks as a function of their transverse momentum in $\sqrt{s} = 8$ TeV proton–proton collisions using the ATLAS detector*, *Phys. Rev. D* **93** (2016) 032009 [[1510.03818](#)]. Cited on p. 113.
- [206] J. M. Butterworth, B. E. Cox and J. R. Forshaw, *WW scattering at the LHC*, *Physical Review D* **65** (2002) . Cited on p. 124.
- [207] J. Thaler and L.-T. Wang, *Strategies to Identify Boosted Tops*, *Journal of High Energy Physics* **2008** (2008) 092. Cited on p. 124.
- [208] ATLAS Collaboration, *Measurement of k_T splitting scales in $W \rightarrow \ell \nu$ events at $\sqrt{s} = 7$ TeV with the ATLAS detector*, *Eur. Phys. J. C* **73** (2013) 2432 [[1302.1415](#)]. Cited on p. 124.
- [209] S. Chekanov and J. Proudfoot, *Searches for TeV-scale particles at the LHC using jet shapes*, *Physical Review D* **81** (2010) 114038. Cited on p. 124.
- [210] ATLAS Collaboration, *Identification of high transverse momentum top quarks in pp collisions at $\sqrt{s} = 8$ TeV with the ATLAS detector*, *JHEP* **06** (2016) 093 [[1603.03127](#)]. Cited on p. 124.
- [211] O. Behnke, K. Kröninger, G. Schott and T. Schörner-Sadenius, *Data Analysis in High Energy Physics*. Wiley-VCH, 2013. Cited on p. 136.
- [212] G. D’Agostini, *A multidimensional unfolding method based on Bayes’ theorem*, *Nucl. Instrum. Meth. A* **362** (1995) 487 . Cited on p. 136.
- [213] T. Adye, *Unfolding algorithms and tests using RooUnfold*, in *Proceedings, 2011 Workshop on Statistical Issues Related to Discovery Claims in Search Experiments and Unfolding (PHYSTAT 2011)*, pp. 313–318, [1105.1160](#), DOI. Cited on p. 136.
- [214] S. Schmitt, *Data Unfolding Methods in High Energy Physics*, *EPJ Web Conf.* **137** (2017) 11008 [[1611.01927](#)]. Cited on p. 143.
- [215] C. Hadfield, *An Astronaut’s Guide to Life on Earth*. Random House of Canada, 2013. Cited on p. 154.

- [216] ATLAS Collaboration, “Early Inner Detector Tracking Performance in the 2015 Data at $\sqrt{s} = 13$ TeV.” ATL-PHYS-PUB-2015-051, 2015. Cited on p. 156.
- [217] ATLAS Collaboration, “ATLAS simulation of boson plus jets processes in Run 2.” ATL-PHYS-PUB-2017-006, 2017. Cited on p. 157.
- [218] <http://hilumilhc.web.cern.ch/>, “High Luminosity LHC project.” Cited on pp. 179 and 261.
- [219] ATLAS Collaboration, “Projections for measurements of Higgs boson cross sections, branching ratios, coupling parameters and mass with the ATLAS detector at the HL-LHC.” ATL-PHYS-PUB-2018-054, 2018. Cited on p. 179.
- [220] ATLAS Collaboration, “Measurement prospects of the pair production and self-coupling of the Higgs boson with the ATLAS experiment at the HL-LHC.” ATL-PHYS-PUB-2018-053, 2018. Cited on p. 179.
- [221] M. Cepeda, S. Gori, P. Ilten, M. Kado, F. Riva, R. A. Khalek et al., <http://arxiv.org/abs/1902.00134>, *Higgs Physics at the HL-LHC and HE-LHC*, Mar., 2019. Cited on p. 179.
- [222] J. Nielsen, *Physics prospects for ATLAS at the HL-LHC*, *Journal of Physics: Conference Series* **1690** (2020) 012156. Cited on p. 179.
- [223] ATLAS Collaboration, *Search for flavour-changing neutral-current interactions of a top quark and a gluon in pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector*, *Eur. Phys. J. C* **82** (2021) 334 [2112.01302]. Cited on p. 180.
- [224] ATLAS Collaboration, *Search for flavour-changing neutral current interactions of the top quark and the Higgs boson in events with a pair of τ -leptons in pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector*, 2208.11415. Cited on p. 180.
- [225] ATLAS Collaboration, “Search for flavor-changing neutral-current couplings between the top quark and the Z boson with LHC Run 2 proton–proton collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector.” ATLAS-CONF-2021-049, 2021. Cited on p. 180.
- [226] ATLAS collaboration, *Search for flavour-changing neutral-current couplings between the top quark and the photon with the ATLAS detector at $\sqrt{s} = 13$ TeV*, *Phys. Lett. B* **842** (2023) 137379 [2205.02537]. Cited on p. 180.

