Characterising the Decays of High-p_{T} Top Quarks and Addressing Naturalness with Jet Substructure in ATLAS Runs I and II

Matthew Edgar LeBlanc
B.Sc., Acadia University, 2011

A Dissertation Submitted in Partial Fulfillment of the
Requirements for the Degree of

DOCTOR OF PHILOSOPHY

in the Department of Physics and Astronomy

Supervisory Committee

Dr. R. McPherson, Supervisor
University of Victoria, Department of Physics and Astronomy

Dr. M. Lefebvre, Member
University of Victoria, Department of Physics and Astronomy

Dr. A. Ritz, Member
University of Victoria, Department of Physics and Astronomy

Dr. S. Dosso, Outside Member
University of Victoria, School of Earth & Ocean Sciences

Abstract

The coupling of the Standard Model top quark to the Higgs boson is $\mathcal{O}(1)$, which leads to large quantum corrections in the perturbative expansion of the Higgs boson mass. Possible solutions to this so-called naturalness problem include supersymmetric models with gluinos and stop squarks whose masses are at the electroweak scale, $\mathcal{O}(1$ TeV). If supersymmetry is realised in nature at this scale, these particles are expected to be accessible with the Large Hadron Collider at CERN. A search for gluino pair production with decays mediated by stop- and sbottom-squark loops in the initial 14.8 fb^{-1} of the ATLAS run 2 dataset is presented in terms of a pair of simplified models, which targets extreme regions of phase space using jet substructure techniques. No excess is observed and limits are set which greatly extend the previous exclusion region of this search, up to 1.9 TeV (1.95 TeV) for gluinos decaying through light stop (sbottom) squarks to the lightest neutralinos. A performance study of top tagging algorithms in the 20.3 fb^{-1} 2012 dataset is also presented, which includes the first measurements of substructure-based top tagging efficiencies and fake rates published by ATLAS, as well as a detailed comparison of tagger performance in simulation. A benchmarking study which compares commercially available cloud computing platforms for applications in High Energy Physics, and a summary of ATLAS liquid argon calorimeter data quality work focused on monitoring and characterising the sporadic phenomena of Mini Noise-Bursts in the electromagnetic barrel calorimeter are also included.

Declaration

There are over 3000 scientific authors on publications made by the ATLAS collaboration. This document is an attempt to collect the work of just one. Collaboration within ATLAS occur without regard for age, affiliation or any other possible barrier. The calibrations applied to leptons in these studies were derived due to the work of many people whom the author will never meet. The cross-checks of jet reclustering presented here could be cited by searches about which the author does not know. It can occasionally seem a bit tricky to disentangle the work of one student from the collective ATLAS effort, though it is necessary to be explicit given the context of this document.

Work which the author played a central role in, presented within this dissertation, includes:

- The benchmarking studies of commercial cloud resources presented in section 2.4.1 formed the basis of my ATLAS authorship qualification task through the computing group, and were carried out within the High Energy Physics Research Computing group at the University of Victoria. These studies were presented at the *21st International Conference on Computing in High Energy and Nuclear Physics (CHEP)* conference held in Okinawa, Japan during April 2015. They were published in the proceedings of that coference [1].

- The studies presented in chapter 4 are a subset of those presented at the *7th International Workshop on Boosted Object Phenomenology, Resonstruction and Searches in HEP (BOOST)* held in Chicago, USA during August 2015. These results were subsequently published in the Journal of High Energy Physics [2]. I was the primary (sole) analyser responsible for studies of the jet substructure-based tagging methods, which includes the comparisons between data and simulation, efficiency and fake rate measurements, as well as the substructure tagger inputs for the direct comparison studies. This work was performed in collaboration with researchers at the University of Edinburgh and the Universität Heidelberg, who studied other top tagging methods also presented in this paper and with whom some centralised software was developed.

- The search detailed in chapters 5, 7 and appendix C was presented as a preliminary ATLAS result at the *38th International Conference on High Energy Physics* held in Chicago, USA during August 2016. I was the main analyser on this iteration of the search, and was responsible for nearly the entire chain of analysis from the initial optimisation of the stop-mediated regions (though not the sbottom-mediated ones) to the final limit setting procedure (for both the stop- and sbottom-mediated results). I also introduced and optimised the reclustered $M_J^{\sum,4}$ observable, which was newly introduced for this iteration of the search as a replacement for reclustered top tagging. A publication which builds upon these preliminary results with the complete run 2 dataset is currently in preparation. This work was performed in close collaboration with researchers from the University of Chicago, l'Université de Montréal, McGill University and the Institut de Física d'Altes Energies (Barcelona).

- The studies of the performance of jet reclustering presented in appendix B originated as cross-checks I performed for the search mentioned above, particularly those studying close-by effects of jets. This work was collected internally but not published with the results of that search. Instead, a publication documenting the performance of this technique is currently under preparation, which will allow these studies to form their own public result. In addition to performing these studies, I am one of three internal editors of this result, and am therefore partially responsible for both documenting it and overseeing its progress through the ATLAS approval procedure.

- The studies of data quality in the liquid argon calorimeter systems in relation to the phenomena of mini noise bursts, presented in appendix A, was carried out in association with the LAr Data Quality coordinator over the course of ATLAS operations in 2016. In particular, I implemented the monitoring code within the core ATLAS software and also performed the offline studies which are described.

Acknowledgements

It's been almost seven years to the day since I first arrived at CERN as a Canadian summer student in 2010. There are many incredible people who have done a lot to support and push me since then – I'm supposed to acknowledge them here. Before naming individuals, I feel that it's important to express how much gratitude I have for the IPP and TRIUMF summer student programmes, which offered me the opportunity to begin this journey in the first place.

Several postdoctoral fellows from many different institutes have directly acted as my mentors at various points of my studies. Alex Martinyiuk, Danilo Enoque Ferreira De Lima, Johannes Erdmann, Frank Berghaus, Max Swiatlowski and Emma Kuwertz are all fine physicists with deeper reserves of patience than anyone else I've ever met.

My supervisor Rob McPherson provided both constant guidance and space for me to develop my own skills and interests as a physicist – this journey would have surely been a more difficult one without his knowledge and experience. I must also gratefully acknowledge Michel Lefebvre and Dave Axen for many bright exchanges and enthusiastic support at different periods over the course of my studies in Victoria and at TRIUMF.

It has been a joy to work within the ATLAS Jet/EtMiss Combined Performance group over the course of my degree. This community as a whole is one of the most welcoming and engaging that I've ever encountered, and I look forward to continuing to collaborate with this outstanding group of people for as long as I am able to do so.

Other teachers of physics and mathematics had an impact on my trajectory even earlier, and I owe them a lot as well. Their guidance sometimes came in different ways – sometimes in the classroom and sometimes just passing in the hallway – but I believe that things would have probably wound up differently if any of them hadn't been there at the right time. Max Turton, Nick Down, Erick Lee and Paul Myers from JLI; Mike Robertson, Svetlana Barkanova and Anna Kiefte from the Acadia physics department; Richard Karsten, Holger Teismann and Franklin Mendivil from Acadia Mathematics. A pair of math graduate students also had an impact which went beyond first-year calculus: Natasha Mandyrk and Avra Laaraker.

Some people have been supporting me for even longer than those I've already mentioned. My family has always been a universal constant to whom I will never be able to fully express my overwhelming gratitude (thanks Mom, Dad, Nick & everyone else!). Many friends – too many to name – have also been present without fail when times were toughest, no matter where I happened to be on the planet. Kate and Emma are two with whom I cross paths with more often than most, and so they've likely had to put up with a lot more than the others: thank you!

Once again, to everyone listed here, and to everyone who isn't but should be: thank you, thank you, thank you. I'll do my best to make you proud.

Contents

List of Figures

List of Tables

1

Introduction

1.1 The Standard Model of Particle Physics

According to our present scientific understanding of the universe, all matter interacts via one or more of four fundamental forces:

I **Gravity**

II **Electromagnetism**

III **The Weak Force**

IV **The Strong Force**

Gravity is well-described by Einstein's classical theory of General Relativity. The remaining three forces are described by a single quantum theory, called the **Standard Model of Particle Physics** ('The Standard Model', or SM), to high precision. No attempt at describing gravity with a quantum theory has been successful to-date, and so it has not been incorporated within the Standard Model. The impact of this omission at the energy scales presently examinable in the laboratory is negligible. The effects of gravitation become relevant within the Standard Model at the Planck scale, which corresponds to energies of $E_{\text{Planck}} = \sqrt{\hbar c^5 / G_{\text{Newton}}} \sim 1.2 \times 10^{19}$ GeV – well beyond the reach of any modern technology.

Several particles are described by the Standard Model, which may be classified generally into two groups. **Fermions** compose all matter in the universe, possess a half-integer spin quantum number, and interact via the fundamental forces listed above. **Bosons** are associated with each force, posses an integer spin quantum number, and mediate the interactions between fermions: the photon (γ) governs the electromagnetic interaction, the $W^{\pm}$ and Z bosons govern the weak interaction, and the gluon (g) governs the strong interaction. Not

all fermions participate in every interaction: **quarks** interact strongly and electroweakly, though **leptons** only experience the latter force. One additional particle, the Higgs boson, is a remnant of the mechanism which endows the other particles with their masses. The existence of all particles predicted by the SM has been experimentally verified following the discovery of the Higgs boson in 2012 by the ATLAS and CMS collaborations at CERN's Large Hadron Collider, which allowed the 2013 Nobel Prize in Physics to be awarded to Peter Higgs and François Englert [3, 4, 5]. A summary of the familiar particle content of the SM is provided in figure 1.1.

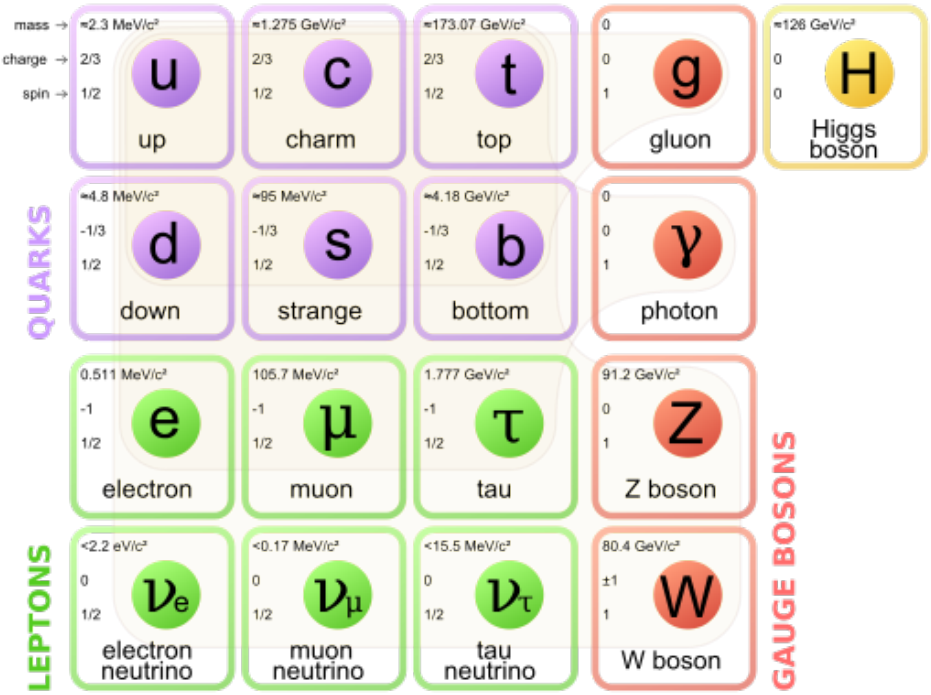

Figure 1.1: Summary of the particle content of the Standard Model of Particle Physics. The information presented has been provided by the Particle Data Group [6].

The Standard Model is formally described with a pair of Yang-Mills quantum field theories (QFTs), based on the overarching symmetry group $SU(3)_C \times SU(2)_L \times U(1)_Y$. The Glashow-Weinberg-Salam Electroweak Theory describes the dynamics of the unified electromagnetic and weak interactions, and is associated with the symmetry group $SU(2)_L \times U(1)_Y$. The theory of Quantum Chromodynamics (QCD) describes the strong interaction, and is associated with the symmetry group $SU(3)_C$. The predictions of the SM have been verified over a wide range of energy scales and by many independent experiments, often to startling precision. Even so, the SM is known to be at least partially incomplete, and some of its aspects are considered problematic by members of the scientific community.

Based on the symmetry group, a Lagrangian density may be developed which describes the dynamics of a theory. For the sake of clarity and brevity, It is helpful to highlight only a handful of the relevant terms in further detail; more in-depth discussions of the SM at large are available elsewhere [6].

1.2 Electroweak Symmetry Breaking and Naturalness

The Glashow-Weinberg-Salam model of electroweak interactions may be described by a Lagrangian density with $SU(2)_L \times U(1)_Y$ symmetry. The kinetic portion of the Lagrangian introduces the $SU(2)$ and $U(1)$ gauge fields $W_\mu^{i=1,2,3}$ and B_μ through their respective field strength tensors, $W^{\mu\nu}$ and $B^{\mu\nu}$:

$$\mathcal{L}^{\mathbf{W},B}_{SU(2)_L\times U(1)_Y} = -\frac{1}{4}\mathbf{W}_{\mu\nu}\cdot\mathbf{W}^{\mu\nu} - \frac{1}{4}B_{\mu\nu}B^{\mu\nu} \tag{1.1}$$

where

$$\mathbf{W}^{\mu\nu} = \partial^\mu\mathbf{W}^\nu - \partial^\nu\mathbf{W}^\mu - g\mathbf{W}^\mu\times\mathbf{W}^\nu, \tag{1.2}$$

$$B^{\mu\nu} = \partial^\mu B^\nu - \partial^\nu B^\mu \tag{1.3}$$

and g is the $SU(2)$ gauge coupling constant. Linear combinations of the gauge fields $W_\mu^{i=1,2,3}$ and B_μ will become associated with the SM $W^\pm$, Z and photon. The mechanism of **electroweak symmetry breaking** allows for the transformation of these fields into the familiar physical particles, by introducing a doublet of complex scalar fields – the Higgs doublet $H(x)$ – and its corresponding potential:

$$\mathcal{L}^{\text{Higgs}}_{SU(2)_L\times U(1)_Y} = [D_\mu H(x)][D^\mu H(x)] - V(x), \tag{1.4}$$

$$H(x) = \begin{pmatrix} H^+(x) \\ H^0(x) \end{pmatrix} \tag{1.5}$$

where the covariant derivative is taken to be

$$D^\mu = \partial^\mu + i\frac{g}{2}\tau\cdot\mathbf{W}^\mu + i\frac{g'}{2}yB^\mu. \tag{1.6}$$

Here, τ are the Pauli matrices and g' is the $U(1)$ gauge coupling constant. The potential $V(x)$ corresponds to the mass- and self-interaction terms of this scalar field:

$$V(x) = \mu^2 H^\dagger(x)H(x) + \lambda[H^\dagger(x)H(x)]^2. \tag{1.7}$$

The shape of the Higgs potential is determined by the parameters μ and λ. If $\lambda > 0$, then the potential will possess some stable ground state. The choice of μ^2 is more interesting: if $\mu^2 > 0$, then the potential will possess a unique minimum at $H^\dagger H = 0$; however, if $\mu^2 < 0$, then the quartic potential will instead produce a set of identical minima with a value of

$$H^\dagger H = -\frac{\mu^2}{2\lambda} \equiv \frac{v^2}{2}. \tag{1.8}$$

The quantity v is the **vacuum expectation value** (vev) of the Higgs field, whose value dictates the electroweak scale. The Higgs vev is the only dimensionful parameter of the Standard Model, and has been measured to highest precision by the MuLan collaboration, who quote a value of ∼242 GeV with a precision of 0.6 ppm [7]. This degenerate set of ground states provides a mechanism with which to spontaneously break the $SU(2)_L \times U(1)_Y$ symmetry while maintaining the Lagrangian's gauge invariance, ultimately granting mass to the $W^\pm$ and Z vector bosons and producing the physical Higgs boson.

Fermion masses are also generated via the Higgs mechanism, by manually introducing Yukawa interactions to the Lagrangian which imbue each fermion with its experimentally-determined mass. These terms take the form

$$\mathcal{L}_{fH} = -\lambda_f[\bar{f}_L H f_R + \bar{f}_R H^\dagger f_L] \tag{1.9}$$

where λ_f is the **Yukawa coupling** of the particular fermion, and f_L, f_R and H are respectively the fermion left-handed doublet, right-handed singlet and Higgs scalar doublet, which may be inserted explicitly:

$$\begin{aligned} L_{fH} &= -\frac{\lambda_f}{\sqrt{2}}[v(\bar{f}_L f_R + \bar{f}_R f_L) + \bar{f}_L h f_R + \bar{f}_R h f_L] \\ L_{fH} &= -\frac{\lambda_f v}{\sqrt{2}}\left(f\bar{f} + 2hf\bar{f}\right). \end{aligned} \tag{1.10}$$

The first term is identified as the fermion mass term, and the second term represents the fermion-Higgs interaction. The mass term and coupling are:

$$m_f = -\frac{\lambda_f v}{\sqrt{2}} \tag{1.11}$$

$$h\bar{f}f : \; -\frac{i\lambda_f}{\sqrt{2}} \tag{1.12}$$

which implies that the Yukawa couplings of *most* fermions are much smaller than 1.

Of particular interest in this context is the top quark (t), which was discovered in 1995

by the CDF and D0 collaborations at Fermilab's Tevatron. The top has an electric charge of $+2/3$ and a mass of 173.3 GeV, making it the most massive known fundamental particle. This large mass implies a short lifetime; top quarks decay within $\sim 0.5 \times 10^{-24}$ s of their creation. This lifetime is shorter than the characteristic timescale of hadronisation (section 1.4), and so top quarks are the only quarks which decay while 'bare,' nearly always[1] into a b-quark and W boson. As the top quark couples so strongly to the Higgs boson, it is fundamentally connected to the properties of electroweak symmetry breaking, and physics beyond the Standard Model within this sector could manifest itself through anomalous properties or production of top quarks.

The top quark is also linked to an aesthetic problem of the Standard Model, known as the **hierarchy problem**, which is rooted in the enormous difference in energy scales between the electroweak and gravitational sectors[2]. An explicit sensitivity to the scale to which the SM is valid enters the theory through the mass of any fundamental scalar particle, such as the Higgs boson. The masses of scalar particles receive higher-order corrections in perturbation theory due to fermion loops, such as the diagram shown in figure 1.2(a). These corrections take the form

$$\delta m_h^2 = -\frac{\lambda_f^2}{8\pi^2}\Lambda^2 + ... \tag{1.13}$$

where Λ is the energy scale used to regularise the loop integral. This scale is generally taken to be the scale to which the SM is known to be valid, which could be as large as M_{Planck}! The Yukawa coupling of each fermion also enters into these corrections due to the fermion-Higgs vertices, and so the dominant correction is made by the top quark, for which $\lambda_f \sim \mathcal{O}(1)$. There are two possible interpretations of this problem:

1. The value of the Higgs mass happens to be fine-tuned, resulting in the measured value.

2. There is another, intermediate, energy scale below which the SM is valid. At this scale, new physics could help mitigate these large corrections.

The case that nature is simply fine-tuned is possible, though does not lead to conclusions which may be pursued further at the present time. The second possibility is more intriguing, particularly after noting that higher-order contributions to the value of the Higgs mass which arise from new bosons would enter into the series with an opposite sign. In the event that these new particles shared the Yukawa couplings of the SM fermions, these large quantum corrections would neatly remove each other from the sum. Such an argument is known as an

[1]The suppressed decays $t \to Ws$ and $t \to Wd$ account for about $\sim 0.1\%$ and $\sim 0.01\%$ of the total branching ratio, the remainder being associated with the $t \to Wb$ decay.

[2]Recall that $\Lambda_{\text{Planck}}/\Lambda_{\text{Electroweak}} \sim \mathcal{O}(10^{17})$!

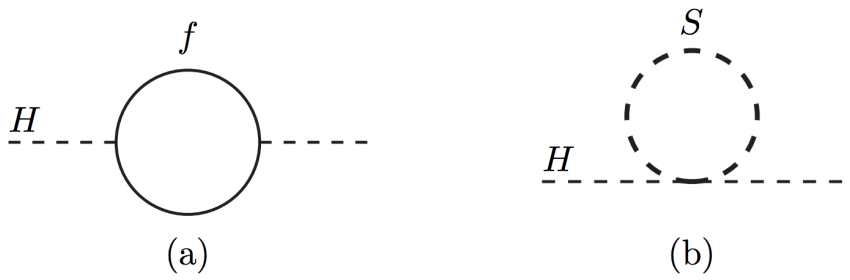

Figure 1.2: Loop diagrams which contribute quantum corrections to the mass of a scalar particle such as the Standard Model Higgs boson, originating from (a) a fermion and (b) a boson. These diagrams contribute with different signs in the series of corrections, and so could partially compensate for each other. From [8].

appeal to **naturalness**, and relies on an underlying belief that fundamental theories which properly describe nature should not be afflicted with dramatically fine-tuned parameters.

Natura valde simplex est et sibi consona.

1.3 Supersymmetry

Today, the most popular theories which extend the Standard Model in a way which addresses this problem of naturalness are models which exhibit **supersymmetry**, in which new particles are introduced which form a correspondence between fermions and bosons. In the simplest supersymmetric models, each particle in the SM receives a superpartner whose spin differs by 1/2, but otherwise possesses the same quantum numbers. It follows, then, that the superpartners of SM fermions are bosons, and vice-versa: the hierarchy problem solves itself! Building on the earlier example, contributions from the new scalar particle loops (figure 1.2(b)) will contribute Higgs-mass corrections with size

$$\delta m_h^2 = \frac{\lambda_s^2}{16\pi^2}\Lambda^2 + ... \tag{1.14}$$

It is important to note that both the right- and left-handed SM fermions receive partners under SUSY. Two positive bosonic loops arise for each negative fermionic loop, and the contributions from each perfectly balance!

One caveat to the above remarks is that supersymmetry has not yet been observed in nature. This implies that not all of the quantum numbers may be shared between the SM particles and their superpartners: the masses of the superpartners must be sufficiently larger than their SM counterparts that they would have avoided detection to-date. This

requirement may be satisfied by including an additional term in the SUSY Lagrangian which breaks the new symmetry in such a way that the masses of the superpartners are increased. This process is referred to as 'soft' symmetry breaking, as it is performed with some care in order to avoid the introduction of additional strongly-divergent terms to the theory. An additional correction to the Higgs mass value which arises from this new sector will take the form

$$\delta m_h^2 = m_{SOFT}^2 \left[\frac{\lambda_f^2}{16\pi^2} \ln(\frac{\Lambda}{m_{SOFT}}) \right], \tag{1.15}$$

where m_{SOFT} is the characteristic scale at which the SUSY breaking occurs, which sets the possible scale of superpartner masses. In practice, if this scale becomes too large, these corrections will become larger than the problem they are meant to solve. An estimate of m_{SOFT} can be made by taking Λ as the Planck scale, and λ_f to be about 1: in this case, m_{SOFT} ought to be at the TeV-scale – with luck, within reach with modern technology. A light (TeV-scale) stop squark is thus the most crucial ingredient in a natural explanation of the light SM Higgs mass.

One downside of light scalar partners introduce their own sets of potentially-large quantum corrections, and in the case of the stop squark, it is the contribution from the gluino ($\tilde{g}$) which plays the dominant role. The gluino couples strongly to squarks due to their mutual colour charge, and this strong coupling causes this correction to inflate the mass of the squark. Ultimately, the gluino must also take a mass near the TeV-scale, in order to avoid the introduction of large corrections elsewhere in the theory.

Certain supersymmetric models offer other appealing properties. A certain class of models requires that superparticles are produced and annihilated in pairs, due to an imposed symmetry known as R-pairity. In such R-parity conserving (RPC) models, the lightest supersymmetric particle (LSP) is stable, and may form a good candidate for WIMP dark matter if it is sufficiently light [9, 10, 11].

1.4 Quantum Chromodynamics (QCD)

The dynamics of quarks and gluons, which participate in the strong interaction, are described by the theory of Quantum Chromodynamics (QCD). These particles possess a nonzero colour quantum number, which may take one of three values commonly referred to as red, blue or green. Quarks possess a single colour charge, antiquarks an anticolour charge, and gluons simultaneously both colour and anticolour charges.

The Lagrangian associated with the theory of Quantum Chromodynamics corresponds

with the $SU(3)_C$ symmetry group:

$$\mathcal{L}_{SU(3)_C} = \bar{q}(i\gamma_\mu D^\mu - m_q)q - \frac{1}{4}G_a^{\mu\nu}(G_a)_{\mu\nu} \tag{1.16}$$

where the covariant derivative is

$$D^\mu = \partial^\mu + ig_s\frac{\lambda_a}{2}G_a^\mu, \tag{1.17}$$

and $G_a^{\mu\nu}$ is the $SU(3)_C$ field strength tensor,

$$G_a^{\mu\nu} = \partial^\mu G_a^\nu - \partial^\nu G_a^\mu + g f_{abc} G_b^\mu G_c^\nu, \tag{1.18}$$

where g_s and f_{abc} are the $SU(3)_C$ gauge coupling constant and structure constants. The index a runs over the eight gluons which play the role of the force mediators in the theory. The first term of the Lagrangian describes the interactions of quarks and gluons with each-other, while the second term describes the self-interactions of the gluons amongst themselves.

The behaviour of the $SU(3)_C$ gauge coupling $g_s(\mu)$, or equivalently the strong coupling constant $\alpha_s(\mu)$, as a function of the interaction energy scale μ plays a central role in QCD. The coupling $\alpha_s(\mu)$ runs with energy as

$$\alpha_s(\mu) = \frac{4\pi}{\beta_0 \ln(\mu^2/\Lambda_{QCD}^2)} \tag{1.19}$$

where β_0 and Λ_{QCD} are measured constants. The inverse logarithmic dependence on the energy scale causes QCD to be strongly coupled ($\mathcal{O}(1)$) at low energy scales, and weakly coupled ($\alpha_s << 1$) at high energy scales. The strength of the α_s coupling has been measured at many energy scales by many experiments; a summary of their measurements is given in figure 1.3 along with the theoretical prediction.

Two distinctive traits of QCD emerge due to this dependence. At large energy scales (equivalently, small distance scales), the coupling is small and so the theory is well-described by perturbative methods. This phenomenon is known as **asymptotic freedom**; it allows the calculation of observables related to quarks and gluons, such as their production rates, using perturbative approaches similar to those used in electroweak theory. At small energy scales (or, large distance scales), α_s becomes very large, rendering the theory nonperturbative. Due to this strong coupling, quarks and gluons are subject to the phenomenon of **confinement**. As the distance between two particles which carry a colour charge increases, the strength of the force acting between them also increases. At a certain point, the binding energy between the two particles exceeds the energy threshold for quark-antiquark production, and

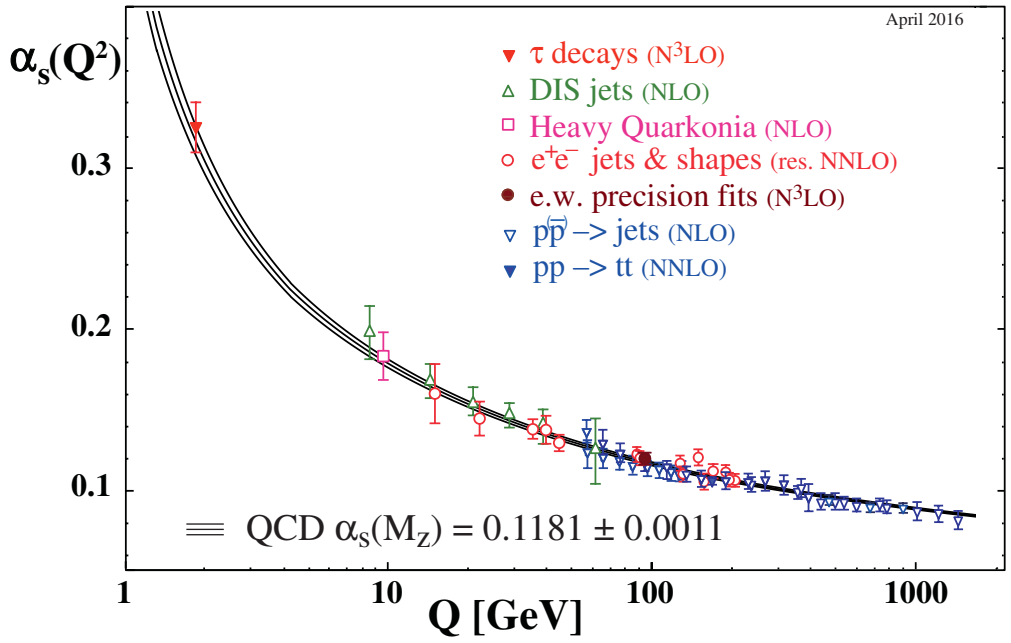

Figure 1.3: The dependence of the strong coupling α_s on the measured energy scale, here labeled Q. At low values of Q, the coupling grows large, resulting in the phenomena of confinement, which causes strongly-interacting particles to be observable only in colour-singlet states. At large values of Q, the coupling decreases, resulting in asymptotic freedom – in this regime, perturbative approaches are applicable to QCD. From [6].

new quark and gluon pairs are created in a process known as **hadronisation**. These new particles subsequently undergo **fragmentation**, during which they couple with the original pair and each other in order to produce a set of colour-singlet hadrons, each of which shares some of the original particle's momentum. These **jets** of hadrons are the experimentally-observable remnants of the quarks or gluons produced in an interaction at a high energy scale, such those produced by proton-proton collisions within the Large Hadron Collider.

1.5 Simulation

The strong coupling of QCD renders a number of processes relevant to collider physics non-perturbative. Already mentioned is the process of parton showering, during which particles which experience the strong interaction hadronise and fragment to form the physical jets of

colour-singlet hadrons which we observe experimentally. The dynamics which occur within hadrons also take place within the nonperturbative regime of QCD, and are critical to understand when considering, for example, the consequences of colliding a pair of hadrons head-on, at high energy.

To this end, studies of modern high energy physics rely heavily on simulation of these processes using various techniques and models in order to obtain an expectation which may be compared with experimental data. These simulations are typically reliant on random samplings, or **Monte Carlo** (MC) techniques, in order to produce statistical predictions for a given process. Many different MC generators are applied in various situations – some may be more suitable for the modelling of one process than another, or may be more computationally efficient when a large sample is required. Often, predictions produced by different simulations are compared to each other as well as to data, in order to evaluate the impact of these choices of generator and parton shower algorithms, along with any uncertainties they introduce.

Two ubiquitous generator choices are Pythia [12] and Herwig [13], which vary particularly in their models of parton showering. The evolution of the parton shower splitting probabilities is ordered by p_{T} in Pythia and by angle in Herwig, prioritising high-energy or high-angle emissions, respectively. The treatment of the binding energy present between strongly-interacting particles during hadronisation also differs between the two generators. Other common generator choices include Sherpa [14], MC@NLO [15, 16] and MadGraph [17], the use of which depends somewhat on the process to be simulated.

The protons which are accelerated and made to collide within the LHC are colour-singlet objects, themselves composed of quarks and gluons. Within each proton, two u and one d quark occupy the 'valence' positions and define the proton's quantum numbers (electromagnetic charge, spin, etc.). These valence quarks are bound by a complex tapestry of gluons and virtual quarks which compose the 'quark sea.' Any of these valence or sea particles – collectively referred to as **partons** – in either proton could be the particles which produce the 'hard scattering' event during a collision, and each of these partons carries a different fraction of the proton's total energy. The probability that a given parton may participate in a scattering event is summarised by **parton distribution functions** (PDFs), which provide these likelihoods as a function of the hard-scattering energy. PDF sets are produced by various collaborations as a function of the energy fraction x of each parton within the proton, and the energy scale of the scattering event, Q^2. An example PDF set from the NNPDF collaboration is shown in figure 1.4, at two different Q^2 values. At low Q^2, the valence quarks carry most of the proton's energy and are more likely to interact during a collision than other partons. As the energy scale increases, however, the sea quarks (even those of heavy flavour) and gluons become more likely to play a role.

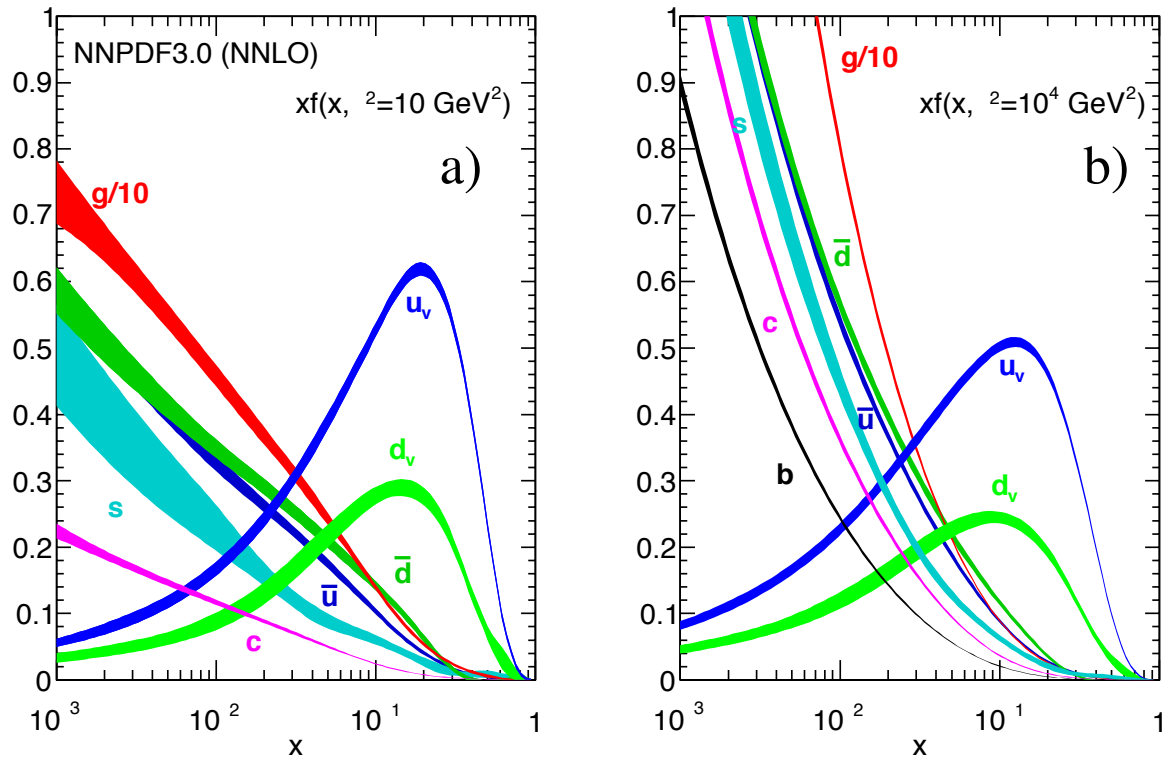

Figure 1.4: An example PDF set from the NNPDF collaboration, at two different Q^2 values (energy scales). At low Q^2, the valence quarks carry most of the proton's energy and are more likely to interact during a collision than other partons. As the energy scale increases, however, the sea quarks and gluons become more likely to play a role. From [6, 18].

2

Infrastructure & Apparatus

2.1 CERN, the European Organisation for Nuclear Research

Physics saw its first large-scale collaborations over the course of the Second World War. During this time, unprecedented groups of of physicists and engineers had been assembled to pursue focussed research programmes with unprecedented levels of funding. Subsequently, the idea of large-scale collaborations for non-military purposes became appealing, particularly given recent developments in the field of accelerator physics. At the Berkeley Radiation Laboratory, an ambitious accelerator programme was proposed by E. O. Lawrence, which included a synchro-cyclotron built from converted wartime equipment in its first phase. This machine was operational by 1946 and upgraded in 1949 to be capable of producing mesons in the laboratory for the first time. Previously the sole domain of cosmic ray physicists, now mesons could be created with high rates in reproducible conditions, allowing for more detailed study of the properties of these particles.

On the theoretical frontier, the conceptual difficulties of QED were being tidied up by the likes of Tomonaga, Schwinger and Feynman during the late fourties. Theorists began to set their sights beyond the description of electrons and photons, towards the plethora of new particles which had been recently discovered by cosmic ray physicists, whose precision study would now be enabled for the first time by the new Berkeley accelerator. Feynman acknowledged the importance of the interplay between experimental physicists and the theoretical community in 1948 following a theory conference held in Pocono, Pennsylvania; writing in his report on the meeting [19] that

> "The conference showed that just as we were apparently closing one door, that of the physics of electrons and photons, another was being opened wide by the

experimenters, that of high-energy physics. The remarkable richness of new particles and phenomena presents a challenge and a promise that the problems of physics will not all be solved for a very long time to come."

The situation in Europe following the war had not been as fruitful for the development of high-energy physics as that in America. Individual continental states lacked the resources to construct the costly infrastructure necessary to perform high-energy physics research[1]. The foundation of the UN Atomic Energy Comission (AEC) in 1946 provided an international forum where diplomats and their scientific advisors – the likes of Pierre Auger and J. R. Oppenheimer – could discuss the future of European physics research. At the time, an initiative to begin a joint European laboratory dedicated to the study of nuclear physics seemed to be the most likely means to perform competitive high-energy science in the region, by pooling available resources. Interest was also taken by pro-European parties, hoping to foster a sense of post-war unity through undertaking in fundamental science.

These discussions soon spread beyond the AEC. A letter from Louis de Broglie to the European Cultural Conference in Lausanne during December 1949 set the stage for the first official large-scale international discussions about how to practically realise this dream, spurred on by a resolution drafted by American physicist Isidor Rabi which was adopted at the UNESCO General Conference in Florence the following June which authorised UNESCO to "encourage the creation of regional research laboratories in order to increase international collaboration." In December 1951, discussions were held in Paris which lead to the establishment of a provisional council to manage the formation of the laboratory – the *Conseil Européen pour la Recherche Nucléaire* (CERN) – consisting of delegates from eleven countries[2]: Belgium, Denmark, France, Greece, the Federal Republic of Germany, Italy, the Netherlands, Norway, Sweden, Switzerland and Yugoslavia. The laboratory's site was chosen to straddle the Franco-Swiss border near Geneva, where construction began in May of 1954. By the end of that year, the convention which formally established the multinational laboratory had been ratified by its twelve founding members, and the provisional council was disbanded: CERN as we know it today officially came into being.

During the sixty years since, CERN has been the leading example of international scientific collaboration at a large scale. Today, it represents the combined interests of 22 member states, hosts users from over 600 universities around the world and is an observer of the UN

[1]Several high-energy accelerators had been proposed and built in the United Kingdom in the post-war period. The largest in continental Europe during this period was built in Uppsala, Sweden, at the Gustaf Werner Institute for Nuclear Chemistry. This synchro-cyclotron circulated its first beam in 1951. Its energy did not exceed the meson production threshold.

[2]The United Kingdom officially joined the provisional council later, though a British delegation did

General Assembly. The primary concern of the laboratory is to "provide for collaboration among European states in nuclear research of a pure scientific and fundamental character [...] (CERN) shall have no concern with work for military requirements." These provisions primarily take the form of a complex accelerator chain and the staff necessary to maintain and operate it, which supplies high-energy protons and heavy ions for a variety of experiments managed separately by international collaborations. A wide programme of research in fundamental physics is pursued with a varied system of many accelerators (figure 2.1), the current cornerstone of which is the Large Hadron Collider (LHC).

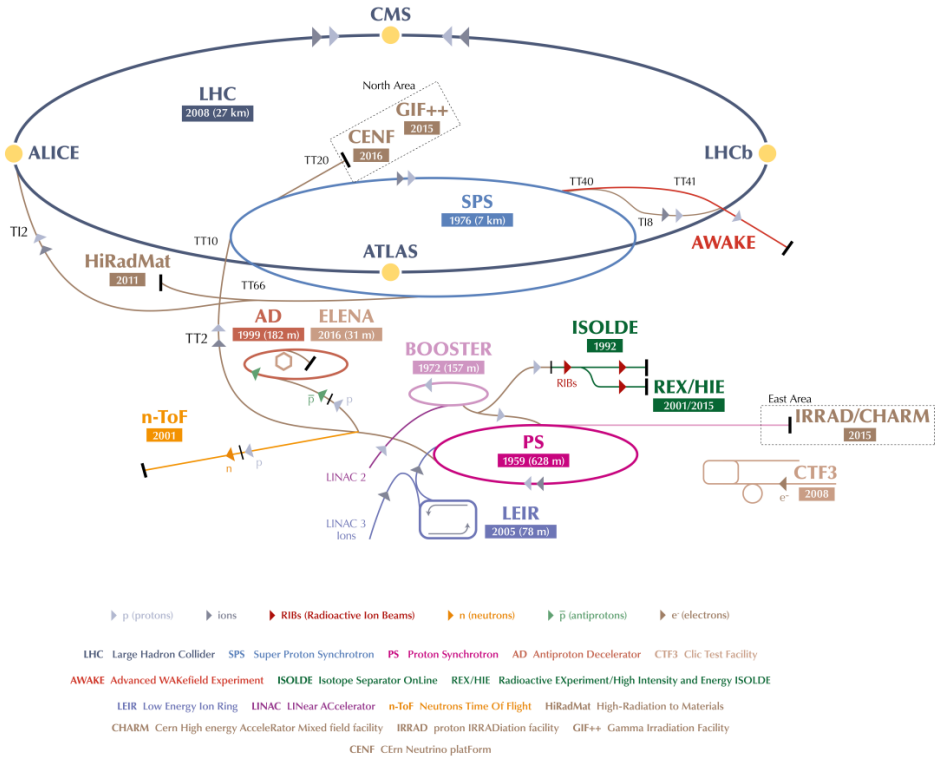

Figure 2.1: The CERN accelerator complex. From [20].

All protons at CERN begin their journey from a source canister of hydrogen gas, which are ionized and then accelerated to an initial beam energy of 50 MeV by a linear accelerator, LINAC2. This beam is transferred to the Proton Synchrotron (PS) Booster, a set of four stacked synchrotron rings which increase the beam energy to 1 or 1.4 GeV before sending the beam to the PS, which provides further acceleration to 25 GeV before passing it again

participate in these formative discussions.

into the Super Proton Synchrotron (SPS).

The energy of proton beams in the SPS is increased to 450 GeV, the minimum beam energy which the Large Hadron Collider is capable of accepting. During the 1970's and 80's, this accelerator (then, the SP$\overline{\text{P}}$S) accepted both protons and antiprotons produced further upstream, and delivered $p\bar{p}$ collisions to the UA1 and UA2 experiments. These collaborations discovered the W and Z bosons in 1983, for which the 1984 Nobel prize in physics was partially awarded. Today, the SPS energy is the minimum beam energy which may be accepted by the Large Hadron Collider, and it provides the final link in the injection chain.

2.2 The Large Hadron Collider

A main focus of CERN's contemporary efforts is the operation and management of the **Large Hadron Collider** (LHC), a proton synchrotron with a 26.7 km circumference installed in the tunnel which formerly housed the Large Electron Positron (LEP) collider, at a 100 m depth beneath the border of Switzerland and France. The LHC accepts proton beams injected from the SPS at 450 GeV and accelerates them up to a design energy of 7 TeV before colliding them in several **interaction points** (IPs) around the LHC ring where the proton beams may be made to intersect. The collisions produced by the LHC may reach a design center-of-mass energy of $\sqrt{s} = 14$ TeV, the highest ever produced in the laboratory setting[3].

Once the 450 GeV SPS beam has been passed to the LHC, it is *locked* to the machine by operating the RF cavities, which drive acceleration of the proton beam, at a harmonic of the accelerator's revolution frequency:

$$f_{\text{RF}} = h f_{\text{rev}}. \tag{2.1}$$

This locking procedure creates stable regions of phase space known as *buckets* along the beam's orbit within which the beam is not accelerated, and around which the protons become localised. As the magnetic field of the synchrotron is increased, the stable momentum of these buckets increases, accelerating the orbiting protons. At the LHC, the revolution frequency is $f_{\text{rev}} = 11.2455$ kHz and the RF operates at 40 MHz, resulting in $h = 35640$ stable regions along the beam. One in ten of these buckets may be filled with protons, and typically several more are left empty in order to provide sufficient time for the LHC beam to be extracted safely during a fill. At design specifications, 2808 of the available 3564 buckets are filled,

[3]Fermilab's Tevatron, the collider which previously held this record, produced proton-antiproton ($p\bar{p}$) collisions at centre-of-mass energies between 1.6 and 1.96 TeV, from 1985 until its shutdown in 2011.

grouped into larger structures called *trains*. Within a train, the spacing between bunches may be as little as 25 ns (corresponding to the nominal 40 MHz RF frequency), while the spacing between trains may be several bunch crossings long in order to accommodate the LHC beam dump kicker rise time (3 μs) and specifications of the SPS injection kicker (a rise time of 0.95 μs, and a flat-top which may not exceed 7.86 μs). The nominal filling scheme for 25 ns operations is written as

$$\begin{aligned} 3564 \;=\; & [2 \times (72b + 8e) + 30e] \\ & + [3 \times (72b + 8e) + 30e] + [4 \times (72b + 8e) + 30e] \\ & + 3 \times \{2 \times [3 \times (72b + 8e) + 30e] + [4 \times (72b + 8e) + 31e]\} \\ & + 80e \end{aligned} \tag{2.2}$$

where b denotes a bucket filled with protons, and e an empty bucket. The filling scheme may also be written in terms of the SPS cycles and the number of proton batches provided from the pre-injection chain during each cycle:

234 334 334 334

where the first cycle contains a pair of 72-bunch batches from the pre-injectors, the second cycle a triplet, and so-on for the 12 injection cycles required to fill each LHC beam. Each SPS injection cycle takes 21.6 seconds to perform, and so the nominal filling time of each LHC beam is about 4 minutes. Following the filling procedure, the magnetic field of the LHC may be increased adiabatically by *ramping* from an initial field strength of 0.54 T up to a maximum design level of 8.33 T, a process which may take more than half an hour.

Once the LHC has been filled and the proton beams have been brought to the desired energy, collisions are provided to each of the LHC experiments. The rate of collisions, R, is typically expressed as the product of the instantaneous luminosity $\mathcal{L}_{\text{inst}}$, and the total proton-proton interaction cross-section:

$$R = \mathcal{L}_{\text{inst}} \sigma_{pp}. \tag{2.3}$$

The value of σ_{pp} is chiefly measured at the LHC by the TOTEM collaboration and ATLAS ALFA detectors, using roman pot detectors and particle telescopes installed close to the beamline within the 250 metres up- and down-stream of the CMS and ATLAS detectors [21, 22]. The instantaneous luminosity, dependent on conditions of the machine and colliding beams, may be expressed as

$$\mathcal{L}_{\text{inst.}} = f_{\text{rev}} \frac{n_c N_1 N_2}{4\pi\sigma^2} \tag{2.4}$$

where n_c is the number of colliding bunch pairs in the current LHC filling scheme, N_1 and N_2 are the number of protons per-bunch in the respective LHC beams, and σ is the transverse size of the beam at the IP. A careful measurement of N and σ allows for a calibration of the instantaneous luminosity of the LHC to be performed at a particular IP of the LHC ring. The former quantity is constantly measured using dedicated beam-current monitors installed along the LHC ring. The latter may be factorised further into the x and y components of the beam size:

$$\sigma^2 = \sigma_x \sigma_y. \tag{2.5}$$

These quantities are measured for each IP using so-called van der Meer scans [23] performed periodically during routine LHC operations, during which the LHC beams are slightly displaced and scanned across one-another first horizontally, then vertically, in order to accurately determine the beam size. Scans of this nature were first developed by Simon van der Meer and applied to measure the size of the proton beams at CERN's Intersecting Storage Rings during the 1970s [24].

Due to the bunched structure of the beams, there is a large probability for more than a single proton to interact during each bunch-crossing. This phenomena is known as **pile-up**, and represents a major experimental challenge for physics performed at the LHC. These additional pp interactions generally produce many additional, low-p_{T} jets associated with different vertices than the objects which originate from the hard scattering event. These extra jets produce spurious noise in detector systems which must be accounted for. The average number of interactions per bunch-crossing, μ, is shown in figure 2.2 for 2011/2012 and for 2015/2016 operations.

In total, seven experiments reside within the experimental caverns positioned around the LHC ring. A pair of specialized detectors, ALICE and LHCb, study the collisions of relativistic lead ions produced by the LHC, and the physics of b-hadrons and CP violation. The MOEDAL experiment, situated within the same cavern as LHCb, searches for magnetic monopoles and highly-ionising stable massive particles. A pair of general-purpose detectors, ATLAS and CMS, were designed to both discover the Standard Model Higgs boson and to maintain broad sensitivity to other TeV-scale new physics scenarios, such as natural supersymmetry or universal extra dimensions. Two additional experiments are installed alongside these general-purpose detectors: the LHCf experiment shares an interaction point with the ATLAS detector, and studies the origin of ultra-high-energy cosmic rays by examining cascades of particles produced at extremely low angles by LHC collisions. The previously-mentioned TOTEM experiment [25], installed on either side of the CMS interaction point, provides precise measurements of the total proton-proton interaction cross section (σ_{pp} in equation 2.3).

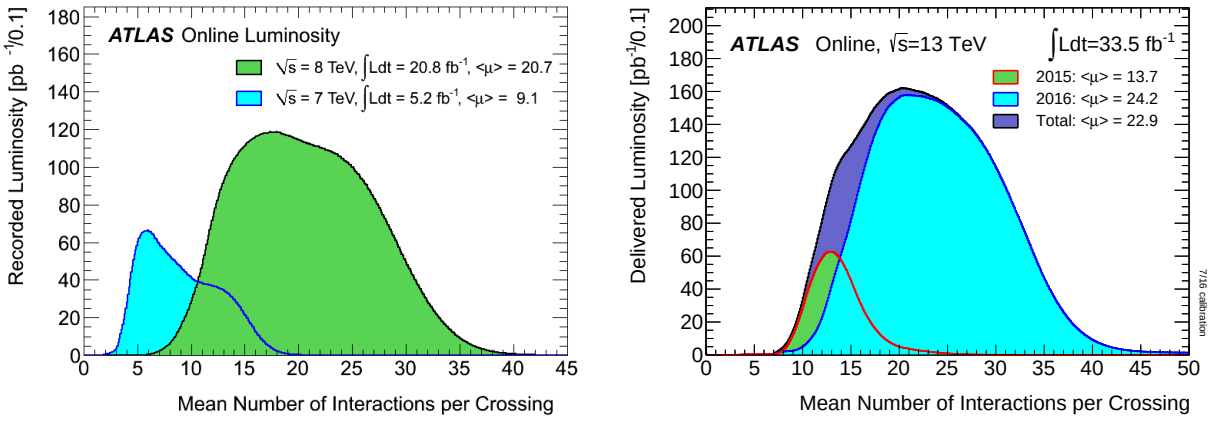

(a) (b)

Figure 2.2: The average number of interactions per bunch crossing, μ, during (a) run I and (b) run II operations. During 2012, the average μ value was 20.7, which has increased to 24.2 during 2016. The maximum observed μ has also increased, from > 40 during run I to > 50 during run II.

2.3 The ATLAS Detector

Located at interaction point 1 of the LHC, near CERN's Meyrin campus, the ATLAS[4] detector boasts a length of 44 metres and a 25 metre diameter, and is the largest experiment installed at the facility. ATLAS was designed with a cylindrical geometry centered around the LHC interaction point[5]. Despite its titanic scale, the detector is surprisingly light: weighing a mere 7000 tonnes, ATLAS would float in water, due mostly to the low density of the air-cored system of superconducting toroidal magnets after which the detector and collaboration are named. A computer-generated illustration of the ATLAS detector is shown in figure 2.3, which includes a pair of average-sized physicists for scale.

The collaboration was founded with two primary goals: to discover the Higgs boson

[4] **A** **T**oroidal **L**HC **A**pparatu**s**

[5] ATLAS uses a right-handed coordinate system, where the positive $\hat{x}$ direction is oriented towards the middle of the ring. The positive $\hat{y}$ direction is 0.704° off-vertical due to the slanted cavern floor. The positive $\hat{z}$ axis points counter-clockwise along the LHC beamline. The pseudorapidity is defined as

$$\eta = -\ln\left(\tan\frac{\theta}{2}\right), \tag{2.6}$$

where θ is the polar angle measured from the positive $\hat{z}$ axis. η is an often-used observable in collider physics because differences in η are invariant under longitudinal Lorentz transformations. The azimuthal angle ϕ is measured about the $\hat{z}$ axis.

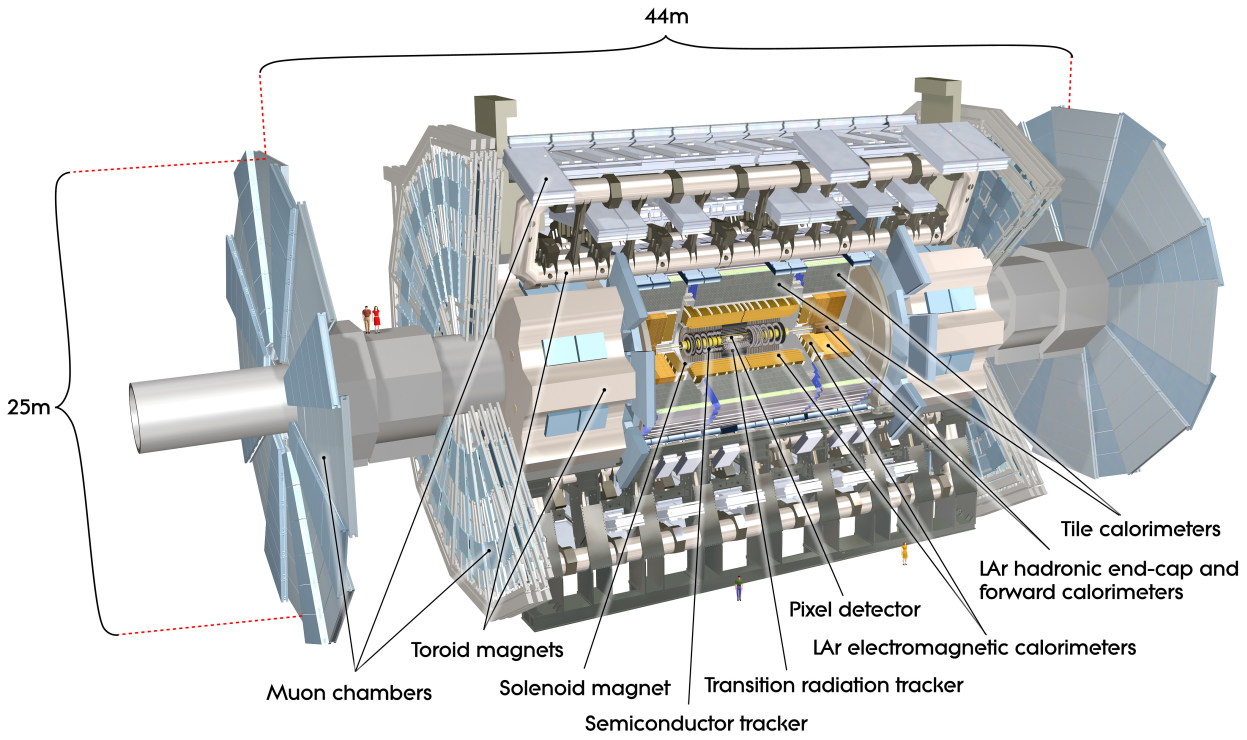

Figure 2.3: A computer-generated, cut-away view of the ATLAS detector. From [26].

predicted by the Standard Model, and to search for any signs of new physics which may present themselves within the LHC data.

The ATLAS detector is typically discussed in terms of three primary subsystems, nested within one another. The inner detector (ID) [27, 28] (figure 2.4), immersed within a 2 T magnetic field, provides precision measurements of the transverse momentum of charged particle tracks, and charged particle identification. The calorimeters [29] (figure 2.5), consisting of both liquid argon and scintillating tile components, provide containment and energy measurements of electromagnetic and hadronic showers. The muon spectrometer (MS) [30] (figure 2.9), mounted within the toroidal magnet system, allows for precise reconstruction of high-p_{T} muons produced in LHC collisions. The ATLAS trigger uses a multi-level approach with both hardware- and software-based triggers to collect data for analysis at a manageable rate of 1 kHz, greatly reduced from the nominal LHC collision frequency of 40 MHz. These data are processed and stored using a complex distributed computing infrastructure known as the Worldwide LHC Computing Grid (WLCG).

2.3.1 Inner Detector

The pixel detector, the closest detector system to the LHC interaction point, recieved the largest hardware change in the ATLAS detector between runs 1 and 2. Initially, three concentric cylinders of pixel detector modules at $\langle R \rangle = 50.5$ mm, 88.5 mm and 122.5 mm formed the barrel pixel detector, while three disk-shaped arrays upstream and downstream at $\langle Z \rangle = 495$ mm, 580 mm and 650 mm provide further angular coverage. The modules from which these detectors are composed are 250 μm thick silicon sensors which each contain 47232 individual pixels measuring either 50 μm $\times$ 400 μm (most pixels – more than 90% – are this smaller size) or 50 μm $\times$ 600 μm. These silicon sensors are bump-bonded directly to the pixel detector front-end electronics, which read out 2880 channels per module. The barrel pixel layers respectively contain 286, 494 and 676 modules, while the six (three per side) disks contain 48 apiece. Altogether, the original ATLAS pixel detector possessed 80.4 million readout elements.

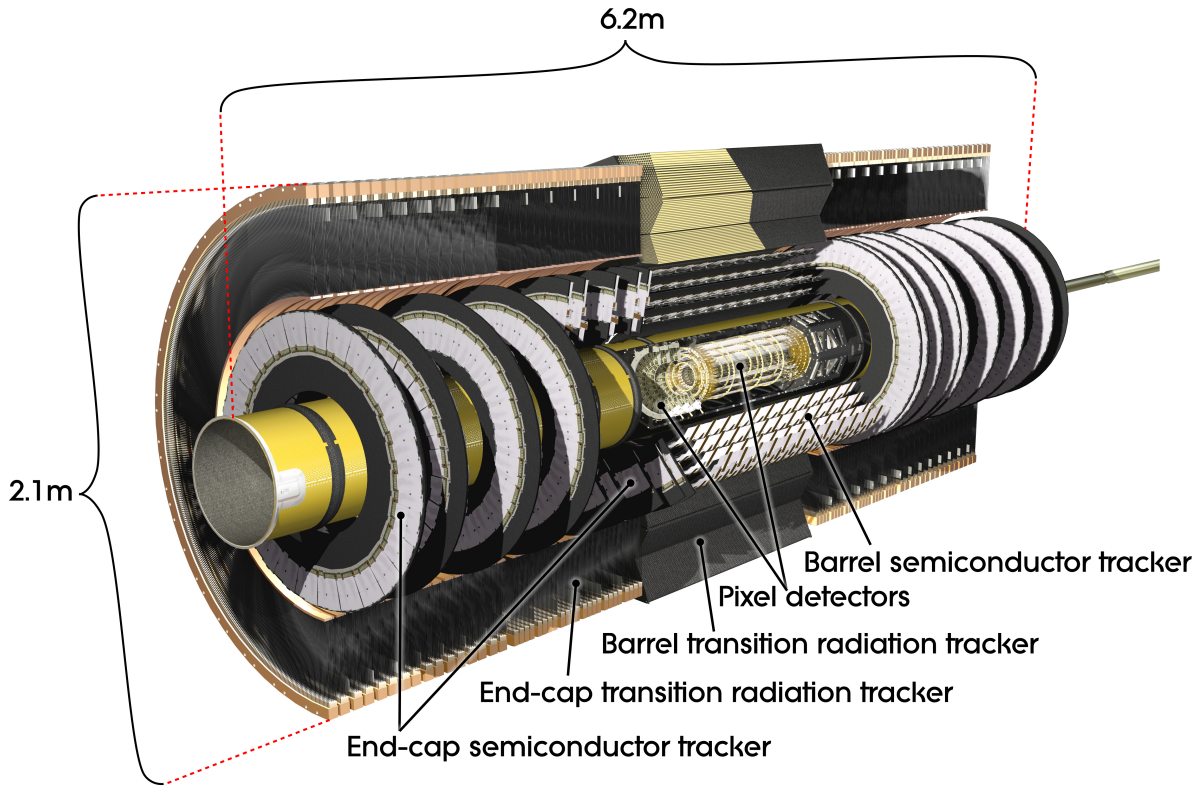

Figure 2.4: A computer-generated, cut-away view of the ATLAS inner detector systems. From [26].

