Cross-section measurements of the electroweak production of a single top quark in association with a W boson at $\sqrt{s} = 13\,\text{TeV}$ with the ATLAS detector

Carl J. E. Suster

A thesis submitted in fulfillment of the requirements for the degree of

Doctor of Philosophy

2019

Faculty of Science

The University of Sydney

Cross-section measurements of the electroweak production of a single top quark in association with a W boson at $\sqrt{s} = 13\,\text{TeV}$ with the ATLAS detector

Carl J. E. Suster

Abstract

The top quark is the most massive fundamental particle within the Standard Model of particle physics, discovered at Fermilab in 1995. In *pp* collisions at the Large Hadron Collider at CERN, large quantities of top quarks are produced, offering the first opportunity to explore its properties in rarer production modes in detail. This thesis presents measurements of the cross-section of the electroweak production of a single top quark in association with a W boson, using data collected at 13 TeV centre-of-mass energy with the ATLAS detector. The total tW production cross-section is measured in events with two oppositely-charged electrons or muons. Differential tW cross-sections are measured with respect to several observables, providing the first detailed validation of the modelling of the process. The measurements are consistent with theoretical predictions, and contribute to future improvements in the modelling of tW production.

Statement of originality

This is to certify that to the best of my knowledge, the content of this thesis is my own work. This thesis has not been submitted for any degree or other purposes. I certify that the intellectual content of this thesis is the product of my own work and that all the assistance received in preparing this thesis and sources have been acknowledged.

Authorship attribution statement

The research presented in this thesis was undertaken within the context of the ATLAS Collaboration at CERN, an international effort of thousands of physicists working together to design and operate the ATLAS experiment and to analyse the data it collects. Chapter 1 outlines the theoretical framework. Chapters 2 and 3 describe the experiment and the preparation of data for analysis; while I contributed to the operation of the experiment—and the exposition and selection of material is my own—none of the work described in these chapters was performed by me and it is included to provide context to discussions of decisions in the design of the analysis. The experimental data were collected by the collaboration as a whole and my specific contributions included data quality shifts in the experiment's control room at CERN, offline reconstruction of the beam spot, and review of reconstruction software. The simulated datasets were prepared by others in the collaboration.

The main work of this thesis corresponds to two analyses of tW production performed by a small group of analysers working within the wider collaboration. The common analysis strategy is described in Chapter 4. The first analysis was published as Ref. [50] and is described in Chapter 5. The second analysis was published as Ref. [47] and is described in Chapter 6. As a member of the small team that produced this work, I contributed to the analysis design, analysis software architecture and implementation, data processing, validation studies, BDT optimisation, and the preparation of the publications. Each individual on the team has contributed to different aspects of the work over time; this is a very typical way of working in large collaborations since it is impractical for a single physicist to manage all of the technical details of the analysis and the experiment. Figures within these chapters that are marked as official ATLAS plots—and the results themselves—are taken from the relevant publications, and were originally prepared by me and my collaborators.

Chapter 7 describes the context of the results and related measurements performed by others. Throughout this thesis, all figures that are not explicitly attributed to others are my own work.

In addition to the statements above, in cases where I am not the corresponding author of a published item, permission to include the published material has been granted by the corresponding author.

Carl J. E. Suster
December 11, 2019

As supervisor for the candidature upon which this thesis is based, I can confirm that the authorship attribution statements above are correct.

Prof. Kevin Varvell
December 11, 2019

Contents

Introduction

In the pursuit of new insights into the physics governing the full diversity of macroscopic interactions between familiar species of matter, we have sought explanations in terms of a small number of fundamental interactions between elementary components of matter. The new field of nuclear physics emerged with the discoveries of radiation and the nucleus at the dawn of the twentieth century. Over the subsequent decades a theory of the interactions between nucleons was developed, explaining and predicting features of the chemical elements using only a handful of more elementary components. Through the study of high-energy phenomena with cosmic rays and particle accelerators, completely new fundamental particles were discovered that did not appear necessary to account for the interactions of usual matter, and the population of the known subatomic realm grew steadily. Patterns in the taxonomy of new particles suggested that some of them were in fact bound states of more fundamental particles, quarks. The Standard Model of particle physics was developed in the latter part of the twentieth century, organising this new complexity into an elegant quantum field theory describing a smaller number of (apparently) truly fundamental particles and the symmetries governing their interactions. Nuclear physics was re-conceived as a low-energy effective theory. We now believe that the Standard Model, despite its successes, is itself an effective theory requiring extension or replacement at higher energies to explain all observed phenomena in cosmology and fundamental physics[1].