- [227] ATLAS Collaboration, “Expected performance of the ATLAS detector at the High-Luminosity LHC.” ATL-PHYS-PUB-2019-005, 2019. Cited on p. 181.
- [228] ATLAS Collaboration, “Prospects for searches for staus, charginos and neutralinos at the high luminosity LHC with the ATLAS Detector.” ATL-PHYS-PUB-2018-048, 2018. Cited on p. 181.
- [229] ATLAS Collaboration, “Measurement of the charge asymmetry in top quark pair production in association with a photon with the ATLAS experiment.” ATLAS-CONF-2022-049, 2022. Cited on p. 181.
- [230] ATLAS Collaboration, *Measurement of single top-quark production in the s-channel in proton–proton collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector*, [arXiv:2209.08990](https://arxiv.org/abs/2209.08990). Cited on p. 181.
- [231] ATLAS Collaboration, “Prospect for a measurement of the Weak Mixing Angle in $pp \rightarrow Z/\gamma^* \rightarrow e^+e^-$ events with the ATLAS detector at the High Luminosity Large Hadron Collider.” ATL-PHYS-PUB-2018-037, 2018. Cited on p. 181.
- [232] ATLAS Collaboration, “Prospects for the measurement of the W-boson mass at the HL- and HE-LHC.” ATL-PHYS-PUB-2018-026, 2018. Cited on p. 181.
- [233] CDF Collaboration, *High-precision measurement of the W boson mass with the CDF II detector*, *Science* **376** (2022) 170. Cited on p. 181.
- [234] O. Brüning, H. Gray, K. Klein, M. Lamont, M. Narain, R. Polifka et al., *The scientific potential and technological challenges of the High-Luminosity Large Hadron Collider program*, *Reports on Progress in Physics* **85** (2022) 046201. Cited on p. 184.
- [235] “Event displays from upgrade physics simulated data.” <https://twiki.cern.ch/twiki/bin/view/AtlasPublic/UpgradeEventDisplays>. Cited on p. 185.
- [236] ATLAS Collaboration, “Expected tracking and related performance with the updated ATLAS Inner Tracker layout at the High-Luminosity LHC.” ATL-PHYS-PUB-2021-024, 2021. Cited on pp. 186, 187, and 194.
- [237] J. S. Keller, *The ATLAS ITk strip detector system for the High Luminosity LHC upgrade*, *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* **958** (2020) 162053. Cited on p. 185.

- [238] L. Rossi, P. Fischer, T. Rohe and N. Wermes, *Pixel Detectors - From Fundamentals to Applications*. Springer, 1st ed., 2006. Cited on pp. 190 and 261.
- [239] S. Möbius, *Module development for the ATLAS ITk pixel detector*, *Journal of Instrumentation* **17** (2022) C03042. Cited on pp. 191, 192, and 193.
- [240] ATLAS Collaboration, <https://cds.cern.ch/record/2285585>, *ATLAS Inner Tracker Pixel Detector: Technical Design Report*, 2017. Cited on pp. 193, 200, and 262.
- [241] RD53 Collaboration, https://cds.cern.ch/record/2287593/files/RD53A_Manual_V3-5.pdf, *RD53A Manual*, Apr., 2019. Cited on p. 192.
- [242] R. Diener, J. Dreyling-Eschweiler, H. Ehrlichmann, I.-M. Gregor, U. Kötz, U. Krämer et al., *The DESY II Test Beam Facility*, *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* **922** (2019) 265. Cited on pp. 195, 196, and 262.
- [243] J. Wenninger, “Sps operation.” <https://jwenning.web.cern.ch/SPS.html>. Cited on pp. 196 and 262.
- [244] S. Spannagel, K. Wolters, D. Hynds, N. Alipour Tehrani, M. Benoit, D. Dannheim et al., *Allpix2: A modular simulation framework for silicon detectors*, *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* **901** (2018) 164. Cited on p. 197.
- [245] H. Jansen, S. Spannagel, J. Behr, A. Bulgheroni, G. Claus, E. Corrin et al., *Performance of the EUDET-type beam telescopes*, *EPJ Techniques and Instrumentation* **3** (2016) 7. Cited on p. 198.
- [246] A. Himmi et al., “Mimosa26 User Manual.” http://www.iphc.cnrs.fr/IMG/pdf/M26_UserManual_light.pdf, December, 2008. Cited on p. 198.
- [247] Karlsruhe Institute of Technology, “Kit proton irradiation.” <https://www.etp.kit.edu/english/264.php>. Cited on p. 199.
- [248] RD53 Collaboration, “The RD53A Integrated Circuit.” CERN-RD53-PUB-17-001, 2017. Cited on p. 201.