During the long shutdown period between LHC run 1 and 2, an additional pixel layer was

installed between the original b-layer and the beam. This new module, called the Insertible b-Layer (IBL), was designed to improve tracking efficiency, flavour tagging performance and primary vertex finding in the busy LHC environment. The IBL is located at $\langle R \rangle = 33$ mm, so close to the beam that it was also necessary to install a new, smaller beampipe within the ATLAS sector (the new beampipe has an inner radius of 25 mm, reduced from the original 29 mm design). Fourteen carbon-fibre staves support the new pixel modules which compose the IBL, which are mounted directly on the outer beampipe wall.

Two types of silicon pixel modules are used in the IBL. The planar module sensors are 200 μm thick and have a granularity of $\Delta\phi \times \Delta z = 50$ μm $\times$ 250 μm, representing an improvement on the original ATLAS silicon pixel sensor design. New 3D silicon sensors are 230 μm thick with the same pixel granularity as the planar modules, fabricated with a combination of MEMS and VLSI technologies to etch and dope electrodes across the entire width of the sensor. This design allows for more active area along the edges of the sensor, and renders them more radiation-hard than traditional planar pixel sensors. Twelve double-sensor planar modules populate the centre of each IBL staff, while an additional four single-sensor 3D modules are mounted on each end, resulting in 12 million additional pixel detector readout channels.

Surrounding the pixel detector, the semiconductor tracker (SCT) is composed of silicon microstrip detectors measuring 6.36 cm $\times$ 6.40 cm, each with 768 readout strips measuring 80 μm $\times$ 6.40 cm. A set of concentric barrel cylinders and end-cap disks were also chosen for the SCT design, which utilise a slightly different module design. In the barrel SCT, each module is built from four silicon detectors, wire-bonded lengthwise in pairs to form 12.8 cm strips which are then glued back-to-back at a stereo angle of 40 mrads. Four barrel layers are built with 384, 480, 576 and 672 individual modules at respective radii of $\langle R \rangle = 299$ mm, 371 mm, 443 mm and 514 mm. The endcap SCT detectors use similarly constructed modules but whose microstrips are tapered, with one set of strips aligned radially, to build 9 disks at distances of $\langle Z \rangle$ between 853.8 mm – 2720.2 mm from the centre of the ATLAS detector.

When a charged particle passes through a silicon sensor, whether it is a pixel or a microstrip, the readout procedure is similar: electrons and holes are created at the surface of a silicon sensor or a group of sensors (referred to as a pixel 'cluster'), causing an increase in the current reported by the sensor for some time interval. The electrode voltage is compared to a threshold value; should the sensor's time-over-threshold be sufficiently long, the sensor or cluster is registered as a *hit* – the individual points through which tracks are extrapolated by fitting algorithms.

The transition radiation tracker (TRT) is a straw-tube detector which surrounds the SCT and extends to the barrel solenoid. Each 4 mm polyimide tube contains a thin aluminium

layer serving as the detector's cathode, and is reinforced structurally with carbon fibre using a unique winding procedure and machine developed at CERN. These tubes are strung with anode wires made from a gold-tungsten alloy. Particles traversing the barrel TRT, within $|Z| < 720$ mm, produce an average of 36 additional tracking hits in the region of $\langle R \rangle =$ 560-1070 mm. The endcap TRT modules provide coverage to $|\eta| < 2.5$, where 18 individual wheels of radially-aligned tubes provide forward tracking hits between $|Z| = 830$ mm to 3400 mm. The TRT is nominally filled with a 70% Xe / 27% CO_2 / 3% O_2 gas mixture. The space between TRT straws is filled with a matrix of polymer fibres and foils, whose dielectric properties differ from those of the tubes themselves. This dielectirc boundary causes charged particles crossing the interface to emit some amount of low-energy radiation, proportional to their mass. As the momentum of a charged particle may be measured using other detector subsystems, the TRT provides a standalone method of discriminating electrons from heavier charged hadrons based on their transition-radiation signatures [27, 28].

Tracks are composed from the set of ID hits using global χ^2 minimisation and Kalman fitting algorithms [31]. Initially, tracks are seeded using hits from the pixel detector and the first layer of the SCT, then extended through the remaining SCT layers and TRT. After track-finding, a vertexing algorithm is used to determine the primary vertices originating from any energetic proton-proton interactions in the event. Secondary vertex reconstruction and the identification of displaced vertices originating from the decays of long-lived particles, such as b-hadrons, are also crucial components of vertexing in which performance translates directly to improved heavy-flavour identification potential. The b-tagging algorithms of ATLAS (section 3.2.5), used extensively when selecting events with top quarks (chapter 4) or in searches for new physics decaying to third-generation particles (chapters 5-7) rely heavily on this information.

2.3.2 Calorimeters

Calorimetry systems within ATLAS aim to fully contain and measure the energy of showers produced by electromagnetically- and hadronically-interacting particles as they interact with material within the detector. Showers produced by these different interactions have distinct characteristics within the non-compensating ATLAS calorimeters: electromagnetic cascades tend to be more compact than hadronic ones, and some portion of energy deposition by hadronic showers is invisible to the calorimetry, leading to a lower observed response for a showering hadronic particle than an electromagnetic one with the same initial energy.

ATLAS contains two primary calorimeter subsystems which target these different types of interactions: within the central region of the detector, the liquid argon (LAr) based electromagnetic calorimeter (section 2.3.2) is located just beyond the barrel solenoid and provides

coverage within $|\eta| < 1.475$, while the scintilating-tile hadronic calorimeter (section 2.3.2) resides between the solenoid return yoke and the muon spectrometer, covering $|\eta| < 0.8$ with its main barrel module and to $|\eta| < 1.7$ via the TileCal extended barrel.

Further forward, additional liquid argon systems extend calorimeter coverage as far as $|\eta| < 4.9$. The LAr electromagnetic and hadronic end-cap (EMEC and HEC) systems cover the region between $1.5 < |\eta| < 3.2$ using technologies similar to that of the electromagnetic barrel calorimeter. The remaining coverage is provided by the LAr forward calorimeter (FCal), specifically designed to withstand the high particle fluxes and increased radiation exposure near the beam.

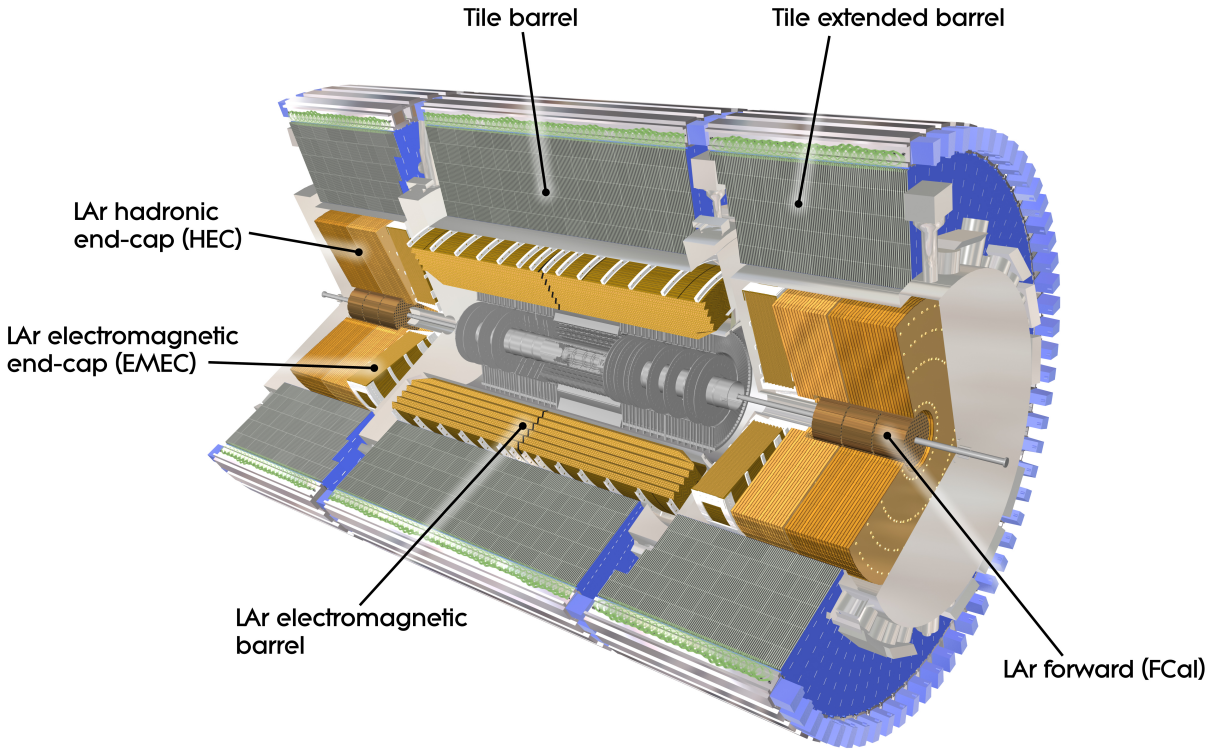

Figure 2.5: A computer-generated, cut-away views of the ATLAS calorimetry systems. From [26].

Liquid Argon Calorimeters

The liquid argon calorimeters form a crucial part of the ATLAS detector infrastructure, and provide energy measurements which are used in the reconstruction of electrons, photons,

jets and missing transverse momentum. The LAr calorimeter's trigger systems allow for the rapid selection of events containing these objects during data taking.

The electromagnetic barrel calorimeter (EMB) is positioned directly behind the barrel solenoid magnet within the main barrel cryostat, extending outward from $\langle R \rangle = 1.15$ m to 2.25 m and covering pseudorapidities within $|\eta| < 1.475$. The EMB is a sampling calorimeter with layered lead and stainless steel absorbers which zig-zag in the azimuthal direction, resulting in the accordion-like geometry shown in figure 2.6. This design results in > 22 radiation lengths (X_0[6]) of electromagnetic coverage for the barrel region. The distinct geometry allows for continuous azimuthal coverage, and has become a hallmark of the ATLAS LAr calorimeter. Each half-barrel of the EMB is constructed from 1024 absorber layers, each spaced 4.2 mm apart. Between each pair of absorbers, an electrode made of alternating copper and polyimide sheets is positioned using 2.1 mm honeycomb spacers. These electrodes are held at 2000 V, collecting electrons knocked free of argon atoms by the passage of electromagnetically-interacting particles with an average drift time of 450 ns.

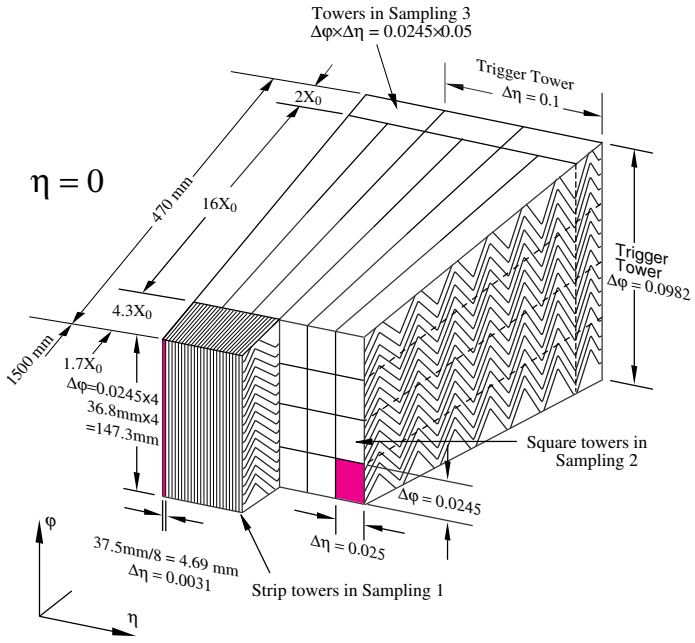

Figure 2.6: The 'accordion' design of the electromagnetic barrel calorimeter, shown in detail. The relative depth of each layer, in electromagnetic radiation lengths X_0, is provided. From [29].

[6]One radiation length is defined as the distance through which an electron must pass in order to lose $1 - 1/e \sim 63\%$ of its initial energy by bremtrahlung.

The LAr calorimeters are divided into transverse segments called **cells**, which are arranged in several layers of varying granularity to form its basic readout channels. The EMB is divided into three layers within $|\eta| < 1.35$ and two layers between $1.35 < |\eta| < 1.475$. Inside of $|\eta| < 1.52$, there exists enough material in front of the calorimeter[7] to necessitate an initial presampling layer with a strip-like geometry and η–ϕ granularity of $\Delta\eta \times \Delta\phi = 0.025 \times \pi/32$, resulting in 7808 individual cells. The first EMB layer segments more finely in η within $|\eta| < 1.4$ and in ϕ within $1.4 < |\eta| < 1.475$, resulting in respective granularities within these regions of $\Delta\eta \times \Delta\phi = 0.025/8 \times \pi/32$ and $\Delta\eta \times \Delta\phi = 0.025 \times \pi/128$. The second layer granularity becomes finer within $1.4 < |\eta| < 1.475$ and coarser within $1.4 < |\eta| < 1.475$, respectively $\Delta\eta \times \Delta\phi = 0.025/8 \times \pi/32$ and $0.075 \times \pi/128$. The third layer covers $|\eta| < 1.35$ with uniformly coarser granularity, $\Delta\eta \times \Delta\phi = 0.050 \times \pi/128$. Altogether, the EMB has 109 568 individual readout channels.

Beyond $|\eta| = 1.375$, the LAr endcap calorimetry systems provide forward coverage until $|\eta| = 3.2$. The electromagnetic and hadronic endcap calorimeters (EMEC and HEC), along with the forward calorimeters (FCal), are situated in separate cryostats than the EMB, which are located as closely as possible to the barrel module. The endcap calorimeters form three large wheels, where the innermost is the EMEC, and the latter pair compose the HEC. The EMEC follows the EMB's choice of an accordion-shaped lead and stainless steel absorber, though in the endcap design these absorbers are arranged radially[8]. The EMEC was manufactured as two separate rings which respectively cover $1.4 < |\eta| < 2.5$ and $2.5 < |\eta| < 3.2$. Readout is again segmented into three layers of cells, whose granularity changes between $\Delta\eta \times \Delta\phi = 0.025 \times \pi/128$ to $\Delta\eta \times \Delta\phi = 0.10 \times \pi/32$ depending on the particular location within the detector. The EMEC contains 63 744 separate readout channels. The EMEC provides $> 24X_0$ worth of coverage to the endcap region.

Hadronic calorimetry in the endcap region is provided by the hadronic endcap calorimeters (HEC), each located immediately behind the corresponding EMEC wheel within the endcap cryostats. The HEC utilises flat copper absorbers with active liquid argon gaps, each divided into a pair of modules (HEC1 and HEC2) whose design slightly varies. The thickness of the absorber in HEC1 is 25 mm, doubling to 50 mm in HEC2 – although, the first and last plates in each calorimeter are of half-thickness. HEC1 contains 25 absorber plates, while HEC2 contains just 17. The outer radius of all plates is a constant 2030 mm from $z = 0$, though the inner radius varies from 372 mm for the first 9 plates of HEC to 475 mm for

[7]On average, there exists between 2-4 X_0 of material before the calorimeter within this region, and in some places there is over 6 X_0 in front of the end-cap calorimeters. The barrel solenoid accounts for just 0.63 X_0 – the ATLAS magnet is very thin!

[8]This arrangement is occasionally referred to as a 'Spanish fan' geometry.

all remaining plates, in order to provide space for the forward calorimeter. The LAr gap between plates is 8.5 mm, which is instrumented with a three-layer kapton-copper electrode with 2.1 mm spacing, held at 2000 V relative to the absorbers to provide the signal readout. All gaps – between the three electrode layers, and between the outer layers of the electrode and the absorber – are maintained using 1.816 mm aramid resin honeycomb spacer mats. Cell granularity of the HEC is constant within given ranges of $|\eta|$: for the region between $1.5 < |\eta| < 2.5$, $\Delta\eta \times \Delta\phi = 0.1 \times \pi/32$ and doubles within the region from $2.5 < |\eta| < 3.2$, to $\Delta\eta \times \Delta\phi = 0.2 \times \pi/16$. This results in a total of 5 632 readout channels.

The LAr forward calorimeter modules are the closest calorimeter subsystems to the beamline, and extend hermetic coverage to $|\eta| < 4.9$. The FCal consists of three modules, one to measure electromagnetic activity, and two to measure hadronic activity. The electromagnetic module consists of a hexagonal matrix of holes 7.5 mm apart, bored through a 450 mm thick series of copper disks with 455 mm outer radii, within which copper rods are positioned such that a 0.269 mm gap is left surrounding the rod, which is filled with liquid argon. These rods are the electrodes of the FCal, and are held at 250 V relative to the bulk copper matrix. Electrons released from argon atoms ionised by passing particles collect on the electrode's surface, producing a triangular current pulse of 1 ns rise time and 61 ns linearly decreasing tail due to the electron drift time within the gap. The hadronic modules follow a similar design, but utilise tungsten absorber layers built from slugs positioned around copper tubes filled with tungsten electrode rods, with solid copper end-cap plates. The spacing between electrodes in the inner (outer) hadronic FCal module is 8.18 mm (9.00 mm), the size of the LAr gap is 0.376 mm (0.508 mm) and its electrodes are held at 375 V (500 V), producing average electron drift times of 81.7 ns (113.1 ns). Electrodes are ganged in fours across each module to be read out as a single cell, resulting in a total of 3 524 FCal readout channels. At the outside end of the forward calorimeter, a passive brass plug with the same dimensions as one of the FCal modules has been placed in order to provide additional shielding for the muon systems.

The containment capability of the hadronic calorimetry systems is measured in hadronic interaction lengths, λ_0, defined as the average distance a hadronically-interacting particle may travel through material without experiencing a nuclear interaction. The total number of λ_0's varies as a function of pseudorapidity, as shown in figure 2.7. There typically exists $> 10\lambda_0$ across the entire range of calorimeter coverage, roughly $2\lambda_0$ of which is provided by the EM calorimetry systems. The complete containment of hadronic showers is desirable in order to prevent the punch-through of particles into the muon systems (sec. 3.2.2). In this context, even non-instrumented structures provide additional shielding, and are represented as the unlabelled cyan band on this histogram.

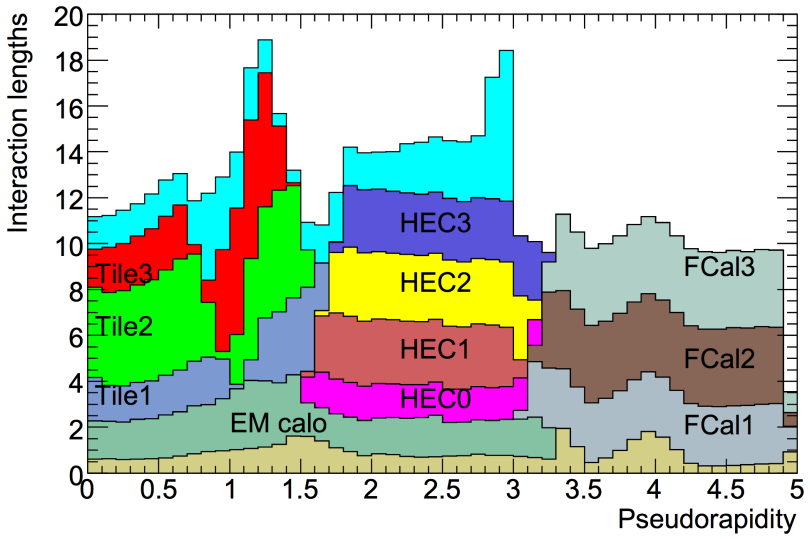

Figure 2.7: The number of hadronic interaction lengths provided by ATLAS calorimetry systems, as a function of pseudorapidity (η). There typically exists $> 10\lambda_0$ across the entire range of calorimeter coverage, roughly $2\lambda_0$ of which is provided by the EM calorimetry systems.Even non-instrumented structures provide effective shielding for the muon systems, represented as the unlabelled cyan band in this histogram. From [32].

LAr Calorimeter Readout

When charged particles traverse the active layers of the liquid argon calorimeter, they ionise the liquid argon atoms as they pass. Electrons freed by ionisation drift to the electrodes, held at high voltage, and induce a current: the signal of the LAr calorimeter. The current induced by ionisation electrons is the triangular pulse shown in figure 2.8, which has a rapid initial rise time and a slow, linear descent over the next ∼450 ns (the typical drift time of the 2 mm ionisation gaps found throughout the calorimeter). The signals from electrodes in a single readout cell are passed into the Front End (FE) system of electronics, mounted on the ATLAS cryostats. Within this FE system, 1524 front end boards (FEBs) handle the shaping, digitisation and readout of the calorimeter systems for up to 128 channels apiece.

Bipolar shaping is applied to the analogue signal from a calorimeter cell by the corresponding FEB, resulting in the typical pulse shape shown in figure 2.8, which is sampled at the 40 MHz rate of the LHC clock. In the event that an L1 trigger accept signal is received, a predetermined number of sampled values are read out and digitised by a 12-bit analogue-to-digital converter on the FEB. Five samplings were read out per-FEB until run 2, when the number of samplings was reduced to four. This reduction was necessary due to the increased instantaneous luminosity of the LHC, in order to maintain the capability of processing the LHC collision data in real time.

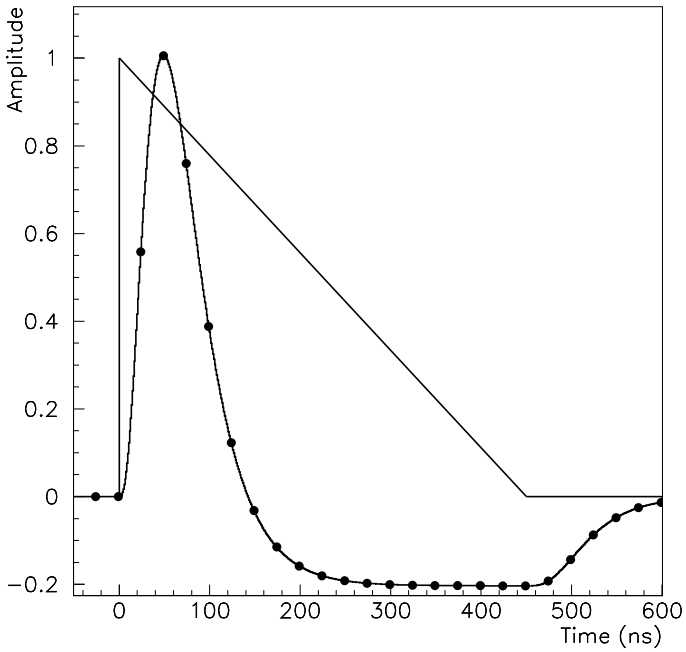

Figure 2.8: The ATLAS liquid argon signal pulse, before and after shaping. Sampling points are shown as dots on the shaped curve. From [29].

Samples from the ADC are passed along an optical link to the Read-Out Driver (ROD), located within the Back-End (BE) electronics housed in the USA15 cavern adjacent to the ATLAS experimental hall. The RODs reconstruct the cell energy for the selected bunch crossing from the digitised signal, s_j, by reconstructing the pulse amplitude in ADC counts,

$$A = \Sigma_{j=1}^{N_{\text{samples}}} a_j (s_j - p), \tag{2.7}$$

where p, the **pedestal**, is the expected mean value of s_j (known from regular calibration of the FE electronics during operations), and a_j are the **Optimal Filtering Coefficients** (OFCs) for the particular cell, which filter the impact of electronic noise and out-of-time pileup in the energy reconstruction process. The OFCs are computed based on the expected pulse shape, properties of the known electronic noise in the cell, and the expected instantaneous luminosity of the operational period. Once the signal amplitude A is reconstructed for a given gain, G, the cell energy may be computed as

$$E_{\text{cell}} = F_{\mu\text{A}\rightarrow\text{ MeV}} \times F_{\text{DAC}\rightarrow\mu\text{A}} \times \frac{1}{\frac{M_{\text{phys}}}{M_{\text{cali}}}} \times G \times A. \tag{2.8}$$

Several factors are required to make this conversion: $F_{\text{DAC}\rightarrow\mu\text{A}}$ and $F_{\mu\text{A}\rightarrow\text{ MeV}}$ perform the

conversion from ADC counts to μA and then to MeV. $M_{\text{phys}}/M_{\text{cali}}$ is the ratio of the response to a calibration or ionisation signal of the same current in the same cell. These factors are determined from regular calibration runs of the LAr systems, simulation and from ATLAS test beam data [33].

Once the cell energy has been obtained, the reconstruction of topological clusters and higher-level objects such as jets and missing transverse energy from this fundamental measurement becomes possible (chapter 3.7).

The LAr calorimetry systems, like all ATLAS subsystems, are continually monitored in order to ensure the quality of data taken during operations is suitable for use in physics analysis. More details about monitoring and mitigating sporadic bursts of coherent electronic noise which occur within the calorimeter read-out electronics during physics datataking are given in appendix A.

Scintilating Tile Calorimeters

In the ATLAS barrel (within $|\eta| < 0.8$) and extended barrel (between $0.8 < \eta < 1.7$) at radii between $\langle R \rangle = 2280$ mm to 4250 mm, additional calorimetry to measure the energy of hadronically interacting particles passing beyond the LAr systems is provided by the Scintilating Tile Calorimeter (TileCal) [34]. This calorimeter utilises $4 - 5$ mm steel tiles as absorber layers, interspersed with 3 mm active tiles built from doped polystyrene scintillators attached to photomultiplier tubes via wavelength-shifting optical fibres. The TileCal is constructed in radially projective towers which provide three layers of measurement. These three layers of TileCal were also designed with varying granularity: the cell size of the inner two layers is $\Delta\eta \times \Delta\phi = 0.1 \times 0.1$, while that of the third layer is $\Delta\eta \times \Delta\phi = 0.1 \times 0.2$. Altogether, the TileCal is between 8-12 interaction lengths (λ) thick, depending on the pseudorapidity of the incident particle.

With time, radiation damage to the scintillating tiles will cause the amount of scintillating light they produce to decrease. Damage may also affect the readout fibres, decreasing their transmissibility and therefore decreasing the amount of light reaching a PMT. In order to provide a constant reference against which to calibrate the TileCal signal, regular cesium scans are performed during which a known radioactive sample is moved throughout the calorimeter systems using a hydraulic system. A laser system is also used to calibrate the PMTs directly, by simulating producing pulses of light similar to calorimeter signals.

2.3.3 Muon Spectrometer

One barrel and two endcap superconducting toroidal magnets, each built from eight flat coils, form the backbone of the ATLAS Muon Spectrometer (MS) [30]. These air-core magnets are situated behind the calorimetry systems, and their magnetic field bends muons in the plane perpendicular to their previous deflection by the barrel solenoid magnet. This allows for an independent measurement of their momentum, which may be combined with information from the ID in order to precisely reconstruct the muons produced by LHC collisions. These magnets provide a non-uniform field, with an average bending power of 3 Tm in the barrel and 6 Tm in the endcap regions.

The MS is built from several concentric cylinders of muon chambers which provide coverage within the barrel region of ATLAS. These chambers are staged in three layers around the barrel toroid, at $\langle R \rangle = 5000$ mm, 7500 mm and 10000 mm. Coverage for the endcap regions is obtained using additional disks of chambers located both up- and down-stream of ATLAS at $\langle Z \rangle = 7000$ mm, 10000 mm, 14000 mm and between 21000-23000 mm. Four types of chambers are used, two for the precise measurement of muon hits as they traverse ATLAS, and two which provide dedicated, rapid event selection. An annotated, cut-away illustration of the MS is provided in figure 2.9, which highlights the locations of the different types of muon chambers.

Monitored drift tubes (MDTs) are installed throughout the barrel and endcap regions of the MS, forming the backbone of the precision measurement subsystems. MDTs provide precision coverage for all muon stations within $|\eta| < 2.7$; except for the inner ring of the first end-cap station; altogether, over 5500 m^2 of coverage is provided by this type of chamber. MDTs follow a traditional drift-tube detector design, consisting of a 30 mm-thick aluminium tubes whose length may vary between 700 mm to 6300 mm, and a cathode wire made of a tungsten-rhenium alloy (97%–3%). These wires are held at a potential of 3270 V, and the tubes are filled with a non-flammable gas mixture of 91% Ar, 4% N_2, and 5% CH_4. Tubes are arranged in groups of three or four layers on each side of so-called 'spacer frame' structures, which provide structural support for both the MDTs and the monitored optical alignment systems for the 1194 individual chambers after which these chambers are named.

Cathode strip chambers (CSCs) provide precision coverage within the innermost ring of the forward end-cap region (between $2.0 < |\eta| < 2.7$). In this region, due to increased particle fluxes, a precision detector with a smaller drift time than the MDTs (≤ 30 ns, an order of magnitude less than the MDT drift time of 480 ns) and lower neutron sensitivity is required. The CSCs are multiwire proportional chambers strung by anode wires made from the same W-Re alloy as those in the MDT, with a 2.54 mm pitch. Copper cathode strips are etched into FR4 laminate boards with a pitch of 5.08 mm which surround the wire chamber

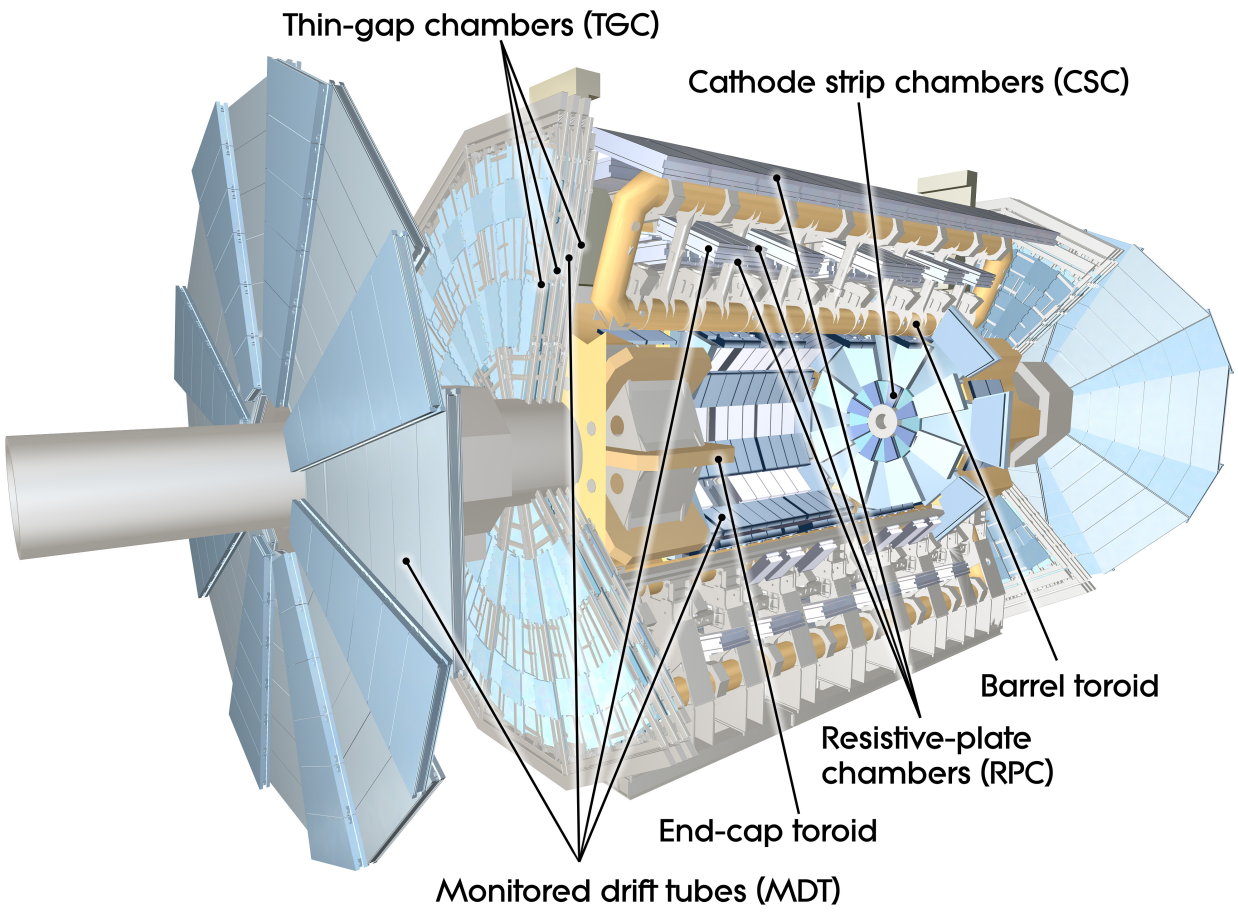

Figure 2.9: A computer-generated, cut-away view of the ATLAS muon spectrometer systems. From [26].

using a photolithographic process. Strips in one plane are oriented perpendicular to the chamber wires in order to measure the precision coordinate (r in the barrel region or z in the endcap region), while those in the other plane are oriented parallel to the chamber wires in order to provide a measurement of the transverse coordinate. The CSCs are operated with a gas mixture of 30% Ar, 50% CO_2 and 20% CF, which has a low sensitivity to neutrons due to the absence of hydrogen.

The ATLAS muon trigger systems must be capable of selecting events with high-energy muons with a time resolution capable of identifying the LHC bunch-crossing at the nominal rate of 25 ns. As the MDT drift times are too long for this purpose, a separate set of chambers dedicated for triggering has also been installed within the MS. In the barrel region, three layers of resistive plate chambers (RPCs) have been positioned on either side of the central layer, and either above or below the outermost layer of MDTs. These chambers contain a 2 mm gas gap contained by resistive bakelite sheet electrodes, coated with thin layers of

graphite and linseed oil. This gap is filled with a mixture of 97% $C_2H_2F_4$ and 3% C_4H_10, and a uniform electric field of 4.5 kV/mm is produced by the electrodes. A digital signal is produced by the passage of charged particles, which initiate an avalanche of electrons between the chamber electrodes. This charge accumulation induces a signal on the cathode, which may be read out for use in trigger decisions with a time resolution of 1.5 ns.

Three layers of thin gap chambers (TGCs), located behind the first endcap wheel, provide dedicated triggering in the forward region. These multi-wire proportional chambers are characterised by the distance between the cathode planes and gold-plated tungsten anode wires (1.4 mm), which is smaller than the distance between adjacent wires (1.8 mm). TGCs are operated with a gas mixture of 55% CO_2 and 45% n-pentane, which rapidly quenches showers formed by the passage of charged particles. Due to the gas mixture and high electric field – the annode wires are held at a high voltage of 3.1 kV – the TGCs operate in a saturated proportional mode, which provides a high signal to noise ratio and a typical signal rise time under 5 ns.

2.3.4 Trigger and Data Acquisition

The nominal LHC collision rate of 40 MHz is too large to write all events to disk for analysis. The ATLAS trigger system [35] reduces the incoming stream of data to a more manageable final rate of about 1 kHz which are written to disk for complete reconstruction and later, analysis use. An inital hardware-based trigger selects events based on information from the calorimeters, muon systems and forward detectors, then passes interesting events to the *high-level trigger* (HLT), a software-based system which make the final decision.

The Level 1 (L1) trigger is implemented on hardware located in the USA15 cavern. Dedicated systems in the calorimeters (L1Calo), muon spectrometer (L1Muon) and forward region (the minimum bias trigger scintilators, MBTS) are monitored for signs of interesting events: the presence of energetic jets, electrons or photons may be inferred from a localised deposition of energy in the EM calorimeter, or energetic muons from a signal from the MS TGCs. A new L1 system, the L1Topo trigger, was implemented between runs 1 and 2 and selects events based on their global quantities, such as the invariant mass of dijet pairs, or the event's missing transverse momentum. When the L1 systems are triggered, a decision to accept or reject the event must be made within 2.5 μs. Should the L1 system decide to accept, information from all detector subsystems is read out to buffers in the data acquisition systems designed to hold event information until the subsequent HLT decision has occurred. The L1 trigger is designed to reduce the event rate from the nominal collision frequency to about 75 kHz.

The HLT is a software-based trigger system which receives candidate events from the L1

triggers and performs a limited reconstruction in order to refine the selections which may be applied before deciding whether to write events to tape, further reducing the overall rate of event retention. This reconstruction is initially localised within *regions of interest* (RoIs): windows in $\eta - \phi$ space surrounding the trigger object(s). The average HLT processing time is 235 ms, which sets the maximum rate of event acceptance for physics use to be 1 kHz due to the available computing resources at the HLT farm. Following an event's acceptance by the HLT, the event is transferred from the data acquisition system buffer to the ATLAS Tier-0 computing centre at CERN, in order to be processed for eventual long-term storage.

2.4 ATLAS Distributed Computing

Events are output by the ATLAS high-level trigger in 'RAW' format, at a rate of 1 kHz. In this format, an individual event is roughly 1-2 MB, depending on detector activity and pileup conditions, and is not easily analysible. RAW events are transferred from the HLT farm to the CERN Tier-0 computing center, which is responsible for both maintaining a local copy of all RAW data, for processing this data into more manageable derived formats (such as the intermediate Event Summary Data (ESD) format, or the Analysis Object Data (AOD) format from which most users begin their own analysis) and for distributing to Tier-1 facilities the raw data for archival and the derived formats for downstream accessibility. Distribution of data from the Tier-0 to Tier-1 sites is facilitated by 10-100 Gbps optical links[9]. Typical users are unable to access the Tier-0 facility directly, and so the distribution of data to remote sites is necessary before any downstream analysis may begin in earnest.

Ten Tier-1 sites across the globe[10] collectively provide long-term curation of a second copy of all RAW format data collected by the ATLAS detector, in order to protect against possible site failures which would result in the loss of one stored copy of the RAW files. Tier-1 facilities may also produce derived formats of RAW data located on-site, and store the most recently derived versions of these datasets on-disk in order to provide replicas to Tier-2 computing sites for end-user analysis. Tier-1 sites also receive simulated samples produced by Tier-2 sites for archival storage.

Tier-2 computing sites are more numerous, located at 155 collaborating universities[11] and laboratories. These sites are the primary location for the storage of the current versions of derived datasets in AOD format which are easily accessible to users for physics analysis,

[9] From the LHC Optical Private Network (LHCOPN).

[10] One of the ten Tier-1 computing sites is located at TRIUMF, Canada's national lab for nuclear and particle physics located in Vancouver, BC.

[11] Including the University of Victoria!

and perform most analysis jobs requested by ATLAS physicists. Typically, two copies of the current AOD samples are stored across the Tier-2 sites, in order to ensure reliable and prompt access for all members of the collaboration. The Tier-2 sites are also the computing resources used for the generation of most of the Monte Carlo simulation required by ATLAS analyses, a copy of which is subsequently transferred upstream to Tier-1 centres for long-term storage and possible reprocessing.

The typical ATLAS physicist interacts most directly with their local Tier-3 site, usually a local computing cluster which provides both adequate storage and batch computing resources to perform some routine tasks, and a connection to the Worldwide LHC Computing Grid (WLCG, or simply *the grid*). For large-scale data analysis, users may submit tasks to the grid which are performed remotely at Tier-2 sites, and then subsequently transfer the (highly reduced) output back to their Tier-3 site. Final statistical analyses and visualisation is often performed directly on the user's local Tier-3, using common analysis tools such as ROOT which are provided and maintained by the site.

2.4.1 Benchmarking Cloud Resources for ATLAS

The multi-tiered computing infrastructure was laid out by ATLAS in 2005, and remains to this day the model followed by the collaboration. There are advantages to having multiple copies of the ATLAS samples spread amongst data centres which span the globe: there is a quick turnaround time for user jobs, the availability of samples to users is high, and users based at institutes which are geographically distant from CERN may still be nearby to a Tier-1 or Tier-2 computing centre, leading to fewer potential network problems. There are drawbacks with this system as well: all ATLAS computing sites must be staffed by personnel familiar with the ATLAS computing infrastructure, hardware must be maintained and software environments may not always be identical from site to site. The amount of available resources is finite, which may become problematic to users attempting to perform data analysis at times when there is a large demand on Tier-2 resources due to a reprocessing or simulation campaign.

One potential solution to many of these challenges could be the application of cloud computing techniques within the ATLAS computing infrastructure. If analysis and simulation jobs are performed in virtual machines (VMs) rather than on the 'bare metal,' the task of managing software at grid sites becomes a matter of ensuring that the proper images are used when creating the VMs. Archival of analysis environments also becomes easier – an image could be produced by analyzers containing their exact code and software configuration, which would be used in the future to instantiate a VM within which their results could be exactly replicated.

Some cloud-based services exist already on academic computing resources, and are used by the ATLAS collaboration. In recent years, many commercial options for *Infrastructure-as-a-Service* (IaaS) clouds, such as Amazon's EC2, Microsoft's Azure and the Google Compute Engine (GCE) have come to market. These commercial clouds allow for dynamic resource acquisition by ATLAS during periods of high demand: when the bare-metal grid becomes saturated, it could become desirable to purchase more resources from a commercial provider and instantiate more ATLAS environments, in order to reduce job queues back to sustainable levels more quickly.

In an effort to survey the suitability of different commercially available IaaS cloud providers for typical ATLAS purposes, a benchmarking study [1] was performed which examined several different available VM types from both the Amazon EC2 and GCE platforms, and also compared the performance of these clouds to academically available cloud resources managed by Compute Canada (CC-East and CC-West, located in Sherbrooke, QC and Victoria, BC). The VM types considered possess different numbers of vCPUs and virtual memory, leading to different performance. This study represented the first effort to perform detailed resource accounting of the commercially available cloud platforms by ATLAS, a process which allows for well-informed desicisons to be made when procuring commercial cloud resources. Not all clouds are created equal: VM performance of similar types (with the same number of vCPUs, or the same amount of virtual memory) from different providers may vary significantly depending on the underlying hardware hosting the cloud.

Comparisons were made using the HEP-SPEC06 (HS06) benchmark[12] [37]. HS06 performs a variety of integer and floating-point calculations in succession which are representative of typical HEP workflows[13]. The hardware used to provide the cloud infrastructure differs between providers, and also between availability regions of the same cloud provider. For the GCE, the processor type is homogeneous across the zone hosting the virtual machines: these studies examined only the `us-central1-b` zone, and so all benchmarks in the GCE environment have occurred on physical Intel Xeon E5-2670 (Sandy Bridge) 2.6 GHz processors. Amazon EC2 resources are hosted on either physical Intel Xeon E5-2670 v2 (Ivy Bridge) or Xeon E5-2670 (Sandy Bridge) 2.6 GHz processors. Compute Canada resources are deployed on Intel Xeon E5-2650v2 CPUs (Sandy Bridge) 2.6 GHz physical processors.

Virtual machines were instantiated using a standard Cern μVM image [38][39]. Following their creation, HS06 was transferred to the VM via condor jobs which also managed running the benchmark on the VM and reporting the results before terminating the instance.

[12]HS06 is based on the SPEC CPU 2006 benchmarking suite [36].

[13]The `all_cpp` benchmark set was selected, consisting of the individual SPEC benchmarks `444.namd`, `447.dealII`, `450.soplex`, `453.povray`, `471.omnetpp`, `473.astar` and `483.xalancbmk`.

HS06 was configured to utilize all available cores of a VM's vCPU. Fifty benchmarks were completed for each VM type, in order to examine possible variance in performance caused by the subscription of a single physical node to multiple virtual machines.

The resources which have been evaluated are summarized in table 2.1, along with their HS06 scores, which are also illustrated in figure 2.10. Standard virtual machine types in each cloud were selected. Additionally, a specialised set of GCE virtual machine types with a higher ratio of virtual CPUs to memory (the `highcpu` series) were benchmarked. Variation in the amount of memory available to the virtual machines did not significantly impact the final scores: provided that the machines possess enough memory to run their required tasks, additional memory is not beneficial and often incurs additional costs. Commercial clouds provide an overall more consistent benchmarking result than the CC resources, whose scores for VMs of the same type varied widely: a possible reflection of differences in the resource subscription policies of the various resources. CC clouds reported higher average HS06 scores than IaaS clouds, a reflection of the newer hardware used to host these resources. The GCE `n1-standard-1` instance type was found to be unexpectedly performant, likely a result of being assigned one physical hyperthreaded processor per single-cored virtual machine. This performance is contingent on low occupancy of the cloud, as multiple VMs could be assigned to share same hyperthreaded core under heavier use of the cloud, and so it would not necessarily indicate that the use of this VM type would be a more efficient use of ATLAS resources.

Cloud	CPU Model	VM Type	vCPUs	RAM (GB)	HS06 Score
Google (GCE)	E5-2670	`n1-standard-1`	1	3.75	18.97 ± 0.66
		`n1-standard-2`	2	7.50	23.66 ± 0.62
		`n1-standard-4`	4	15.0	45.46 ± 0.55
		`n1-standard-8`	8	30.0	85.46 ± 0.92
		`n1-standard-16`	16	60.0	149.40 ± 2.95
		`n1-highcpu-2`	2	1.80	23.48 ± 0.54
		`n1-highcpu-4`	4	3.60	44.88 ± 1.12
		`n1-highcpu-8`	8	7.20	86.31 ± 1.09
		`n1-highcpu-16`	16	14.4	149.61 ± 3.05
Amazon EC2	E5-2670, E5-2670v2	`m3.large`	2	7.5	25.20 ± 0.20
		`m3.xlarge`	4	15	48.26 ± 1.12
		`m3.2xlarge`	8	30	91.31 ± 2.46
Victoria	E5-2650v2	`c2.low`	2	7.68	38.64 ± 15.81
		`c4.low`	4	15.36	61.83 ± 5.98
		`c8.low`	8	30.72	119.71 ± 24.39
		`c16.low`	16	61.44	147.24 ± 17.64
Sherbrooke	E5-2650v2	`c4.low`	4	15.36	70.61 ± 5.92
		`c8.low`	8	30.72	129.03 ± 13.78
		`c16.low`	16	61.44	198.24 ± 27.55

Table 2.1: A summary of the properties of the different clouds and their virtual machines surveyed in the cloud benchmarking study, along with the measured HS06 scores of these virtual machine types (larger scores correspond with better performance). Google Compute Engine (GCE) resources were located in the us-central1-b region. From [1].

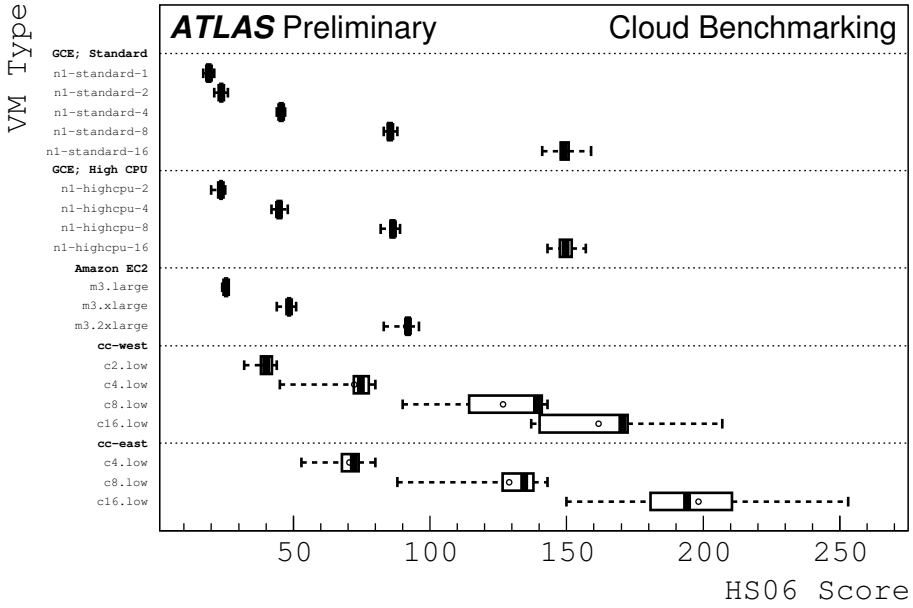

(a)

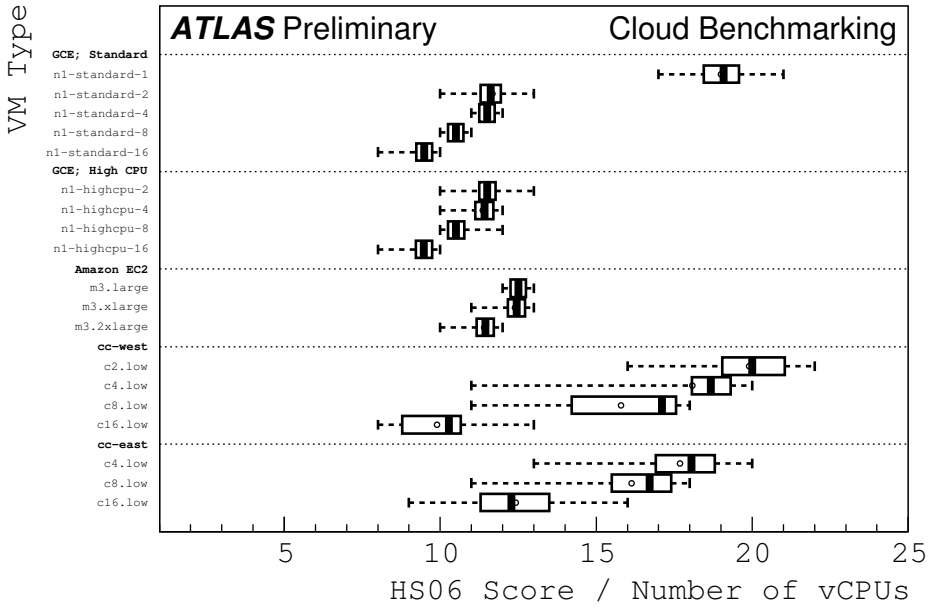

(b)

Figure 2.10: HS06 benchmarking scores for GCE, Amazon EC2 and Compute Canada cloud resources. Below, the scores have been normalised by the number of available vCPUs within the particular virtual machine type. From [1].

3

Reconstructing Objects in ATLAS

After an event is accepted by the ATLAS trigger systems for use in data analysis, the objects which were present must be reconstructed from the low-level detector signals. These objects are complex, and often built from lower-level inputs such as the energy read out from calorimeter cells, or tracking hits in the inner detector or muon spectrometer. Familiar particles such as electrons, photons and muons are reconstructed directly from these low-level quantities, each leaving a distinct signature within the detector. The collimated sprays of many particles known as jets are built from intermediate objects called **topological clusters** (or, topoclusters), which attempt to collect the energy from one particle's interaction with the calorimeter in a form which is robust against noise in the calorimeter cells due to electronic effects and pile-up. The missing transverse momentum ($E_{\mathrm{T}}^{\mathrm{miss}}$) of an event, from which information regarding the presence of particles in an event which did not interact with ATLAS may be inferred, is the most complex object which is reconstructed. All other reconstructed objects are utilised in the calculation of the $E_{\mathrm{T}}^{\mathrm{miss}}$, along with soft energy depositions in the calorimeter or isolated tracks in the inner detector which are not otherwise associated with a reconstructed object.

Once reconstructed, the measured properties of these objects may be calibrated to a particular energy scale. Due to the non-compensating nature of the ATLAS calorimetry systems, electromagnetic and hadronic particles interacting with the calorimeter with the same initial energy will deposit energy differently within the detector. A classification algorithm provides weights for topological clusters based on their compatibility with electromagnetic or hadronic decays, and sophisticated calibration chains are subsequently applied to the objects which these clusters are used to construct.

3.1 Topological Clusters

Energy from a shower of particles (section 1.4) is deposited in the calorimeter and sampled by its active layers (section 2.3.2). This signal is read out from each calorimeter cell and is used to construct the primary cell-level observable used for the formation of topological clusters, the **cell significance**, defined as the ratio of the cell signal to the expected average noise in the cell, both measured at the electromagnetic (EM) scale:

$$\varsigma_{\text{cell}} = \frac{E^{\text{EM}}_{\text{cell}}}{\sigma^{\text{EM}}_{\text{noise,cell}}}. \tag{3.1}$$

In early ATLAS operations (prior to 2011 data-taking), the cell noise was almost entirely electronic noise. After the switch to 50 ns operations in 2011, noise originating from signal fluctuations due to pile-up became the dominant source, and the noise thresholds were redefined:

$$\sigma^{\text{EM}}_{\text{noise,cell}} = \begin{cases} \sigma^{\text{electronic}}_{\text{noise}} & \text{(2010 \& prior)}, \\ \sqrt{\left(\sigma^{\text{electronic}}_{\text{noise}}\right)^2 + \left(\sigma^{\text{pile-up}}_{\text{noise}}\right)^2} & \text{(2011 \& after)}. \end{cases} \tag{3.2}$$

The particular value of $\sigma^{\text{pile-up}}_{\text{noise}}$ is selected for a given period of operations based on the expected amount of average in-time pile-up, μ during this run period: it was chosen based on values of $\langle\mu\rangle = 8$ for 2011 and $\mu = 30$ for 2012 and 2015 running.

Topological clusters are built from calorimeter cells through the application of a growing-volume algorithm whose performance is governed by three parameters, S, N and P:

$$\left|E^{\text{EM}}_{\text{cell}}\right| > S\sigma^{\text{EM}}_{\text{noise,cell}} \Rightarrow |\varsigma_{\text{cell}}| > S \qquad \text{(primary seed threshold)}; \tag{3.3}$$

$$\left|E^{\text{EM}}_{\text{cell}}\right| > N\sigma^{\text{EM}}_{\text{noise,cell}} \Rightarrow |\varsigma_{\text{cell}}| > N \qquad \text{(threshold for growth control)}; \tag{3.4}$$

$$\left|E^{\text{EM}}_{\text{cell}}\right| > P\sigma^{\text{EM}}_{\text{noise,cell}} \Rightarrow |\varsigma_{\text{cell}}| > P \qquad \text{(principal cell filter)}. \tag{3.5}$$

These parameters were chosen to be $S = 4$, $N = 2$ and $P = 0$, following studies of ATLAS test-beam data [40]. This choice of parameters dictates that topological clusters must be seeded by cells whose significance is $> 4\sigma$ above the noise threshold of the given run period. Cells which directly neighbour[1] the seed cells and whose own significance is $> 2\sigma$ above threshold are added to the cluster seeds to form so-called protoclusters, and then all cells

[1] *n.b.* topological clustering proceeds in three dimensions within the ATLAS calorimetry systems, and topoclusters may span the boundaries between different subsystems.

which directly neighbour these protoclusters are also included (due to the choice of $P = 0$). Protoclusters with a negative energy are discarded at this step, as they are associated with out-of-time pileup rather than signal energy depositions.

Following the formation of these protoclusters, a splitting algorithm is applied to divide protoclusters with two or more local energy maxima into distinct topological clusters, improving the spatial resolution of the energy density within the calorimeter. Protoclusters which contain cells with a local maximum above 500 MeV and at least 4 direct neighbours are split into multiple clusters, until each cluster possesses only one maximum above the energy requirement. The splitting algorithm is applied twice, each time restricting the cells which are examined for local maxima to different regions of the calorimeter. First, layers two and three of the electromagnetic barrel and endcap, and module 0 of the forward calorimeter are considered. Subsequently, the splitting is refined by including layer 1 of the electromagnetic barrel and endcap, and all layers of the hadronic calorimetry systems. During the splitting steps, some cells may find themselves associated to more than one protocluster. The energy from such cells is split between the energies of the two most energetic clusters they are associated with $E^{EM}_{\text{clus},1}$ and $E^{EM}_{\text{clus},2}$, based on the distances from the cell to the centre-of-energy of the two clusters, d_1 and d_2. This splitting is expressed in terms of a pair of weights, $w^{\text{geo}}_{\text{cell},1}$ and $w^{\text{geo}}_{\text{cell},2}$:

$$w^{\text{geo}}_{\text{cell},1} = \frac{E^{\text{EM}}_{\text{clus},1}}{E^{\text{EM}}_{\text{clus},1} + rE^{\text{EM}}_{\text{clus},2}}\,, \tag{3.6}$$

$$w^{\text{geo}}_{\text{cell},2} = 1 - w^{\text{geo}}_{\text{cell},1}\,, \tag{3.7}$$

$$r = \exp(d_1 - d_2)\,. \tag{3.8}$$

3.1.1 Local Hadronic Calibration

Due to the non-compensating nature of the ATLAS calorimetry systems (see section 2.3.2), the signal from hadronically-interacting particles will always be smaller than the signal from electromagnetically-interacting particles depositing the same amount of energy. In order to account for this difference, as well as to partially correct for signal losses due to dead material and due to inefficiencies in the clustering procedure itself, a local cluster calibration is derived. A weight is assigned to each cluster based on its probability to have been generated by an electromagnetic shower ($\mathcal{P}^{\text{EM}}_{\text{clus}}$), which is evaluated as a function of the cluster energy, pseudorapidity, the cluster depth, λ_{clus}, and the average cell energy of the cluster, ρ_{clus}. Signal weights are applied to each individual cell in each topological cluster from both

the electromagnetic and hadronic calibration procedures in a linear combination, which is weighted by this probability:

$$w_{\text{cell}}^{\text{cal}} = \mathcal{P}_{\text{clus}}^{\text{EM}} \cdot w_{\text{cell}}^{\text{em-cal}} + (1 - \mathcal{P}_{\text{clus}}^{\text{EM}}) \cdot w_{\text{cell}}^{\text{had-cal}}. \tag{3.9}$$

The electromagnetic signal weight, $w_{\text{cell}}^{\text{em-cal}}$, is equal to unity by construction due to the default reconstruction of all cell signals at the EM scale. The hadronic signal weight, $w_{\text{cell}}^{\text{had-cal}}$, is defined as the ratio of the energy deposited per-cell to the same energy reconstructed at the EM scale:

$$w_{\text{cell}}^{\text{had-cal}} = \frac{E_{\text{cell}}^{\text{dep}}}{E_{\text{EM}}^{\text{cell}}}. \tag{3.10}$$

$E_{\text{cell}}^{\text{dep}}$ is determined from lookup tables built using simulations of single charged pions with a wide range of energies entering the calorimeter from all directions. Following the calibration of cluster energy scales, additional cluster-based corrections are derived separately for electromagnetic- and hadronic-scale clusters which account for out-of-cluster losses and the effects of dead material in certain locations within the detector.