The proton that had once appeared elementary revealed an increasing complexity as it was probed at higher energies, corresponding to smaller distance scales. By colliding protons at high energies, the partons within them interact to produce new particles; amongst the particles discovered in this way is the top quark, a heavier relative of the up quark. Whereas up quarks bind with down quarks to form nucleons, the top quark is so unstable that it decays on time scales faster than the interactions that would bind it to other quarks. Due in part to its relatively recent discovery, there is still much to learn about the nature of the top quark and its interactions. For this reason, and because of its peculiarities, the top quark presents an avenue for testing the boundaries of the Standard Model.

This thesis explores the production of top quarks in collisions between protons at the Large Hadron Collider at CERN (Conseil européen pour la recherche nucléaire, the European Organisation for Nuclear Research). Whereas the dominant production modes are governed by the strong interaction and result in a pair of top quarks, our aim here is to isolate a sample of collisions that resulted in the production of a single top quark by electroweak interactions. Specifically, the production of a single top quark in association with a W boson (the tW process) is the target of investigation. This is not because it is inherently more or less deserving of attention than other processes, but simply because the challenges entailed in the analysis of top quark production differ sufficiently between processes so

[1] This brief history implies a teleological clarity that only exists in hindsight.

as to warrant separate investigations; the differences in the physics of those processes demands that we study each of them. The measurements of the cross-section of the tW process presented in this thesis allow verification of its theoretical modelling at a new collision energy, and present approaches for handling similar measurements in the future.

Chapter 1 outlines the theoretical background with a particular emphasis on issues in the modelling of proton collisions related to top quark production. The Large Hadron Collider is introduced in Chapter 2 along with the ATLAS detector used to collect the data analysed here. The detector produces electronic responses to the passage of particles through its volume, but that volume is larger than the distances many particles of interest can travel before decaying. The process of reconstructing information about particles from the detector response is discussed in Chapter 3. The tools used to isolate a sample of single top quarks are introduced in Chapter 4, along with an outline of the general analysis approach and the treatment of systematic uncertainty in the measurements. The first measurement of the cross-section of tW production at a collision energy of $\sqrt{s} = 13\,\text{TeV}$ is presented in Chapter 5. The first measurements of the differential cross-sections of tW production as a function of several observable quantities is presented in Chapter 6. The current status and future of experimental investigation into tW production is briefly discussed in Chapter 7.

Chapter 1

The top quark and the Standard Model

The top quark is one of six flavours of quarks, fundamental particles of matter, described by the Standard Model (SM) of particle physics. While the SM does not predict the masses of these particles, the top quark has been observed to have a mass so large that it is comparable to the mass of a tungsten atom ($Z = 74$). On a theoretical level this large mass suggests that the top quark might play a role in the mechanism of electroweak symmetry breaking. Experimentally, there are a number of interesting consequences of the top-quark mass, the first of which being that the top quark was not discovered until as late as 1995 due to the high collision energies necessary to produce it at particle colliders. Whereas the lighter quarks are confined by the nature of the strong interaction to exist in bound states, the very short lifetime of the top quark as a consequence of its mass[1] means that it decays before such states can form. Because of this, information about the spin of the top quark is preserved relatively intact in angular correlations of its decay products, and its intrinsic width is sensitive to higher-order perturbative corrections [161]. This feature of its decay leads to the description of the top quark as the only "pseudo-bare" quark[2], permitting a richer set of measurements that are not possible with the other quarks.

At the Large Hadron Collider (LHC) top quarks are produced at a rate of several pairs per second, and samples of millions of collisions involving top quarks have been gathered for analysis. In addition to the information to be gained through study of the top quark itself, precise knowledge of its production cross-section is necessary in order to observe rare processes and potential physics beyond the Standard Model (BSM).

1.1 A brief overview of the Standard Model

The SM is a theory expressed in the language of quantum field theory (QFT), based on a particular choice of fields and gauge symmetries. A unique SM Lagrangian is determined by these choices and the restrictions that the theory be renormalisable, gauge-invariant, and local. The phenomenological consequences of the SM—the particles and their interactions—then arise from the underlying group

[1] Being substantially more massive than a W boson, its decay via the weak interaction is not suppressed, unlike other quarks, and it therefore has a short lifetime of about 0.5×10^{-24} s [205, §67.2].

[2] The top quark is not free from nontrivial structure due to strong interactions. Calculations in the context of future lepton colliders have considered predictions for the process $e^- e^+ \to b\bar{b}\, W^+ W^- X$ near the $t\bar{t}$ production threshold (i.e. around 345 GeV, twice the top quark mass). The top threshold dynamics are dominated by strong interactions in a technically difficult regime of QCD with contributions from EW and nonresonant effects [197, 92].

Table 1.1: Matter particle content of the Standard Model. The generation (Gen.), electric charge (Q) and mass of each particle is shown, except for the neutrinos which do not have well-defined masses when written in the flavour basis [205].