- [249] H. Perrey, *EUDAQ and EUTelescope: Software Frameworks for Test Beam Data Acquisition and Analysis*, in *Proceedings of Technology and Instrumentation in Particle Physics 2014 — PoS(TIPP2014)*, (Amsterdam, the Netherlands), p. 353, Sissa Medialab, July, 2015, DOI. Cited on p. 203.
- [250] <http://eutelescope.web.cern.ch/>, *EUTelescope*, 2019. Cited on p. 203.
- [251] C. Kleinwort, *General Broken Lines as advanced track fitting method*, *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* **673** (2012) 107. Cited on p. 205.
- [252] Center for High Throughput Computing, <https://htcondor.readthedocs.io/en/stable/>, *HTCondor Version 8.8.9 Manual*, May, 2020. Cited on p. 208.
- [253] F. Maltoni, M. Selvaggi and J. Thaler, *Exposing the dead cone effect with jet substructure techniques*, *Physical Review D* **94** (2016) 054015. Cited on p. 222.
- [254] G. Gorham, *Philosophy of Science*, Beginners Guides. Oneworld Publications, 2009. Cited on p. 225.
- [255] A. Rennie and ATLAS Collaboration, “ATLAS Virtual Visits - Take part from anywhere in the world.” <https://cds.cern.ch/record/2800125>, 2022. Cited on pp. 230 and 263.
- [256] A. Buckley, “The hepthesis L^AT_EX class.” Cited on p. 232.

List of figures

1.1. The Standard Model	7
1.2. Vertices for different SM interactions.	9
1.3. $t\bar{t}$ final states.	13
1.4. $t\bar{t}$ production.	13
1.5. $t\bar{t}$ production at LO and NLO.	13
1.6. ATLAS cross-section measurements of SM processes.	14
1.7. Simulation of a $t\bar{t}$ event.	17
2.1. The CERN accelerator complex.	22
2.2. ATLAS during construction.	23
2.3. Schematic of the ATLAS detector.	24
2.4. ATLAS operating status over the course of Run 2.	25
2.5. Interactions of particles with ATLAS.	25
2.6. The ATLAS Inner Detector.	27
2.7. ATLAS ID SCT module.	29
2.8. ATLAS LAr calorimeter.	32
2.9. The ATLAS Muon Spectrometer.	34
2.10. Exploded view of ATLAS MS.	35
2.11. ATLAS magnet system	37

2.12. ATLAS TDAQ system.	38
3.1. ℓ +jets channel for boosted $t\bar{t}$	44
3.2. Example of IR unsafe behaviour.	46
3.3. Example of collinear unsafe behaviour.	46
3.4. Comparison of jet clustering algorithms.	51
3.5. Jet resolution for LC jets and PFlow jets.	53
3.6. Pileup behaviour of LC jets and PFlow jets.	53
3.7. Particle flow cell subtraction procedure.	57
3.8. Particle flow algorithm.	57
3.9. Scale uncertainties for RC and trimmed LC jets.	58
3.10. ATLAS JES calibration.	59
3.11. b -jet properties.	71
3.12. Performance of b -tagging algorithms.	72
4.1. Migrations for cluster-based and track-based substructure, for both LCTopo and RC large- R jets.	86
4.2. Comparison of τ_{32} for RC and LCTopo jets.	87
4.3. Δp_T for matched LCTopo and RC jets.	88
4.4. Δp_T for matched particle-level LCTopo and RC jets.	88
4.5. Top quark mass spectrum for LCTopo and RC jets.	89
4.6. p_T distribution for jets failing p_T cut.	90
4.7. p_T distribution for jets failing mass window cut.	91
4.8. Pileup dependence of LCTopo and RC jets.	92
5.1. ℓ +jets boosted $t\bar{t}$ topology.	93
5.2. Jet mass distribution prior to final selection criteria.	97