After these corrections have been applied, clusters are ready for use in jet reconstruction. Jets in ATLAS are typically built from clusters calibrated at the electromagnetic scale (**EMtopo** jets, the default jet collection during run II, for which $\mathcal{P}_{\text{clus}}^{\text{EM}}$ is taken to be unity), or at the local cluster weighted scale (**LCTopo** jets, the default jet collection during run I, for which $0 < \mathcal{P}_{\text{clus}}^{\text{EM}} < 1$).

3.2 Jets

3.2.1 Jet Reconstruction Algorithms

The shower of particles within a detector originating from the fragmentation and hadronization of a strongly-interacting particle produced in a collision is reconstructed experimentally as an object called a **jet**. Jets are not uniquely defined: their properties depend strongly on the particular algorithm which is used for their reconstruction. The consideration of which jet reconstruction algorithms were the most desirable started early – Ellis *et al.* had already provided a sort of physicist's 'wish list' by 1989 [41], which mentioned that algorithms should be:

I Simple to implement in an experimental analysis;

II Simple to implement in the theoretical calculation;

III Defined at any order of perturbation theory;

IV Yields finite cross section at any order of perturbation theory;

V Yields a cross section that is insensitive to hadronisation.

Prior to run II of the Tevatron at Fermilab, cone jet algorithms were the most widely-used option at hadron colliders [42]. These algorithms attempt to collect all particles within some stable cone defined by a distance metric R, which determines the cross-sectional area of the cone in the $\eta - \phi$ plane. The placement of each cone is determined by choosing a 'seed' object – for example, the input object with the largest transverse momentum – and grouping all particles within R of the seed together as a jet, then repeating for the next yet-unassociated seed.

Although they were used with some success by early collaborations, cone algorithms inherently fail to be **infrared-** and **collinear-safe** (IRC-safe): the procedure of fitting a cone to a shower of particles is inherently sensitive to the effects of soft and collinear radiation present in their vacinity. Should the original seed particle have experienced a collinear splitting, a different seed may have been selected, resulting in a different final set of jets. Similarly, the presence of a soft (infrared) emission could bias the positioning of cones, or cause two previously resolved jets to be reconstructed as one. These problems are illustrated in figure 3.1. The ultimate pitfall of IRC-sensitive jet algorithms is that they produce results which are overly-sensitive to the details of the parton-shower modeling which is used, resulting in undesirable behaviour – the requirements of the 'wish list' were not satisfied.

These theoretical problems were mitigated through the application of successive combination algorithms previously used at e^+e^- colliders in the hadron collider environment [44], defined generally in five steps:

1. For each protojet i, define the **jet-beam distance measure**,

$$d_i = p_{T,i}^{2P}. \tag{3.11}$$

   For each pair of protojets $(i, j; i \neq j)$, define the **jet-jet distance**,

$$d_{i,j} = \min(p_{Ti}^{2P}, p_{Tj}^{2P}) \left(\frac{(y_i - y_j)^2 + (\phi_i - \phi_j)^2}{R^2} \right). \tag{3.12}$$

   The parameter R in this metric is a parameter selected by the user, which controls the ultimate size of the jets which are produced. Typical choices of R within the ATLAS collaboration are $R = 0.4$, and for large jets, $R = 1.0$. P is another parameter designated by the user, discussed further below.

2. Minimize the set $\{d_i \cup d_{i,j}\}$, and label the minimum element d_{min}.

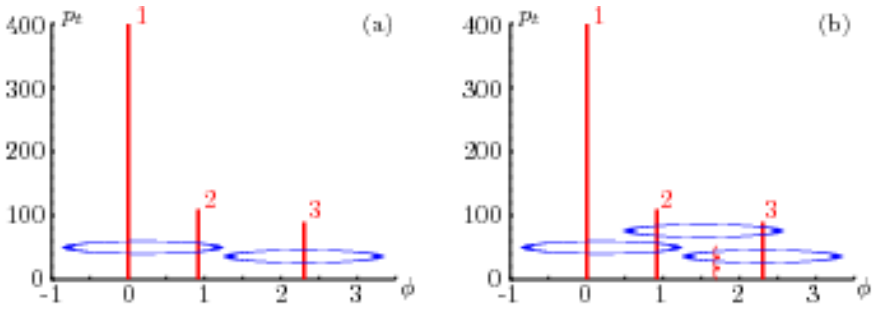

(a) An infrared (IR) instability.

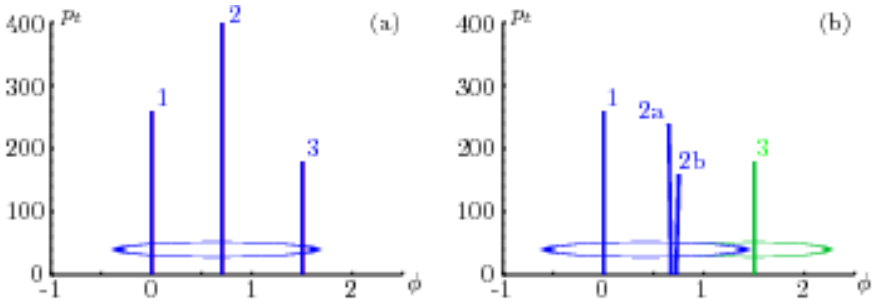

(b) A collinear (CL) instability.

Figure 3.1: Examples of (a) infrared and (b) collinear instabilities which may occur during jet reconstruction. From [43].

3. Should d_{min} belong to $\{d_{i,j}\}$, then the ith and jth protojets are removed from the set and replaced with a new protojet k, obtained by adding the four-vectors of i and j.

4. Should d_{min} belong to $\{d_i\}$, then the ith protojet is not merge-able. This element is removed from the set of d_i's and defined as a jet.

5. Repeat until there are no remaining protojets.

The choice of parameter P in this algorithm ultimately corresponds with the choice of the particular algorithm which is applied:

- $P = 1$, the k_t **algorithm** [45][44], first clusters the softest protojets into harder ones, then clusters the hardest protojets together. This ordering encodes information about the hard structure of a jet within its clustering history. The prioritised treatment of soft radiation can make the k_t algorithm more susceptible to energy density fluctuations from pile-up and detector noise. The k_t algorithm may also produce jets of irregular shape: protojets for which $\Delta R > R$ may be included within a jet. This non-uniformity may cause complications with calibration when these jets lie near poorly-instrumented regions of the detector.

- $P = 0$, the **Cambridge-Aachen (C/A) algorithm** [46], omits the p_{T} dependence of clustering and results in a purely angularly-ordered clustering of the jet constituents. This angular hierarchy leads to desirable properties, due to the angular ordering of QCD: the jet axis of a C/A jet tends to be aligned with its hardest subjet, and the subjets within a C/A jet are typically separated by distance scales on the order of the subjet radius.

- $P = -1$, the **anti-k_t algorithm** [47], is the default choice for jet reconstruction in ATLAS. The inverted power law dependence on protojet p_{T} clusters the hardest constituents together first, and then soft radiation in the final stages of recombination. The resulting jets possess a characteristic circular shape and tend to be centered around hard deposits of radiation. The size of an anti-k_t jet will not exceed R. Should two energetic deposits lie within a distance of $2R$ from each other, two jets will be reconstructed; the harder jet with a circular shape, and the softer with the lens-shaped overlapping portion excised.

Jets in ATLAS are reconstructed with the FastJet software package [48][49].

3.2.2 Jet Energy Calibration

A jet constructed from topological clusters does not necessarily reflect the properties of the showering particle created by the hard-scattering event. This difference is quantified by the **jet response**:

$$\mathcal{R}_{p_{\mathrm{T}}} = \frac{p_{\mathrm{T}}^{\mathrm{calo}}}{p_{\mathrm{T}}^{\mathrm{truth}}}, \tag{3.13}$$

$$\mathcal{R}_{m} = \frac{m^{\mathrm{calo}}}{m^{\mathrm{truth}}}. \tag{3.14}$$

The response of a jet may differ from unity due to many factors, such as the non-compensating nature of the ATLAS calorimetry systems (see section 2.3.2), dead material in the detector, losses due to energy deposits which fall 'out-of-cone' or 'punch through' the calorimetry systems into the muon spectrometer, or due to an inefficiency to occasionally form clusters due to the noise thresholds used in the topoclustering algorithm ($4\sigma/2\sigma/0\sigma$). In order to correct for these effects, a dedicated jet energy scale (JES) calibration procedure must be developed at the level of the jets themselves, and applied before these objects may be used in physics analysis. Calibrations may be developed for jets built using input topological clusters which have been calibrated at the electromagnetic scale (EM) or using a local cell weighting (LCW) algorithm. These jet calibrations are derived in several steps using both simulation-based

and *in-situ* analyses, outlined below.

Large-R jets, most commonly used to reconstruct electroweak bosons and top quarks with large transverse momenta, are treated separately. The calibration procedure for these jets is outlined in section 3.3.2.

Jet origin correction

Topoclusters are first reconstructed using the geometric center of the ATLAS detector (*i.e.* (0,0,0)) as their origin, and so jets also initially point to this location when reconstructed. The origin of topological clusters is shifted to point at the primary vertex of the event once it has been identified, and so the jet four-vector must be recalculated based on these origin-corrected inputs.

Pile-up correction

The effect of energy depositions from pile-up is parameterised event-by-event based on the average energy density, calculated using k_t $R = 0.4$ jets as

$$\rho = \text{median} \left(\frac{p_{\text{T}}}{A} \right) \Big|_{\text{k}_\text{t} \text{ jets}} \tag{3.15}$$

where the area, A, is the area associated with a given jet. The p_{T} of each jet is corrected by a factor of $\rho \times A$, then a simulation-based residual correction is made based on the jet's pseudorapidity, the number of reconstructed primary vertices in the event (N_{PV}), and the average number of interactions per bunch crossing in a given luminosity block (μ):

$$p_T^{\text{corr}} = p_T^{\text{EM}} - \rho \times A - \alpha(\eta) \times (N_{\text{PV}} - 1) - \beta(\eta) \times \mu. \tag{3.16}$$

Correcting the jet p_{T} as a function of both N_{PV} and $\langle \mu \rangle$, separately parameterised by jet pseudorapidity using the functions $\alpha(\eta)$ and $\beta(\eta)$ (determined from simulation), allows for both in-time and out-of-time pileup effects to be accounted for.

Four sources of uncertainty are propagated to the total JES systematic uncertainty from the pile-up correction, accounting for uncertainties on N_{PV} and $\langle \mu \rangle$, an uncertainty due to the p_{T}-dependence of the correction itself, and an uncertainty introduced by choosing a particular definition of the energy density ρ used to compute the correction.

Energy and η calibration

Following the pile-up and jet origin corrections, a simulation-based energy calibration is derived which restores the central value of the reconstructed jet energy response to that at truth level. The jet energy response of isolated, reconstructed jets in QCD dijet simulation is binned according to the energy of their associated truth-jet, E_T^{true}, and pseudorapidity in the detector frame, η_{det}. In each bin, the average jet response, $\langle \mathcal{R}_{\text{EM}}^{\text{jet}} \rangle$ is determined by fitting the jet response with a gaussian distribution and taking the peak value of the fit. The average jet energy, $\langle E_{EM}^{\text{jet}} \rangle$, is taken as the mean of the $E_{\text{truth}}^{\text{jet}}$ binned jet energy distribution.

A fit is performed in each η_{det} bin to $\langle \mathcal{R}_{\text{EM}}^{\text{jet}} \rangle$ and $\langle E_{EM}^{\text{jet}} \rangle$ as a function of $E_{\text{truth}}^{\text{jet}}$, resulting in the jet calibration function $\mathcal{F}_{\text{calib}}(E_{EM}^{\text{jet}})$. This function is inverted numerically to obtain the correction leading to the calibrated jet energy at EM+JES scale:

$$E_{\text{EM+JES}}^{\text{jet}} = \frac{E_{\text{EM}}^{\text{jet}}}{\mathcal{F}(E_{\text{EM}}^{\text{jet}})} = c_{\text{EM+JES}} \times E_{\text{EM}}^{\text{jet}}. \tag{3.17}$$

Global sequential calibration (GSC)

The jet energy scale is still dependent on longitudinal and transverse features of jets following the EM+JES calibration, as well as on the fraction of a jet's energy which is not absorbed by the ATLAS calorimetry systems ("punch-through"). An additional calibration which accounts for these dependencies is applied afterwards, known as the Global Sequential Calibration (GSC). Sequential corrections are made by fitting the five following uncorrelated variables:

1. The fraction of the total energy of the jet deposited in the first layer of the scintillating tile calorimeter.
2. The fraction of the total energy of the jet deposited in the third layer of the electromagnetic calorimeter.
3. The number of tracks associated with the jet.
4. The width of the tracks associated with the jet, weighted by their p_{T}.
5. The number of muon segments associated with the jet.

Residual *in-situ* calibration

To account for differences between data and simulation effecting the p_{T} measurement of jets, an *in-situ* calibration is derived. First, probe jets in forward regions ($0.8 < |\eta_{\text{det}}| < 4.5$) are

calibrated using those in the central detector ($|\eta| < 0.8$) as reference objects, within dijet events. The asymmetry between these jets, defined as

$$\mathcal{A} = \frac{p_{\mathrm{T}}^{\mathrm{probe}} - p_{\mathrm{T}}^{\mathrm{ref}}}{p_{\mathrm{T}}^{\mathrm{avg}}}, \tag{3.18}$$

characterises the relative calorimeter response of the probe region with respect to the reference region. A correction factor, c, is defined in terms of the average value of $\mathcal{A}$, computed in bins of $p_{\mathrm{T}}^{\mathrm{avg}}$ and η_{det} of the probe jet:

$$\frac{1}{c} = \langle \frac{p_{\mathrm{T}}^{\mathrm{probe}}}{p_{\mathrm{T}}^{\mathrm{ref}}} \rangle \sim \frac{2 + \langle \mathcal{A} \rangle}{2 - \langle \mathcal{A} \rangle}. \tag{3.19}$$

After the eta intercalibration has been performed, the pseudorapidity-dependence of reconstructed jet responses is mitigated. In order to correct for remaining differences between reconstructed jets in data and simulation, well-calibrated reference objects are balanced against isolated jets, and the double-ratio of the reference object p_{T} to that of the jet, in data and MC,

$$r = \langle \frac{p_{\mathrm{T}}^{\mathrm{jet}}}{p_{\mathrm{T}}^{\mathrm{ref}}} \rangle_{\mathrm{data}} \Big/ \langle \frac{p_{\mathrm{T}}^{\mathrm{jet}}}{p_{\mathrm{T}}^{\mathrm{ref}}} \rangle_{\mathrm{MC}} \tag{3.20}$$

is taken as a residual *in-situ* correction applied to jets in data. The particular reference object used to construct r depends on the p_{T} of the jet to be calibrated. For jets with p_{T} between 20 GeV and 250 GeV and $|\eta_{\mathrm{det}}| < 0.8$, a Z-jet balance is performed by selected leptonically-decaying Z bosons. Similarly, a photon-jet balance is performed for jets with p_{T} between 30 GeV and 800 GeV and $|\eta| < 0.8$. High-p_{T} jets are calibrated by creating a balance against recoiling systems of lower-p_{T} jets which have already been calibrated by other aspects of the *in-situ* calibration. This multi-jet balance technique provides calibrations for jets with p_{T} between 300 GeV and 1700 GeV.

A large number of the different uncertainties on the total JES are associated with the various *in-situ* techniques used: 23 systematic and 35 statistical sources of uncertainty are considered in the context of these different methods.

Single hadron response

A limitation of the *in-situ* calibration is its p_{T} reach, which is restricted by the availability of high-p_{T} photons and multi-jet systems against which to balance the jet. One method used to extend the jet energy scale calibration and its uncertainties to transverse momenta beyond that statistically accessible by the *in-situ* balance techniques is to treat jets as a collection of individual particles, the responses of which are known from ATLAS test beam

data. The uncertainty associated with the single-hadron response method is considerably larger than that of the direct balance analyses, and so this technique is only used to provide calibrations and uncertainties for jets with $p_{\mathrm{T}} > 1500$ GeV, where other approaches become unsuitable. It is propagated to the total JES uncertainty as a single additional source of systematic uncertainty.

3.2.3 The total JES Uncertainty

A total of 63 sources of systematic uncertainty arise from various aspects of the calibration procedure, which are combined with two additional sources of uncertainty which account for differences in the jet response due to flavour-dependent effects. The final uncertainty on the jet energy scale computed at 8 TeV during run I is shown in figure 3.2, and the total uncertainty as computed at 13 TeV during run II is shown in figure 3.3. In both figures, the total uncertainty is computed by adding the individual components in quadrature.

3.2.4 The Jet Energy Resolution Uncertainty

In order to account for differences in the width of the jet energy response between data and simulation, an *in-situ* measurement of the jet energy resolution (JER) is made using two methods, and compared to particle-level jets in simulation.

The dijet balance method constructs the asymmetry relationship used in the η-intercalibration procedure (section 3.2.2, equation 3.18). The width of the asymmetry distribution is given by

$$\sigma(\mathcal{A}) = \frac{\sqrt{(\sigma(p_{\mathrm{T}}^{\mathrm{ref}})^2 + \sigma(p_{\mathrm{T}}^{\mathrm{probe}})^2)}}{p_{\mathrm{T}}^{\mathrm{avg}}}, \tag{3.21}$$

and may be measured directly as the width of a Gaussian distribution fit to its core. The dijet balance method used for the determination of the JER uses the same reference region as that used in the η-intercalibration, and provides a measurement in bins of $p_{\mathrm{T}}^{\mathrm{avg}}$ and η_{det}

3.2.5 Flavour Tagging

The ability to discriminate well between heavy-flavour and light-flavour (b or c vs. light quark or gluon) jets plays a crucial role in the physics programme of the general-purpose LHC experiments. The SM Higgs decay channel $H \to b\bar{b}$ possesses a large branching fraction, and top quarks produce final states with b-jets nearly 100% of the time. The heavy flavour content of many BSM scenarios is rich, particularly in natural supersymmetry models where light squarks are decoupled while the stop or sbottom are within reach of the LHC (chapter

1). A b-quark which is produced by the LHC will hadronise into a B-hadron: a bound state of a b- and light-quark with a average lifetime around 1 ps, depending on its exact content. These hadrons carry most of the energy of the b quark and fly a short distance away from the primary vertex before decaying into additional hadrons, electrons or muons – a process which is key in the identification of these jets.

Three different methods to tag b-jets are applied simultaneously within ATLAS, which are in-turn taken as the inputs to a multivariate algorithm which provides the final b-tagging discriminant for each jet. Impact parameter (IP) taggers extrapolate tracks back to the primary vertex, and compare their points of closest approach: tracks originating from secondary vertices possess a systematically larger impact parameter than those originating from the primary vertex itself. Secondary vertex (SV) taggers reconstruct the most likely secondary vertex candidate in each jet, and examine properties such as its invariant mass and the relative energy of tracks associated with the candidate secondary vertex compared to other tracks associated with the jet for compatibility with b-hadron decay. The JetFitter algorithm attempts to exploit the decay properties of b-hadrons, which decay through D-hadrons and create additional secondary vertices along the track's trajectory. This approach fits tracks passing through sequential vertices and extrapolates the fit back to the primary vertex, developing a discriminant based on compatibility with a b- or c-quark decay.

The IP, SV and JetFitter results are provided as inputs to a multivariate algorithm, though the details of this procedure differ between runs 1 and 2. During run 1, the MV1 algorithm was a neutral-network-based tagger which received preprocessed combinations of the three baseline taggers, and utilised other multivariate techniques in the preprocessing [52]. MV1 was trained to discriminate b-jets from light-flavour jets, and provided adequate performance during the initial run, and was utilised in the studies presented in chapter 4. In run 2, the baseline tagger outputs are fed directly into the MV2 algorithm, which utilises a boosted decision tree (BDT) to directly combine their information without intermediate processing steps. The MV2 algorithm is trained to discriminate b-jets from a mixture of c and light-flavour jets. For the results presented in chapters 5-7, the MV2c10 algorithm was applied, which includes a mixture of 10% c-jets and 90% light-flavour [53, 54].

The performance of these taggers is illustrated in figure 3.4, where the c- and light-flavour jet rejection is shown as a function of the b-tagging efficiency for several taggers considered by ATLAS during analysis in run 1 and 2, including the MV1 and MV2c10 algorithms. The performance of b-tagging between run 1 and 2 has improved by more than a factor of 2 in c-jet rejection, and by nearly 40% in light-jet rejection for the same 77% b-tagging efficiency working point. A direct comparison between a run 1 and a run 2 b-tagging algorithm is shown in figure 3.5, where the impact of the IBL on b-tagging performance may be seen as

the improved rejection of c- and light-flavour jets at lower values of p_{T} [53].

3.3 Large-R Jets

When a heavy object is produced in LHC collisions with a significant Lorentz boost, its decay products become collimated in the direction of its boost. Massive electoweak bosons and top quarks may be reconstructed as single, large-radius (large-R) jets rather than by combining distinct objects. This method of reconstruction eliminates many combinatoric factors which otherwise complicate hadronic final states, at the cost of introducing a new challenge: one must be capable of differentiating the large-R jets originating from the signal processes of boosted bosons and top quarks from those originating from the background process of large jets initiated by light quarks and glons with wide-angle hadronic activity. This classification is performed by examining the internal energy density distributions of the large-R jet, and identifying substructure within it which originates from the heavy particle decay.

The conventional large-R jets used in ATLAS physics analysis are built from topological clusters calibrated with the LCW scheme, using the anti-k_t algorithm and a distance parameter of $R = 1.0$. Due to the large size of these jets, the impact of pile-up on their reconstruction is more severe than for standard jets: a jet grooming technique is applied, which trims away areas of low energy density within the large-R jet. After this trimming procedure has been applied, the energy and mass scales of these large-R jets are calibrated using dedicated procedures, in order to correct for the detector response of the non-compensating ATLAS calorimeter. Once trimmed and calibrated, the substructure of large-R jets may be examined for compatibility with the decays of heavy boosted objects, such as electoweak bosons or top quarks.

3.3.1 Jet Trimming

The effect of pileup on jets is exacerbated when dealing with large-R jets, as contamination from pileup increases quadratically as a function of the jet algorithm's distance parameter R [55]. Different treatments for large-R jets exist, collectively known as jet grooming algorithms, which endeavour to minimise the impact of pileup on the jet four-vector and substructure observables. A comprehensive study carried out by ATLAS during run I found that **jet trimming** [56] outperformed other grooming techniques for values of μ between 1 and 15 [57].

Jet trimming is a four-step algorithm which removes soft components of radiation from jets, proceeding by:

1. Clustering the event using a standard jet algorithm to obtain 'seed jets.'

2. Within each seed jet, reclustering the input constituents with a smaller distance parameter, R_{min}, than that used to initially construct the seed jet.

3. Discarding subjets which fail to satisfy the trimming condition,

$$\frac{p_T^{\text{subjet}}}{p_T^{\text{jet}}} > f_{\text{cut}}.$$

4. Combining the remaining subjets to form the trimmed jet.

Standard parameter values of R_{min} and f_{cut} used by ATLAS were $R_{\text{min}} = 0.3$ and $f_{\text{cut}} = 0.05$ during run I, and $R_{\text{min}} = 0.2$ and $f_{\text{cut}} = 0.05$ during run II.

The benefits of trimming go beyond resiliency to pileup: the mass resolution of jets originating from the boosted decays of bosons and top quarks is generally improved by the application of grooming techniques. Large-R jets originating from the wide-angle decays of gluons and light quarks tend to lack hard structure, and so are generally more affected by the trimming procedure. This results in a relatively smaller trimmed mass than for background jets, leading to better performance when using the trimmed jet mass to tag boosted objects.

There are some small caveats to this additional performance: when gluon and light-quark jets do possess some hard structure, it can be enhanced by the trimming procedure. This results in decreased discrimination power for some jet substructure observables (discussed in section 3.3.3), notably those which are more sensitive to soft structure such as the N-subjettiness. The N-subjettiness values used by ATLAS are occasionally calculated on the initial untrimmed jet and then associated with the final trimmed one for this reason.

3.3.2 Large-R jet calibration

For the same reasons as those given for standard jets outlined in chapter 3.7, large-R jets must also be calibrated before they may be used in analysis. As both the energy and mass scales are of interest when reconstructing massive objects at high p_{T}, a calibration for each quantity must be derived.

Large-R jet Energy Scale

A Monte Carlo-based calibration is derived to correct the energy scale of reconstructed trimmed, anti-k_t $R = 1.0$ jets back to the scale at truth-level. Analogously to the calibration derived for standard jets, this correction is made by scaling the central value of the jet energy response in bins of $p_{\text{T}}^{\text{true}}$ and η_{det} of the large-R jet. There is no jet origin or pileup correction applied: dependence on pileup is mitigated by the application of the jet trimming procedure.

Uncertainties on the Large-R JES are provided by making a comparison to jets built from charged particle tracks in the Inner Detector which are geometrically associated with calorimeter jets, using the so-called r_{trk} method. As systematic effects between the calorimetry and tracking subsystems are uncorrelated, comparisons between calorimeter and track jets provide a reliable method to factorise effects from detector and physics processes. The ratio of the calorimeter to track jet p_{T} is formed, and compared between data and Monte Carlo by constructing a double ratio:

$$r_{\mathrm{track}}^{p_{\mathrm{T}}} = \frac{p_{\mathrm{T}}^{\mathrm{calo}}}{p_{\mathrm{T}}^{\mathrm{track}}}, \tag{3.22}$$

$$R_{\mathrm{track}}^{p_{\mathrm{T}}} = \frac{r_{\mathrm{track}}^{\mathrm{p_T,data}}}{r_{\mathrm{track}}^{\mathrm{p_T,MC}}}. \tag{3.23}$$

The double ratio $R_{\mathrm{track}}^{p_{\mathrm{T}}}$ is taken, as a function of p_{T}, to be the uncertainty on the large-R jet energy scale calibration. This uncertainty was typically between 4-8% of the large-R jet p_{T} during run I, and is between 2-4% in run II following the installation of the IBL.

During run I, an *in-situ* residual calibration was derived for large-R jets using photons as reference objects for the balancing technique outlined in section 3.2.2. Photons balanced against large-R jets provide a well-calibrated constraint on the jet energy scale in data for large-R jets with $p_{\mathrm{T}} < 800$ GeV and $|\eta| < 2.1$, beyond which the method is statistically limited by the number of high-p_{T} photons available within the dataset. An additional uncertainty is derived with this method, which accounts for differences in calorimeter response due to the event topology (QCD dijet events vs. $t\bar{t}$ pair production). Uncertainties provided by the γ-large-R jet balance tend to be smaller than those provided by the r_{track} method: about $1-3\%$ relative to the large-R jet energy scale.

Large-R Jet Mass Scale

Following the calibration of the large-R jet energy scale, a similar calibration to correct the average reconstructed large-R jet mass to particle level is performed. The jet mass response is binned in terms of the truth-level large-R jet p_{T}, the reconstructed large-R jet pseudorapidity without origin correction η_{det}, and the truth-level large-R jet mass m^{true}. In each bin, the difference between the mean of the response distribution and unity is taken as a correction factor, obtained using the same numerical inversion technique described in section 3.2.2:

$$m_{\mathrm{LC+JES+JMS}}^{\mathrm{jet}} = \frac{m_{\mathrm{LC+JES}}^{\mathrm{jet}}}{\mathcal{F}(m_{\mathrm{LC+JES}}^{\mathrm{jet}})} = c^{\mathrm{JMS}} m_{\mathrm{LC+JES}}^{\mathrm{jet}}, \tag{3.24}$$

where $1/\mathcal{F}(m_{\mathrm{LC+JES}}^{\mathrm{jet}})$ is the numerically inverted mass calibration function.

As for the large-R JES calibration, uncertainties are defined on the JMS calibration by comparing track and calorimeter jets in data and Monte Carlo, using the r_{track} ratios

$$r^{m}_{\text{track}} = \frac{m^{\text{calo}}}{m^{\text{track}}}, \tag{3.25}$$

$$R^{m}_{\text{track}} = \frac{r^{\text{m,data}}_{\text{track}}}{r^{\text{m,MC}}_{\text{track}}}. \tag{3.26}$$

Uncertainties on the large-R jet mass scale defined using this technique were typically between $4-8\%$ during run I and $3-6\%$ during run II.

3.3.3 Jet substructure observables

Large-R Jet Mass

The large-R jet mass is the cornerstone upon which most jet substructure techniques are built. Many other substructure observables, particularly those which examine the hard substructure of a jet, are highly correlated with the jet mass, and so tend to provide modest gains in discriminating power when compared to that obtained by using the jet mass itself as a discriminant. The mass of a jet is defined to simply be the mass of the sum of its constituent four-vectors, whether those constituents are truth particles, topological clusters or calibrated jets built using a smaller R parameter; a large-R jet which contains the decay products of a high-p_{T} top quark will have the top quark's mass (figure 3.6). When jets are built from more than a single constituent, the jet mass is dominated by the angular separation of the most energetic constituents. In the event that a large-R jet is built from a single constituent, as is occasionally the case when reconstructing boosted objects by reclustering smaller jets (a procedure detailed in chapter B), the mass of the large-R jet is simply the mass of this input object.

k_t Splitting Scales

One feature of the k_t clustering algorithm discussed in section 3.2.1 is that it combines constituents sequentially from the softest to hardest splittings. By re-clustering the constituents which were used to build a large-R jet using the k_t algorithm, one may examine the last n combinations of the clustering history for signs of hard substructure within the jet. The k_t **splitting scale** between protojets i and j is defined as

$$\sqrt{d_{ij}} = \min(p_{Ti}, p_{Tj}) \times \Delta R_{ij} \tag{3.27}$$

where

$$\Delta R_{ij} = \sqrt{\Delta\phi_{ij}^2 + \Delta\eta_{ij}^2}. \tag{3.28}$$

When the last step of the k_t combination combines two protojets of roughly equal p_T which are well distributed within the jet, the k_t splitting scale associated with this combination will be large, as shown in figure 3.7 where the values of d_{12} and d_{23} are shown for boosted top quarks and background processes. These characteristics are typical of a heavy object decaying into a pair of objects with similar masses; for example, the decay of a Higgs boson into a $b\bar{b}$ pair. In such cases, the k_t splitting scale associated with this hard splitting will peak at roughly half of the decaying state's mass. When splittings are instead soft or collinear, the corresponding splitting scale will be characteristically small.

N-Subjettiness

The 'pronginess' of a jet's radiation may be quantified according to the **N-subjettiness** (τ_N), which quantifies how well-described a jet's radiation would be by an N-subjet-like configuration. This process requires the definition of a candidate N-pronged structure for the jet, which is achieved by reclustering the jet's constituents with the k_t algorithm, but truncating the sequence prematurely when exactly N subjets remain. The momenta of the jet constituents are then projected onto the nearest axis such that

$$\tau_N = \frac{1}{d_0} \sum_k p_{T,k} \cdot \min\{\Delta R_{1,k}, \Delta R_{2,k}, ..., \Delta R_{N,k}\}, \tag{3.29}$$

where $\Delta R_{i,k}$ is the distance in $\eta - \phi$ space from the kth jet constituent to the ith protojet axis, and d_0 is the normalisation factor

$$d_0 = \sum_k p_{Tk} \times R. \tag{3.30}$$

If the value of τ_N is small, then the jet is well described by an N-pronged hypothesis. Should a large portion of the jet's constituent radiation lie far from any of the axes, the structure would not be well-described by the given hypothesis and the corresponding value of τ_N would be large. Ratios of N-subjettiness variables are particularly effective in distinguishing the decays of heavy objects; for instance, the ratio $\tau_{32} = \tau_3/\tau_2$, shown in figure 3.8, has been studied by various ATLAS analyses in the context of distinguishing the decays of boosted top quarks from QCD backgrounds.

3.4 Electrons and Photons

Both electrons and photons are reconstructed based on the electromagnetic shower they produce in the EM calorimeter (section 2.3.2). A 3×5 cell sliding-window is used to scan the energy depositions in the LAr calorimeter for areas of local maxima, which are subsequently examined for their compatibility with the deposition of an unconverted photon, a converted photon or a prompt electron produced by the LHC collision. This distinction is primarily made based on the extrapolation of inner detector tracks to the calorimeter deposition. Photons carry no electric charge, and so pass through the ID without producing hits: if there is no track extrapolated to the calorimeter deposition, the object is labelled an unconverted photon. Roughly 30% of photons produced in ATLAS convert into an electron-positron pair within the tracker, and so if the matched track is extrapolated from a secondary vertex of the reconstructed event, the deposition is labelled as a converted photon. If the track is extrapolated from the primary vertex, it was likely that an electron was produced in the hard scatter, and the deposition is labeled as an electron instead. Following this classification, the final electron and photon candidate objects are built using windows of different sizes in the EM barrel calorimeter: a window of 3×7 cells is used for electrons and converted photons, while a window of 3×5 cells is used for unconverted photons. All electron and photon candidates in the EMEC are reconstructed using a window of 5×5 cells.

Electrons and photons are calibrated using a combination of MC-based and *in-situ* techniques, including corrections to data which create uniformity in the electromagnetic response across separate regions of the detector. Scale factors and resolution smearing factors are derived as a function of the energy, pseudorapidity and azimuthal angle of each electron/photon by studying the differences between $Z \to ee$ events in data and simulation.

3.5 Muons

Four different muon classifications are defined, based on combinations of information available from various detector subsystems. Quality requirements are placed on the ID and MS tracks which are used to reconstruct muons: a muon traversing the ID will typically leave 3 pixel, 8 SCT and 30 TRT hits. At least 5 SCT and 9 TRT hits are required to include the track, in addition to a requirement that no more than 2 pixel or SCT modules were traversed by the particle without registering a hit [58]. MS tracks are classified as '3-station' or '2-station' tracks depending on whether the muon traversed two or three layers of MDTs. The process of reconstructing the momentum of these tracks varies, but in general 3-station MS tracks give better p_{T} resolution, and are the most common type.

Standalone muons (SA) are reconstructed from tracks in the MS, but are not extrapolated to match a track in the ID. Segment-Tagged (ST) muons are reconstructed from tracks in the ID, which are extrapolated to match MS hits, but not a reconstructed MS track. Calo-Tagged (CT) muons are similar to ST muons, but the ID track is extrapolated to a corresponding deposit of energy in the calorimeter rather than to MS hits. Combined (CB) muons are reconstructed from MS tracks which are extrapolated back to matching ID tracks, and are the only class of muon used in the studies presented in chapters 4, and 5-7. Combined muons possess the lowest fake rate of any muon class, and have a better momentum resolution than either ID or MS tracks alone due to the statistical combination of tracks.

Muons are calibrated using the well-studied $J/\Psi \to \mu\mu$, $\Upsilon \to \mu\mu$ and $Z \to \mu\mu$ resonances, which provide sources of muons across a wide range of p_{T}. These resonances are used to derive scale factors for each muon's p_{T}, binned in η and ϕ, which correct simulation to agree with the observed data.

3.6 Taus

Taus are not used directly in the studies presented in chapters 4 and 5-7, though they are long-lived enough to produce a secondary vertex. This makes taus capable of faking a b-tagged jet (section 3.2.5), and so their hadronic decay modes are therefore indirectly relevant to the selections used in these studies.

Hadronically decaying tau leptons are reconstructed using anti-k_t $R = 0.4$ jets calibrated at the local hadronic scale (section 3.1.1), to which an additional tau-specific MC-based correction is applied. A multivariate identification technique using Boosted Decision Trees (BDTs) is applied to construct a discriminant used to determine how likely a jet is to have been produced by the hadronic decay of a tau [59]. This BDT is trained using a collection of 11 variables which describe the distribution of energy within the calorimeter, the relative energy fractions of the ID tracks associated with the tau-candidate, the jet's secondary vertex properties and information based on the reconstructed π^0 mesons within the jet.

3.7 Missing Transverse Momentum ($E_{\mathrm{T}}^{\mathrm{miss}}$)

The missing transverse momentum ($E_{\mathrm{T}}^{\mathrm{miss}}$) is the most complex object reconstructed in ATLAS, taking as inputs all other previously-reconstructed objects – electrons, photons, taus, muons and jets – as well as the topological clusters and tracks which have not already been associated with these physics objects. As electrons, photons and hadronically-decaying tau leptons may be also reconstructed as jets, preference is given to energy to be included within

objects *in this order*. The $E_{\mathrm{T}}^{\mathrm{miss}}$ is constructed as the negative sum of the momenta of these objects, as well as the *soft term*, which accounts for energy unassociated with reconstructed physics objects:

$$-\vec{E}_t^{\mathrm{miss}} = \sum_{\text{electrons}} \vec{p}_t + \sum_{\text{photons}} \vec{p}_t + \sum_{\text{taus}} \vec{p}_t + \sum_{\text{muons}} \vec{p}_t + \sum_{\text{jets}} \vec{p}_t + \sum_{\text{soft}} \vec{p}_t. \tag{3.31}$$

Some selection criteria are applied to all objects included in the $E_{\mathrm{T}}^{\mathrm{miss}}$, in order to ensure that only well-understood objects calibrated to their proper energy scales are included in the calculation. In some cases, selections have changed slightly between run 1 and 2: typical values are provided, while differences are indicated where appropriate.

Electrons are required to have $p_{\mathrm{T}} > 10$ GeV, and be located within $|\eta < 2.47|$ excluding the barrel-endcap transition region from $1.37 < |\eta| < 1.52$, while satisfying basic quality criteria. Photons are required to have $p_{\mathrm{T}} > 10$ GeV (run 2: $p_{\mathrm{T}} > 25$ GeV) and satisfy the "tight" photon quality working point [60]. Both electrons and photons included in the $E_{\mathrm{T}}^{\mathrm{miss}}$ are calibrated at the electromagnetic scale. In the event that the same deposition is reconstructed as both an electron and a photon, the electron is retained and the photon discarded.

Hadronically-decaying taus are accounted for by finding jets with at least $p_{\mathrm{T}} > 10$ GeV (run 2: $p_{\mathrm{T}} > 20$ GeV) within $|\eta| < 2.5$ which satisfy the "medium" tau identification working point [61]. These objects are included in the tau term of the $E_{\mathrm{T}}^{\mathrm{miss}}$, and their energy scale is calibrated with the appropriate tau-specific MC calibration.

Muons are required to be CB within $|\eta| < 2.5$, or SA within $2.5 < |\eta| < 2.7$, beyond the ID coverage. They must satisfy $p_{\mathrm{T}} > 5$ GeV (run 2: $p_{\mathrm{T}} > 10$ GeV) to be included in the $E_{\mathrm{T}}^{\mathrm{miss}}$ computation. To avoid double-counting, the energy deposited by muons in the calorimeter is subtracted from the appropriate term using either an estimate parameterised by the muon four-vector, or through a direct measurement when possible.

Jets are reconstructed from topological clusters which are calibrated at either the electromagnetic or local scale (section 3.1.1), using the anti-k_t algorithm and a distance parameter of $R = 0.4$, with the appropriate JES calibration applied (section 3.2.2). In order to reject those jets which arise from secondary pp interactions in the event (pile-up jets), a requirement on the fraction of the jet's energy associated with the primary event vertex may be imposed. Following calibration, a jet is required to satisfy $p_{\mathrm{T}} > 20$ GeV to be included in the $E_{\mathrm{T}}^{\mathrm{miss}}$ calculation. In the event that the energy depositions which are reconstructed as a jet are also reconstructed as another physics object, energy from the jet may still enter the $E_{\mathrm{T}}^{\mathrm{miss}}$ calculation provided that no more than 50% of the jet's total EM-scale energy is represented by the other object. In this case, the energy from the conflicting object is subtracted from

that of the jet before its inclusion.

The soft term is the final component of the $E_{\mathrm{T}}^{\mathrm{miss}}$. There are several different approaches to the calculation and calibration of this quantity, and so only the two formulations relevant to the studies contained within chapters 4 and 5 will be discussed. When applying the semileptonic $t\bar{t}$ selection necessary to study the performance of boosted top quark identification (chapter 4), the $E_{\mathrm{T}}^{\mathrm{miss}}$ soft term is chosen to be the **cluster soft term** (CST), which is built from energy depositions within the calorimeter which are not associated with reconstructed physics objects. This approach to calculating the soft term is sensitive to the effects of pile-up, and so in run 2 (notably, in the context of the search for gluino pair-production presented in chapter 5), the $E_{\mathrm{T}}^{\mathrm{miss}}$ soft term is defined as the **track soft term** (TST). The TST is constructed from ID tracks which are not within $\Delta R < 0.05$ of an electron or photon, $\Delta R < 0.2$ of a tau, or associated with a combined muon or jet. These ID tracks must also be associated with the primary vertex of the event, a requirement which provides stability within the challenging environment of pile-up during run 2 (figure 3.9).

Systematic uncertainties from the input objects are propagated through the $E_{\mathrm{T}}^{\mathrm{miss}}$ calculation. Uncertainties on the soft term are determined using comparisons between data and simulation with a $Z \to \ell\ell$ selection applied which includes a veto on reconstructed jets in the final state.

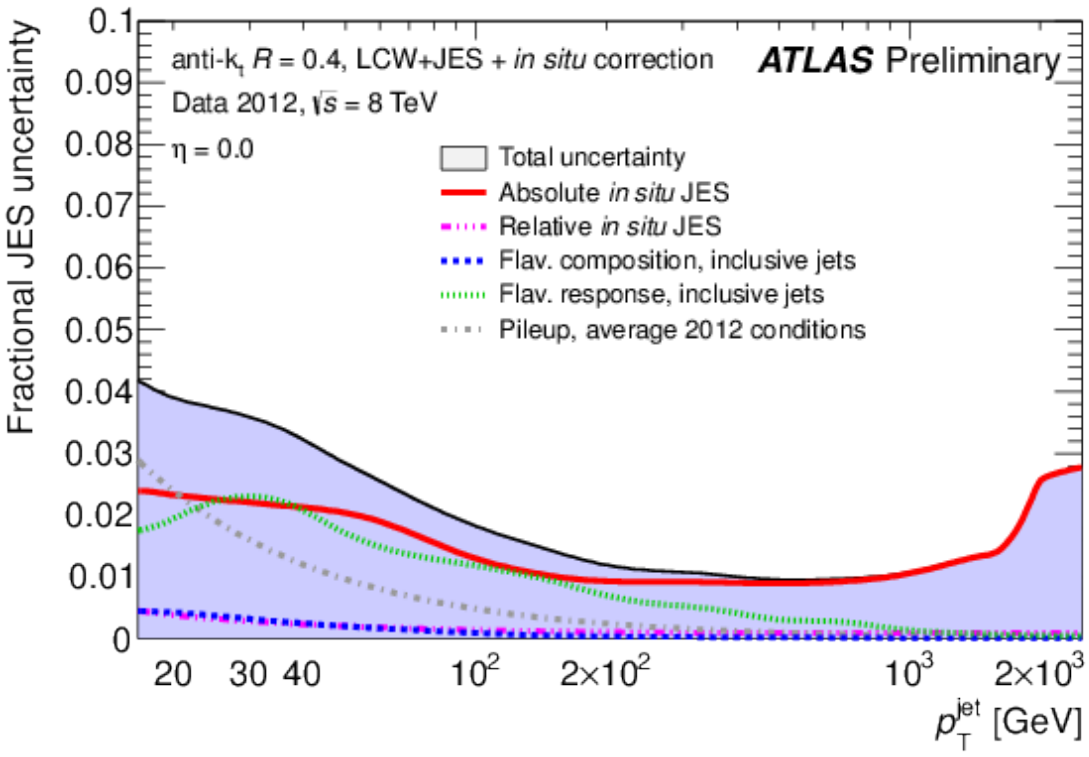

(a)

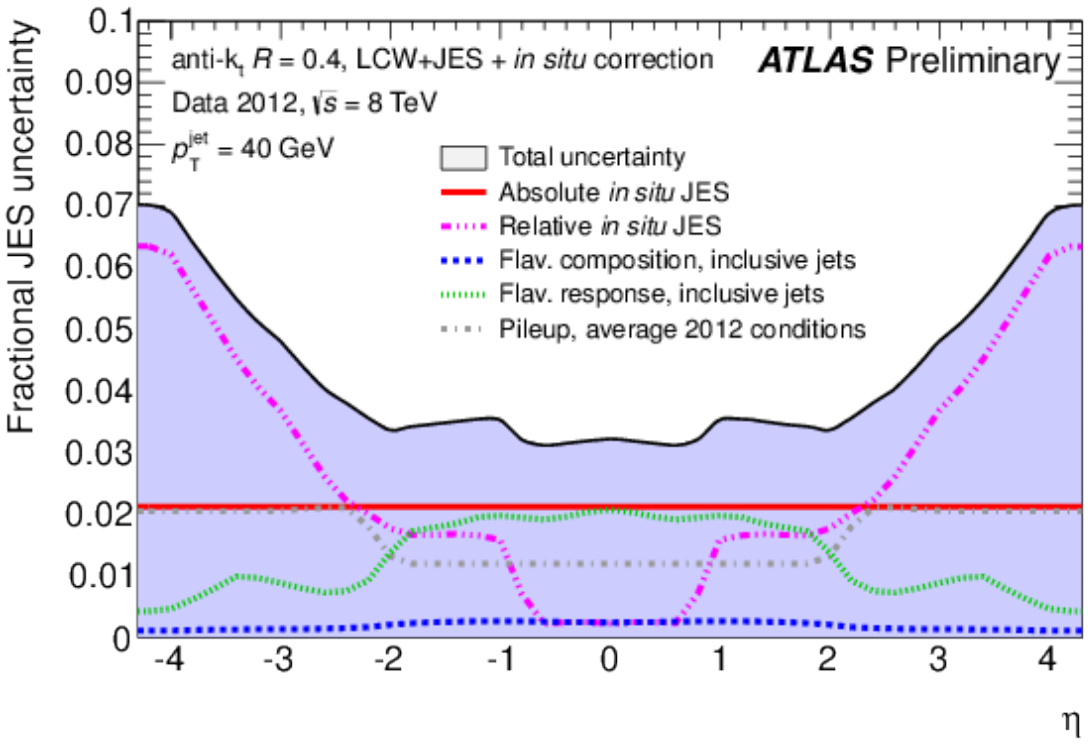

(b)

Figure 3.2: The total ATLAS jet energy scale uncertainty during run I, as a function of (a) transverse momentum and (b) pseudorapidity. From [50].

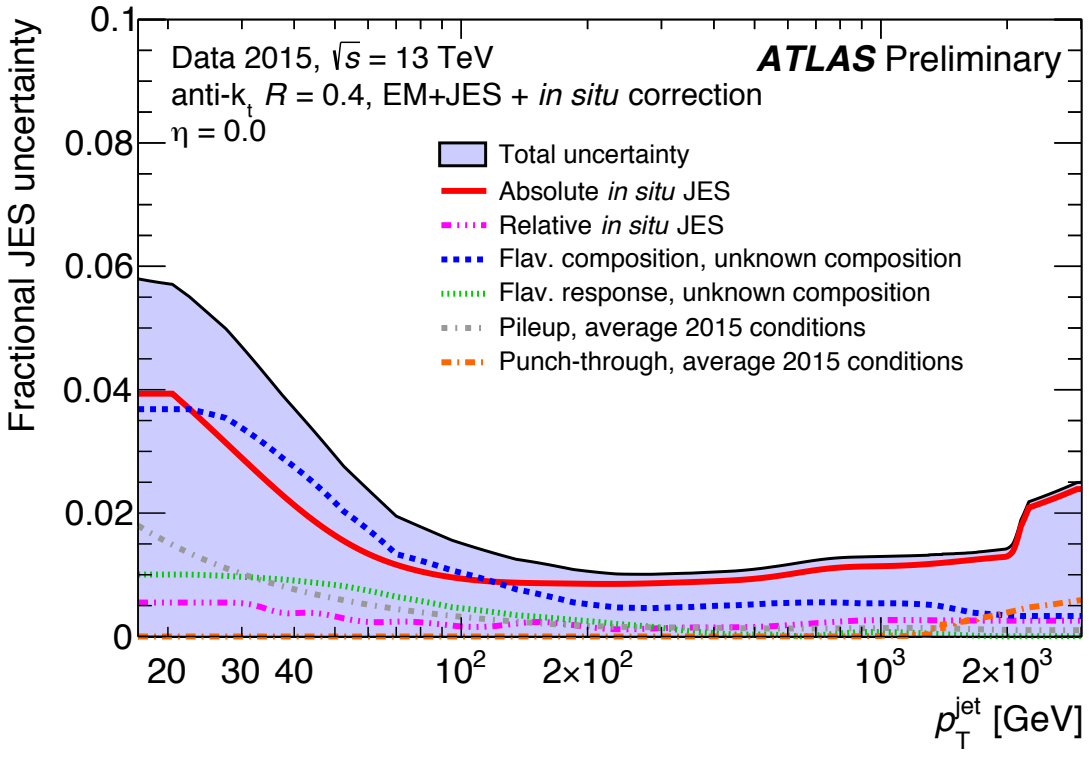

(a)

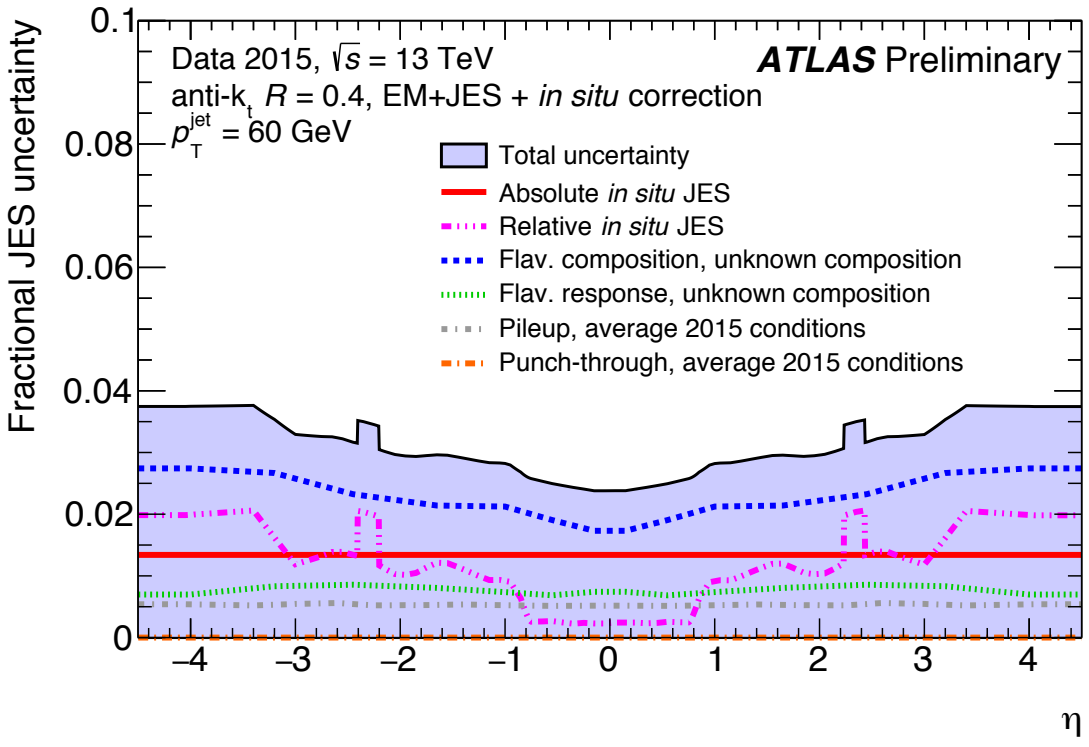

(b)

Figure 3.3: The total ATLAS jet energy scale uncertainty during run II, as a function of (a) transverse momentum and (b) pseudorapidity. From [51].

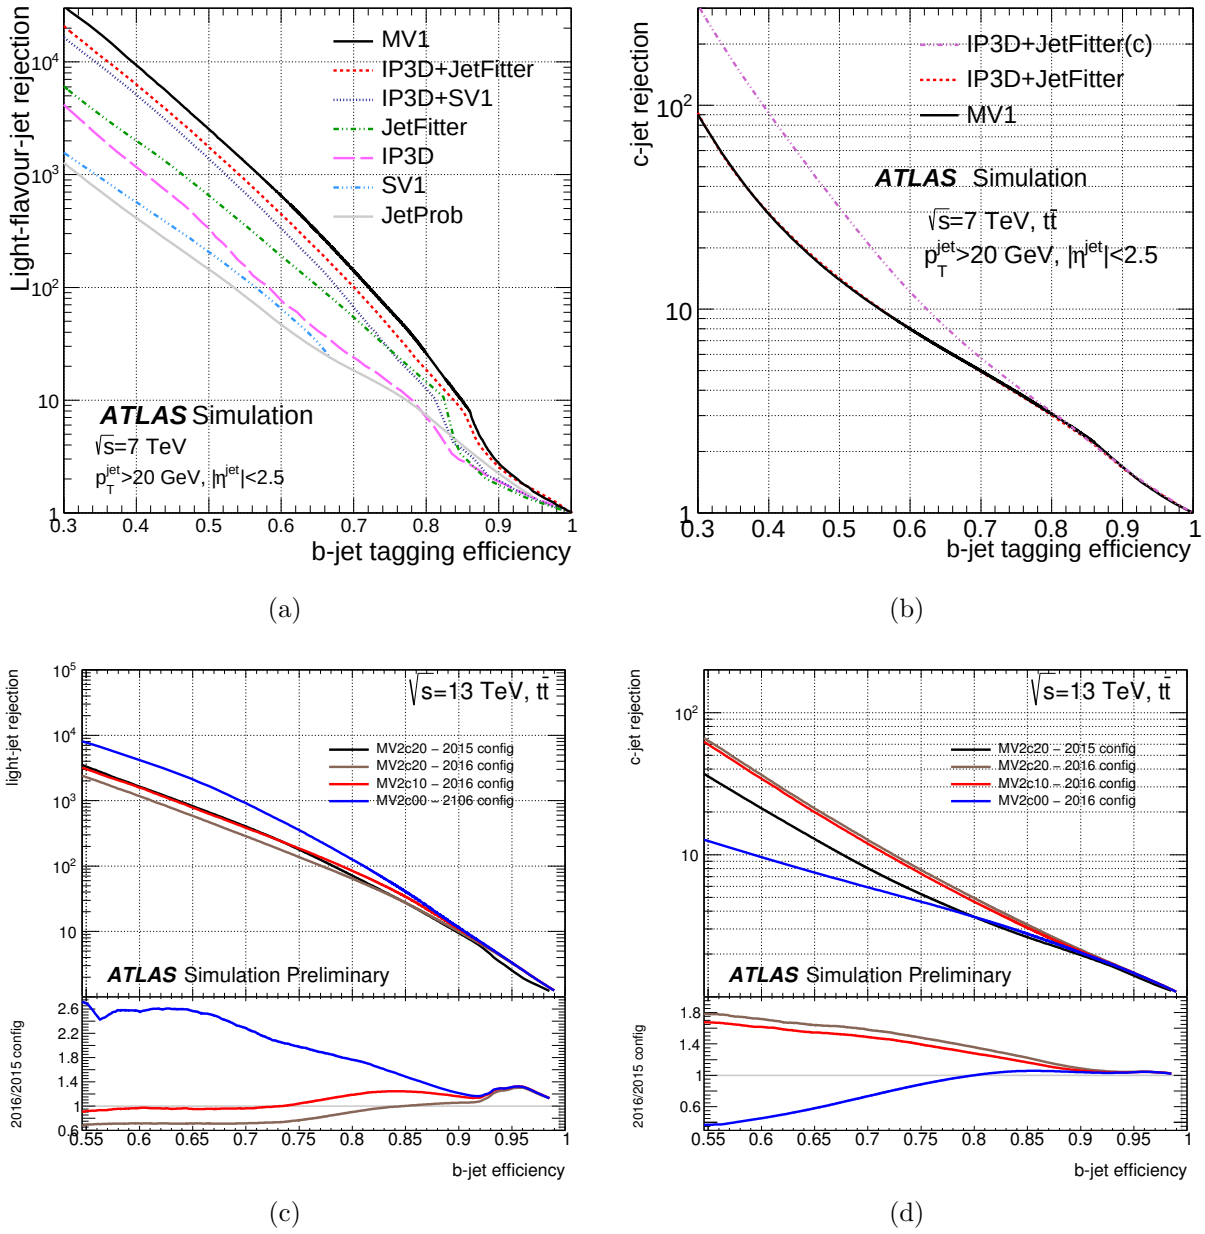

Figure 3.4: Above, the light- and c quark rejection as a function of the b-tagging efficiency for several b-tagging algorithms considered in the context of run 1 ATLAS analyses, including the MV1 tagger which produces a multivariate discriminant based on inputs from the other approaches. Below, The light- and c-quark rejection of the MV2 b-tagging algorithms, shown as a function of the b-jet tagging efficiency. The best overall performance is achieved by the MV2c10 variant, which trains a boosted decision tree to discriminate between b-tagged jets and a mixture of 10% c- and 90% light-flavour jets as background. From [52], [54].

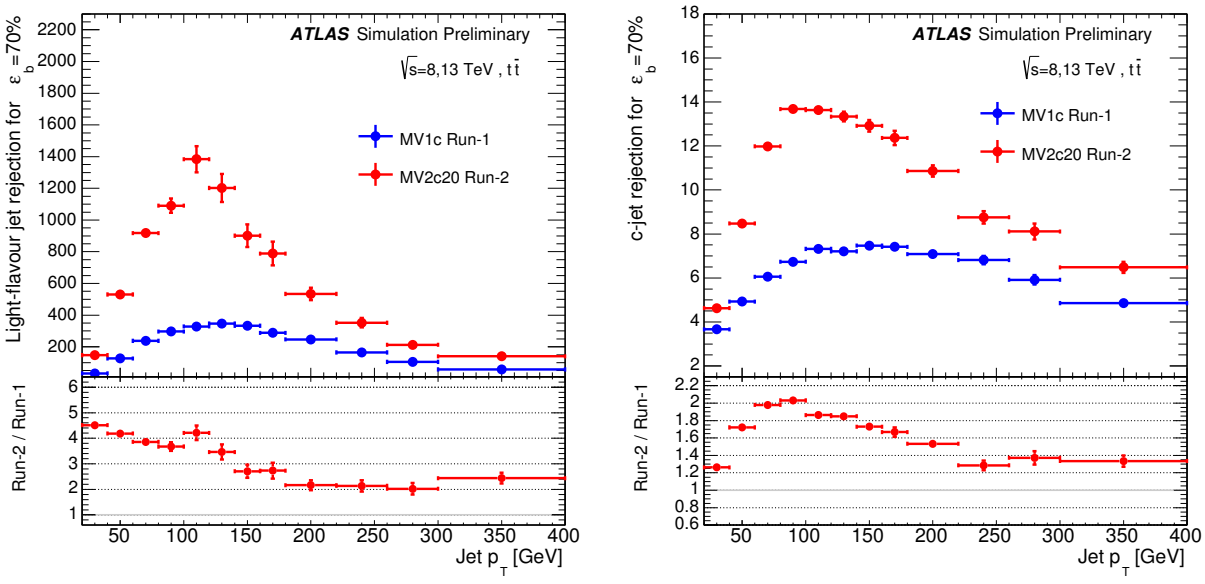

(a) (b)

Figure 3.5: A direct comparison of b-tagging performance between runs 1 and 2, where the c- and light-flavour rejection is shown as a function of jet p_{T}. The large performance gains at lower values of jet p_{T} are attributed with the addition of the IBL, while gains at higher jet p_{T} are associated with algorithmic improvements. From [53].