Particle		Gen.	Q	Mass [GeV/c^2]		
LEPTONS						
electron	e^-	1	-1	5.109 989 461	$\pm$0.000 000 031	$\times 10^{-4}$
muon	μ^-	2	-1	1.056 583 745	$\pm$0.000 000 024	$\times 10^{-1}$
tau	τ^-	3	-1	1.776 86	$\pm$0.000 12	
electron neutrino	ν_e	1	0			
muon neutrino	ν_μ	2	0			
tau neutrino	ν_τ	3	0			
QUARKS						
down quark	d	1	$-1/3$	4.67	${}^{+0.48}_{-0.17}$	$\times 10^{-3}$
up quark	u	1	$+2/3$	2.16	${}^{+0.49}_{-0.26}$	$\times 10^{-3}$
strange quark	s	2	$-1/3$	9.3	${}^{+1.1}_{-0.5}$	$\times 10^{-2}$
charm quark	c	2	$+2/3$	1.27	$\pm$0.02	
bottom quark	b	3	$-1/3$	4.18	${}^{+0.03}_{-0.02}$	
top quark	t	3	$+2/3$	1.729	$\pm$0.004	$\times 10^2$

structure associated with the symmetries in the theory. This particle content of the SM consists of three generations of fermions with spin one half, and of bosons with integer spin.

1.1.1 Matter

The fundamental particles of matter in the SM are the fermions, particles with spin one half, listed in Table 1.1. They occur in three generations, and particles of the same type in successive generations differ only by their masses (which increase with the generation) and flavours, with all other quantum numbers identical. The first generation includes the electron (e), which is stable, and the up and down quarks (u and d) which form bound states as protons and neutrons. These particles, together with the electron neutrino (ν_e), account for the vast majority of the stable matter in the universe. The second and third generation particles are primarily relevant to high energy phenomena including astrophysical sources and collider physics. The heavier quarks and charged leptons are unstable and in general decay rapidly to first-generation particles. The antimatter partners of the quarks and charged leptons differ in the sign of their electric charge.

The neutrinos are neutral leptons in three flavours corresponding to the three flavours of charged leptons. The physical neutrinos that propagate through space are mixtures of these flavoured neutrinos, with the rotation between the propagation (or mass) and flavour (or weak) bases parametrised by the unitary Pontecorvo-Maki-Nakagawa-Sakata [177] matrix. Although the neutrinos described by the SM are massless particles, the observation of neutrino oscillations implies that their mass is in fact nonzero. Each flavour of neutrino has a corresponding antineutrino in the standard description, although it is at present unresolved as to whether the neutrinos are Majorana particles, in which case each would be its own antiparticle partner [176].

The sum of the masses of the propagation eigenstates of the neutrinos—ν_1, ν_2, and ν_3—is

Interaction	Coupling constant		Mediator
EM	$\alpha(Q^2 = m_W^2)$	$\approx 7.8 \times 10^{-3}$	γ
Weak	G_F [GeV^{-2}]	$\approx 1.16 \times 10^{-5}$	$W^{\pm}$, Z
Strong	$\alpha_s(m_Z)$	≈ 0.12	g
Gravity	G_N [GeV^{-2}]	$\approx 6.71 \times 10^{-39}$	?

Table 1.2: The four fundamental forces, or interactions, and their associated coupling constants (at the electroweak scale) and mediators [182, 205].

constrained by experiment to be less than approximately 0.15 eV[3] [205]. The masses of the matter particles range from the neutrinos through to the 173 GeV of the top quark, with the quark masses alone varying over almost five orders of magnitude.

1.1.2 Interactions

Three fundamental interactions between particles are described by the SM: the electromagnetic (EM), weak and strong interactions. Each is associated with vector bosons that mediate the interaction. The relative strength of each interaction can be crudely assessed by comparing the relevant coupling constants at the same energy scale[4], as is done in Table 1.2 where gravity is also included. The fact that gravity is orders of magnitude weaker than the interactions in the SM contributes to the fact that it is irrelevant to the description of particle interactions in the context of high energy physics until considering energy scales substantially higher than the limits of experimental accessibility.

The EM interaction is mediated by the photon, γ, and is experienced by all particles with a non-zero electric charge. Because the γ is massless, the interaction effectively has an infinite range, though its strength falls as the inverse square of distance. The weak interaction is mediated by three bosons: the W^+, W^- and Z. Because these are massive particles, the weak interaction is restricted to short ranges. It is the only interaction through which quarks of different flavours couple. At energies high with respect to the masses of the weak mediators, electromagnetism and the weak interaction can be considered as a single electroweak (EW) interaction.

The strong interaction is mediated by the 8 gluons, g, and its charge is called colour charge to suggest an analogy with visual colour in the way that the charges combine[5]. The quantum field theory describing the strong interaction is quantum chromodynamics (QCD), a non-Abelian gauge theory. Coloured particles (the gluons and the quarks) have the property of asymptotic freedom, meaning that they can be treated as free particles in the asymptotic limit of high momentum transfer. In the low energy regime, however, such particles are always found confined in colourless bound states called *hadrons*. Quark-antiquark pairs with opposite colour charge form *mesons*, while sets of three differently-coloured quarks form colour singlets as *baryons*.