5.3. Effect of b -tagging on τ_{32} distribution.	98
5.4. $m_{\ell b}$ distribution without cut.	98
5.5. τ_{32} distribution compared for different selection criteria.	99
5.6. Signal efficiency compared for $m_{\ell b}$ cut.	100
5.7. Background acceptance compared for $m_{\ell b}$ cut.	100
5.8. $m_{\ell b}$ ROC curve.	101
5.9. $m_{\ell b}$ efficiency double ratio.	102
5.10. Correlation between τ_{32} and $m_{\ell b}$	102
5.11. Background composition.	103
5.12. τ_{32} in data and MC for varied mass window.	104
5.13. Methods for choosing top candidate jet.	105
5.14. Leading-order single top production.	106
5.15. Fake lepton estimate in regions of E_T^{miss}	109
5.16. Fake lepton estimate in regions of $E_T^{\text{miss}} + m_T^W$	110
5.17. Distribution of fake lepton events for substructure observables.	110
5.18. τ_{32} distribution compared for different fake lepton methods.	111
5.19. Lepton $ \eta $ distribution.	112
5.20. Event weights as a function of $ \eta $	112
5.21. Fake lepton distribution in $ \eta $, only small event weights allowed.	113
5.22. Distribution of fake electrons and muons.	114
5.23. Detector-level distribution for $m_{\ell b}$ and $E_T^{\text{miss}} + m_T^W$	115
5.24. Detector-level distribution for number of jets and number of b -jets.	116
5.25. Detector-level distribution for jet $ \eta $ and E_T^{miss}	117
5.26. Detector-level distribution for number of large- R jets and large- R jet $ \eta $	118
5.27. Detector-level distribution for large- R jet ϕ and jet p_T	119

5.28. Detector-level distribution for top jet p_T and top jet mass.	120
5.29. Detector-level distribution for lepton $ \eta $ and lepton ϕ	121
5.30. Detector-level distribution for lepton p_T	122
6.1. Particle- and detector-level comparisons for the τ_{32} substructure observable.	126
6.2. Particle- and detector-level comparisons for the τ_{21} substructure observable.	127
6.3. Particle- and detector-level comparisons for the ECF2 substructure observable.	128
6.4. Particle- and detector-level comparisons for the D_2 substructure observable.	129
6.5. Particle- and detector-level comparisons for the C_3 substructure observable.	130
6.6. Particle- and detector-level comparisons for the $p_T^{d,*}$ substructure observable.	131
6.7. Particle- and detector-level comparisons for the LHA substructure observable.	132
6.8. χ^2 comparison for considered substructure observables with different NLO+PS predictions.	133
6.9. Migration matrices for τ_{32} and τ_{21} substructure observables.	137
6.10. Migration matrices for τ_3 and ECF2 substructure observables.	137
6.11. Migration matrices for C_3 and LHA substructure observables.	138
6.12. Migration matrices for D_2 and $p_T^{d,*}$ substructure observables.	138
6.13. Efficiency and acceptance corrections.	140
6.14. Efficiency and acceptance corrections for double-differential measurement.	141
6.15. Average correlation factor over iterations of unfolding.	144

6.16. Average correlation factor over iterations of unfolding, double-differential.	145
6.17. χ^2 between unfolding iterations.	146
6.18. χ^2 between unfolding iterations, double-differential.	146
6.19. Statistical uncertainties as a function of unfolding iterations.	146
6.20. Results of unfolding closure test for τ_{32} and D_2	148
6.21. Unfolding pull test results for substructure observables.	150
6.22. Unfolding pull test results for substructure observables in regions of the top jet mass and p_T	150
6.23. Unfolding stress functions.	152
6.24. Results of the unfolding stress test.	152
6.25. Results of the unfolding stress test, double-differential.	153
7.1. Distribution of uncertainties for the τ_{32} and τ_{21} substructure observables.	158
7.2. Distribution of uncertainties for the τ_3 and ECF2 substructure observables.	159
7.3. Distribution of uncertainties for the C_3 and LHA substructure observables.	159
7.4. Distribution of uncertainties for the D_2 and $p_T^{d,*}$ substructure observables.	160
7.5. Distribution of uncertainties for the τ_{32} substructure observables in regions of the top jet mass.	161
7.6. Distribution of uncertainties for the τ_{32} substructure observables in regions of the top jet p_T	162
7.7. Distribution of uncertainties for the D_2 substructure observables in regions of the top jet mass.	163
7.8. Distribution of uncertainties for the D_2 substructure observables in regions of the top jet p_T	164
8.1. Unfolded results for the τ_{32} and τ_{21} substructure observables.	165
8.2. Unfolded results for the τ_3 and ECF2 substructure observables.	166