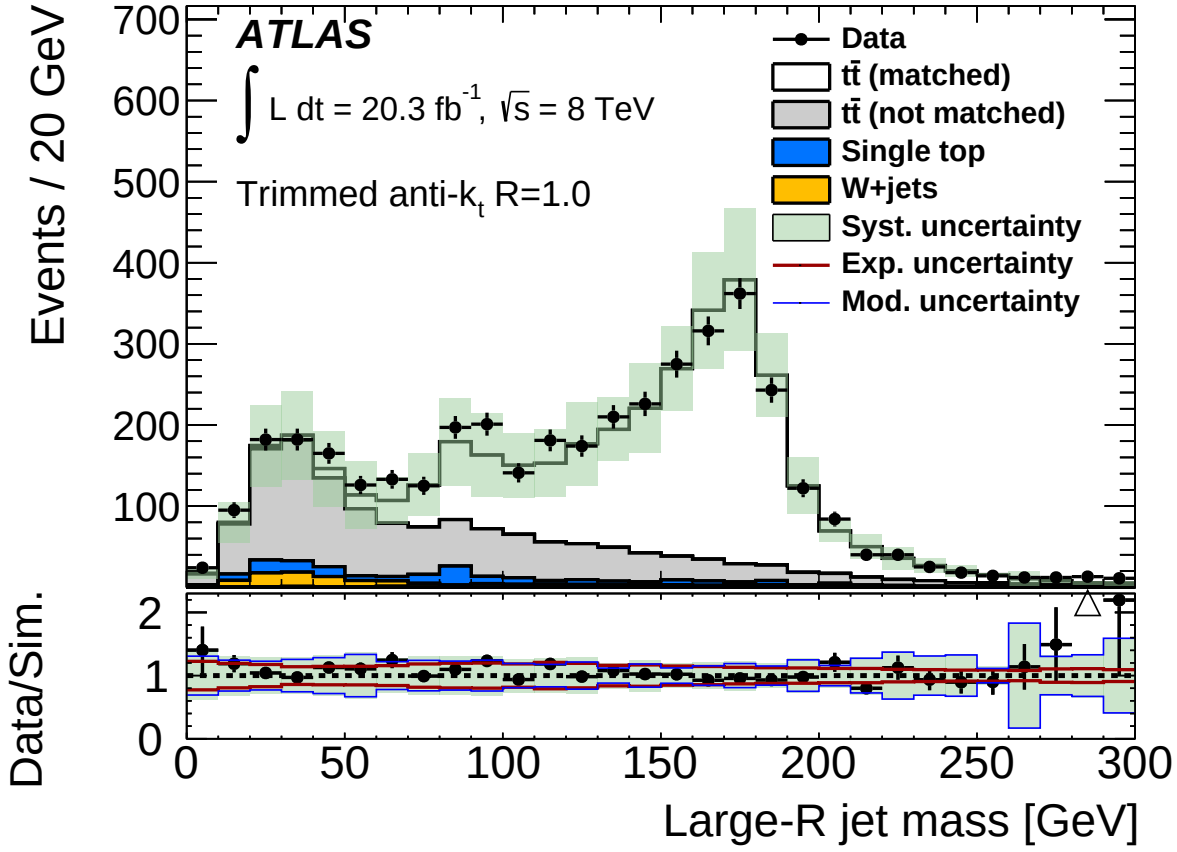

Figure 3.6: The large-R jet mass distribution following a the preselection for semi-leptonic $t\bar{t}$ events described in section 4.2. Peaks at the masses of the top quark and W boson are clearly visible, originating from the decays of boosted top quarks which are respectively fully- and partially-contained within a large-R jet. The structure at lower mass values is the QCD Sudakov peak, sculpted due to requirements placed on the large-R jet p_{T}, which is correlated with the jet mass.

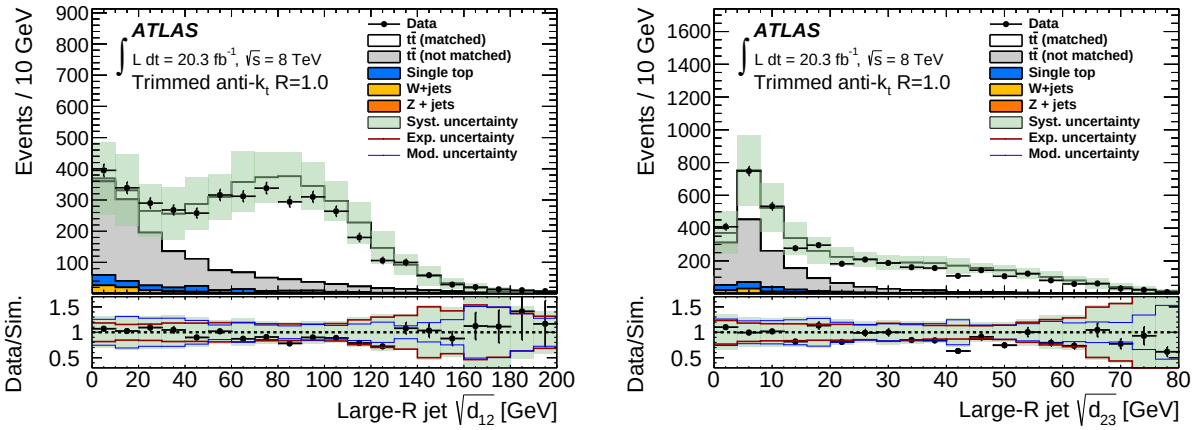

(a) $\sqrt{d_{12}}$ (b) $\sqrt{d_{23}}$

Figure 3.7: Distributions of the (a) first and (b) second k_t splitting scales for trimmed large-R jets following the preselection for semi-leptonic $t\bar{t}$ events described in section 4.2.

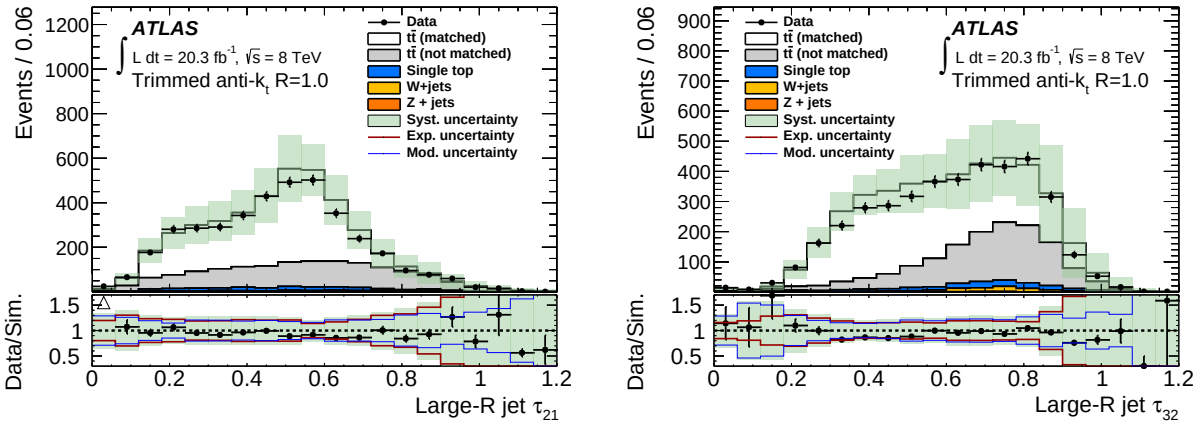

(a) τ_{21} (b) τ_{32}

Figure 3.8: Distributions of the (a) τ_{21} and (b) τ_{32} N-subjettiness ratios for trimmed large-R jets following the preselection for semi-leptonic $t\bar{t}$ events described in section 4.2.

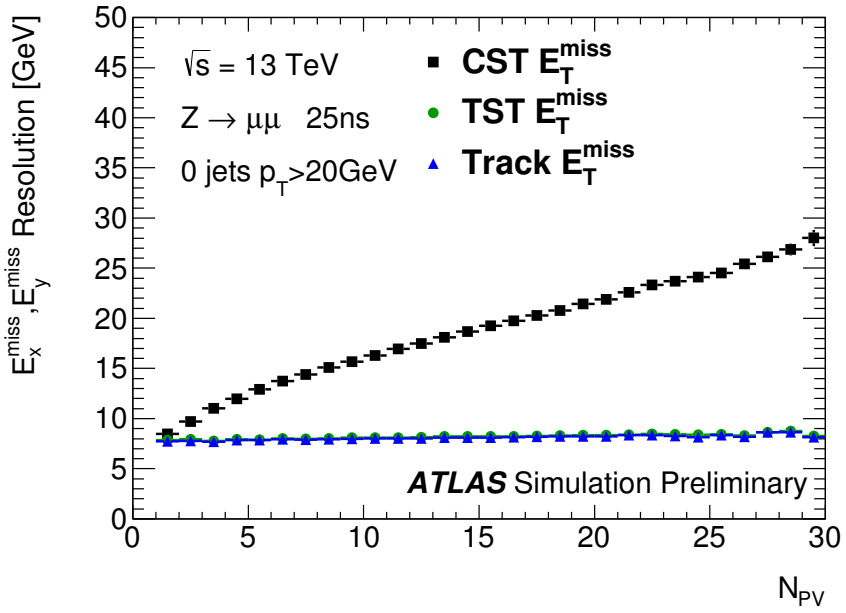

(a) ==0 jet selection.

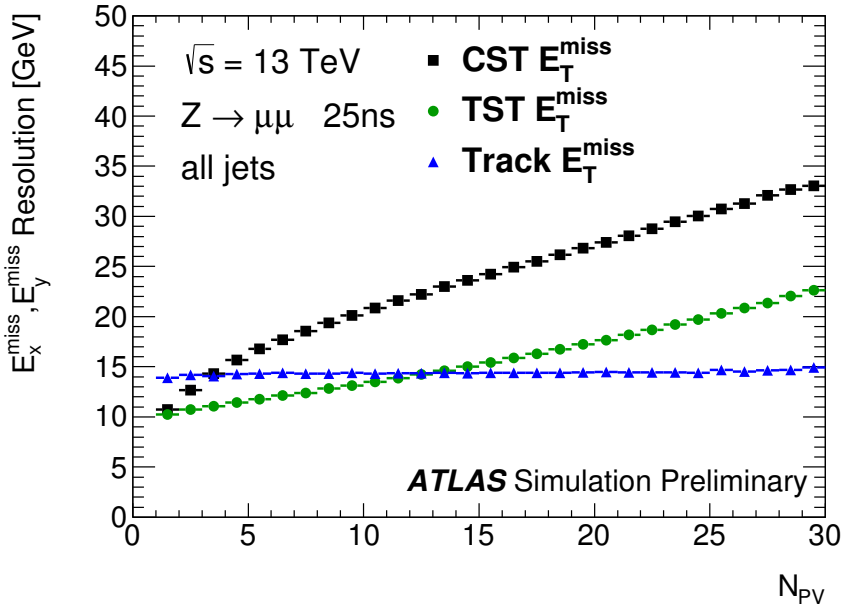

(b) Inclusive-jet selection.

Figure 3.9: A comparison of the $E_{\mathrm{T}}^{\mathrm{miss}}$ resolution, quantified as the RMS of its x and y components, as a function of the reconstructed number of vertices (NPV) in simulation, for selections either (a) requiring no jets or (b) including all jets. A strong dependence on NPV is observed for the cluster-based soft term (CST) in both cases, while the track-based soft term (TST) is more robust. The track-$E_{\mathrm{T}}^{\mathrm{miss}}$ demonstrates very little dependence on NPV, but excludes calorimeter information altogether. From [62].

4

Identifying Boosted Top Quarks

A study of the performance of various top quark tagging algorithms ranging from six simple cut-based jet substructure taggers to the more sophisticated HEPTopTagger [63, 64] and Shower Deconstruction (SD) [65, 66] algorithms was performed during the first long shutdown period of the LHC on data collected by the ATLAS detector during LHC run I [67, 2]. Previous studies of top tagging done by the ATLAS collaboration had not broadly considered the effect of systematic uncertainties on tagger performance, nor had the tagging efficiencies and fake rates previously been measured *in situ*. These measurements in particular represent a key result of the study, and are presented in sections 4.6 and 4.7.

4.1 Data and Simulation

This analysis was performed using 20.3 fb^{-1} of pp collision data at a centre-of-mass energy of $\sqrt{s} = 8$ TeV with a bunch spacing of 50 ns, collected by the LHC over the course of run I operations during 2012. Standard data quality requirements have been imposed on the data, which ensure that all detector subsystems were operational and not impacted by noise fluctuations or spurious effects which could compromise the integrity of the data for use in physics analysis.

The nominal $t\bar{t}$ signal process is simulated with Powheg-Box, interfaced with Pythia 6 with the CT10 set of parton distribution functions (PDFs) [68]. Other generators and parton showering algorithms are used to evaluate systematic uncertainties arising from the modeling of the ttbar process: these samples are outlined in the context of these uncertainties in section 4.4.2.

Single top production in the s- and Wt-channels was simulated with the same combination of generators as the nominal $t\bar{t}$ process. Diagram overlap between the $Wt-$ channel and the $t\bar{t}$ process was treated via the diagram removal scheme [69]. Processes including a vector boson

and additional jets (W+jets, Z+jets; collectively V+jets) were simulated using ALPGEN, interfaced to PYTHIA 6 with the CTEQ6L1 PDF set.

Simulated multijet events were used in the evaluation of tagger fake rates and background rejection factors. These events were generated and showered using PYTHIA 8 with the CT10 set of parton distribution functions [68]. The exotic $Z' \to t\bar{t}$ process is also used to obtain a simulated sample of boosted top quarks with a wide range of transverse momenta. This process is also simulated and showered with PYTHIA 8, using the MSTW2008 68% CL LO PDF set [70].

A summary of these simulated samples is provided in table 4.1

Process	Generator	PDF	Parton Shower
$t\bar{t}$ (nominal)	POWHEG-BOX	CT10	PYTHIA 6
single top	POWHEG-BOX	CT10	PYTHIA 6
V+jets	ALPGEN	CTEQ6L1	PYTHIA 6
multijets	PYTHIA 8	CT10	PYTHIA 8
$Z' \to t\bar{t}$	PYTHIA 8	MSTW2008	PYTHIA 8
$t\bar{t}$ ($\sigma^{\text{ISR/FSR}}$)	ACERMC	CTEQ6L1	PYTHIA 6
$t\bar{t}$ ($\sigma^{\text{MC Gen.}}$)	MC@NLO 6	CT10	HERWIG 6+JIMMY 4
$t\bar{t}$ ($\sigma^{\text{Parton Shower}}$)	POWHEG-BOX	CT10	HERWIG 6+JIMMY 4
$t\bar{t}$ (σ^{PDF})	POWHEG-BOX	HERAPDF	PYTHIA 6

Table 4.1: A summary of the Monte Carlo samples used to study the performance of top tagging during ATLAS run I.

4.2 Object and Event Selections

Two main event selections are used in this analysis: the **signal selection** selects semileptonic ttbar events containing a clean and isolated boosted top candidate, and the **background selection** which selects QCD multijet events in which the fake rates of the top tagging algorithms could be measured. All reconstructed events are required to include at least one primary vertex associated with at least five charged particle tracks with a $p_{\text{T}} > 400$ GeV. In both the signal and background selections, should any anti-k_t $R = 0.4$ jet with $p_{\text{T}} > 20$ GeV fail to satisfy at least the "looser" jet quality designation, the entire event is rejected from consideration.

4.2.1 Signal Selection

Events in the signal selection were triggered using one of four single-lepton triggers: two each for the electron and muon channels of the analysis. A single-muon trigger with an additional isolation requirement is used to provide muons with at least 24 GeV of trigger-level transverse momentum, and an additional trigger without isolation requirement provides muons with at least 36 GeV of trigger-level transverse momentum. Similarly, an electron trigger with an isolation requirement selects events containing a trigger-level electron with transverse momentum as low as 24 GeV, and a second trigger provides electrons without the additional requirement from p_T as low as 60 GeV. Events are required to contain exactly one reconstructed muon with at least 25 GeV of transverse momentum or exactly one reconstructed electron with at least 25 GeV of transverse energy, in order to ensure that the selected triggers are fully efficient offline.

Reconstructed muons may be rejected due to an overlap removal procedure which removes them from consideration when the muon lies with a p_T-dependent distance of an anti-k_t R=0.4 jet with a p_{T} above 25 GeV, namely if

$$\Delta R(\mu, \mathrm{jet}) < \left(0.04 + \frac{10\ \mathrm{GeV}}{p_T^{\mu}}\right). \tag{4.1}$$

If an electron is reconstructed within $\Delta R < 0.4$ of a anti-k_t $R = 0.4$ jet, its momentum is subtracted from that of the jet before jet selection criteria are applied, in order to account for energy from the electron which may be double-counted as contributing to that of the jet. If the jet still satisfies the jet selection criteria, the electron is removed and the original jet four-vector is replaced.

The leptonic decay of a W boson produces missing transverse momentum (E_T^{miss}), and exhibits specific kinematic properties which allow the intermediate state to be partially reconstructed. In particular, the **transverse mass** of the leptonically-decaying W boson, m_T^W, is constrained by the mass of the W boson such that $M_T^W < M^W$. It is defined as

$$m_T^W = \sqrt{2 p_T^{\ell} E_T^{\mathrm{miss}} (1 - \cos\phi)} \tag{4.2}$$

where p_T^{ℓ} is the lepton's transverse momentum, and ϕ is the azimuthal angle between $\vec{p_T^{\ell}}$ and $\vec{E_T^{\mathrm{miss}}}$ in the transverse plane. In order to select events consistent with leptonically decaying W bosons, at least 20 GeV of missing transverse momentum is required, and the sum of E_T^{miss} and M_T^W is required to be at least 60 GeV.

At least two anti-k_t $R = 0.4$ jets are required to have a transverse momentum above 25 GeV and to be b-tagged by the MV1 algorithm [52] at the 70% b-tagging efficiency working

point. One of these b-tagged jets is presumed to originate from each of the top quarks, and so additional geometric requirements are imposed based on the properties of a clean, back-to-back semileptonic $t\bar{t}$ system: one b jet is assumed to come from the leptonically decaying top and must be found within a distance $\Delta R < 1.5$ of the lepton, while the other is assumed to come from the hadronically decaying top and must be found outside of this region.

The leading trimmed anti-k_t $R = 1.0$ jet with $p_{\mathrm{T}} > 200$ GeV and $|\eta| < 2.0$ which satisfies the geometric requirement $\Delta R(\mathrm{jet}, \ell) > 1.5$ is taken to be the candidate hadronically-decaying top quark, and is the object to which various top-tagging algorithms will be applied. In simulation, these large-R jets are assigned to one of two categories based on their proximity to a truth-level top particle: jets within $\Delta R < 0.75$ of a hadronically decaying truth-level top are labeled as **truth-matched** jets, and it is the tagging efficiency and fake rate with respect to this sub-sample of $t\bar{t}$ events which will be measured. Other large-R jets which do not satisfy the truth matching requirement, even those from $t\bar{t}$ events, are labelled as **other** $t\bar{t}$ jets and form an additional background of the study.

4.2.2 Background Selection

In order to study the mis-tag rates of top-tagging algorithms, it is necessary to obtain a sample of the QCD multijet events which are representative of the typical background large-R jets which top tagging is applied to suppress. Obtaining a sample of these jets is complicated by the high thresholds placed on single-jet triggers during ATLAS run I: the lowest trigger available without a prescale has a jet p_{T} threshold of 360 GeV at the trigger level. This threshold biases the p_{T} spectrum of large-R jets in the event, *even if they are well-separated from the trigger object.*

One means by which this bias can be mitigated is to instead select multi-jet events where a low-p_{T} trigger-level jet fakes a trigger-level electron by using a single-electron trigger, available with a much lower threshold of 24 GeV, and then to require that zero reconstructed electrons are present in the event. Large-R jets are required to satisfy a distance requirement of $\Delta R > 1.5$ with respect to the trigger object. The large-R jet with the largest transverse momentum is selected as the candidate jet to which top tagging algorithms will be applied. When applying the background selection to simulated events, it is not necessary to apply the trigger requirements, and therefore the isolation requirements and electron veto are also omitted.

4.3 Substructure-based top taggers

Six cut-based top taggers which utilise the substructure of trimmed anti-k_t $R = 1.0$ jets to differentiate boosted, hadronically-decaying top quarks from the QCD multijet background were studied in order to form a baseline of well-understood taggers against which more sophisticated tagging methods could be compared. To construct these taggers, cuts on various combinations of the jet mass (sec. 3.3.3), k_t splitting scales (sec. 3.3.3) and N-subjettiness ratios (sec. 3.3.3) are defined. The resultant taggers are outlined in table 4.2 and ordered by their relative tagging efficiency and background rejection capability: for instance, tagger I is the most efficient and least rejective. Tagger III, used in the run I ATLAS search for semi-leptonic $t\bar{t}$ resonances [71], represents a good all-around tagger in terms of both its tagging efficiency and background rejection capabilities. The ATLAS W' top tagger, the least efficient and most rejective substructure-based tagger studied, acquired its name following its use in the run I ATLAS search for exotic W' bosons decaying to tb final states [72]. The intermediate taggers, appropriately, populate the intermediate tagging phase space. The effects of the six substructure taggers on the large-R jet p_{T} and mass spectra may be observed in data and simulation in figures 4.1 and 4.2. Direct comparisons between these six taggers and more sophisticated alternatives are made in simulation, and shown in section 4.5.

4.4 Systematic uncertainties

The effects of systematic uncertainties on these studies were subdivided into two categories: those uncertainties arising from experimental effects, and those caused by effects in the theoretical modelling of high-p_{T} top quarks.

4.4.1 Experimental uncertainties

The total luminosity measurement of the ATLAS dataset and its related uncertainty is made by a variety of subdetectors which are separately calibrated using van der Meer scans [23, 24], occasionally performed by the LHC during machine development periods. A combination of the luminosity measurements made by the Beam Conditions Monitor (BCM), LUCID detector, Tile and Forward Calorimeters is made using following the methodology used during the 2011 LHC operations [74], resulting in an uncertainty on the total luminosity measurement used in these studies of 2.8%.

The trigger, reconstruction and identification efficiencies are measured for electrons [75, 76] and muons [58] in events with leptonically-decaying Z bosons. Uncertainties on the

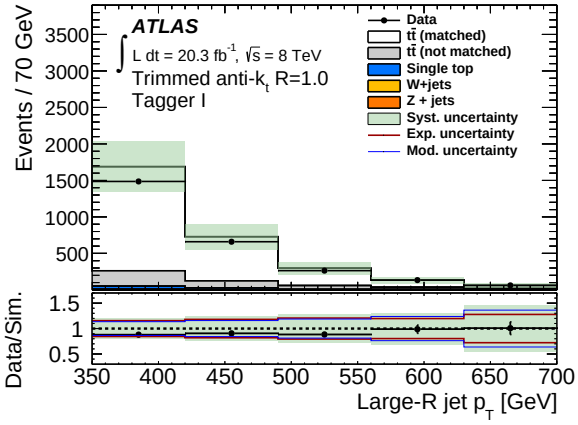

(a) Tagger I

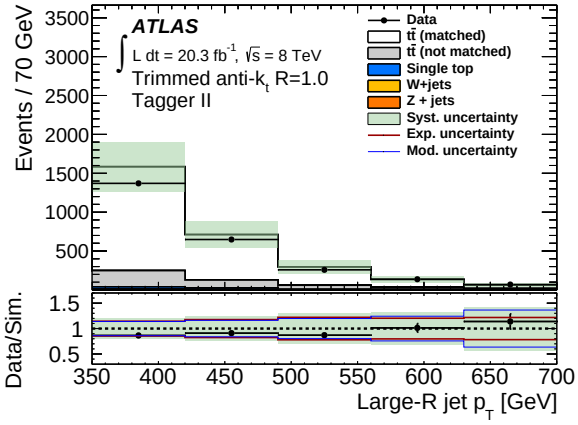

(b) Tagger II

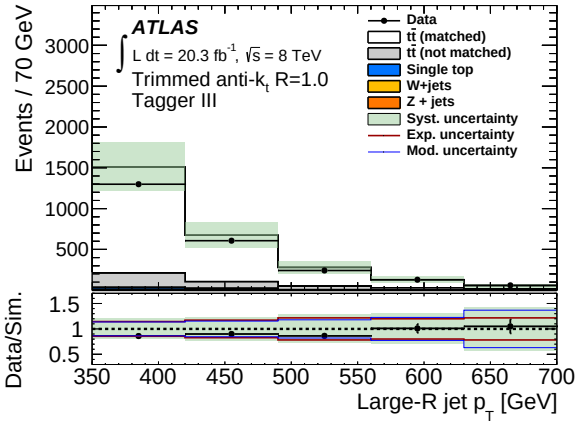

(c) Tagger III

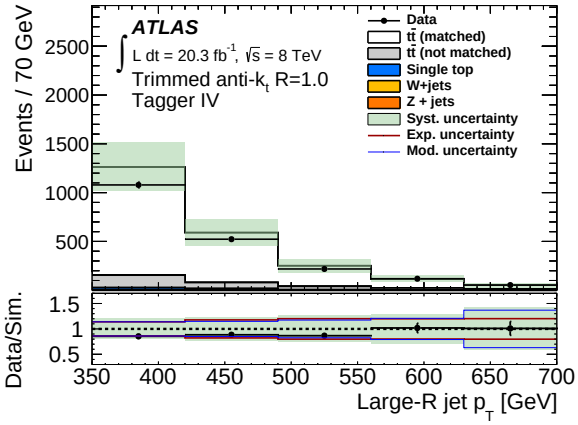

(d) Tagger IV

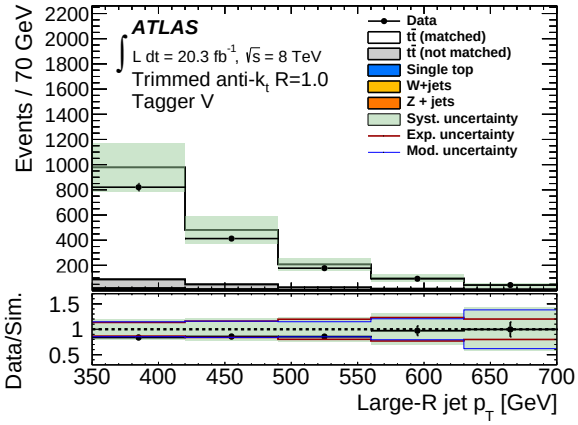

(e) Tagger V

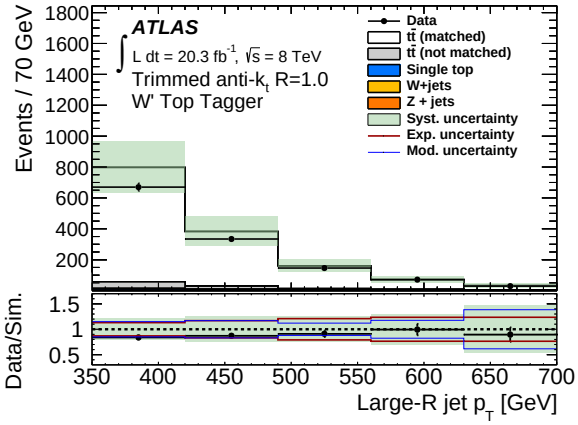

(f) Tagger VI

Figure 4.1: Transverse momentum spectra of large-R jets following the application of substructure-based top taggers I – VI.

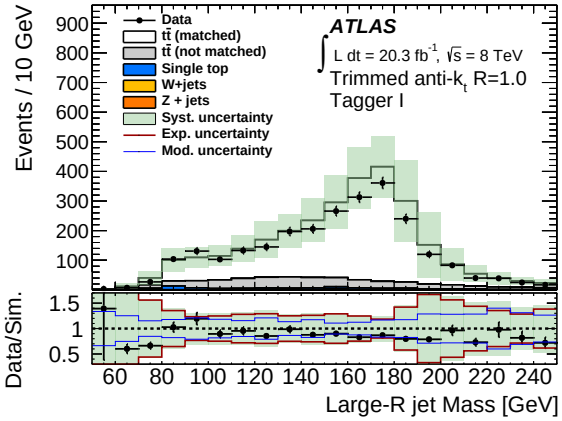

(a) Tagger I

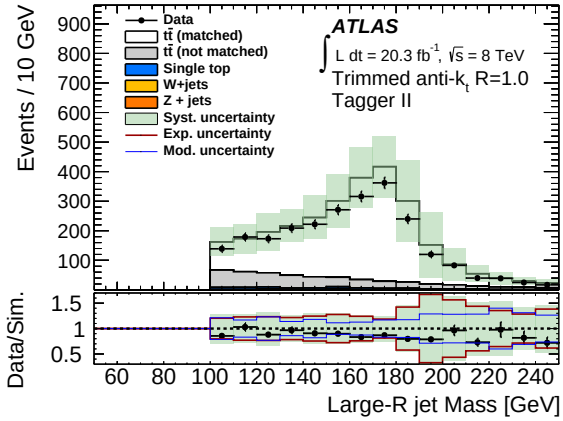

(b) Tagger II

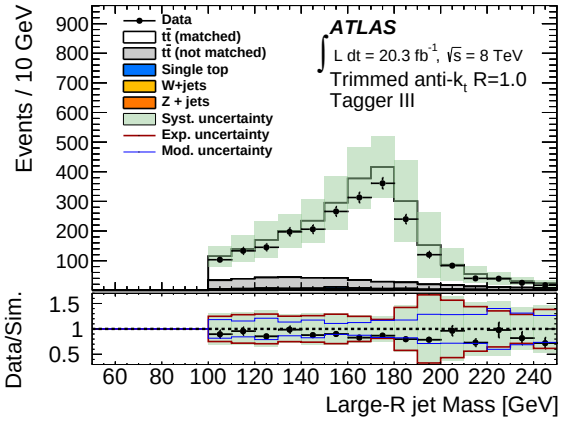

(c) Tagger III

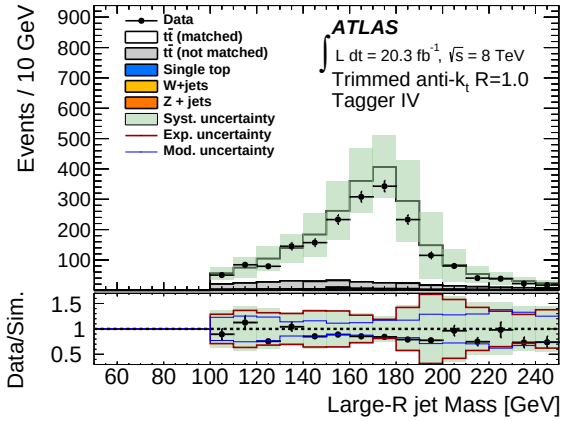

(d) Tagger IV

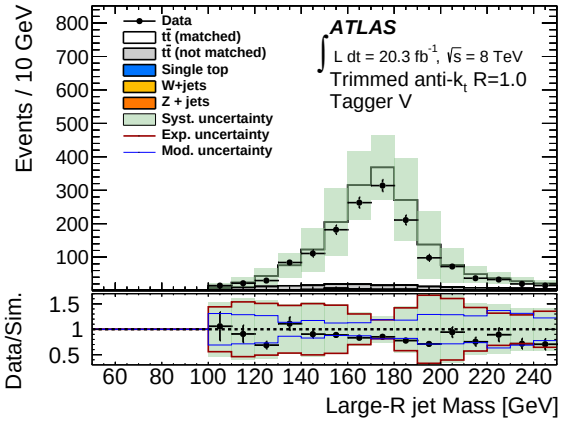

(e) Tagger V

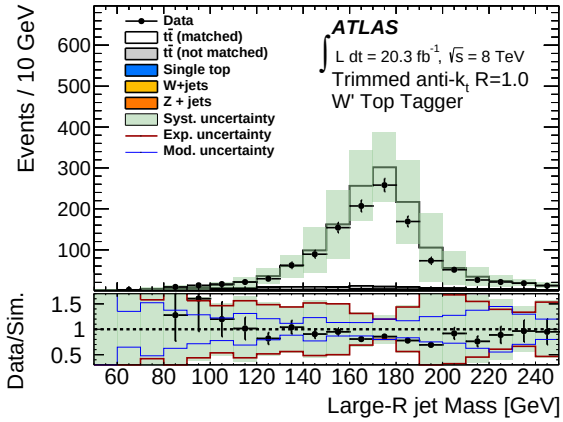

(f) Tagger VI

Figure 4.2: Large-R jet mass spectra following the application of substructure-based top taggers I – VI.

Tagger	Selection	Publication
I	$\sqrt{d_{12}} > 40$ GeV	
II	$m > 100$ GeV	Similar to J. High Energy Phys. 11 118 [73]
III	$m > 100$ GeV $\sqrt{d_{12}} > 40$ GeV	Phys. Rev. D 88, 012004 [71]
IV	$m > 100$ GeV $\sqrt{d_{12}} > 40$ GeV $\sqrt{d_{23}} > 10$ GeV	
V	$m > 100$ GeV $\sqrt{d_{12}} > 40$ GeV $\sqrt{d_{23}} > 20$ GeV	
VI, 'W' top tagger'	$\sqrt{d_{12}} > 40$ GeV $0.4 < \tau_{21} < 0.9$ $\tau_{32} < 0.65$	Eur. Phys. J. C 75: 165 [72]

Table 4.2: The six cut-based top taggers applied to trimmed anti-k_t $R = 1.0$ jets in the run-I top tagging studies. Where possible, the run I ATLAS publications where these taggers were first used have also been provided.

lepton energy scale and resolution for reconstructed leptons were found to be negligible, and so their effects on the final results are not explicitly considered.

Scale factors which correct the b-tagging efficiency in simulation are derived by performing fits to $t\bar{t}$ events and jets with associated muons [52]. A variation of these scale factors provides an estimate of the impact of systematic effects associated with b-tagging.

Studies used to estimate the systematic uncertainty for the jet energy scale and resolution of small-R and large-R jets in run I have been explained previously, respectively in sections 3.2.2 and 3.3.2. The methodology used to estimate the jet mass scale and resolution uncertainty of large-R jets and the uncertainties on the jet substructure observables used to define the substructure taggers have also been outlined in the latter section. As correlations between substructure observables have not yet been determined, the largest variation based on multiple correlation configurations was always selected in these studies.

4.4.2 Modelling uncertainties

The pair-production cross-section for top quarks with mass 172.5 GeV has been calculated at NNLO in QCD to be 253^{+13}_{-15} pb in $p\bar{p}$ collisions at $\sqrt{s} = 8$ TeV [77, 78, 79, 80, 81, 82, 83].

The theoretical uncertainty this calculation was taken along with an additional uncertainty on the $t\bar{t}$ cross-section of $^{+7.6}_{-7.3}$ pb, which arises due to uncertainty of the mass of the top quark, obtained by varying the top quark mass by ± 1.0 GeV. Both uncertainties were added in quadrature and propagated as a relative shift on the total normalisation of the $t\bar{t}$ process amounting to $^{+5.9}_{-6.6}\%$. The impact of increasing or decreasing the QCD renormalisation scale by a factor of two, which can alter the $t\bar{t}$ p_{T} spectrum, is also evaluated in this way and found to have an impact of roughly 1% [16]. Uncertainties from the choice of the CT10 PDF sets used in the nominal $t\bar{t}$ simulations for these studies were evaluated by evaluating the impact of switching to the MSTW2008 68% CL NNLO or NNPDF2.3 5f FFN sets [70, 84], and leads to an uncertainty of between 4-7%.

All other uncertainties arising due to the modeling of the $t\bar{t}$ process are estimated by comparing the nominal Powheg +Pythia $t\bar{t}$ samples used in the analysis to samples which have simulated the same process using a different generator, parton showering algorithm, PDF set, etc. For each comparison, the cross section of the nominal sample is used in the normalisation, in order to avoid multiple-counting of theoretical uncertainties related to the cross section itself.

The effect of initial and final state radiation (ISR, FSR) modelling is examined by comparing a pair of $t\bar{t}$ samples generated at leading-order with AcerMC and showered with Pythia which are tuned to include increased and decreased amounts of radiation relative to the nominal sample. An uncertainty is evaluated by comparing the predictions from the varied samples to each other, symmetrising the difference and applying it to the nominal sample:

$$\sigma^{\mathrm{ISR/FSR}} = N_{\mathrm{nominal}} \times \frac{N^{\mathrm{up}}_{\mathrm{AcerMC+Pythia}} - N^{\mathrm{down}}_{\mathrm{AcerMC+Pythia}}}{N^{\mathrm{up}}_{\mathrm{AcerMC+Pythia}} + N^{\mathrm{down}}_{\mathrm{AcerMC+Pythia}}}. \tag{4.3}$$

The impact of the choice of Monte Carlo generator is evaluated by comparing the $t\bar{t}$ sample to a pair of samples generated with MC@NLO and Herwig, but both showered in Herwig. This uncertainty is taken to be

$$\sigma^{\mathrm{MC\ Gen.}} = N_{\mathrm{nominal}} \times \frac{N_{\mathrm{MC@NLO+Herwig}}}{N_{\mathrm{Powheg+Herwig}}}. \tag{4.4}$$

The impact of the choice parton showering algorithm is evaluated by comparing the nominal $t\bar{t}$ sample to a pair of samples generated with Powheg, but showered with either Pythia or Herwig.

$$\sigma^{\mathrm{PS}} = N_{\mathrm{nominal}} \times \frac{N_{\mathrm{Powheg+Pythia}}}{N_{\mathrm{Powheg+Herwig}}}. \tag{4.5}$$

4.5 MC-based comparison of tagger performance

In order to compare the performance of different taggers to each other, a study in simulation was made which allows access to kinematic regions otherwise inaccessible using the ATLAS run I dataset. Simulated exotic $Z' \to t\bar{t}$ samples with masses from 400 GeV to 3 TeV provide a source of boosted top quarks over a wide p_{T} range[1]. Reconstructed, trimmed anti-k_t $R = 1.0$ jets are matched to particle-level large-R jets by requiring the jet axes to lie within $\Delta R < 0.75$ of one another. Background jets are similarly selected from Pythia multijet samples (sec. 4.2).

Events are reweighted such that the truth-level p_{T} spectrum of signal jets matches that of the background, in order to remove dependence on the Z' model from which the signal is obtained. Plots of the background rejection factor (defined as the reciprocal of the background tagging efficiency) as a function of the tagging efficiency are presented in bins of truth-jet p_{T} in figure 4.3, for truth-jet p_{T} values ranging from 350 GeV to 1.5 TeV. Individual substructure taggers are represented as distinct points in this space, demonstrating their expected ordering in tagging efficiency across the kinematic range. Curves obtained by scanning the threshold value of one component of these taggers are also drawn, which illustrate the relative performance of these variables at other points in tagging parameter space.

The trimmed mass and first k_t splitting scale are the most performant variables at high efficiency values, while the N-subjettiness requirements of the W' top tagger result in better performance at high values of background rejection. At low p_{T}, the substructure-based taggers are competitive with the more sophisticated tagging methods presented in this result – for instance, the W' tagger scan overlaps with the scan over the shower deconstruction discriminant. At high p_{T}, more variation is apparent across the parameter space, with the W' top tagger being the most rejective tagger of the substructure-based approaches at low efficiencies. It is relatively more difficult for the complex taggers to achieve high efficiency at high-p_{T}: for searches which do not require much background rejection and whose signal possesses a small cross-section (typical of searches for natural supersymmetry), a substructure-based tagger may be the best option across the entire p_{T} range.

[1] A simplified model is used in which the Z' possesses the properties of the standard-model Z boson, but a different mass.

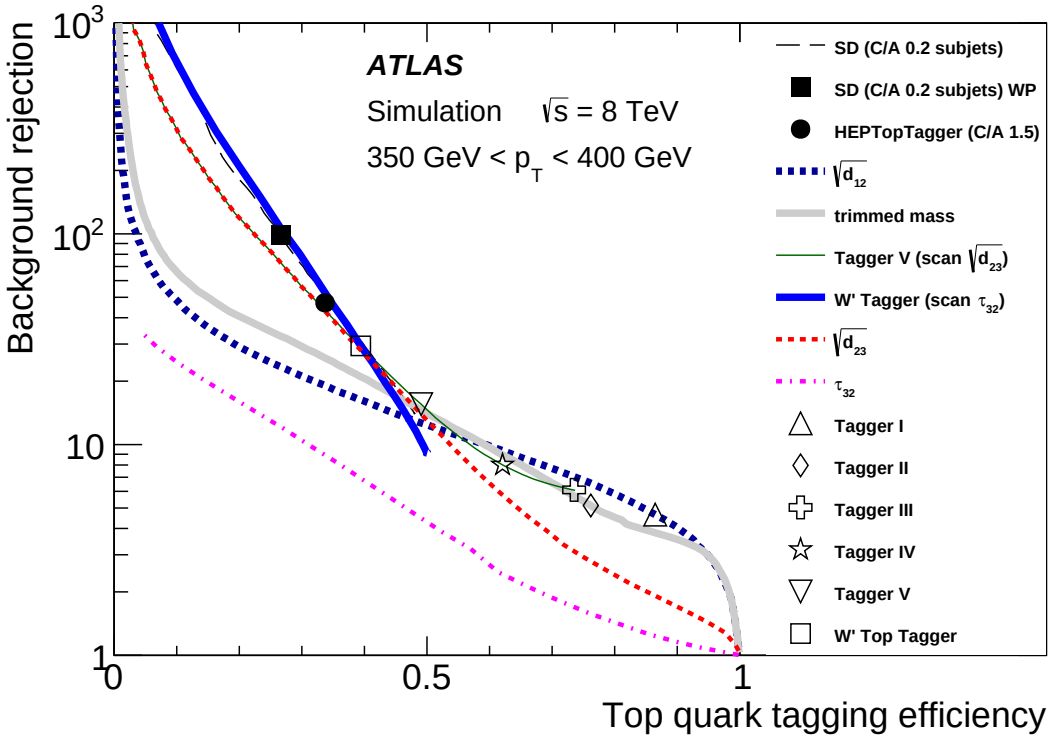

(a) 350 GeV $< p_T <$ 400 GeV

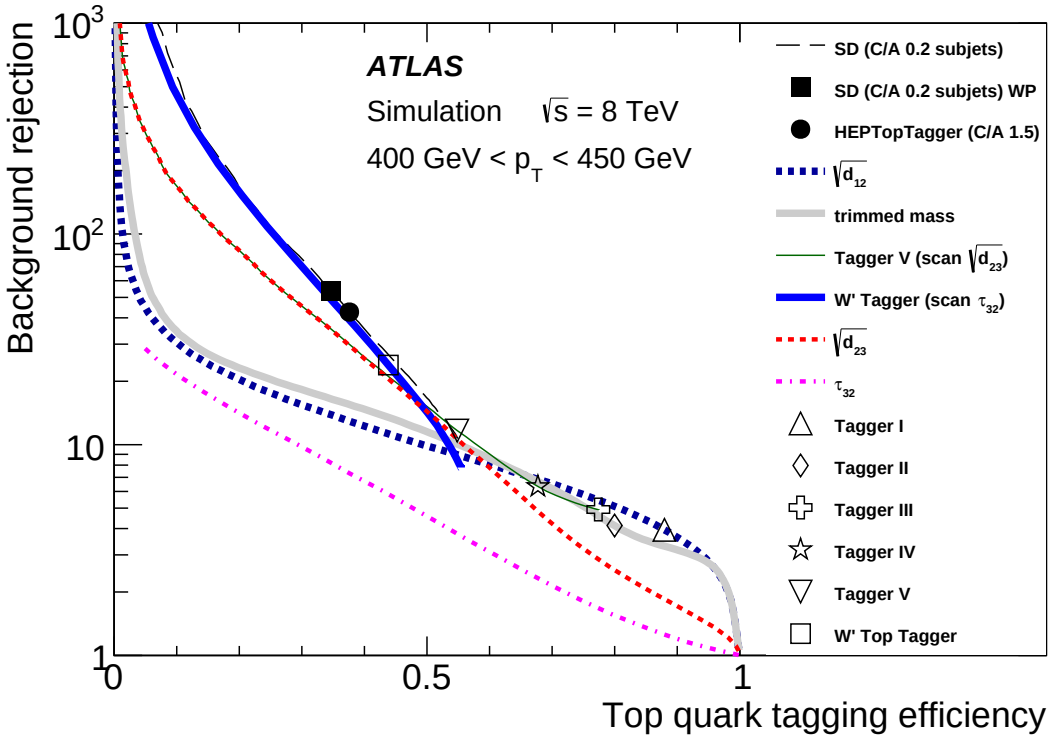

(b) 400 GeV $< p_T <$ 450 GeV

Figure 4.3: Comparison of the performance of various top-taggers in simulation, in terms of their boosted top-quark tagging efficiency and background rejection rate.

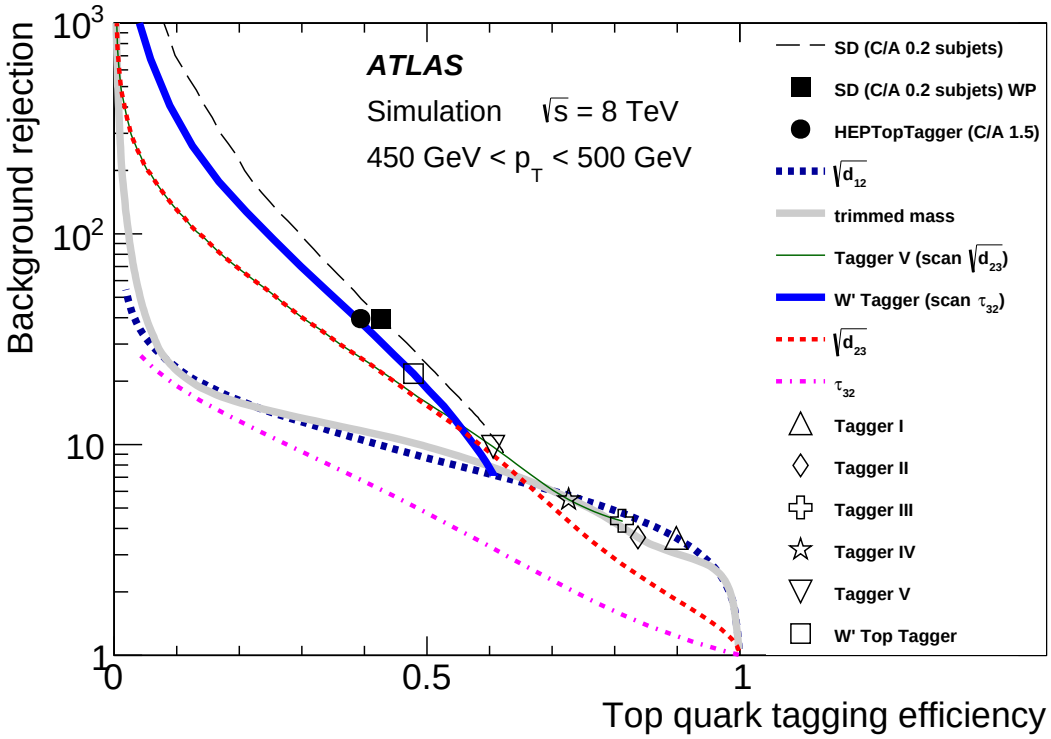

(c) 450 GeV $< p_T <$ 500 GeV

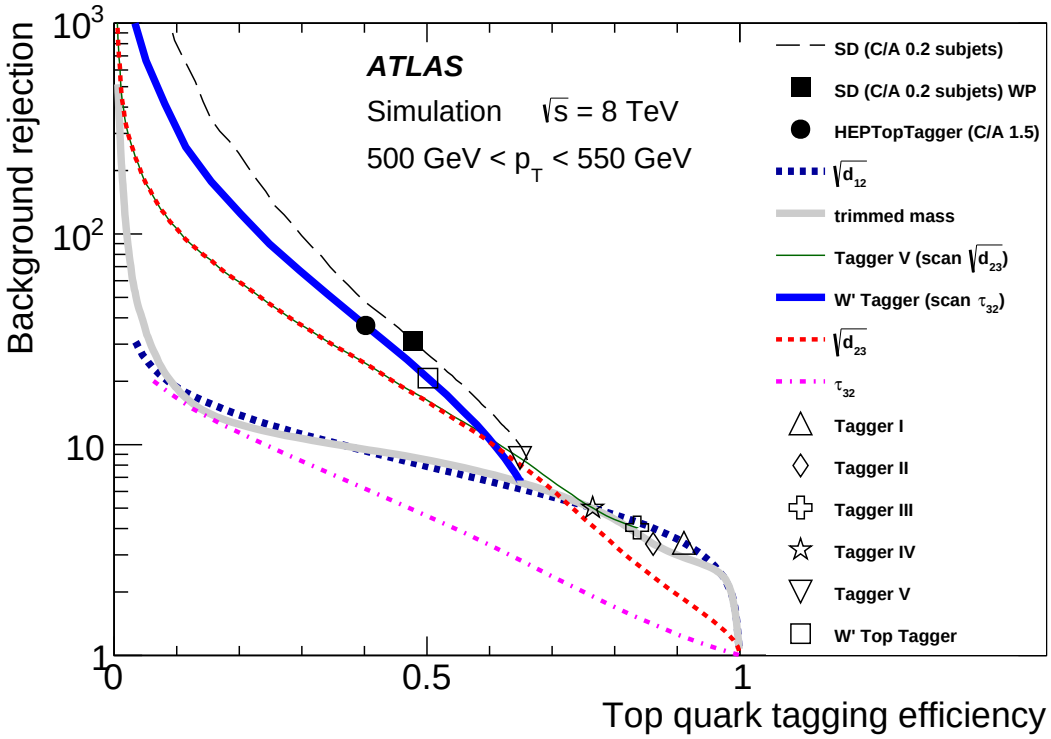

(d) 500 GeV $< p_T <$ 550 GeV

Figure 4.3: Comparison of the performance of various top-taggers in simulation, in terms of their boosted top-quark tagging efficiency and background rejection rate.

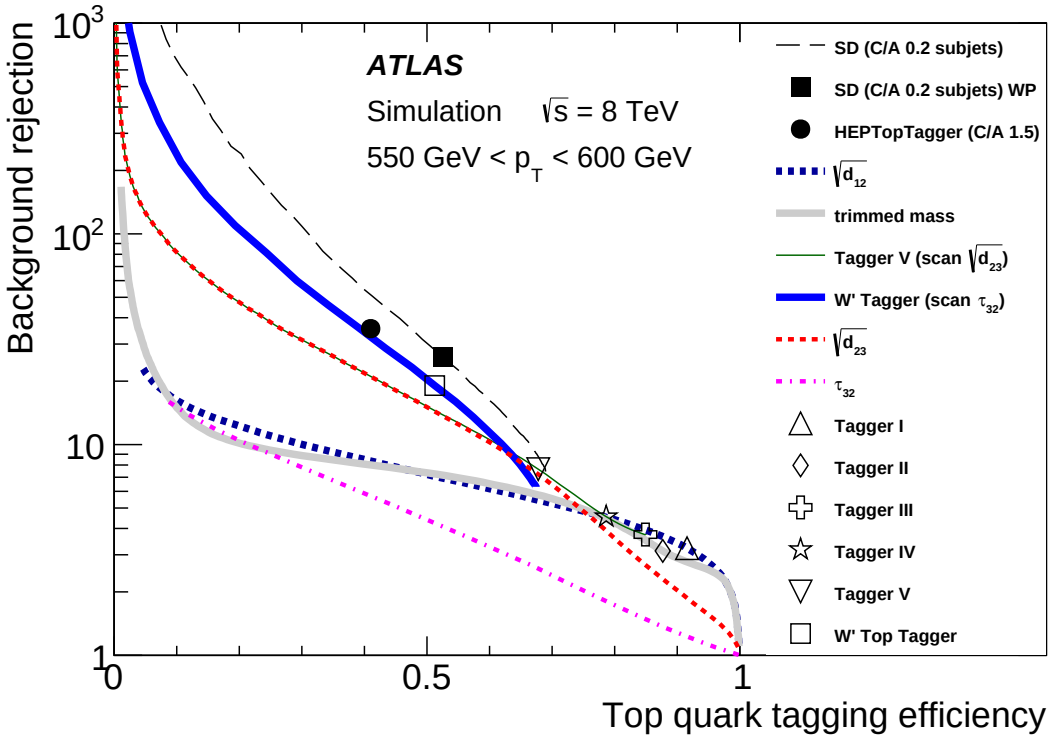

(e) 550 GeV $< p_T <$ 600 GeV

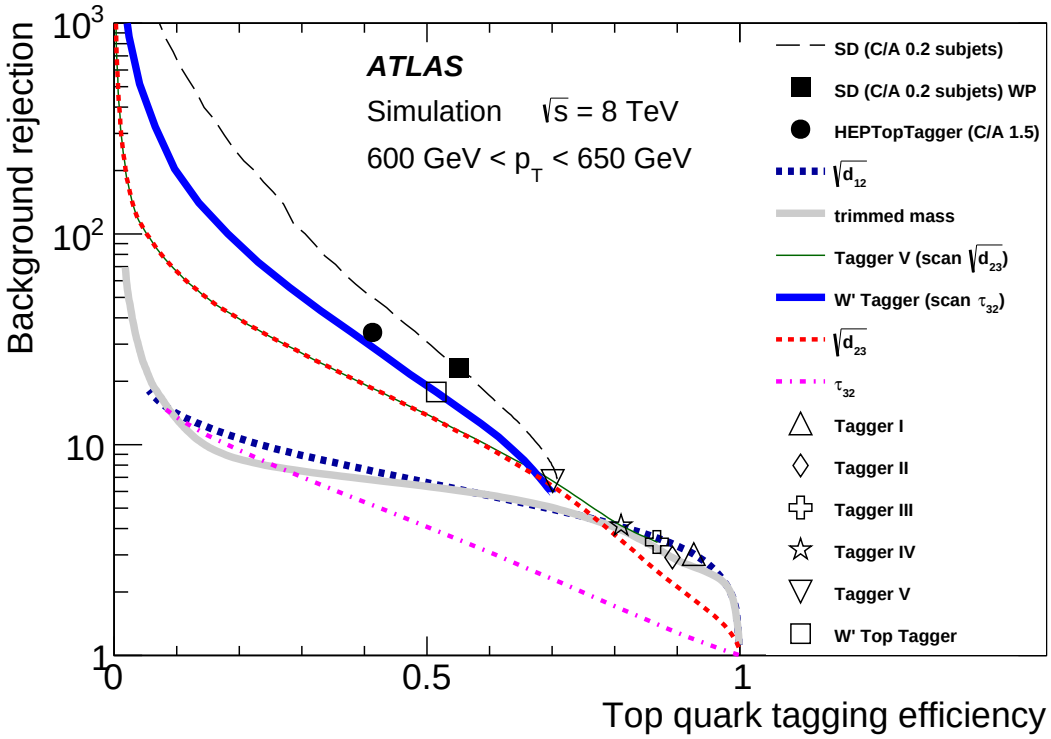

(f) 600 GeV $< p_T <$ 650 GeV

Figure 4.3: Comparison of the performance of various top-taggers in simulation, in terms of their boosted top-quark tagging efficiency and background rejection rate.

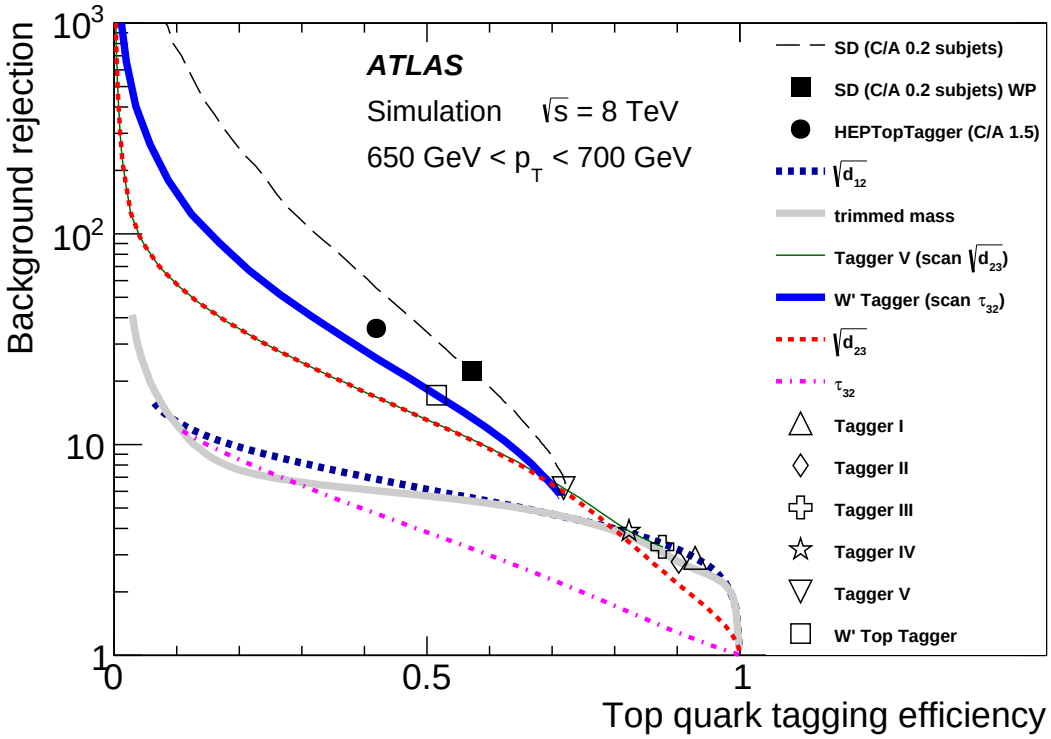

(g) 650 GeV $< p_T <$ 700 GeV

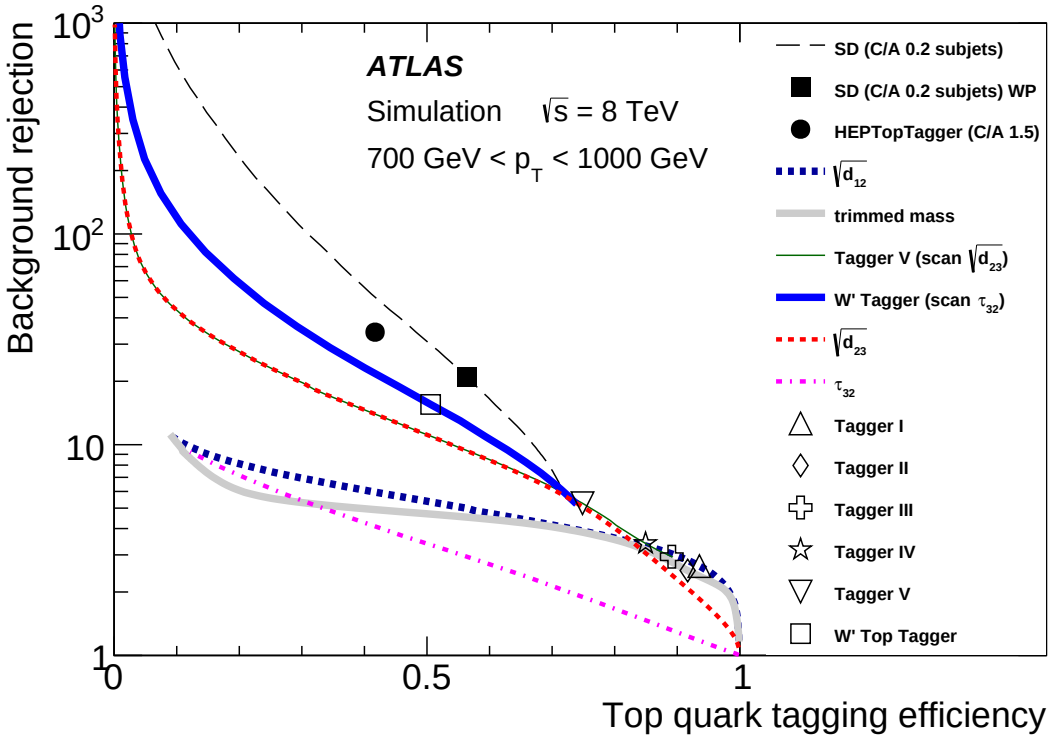

(h) 700 GeV $< p_T <$ 1000 GeV

Figure 4.3: Comparison of the performance of various top-taggers in simulation, in terms of their boosted top-quark tagging efficiency and background rejection rate.