The interactions and field content of the SM are encoded into its Lagrangian density. In addition to respecting symmetries associated with the rotation and translation of spacetime and boosting

[3]Throughout this thesis natural units are used where $c = \hbar = 1$. In this system of units, mass, energy, and momentum are all expressed in units of eV.

[4]The fine structure constant α and the strong coupling constant α_s are functions of the energy scale of the interaction. Compared to EM, at scales relevant to quarks near 10^{-18} m the strong interaction is about 25 times stronger, the weak interaction is of the same order as EM, and gravity is 41 orders of magnitude smaller.

[5]For example a mixture of red, green, and blue light gives white light, and a mixture of red, green, and blue charges results in a colourless object. Red and green light mix to give yellow light, which is in a sense the opposite of blue; and a mixture of red and green charges functions as an anti-blue charge.

Table 1.3: Quantum numbers of the SM fermion and Higgs fields under the three gauge transformations. In each case i runs over the three generations. Q_L^i are the left-handed quarks in doublets of up- and down-type quarks. u_R^i and d_R^i are the right-handed quark singlets for up- and down-type quarks respectively. L_L^i are the left-handed leptons in doublets of the neutrino and charged lepton. e_R^i and ν_R^i are the right-handed leptons in singlets. ϕ is the Higgs field.

Field	$SU(3)_C$	$SU(2)_L$	$U(1)_Y$
Q_L^i	3	2	1/6
u_R^i	3	1	2/3
d_R^i	3	1	−1/3
L_L^i	1	2	−1/2
e_R^i	1	1	−1
ν_R^i	0	0	0
ϕ	1	2	1/2

between reference frames, the Lagrangian obeys a set of internal symmetries. These are related to the gauge group of the SM: $SU(3)_C \times SU(2)_L \times U(1)_Y$. The $SU(3)_C$ subgroup, where the C indicates colour, is associated with the strong interaction, and its 8 generators are the gluons. The $SU(2)_L \times U(1)_Y$ group is associated with the unified EW interaction. The L indicates that under $SU(2)_L$ transformations the left-handed chiral fields are doublets whereas the right-handed fields are singlets[6]. Y is the weak hypercharge quantum number, related to the electric charge and weak isospin[7]. The fields of the SM and their quantum numbers under each of these groups is shown in Table 1.3.

The structure of these gauge groups results in a ground state of the theory in which all gauge bosons are massless, inconsistent with the observed masses of the Z and W bosons. In order to give mass to these bosons, the EW symmetry group $SU(2)_L \times U(1)_Y$ is broken, and only the $U(1)_{\text{EM}}$ symmetry associated with the EM interaction remains. The Brout-Englert-Higgs mechanism [152, 135] accomplishes this by introducing a new complex scalar field ϕ, a doublet under $SU(2)_L$. As the universe cooled, the Higgs field adopted a nonzero expectation value from a sphere of possible minima in its the potential, breaking the EW symmetry. Three of the four degrees of freedom in ϕ are mixed with the generators of the EW symmetry group to produce the three massive weak bosons and the massless photon. The remaining degree of freedom corresponds to the physical Higgs boson. The Yukawa coupling of the Higgs to the left- and right-handed fermions mix, and as a result of the symmetry breaking this gives rise to fermion mass terms in the Lagrangian. The top quark has the largest Yukawa coupling of the fermions.

The top quark is charged under the strong interaction, meaning that it couples to gluons. However, since the emission of a gluon is never able to change the flavour of a fermion, only couplings of the form $gt\bar{t}$ are permitted. The top quark also has an electric charge and a weak hypercharge, so couples with all of the EW gauge bosons. Neutral current interactions—those mediated by the photon and Z boson—describe its EM interactions, and are also unable to change the flavour of the top quark. Charged current interactions mediated by the W bosons, however, are able to violate the conservation of fermion flavour and couple the top quark to other quark flavours. The couplings between pairs of quark flavours are denoted V_{ij} where i and j run over the flavours. These terms are the elements of the Cabibbo-Kobayashi-Maskawa (CKM) [170] matrix, a 3×3 unitary matrix describing the rotation between the basis in which quarks propagate and the basis in which they participate in the weak interaction. Empirically this matrix is found to be diagonal to first order, such that the coupling

[6] The fermion fields in the SM Lagrangian are decomposed into their chiral components since the gauge symmetries act differently on the left and right chiral fields. For example, an $SU(2)_L$ transformation can rotate a left-handed electron into a left-handed electron neutrino, but is not able to transform a right-handed electron field nontrivially.

[7] $Q = Y + T_3$, where T_3 is the third component of weak isospin.

of quarks between generations is highly suppressed. There is a further hierarchy between the size of the off-diagonal terms; for the couplings to the top quark we have $|V_{tb}| \gg |V_{ts}| > |V_{td}|$ where $|V_{tb}| \approx 0.9991$.