8.3. Unfolded results for the C_3 and LHA substructure observables.	167
8.4. Unfolded results for the D_2 and $p_T^{d,*}$ substructure observables.	168
8.5. Unfolded results for the τ_{32} substructure observable in regions of the top jet mass and p_T	171
8.6. Unfolded results for the D_2 substructure observable in regions of the top jet mass and p_T	172
8.7. Unfolded results for the τ_{32} substructure observable in regions of the top jet mass.	172
8.8. Unfolded results for the τ_{32} substructure observable in regions of the top jet p_T	173
8.9. Unfolded results for the D_2 substructure observable in regions of the top jet mass.	173
8.10. Unfolded results for the D_2 substructure observable in regions of the top jet p_T	174
9.1. Schedule for the HL-LHC upgrade [218].	179
9.2. ATLAS SM cross-section measurements.	180
9.3. HL-LHC magnet configuration at ATLAS IP.	183
9.4. Effect of the RF crab cavities on β^* at the HL-LHC.	184
9.5. HL-LHC $t\bar{t}$ event display.	185
9.6. ATLAS ITk layout.	186
9.7. Model of the ATLAS ITk.	187
9.8. Hits per track in silicon layers for ITk.	187
9.9. Schematics showing the general makeup of a hybrid pixel detector [238].	190
9.10. Expected fluence in ATLAS during HL-LHC.	191
9.11. Diagrams of ITk pixel sensor technologies.	193
9.12. ITk pixel quad module.	193

10.1. The testbeam facility at the CERN SPS [243].	196
10.2. The testbeam facility at DESY [242].	196
10.3. Simulation of DESY telescope configuration.	197
10.4. Photograph of DESY beam telescope.	197
10.5. RD53A-based quad module.	199
10.6. Possible biasing structures for a pixel detector sensor [240].	200
10.7. Flowchart for testbeam reconstruction and analysis.	202
10.8. Pixel cluster size.	204
10.9. Correlations between DUT and telescope planes.	205
10.10 Track fit quality for telescope tracks.	206
10.11 Residuals between DUT cluster and telescope tracks.	206
10.12 Example in-pixel efficiency map.	210
10.13 Efficiency vs. threshold for no PTB at $V_{\text{bias}} = 600 \text{ V}$	213
10.14 In-pixel efficiency map for differential FE, no PTB, threshold 1200– 1600 e^- , $V_{\text{bias}} = 600 \text{ V}$	213
10.15 In-pixel efficiency map for differential FE, no PTB, threshold 1158– 2136 e^- , $V_{\text{bias}} = 600 \text{ V}$	214
10.16 Efficiency vs. threshold for zigzag and no PTB at varied V_{bias}	215
10.17 In-pixel efficiency map for linear FE, zigzag PTB, threshold 900–1200 e^- , $V_{\text{bias}} = 600 \text{ V}$	215
10.18 Device efficiency vs V_{bias} , both no PTB at $\text{THL} = 1027 e^-$ and zigzag PTB at $\text{THL} = 1200 e^-$	217
10.19 In-pixel efficiency map for zigzag PTB, linear FE, $\text{THL} = 1200 e^-$, $V_{\text{bias}} = 100 - -600 \text{ V}$	218
10.20 In-pixel efficiency map for no, linear FE, $\text{THL} = 1027 e^-$, $V_{\text{bias}} = 200 -$ -600 V	219

11.1. Trends in publications on jet substructure. 223

A.1. A virtual visit to the ATLAS detector in progress, March 2021 228

A.2. Poster on the ATLAS Virtual Visit program [255]. 230

List of tables

2.1. ATLAS calorimeter resolution.	33
4.1. Effect of RC jets on event yield.	91
5.1. Event yields.	94
5.2. Summary of the event selection procedure in the lepton + jets channel.	96
5.3. Results of the fitting and scaling process for the fake template method.	108
6.1. Groups of substructure observables considered.	124
6.2. Results of observable selection.	135
6.3. Optimised number of unfolding iterations.	144
6.4. Results of closure test.	148
8.1. χ^2 comparison of unfolded results, including nominal and alternative predictions.	175
8.2. χ^2 comparison of unfolded results, including nominal and variations of the nominal prediction.	175
9.1. Comparison of the pixel systems in the ID and ITk.	189
10.1. Silicon sensor designs tested.	200