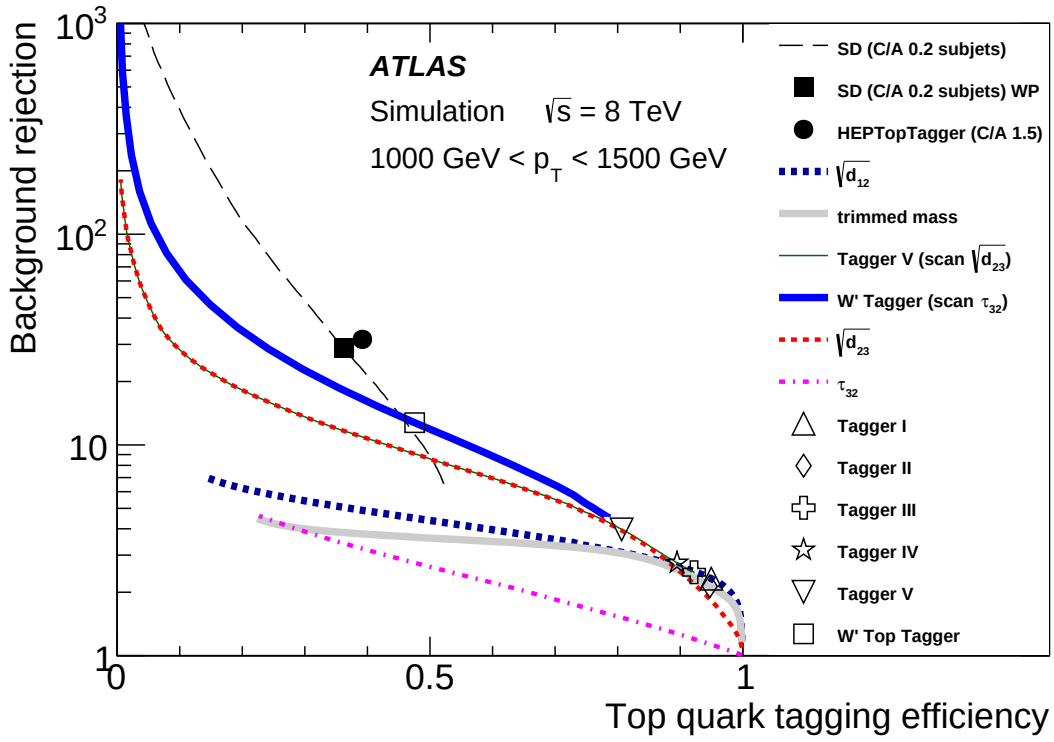

(i) 1000 GeV $< p_T <$ 1500 GeV

Figure 4.3: Comparison of the performance of various top-taggers in simulation, in terms of their boosted top-quark tagging efficiency and background rejection rate.

4.6 Efficiency measurements in data

The tagging efficiency of each tagger under study is defined as the ratio of the number of tagged, truth-matched jets to the total number of truth-matched jets, following the signal selection. When measuring the efficiency in simulation, direct access to the truth-matching information is possible, and so the efficiency is simply defined to be the ratio:

$$f_{\mathrm{MC}} = \left(\frac{N_{\mathrm{MC}}^{\mathrm{tagged}}}{N_{\mathrm{MC}}} \right) \ . \tag{4.6}$$

When measuring the efficiency *in situ*, the contribution from unmatched $t\bar{t}$ and other backgrounds is statistically subtracted from the data directly using simulation, and so the efficiency measured in data is defined to be:

$$f_{\mathrm{data}} = \left(\frac{N_{\mathrm{data}}^{\mathrm{tagged}} - N_{t\bar{t}\text{ not matched}}^{\mathrm{tagged}} - N_{\text{non-}t\bar{t}}^{\mathrm{tagged}}}{N_{\mathrm{data}} - N_{t\bar{t}\text{ not matched}} - N_{\text{non-}t\bar{t}}} \right) \ . \tag{4.7}$$

The tagging efficiency has been measured as a function of both transverse momentum and pseudorapidity. Two η bins, $0.0 < |\eta| < 0.7$ and $0.7 < |\eta| < 2.0$, were designed to have roughly equal statistics. Bins extend in p_{T} between 350 and 600 GeV, until the statistical uncertainty from the 8 TeV dataset limits the measurement. The corresponding measured efficiencies are presented in figures 4.4 and 4.5, and range from $80\% - 95\%$ for tagger I in the central bin to $40\% - 60\%$ for the ATLAS W' tagger in the forward bin.

A breakdown of systematic uncertainties on the efficiency measurement is presented in table 4.3. Dominant sources of experimental systematic uncertainty include the large-R jet energy scale and b-tagging efficiency uncertainties. Larger uncertainties arise due to the modeling of physics processes, particularly within the unmatched component of the $t\bar{t}$ background. All efficiencies measured in data agree with those calculated in simulation, within the stated uncertainties.

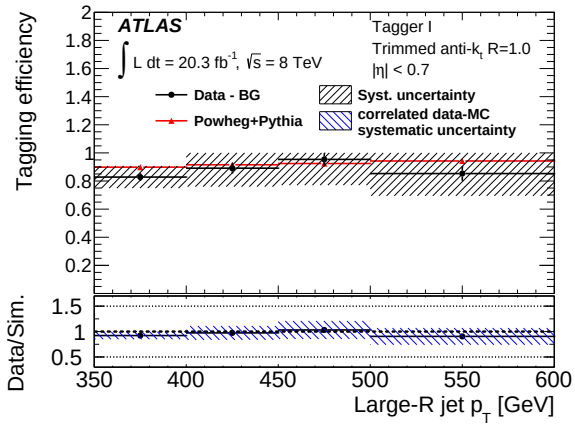

(j) Tagger I

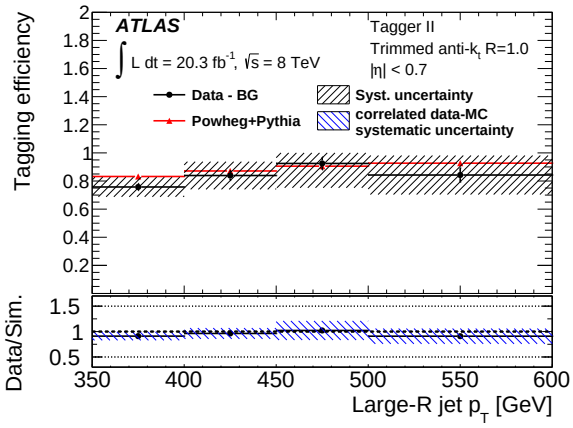

(k) Tagger II

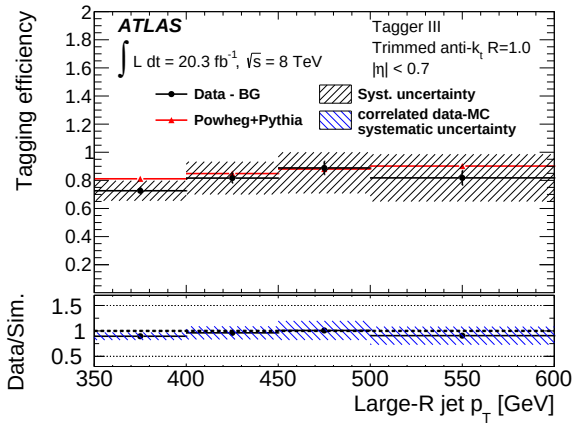

(l) Tagger III

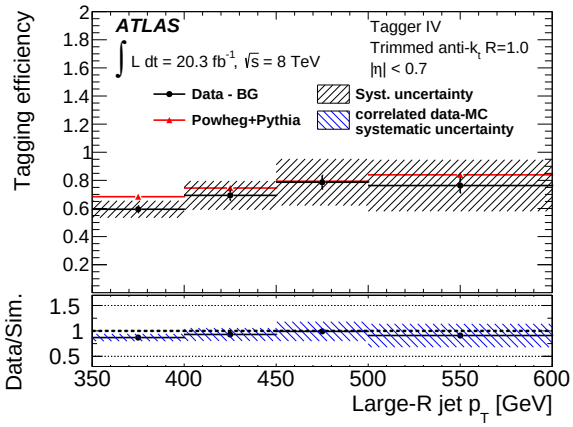

(m) Tagger IV

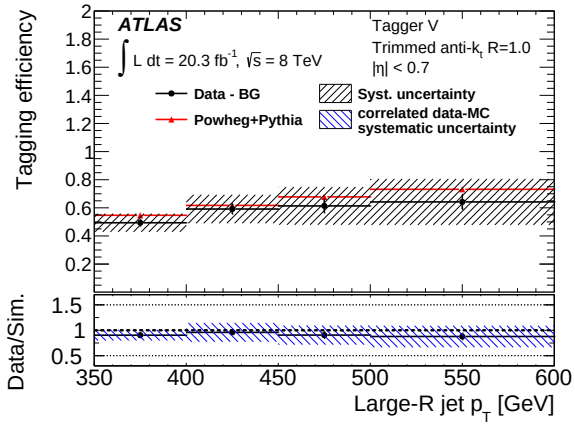

(n) Tagger V

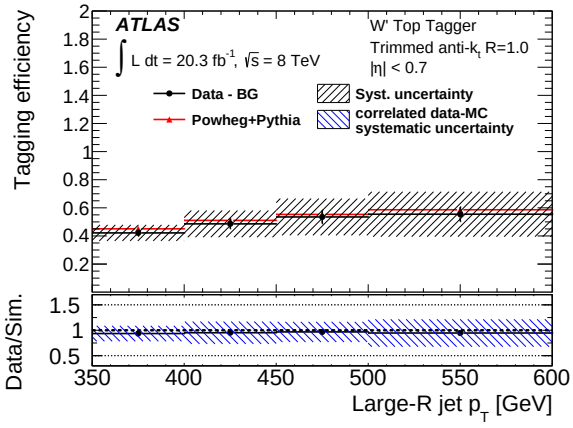

(o) Tagger VI

Figure 4.4: Top quark tagging efficiencies for substructure-based taggers in the central detector region ($|\eta| < 0.8$).

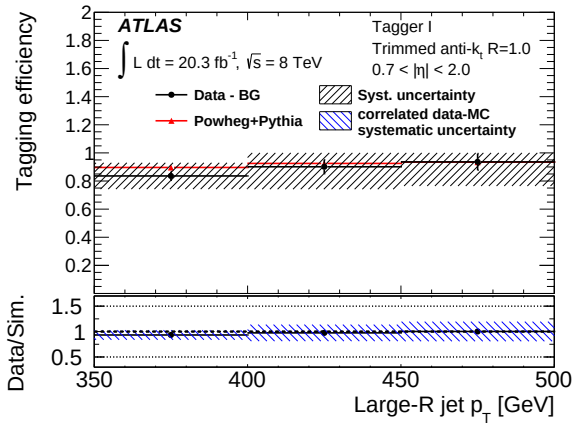

(a) Tagger I

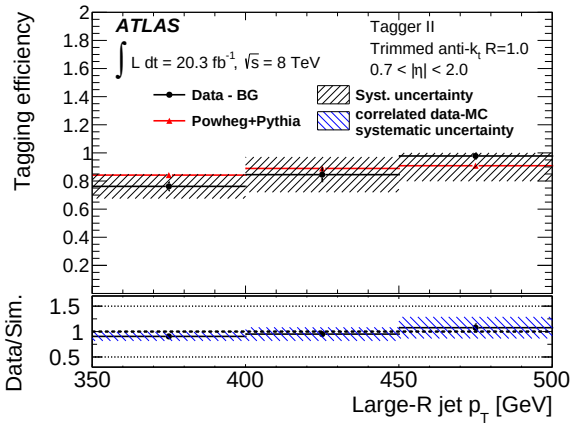

(b) Tagger II

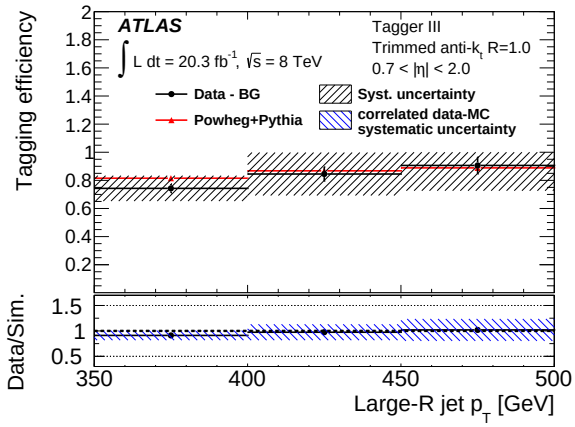

(c) Tagger III

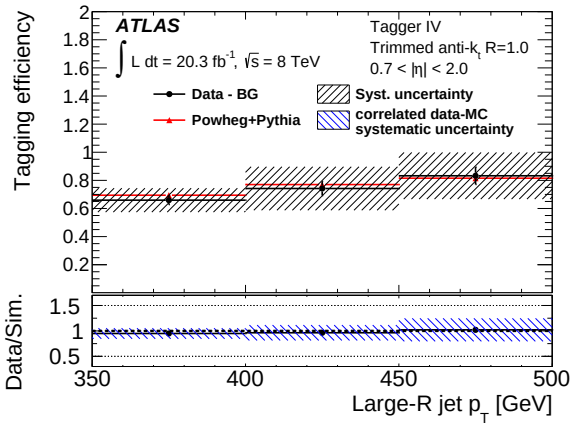

(d) Tagger IV

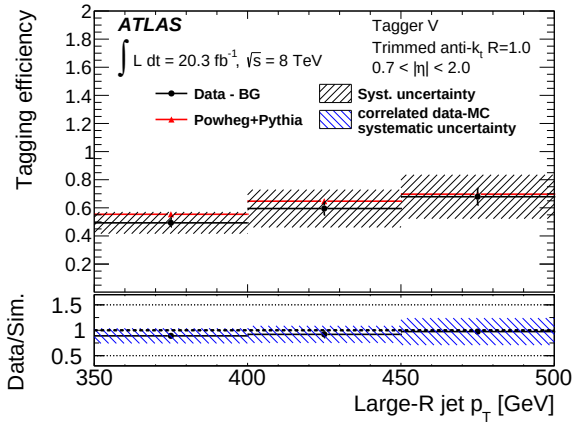

(e) Tagger V

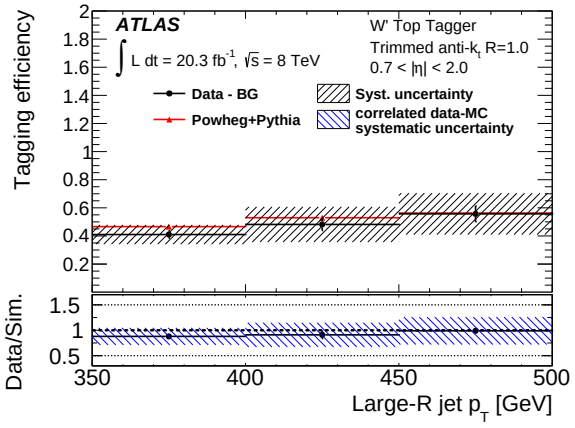

(f) Tagger VI

Figure 4.5: Top quark tagging efficiencies for substructure-based taggers in the forward detector region ($0.8 < |\eta| < 2.0$).

source	relative uncertainty of top tagging efficiency (%)					
	Tagger I	Tagger II	Tagger III	Tagger IV	Tagger V	W′ Tagger
large-R JES	4.4	4.5	4.8	5.3	5.8	6.0
large-R JER	<0.1	0.1	0.1	0.2	0.2	0.3
luminosity	1.0	1.0	1.1	1.2	1.3	1.4
b-tagging efficiency	2.7	2.6	2.9	3.1	3.5	3.7
lepton reco. efficiency	0.5	0.5	0.5	0.6	0.7	0.8
$t\bar{t}$ cross section	1.9	1.8	1.9	2.1	2.4	2.6
$t\bar{t}$ ISR/FSR	1.4	1.3	1.4	0.5	1.6	1.6
$t\bar{t}$ generator	10.3	9.2	10.9	12.0	14.5	16.2
$t\bar{t}$ parton shower	4.8	4.1	4.6	4.8	4.6	5.1
$t\bar{t}$ PDF uncertainty	4.4	3.8	4.5	4.2	5.2	6.8
$t\bar{t}$ renormalisation scale	0.8	0.8	0.8	0.9	1.0	1.1
trimmed large-R JMS	-	1.5	0.8	0.6	0.2	-
trimmed large-R JMR	-	0.1	0.1	0.1	<0.1	-
$\sqrt{d_{12}}$	1.2	-	0.6	0.5	0.4	0.5
$\sqrt{d_{23}}$	-	-	-	0.7	1.1	-
τ_{21}	-	-	-	-	-	0.6
τ_{32}	-	-	-	-	-	1.4
total	13.6	12.4	14.2	15.2	17.9	20.1

Table 4.3: Total systematic uncertainty, in percent, on the inclusive top tagging efficiency measurement.

4.7 Fake rate measurements in data

The background tagging efficiency, or tagging fake rate, is measured using jets arising from gluons and light quarks which are obtained following the background selection outlined in section 4.2.2. The fake rate is defined as the fraction of jets which are tagged by a particular tagger to the total number of jets in the background selection,

$$f_{\text{data}}^{\text{mistag}} = \left(\frac{N_{\text{data}}^{\text{tagged}}}{N_{\text{data}}} \right) \ , \tag{4.8}$$

$$f_{\text{MC}}^{\text{mistag}} = \left(\frac{N_{\text{MC}}^{\text{tagged}}}{N_{\text{MC}}} \right) \ . \tag{4.9}$$

For all substructure-based taggers studied, the background selection contamination due to $t\bar{t}$ events after applying a top tagger is less than 1.6%, which is assumed to be negligible. The measured fake rates are shown in figure 4.6, along with the predicted fake rates from simulation. The highest mistag rates are observed for the least rejective taggers: up to 36% for taggers I and II. The W' tagger provides the tightest substructure-based working point and the lowest fake rate, about 4% in the highest p_{T} bin. While fake rates for the loose taggers may seem large, their primary applications in analyses which have multiple, powerful options for background control (such as searches for supersymmetry with large missing transverse momentum and/or effective mass, see also section 5.3) somewhat mitigates this undesirable quality.

The minimal selection applied to study the mistag rate results in relatively few sources of systematic uncertainty. Uncertainties arise from the luminosity measurement, the large-R jet energy scale and energy resolution, and the substructure variables required by each tagger ($\sqrt{d_{12}}$ for tagger I, the large-R jet mass and mass resolution for tagger II, etc.). Uncertainties on substructure observables are applied in the same fully-correlated way as when performing the tagging efficiency measurement; due to this approach, uncertainties are largest for those taggers which place multiple jet substructure requirements. For tagger V and the W' tagger, systematic uncertainty on the mistag rate reaches respective maxima of $\sim 10\%$ and $\sim 20\%$ in the $600 - 700$ GeV bin.

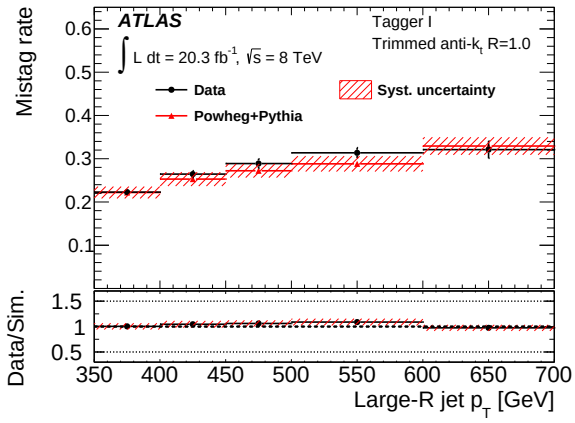

(a) Tagger I

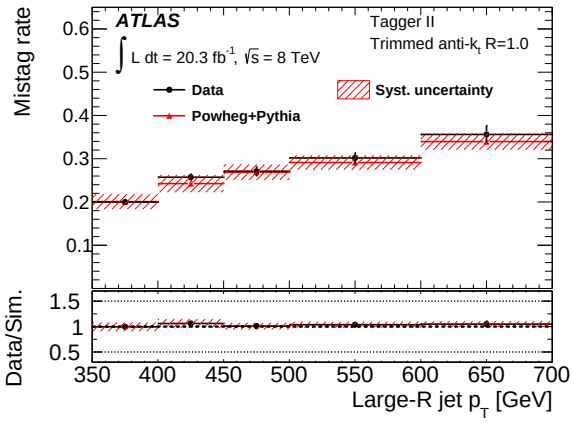

(b) Tagger II

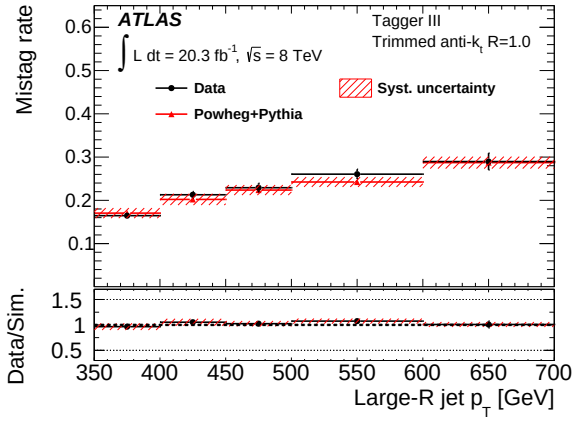

(c) Tagger III

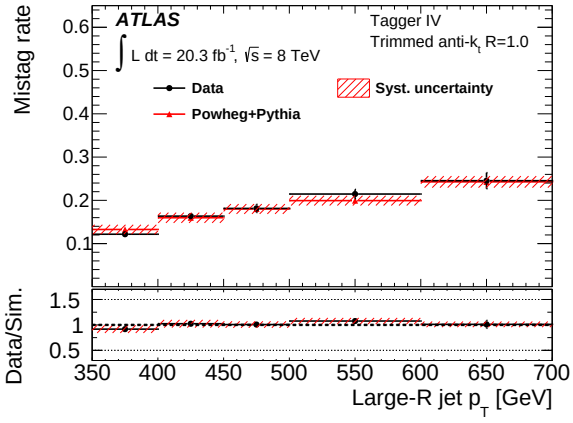

(d) Tagger IV

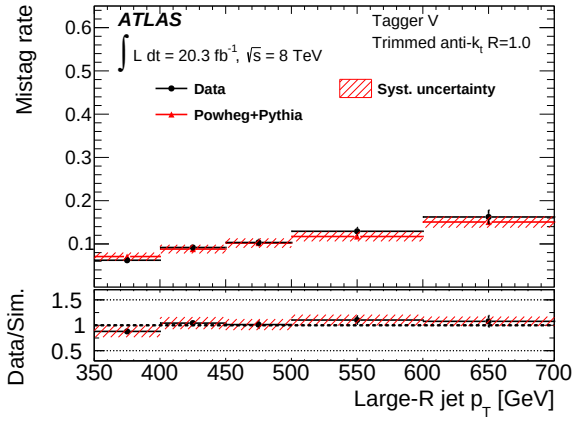

(e) Tagger V

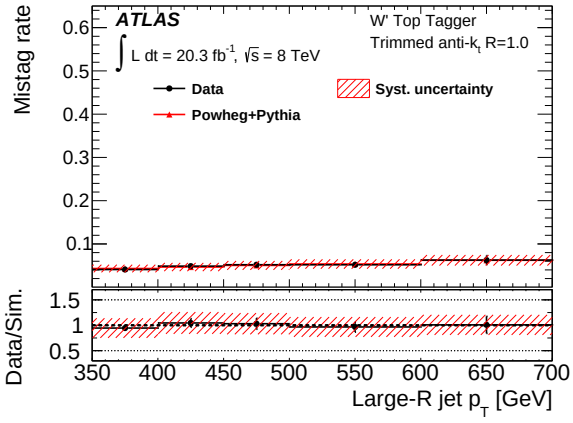

(f) Tagger VI

Figure 4.6: Mistag rates in data and simulation for the substructure-based top taggers, measured over a p_{T} range of 350-700 GeV.

5

Searching for Stop- and Sbottom-Mediated Gluino Production

Searches for natural supersymmetry are among the most well-motivated to perform during the early stages of LHC Run II (chapter 1). The predicted cross-section of the gluino increases dramatically as a function of the collider's centre-of-mass energy, while the cross-section of $t\bar{t}$ production, the dominant background, increases only moderately in comparison. To this end, a search for the pair production of gluinos which decay exclusively through virtual stop and sbottom squarks was performed in the combined 2015 + early 2016 ATLAS dataset, totaling 14.8 fb^{-1} [85]. This search builds on a previous ATLAS run II result which searched for pair-produced gluinos decaying in this manner in the 2015 dataset, corresponding to an integrated luminosity of 3.2 fb^{-1} [86][87].

In complete supersymmetric models, the parameter space is too large to thoroughly scan in every search performed at the LHC. Occasionally, a specific model will be targeted in the context of a given search, though more commonly ATLAS results are optimised for and interpreted using so-called **simplified models** which assume a particular production mechanism and final-state topology. The simplified models illustrated in figure 5.1 are those considered as signal processes for this search. Figure 5.1(a) is referred to as the Gbb model, where all gluinos decay through sbottom squarks into final states with b-jets and missing transverse momentum. Figure 5.1(b) is referred to as the Gtt model, in which all gluinos decay through stop squarks into final state consisting of many top quarks and missing transverse momentum. In both models, the mass of the intermediate squark is assumed to be off-shell[1], effectively resulting in three-body gluino decays and leaving the only free parameters in these simplified models to be the mass of the gluino and $\tilde{\chi}_1^0$.

[1] $m_{\tilde{t}} = 5$ TeV.

Very high missing transverse momentum and effective mass selections are hallmarks of searches for inclusive gluino production, though this analysis is distinct from others performed by ATLAS due to its requirements of a high b-tagged jet multiplicity (> 3) and significant event substructure (section 5.2). Two signal regions are defined which target Gbb gluino decays, and five target Gtt decays – two searching for final states with zero leptons, and three for final states with one lepton or more.

This search is designed such that the dominant background in all signal regions is Standard Model $t\bar{t}$ production. This background is normalised using a semi-data-driven technique which extrapolates the background estimate from a dedicated control region for each signal region using a profile likelihood fit. Extrapolations over the $t\bar{t}$ background flavour composition are minimised, in order to avoid complications due to the potential mismodelling of the heavy flavour content in simulation: extrapolations over the number of b-tagged jets are carefully avoided, but also the total flavour composition of the $t\bar{t}$ background is held as constant as possible across each control-signal region grouping. The performance of this background extrapolation is vetted in dedicated validation regions for each control-signal region pairing, whose flavour content is also designed to be as close as possible to that of the signal region.

No statistically significant excess is observed in any signal region, and so limits are placed on simplified models in the gluino-neutralino mass plane, and model-independent upper limits on the gluino production cross section are calculated (chapter 7).

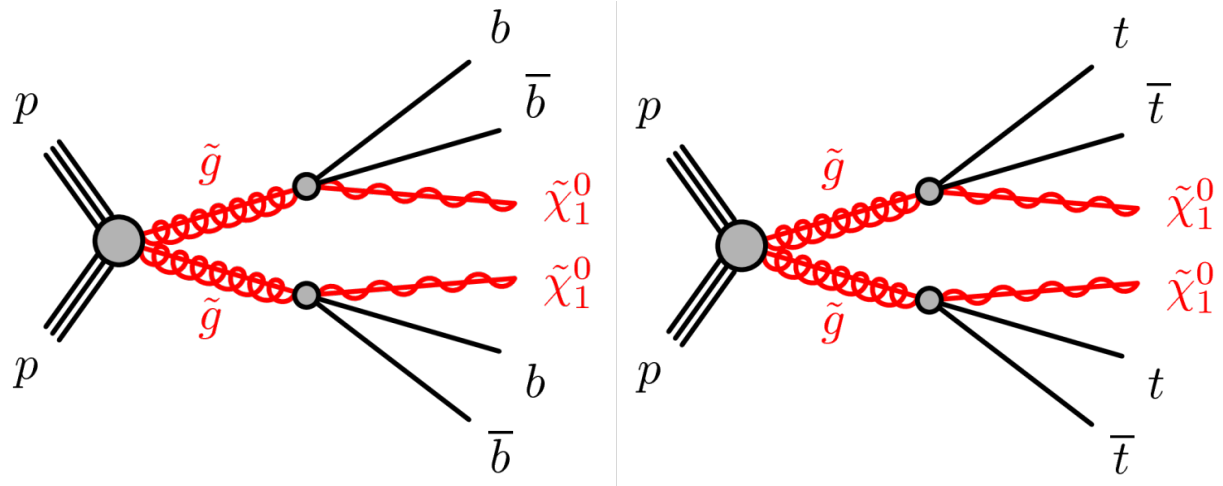

(a) (b)

Figure 5.1: Feynman diagrams representing simplified (a) sbottom- and (b) stop-mediated gluino pair-production models under consideration.

5.1 Data and Simulation

Data collected by the LHC during $\sqrt{s}$ =13 TeVoperations is used in this search. A dataset of 3.2 fb^{-1} collected between August and December 2015 is combined with 12.6 fb^{-1} collected in 2016 before 16 July, resulting in a total dataset of 14.8 fb^{-1} for analysis. This dataset has been collected with an LHC bunch crossing time of 25 ns, and events within it include an average of 14-21 additional pp interactions each.

The simplified Gtt and Gbb models are simulated at leading order with two additional final-state partons using MadGraph5_aMC@NLO v2.2.2, with CTEQ6L1 parton density functions. Generated events are passed to Pythia 8 for showering and hadronisation. A grid of signal points has been generated for both the Gbb and Gtt models, respectively decaying through off-shell sbottom- or stop-loops (the mass of the third generation squark has been set to 5 TeV in both cases). These signal grids are produced as a function of the mass of the gluino and $\tilde{\chi}_1^0$, and provide a means of both optimising the search and producing exclusion limits within this parameter space in the event of a null result.

Top quark pair production is simulated at next-to-leading order using Powheg-Box, linked to Pythia 8 to process showering and hadronisation. This sample includes the CT10 PDF set, as do all others unless otherwise noted. The production of top quark pairs in association with electroweak vector bosons forms a smaller background for this search, simulated with MadGraph5_aMC@NLO and showered in Pythia 8. The SM production of four top quarks is also simulated and showered using the same generator configuration. Top quark pairs produced in association with a Higgs boson are generated in MadGraph5_aMC@NLO and showered in Herwig++. Single-top production in the $Wt-$ and s-channels is generated in Powheg-Box interfaced with Pythia 8, though t-channel single top production is instead showered with Pythia 6 using a four-flavour scheme and appropriately modified PDFs (CT10f4). The production of electroweak vector bosons associated with extra jets ($V + j$) is simulated in Sherpa 2.2, though the production of pairs of electroweak bosons (VV) is simulated with Sherpa 2.1; both of these processes include CT10 parton distribution functions.

All simulation except the Gbb signal grid is reconstructed using a complete ATLAS detector simulation implemented in Geant 4. The Gbb grid is instead reconstructed using a faster simulation which parameterises the response of the calorimeters to showering particles [88, 89], an approach which is more computationally efficient and applicable when no jet substructure techniques are considered.

A summary of the various MC samples used to estimate the signal, background and uncertainties for this search is provided in table 5.1.

Process	Generator	PDF	Parton Shower
$t\bar{t}$ (nominal)	POWHEG-BOX 2	CT10	PYTHIA 8
single top (Wt-, s-channel)	POWHEG-BOX 2	CT10	PYTHIA 8
single top (t-channel)	POWHEG-BOX 2	CT10f4	PYTHIA 6
V+jets	SHERPA 2.2	CT10	SHERPA 2.2
VV	SHERPA 2.1	CT10	SHERPA 2.1
$4t$	MADGRAPH 5 + aMC@NLO 2	NNPDF 2.3	PYTHIA 8
$t\bar{t}+V$	MADGRAPH 5 + aMC@NLO 2	NNPDF 2.3	PYTHIA 8
$t\bar{t}+H$	MADGRAPH 5 + aMC@NLO 2	CT10	Herwig++
Gtt	MADGRAPH 5 + aMC@NLO 2	CTEQ6L1	PYTHIA 8
Gbb	MADGRAPH 5 + aMC@NLO 2	CTEQ6L1	PYTHIA 8
$t\bar{t}$ ($\sigma^{\text{ISR/FSR}}$)	POWHEG-BOX 2	CT10	PYTHIA 6*
$t\bar{t}$ ($\sigma^{\text{MC Gen.}}$)	MADGRAPH 5 + aMC@NLO 2	CT10	Herwig++
$t\bar{t}$ ($\sigma^{\text{Parton Shower}}$)	POWHEG	CT10	Herwig++
single top ($\sigma^{\text{ISR/FSR}}$)	POWHEG-BOX 2	CT10	PYTHIA 6*

Table 5.1: A summary of the Monte Carlo samples used to search for stop- and sbottom-mediated gluino pair production during ATLAS run II. Samples used to evaluate the $\sigma^{\text{ISR/FSR}}$ uncertainty are showered using alternative PYTHIA 6 tunes, Perugia2012radHi & Perugia2012radLo.

5.2 Object Definitions

Jets used in this analysis are built from three-dimensional topological clusters locally calibrated at the electromagnetic scale using the anti-k_t algorithm with distance parameter $R = 0.4$. The energy scale and resolution of these jets are calibrated to particle level using the EM+JES+GSC scheme and residual *in-situ* techniques outlined in section 3.2.2. Jets considered in the overlap removal procedure outlined below are required to possess $p_{\mathrm{T}} > 20$ GeV and $|\eta| < 2.8$, and must additionally satisfy a requirement on the fraction of their energy which is associated with the event's primary vertex if the jet's p_{T} is below 60 GeV in order to suppress jets originating from secondary pp interactions during the event [90]. Jets may be b-tagged using the MV2c10 algorithm [53, 54] at the 77% b-tagging efficiency working point. Those jets which remain following the overlap removal must subsequently satisfy a cut of $p_{\mathrm{T}} > 30$ GeV in order to be selected for use in the analysis.

Selected jets are in-turn used as inputs to the anti-k_t algorithm with distance parameter $R = 0.8$, in order to produce large-R **reclustered jets** (appendix B). These jets are trimmed with an f_{cut} parameter of 10%, and the trimmed jets are required to satisfy $p_{\mathrm{T}} > 100$ GeVand $|\eta| < 2.0$. Previously used for boosted top identification in searches for gluino pair production [87], these reclustered jets are now used in the construction of a continuous, structure-sensitive observable newly introduced for this iteration of the search, the total reclustered jet mass ($M_J^{\sum,4}$, section 5.3). The reclustering parameters selected for this analysis were optimised in the context of its final discovery sensitivity, from choices of $R_{\mathrm{sub}} = 0.8,\ 1.0,\ 1.2$ and $f_{\mathrm{cut}} = 5\%,\ 10\%,\ 15\%$.

Electrons are initially required to satisfy the `loose` likelihood identification criteria, be associated with a tracking hit in the pixel b-layer, and have $p_{\mathrm{T}} > 20$ GeV within $|\eta| < 2.47$. This selection, defined as the *baseline electron* selection, is used to perform object overlap removal. Following the overlap removal procedure, remaining electrons must satisfy the `medium` likelihood identification criteria and `looseTrackOnly` electron isolation criteria [91], as well as additional requirements on its impact parameters: $|z_0 \sin(\theta) < 0.5|$ mm and $d_0/\sigma(d_0) < 5$. These final electrons are defined as the *signal electrons* in the context of the analysis.

Muons must satisfy the `medium` identification working point [92] and have $p_{\mathrm{T}} > 20$ GeV within $|\eta| < 2.5$ to be retained as a *baseline muon* object in the search, which are used in the overlap removal procedure. After the overlap removal has been applied, remaining muons must pass additional requirements imposed on their impact parameters $|z_0 \sin(\theta) < 0.5|$ mm and $d_0/\sigma(d_0) < 3$ in order to be included in the search as a *signal muon.*

Energy depositions in the ATLAS detector may be reconstructed as more than a single

physics object. For this reason, it is important to define an **overlap removal** procedure which will ensure that objects of good quality are included in analysis, and that the double-counting of energy does not occur. The overlap removal procedure applied by this search has been updated since the previous result [87] in order to better reflect the nature of the highly energetic final-state topologies which the signal could possess:

1. Electrons and muons which share tracks are compared. If the muon is a calo-tagged muon (section 3.5), the muon is removed and the electron retained. If not, the electron is removed and the muon retained.

2. Leptons and jets within $\Delta R < 0.2$ are compared. If the jet is b-tagged, it is retained and the lepton is removed. Otherwise, the lepton is retained and the jet is removed.

3. Electrons and jets are compared. If any pair satisfies $\Delta R < \min(0.4, 0.04{+}10\ \text{GeV}/p_T^e)$, the electron is removed and the jet is retained. This criteria is designed to provide better efficiency for retaining electrons in boosted topologies, where they may be located closer to a b-jet in a leptonically decaying top than in other final states.

4. Muons and jets are compared. For pair satisfying $\Delta R < \min(0.4, 0.04 + 10\ \text{GeV}/p_T^\mu)$, the jet is removed and the muon is retained *if* the jet is associated with less than three inner detector tracks of 500 MeV or more. Otherwise, if the distance criteria is satisfied and the jet is associated with more than three ID tracks, the muon is rejected. The cap of $\Delta R = 0.4$ introduced in this step allows for muons in less energetic topologies to be retained with greater efficiency, as the p_{T}-dependent measure alone becomes very large for soft leptons.

The **missing transverse momentum** ($E_{\text{T}}^{\text{miss}}$) is calculated using the prescription for run II (section 3.7) with the calibrated jets, electrons and muons described above. The jet collection before overlap removal is used in the calculation, while the leptons are used post-overlap-removal. An additional term is included which accounts for energy in the event which is not reconstructed as a jet or lepton, based on tracking information available from the inner detector (the *track-based soft term*, or TST).

5.3 Event-Level Observables

The **missing transverse momentum** ($E_{\text{T}}^{\text{miss}}$) is a cannonical observable used in searches for R-parity conserving (RPC) supersymmetry. In such cases, the **lightest supersymmetric particle** (LSP) is neutral and stable, traversing ATLAS without interacting with any

detector subsystems. The existence of these neutral particles creates an imbalance in the observed momentum transverse to the beamline, and their presence can be inferred through this imbalance. The $E_{\mathrm{T}}^{\mathrm{miss}}$ spectrum following the analysis pre-selection is shown in figure 5.3(a).

In fully-hadronic final states, missing transverse momentum may be artificially induced in the event due to fluctuations in the calorimeter response when measuring the energy of jets. In such cases, this 'fake' $E_{\mathrm{T}}^{\mathrm{miss}}$ tends to align itself along the axis of a fluctuating jet, and so these events may be removed by requiring that the $E_{\mathrm{T}}^{\mathrm{miss}}$ is always well-separated from hadronic objects in the event. The measure $\Delta\phi_{\mathrm{min}}^{4j}$ is defined as:

$$\Delta\phi_{\mathrm{min}}^{4j} = \min(|\phi_{\mathrm{jet}_1} - \phi_{\vec{p_{\mathrm{T}}}^{\mathrm{miss}}}|, ..., |\phi_{\mathrm{jet}_4} - \phi_{\vec{p_{\mathrm{T}}}^{\mathrm{miss}}}|), \tag{5.1}$$

which is required to satisfy $\Delta\phi_{\mathrm{min}}^{4j} > 0.4$ in any selection where no leptons are present. This requirement greatly reduces the amount of multijet backgrounds to this search, which is subsequently rendered negligible following b-tagging requirements.

The inclusive **effective mass** ($\mathrm{m}_{\mathrm{eff}}^{\mathrm{incl}}$) is defined as the scalar sum of the p_{T} of all signal jets and leptons present in the event, and the $E_{\mathrm{T}}^{\mathrm{miss}}$. An exclusive variant ($\mathrm{m}_{\mathrm{eff}}^{\mathrm{4j}}$) which sums only the p_{T} of the four leading jets and the $E_{\mathrm{T}}^{\mathrm{miss}}$ is defined for the Gbb signal regions, as no additional jets and no leptons are expected from the signal in this case. The inclusive effective mass distribution is shown in figure 5.3(b) following the analysis pre-selection.

The **total reclustered jet mass** ($M_J^{\sum,4}$) leverages information about jet substructure present in the stop-mediated signal channel due to the large top quark multiplicity. Defined as the scalar sum of the jet masses of the leading four reclustered jets, this variable allows requirements to be placed on the structure of an event in a continuous way, unlike traditional top-tagging approaches. In addition to the cut made on this observable, both the large-R jet radius and the trimming parameter used in the reclustering procedure may be optimised for the search, providing a large amount of flexibility. Figure 5.2 shows the $M_J^{\sum,4}$ distribution for a signal point in the high-mass-splitting region of phase space, reconstructed with R values ranging from 0.8-1.2 and f_{cut} values of 0.05 or 0.10. For this search, a large-R jet radius of $R = 0.8$ and a trimming parameter of $f_{\mathrm{cut}} = 10\%$ were found to be optimal, by re-performing the initial analysis optimisation procedure (chapter 6) several times using various reclustering parameters. The $M_J^{\sum,4}$ distribution, when constructed with these optimal parameters, is shown in figure 5.4(d).

The **transverse mass** of an event with at least one lepton provides a way to differentiate W+jets events from other background processes. It is defined by relating the lepton p_{T} to the missing transverse momentum as

$$m_T = \sqrt{2p_T^{\ell}E_{\mathrm{T}}^{\mathrm{miss}}\{1 - \cos[\Delta\phi(E_{\mathrm{T}}^{\mathrm{miss}}, \ell)]\}} \tag{5.2}$$

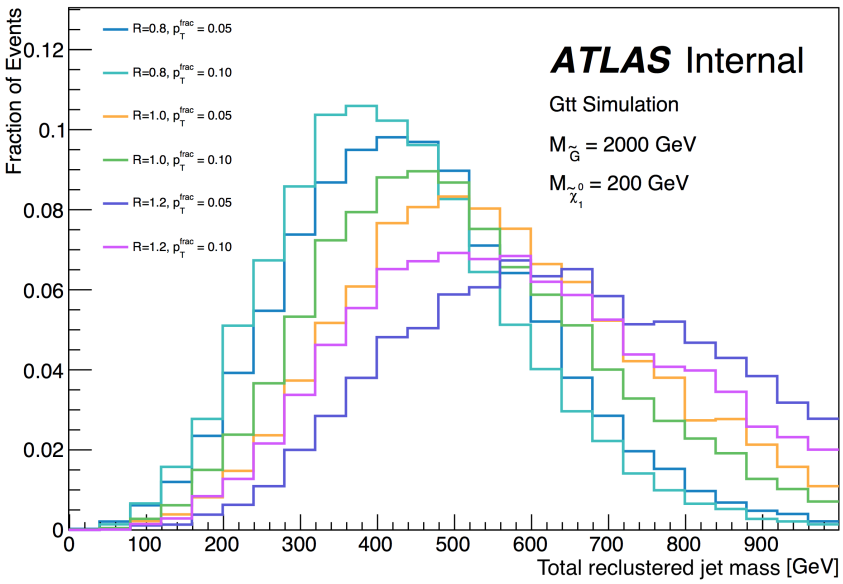

Figure 5.2: The $M_J^{\sum,4}$ distribution for a high-mass-splitting signal point ($M_{\tilde{G}} = 2$ TeV, $M_{\tilde{\chi}_1^0} = 200$ GeV), reconstructed with various choices of anti-$k_t$distance parameter R and trimming threshold f_{cut}.

A related variable which is useful for selecting semi-leptonic $t\bar{t}$ events from other backgrounds is the $m_{T,\text{min}}^{b-\text{jets}}$, which is defined analogously to the classical m_T but using the b-tagged jet which gives the lowest value of the observable in place of the lepton:

$$m_{T,\text{min}}^{b-\text{jets}} = \min_{b-\text{jets}} \left(\sqrt{2p_T^{b-\text{jet}} E_{\text{T}}^{\text{miss}} \{1 - \cos[\Delta\phi(E_{\text{T}}^{\text{miss}}, b-\text{jet})]\}} \right) \tag{5.3}$$

The m_T and $m_{T,\text{min}}^{b-\text{jets}}$ distributions following the analysis pre-selection are shown in figures 5.3(d) and 5.3(c), respectively. These variables provide a powerful means by which the background composition of a selection may be controlled, allowing for the design of the highly-$t\bar{t}$-enriched regions upon which this search depends to be defined.

A pre-selection for this analysis has been defined, in order to examine the modeling of these selected observables. At least 200 GeVof $E_{\text{T}}^{\text{miss}}$, and $\Delta\phi_{\text{min}}^{4j} > 0.4$ are required, along with the presence of at least four signal jets, two of which jets must be b-tagged. The distributions of event-level observables $E_{\text{T}}^{\text{miss}}$, m_{eff}, m_{T} and $\text{m}_{\text{T,min}}^{b-\text{jets}}$ are shown in figure 5.3, and are well-modelled in the regions of interest for this analysis. Object-related quantities such as the jet multiplicity, b-jet multiplicity, signal lepton multiplicity and $M_J^{\sum,4}$ distribution are shown in figure 5.4. Possible mis-modeling of the heavy flavour content in the $t\bar{t}$ simulation leads to the trend observed between data and MC in the b-jet multiplicity distribution, motivating the flavour-content-aware background estimation approach applied in this search.

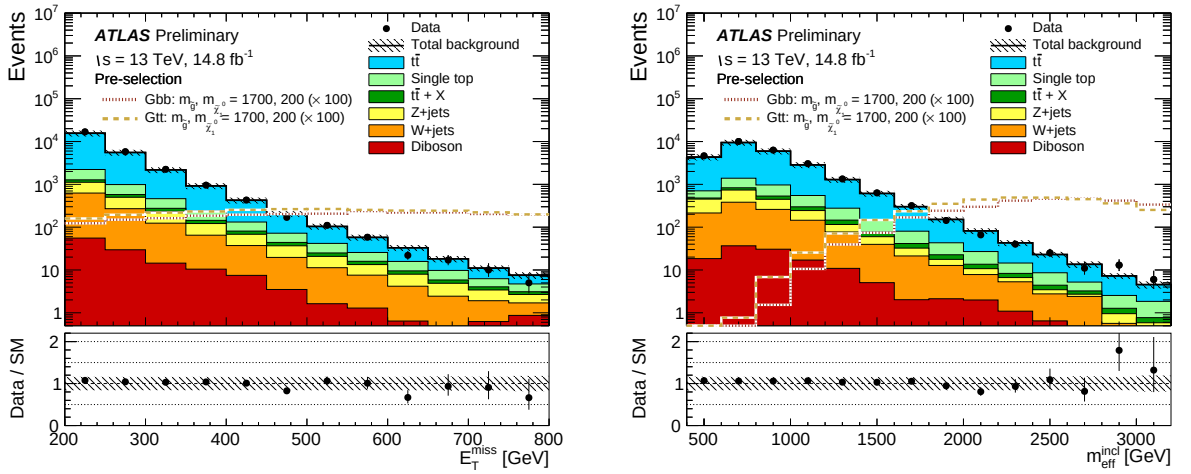

(a) (b)

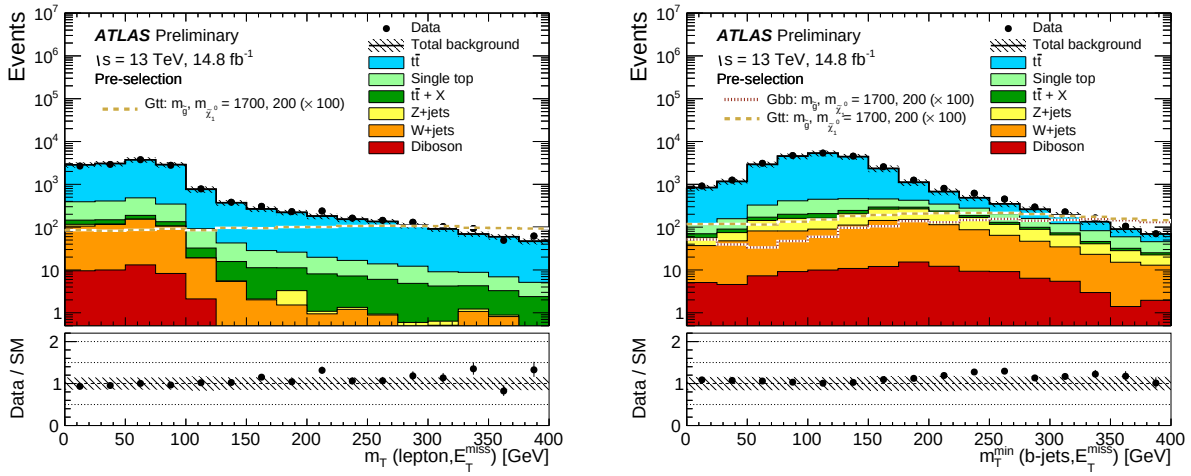

(c) (d)

Figure 5.3: Distributions of (a) E_T^{miss}, (b) $\text{m}_{\text{eff}}^{\text{incl}}$, (c) m_T^W and (d) $m_T^{b,\text{min}}$ following the analysis pre-selection.

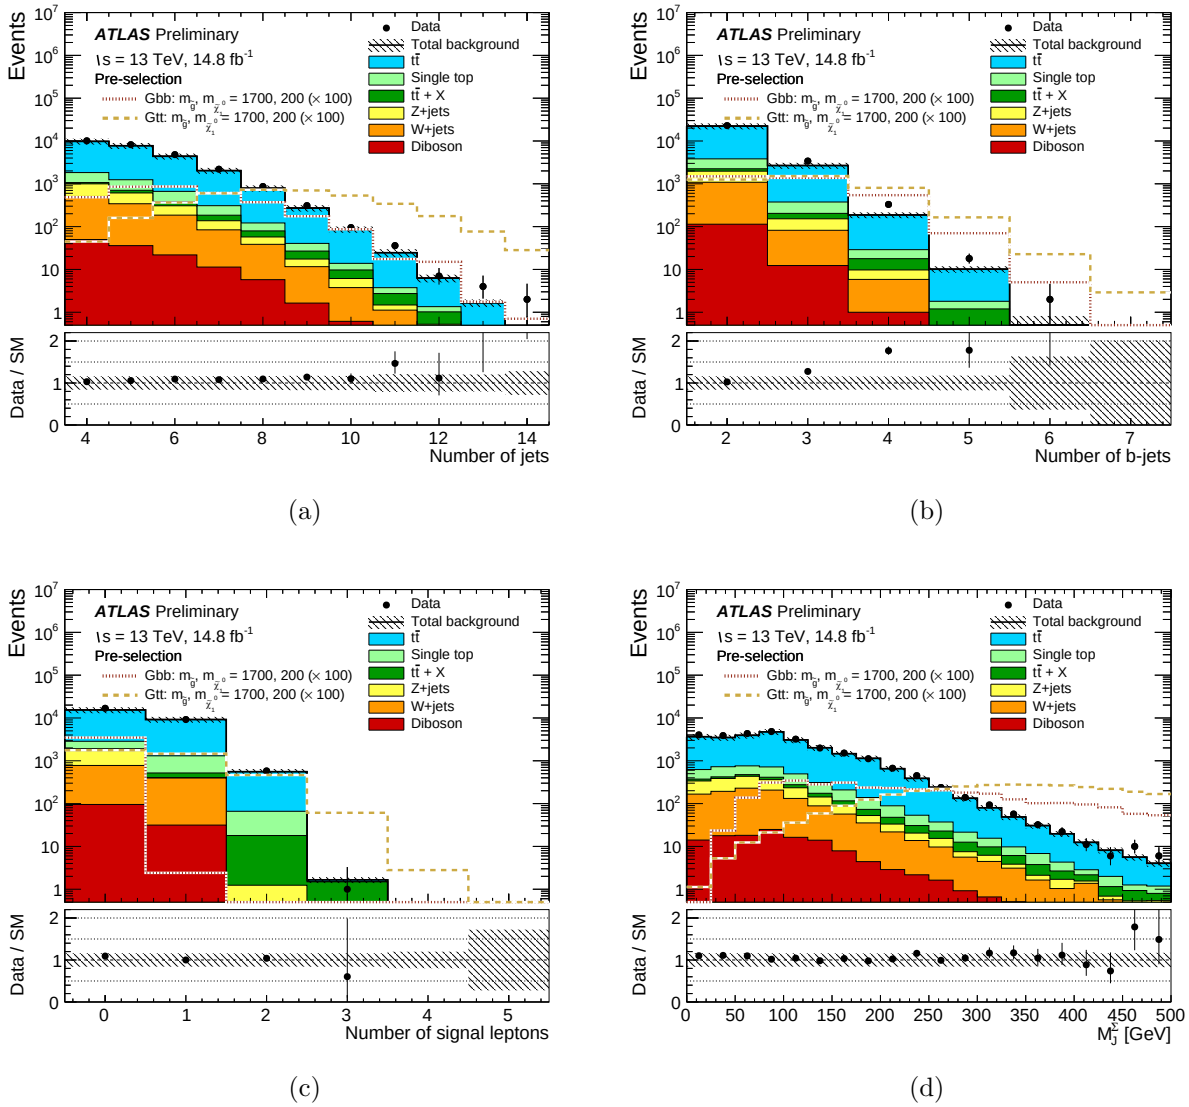

(a) (b)

(c) (d)

Figure 5.4: (a) Jet, (b) b-tag and (c) signal lepton multiplicities following analysis preselection. (d) $M_J^{\Sigma,4}$ distribution following analysis pre-selection.

5.4 Systematic Uncertainties

Systematic uncertainties which are considered in the context of this analysis may be divided into two categories: experimental systematic effects caused by the response and resolution of the ATLAS detector, and theoretical systematic effects which arise due to the theoretical understanding of the various physics processes. The methods used to calculate some systematics has not greatly changed between run I and II: those described previously in section 4.4 will be only briefly mentioned again.

5.4.1 Experimental Uncertainties

The jet energy scale and resolution of the anti-$k_t R = 0.4$ jets used in this analysis contribute two of its largest sources of systematic uncertainties. These uncertainties are determined using the *in-situ* techniques described in section 3.2.2, which result in a set of 56 potential systematic variations which are combined into a reduced set of 17 nuisance parameters. This set is then *strongly reduced*, resulting in four sets of three parameters each, which encode different information about the correlations between the full set of parameters [93]. This analysis was not found to be sensitive to the choice of the strongly reduced parameter set. The default choice was applied, resulting in a 13-25% uncertianty on the background yield as a result of the JES uncertainty. An additional uncertainty of between 6-16% on background yields results from the JER systematic, based on measurements of the resolution at 8 TeV which have been extrapolated for used in run II based on simulation [94].

Scale factors to correct the b-tagging efficiency in simulation are derived from di-leptonic $t\bar{t}$ events as a function of the jet p_{T}, η and truth-level flavour. An additional scale factor accounts for the extrapolation of these SFs to high p_{T} jets. These scale factors are varied in order to provide an estimate of the uncertainty in the tagging efficiency and fake rate measurements on the background yield of this analysis, resulting in an overall uncertainty of between 10-20%.

Due to the application of jet reclustering to construct large-R jets for the calculation of the $M_J^{\sum,4}$, this search is potentially sensitive to the mass scale uncertainty of anti-k_t $R = 0.4$ jets. In order to survey the impact of this uncertainty on the analysis, the final background-only and exclusion fits (chapter 7) were repeated four times, with different levels of smearing between 5 – 20% applied to the masses of the input jets. These variations caused negligible changes in the sensitivity of the analysis, and so this uncertainty was not considered in the final results.

Uncertainties on other physics objects in the event are propagated to a recalculation of the $E_{\mathrm{T}}^{\mathrm{miss}}$, and additional variations of the track soft term are included. Even combined, the

$E_{\mathrm{T}}^{\mathrm{miss}}$ systematic variations have a negligible impact on the analysis sensitivity, and are not included in the final fits.

The total integrated luminosity collected by ATLAS during run II has an associated uncertainty of 3%, resulting in a negligible uncertainty on the background in this analysis.

5.4.2 Theoretical Uncertainties

Several important sources of uncertainty arise from the modelling of the $t\bar{t}$ + heavy flavour processes which form the main background of this search. Uncertainties arising from the simulation of the $t\bar{t}$ process itself are evaluated by comparing the nominal Powheg-Box + Pythia 6 sample to other generator and parton shower configurations. All modelling uncertainties are implemented in the fitting procedure as a relative uncertainty on the control $\rightarrow$ signal or control $\rightarrow$ validation region transfer factor, and so only differences in the total yield in these regions leads to an uncertainty on the final result.

Variations of the nominal $t\bar{t}$ sample were produced using the `Perugia2012RadHi` and `Perugia2012RadLo` Pythia tunes [95], which influence the amount of ISR/FSR radiation which is produced in association with the top pair. Other Pythia parameters were also altered when producing these samples, leading to additionally decreased or increased radiation. Systematic uncertainty resulting from the choice of parton showering algorithm is evaluated by comparing the nominal sample to one generated with Powheg-Box but showered in Herwig++. Uncertainty arising due to the choice of the MC generator itself was evaluated by comparing the nominal sample to one generated with aMC@NLO but showered with the nominal Pythia 6 configuration.

Based on the result of the ATLAS run 1 measurement of the production cross section for $t\bar{t}$ in association with heavy flavour jets [96], an uncertainty of 30% is taken on the fraction of $t\bar{t}$ + heavy flavour ($t\bar{t}$ $+b\bar{b}$ and $t\bar{t}$ $+c\bar{c}$) events observed in simulation following the event selection applied by this analysis.