1.1.3 Status of the Standard Model

The SM has proved to be enormously successful. The spectrum of additional particles and their properties predicted by the SM were subsequently found and confirmed, culminating in the observation of the Higgs boson in 2012 [57, 112]. Despite its successes, the SM is an incomplete fundamental description of nature: there is no clear way to incorporate a consistent quantum description of gravity, and there are no viable dark matter candidates among the SM's particle content. The SM also offers no explanation for cosmological phenomena including the accelerating expansion of the universe or the apparent scarcity of antimatter with respect to matter. While neutrino masses can be accommodated by the SM, it does not provide a fundamental explanation for their origin, unlike the other fermions.

There are also some features of the SM that are aesthetically unsatisfying, including its relatively large number of free parameters and the fact that the strong, weak, and EM forces are not unified at high energies. There is also the issue of the "unnatural" fine-tuning of quadratic corrections to the Higgs mass, in which the top quark is implicated since it is the only quark that couples to the Higgs boson with a strength of order unity. Such features are not necessarily failures of the SM, which is viewed as a low-energy effective theory, however attempts to resolve them often guide investigation into BSM theories. In particular, attempts to address unnaturalness issues associated with the Higgs field generally introduce new physics related to the top quark.

Experimentally, study of the SM generally falls into two areas. The first is BSM searches, in which particular models or classes thereof are tested by searching for evidence of new particles or interactions that they predict. These searches are undertaken at both the high energy frontier where new energy scales are explored, and at the luminosity frontier where large samples facilitate the observation of rare interactions. Such searches at colliders, along with searches in cosmology and other fields of fundamental physics, have succeeded in constraining the parameter space of proposed models and altering the direction of theoretical investigation, but have yet to yield evidence in support of any particular BSM model. The second area of study is precision measurements in which predictions made by the SM are carefully tested to search for discrepancies. Several unexplained results in fundamental physics have been observed including the anomalous magnetic dipole moment of the muon [183] and excesses in the cross-sections of certain decay modes of bound states containing a bottom quark [84, 174, 89]. The goals of precision measurements in collider physics are to validate detailed predictions and calculations made within the SM framework, both to improve the modelling and to identify any phenomena that are not well-described by current models.

1.2 Modelling of proton collisions

Hadrons are bound states of quarks under the strong interaction. While a proton can be said to be composed of two up quarks and one down quark, this is at best an approximation because the quarks constantly radiate and re-absorb gluons that themselves split into further gluons and quarks. The residual force of these interactions manifests as the nuclear force in the low-energy regime, but at high energies[8] a more fundamental description of hadron interaction phenomenology must be understood and modelled in order to make meaningful predictions and observations.

High-energy collisions between protons can be characterised as elastic collisions that result in two outgoing protons, or inelastic collisions in which constituents of one or both protons interact and new particles are produced. The probability of a given interaction is related to its *cross-section*, the equivalent target area presented by the proton under the interaction, measured in barns ($1\,\mathrm{b} = 10^{-28}\,\mathrm{m}^2$). The total proton-proton cross-section $\sigma_{\mathrm{tot.}}$ is then the sum of the elastic and inelastic cross-sections $\sigma_{\mathrm{el.}}$ and $\sigma_{\mathrm{inel.}}$ [8]. In the 1970s measurements at the Intersecting Storage Rings at CERN established that $\sigma_{\mathrm{tot.}}$ increases as a function of the energy available to the proton-proton interaction [9], which is best described by the collision centre-of-mass energy $\sqrt{s}$. Experiments investigating the diffractive scattering of protons are able to measure $\sigma_{\mathrm{tot.}}$ and the elastic fraction. The TOTEM experiment at CERN, using proton-proton interactions with $\sqrt{s} = 13\,\mathrm{TeV}$, recently measured $\sigma_{\mathrm{tot.}} = (110.6 \pm 3.4)\,\mathrm{mb}$ and $\sigma_{\mathrm{inel.}} = (79.5 \pm 1.8)\,\mathrm{mb}$, putting the elastic fraction at $\sigma_{\mathrm{el.}}/\sigma_{\mathrm{tot.}} = 0.281 \pm 0.009$ [206].

The inelastic interactions between protons, comprising the remaining 72 % of the total hadronic cross-section, are the subject of the majority of investigations at hadron colliders. The processes of interest are the hard interactions between constituents of the protons bearing significant fractions of the overall momentum, and these are well-described by perturbative models in the underlying field theory. The overall inelastic cross-section is then divided amongst the multitude of such hard processes allowed by the theory.

The hard process cannot be observed in isolation from a range of soft processes involving its hadronic environment. The full complexity of hadronic interactions is illustrated in Fig. 1.1. Borrowing from the conventions of Feynman diagrams, fermion propagators are represented by straight lines, and gluons by curly lines, with the intersection of lines indicating an interaction vertex between three particles. The hard process in Fig. 1.1 is the scattering of a fermion originating from one colliding proton off a gluon originating from the other. An extensive network of interactions connects the hard process to the initial proton constituents and remnants including the spectator quarks, and the outgoing particles of the hard process evolve into a complex arrangement of hadrons. Only subsequent to these interactions do the final products decay to particles that are sufficiently stable to propagate to a detector. The soft processes in the hadronic interaction are for the most part non-perturbative and so are intrinsically more difficult to model.

The QCD description of the hard process leads to a singularity occurring when a gluon is emitted parallel to a quark, called the collinear divergence. The theoretical insight that allows this to be managed is that the interaction can be *factorised* into short-ranged and long-range effects by setting a factorisation scale, μ_{F}^2. The collinear divergence is absorbed into the bare quark below this scale, whereas above this scale the quark is convoluted with the full structure of the proton to form the long-range observable object. By the factorisation scheme the long-range character of quarks can be described in the perturbation theory, while the complexity of the soft QCD processes related to the environment of the proton is moved into the empirical description of the proton structure. At sufficiently high orders in the perturbative expansion, results are not strongly dependent on the choice of factorisation scale.

[8] Anything above a few GeV is considered high energy in this context, and the energies relevant to this thesis are orders of magnitude higher still, in the TeV range.

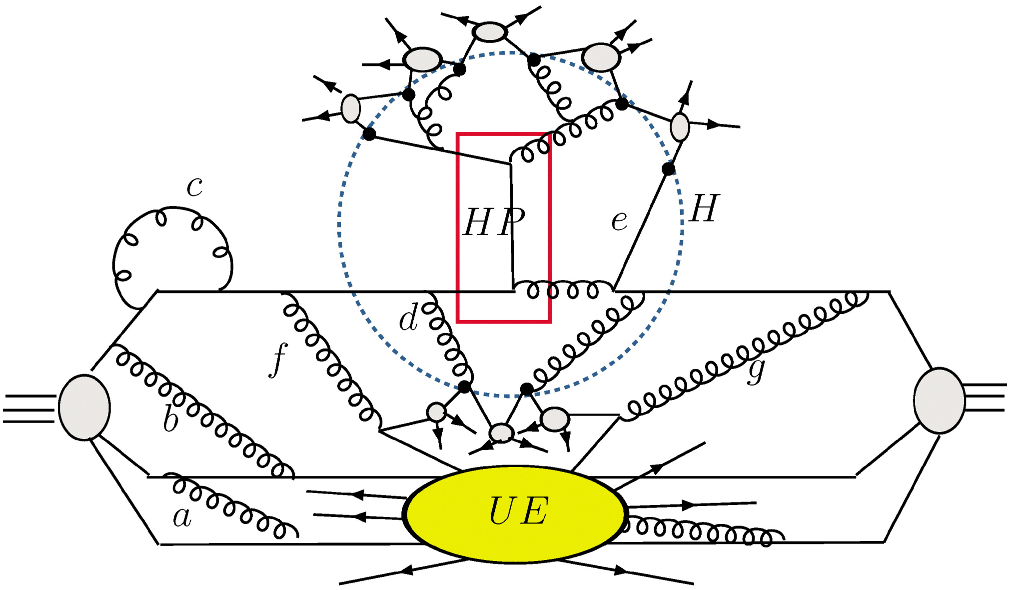

Figure 1.1: Representation of the full complexity of a collision between two protons. The protons are represented by the grey ellipses on the far left and right of the diagram with three lines exiting them representing the constituent valence quarks. These valence quarks constantly exchange soft gluons (a and b) and radiate and re-absorb harder virtual gluons (c). Processes inside the dashed circle H can be described perturbatively, including the hard process (HP) and radiation from the initial states (d and e). The final states of the hard process undergo hadronisation, represented by the smaller grey ellipses. Radiation from early in the perturbative evolution (f and g) can also hadronise and interact with the spectator quarks from the proton in what is called the underlying event (UE). Taken from Ref. [178]

Due to infinities that emerge in such high-order calculations, a renormalisation scale μ_{R}^2 must also be chosen. Commonly the two scales are chosen to have the same value $Q^2 = \mu^2$. These calculations are then supplemented with calculations of the soft processes in the hadronic interaction which are based on empirical parametrisations. The overall calculation is thus broken into a sequence of approximately independent stages and becomes tractable. In fact it is not completely reasonable to treat the stages as independent, for example in Fig. 1.1 there is some colour connection between the initial state and the final state via the quark labelled e. Despite these imperfections the approach holds up empirically depending on the level of detail required in the kinematic description [133, 191].