Modelling uncertainties effecting the subdominant single top background are also considered. A generator comparison is performed to evaluate the impact of radiation parameters on the single top simulation, comparing variations of the nominal Powheg-Box +Pythia 6 sample re-showered with the `Perugia2012RadHi` and `Perugia2012RadLo` Pythia tunes and altered parton shower parameters [95]. The impact of theoretical uncertainty arising due to interference effects between single top production in the Wt channel and $t\bar{t}$ production is evaluated using Powheg-Box +Pythia 6 single top samples which implement both the interference effects using either diagram subtraction (DS) or diagram removal (DR) schemes [69]. As the contribution from single top events is small, these uncertainties have only a minor impact on the ultimate analysis sensitivity.

6

Optimisation

6.1 Signal Regions

The figure of merit used to optimize the sensitivity of this search to the Gbb and Gtt simplified models which will be considered is the expected significance, based on the `BinomialExpZ` function of the `RooStats` statistical package [97].

In the initial optimisation of this search, all available signal points were independently optimised to find the best combination of cuts at each signal point. Optimisations for the Gbb, Gtt-0l and Gtt-1l channels were carried out separately, and the Gtt-1l channel was further decomposed into 3- and 4-b-tag channels (Gtt-1l-3b and Gtt-1l-4b) in order to more specifically target the more challenging region of phase space with small gluino-neutralino mass splitting near the diagonal kinematic boundary. An integrated luminosity of 10 fb^{-1} was assumed for this fine-grained optimisation, the amount initially forecast for inclusion in ATLAS summer results during 2016. The initial scan over observables for use in the search is summarised in table 6.1. Some variables were not included in the optimisation scan due to their ability to change the background composition of a given selection: the jet multiplicity, $\mathrm{m_T}$ and $\mathrm{m_{T,min}^{b-jets}}$ are correlated with $t\bar{t}$ composition, and so were left fixed at values known to preferentially suppress all other background processes. A flat 30% background uncertainty was included for all significance calculations.

Figures 6.1 and 6.2 show the optimal cut value for each scanned variable, for each signal point, in the Gtt-0l and Gtt-1l channels. Higher cuts on $E_{\mathrm{T}}^{\mathrm{miss}}$ and $\mathrm{m_{eff}}$ are favoured as the mass splitting between the gluino and $\tilde{\chi}_1^0$ increases. The $M_J^{\Sigma,4}$ adds additional sensitivity in the region of very high mass splitting, where topologies are most likely to contain boosted top quarks, though it also performs well across the bulk of the Gtt-0l phase space. The significances corresponding to these optimal cut values are shown for each signal point in figures 6.3. Good sensitivity is possible in this ideal case up to gluino masses of nearly 2 TeV for

massless $\tilde{\chi}_1^0$ across both the 0l and 1l channels. The additional Gtt-1l-4b optimisation provides better sensitivity in the softer regions of phase space, providing one straightforward option to target this otherwise difficult region.

Observable	**Range (Inclusive)**	**Stepsize**
N_ℓ (Signal)	0 or ≥ 1	–
$\mathrm{m}_{\mathrm{T,min}}^{b-\mathrm{jets}}$	≥ 80 GeV	–
$\Delta\phi_{\mathrm{min}}^{4j}$	≥ 0.4	–
$E_{\mathrm{T}}^{\mathrm{miss}}$	[200, 600] GeV	50 GeV
$\mathrm{m}_{\mathrm{eff}}^{\mathrm{incl}}$	[0, 2500] GeV	500 GeV
$M_J^{\sum,4}$	[0, 600] GeV	50 GeV
N_{jets}	≥ 8	–
$N_{\mathrm{b-jets}}$	[3, 4] or (3 or 4)	1 or –

Table 6.1: Summary of the optimisation scan performed targetting Gtt models in the 0- and 1-lepton channels of the search for gluino pair production.

Using this fine-grained optimisation as guidance, several signal regions were defined which aim to maintain as much of the initial optimisation's sensitivity to the Gtt phase space as possible. It is preferable to cover the phase space using the fewest possible signal regions, in order to simplify the later background estimation and combination processes as much as possible. Two regions targeting the Gtt-0l channel and three targeting the Gtt-1l channel are defined, summarised in table 6.2. Their performance is illustrated in figures 6.4 and 6.5, where the optimal signal region at each point and its corresponding significance are shown. The performance degradation incurred by moving to discrete regions rather than using the optimal set of cuts for each signal point is observed to be small, except for some signal points in the diagonal region – suggesting that more tailored approaches to target this phase space may lead to better future sensitivity.

Signal regions for the Gbb simplified models were defined in the same manner as those targeting Gtt models, though these details have been omitted here. The signal regions established by the Gbb optimisation are summarised in table 6.2.

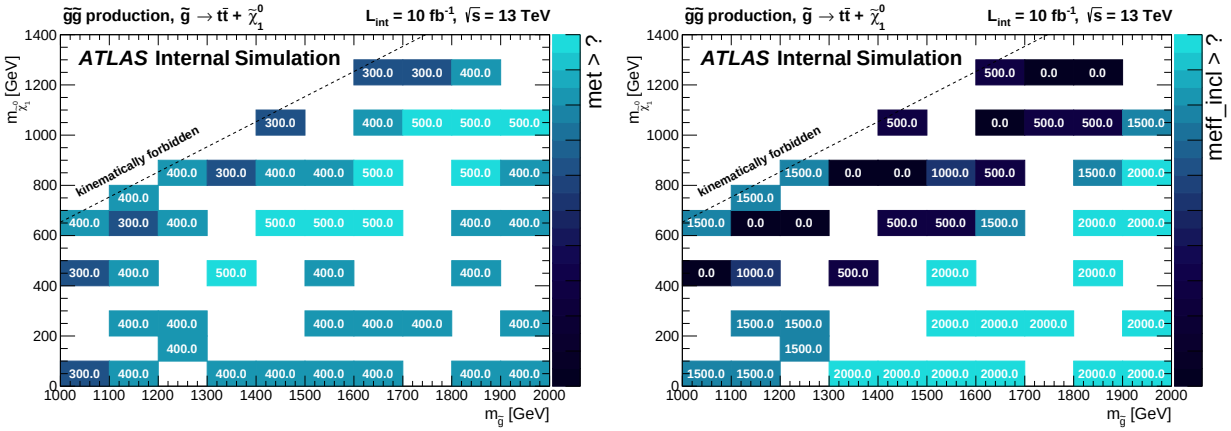

(a)

(b)

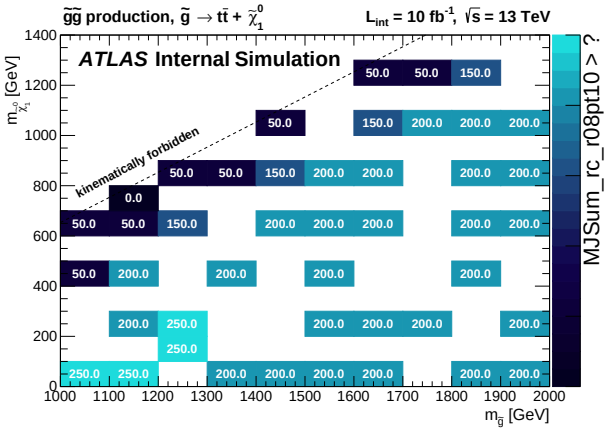

(c)

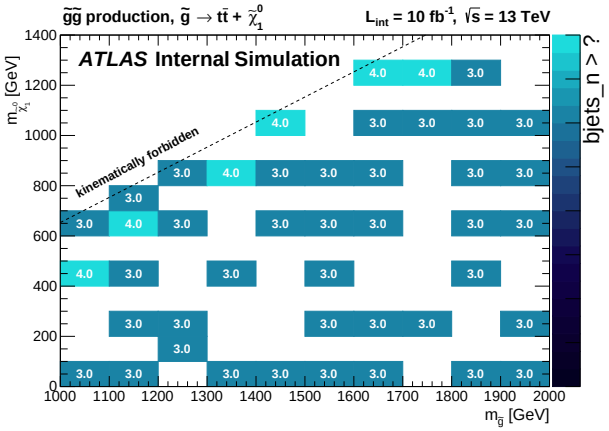

(d)

Figure 6.1: For each signal point in the Gtt mass plane, for the 0l channel, the optimal cut on (a) $E_{\mathrm{T}}^{\mathrm{miss}}$, (b) $\mathrm{m_{eff}}$, (c) $M_J^{\Sigma,4}$ and (d) the b-tagged jet multiplicity is shown.

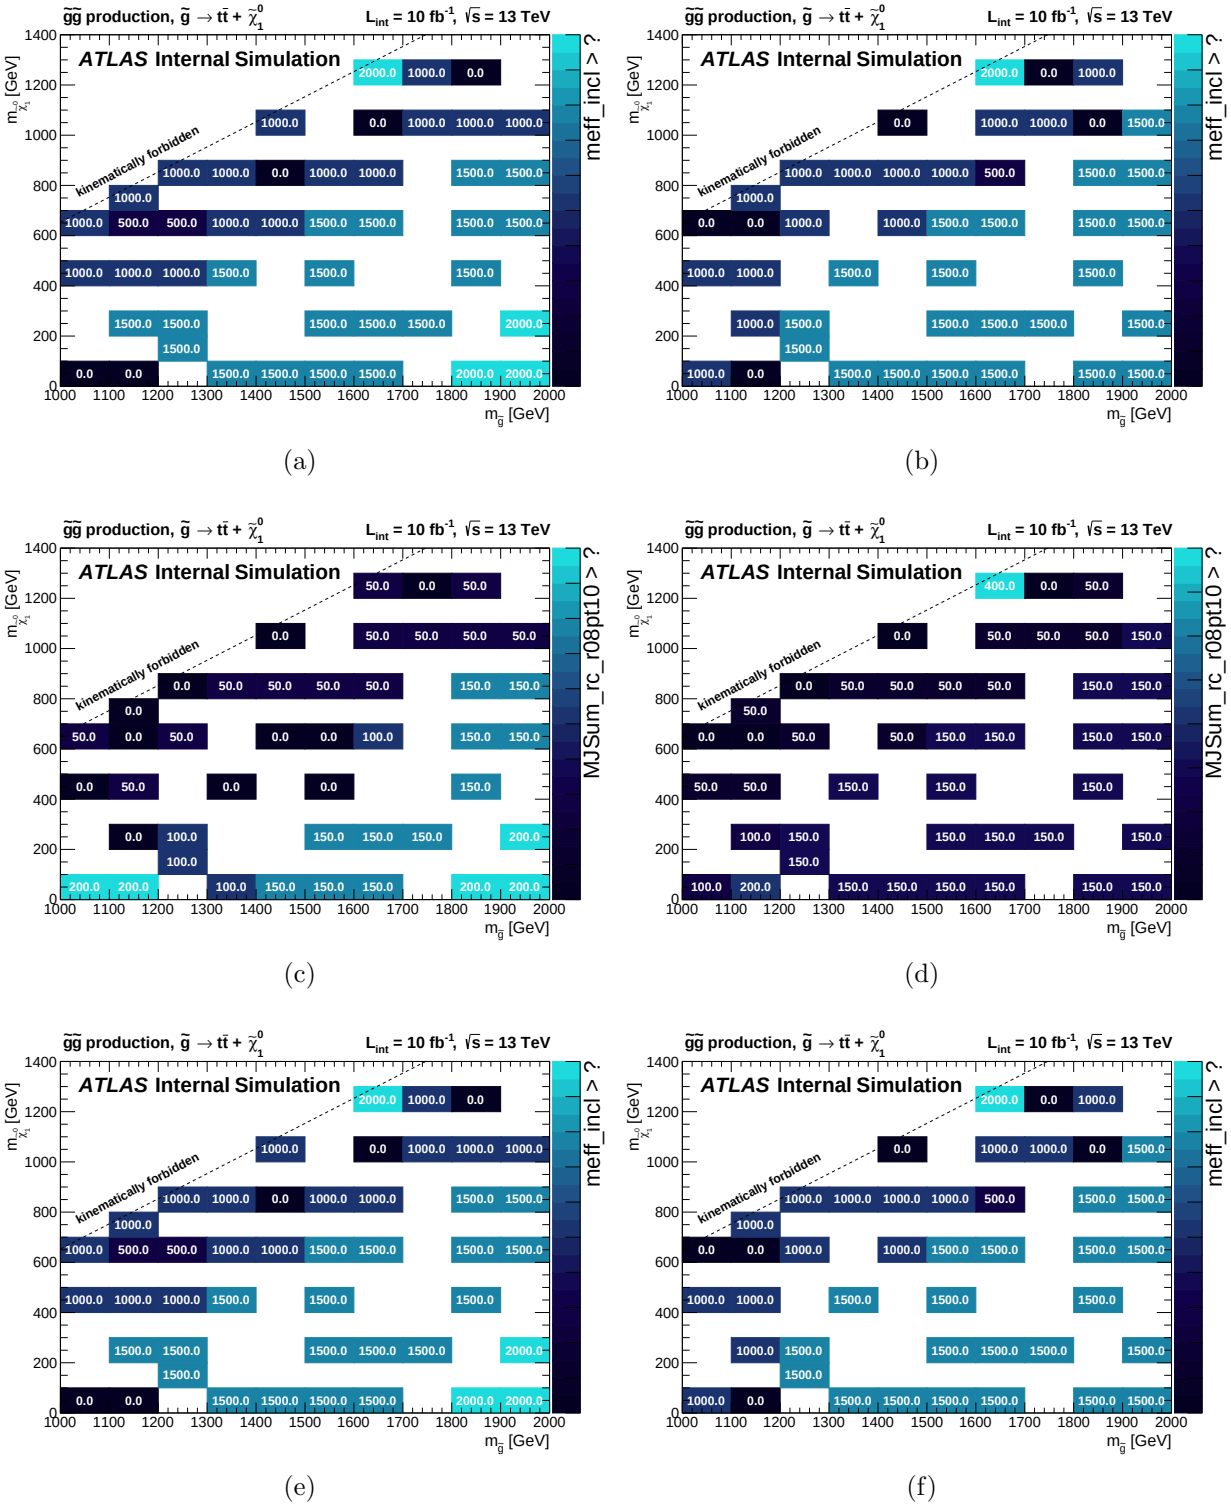

Figure 6.2: For each signal point in the Gtt mass plane, for the Gtt-1l channel, the optimal cut on (a,b) m_{eff}, (c,d) $M_J^{\Sigma,4}$ and (e,f) $E_{\text{T}}^{\text{miss}}$ is shown for the (left) 3b and (right) 4b optimisation stream.

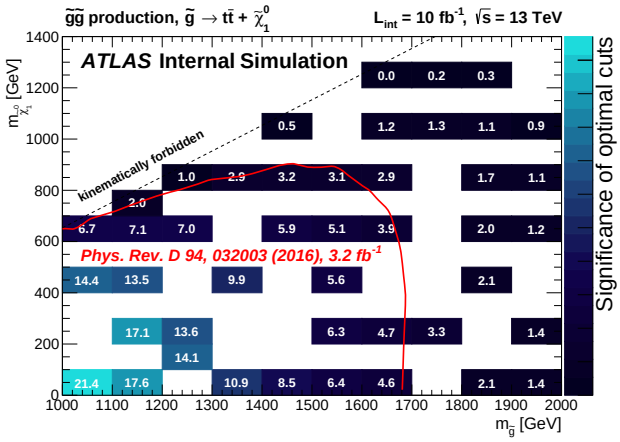

(a)

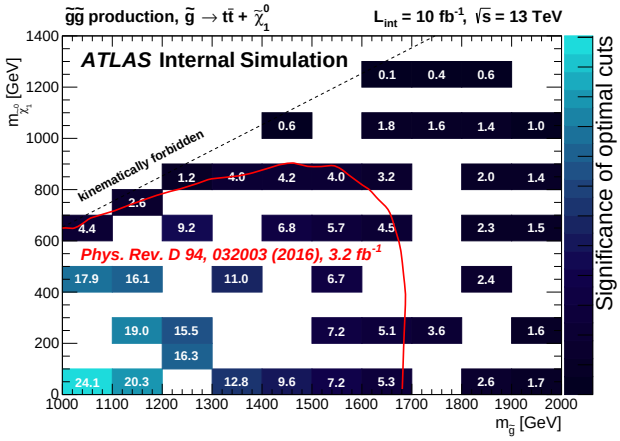

(b)

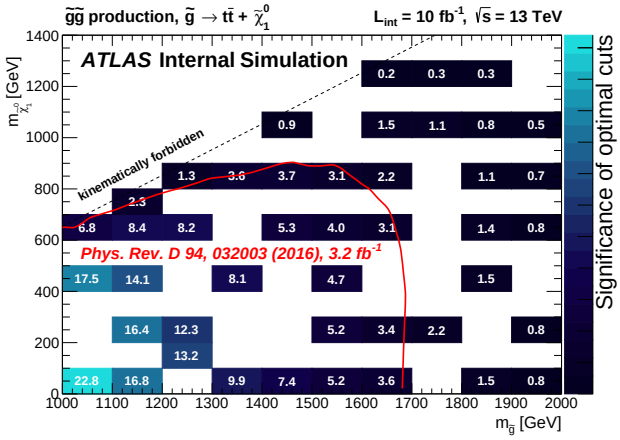

(c)

Figure 6.3: The significances corresponding to the optimal combination of cuts for each signal point in the Gtt mass phase space, for the (a) Gtt-0L, (b) Gtt-1l3b and (c) Gtt-1l4b channels.

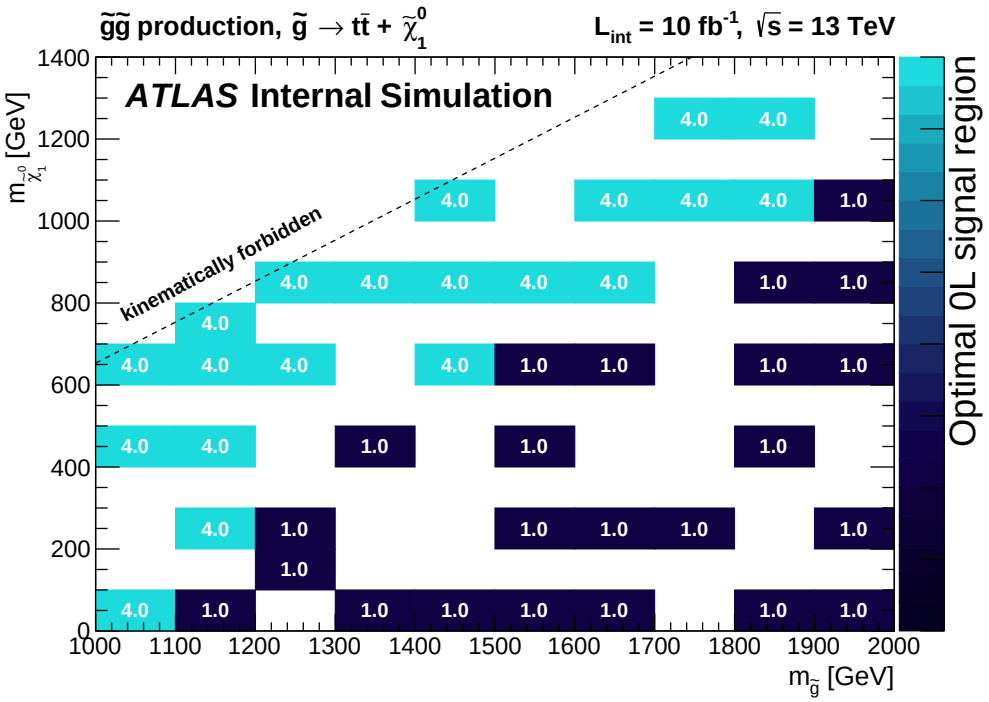

(a)

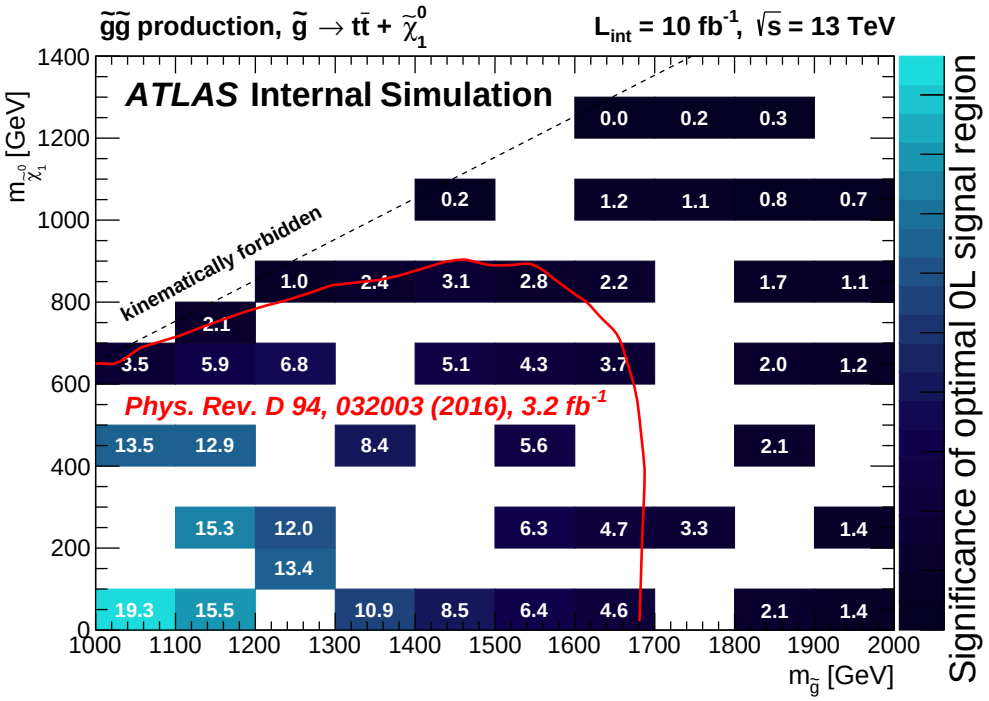

(b)

Figure 6.4: The optimal (a) Gtt-0l signal region and (b) its corresponding significance is shown for each signal point in the Gtt mass phase space. Here, points labelled 1.0 correspond to SR-Gtt-0l-A, while those labelled 4.0 correspond to SR-Gtt-0l-B.

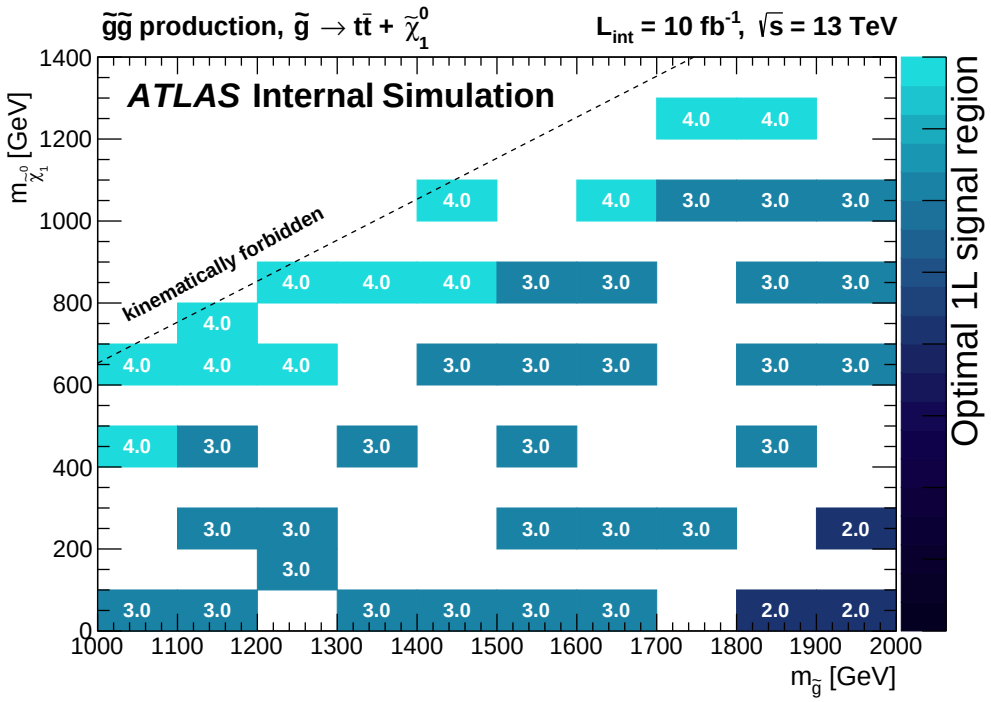

(a)

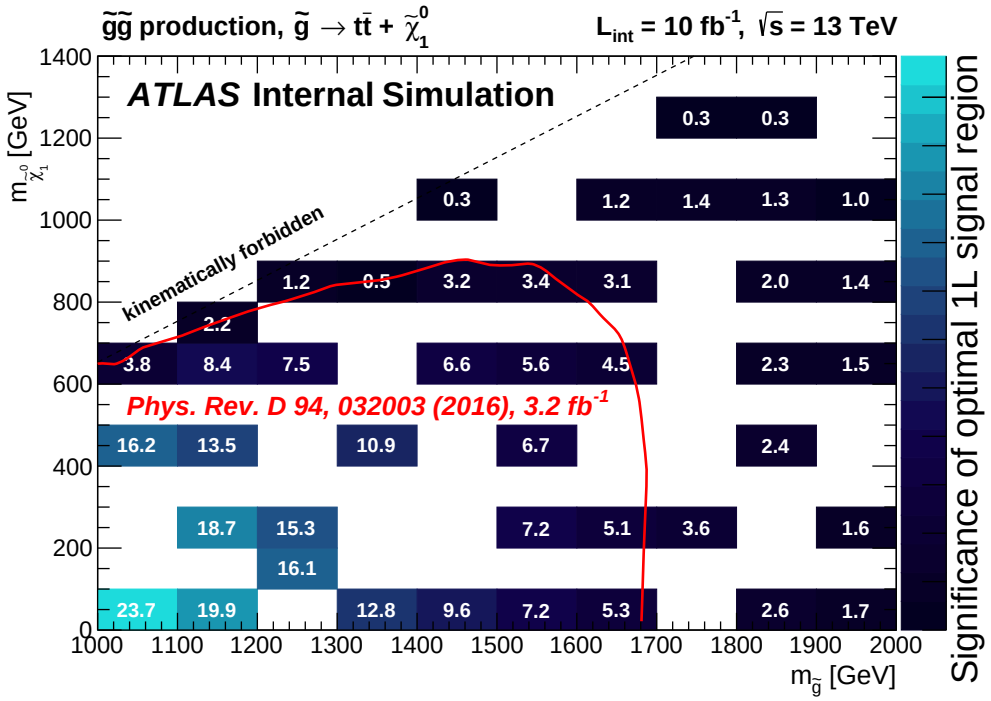

(b)

Figure 6.5: The optimal (a) Gtt-1l signal region and (b) its corresponding significance is shown for each signal point in the Gtt mass phase space. Here, points labelled 2.0 correspond to SR-Gtt-1l-A, points labelled 3.0 correspond to SR-Gtt-1l-B and those labelled 4.0 correspond to SR-Gtt-1l-C.

	Variable	SR-Gtt-0l-A	SR-Gtt-0l-B	SR-Gtt-1l-A	SR-Gtt-1l-B	SR-Gtt-1l-C	SR-Gbb-A	SR-Gbb-B
Selection	$N^{\text{Signal Lepton}}$	$= 0$	$= 0$	≥ 1	≥ 1	≥ 1	$= 0$	$= 0$
	N^{jet}	≥ 8	≥ 8	≥ 6	≥ 6	≥ 6	≥ 4	≥ 4
	$N^{b-\text{tag}}$	≥ 3	≥ 3	≥ 3	≥ 3	≥ 4	≥ 3	≥ 4
	$p_{\text{T}}^{\text{jet}}$ (GeV)	≥ 30	≥ 30	≥ 30	≥ 30	≥ 30	≥ 70	≥ 70
	$\Delta\phi_{\text{min}}^{4j}$	> 0.4	> 0.4	–	–	–	> 0.4	> 0.4
	$\text{m}_{\text{T,min}}^{b-\text{jets}}$ (GeV)	> 80	> 80	> 120	> 120	> 80	–	–
	m_{T} (GeV)	–	–	> 200	> 200	> 150	–	–
	$E_{\text{T}}^{\text{miss}}$ (GeV)	> 400	> 400	> 200	> 350	> 200	> 450	> 350
	$\text{m}_{\text{eff}}^{\text{incl}}$ (GeV)	> 2000	> 1250	> 2000	> 1500	> 500	–	–
	$\text{m}_{\text{eff}}^{\text{4j}}$ (GeV)	–	–	–	–	–	> 1900	> 1400
	$M_J^{\Sigma,4}$ (GeV)	> 200	–	> 200	> 150	–	–	–
Expected SM yield		1.1	6.7	0.91	0.92	4.6	1.88	11.0
Expected % $t\bar{t}$		65%	66%	70%	66%	75%	57%	75%

Table 6.2: A summary of the signal regions defined in the search for gluino pair production. The expected yield of simulated SM events in each region is provided (pre-fit), along with the expected $t\bar{t}$ fraction.

6.2 Developing Control and Validation Regions

6.2.1 Control Regions

Once the signal regions have been defined, their dominant $t\bar{t}$ background is estimated from data using a semi-data-driven approach. An orthogonal control region (CR) is defined for each signal regions (SR), from which the background normalisation is extracted. Each CR is designed to be as kinematically compatible with the signal region as possible, enhanced in $t\bar{t}$, and depleted in signal events. A scale factor is derived from this region which normalises simulation to data by allowing the integral of the $t\bar{t}$ background to float in a likelihood fit. This scale factor is subsequently applied to the simulated $t\bar{t}$ component of the background in the SR in order to obtain a normalised background estimate.

The dominant background process is, specifically, semi-leptonically decaying $t\bar{t}$ production where the lepton either fails an identification requirement or falls outside of the kinematic acceptance. These events produce real $E_{\mathrm{T}}^{\mathrm{miss}}$ due to the presence of a neutrino in the final state, and when the $t\bar{t}$ system is produced in association with additional heavy-flavour jets, may enter the signal selection. The particular approach to defining control regions varies depending on the channel. For the 0l signal regions – both for the Gbb model and the 0l channel of the Gtt model – control regions which require the presence of a signal lepton are required, and an upper bound is placed on the $\mathrm{m_T}$ which maintains orthogonality between the 0l CRs and 1l SRs. This low-$\mathrm{m_T}$ region is highly enriched in $t\bar{t}$ and signal-depleted, resulting in a desirable background composition and low signal contamination within the 0l CRs. Other cuts on event-level observables are relaxed in order to ensure adequate statistics within the CR in order to accurately extract the scale factor. For the Gtt-1l channel, the minimum $\mathrm{m_T}$ requirement from the signal region is inverted, mitigating signal contamination and enriching the region with semi-leptonic $t\bar{t}$, events similarly to the approach used in the 0l CRs. The signal contamination within these regions is quantified by the ratio of the expected number of signal events to background events (S/B). The S/B value for each signal point, in each region, is given in figure 6.6 for the Gtt-0l channel, and in figure 6.7 for the Gtt-1l channel.

In order to maintain orthogonality between the 1l- and 0l- control regions for the purposes of the final statistical combination of the Gtt-0l and Gtt-1l channels, the jet multiplicity in the 1l CRs is required to be exactly six. Other requirements on event-level observables are loosened, which provides enough $t\bar{t}$ acceptance to perform the fit for background estimation.

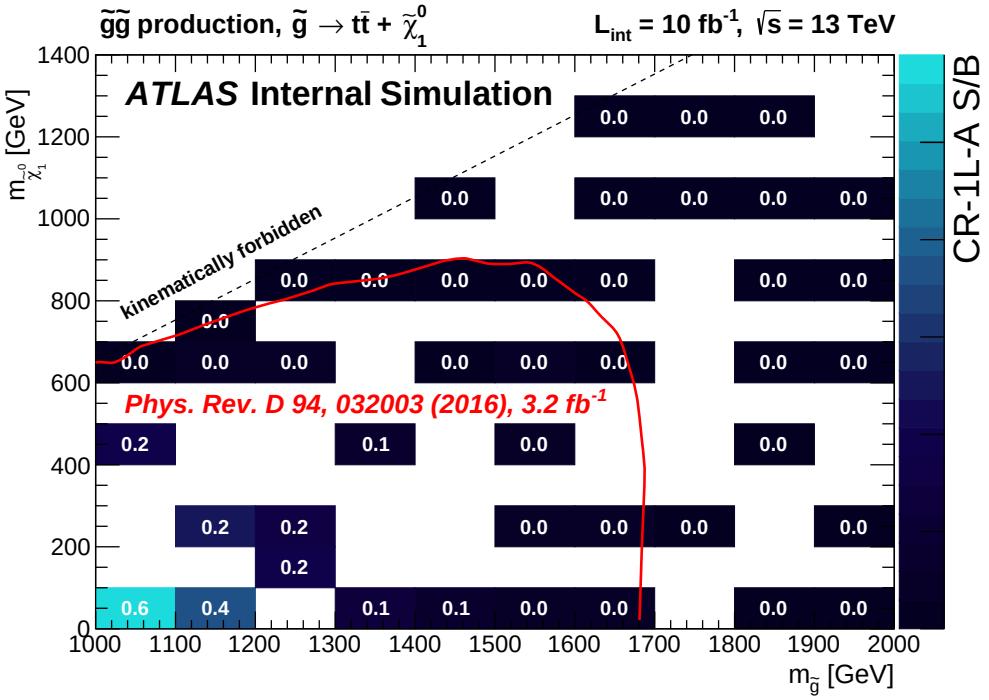

(a)

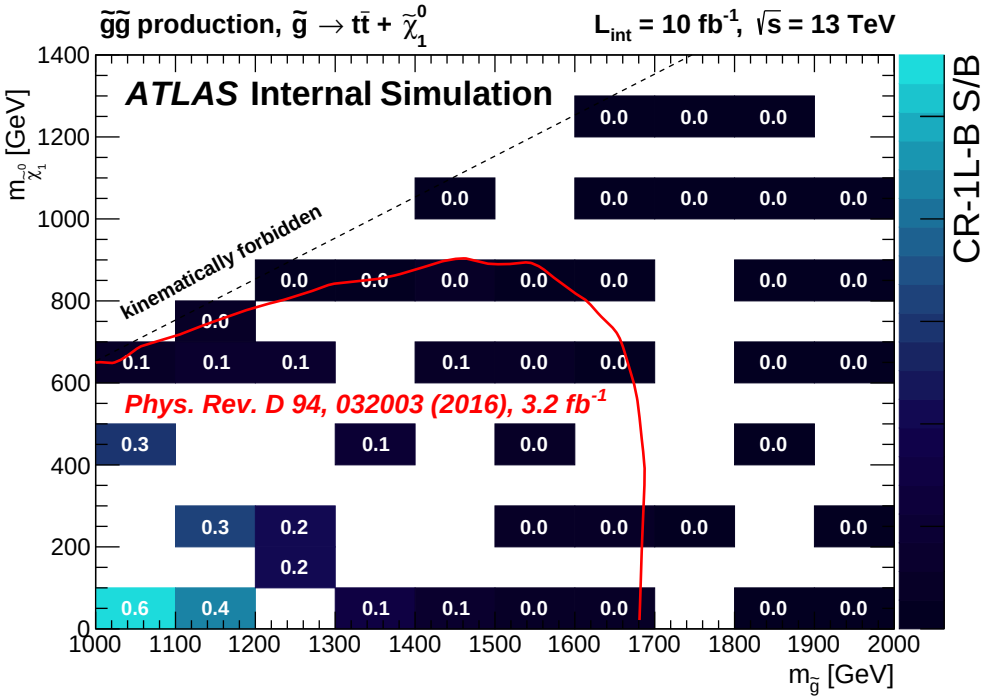

(b)

Figure 6.6: The expected ratio of signal events to background events for each Gtt signal point, for each Gtt-0L control region.

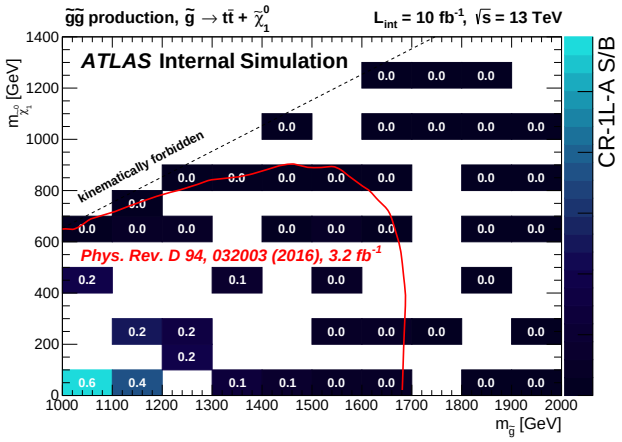

(a)

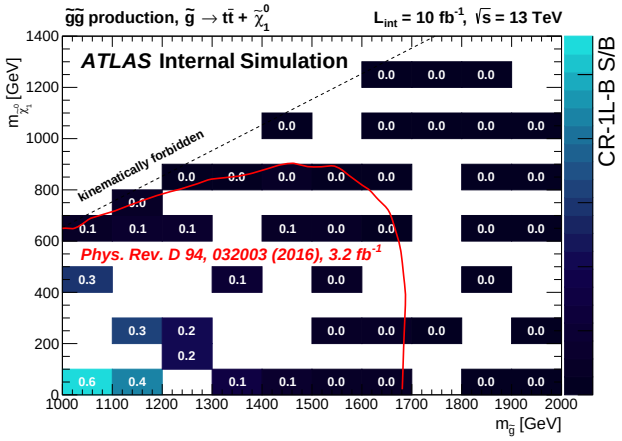

(b)

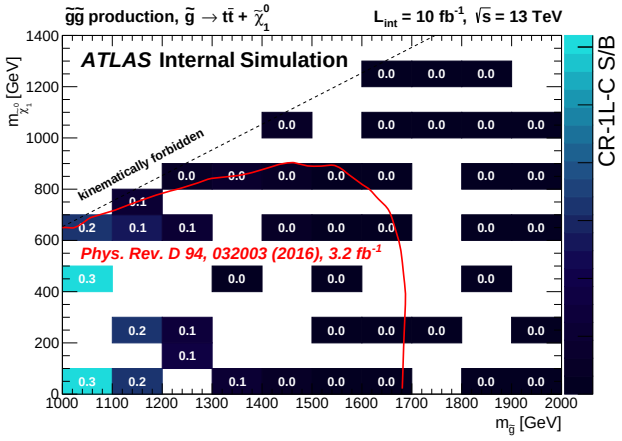

(c)

Figure 6.7: The expected ratio of signal events to background events for each Gtt signal point, for each Gtt-1l control region.

6.2.2 Validation Regions

Intermediate regions are also defined within portions of phase space between to the signal and control regions, within which the normalisation may be validated. These validation regions (VRs) target specific extrapolations, corresponding to the channel they are associated with. For the Gbb simplified model, a set of VRs are defined in table 6.3 which validate the extrapolation from small to large values of m_T, kept orthogonal from the SR by reversing the m_{eff}^{4j} requirement.

For the Gtt-0l channel, two VRs are defined in table 6.4 for each SR/CR pair: one validates the extrapolation across $m_{T,min}^{b-jets}$ (VR1), while the other validates the extrapolation across lepton multiplicity and m_T. Orthogonality from the Gtt-0l signal regions is maintained by either requiring the presence of a signal lepton (as required for the Gtt-0l CRs), or by altering cuts on event-level observables: the requirement on $M_J^{\Sigma,4}$ is reversed for Gtt-0l-VR0L-A, while a sideband requirement at low E_T^{miss} is defined for Gtt-0l-VR0L-B.

In the Gtt-1l channel, two VRs are also defined in table 6.5 per-SR/CR-pair. The first set validate the extrapolation over m_T imposed by reversing the cut on this variable when defining the Gtt-1l control regions, and are kept orthogonal from the SR by instead reversing the requirements on $M_J^{\Sigma,4}$ (Gtt-1l-VRm_T-A and Gtt-1l-VRm_T-B) or $m_{T,min}^{b-jets}$ (Gtt-1l-VRm_T-C). The second set validate the extrapolation over $m_{T,min}^{b-jets}$ caused by the removal of this requirement in the CR definitions, while maintaining the corresponding reversed m_T cut. All Gtt-1l VRs place an inclusive requirement on the jet multiplicity, and so also check the behaviour of the extrapolation made from the exclusive requirement placed on the Gtt-1l CRs.

The signal contamination was evaluated for each VR in the same manner as the CRs, and was found to be between 5-20% for signal points which have not previously been excluded by earlier iterations of this search.

6.2.3 Background and Flavour Composition

In order for the background estimation techniques to remain robust against the possible flavour-content mis-modelling observed in the $t\bar{t}$ simulation (figure 5.4(b)), particular care was taken in the design of the signal and control region pairs to maintain as large a fraction of $t\bar{t}$ relative to the total background in the regions as possible, and to maintain a roughly-constant flavour composition within the $t\bar{t}$ background for these pairs of regions. Figures 6.8 - 6.10 provide a breakdown of both the fractional background and flavour compositions within these regions. In general, acceptable consistency is observed, largely as a result of never performing an extrapolation over the b-tag multiplicity between a given CR to an SR. Across all regions defined for the search, the lowest $t\bar{t}$ fraction is observed in SR-Gtt-1l-B, where the relative fraction of $t\bar{t}$ is 60.7%. The largest differences in flavour content are also observed for the Gtt-1l regions, where the relative fraction of $t\bar{t}$ + heavy flavour may differ by up to 20-30% between the SR and CR.

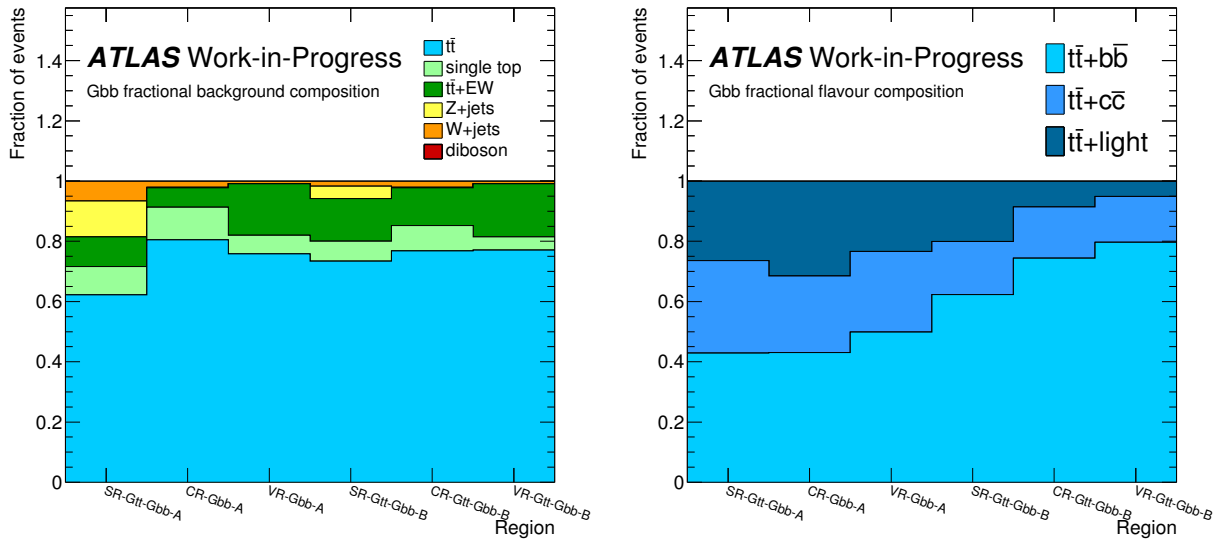

(a) (b)

Figure 6.8: The (a) fractional background composition and (b) fractional flavour composition of the $t\bar{t}$ background within the Gbb signal, control and validation regions.

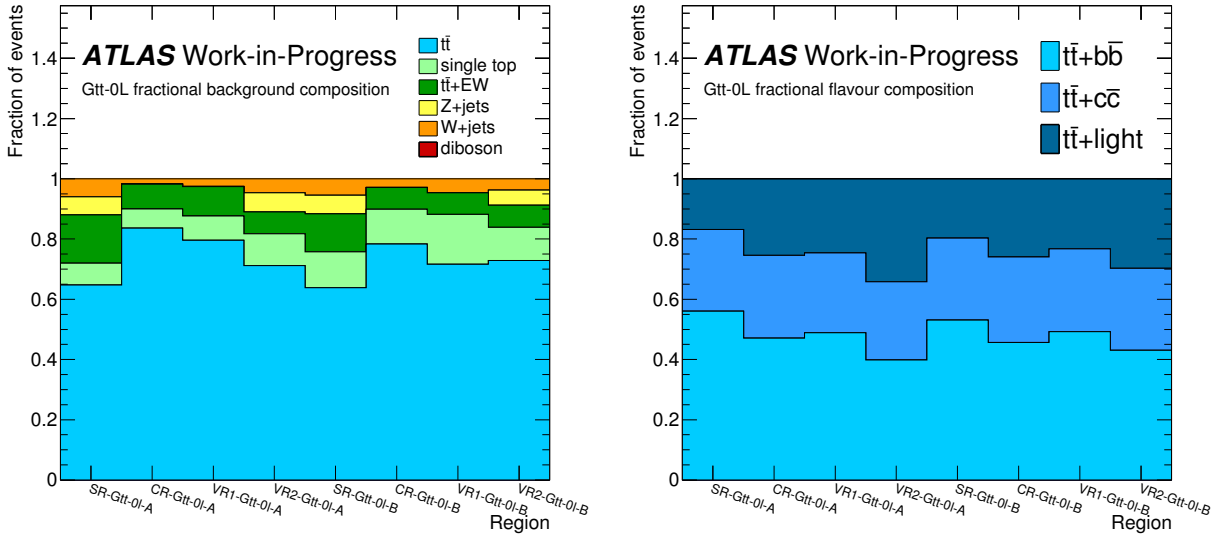

(a) (b)

Figure 6.9: The (a) fractional background composition and (b) fractional flavour composition of the $t\bar{t}$ background within the Gtt-0l signal, control and validation regions.

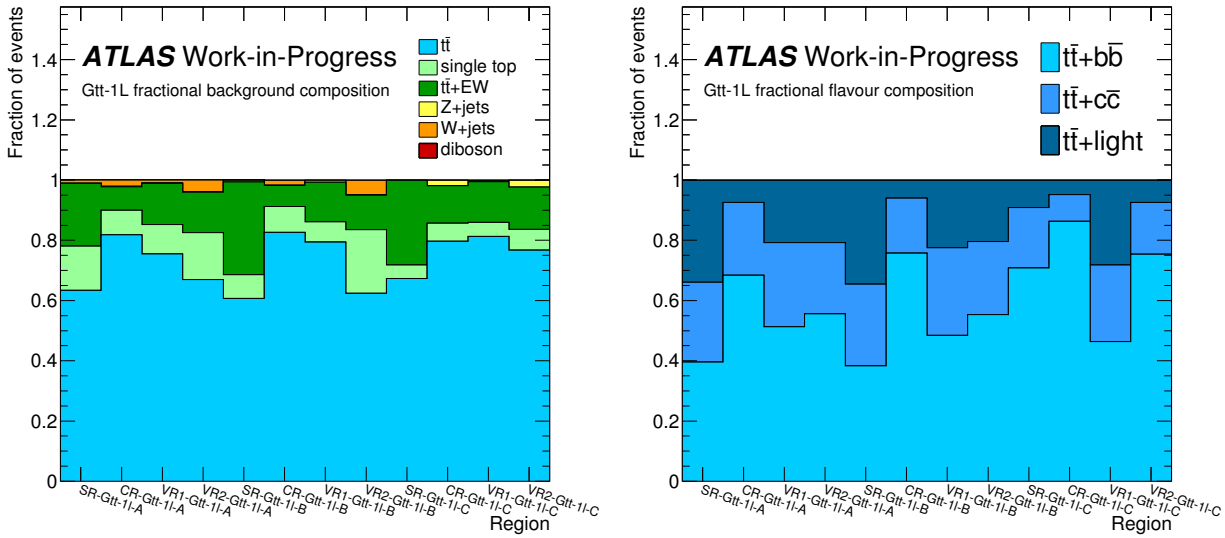

(a) (b)

Figure 6.10: The (a) fractional background composition and (b) fractional flavour composition of the $t\bar{t}$ background within the Gtt-1l signal, control and validation regions.

	Variable	SR	CR	VR
Common Criteria	$N^{\text{Candidate Lepton}}$	$= 0$	$-$	$= 0$
	$N^{\text{Signal Lepton}}$	$-$	$= 1$	$-$
	N^{jet}	≥ 4	≥ 4	≥ 4
	$\Delta\phi^{4j}_{\text{min}}$	> 0.4	$-$	> 0.4
	m_{T}	$-$	< 150	$-$
	m_{T}	$-$	< 150	$-$
Region A (Large mass splitting)	${p_{\text{T}}}^{\text{jet}}$	> 70	> 70	> 70
	$N_{\text{b-jet}}$	≥ 3	≥ 3	≥ 3
	$E_{\text{T}}^{\text{miss}}$	> 450	> 350	> 450
	$\text{m}_{\text{eff}}^{\text{4j}}$	> 1900	> 1750	< 1900
Region B (Small mass splitting)	${p_{\text{T}}}^{\text{jet}}$	> 30	> 30	> 30
	$N_{\text{b-jet}}$	≥ 4	≥ 4	≥ 4
	$E_{\text{T}}^{\text{miss}}$	> 300	> 300	> 275
	$\text{m}_{\text{eff}}^{\text{4j}}$	> 1000	> 1000	< 1000

Table 6.3: Definitions of the Gbb signal, control and validation regions. The unit of all kinematic variables is GeV except $\Delta\phi^{4j}_{\text{min}}$, which is in radians. The jet p_{T} requirement is also applied to b-tagged jets.

	Variable	SR	CR	VR1l	VR0l
Common Criteria	$N^{\text{Signal Lepton}}$	$= 0$	$= 1$	$= 1$	$= 0$
	$N_{\text{b-jet}}$	≥ 3	≥ 3	≥ 3	≥ 3
	${p_{\text{T}}}^{\text{jet}}$	> 30	> 30	> 30	> 30
	$\Delta\phi_{\text{min}}^{4j}$	> 0.4	$-$	$-$	> 0.4
	$\text{m}_{\text{T,min}}^{b-\text{jets}}$	> 80	$-$	> 80	> 80
	m_{T}	$-$	< 150	< 150	$-$
Region A (Large mass splitting)	N^{jet}	≥ 8	≥ 7	≥ 7	≥ 6
	$E_{\text{T}}^{\text{miss}}$	> 400	> 250	> 200	> 300
	$\text{m}_{\text{eff}}^{\text{incl}}$	> 2000	> 1750	> 1750	> 1300
	$M_J^{\sum,4}$	> 200	> 200	> 200	< 200
Region B (Small mass splitting)	N^{jet}	≥ 8	≥ 7	≥ 7	≥ 7
	$E_{\text{T}}^{\text{miss}}$	> 400	> 400	> 325	$\in (300, 400)$
	$\text{m}_{\text{eff}}^{\text{incl}}$	> 1250	> 1250	> 1250	> 1200

Table 6.4: Definitions of the Gtt 0-lepton signal, control and validation regions. The unit of all kinematic variables is GeV except $\Delta\phi_{\text{min}}^{4j}$, which is in radians. The jet p_{T} requirement is also applied to b-tagged jets.

	Variable	SR	CR	VR-m_T	VR-$m_{T,min}^{b-jets}$
Common Criteria	$N^{\text{Signal Lepton}}$	≥ 1	≥ 1	≥ 1	≥ 1
	${p_T}^{\text{jet}}$	> 30	> 30	> 30	> 30
Region A (Large mass splitting)	N^{jet}	≥ 6	$= 6$	≥ 5	≥ 6
	$N_{\text{b-jet}}$	≥ 3	≥ 3	≥ 3	≥ 3
	m_T	> 200	< 200	> 200	< 200
	$m_{T,min}^{b-jets}$	> 120	$-$	$-$	> 120
	E_T^{miss}	> 200	> 200	> 200	> 200
	m_{eff}^{incl}	> 2000	> 1500	> 1350	> 1500
	$M_J^{\sum,4}$	> 200	> 200	< 200	> 200
Region B (Moderate mass splitting)	N^{jet}	≥ 6	$= 6$	≥ 5	≥ 6
	$N_{\text{b-jet}}$	≥ 3	≥ 3	≥ 3	≥ 3
	m_T	> 200	< 200	> 200	< 200
	$m_{T,min}^{b-jets}$	> 120	$-$	$-$	> 120
	E_T^{miss}	> 350	> 300	> 250	> 300
	m_{eff}^{incl}	> 1500	> 1250	> 1100	> 1500
	$M_J^{\sum,4}$	> 150	> 150	< 150	> 150
Region C (Small mass splitting)	N^{jet}	≥ 6	$= 6$	≥ 6	≥ 6
	$N_{\text{b-jet}}$	≥ 4	≥ 4	≥ 3	≥ 4
	m_T	> 150	< 150	> 150	< 150
	$m_{T,min}^{b-jets}$	> 80	$-$	< 80	> 80
	E_T^{miss}	> 200	> 200	> 200	> 200
	m_{eff}^{incl}	> 500	> 500	> 500	> 500

Table 6.5: Definitions of the Gtt 1-lepton signal, control and validation regions. The unit of all kinematic variables is GeV. The jet p_T requirement is also applied to b-tagged jets.

7

Interpretation

Once the signal and control regions are defined and a validation procedure is put in place, the data may be examined for signs of potential new physics.

The statistical analysis for this search is performed in several steps. A background-only fit is performed (section 7.1), which is used to evaluate the yield of events in data within each signal region, which are compared with the expected yield of background events. In the absence of an obvious discovery, an exclusion fit is performed (section 7.2) which derives a model-independant limit on the possible cross-sections of new physics within each signal region, as well as a set of model-dependent exclusion contours for the Gbb and Gtt simplified models, as a function of the gluino and χ_0^1 masses which parameterise them.

The statistical analysis for this search, as for all searches performed by the ATLAS SUSY Physics Group so-far during run 2, was performed within the HistFitter statistical analysis framework [98].

7.1 Background-only Fit

Within each control region, an estimate of the $t\bar{t}$ contribution from data must be derived and propagated to the corresponding validation and signal regions. This estimation is performed by deriving a scale factor for each control region which normalises the predicted background yield to the observed number of signal events with a profile likelihood fit. The information within a given control region is parameterised by the raw yield of each background process from simulation and the observed number of events in data within the region. The number of events predicted by simulation to lie within a region, n_{pred}, may be decomposed into the number of expected signal and background events, such as

$$n_{\text{pred}} = \mu_{\text{sig}} s_{\text{pred}} + b_{\text{pred}} \tag{7.1}$$

where s and b are the number of signal and background events, and the parameter μ is the **signal strength**. The background prediction b_{pred} may be decomposed further; it is helpful to note that the $t\bar{t}$ contribution and its own normalisation factor, $\mu_{t\bar{t}}$, could be written separately from the remaining background prediction:

$$b_{\text{pred}} = \mu_{t\bar{t}} b_{\text{pred}}^{t\bar{t}} + b_{\text{pred}}^{\text{other}}. \tag{7.2}$$

Varying μ_{sig} corresponds to changing the hypothesis which is tested by our statistical analysis: assumptions that the signal is absent or nominally present during a fitting procedure would correspond to values of $\mu = 0$ or $\mu = 1$, respectively. The probability density function (PDF) which describes a single analysis region's contents is a Poisson distribution with expectation value n_{pred}:

$$\begin{aligned}
P_{\text{R}}(n_{\text{obs}}|n_{\text{pred}}) &= \frac{(n_{\text{pred}})^{n_{\text{obs}}}}{n_{\text{obs}}!} e^{-n_{\text{pred}}} \\
P_{\text{R}}(n_{\text{obs}}|\mu_{\text{sig}}, b_{\text{pred}}) &= \frac{(\mu_{\text{sig}} s_{\text{pred}} + b)^{n_{\text{obs}}}}{n_{\text{obs}}!} e^{-(\mu_{\text{sig}} s_{\text{pred}} + b_{\text{pred}})} \\
P_{\text{R}}(n_{\text{obs}}|\mu_{\text{sig}}, \mu_{t\bar{t}}, b_{\text{pred}}^{\text{other}}) &= \frac{(\mu_{\text{sig}} s_{\text{pred}} + \mu_{t\bar{t}} b_{\text{pred}}^{t\bar{t}} + b_{\text{pred}}^{\text{other}})^{n_{\text{obs}}}}{n_{\text{obs}}!} e^{-(\mu_{\text{sig}} s_{\text{pred}} + \mu_{t\bar{t}} b_{\text{pred}}^{t\bar{t}} + b_{\text{pred}}^{\text{other}})}.
\end{aligned} \tag{7.3}$$

Systematic uncertainties may be accounted for by constructing a PDF which is a product of Gaussian distributions each with unit width, corresponding to particular variations:

$$C_{\text{syst}}(\boldsymbol{\theta}^0|\boldsymbol{\theta}) = \prod_{j \,\in\, \text{systematics}} \text{Gaussian}(\boldsymbol{\theta}_j^0 - \boldsymbol{\theta}_j) \tag{7.4}$$

where θ^0 is the vector of nuisance parameters which define the width of each Gaussian (obtained from prior auxiliary measurement), and θ is a vector of parameters which dictate the portion of each systematic to apply to the full likelihood. A value of 1, which applies the entire systematic, is used in most cases.

The predicted $t\bar{t}$ event counts and the observed event counts in data are used to construct PDFs whose expected values thus depend on μ_{sig}, $b_{t\bar{t}}$. and $\boldsymbol{\theta}$, a vector representing the sources of systematic uncertainty as Gaussian-distributed nuisance parameters with widths corresponding to the magnitude of each source. For a particular SR/CR pairing, the likelihood function for a given CR could be written as the product of individual PDFs:

$$L(\mu_{\text{sig}}, \mu_{t\bar{t}}, b_{\text{pred}}^{\text{other}}, \boldsymbol{\theta}|n_{\text{obs}}, \boldsymbol{\theta}^0) = P_{\text{CR}} \times C_{\text{syst}}. \tag{7.5}$$

To check for a potential discovery, a **background-only fit** is performed, in which the

likelihood is maximized assuming that there is no contribution from signal processes in any region (*i.e.* that $\mu_{\text{sig}} = 0$). Only information from the control region is used in the fit, and so the resultant likelihood to be maximized is that listed in 7.5. The normalisation factor of the $t\bar{t}$ process, $\mu_{t\bar{t}}$, is taken as the free parameter in the fit, which maximizes the likelihood function. The value of $b_{\text{pred}}^{\text{other}}$, the total contribution from non-$t\bar{t}$ backgrounds, is obtained directly from simulation, and the impact of the systematic uncertainties θ is determined *a priori* using auxiliary measurements.

Initially, the agreement between data and the post-fit prediction from simulation is assessed in each of the validation regions. This comparison is illustrated in figure 7.1. In the bottom pane of this figure, a quantity known as the **pull** is introduced:

$$\text{pull} = \frac{(n_{\text{obs.}} - n_{\text{pred.}})}{\sigma_{tot}}. \tag{7.6}$$

The pull quantifies the statistical tension between the observed number of data events and the predicted background yield, as normalised by the total uncertainty of the background prediction. In this case, none of the observed VR pulls exceed $2\sigma_{tot}$ – there is good agreement observed across the entire range of Gbb and Gtt validation regions. Following this successful validation of the the background estimation procedure, the observed yield of data within the signal regions may be examined for the first time.