1.2.1 The parton model

The structure of hadrons is described by sets of experimentally-constrained probability distributions called parton distribution functions (PDFs), which encode the probability of finding a *parton* with a fraction x of the hadron's total momentum when probed at a given momentum transfer Q^2. For the protons, the statement that we expect to find on average two up quarks and one down quark corresponds to the summation rules for its up- and down-quark PDFs $u(x, Q^2)$ and $d(x, Q^2)$ requiring that

$$\int_0^1 dx \left[u(x, Q^2) - \bar{u}(x, Q^2)\right] = 2 \qquad \text{and} \qquad \int_0^1 dx \left[d(x, Q^2) - \bar{d}(x, Q^2)\right] = 1.$$

In addition to these two valence quarks, other *sea quarks* are part of the proton's structure. Empirically when the proton PDFs for each of the quarks, $q_i(x, Q^2)$ (where the index i runs over the six quark flavours), are measured from data we find that

$$\sum_i \int_0^1 dx \left[q_i(x, Q^2) + \bar{q}_i(x, Q^2)\right] \approx 0.5,$$

i.e. the quarks carry only half of the proton's total momentum. The other half of the momentum is mostly accounted for by gluons[9].

Figure 1.2b shows one particular set of proton PDFs calculated by fitting to a range of experimental data. The valence quark PDFs peak at approximately 10^{-1}, signifying that when probed at $Q^2 = 10^4\,\mathrm{GeV}^2$ the expectation of the momentum fraction carried by each individual valence quark is 10^{-1}. The sea parton PDFs fall steeply with the momentum fraction, corresponding to a lower probability of finding them individually with a significant fraction of the proton's momentum.

Quarks with masses well above the energy scale under consideration are not intrinsic constituents of the proton since they are subject to the boundary condition that their PDFs must vanish below their mass. For the quarks that *are* intrinsic to the proton at a given scale, their PDFs are non-perturbative and must be determined experimentally by fitting to data. The heavier quarks then arise perturbatively through gluons and the light quarks. PDF sets are produced by fitting to data at a relatively low scale (as shown in Fig. 1.2a) and then using Dokshitzer-Gribov-Lipatov-Altarelli-Parisi QCD equations—along with further corrections—to evolve them up to the scale of interest (as shown in Fig. 1.2b) [133, Ch. 4]. Intrinsic charm content was not typically included in the PDF sets produced by most of the major collaborations until recently [1] and intrinsic top content is not relevant at

[9] Ew interactions further complicate this picture to the extent that some fraction of the momentum of the proton can be carried by just about any fundamental particle with some non-zero probability, depending on the energy with which its structure is probed.