The unblinded search results are presented in figure 7.2, where the number of observed events in data for each Gbb and Gtt signal region is compared to the predicted background yield. No statistically significant excess or deficit is observed in any search region. For all regions but SR-Gtt-0l-B, very good agreement is observed, and pull values are less than $1\sigma_{tot.}$. There is higher tension in SR-Gtt-0l-B, where a slight excess in data is observed, corresponding to a statistical significance of 2.5 Gaussian standard deviations. The $E_{\text{T}}^{\text{miss}}$ distributions of signal events within SR-Gbb-A, SR-Gtt-0l-A and SR-Gtt-1l-A are shown in figure 7.3. A detailed breakdown of the observed data and background prediction before and after the fitting procedure is given in table 7.1 for all Gbb and Gtt signal regions. The values of $\mu_{t\bar{t}}$ are also provided: in most cases, the $t\bar{t}$ background is scaled down slightly by the fit. In regions which are enriched in heavy flavour content such as SR-Gbb-B and SR-Gtt-1L-C; however, the value of $\mu_{t\bar{t}}$ can be nearly 2: this large value is consistent with the underestimation of $t\bar{t}$ production with associated heavy flavour jets observed in figure 5.4(b).

In the absence of a statistically significant excess, model-independent and -dependent exclusion limits may now be calculated based on the observed results.

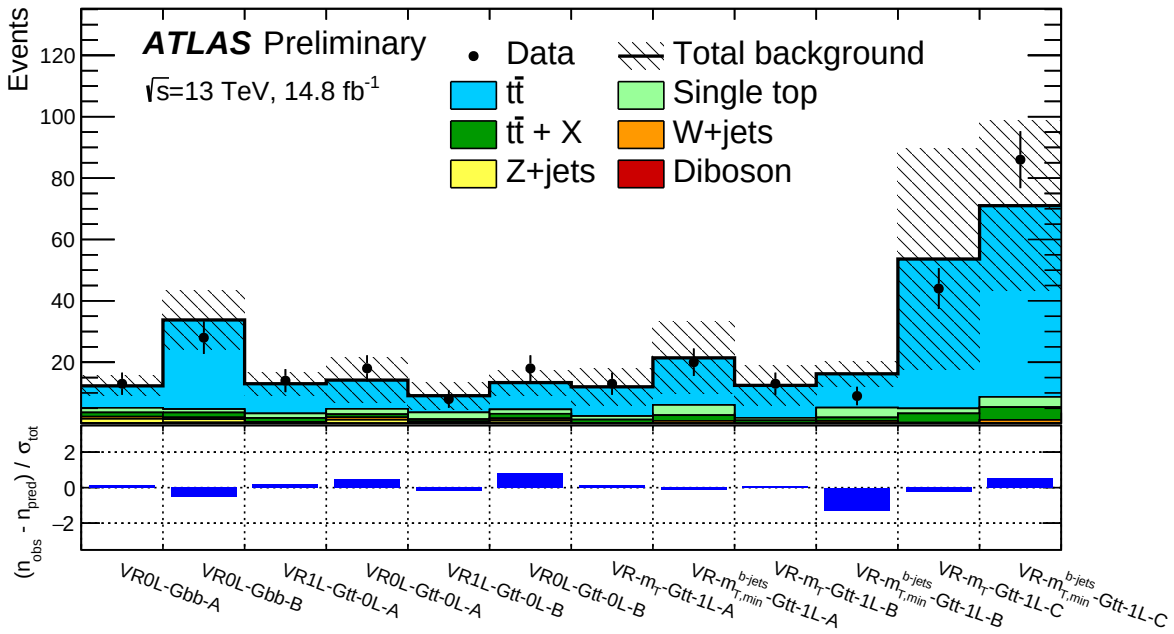

Figure 7.1: Validation region pull plot following background-only fit.

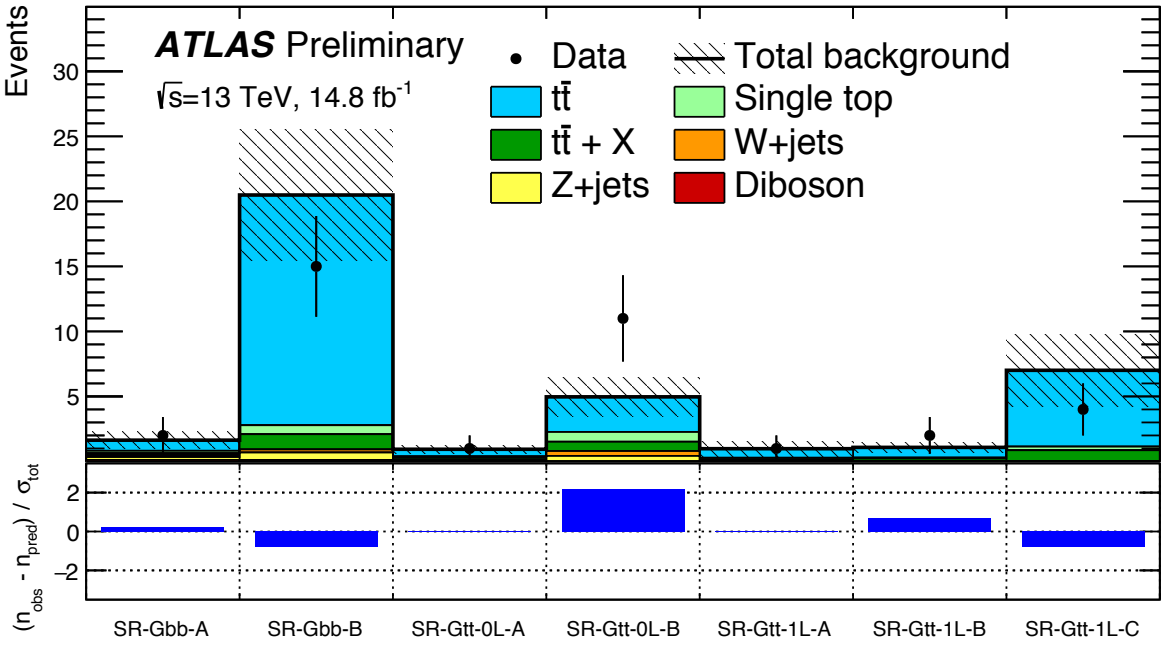

Figure 7.2: Signal region pull plot following background-only fit.

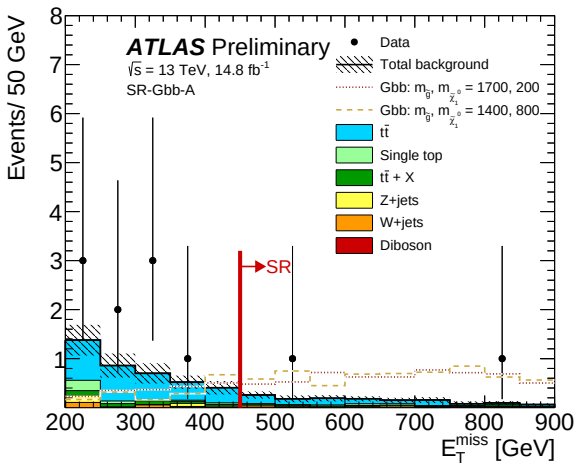

(a) SR-Gbb-A

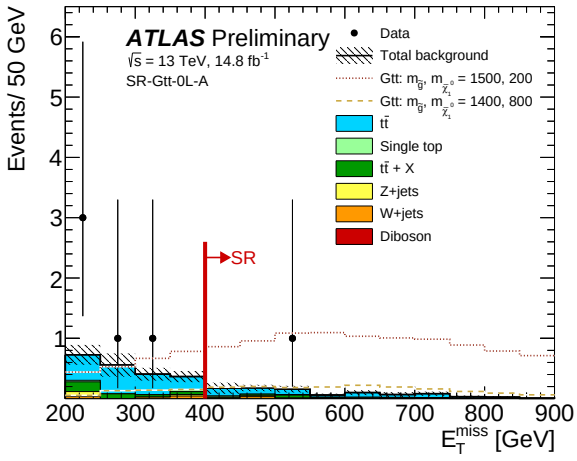

(b) SR-Gtt-0l-A

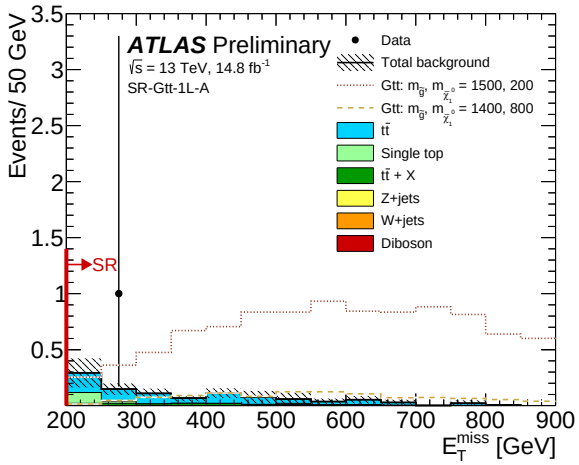

(c) SR-Gtt-1l-A

Figure 7.3: A comparison of $E_{\mathrm{T}}^{\mathrm{miss}}$ distributions in data and simulation, post-fit, in various signal regions: (a) SR-Gbb-A, (b) SR-Gtt-0l-A, (c) SR-Gtt-1l-A.

	SR-Gbb-A	SR-Gbb-B	SR-Gtt-0l-A	SR-Gtt-0l-B	SR-Gtt-1l-A	SR-Gtt-1l-B	SR-Gtt-1l-C
Observed events	2	15	1	11	1	2	4
Fitted BG events	1.6 ± 0.7	21 ± 5	0.93 ± 0.31	5.0 ± 1.5	1.0 ± 0.6	1.1 ± 0.4	7.0 ± 2.8
$t\bar{t}$	0.82 ± 0.34	18 ± 5	0.55 ± 0.20	2.7 ± 1.1	0.7 ± 0.6	0.79 ± 0.30	5.9 ± 2.7
Single-top	0.18 ± 0.14	0.69 ± 0.34	0.08 ± 0.06	0.8 ± 0.5	$0.12^{+0.13}_{-0.12}$	0.07 ± 0.07	0.2 ± 0.22
$t\bar{t} + X$	0.15 ± 0.08	1.2 ± 0.6	0.16 ± 0.09	0.7 ± 0.4	0.15 ± 0.08	0.22 ± 0.12	0.9 ± 0.5
Z+jets	$0.2^{+0.5}_{-0.2}$	0.6 ± 0.5	$0.07^{+0.15}_{-0.07}$	$0.4^{+0.6}_{-0.4}$	< 0.01	< 0.01	< 0.1
W+jets	$0.15^{+0.16}_{-0.15}$	0.24 ± 0.14	$0.08^{+0.09}_{-0.08}$	0.3 ± 0.24	< 0.01	< 0.01	< 0.1
Diboson	0.12 ± 0.08	$0.11^{+0.13}_{-0.11}$	< 0.01	< 0.01	< 0.01	< 0.01	< 0.01
MC-only BG prediction	1.88	11.0	1.1	6.7	0.91	0.92	4.6
$\mu_{t\bar{t}}$	0.77 ± 0.33	2.1 ± 0.7	0.75 ± 0.29	0.61 ± 0.29	1.1 ± 0.3	1.3 ± 0.4	1.7 ± 0.5

Table 7.1: A summary of the yields in data and simulation for each Gbb and Gtt signal region in the search for gluino pair production, following the background-only fit. Statistical and systematic uncertainties are combined in the quoted error. The $t\bar{t}$ scale factor produced by the fit, $\mu_{t\bar{t}}$, is also listed for each region.

7.2 Exclusion Fit

When calculating limits with the so-called **exclusion fit**, the information within the signal regions may be included in the maximized likelihood, as no signal was observed in any Gbb or Gtt region. The likelihood is therefore taken as

$$L(\mu_{\text{sig}}, \mu_{t\bar{t}}, b_{\text{pred}}^{\text{other}}, \boldsymbol{\theta} | n_{\text{obs}}, \boldsymbol{\theta}^0) = P_{\text{SR}} \times P_{\text{CR}} \times C_{\text{syst}}. \tag{7.7}$$

In order to place limits on new physics, a log-likelihood-ratio is constructed which will act as the test statistic for a given value of μ_{sig}[1]:

$$q(\mu_{\text{sig}}) = -2 \log \left(\frac{L(\mu_{\text{sig}}, \hat{\hat{\boldsymbol{\theta}}})}{L(\hat{\mu_{\text{sig}}}, \hat{\boldsymbol{\theta}})} \right), \tag{7.8}$$

where $\hat{\mu_{\text{sig}}}$ and $\hat{\theta}$ are selected to maximize the overall likelihood, and $\hat{\hat{\theta}}$ maximizes the value of the likelihood for a specific value of μ_{sig}. A hypothesis test may then be performed on the distribution $f(q(\mu_{\text{sig}})|\mu_{\text{sig}}, \theta)$ to obtain the p-value that a particular model has been discovered or excluded: in practice, the distribution f must be calculated by computing many pseudoexperiments, which find the value of $q(\mu_{\text{sig}})$ for varied quantities of n_{obs} and θ. Due to the large number of nuisance parameters which are included in typical physics analyses, this approach is computationally unfeasible. Instead, the values of θ^0 are **profiled**, or fit to the data during the likelihood fitting procedure. This procedure reduces the dimensionality of the space which must be sampled by pseudoexperiments in order to generate the distribution of the test statistic, and so allows its computation.

In order to exclude a hypothetical signal model, it is necessary to compare the p-values corresponding to the 'signal + background' (p_{s+b}) and 'background only' (p_b) hypotheses:

$$p_{s+b} = \int_q^\infty f\left(q(\mu_{\text{sig}} = 1)\right) dq \tag{7.9}$$

$$p_b = 1 - \int_q^\infty f\left(q(\mu_{\text{sig}} = 0)\right) dq. \tag{7.10}$$

The explicit use of p_{s+b} when placing limits on specific models is limited due to the fact that it may take small values in the event of a downward statistical fluctuation of n_{obs}, particularly in cases where $n_{\text{obs}} < b_{\text{pred}}$. These low values could result in artificially strong limits, or even the erroneous exclusion of signal points to which there is low sensitivity. The

[1] This ratio is known by the Neyman-Pearson Lemma [99] to be the most efficient possible discriminating variable.

p-value of the background-only hypothesis, p_b, may be used to regulate this behaviour as it will become large in this case. The CL_s value is defined as a quotient of these two p-values:

$$\mathrm{CL}_s = \frac{p_{s+b}}{1 - p_b} = \frac{\mathrm{CL}_{s+b}}{\mathrm{CL}_b}. \tag{7.11}$$

The CL_s value, not the raw p_{s+b} value, is used to set exclusion limits on specific models. A signal point is excluded if its corresponding CL_s value is smaller than a predetermined threshold value, α. Conventionally, α is taken as 0.05 by the ATLAS collaboration, and so corresponding limits are set "at 95% CL."

7.2.1 Model-Independent Cross Section Limits

Using the observed data and background prediction in each signal region, a generic upper limit may be set on the number of BSM events which could be compatible with observation. A signal contibution of $s = 1$ is assumed in each region, and the value of μ_{sig} is increased from 0, evaluating p_{s+b} at each step, until a value of $p_{s+b} = 0.05$ is obtained. The results of performing such a scan for each of the Gbb and Gtt signal regions are summarised in table 7.2. The maximum value of μ_{sig} reached in the scan for each region is reported as $\mathbf{S}^{\mathbf{95}}_{\mathrm{obs}}$, interpreted as the 95% CL upper limit on the number of BSM events. A second upper limit, $\mathbf{S}^{\mathbf{95}}_{\mathrm{exp}}$, is computed for an alternative scenario where the observed data had perfectly agreed with the background prediction. These limits are also translated into an upper limit on the visible cross-section (σ_{vis}) of a BSM process within each region.

For each signal region, the p-value p_0 is provided, which describes the probability that a statistical fluctuation in the expected background could produce the number of events observed in data.

7.2.2 Model-Dependent Exclusion Contours

Exclusion limits may be placed on the same Gbb and Gtt simplified models used to optimise the search, as a function of the parameters used to define the models: the mass of the gluino ($\tilde{g}$), and the mass of the lightest neutralino ($\tilde{\chi}^1_0$). For each available signal point in the Gbb and Gtt grids, an exclusion fit is performed in order to find the value of $\mu_{t\bar{t}}$ which maximizes the exclusion-fit likelihood function. This likelihood includes all information from the control and signal regions, as no excess was observed during the background-only fit. Signal contamination within the control and signal regions is included in the fit with a nominal strength (μ_{sig} =1). After background estimation, the CL_s value is computed for each signal point, and those with $\mathrm{CL}_s < 0.05$ are labelled as 'excluded.' As multiple signal

Signal Region	p_0	σ_{vis} [fb]	S^{95}_{obs}	S^{95}_{exp}
SR-Gbb-A	0.37	0.31	4.6	$4.3^{+1.5}_{-0.9}$
SR-Gbb-B	0.86	0.68	10.0	12^{+5}_{-4}
SR-Gtt-0l-A	0.42	0.25	3.8	$3.9^{+1.1}_{-0.6}$
SR-Gtt-0l-B	0.02	0.90	13.3	$7.1^{+2.9}_{-1.8}$
SR-Gtt-1l-A	0.43	0.26	3.8	$3.9^{+1.2}_{-0.5}$
SR-Gtt-1l-B	0.20	0.33	4.9	$3.6^{+1.5}_{-0.4}$
SR-Gtt-1l-C	0.94	0.38	5.7	$6.9^{+2.9}_{-1.8}$

Table 7.2: Model-independent upper limits (S^{95}_{obs} & S^{95}_{exp}) obtained from the measurements made in each of the signal regions targetting Gbb and Gtt simplified models. An upper limit on the visible cross section for BSM processes within these regions (σ_{vis}), as well as the p-value describing the compatibility of each observation with the background-only hypothesis (p_0) is provided.

regions target each model differently, the region with the best expected sensitivity for a given signal point is always selected to calculate its exclusion.

This fitting procedure is applied directly to the Gbb simplified model, and the resulting exclusion contour is illustrated in figure 7.4, alongside the result previously published with 3.2 fb^{-1} of pp data collected by ATLAS during 2015. The new limits extend the expected exclusion of this search by roughly 200 GeV in the high-mass-splitting region of the signal plane, and by as much as 150 GeV further towards the diagonal region. As no excess was observed in the Gbb signal regions, the observed limit and the expected limit are compatible with each other, and are located beyond the previous limits across the entire mass plane. The most sensitive Gbb signal region at each grid point is given in figure 7.5. As anticipated, SR-Gbb-A is the most performant region throughout the high-mass-splitting and bulk regions of the Gbb phase space, while SR-Gbb-B is more sensitive closest to the diagonal region.

Including more information in the exclusion fit can lead to increased sensitivity. For the Gtt simplified model, all possible combinations of signal regions (and their corresponding control regions) between the 0L and 1L channels are orthogonal, and so may be fit simultaneously. The pair of regions with the best expected sensitivity is then selected to calculate the excluded region. This combination was also performed for the 2015 result, and so both combinations are shown in figure 7.6 for comparison. The inclusion of the 2016 data and the new $M_J^{\sum,4}$ variable have greatly improved sensitivity to the high-mass-splitting region of the signal plane, where an improvement of nearly 275 GeV has been made in the expected limit. For simpified Gtt models with a massless $\tilde{\chi}^0_1$, gluinos with masses up to 1.95 TeV are excluded. Due to the moderate excess observed in one of the Gtt signal regions, the observed

limit is consistently lower than the expected, and closely approaches the expected limit from the 2015 result in the diagonal region of the signal plane. For each of the six possible 0L+1L combinations, the most sensitive pair at each grid point is given in figure 7.7. The combination of SR-Gtt-0l-A and SR-Gtt-1l-A is the most performant for large mass splittings, though it is the combination of SR-Gtt-0l-A and SR-Gtt-1l-B which is most sensitive to the Gtt model throughout most of the excluded phase space. In the diagonal region, combinations of SR-Gtt-0l-B and SR-Gtt-1l-B or SR-Gtt-1l-C are more performant, as these regions were designed to target this phase space more specifically.

Figure 7.8 provides a summary of all results presented at the *38th International Conference on High Energy Physics* (ICHEP) held in Chicago, USA by the ATLAS [100, 101, 102, 85] and CMS [103, 104, 105, 106, 107, 108, 109] collaborations which set limits on the Gbb and Gtt simplified models. The observed limits for the Gbb models are very competitive between the two experiments, and are compatible in most places. For Gtt, both experiments are again competitive in the diagonal region of phase space, though the observed ATLAS limit in the high-mass-splitting region is nearly 150 GeV higher than the best observed limit by CMS. This is largely attributable to two factors: first, the ATLAS results included 2 fb^{-1} more integrated luminosity than the CMS results, which will increase the sensitivity of a luminosity-driven search such as those for strongly-produced SUSY. Secondly, the ATLAS 0L+1L Gtt combination extends the sensitivity to this model by nearly 150 GeV beyond what would be excluded by either channel alone (individual exclusion contours for the 0L and 1L channels of the ATLAS result are provided in appendix C). Such a combination is not feasible in the CMS result, due to the complex design of their multi-bin fit, and so this additional statistical power is not available.

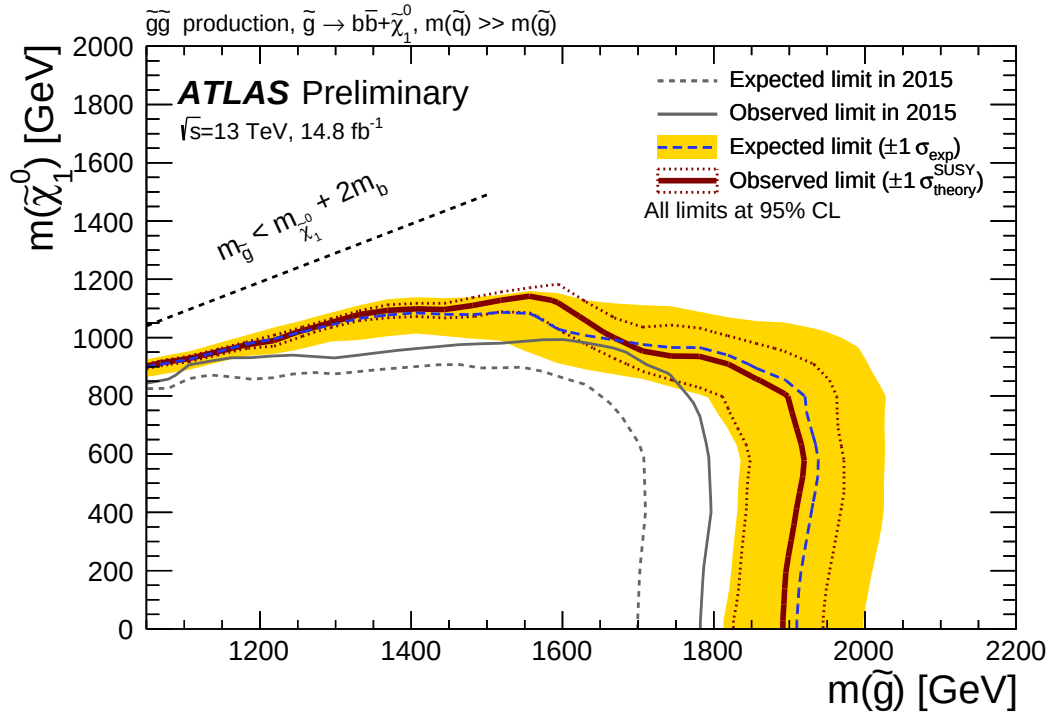

Figure 7.4: Gbb exclusion contour.

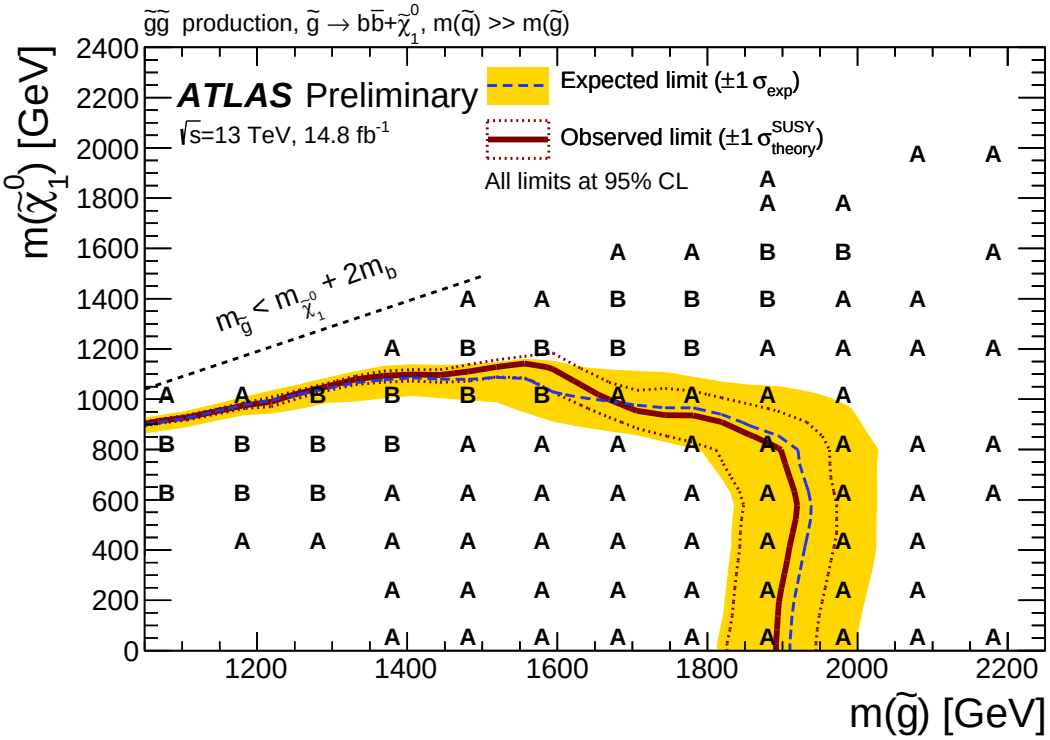

Figure 7.5: Gbb best expected signal region at each signal point.

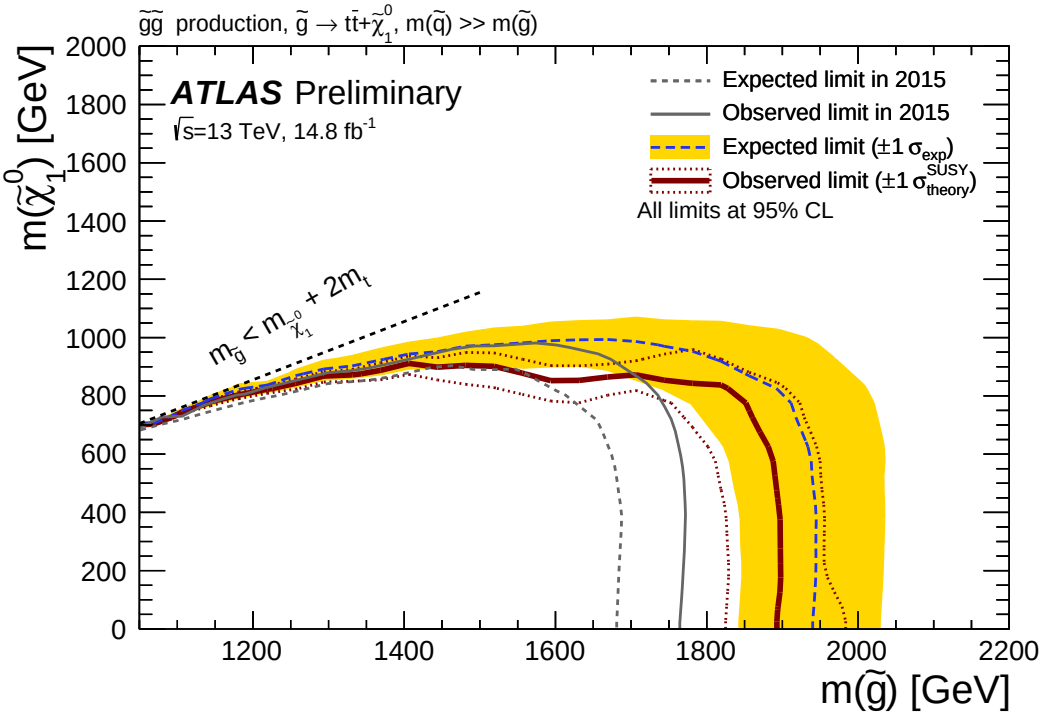

Figure 7.6: Gtt 0l+1l combination exclusion contour.

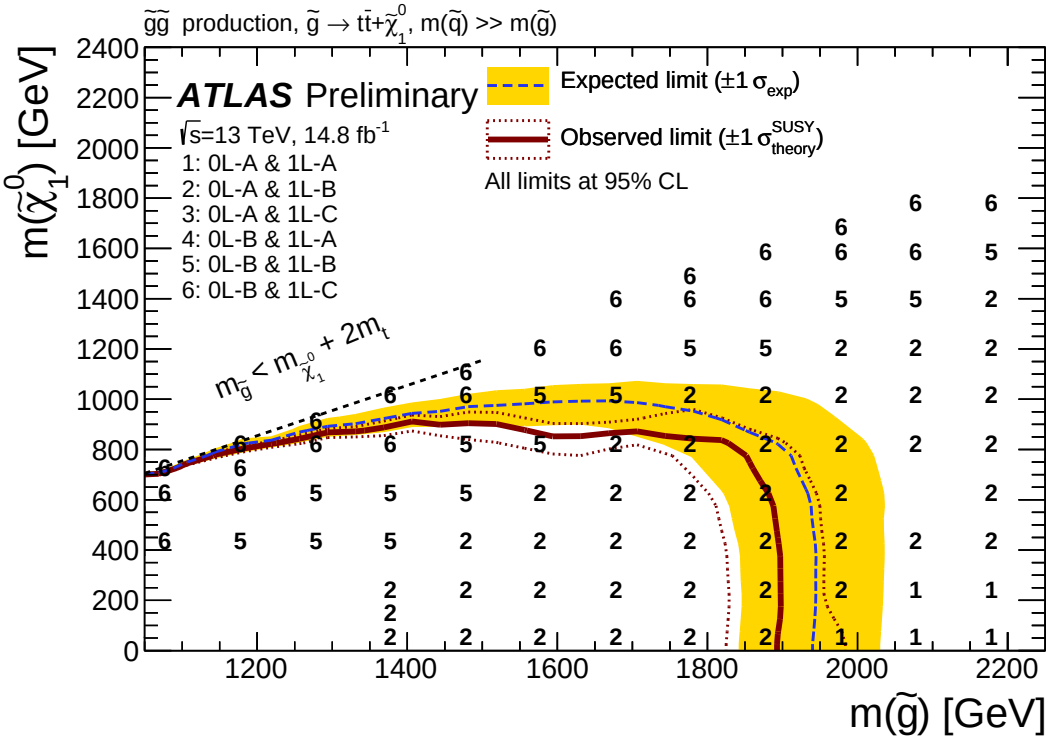

Figure 7.7: Gtt best expected signal region combination at each signal point.

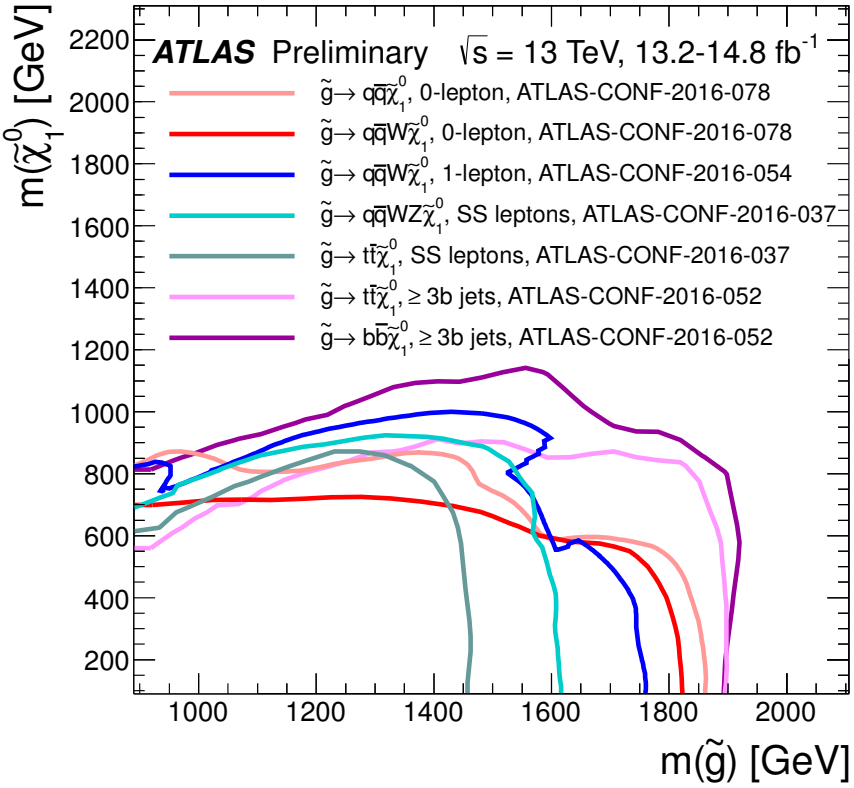

(a) ATLAS Gqq, Gbb, Gtt

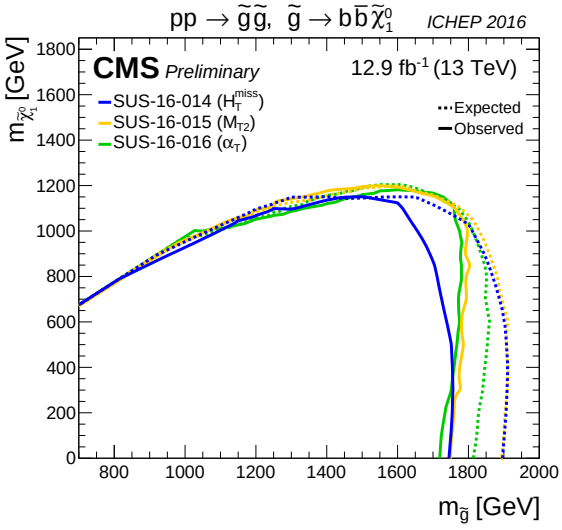

(b) CMS Gbb

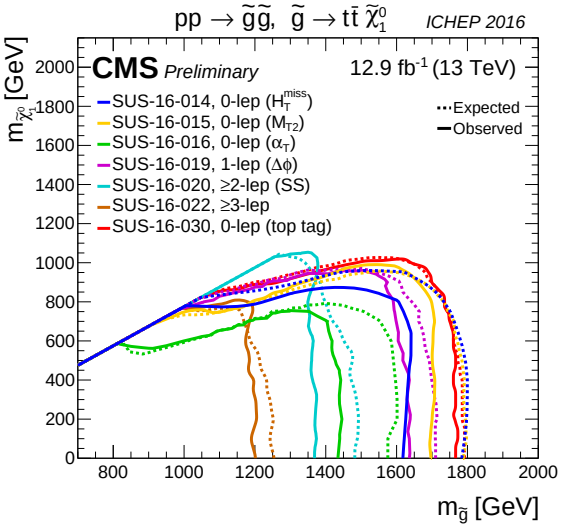

(c) CMS Gtt

Figure 7.8: Summary figures of all ATLAS [100, 101, 102, 85] and CMS [103, 104, 105, 106, 107, 108, 109] strongly-produced SUSY results presented at ICHEP 2016 in Chicago, USA which include searches targetting the Gbb or Gtt simplified models.

8

Concluding Remarks

The results presented in this thesis span the first two operational periods of the LHC, including data taken at both $\sqrt{s} = 8$ TeV and $\sqrt{s} = 13$ TeV, during 2012 and 2015-2016.

The subfield of jet substructure was first brought to mainstream attention as a method to discover a light Higgs boson decaying into $b\bar{b}$ pairs at the LHC [110], though it was born even earlier in the context of high-p_{T} W bosons produced by heavy Higgs bosons [111]. The boosted Higgs decays have still not been observed – the Standard Model Higgs has turned out to be slightly too heavy for the original approach to be optimal [112] – though the concept of reconstructing a hadronically decaying object as a single, large-R jet is more popular than ever. Detailed understanding of jet substructure techniques has risen from novelty to necessity during $\sqrt{s} = 13$ TeV operations, and the area is advancing rapidly. The studies presented in chapter 4 are in some ways more straightforward compared to the state-of-the-art, though they still represent the only published efficiency and fake rate measurements for substructure-based top tagging by the ATLAS experiment, and are still the most complete comparison of tagger performance available. This straightforward approach of identifying the decays of high-p_{T} top quarks has also been the most commonly applied method during run II; substructure-based taggers have been used extensively in the context of early searches for supersymmetry and exotica due to their simplicity and flexibility. More sophisticated techniques may be applied in certain contexts given the full run 2 dataset, though for many analyses which are not limited by the integrated luminosity, a cut-based approach with high efficiency may still be the most appropriate way to tag a top.

It is possible that searches for strongly produced supersymmetry were the most important results to have been published by the LHC during the first half of run 2. The increase in the gluino pair-production cross section between $\sqrt{s} = 8$ TeV and $\sqrt{s} = 13$ TeV provided a more rapid increase in sensitivity than will ever again be possible at the LHC, and early run 2 limits with only the first inverse femtobarns of data had already surpassed those set

with the 2012 dataset by hundreds of GeV. The sensitivity of these searches will continue to increase – both as more luminosity is collected and more sophisticated analysis techniques are introduced – though the 'low-hanging fruit' will have been harvested even by the end of 2016. Other physics will likely take centre-stage for studies performed with the run 2 dataset.

The case for naturalness has become more constrained, though the possibility of natural supersymmetry has certainly still not been excluded. Limits on the gluino mass quoted in this thesis approach 2 TeV, and limits on the direct production of stop squark pairs published concurrently approach 900 GeV [113, 114]. These limits have all been set on simplified models, which act as easily-understood benchmarks for LHC searches, though are not necessarily motivated by realistic SUSY scenarios. Branching ratios have not been treated in a realistic way: limits are set on models where gluinos decay through stop or sbottom loops 100% of the time, and those particles in turn to only final states with top or bottom quarks, respectively. Nature has (unsurprisingly) shown herself not to be this simple, and a more sophisticated treatment of these branching fractions in the future could allow for a better capacity to transfer these limits into statements about naturalness and electroweak symmetry breaking. Events like the one pictured in figure 8.1 are still intriguing in the face of initial null results. More luminosity will certainly make more sophisticated interpretations of the ATLAS data possible, and so there will always be new paths for strong SUSY searches to follow still at 13 TeV.

The top sector remains a promising area to probe, and the puzzle of the large top-Higgs Yukawa coupling is ripe for investigation at the LHC. Precision measurements of quantities in the top sector – differential cross-section measurements and the initial observations of rare Standard Model processes such as the simultaneous production of four top quarks – will provide no shortage of work to do before the High-Luminosity LHC project begins.

The results of the second LHC run will likely shape the future of high energy physics for the forseeable future. While no major discovery has yet been made, the possibility is still strong. I look forward to what there is in store.

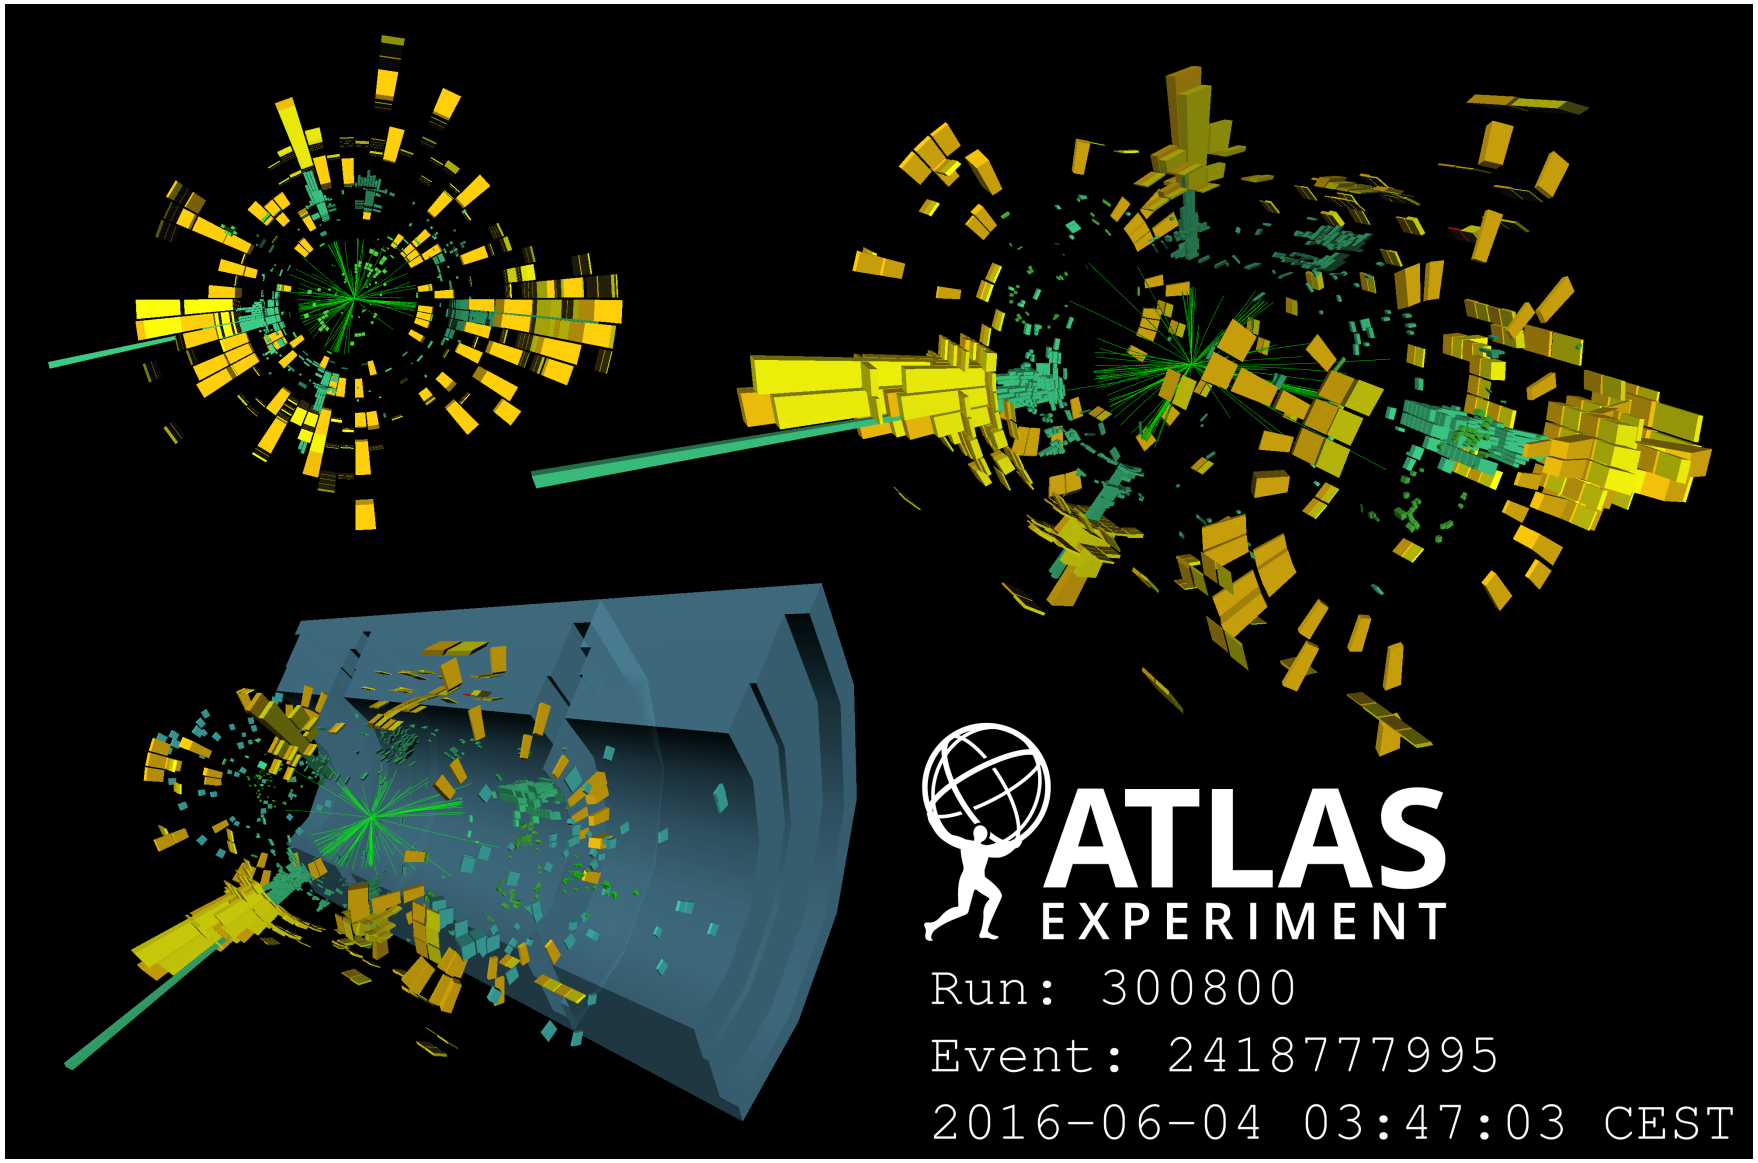

Figure 8.1: Visualisation of an event satisfying the signal criteria outlined by the search for gluino pair production, particularly SR-Gbb-A, SR-Gtt-0l-A and SR-Gtt-0l-B. This event contains 3 b-tagged jets, $E_{\mathrm{T}}^{\mathrm{miss}} > 500$ GeV and $\mathrm{m}_{\mathrm{eff}}^{\mathrm{incl}} > 2.2$ TeV.

Appendix A

LAr Data Quality: Mini-Noise Bursts

During LHC datataking, the liquid argon (LAr) calorimeter readout systems are intermittently effected by large-scale, coherent bursts of electronic noise. This phenomena manifests itself differently within the barrel and endcap calorimeters, although the strategy to ensure that data collected for use in physics analyses is uncontaminated by such bursts is similar in both cases. An important quantity related to the discussion of noise in the LAr calorimeter is the **quality factor** (Q-factor) of a cell, a number between 0 and 65535 computed during the signal digitisation procedure which describes how closely the pulse read out from the LAr systems resembles the expected pulse shape. For physics pulses this number tends to be small; it is below 4000 for over 99% of pulses processed by the LAr read-out.

The electromagnetic end-cap calorimeters are occasionally effected by **noise bursts** (NBs), which are a well-known phenomena previously described in detail by the ATLAS collaboration [115]. NBs effect large portions of the calorimeter when they occur, and tend to be short in duration (10 μs). Events which are contaminated by an NB are removed from the good-quality dataset through the application of a **time-window veto**, which sets a flag signifying the presence of a LAr data quality defect on a per-event level, in order to maximize the amount of data retained for physics. The overall amount of data removed by such a veto is small ($\ll$1% during typical data-taking conditions).

Within the electromagnetic barrel calorimeter, more localised bursts of noise take place. These bursts, so-called **mini-noise bursts** (MNBs), are known to effect only the specific FEBs which are listed in table A.1. The duration of an MNB tends to be shorter than 10 μs, though they may occur frequently during datataking operations (as often as once per minute). A two-tiered flagging scheme is used to identify MNBs during datataking:

- **Loose MNB flags** are set if more than 5 channels associated with one of the candidate FEBs have a Q-factor $>$ 4000.
- **Tight MNB flags** are set if more than 17 channels associated with one of the candidate

FEBs have a Q-factor > 4000.

If two flagged events occur within a 10 ms coincidence window and at least one of the flags is tight, a time-veto window is defined which includes the flagged events and all events within 25 ms before and after the flags. Figure A.1 illustrates the effects of applying this veto procedure on a portion of the C-side EMB calorimeter where a FEB is exhibiting noisy behaviour during a run; once applied, the time-window veto mitigates most of the noisy behaviour observed on the problematic FEB.

Partition	**Slot**	**Feedthrough**
EMB A	10	10
EMB A	14	10
EMB A	14	31
EMB A	2	9
EMB C	14	27
EMB C	5	27
EMB C	7	22

Table A.1: The electromagnetic barrel calorimeter front-end boards (FEBs) which are known to exhibit mini-noise bursts (MNBs).

It is necessary to monitor the behaviour of the time-window veto algorithm and the MNBs themselves, in order to ensure that only the minimal amount of data is removed by the veto to ensure good data quality. Monitoring histograms are published in the ATLAS data quality monitoring web interface after each run is processed by the tier 0 computing center, which allow MNBs to be identified and the performance of the time-veto algorithm to be verified. An example of such a plot is given in figure A.2, which shows the fraction of events with a loose or tight flag set for FEBs in the A-side of the EMB calorimeter during ATLAS run 307732.

Figure A.3 shows the percentage of loose and tight flags which are set by each of the problematic FEBs, as a function of the ATLAS run number during 2016 operations. Based on these plots, it is possible to conclude that very few (one or two) noisy FEBs tend to drive the flagging procedure, which sets veto periods when other FEBs coherently exhibit noisy behaviour. In particular, the FEBs in EMBC feedthrough 27 / slot 14 and EMBA feedthrough 10 / slot 14 produce the majority of MNB flags, though usually not simultaneously over the course of operations.

A comparison of cell-level observables was made for these two FEBs and the EMBC FEB in feedthrough 10 / slot 05, which is not known to produce noisy behaviour, across three

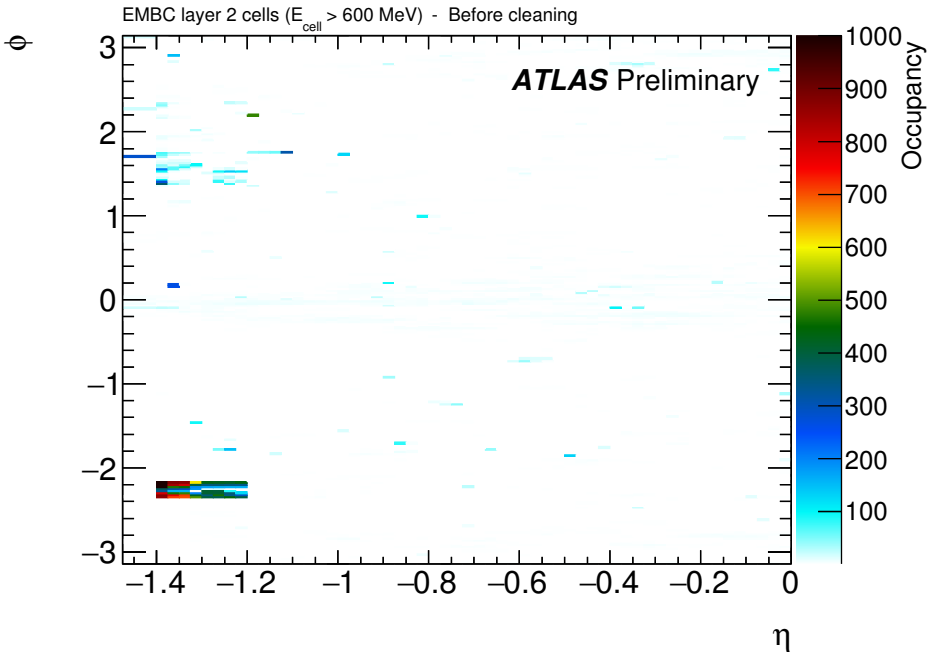

(a)

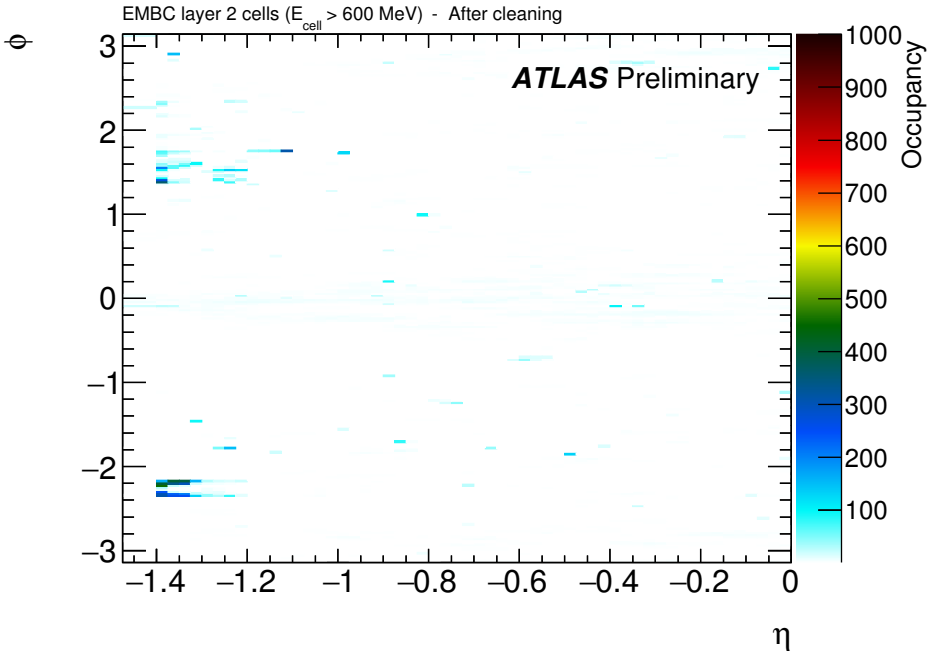

(b)

Figure A.1: The occupancy of cells in layer 2 of the C-side electromagnetic barrel calorimeter with an energy above 600 GeV during ATLAS run 306310 (455 pb^{-1}). In (a), the effects of a mini noise burst on a specific FEB are visible. After the time-window veto procedure has been applied, (b) shows that most of the noise has been mitigated. Some residual noise persists due to the flagging thresholds selected for the time window veto procedure.

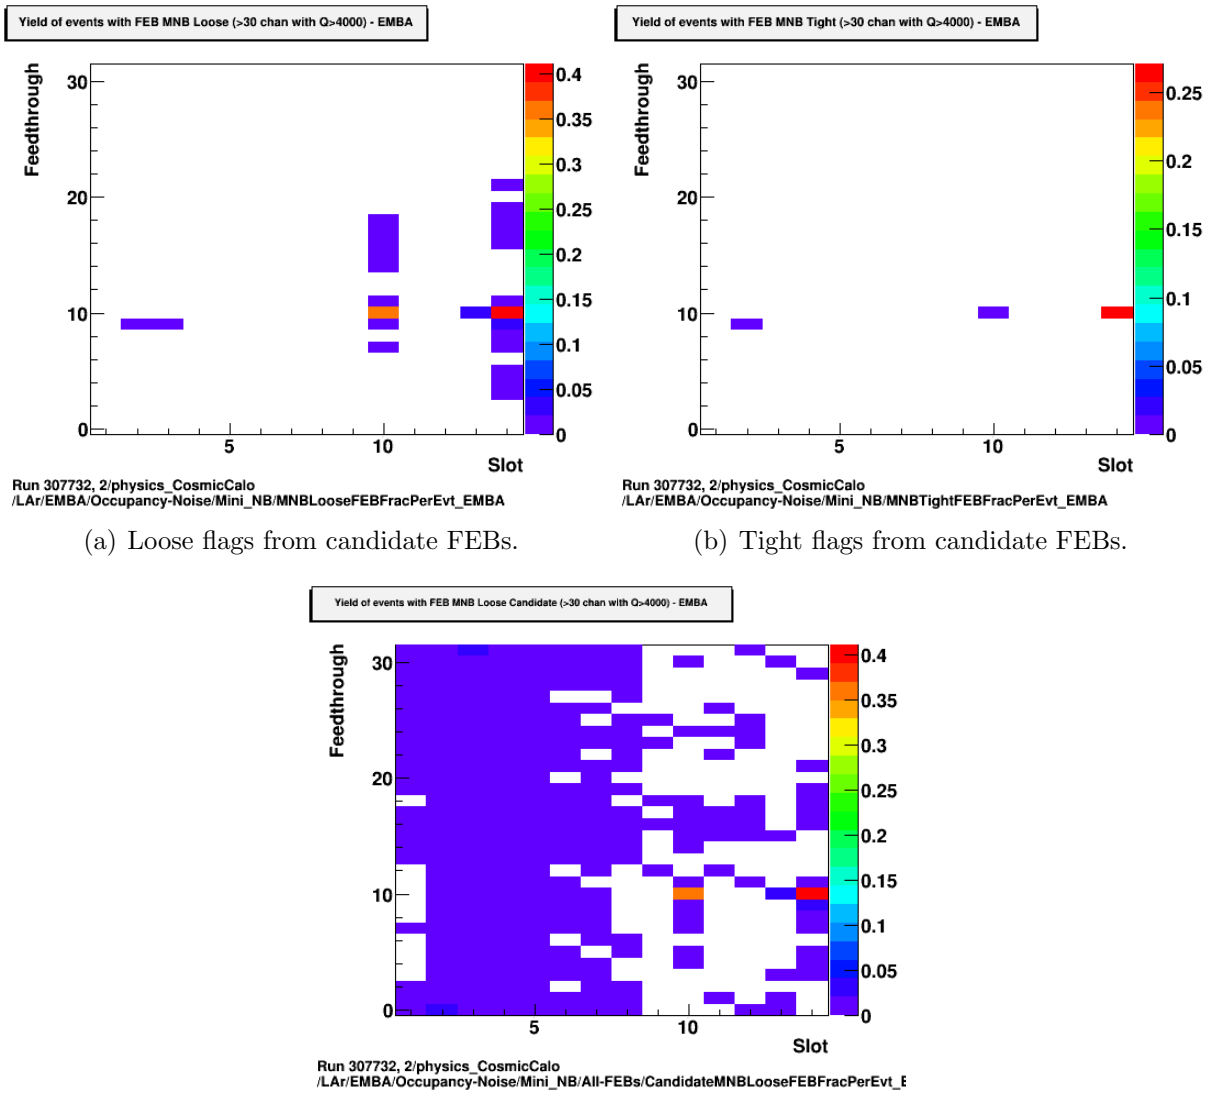

(a) Loose flags from candidate FEBs.

(b) Tight flags from candidate FEBs.

(c) Loose flags from all FEBs.

Figure A.2: Candidate FEBs flagged as exhibiting (a) loose and (b) tight mini-noise burst (MNB) activity, during ATLAS run 307732. Figure (c) shows the distribution of all FEBs passing the loose flagging criteria, allowing new candidate FEBs to be identified. Histograms such as these are automatically generated at the CERN tier 0 computing center and uploaded to the ATLAS data quality monitoring web interface after the run has been processed, in order to allow MNB activity in the calorimeter to be monitored.

runs: run 307732, in which both EMBC 27/14 and EMBA 10/14 were noisy; run 305723, in which only EMBC 27/14 was noisy; and run 307716, in which only EMBA 10/14 was noisy. The cell energy, significance and Q-factor are shown in figure A.4. Some differences are observed between the two noisy FEBs: EMBA 10/14 tends to produce long tails at high significance values while EMBC 27/14 does not, and EMBC 26/14 tends to have a larger population of cells with low Q-factors in the run where both FEBs are noisy. Both noisy FEBs exhibit an energy offset of roughly 150 MeV across all three runs, *i.e.* whether or not the FEB is producing MNBs during the run.

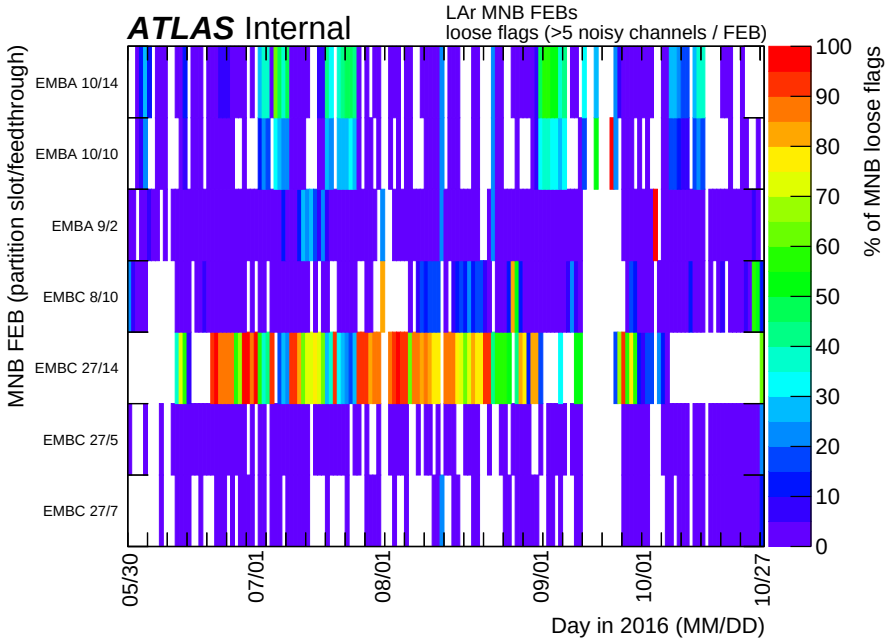

(a)

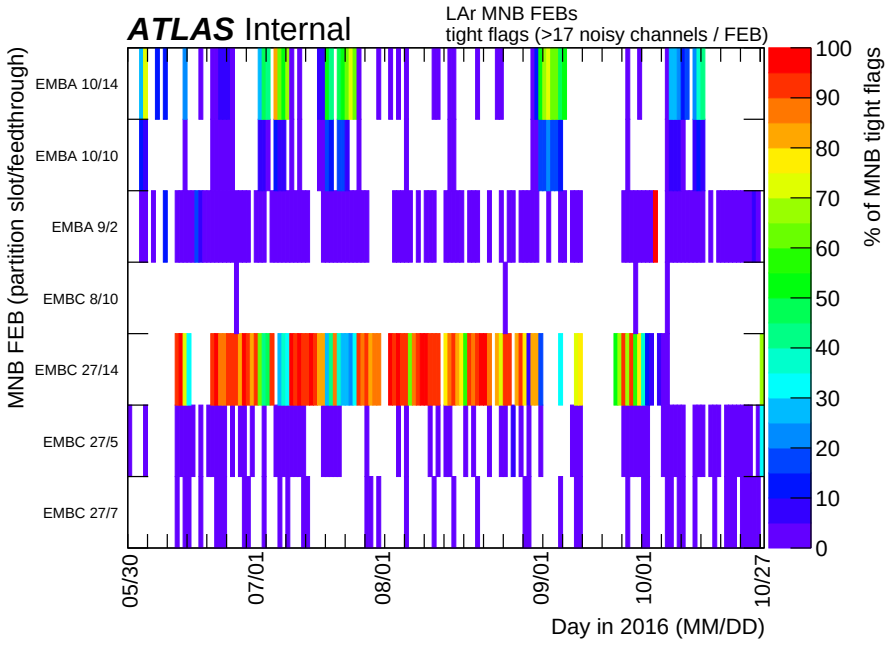

(b)

Figure A.3: Percentage of (a) loose and (b) tight mini-noise burst (MNB) flags set by each problematic front-end board (FEB), as a function of run number during 2016. Certain FEBs appear to drive the flagging of MNBs during datataking, though the FEB which sets the most flags changes with time.

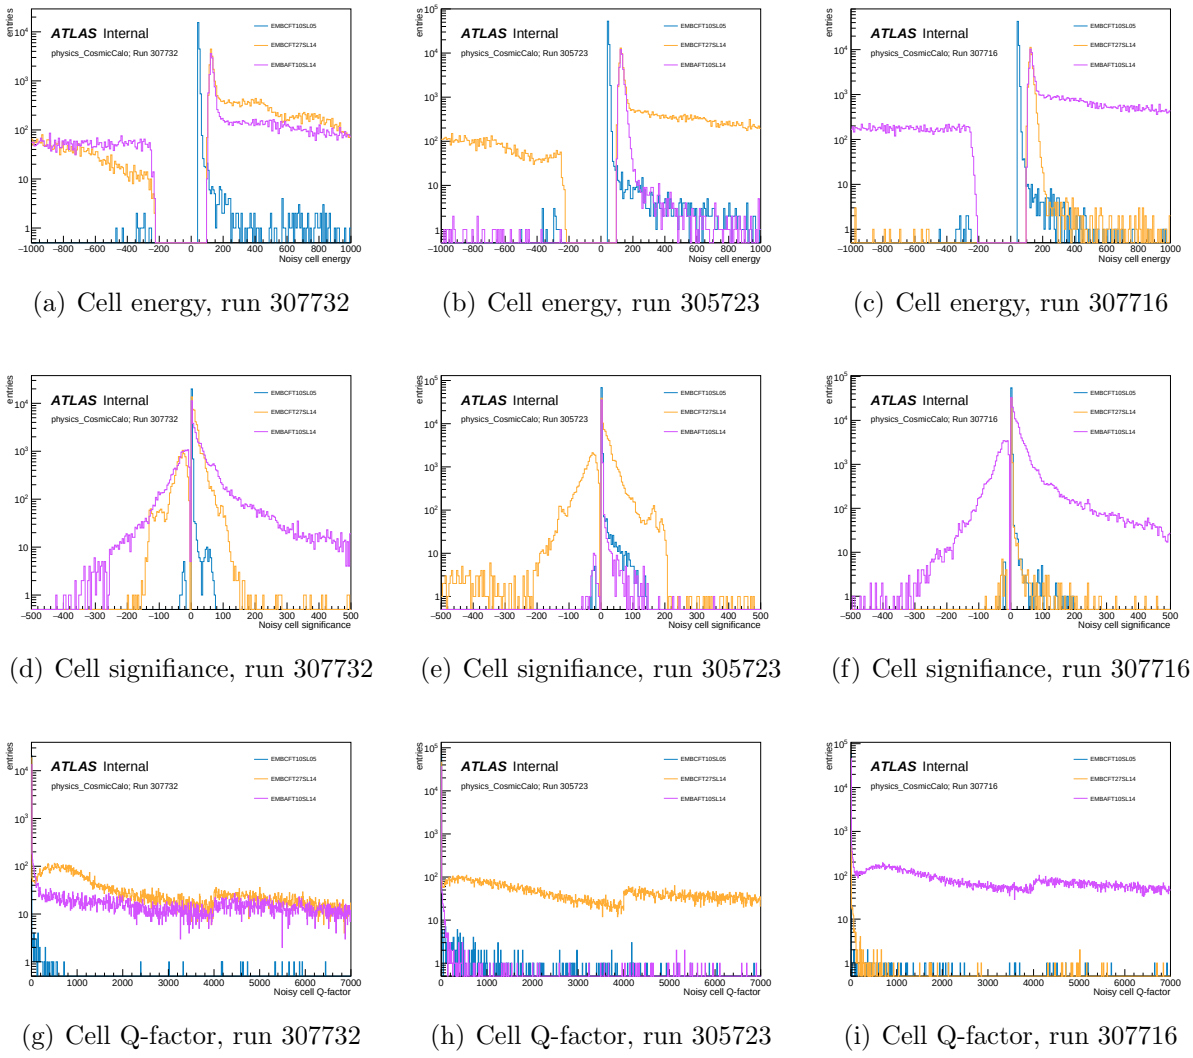

(a) Cell energy, run 307732 (b) Cell energy, run 305723 (c) Cell energy, run 307716

(d) Cell signifiance, run 307732 (e) Cell signifiance, run 305723 (f) Cell signifiance, run 307716

(g) Cell Q-factor, run 307732 (h) Cell Q-factor, run 305723 (i) Cell Q-factor, run 307716

Figure A.4: Cell-level observables for two front-end boards (FEBs) which exhibit mini noise bursts (MNBs), and one which does not. The cell energy, significance and quality factor (Q-factor) are shown across three runs: run 307732, in which both EMBC 27/14 and EMBA 10/14 were noisy; run 305723, in which only EMBC 27/14 was noisy; and run 307716, in which only EMBA 10/14 was noisy.

Appendix B

Jet Reclustering

One desirable property of an ideal jet reconstruction algorithm is some agnosticism about the exact nature of the input constituents. For instance, whether they are groups of four-vectors belonging to generator-level particles, or those of ATLAS topological clusters. It is possible to reconstruct jets using even higher-level inputs: trigger-level towers have been studied in the context of pileup suppression techniques [**?**], and the use of calibrated, small-R jets as inputs to large-R jet reconstruction has been applied in many ATLAS searches for supersymmetry and exotica throughout runs I and II of the LHC []. The latter procedure, known as **jet reclustering** [116], propagates calibrations and uncertainties from the input jet collections through to the final large-R jets, exploiting our detailed understanding of the input objects in order to provide greater flexibility at the physics analysis level. Many jet substructure techniques may be performed with reclustered large-R jets. The mass and hard structure of these objects allows for the reconstruction and identification of boosted electroweak bosons and top quarks using techniques similar to those applied to the large-R jets previously discussed in chapters 3.7 and 4.

The forward-propagation of calibrations and uncertainties to reclustered jets eliminates the need to prepare a specific calibration for each large-R jet collection considered for use in analysis. This allows individual searches and measurements to include the size of their large-R jets and the details of the grooming procedure applied to it in their optimisation procedures. When searching for high-mass resonances or massive supersymmetric particles decaying into high-p_{T} electroweak bosons and top quarks, a jet built with $R = 1.0$ may not be the optimal scale to reconstruct the boosted object.

B.1 Jet reclustering performance

The response of a jet observable, such as its transverse momentum or mass, has been previously discussed in chapter 3.7 (equation 3.13) in the context of jet energy and mass scale calibrations. Once a jet collection has been calibrated, its average p_{T} and mass values should be restored to the particle-level values in each bin used to derive the calibration. In order to verify the validity of the assumption that no additional calibration is necessary for large-R jets produced by reclustering calibrated anti-k_t $R = 0.4$ jets, an explicit study of this **closure** has been carried out in simulation across a wide range of p_{T} and mass phase space for topologies arising from light quarks and gluons, as well as for boosted top quarks. Samples of exotic, hadronically-decaying Z' bosons[1] are utilised in order to provide a population of boosted top quarks across a wide range of p_{T}. The particular samples for these studies are outlined in table B.1.

Process	ME Generator	ME PDF	Fragmentation	UE Tune
QCD multijets	Pythia 8	NNPDF23LO	Pythia 8	A14 [117]
$Z' \to t\bar{t}$	Pythia 8	NNPDF23LD	Pythia 8	A14

Table B.1: Monte Carlo samples used to study the performance of jet reclustering.

Reclustered jets are built from anti-k_t $R = 0.4$ jets with at least 25 GeV of transverse momentum, which have been calibrated at the electromagnetic scale. These small-R jets are used as inputs to the anti-k_t algorithm with $R = 1.0$ in order to produce large-R reclustered jets of the same size as the 'conventional' large-R jets used by ATLAS. These reclustered jets are then trimmed by discarding any constituents (in this case, the input small-R jets) which fail to possess at least 5% of the untrimmed reclustered jet's transverse momentum. Truth-level reclustered jets are built by adding the four-vectors of anti-k_t $R = 0.4$ truth jets, built from stable truth-level particles excluding muons and neutrinos, which are matched to within $\Delta R < 0.3$ of the detector-level constituents of a reclustered jet. Reconstructed reclustered jets are required to be within $\Delta R < 0.6$ of a truth-level reclustered jet, or else are discarded.

Comparisons are drawn directly to the conventional trimmed, anti-k_t $R = 1.0$ large-R jets built directly from locally-calibrated topological clusters. These jets are trimmed using the standard ATLAS prescription, with parameter values $f_{\mathrm{cut}} = 0.05$ and $R_{\mathrm{sub}} = 0.2$. Conventional large-R jets are required to be within $\Delta R < 0.6$ of a truth-level trimmed

[1] A simplified model is used in which the Z' possesses the properties of the standard-model Z boson, but a different mass. Only decays to top quarks are included in these simulated events.

anti-$k_t R = 1.0$ large-R jet in order to be included in this study.

For jets created by high-p_{T}, hadronically decaying top quarks, a particle-level, hadronically-decaying top quark must be located within $\Delta R < 0.75 \times R$. This requirement is equivalent to the truth-matching criterion utilised in chapter 4. All large-R jets are required to be well-isolated at both truth- and detector-level: large-R jets with another large-R truth jet or large-R reconstructed jet within $\Delta R < 2.5$ satisfying $p_{\mathrm{T}} > 150$ GeV and $|\eta| < 2.0$ are not selected for these modelling studies.

Response distributions binned in pseudorapidity are shown as a function of p_{T} and m in figures B.1—B.4. Excellent closure is observed from both reclustered and large-R jets in almost all cases, and is everywhere compatible between the two jet collections. The largest non-closures are observed for low-mass top quark jets, which typically fall below of the jet mass requirement imposed by searches and measurements interested in these topologies. This nonclosure corresponds with the small nonclosures observed at low p_{T}, and likely arise due to some products of the top quark decay falling outside of the acceptance of the $R = 1.0$ jet.

A comparison of the reclustered and large-R jet p_{T} and mass resolution, defined as the 68% inter-quartile range of the response distribution divided by twice the median of the response distribution[2], is shown in figure B.5 is shown as a function of the truth jet transverse momentum for jets produced from the decays of high-p_{T}, hadronically-decaying top quarks. Reclustered jets are observed to have both a smaller p_{T} and mass resolution than conventional trimmed large-R jets across nearly the entire p_{T} range, which suggests that these jets could be more performant than the conventional collection for analyses which identify high-p_{T} top quarks using their mass peak.

[2]In the case of a Gaussian response, this definition is equivalent to the distribution's standard deviation.

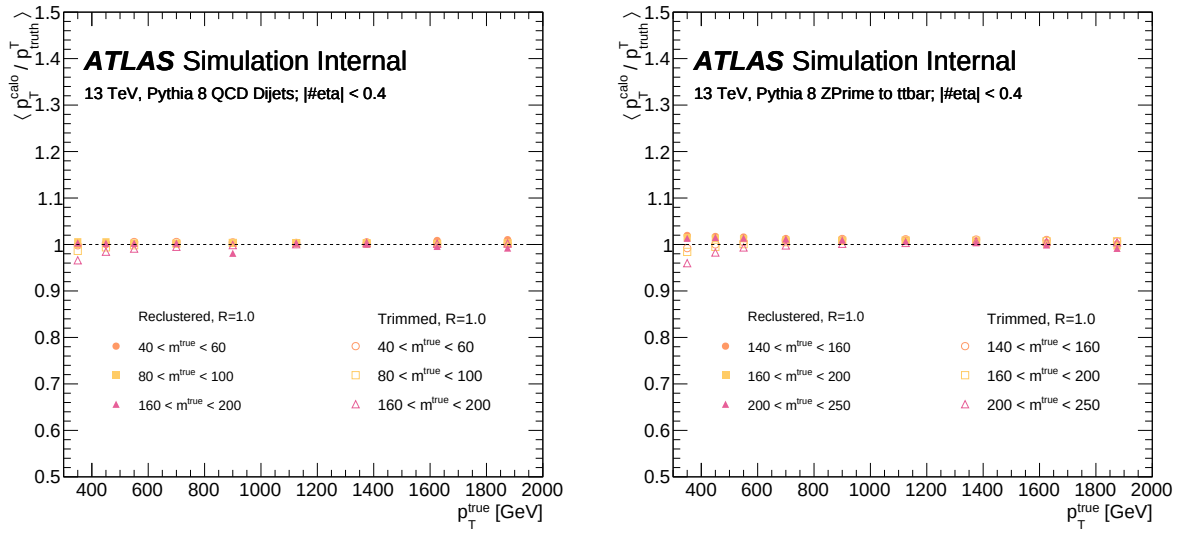

(a) dijets (b) high-p_{T} top quarks

Figure B.1: Reclustered and conventional trimmed large-R jet p_{T} responses, shown as a function of the matched truth-jet transverse momentum in selected mass bins for (a) dijet events and (b) jets created by high-p_{T}, hadronically-decaying top quarks.

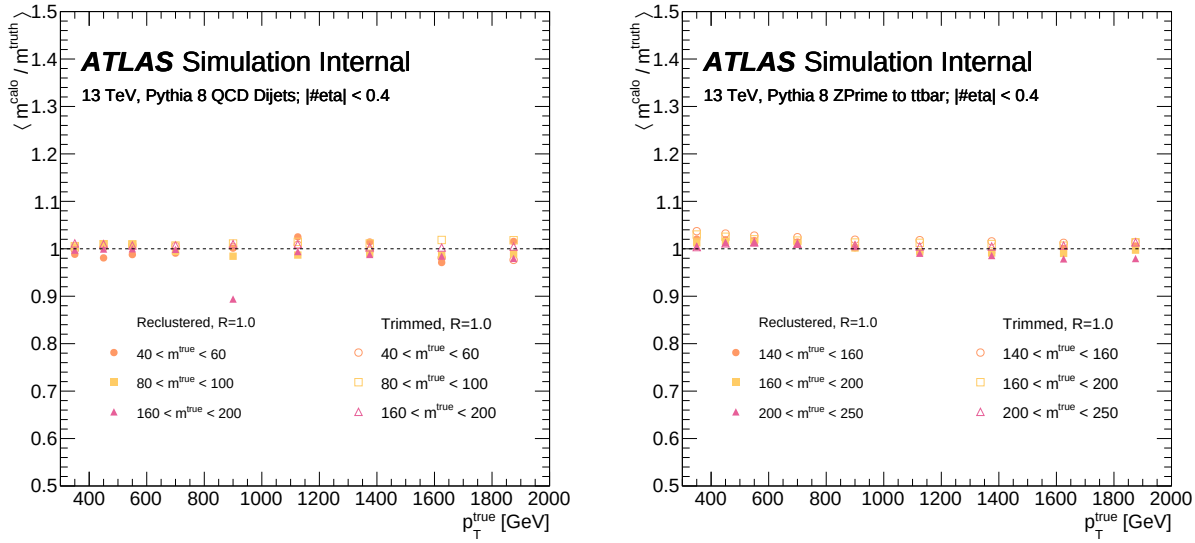

(a) dijets (b) high-p_{T} top quarks

Figure B.2: Reclustered and conventional trimmed large-R jet mass responses, shown inclusively as a function of the matched truth-jet transverse momentum for selected mass bins for (a) dijet events and (b) jets created by high-p_{T}, hadronically-decaying top quarks.

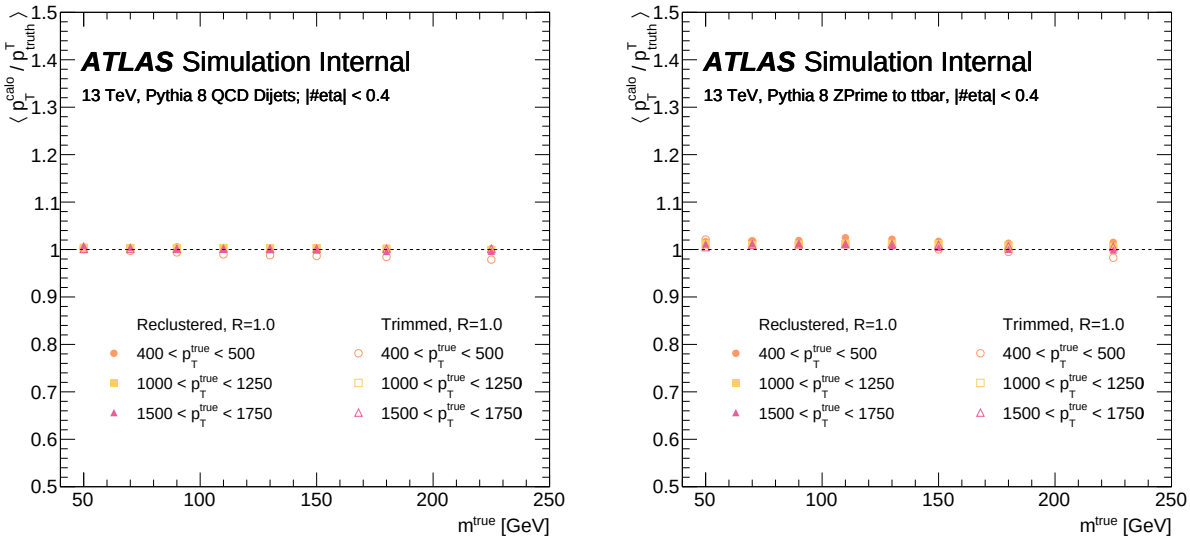

(a) dijets (b) high-p_{T} top quarks

Figure B.3: Reclustered and conventional trimmed large-R jet p_{T} responses, shown as a function of the matched truth-jet mass in selected p_{T} bins for (a) dijet events and (b) jets created by high-p_{T}, hadronically-decaying top quarks.

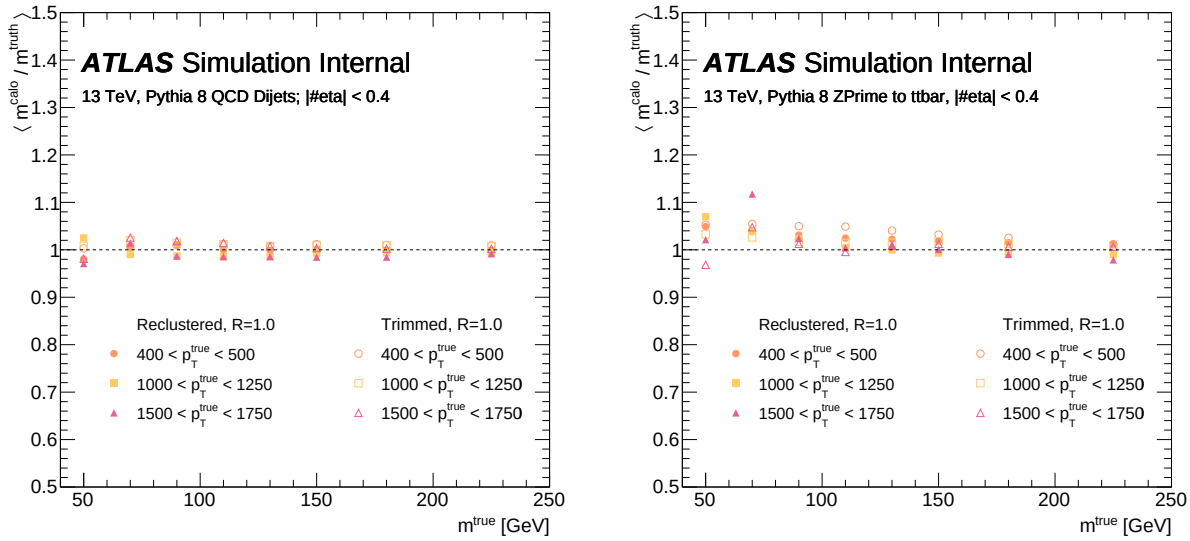

(a) dijets (b) high-p_{T} top quarks

Figure B.4: Reclustered and conventional trimmed large-R jet mass responses, shown inclusively as a function of the matched truth-jet mass in selected p_{T} bins for (a) dijet events and (b) jets created by high-p_{T}, hadronically-decaying top quarks.

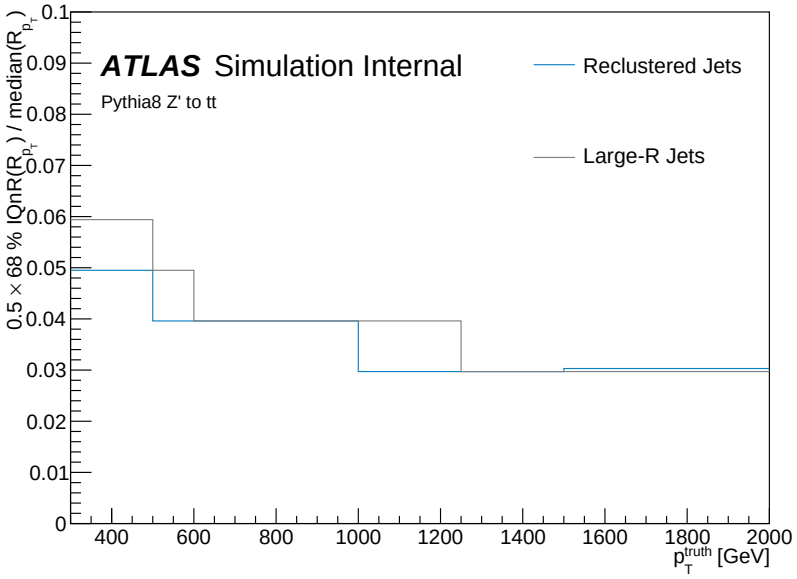

(a) top quark jet p_{T} resolution

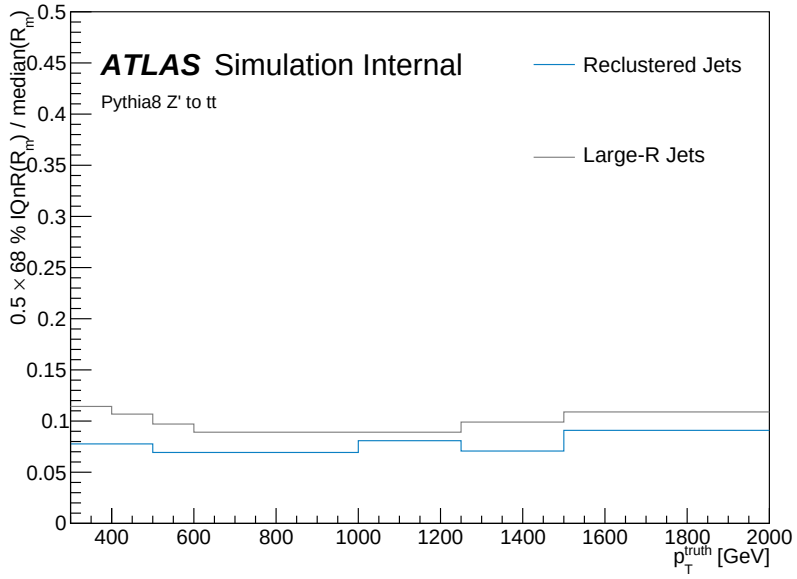

(b) top quark jet mass resolution

Figure B.5: Reclustered and trimmed large-R jet p_{T} and mass resolutions, for jets originating from high-p_{T}, hadronically-decaying top quarks.

Appendix C

Individual Gtt 0L and 1L channel results

Individual exclusion results for the Gtt-0l and 1l channels of the ATLAS search for gluino pair-production in decays through stop- and sbottom-squark loops (chapters 5-7) are provided, for comparison with the results of the statistical combination presented in chapter 7.

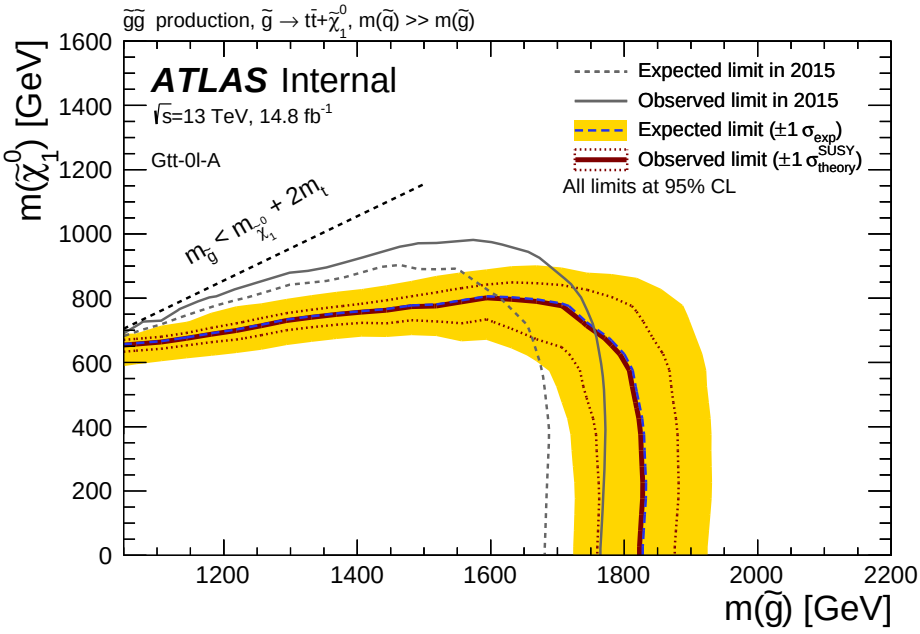

(a) Gtt-0l-A

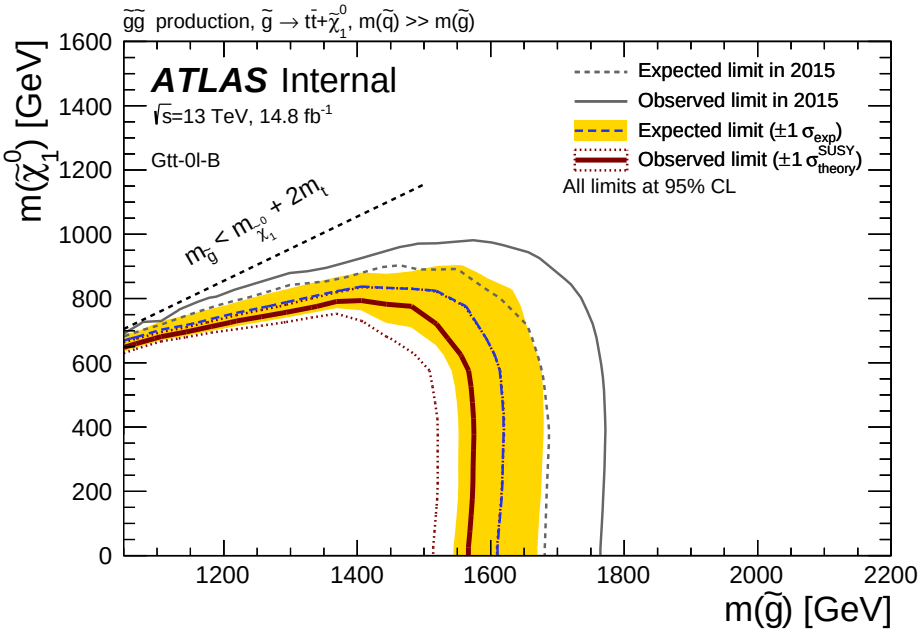

(b) Gtt-0l-B

Figure C.1: Gtt-0l signal region performance.

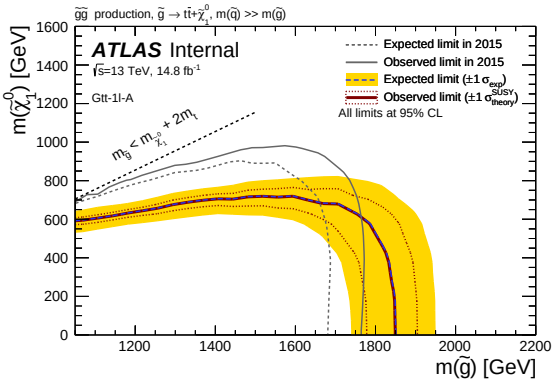

(a) Gtt-1l-A

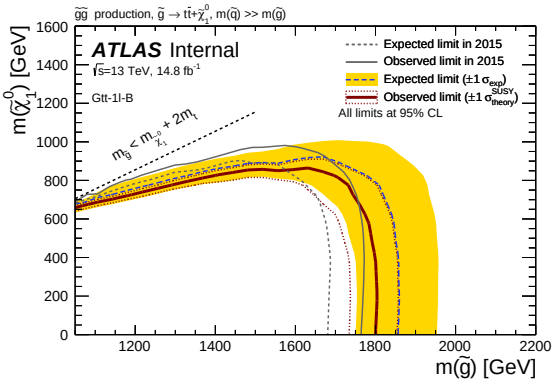

(b) Gtt-1l-B

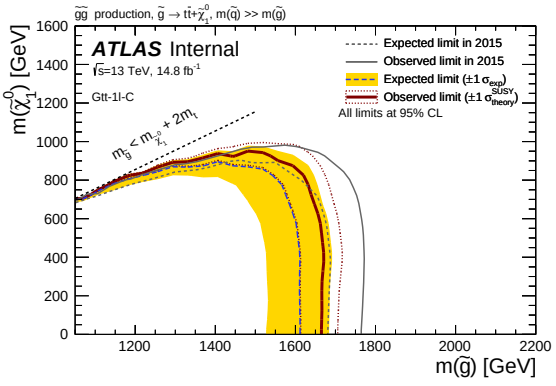

(c) Gtt-1l-C

Figure C.2: Gtt-1l signal region performance.

Bibliography

[1] Ryan P Taylor, Frank Berghaus, Franco Brasolin, Cristovao Jose Domingues Cordeiro, Ron Desmarais, Laurence Field, Ian Gable, Domenico Giordano, Alessandro Di Girolamo, John Hover, Matthew LeBlanc, Peter Love, Michael Paterson, Randall Sobie, and Alexandr Zaytsev. The Evolution of Cloud Computing in ATLAS. *Journal of Physics: Conference Series*, 664(2):022038, 2015.

[2] The ATLAS Collaboration. Identification of high transverse momentum top quarks in pp collisions at $\sqrt{s} = 8$ TeV with the ATLAS detector. *JHEP*, 06:093, 2016.

[3] The 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.*, B716:1–29, 2012.

[4] The ATLAS Collaboration. Evidence for the spin-0 nature of the Higgs boson using ATLAS data. *Phys. Lett.*, B726:120–144, 2013.

[5] The ATLAS Collaboration. Measurements of Higgs boson production and couplings in diboson final states with the ATLAS detector at the LHC. *Phys. Lett.*, B726:88–119, 2013. [Erratum: Phys. Lett.B734,406(2014)].

[6] C. Patrignani et al. Review of Particle Physics. *Chin. Phys.*, C40(10):100001, 2016.

[7] D. M. Webber, V. Tishchenko, Q. Peng, S. Battu, R. M. Carey, D. B. Chitwood, J. Crnkovic, P. T. Debevec, S. Dhamija, W. Earle, A. Gafarov, K. Giovanetti, T. P. Gorringe, F. E. Gray, Z. Hartwig, D. W. Hertzog, B. Johnson, P. Kammel, B. Kiburg, S. Kizilgul, J. Kunkle, B. Lauss, I. Logashenko, K. R. Lynch, R. McNabb, J. P. Miller, F. Mulhauser, C. J. G. Onderwater, J. Phillips, S. Rath, B. L. Roberts, P. Winter, and B. Wolfe. Measurement of the positive muon lifetime and determination of the fermi constant to part-per-million precision. *Phys. Rev. Lett.*, 106:041803, Jan 2011.

[8] Stephen P. Martin. A Supersymmetry primer. *Adv. Ser. Direct. High Energy Phys.*, 18(1), 1998.

[9] Gerard Jungman, Marc Kamionkowski, and Kim Griest. Supersymmetric dark matter. *Phys. Rept.*, 267:195–373, 1996.

[10] John Ellis, J.S. Hagelin, D.V. Nanopoulos, K. Olive, and M. Srednicki. Supersymmetric relics from the big bang. *Nuclear Physics B*, 238(2):453 – 476, 1984.

[11] H. Goldberg. Constraint on the photino mass from cosmology. *Phys. Rev. Lett.*, 50:1419–1422, May 1983.

[12] Torbjorn Sjostrand, Stephen Mrenna, and Peter Z. Skands. PYTHIA 6.4 Physics and Manual. *JHEP*, 05:026, 2006.

[13] G. Corcella, I. G. Knowles, G. Marchesini, S. Moretti, K. Odagiri, P. Richardson, M. H. Seymour, and B. R. Webber. HERWIG 6: an event generator for hadron emission reactions with interfering gluons (including supersymmetric processes). *JHEP*, 01:010, 2001.

[14] T. Gleisberg, Stefan. Hoeche, F. Krauss, M. Schonherr, S. Schumann, F. Siegert, and J. Winter. Event generation with SHERPA 1.1. *JHEP*, 02:007, 2009.

[15] Stefano Frixione and Bryan R. Webber. Matching NLO QCD computations and parton shower simulations. *JHEP*, 06:029, 2002.

[16] A. D. Martin, W. J. Stirling, R. S. Thorne, and G. Watt. Uncertainties on alpha(S) in global PDF analyses and implications for predicted hadronic cross sections. *Eur. Phys. J.*, C64:653–680, 2009.

[17] J. Alwall, R. Frederix, S. Frixione, V. Hirschi, F. Maltoni, O. Mattelaer, H.-S. Shao, T. Stelzer, P. Torrielli, and M. Zaro. The automated computation of tree-level and next-to-leading order differential cross sections, and their matching to parton shower simulations. *Journal of High Energy Physics*, 2014(7):79, 2014.

[18] Richard D. Ball et al. Parton distributions for the LHC Run II. *JHEP*, 04:040, 2015.

[19] R. P. Feynman. Pocono Conference. *j-PHYS-TODAY*, 1(2):8–10, June 1948.

[20] Esma Anais Mobs. The CERN accelerator complex. Complexe des accélérateurs du CERN, Oct 2016. General Photo.

[21] G Antchev, P Aspell, I Atanassov, V Avati, J Baechler, V Berardi, M Berretti, E Bossini, M Bozzo, P Brogi, E Brcken, A Buzzo, F Cafagna, M Calicchio, M G Catanesi, C Covault, T Csrg, M Deile, K Eggert, V Eremin, R Ferretti, F Ferro,

A Fiergolski, F Garcia, S Giani, V Greco, L Grzanka, J Heino, T Hilden, M R Intonti, J Kaspar, J Kopal, V Kundrt, K Kurvinen, S Lami, G Latino, R Lauhakangas, T Leszko, E Lippmaa, M Lokajcek, M Lo Vetere, F Lucas Rodrguez, M Macr, L Magaletti, A Mercadante, S Minutoli, F Nemes, H Niewiadomski, E Oliveri, F Oljemark, R Orava, M Oriunno, K sterberg, P Palazzi, J Prochzka, M Quinto, E Radermacher, E Radicioni, F Ravotti, E Robutti, L Ropelewski, G Ruggiero, H Saarikko, G Sanguinetti, A Santroni, A Scribano, W Snoeys, J Sziklai, C Taylor, N Turini, V. Vacek, M Vtek, J Welti, and J Whitmore. First measurement of the total proton-proton cross section at the LHC energy of $\sqrt{s} = 7$ TeV. *EPL*, 96:21002. 11 p, Oct 2011.

[22] The ATLAS Collaboration. Measurement of the total cross section from elastic scattering in pp collisions at $\sqrt{s} = 8$ TeV with the ATLAS detector. *Phys. Lett.*, B761:158–178, 2016.

[23] S van der Meer. Calibration of the effective beam height in the ISR. Technical Report CERN-ISR-PO-68-31, CERN, Geneva, 1968.

[24] Carlo Rubbia. Measurement of the luminosity of $p\bar{p}$ collider with a (generalized) Van der Meer Method. Technical Report CERN-PPBAR-NOTE-38, CERN, Geneva, Nov 1977.

[25] The TOTEM Collaboration. The TOTEM Experiment at the CERN Large Hadron Collider. *Journal of Instrumentation*, 3(08):S08007, 2008.

[26] The ATLAS Collaboration. *Expected performance of the ATLAS experiment: detector, trigger and physics*. CERN, Geneva, 2009.

[27] The ATLAS Collaboration. ATLAS Inner Detector: Technical Design Report. Vol. 1. 1997.

[28] The ATLAS Collaboration. ATLAS Inner Detector: Technical Design Report. Vol. 2. 1997.

[29] The ATLAS Collaboration. ATLAS Liquid Argon Calorimeter: Technical Design Report. 1996.

[30] The ATLAS Collaboration. *ATLAS Muon Spectrometer: Technical Design Report*. Technical Design Report ATLAS. CERN, Geneva, 1997.

[31] R. Fruhwirth. Application of Kalman filtering to track and vertex fitting. *Nucl. Instrum. Meth.*, A262:444–450, 1987.

[32] The ATLAS Collaboration. The ATLAS Experiment at the CERN Large Hadron Collider. *Journal of Instrumentation*, 3(08):S08003, 2008.

[33] The ATLAS Collaboration. Readiness of the ATLAS liquid argon calorimeter for LHC collisions. *The European Physical Journal C*, 70(3):723–753, 2010.

[34] The ATLAS Collaboration. *ATLAS Tile Calorimeter: Technical Design Report*. CERN, Geneva, 1996.

[35] The ATLAS Collaboration. Performance of the ATLAS Trigger System in 2015. Technical Report arXiv:1611.09661. CERN-EP-2016-241, CERN, Geneva, Nov 2016.

[36] Cloyce D. Spradling. Spec cpu2006 benchmark tools. *SIGARCH Comput. Archit. News*, 35(1):130–134, March 2007.

[37] M Alef and I Gable. HEP specific benchmarks of virtual machines on multi-core CPU architectures. 219(5):052015, 2010.

[38] P Buncic, C Aguado Sanchez, J Blomer, L Franco, A Harutyunian, P Mato, and Y Yao. CernVM a virtual software appliance for LHC applications. *Journal of Physics: Conference Series*, 219(4):042003, 2010.

[39] Michele Michelotto, Manfred Alef, Alejandro Iribarren, Helge Meinhard, Peter Wegner, Martin Bly, Gabriele Benelli, Franco Brasolin, Hubert Degaudenzi, Alessandro De Salvo, Ian Gable, Andreas Hirstius, and Peter Hristov. A comparison of HEP code with SPEC benchmarks on multi-core worker nodes. *Journal of Physics: Conference Series*, 219(5):052009, 2010.

[40] Peter Speckmayer, Tancredi Carli, and Christian Wolfgang Fabjan. *Energy Measurement of Hadrons with the CERN ATLAS Calorimeter*. PhD thesis, Vienna, Tech. U., Vienna, 2008. Presented on 18 Jun 2008.

[41] Stephen D. Ellis, Zoltan Kunszt, and Davison E. Soper. One-jet inclusive cross section at order α_s^3. Gluons only. *Phys. Rev. D*, 40:2188–2222, Oct 1989.

[42] John 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) Snowmass, Colorado, June 25-July 13, 1990*, pages 0134–136, 1990.

[43] Gregory Soyez. The SISCone and anti-k_t jet algorithms. In *Proceedings, 16th International Workshop on Deep Inelastic Scattering and Related Subjects (DIS 2008): London, UK, April 7-11, 2008*, page 178, 2008.

[44] Stephen D. Ellis and Davison E. Soper. Successive combination jet algorithm for hadron collisions. *Phys. Rev.*, D48:3160–3166, 1993.

[45] S. Catani, Yu.L. Dokshitzer, M.H. Seymour, and B.R. Webber. Longitudinally-invariant k-clustering algorithms for hadron-hadron collisions. *Nuclear Physics B*, 406(1):187 – 224, 1993.

[46] Yuri L. Dokshitzer, G. D. Leder, S. Moretti, and B. R. Webber. Better jet clustering algorithms. *JHEP*, 08:001, 1997.

[47] Matteo Cacciari, Gavin P. Salam, and Gregory Soyez. The anti-k_t jet clustering algorithm. *JHEP*, 04:063, 2008.

[48] Matteo Cacciari, Gavin P. Salam, and Gregory Soyez. FastJet User Manual. *Eur. Phys. J.*, C72:1896, 2012.

[49] Matteo Cacciari and Gavin P. Salam. Dispelling the N^3 myth for the k_t jet-finder. *Phys. Lett.*, B641:57–61, 2006.

[50] The ATLAS Collaboration. Public plots: 2013 jes uncertainty. Technical report, CERN, Geneva, August 2016.

[51] The ATLAS Collaboration. Public plots: Jet energy scale uncertainties updated for ichep 2016 using full 13 TeV 2015 dataset. Technical report, CERN, Geneva, August 2016.

[52] The ATLAS Collaboration. Performance of b-Jet Identification in the ATLAS Experiment. *JINST*, 11(04):P04008, 2016.

[53] The ATLAS Collaboration. Expected performance of the ATLAS b-tagging algorithms in run 2. Technical Report ATL-PHYS-PUB-2015-022, CERN, Geneva, Jul 2015.

[54] The ATLAS Collaboration. Optimisation of the ATLAS b-tagging performance for the 2016 LHC Run. Technical Report ATL-PHYS-PUB-2016-012, CERN, Geneva, Jun 2016.

[55] Mrinal Dasgupta, Lorenzo Magnea, and Gavin P. Salam. Non-perturbative QCD effects in jets at hadron colliders. *JHEP*, 02:055, 2008.

[56] David Krohn, Jesse Thaler, and Lian-Tao Wang. Jet Trimming. *JHEP*, 02:084, 2010.

[57] The ATLAS Collaboration. Performance of jet substructure techniques for large-R jets

in proton-proton collisions at $\sqrt{s} = 7$ TeV using the ATLAS detector. *J. High Energy Phys.*, 09(arXiv:1306.4945. CERN-PH-EP-2013-069):076. 93 p, Jun 2013.

[58] The ATLAS Collaboration. Measurement of the muon reconstruction performance of the ATLAS detector using 2011 and 2012 LHC proton-proton collision data. *Eur. Phys. J.*, C74(11):3130, 2014.

[59] The ATLAS Collaboration. Identification and energy calibration of hadronically decaying tau leptons with the ATLAS experiment in *pp* collisions at $\sqrt{s} = 8$ TeV. *Eur. Phys. J. C*, 75:303. 21 p, Dec 2014.

[60] Photon identification in 2015 ATLAS data. Technical Report ATL-PHYS-PUB-2016-014, CERN, Geneva, Aug 2016.

[61] Commissioning of the reconstruction of hadronic tau lepton decays in ATLAS using *pp* collisions at $\sqrt{s} = 13$ TeV. Technical Report ATL-PHYS-PUB-2015-025, CERN, Geneva, Jul 2015.

[62] The ATLAS Collaboration. Performance of missing transverse momentum reconstruction for the ATLAS detector in the first proton-proton collisions at at $\sqrt{s} = 13$ TeV. Technical Report ATL-PHYS-PUB-2015-027, CERN, Geneva, Jul 2015.

[63] Tilman Plehn, Michael Spannowsky, Michihisa Takeuchi, and Dirk Zerwas. Stop Reconstruction with Tagged Tops. *JHEP*, 10:078, 2010.

[64] Tilman Plehn, Gavin P. Salam, and Michael Spannowsky. Fat Jets for a Light Higgs. *Phys. Rev. Lett.*, 104:111801, 2010.

[65] Davison E. Soper and Michael Spannowsky. Finding physics signals with shower deconstruction. *Phys. Rev.*, D84:074002, 2011.

[66] Davison E. Soper and Michael Spannowsky. Finding top quarks with shower deconstruction. *Phys. Rev.*, D87:054012, 2013.

[67] The ATLAS Collaboration. Identification of high transverse momentum top quarks in *pp* collisions at $\sqrt{s} = 8$ TeV with the ATLAS detector. Technical Report ATLAS-CONF-2015-036, CERN, Geneva, Aug 2015.

[68] Hung-Liang Lai, Marco Guzzi, Joey Huston, Zhao Li, Pavel M. Nadolsky, Jon Pumplin, and C. P. Yuan. New parton distributions for collider physics. *Phys. Rev.*, D82:074024, 2010.

[69] Stefano Frixione, Eric Laenen, Patrick Motylinski, Bryan R. Webber, and Chris D. White. Single-top hadroproduction in association with a W boson. *JHEP*, 07:029, 2008.

[70] A. D. Martin, W. J. Stirling, R. S. Thorne, and G. Watt. Parton distributions for the LHC. *Eur. Phys. J.*, C63:189–285, 2009.

[71] ATLAS Collaboration. Search for $t\bar{t}$ resonances in the lepton plus jets final state with ATLAS using 4.7 fb^{-1} of pp collisions at $\sqrt{s} = 7$ TeV. *Phys. Rev.*, D88(1):012004, 2013.

[72] ATLAS Collaboration. Search for $W' \to tb \to qqbb$ decays in pp collisions at $\sqrt{s} = 8$ TeV with the ATLAS detector. *Eur. Phys. J.*, C75(4):165, 2015.

[73] The ATLAS Collaboration. Search for top squark pair production in final states with one isolated lepton, jets, and missing transverse momentum in $\sqrt{s} = 8$ TeV pp collisions with the ATLAS detector. *JHEP*, 11:118, 2014.

[74] The ATLAS Collaboration. Improved luminosity determination in pp collisions at $\sqrt(s) = 7$ TeV using the ATLAS detector at the LHC. *Eur. Phys. J.*, C73(8):2518, 2013.

[75] The ATLAS Collaboration. Electron reconstruction and identification efficiency measurements with the ATLAS detector using the 2011 LHC proton-proton collision data. *Eur. Phys. J.*, C74(7):2941, 2014.

[76] The ATLAS Collaboration. Electron efficiency measurements with the ATLAS detector using the 2012 LHC proton-proton collision data. Technical Report ATLAS-CONF-2014-032, CERN, Geneva, Jun 2014.

[77] Matteo Cacciari, Michal Czakon, Michelangelo Mangano, Alexander Mitov, and Paolo Nason. Top-pair production at hadron colliders with next-to-next-to-leading logarithmic soft-gluon resummation. *Phys. Lett.*, B710:612–622, 2012.

[78] M. Beneke, P. Falgari, S. Klein, and C. Schwinn. Hadronic top-quark pair production with NNLL threshold resummation. *Nucl. Phys.*, B855:695–741, 2012.

[79] Peter Brnreuther, Michal Czakon, and Alexander Mitov. Percent Level Precision Physics at the Tevatron: First Genuine NNLO QCD Corrections to $q\bar{q} \to t\bar{t} + X$. *Phys. Rev. Lett.*, 109:132001, 2012.

[80] Michal Czakon and Alexander Mitov. NNLO corrections to top-pair production at hadron colliders: the all-fermionic scattering channels. *JHEP*, 12:054, 2012.

[81] Michal Czakon and Alexander Mitov. NNLO corrections to top pair production at hadron colliders: the quark-gluon reaction. *JHEP*, 01:080, 2013.

[82] Micha Czakon, Paul Fiedler, and Alexander Mitov. Total Top-Quark Pair-Production Cross Section at Hadron Colliders Through $O(\frac{4}{S})$. *Phys. Rev. Lett.*, 110:252004, 2013.

[83] Michal Czakon and Alexander Mitov. Top++: A Program for the Calculation of the Top-Pair Cross-Section at Hadron Colliders. *Comput. Phys. Commun.*, 185:2930, 2014.

[84] Richard D. Ball et al. Parton distributions with LHC data. *Nucl. Phys.*, B867:244–289, 2013.

[85] The ATLAS Collaboration. Search for pair production of gluinos decaying via top or bottom squarks in events with b-jets and large missing transverse momentum in pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector. Technical Report ATLAS-CONF-2016-052, CERN, Geneva, Aug 2016.

[86] The ATLAS Collaboration. Search for pair-production of gluinos decaying via stop and sbottom in events with b-jets and large missing transverse momentum in $\sqrt{s} = 13$ TeV pp collisions with the ATLAS detector. Technical Report ATLAS-CONF-2015-067, CERN, Geneva, Dec 2015.

[87] The ATLAS Collaboration. Search for pair production of gluinos decaying via stop and sbottom in events with b-jets and large missing transverse momentum in pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector. *Phys. Rev. D*, 94:032003, Aug 2016.

[88] The ATLAS Collaboration. The simulation principle and performance of the ATLAS fast calorimeter simulation FastCaloSim. Technical Report ATL-PHYS-PUB-2010-013, CERN, Geneva, Oct 2010.

[89] The ATLAS Collaboration. Performance of the Fast ATLAS Tracking Simulation (FATRAS) and the ATLAS Fast Calorimeter Simulation (FastCaloSim) with single particles. Technical Report ATL-SOFT-PUB-2014-001, CERN, Geneva, Mar 2014.

[90] The ATLAS Collaboration. Performance of pile-up mitigation techniques for jets in pp collisions at $\sqrt{s} = 8$ TeV using the ATLAS detector. *Eur. Phys. J.*, C76(11):581, 2016.

[91] The ATLAS Collaboration. Electron efficiency measurements with the ATLAS detector using the 2015 LHC proton-proton collision data. Technical Report ATLAS-CONF-2016-024, CERN, Geneva, Jun 2016.

[92] The ATLAS Collaboration. Muon reconstruction performance in early $\sqrt{s} = 13$ TeV data. Technical Report ATL-PHYS-PUB-2015-037, CERN, Geneva, Aug 2015.

[93] The ATLAS Collaboration. A method for the construction of strongly reduced representations of ATLAS experimental uncertainties and the application thereof to the jet energy scale. Technical Report ATL-PHYS-PUB-2015-014, CERN, Geneva, Jul 2015.

[94] The ATLAS Collaboration. Jet Calibration and Systematic Uncertainties for Jets Reconstructed in the ATLAS Detector at $\sqrt{s} = 13$ TeV. Technical Report ATL-PHYS-PUB-2015-015, CERN, Geneva, Jul 2015.

[95] Peter Zeiler Skands. Tuning Monte Carlo Generators: The Perugia Tunes. *Phys. Rev.*, D82:074018, 2010.

[96] The ATLAS Collaboration. Measurements of fiducial cross-sections for $t\bar{t}$ production with one or two additional b-jets in pp collisions at $\sqrt{s} = 8$ TeV using the ATLAS detector. *Eur. Phys. J.*, C76(1):11, 2016.

[97] Kyle Cranmer, George Lewis, Lorenzo Moneta, Akira Shibata, and Wouter Verkerke. HistFactory: A tool for creating statistical models for use with RooFit and RooStats. Technical Report CERN-OPEN-2012-016, New York University, New York, Jan 2012.

[98] M. Baak, G. J. Besjes, D. Cte, A. Koutsman, J. Lorenz, and D. Short. HistFitter software framework for statistical data analysis. *Eur. Phys. J.*, C75:153, 2015.

[99] J. Neyman and E. S. Pearson. On the problem of the most efficient tests of statistical hypotheses. *Philosophical Transactions of the Royal Society of London. Series A, Containing Papers of a Mathematical or Physical Character*, 231:289–337, 1933.

[100] The ATLAS Collaboration. Further searches for squarks and gluinos in final states with jets and missing transverse momentum at $\sqrt{s} = 13$ TeV with the ATLAS detector. Technical Report ATLAS-CONF-2016-078, CERN, Geneva, Aug 2016.

[101] The ATLAS Collaboration. Search for squarks and gluinos in events with an isolated lepton, jets and missing transverse momentum at $\sqrt{s} = 13$ TeV with the ATLAS detector. Technical Report ATLAS-CONF-2016-054, CERN, Geneva, Aug 2016.

[102] The ATLAS Collaboration. Search for supersymmetry with two same-sign leptons or three leptons using 13.2 fb^{-1} of $\sqrt{s} = 13$ TeV pp collision data collected by the ATLAS detector. Technical Report ATLAS-CONF-2016-037, CERN, Geneva, Aug 2016.

[103] The CMS Collaboration. Search for supersymmetry in events with jets and missing transverse momentum in proton-proton collisions at 13 TeV. Technical Report CMS-PAS-SUS-16-014, CERN, Geneva, 2016.

[104] The CMS Collaboration. Search for new physics in the all-hadronic final state with the MT2 variable. Technical Report CMS-PAS-SUS-16-015, CERN, Geneva, 2016.

[105] The CMS Collaboration. An inclusive search for new phenomena in final states with one or more jets and missing transverse momentum at 13 TeVwith the α_T variable. Technical Report CMS-PAS-SUS-16-016, CERN, Geneva, 2016.

[106] The CMS Collaboration. Search for supersymmetry in events with one lepton and multiple jets in proton-proton collisions at $\sqrt{s} = 13$ TeV in 2016. Technical Report CMS-PAS-SUS-16-019, CERN, Geneva, 2016.

[107] The CMS Collaboration. Search for SUSY in same-sign dilepton events at 13 TeV. Technical Report CMS-PAS-SUS-16-020, CERN, Geneva, 2016.

[108] The CMS Collaboration. Search for SUSY with multileptons in 13 TeVdata. Technical Report CMS-PAS-SUS-16-022, CERN, Geneva, 2016.

[109] The CMS Collaboration. Search for supersymmetry in the all-hadronic final state using top quark tagging in pp collisions at $\sqrt{s} = 13$ TeV. Technical Report CMS-PAS-SUS-16-030, CERN, Geneva, 2016.

[110] Jonathan M. Butterworth, Adam R. Davison, Mathieu Rubin, and Gavin P. Salam. Jet substructure as a new Higgs search channel at the LHC. *Phys. Rev. Lett.*, 100:242001, 2008.

[111] Michael H Seymour. Tagging a heavy Higgs boson. Technical Report CAVENDISH-HEP-90-25, Nov 1990.

[112] Jonathan M. Butterworth, Ins Ochoa, and Tim Scanlon. Boosted Higgs $\to b\bar{b}$ in vector-boson associated production at 14 TeV. *Eur. Phys. J.*, C75(8):366, 2015.

[113] The ATLAS Collaboration. Search for the Supersymmetric Partner of the Top Quark in the Jets + $E_{\mathrm{T}}^{\mathrm{miss}}$ Final State at $\sqrt{s} = 13$ TeV. Technical Report ATLAS-CONF-2016-077, CERN, Geneva, Aug 2016.

[114] The ATLAS Collaboration. Search for top squarks in final states with one isolated lepton, jets, and missing transverse momentum in $\sqrt{s} = 13$ TeV pp collisions with the ATLAS detector. Technical Report ATLAS-CONF-2016-050, CERN, Geneva, Aug 2016.

[115] The ATLAS Collaboration. Monitoring and data quality assessment of the ATLAS liquid argon calorimeter. *JINST*, 9:P07024, 2014.

[116] Benjamin Nachman, Pascal Nef, Ariel Schwartzman, Maximilian Swiatlowski, and Chaowaroj Wanotayaroj. Jets from Jets: Re-clustering as a tool for large radius jet reconstruction and grooming at the LHC. *JHEP*, 02:075, 2015.

[117] The ATLAS Collaboration. ATLAS Run 1 Pythia8 tunes. Technical Report ATL-PHYS-PUB-2014-021, CERN, Geneva, Nov 2014.