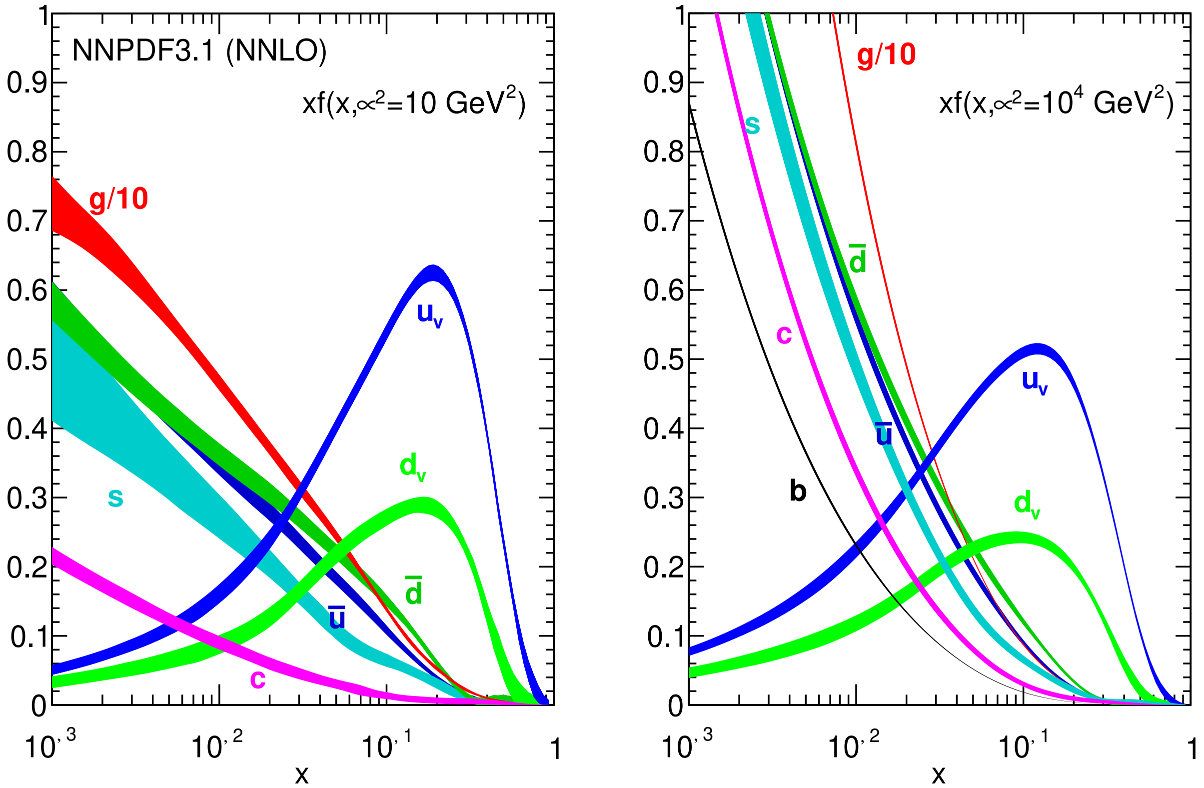

(a) PDF fit scale, $Q^2 = 10\,\mathrm{GeV}^2$

(b) LHC scale, $Q^2 = 10^4\,\mathrm{GeV}^2$

Figure 1.2: NNPDF3.1 NNLO PDFs at different scales $\mu^2 = Q^2$. The horizontal axis corresponds to the fraction x of the total proton momentum, and the vertical axis shows xf, the product of x and the PDF f for a given parton. The width of the bands corresponds to the uncertainty on the PDF. Taken from Ref. [85].

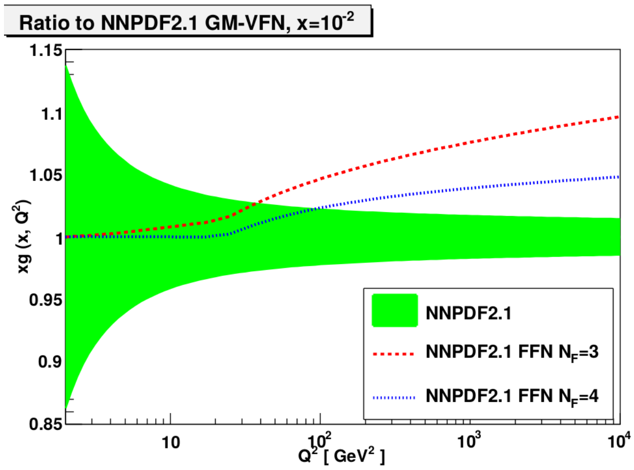

Figure 1.3: The ratio of the gluon PDFs in the three- and four-flavour restricted scheme to the default 5FS in NNPDF2.1. The schemes agree at scales below the matching scale, chosen to be near the mass of the heaviest included quark; for example the 4FS line approaches unity as the energy scale decreases past the b mass $m_b^2 \approx 16\,\text{GeV}^2$ indicating that it agrees perfectly with the nominal scheme in that region. Taken from Ref. [87].

currently accessible values of Q^2.

Figure 1.2 also shows that as Q^2 increases a greater fraction of the proton's momentum is carried by the sea partons and hence by increasing the energies at which protons are collided, finer structure of the proton becomes more relevant. In particular the gluon PDF grows rapidly at smaller momentum fractions so that gluon processes tend to dominate in higher-energy hadron collisions.

In the simple parton model, the partons are taken to be collinear with the proton's direction, carrying a fraction of its overall momentum, and so do not possess any component of momentum orthogonal to the proton's momentum. The non-perturbative effects do in fact lead to an intrinsic transverse momentum that manifests in experimentally observable discrepancies in the tails of distributions of the transverse momenta of outgoing particles in the hard scatter. Correcting for these effects introduces new divergences in the theory in the form of logarithmic enhancements to each order of α_s [118]. In the case of a single such term at each order, this can be handled by resummation at all orders, but in the common case where there are two powers of logarithmic enhancements at each order accuracy is more difficult to achieve. In these cases resummation accounts for the leading logarithm—the term with a power of twice the power of α_s—and further refinements to the calculation can also account for the subleading power to give cross-section predictions with next-to-leading logarithm corrections, and so on for subsequent subleading powers [205, §41.1.4].

1.2.2 Restricted flavour schemes

Instead of allowing the number of quark flavours to increase with the energy scale as described in the previous section, many PDF sets freeze the number of flavours at some matching scale μ^2_{match} and then do not introduce further flavours even as the scale increases beyond the mass of the remaining quarks. The matching scale is typically chosen by considering the masses of the relevant quarks. The five-flavour scheme (5FS) is commonly used by PDF collaborations, meaning that all but the top quark contribution to the proton are included normally. Since the top mass is so large this means that the matching scale is well above energies accessible to the LHC and so this restricted scheme is identical to the full scheme at experimentally relevant energies. PDF sets using the four-flavour scheme (4FS) are also routinely produced, and these agree up until scales above the b mass as shown in Fig. 1.3.

All restricted flavour schemes are consistent in the full perturbative expansion, but not necessarily at fixed finite orders. For this reason it is necessary to model the evolution of the strong coupling

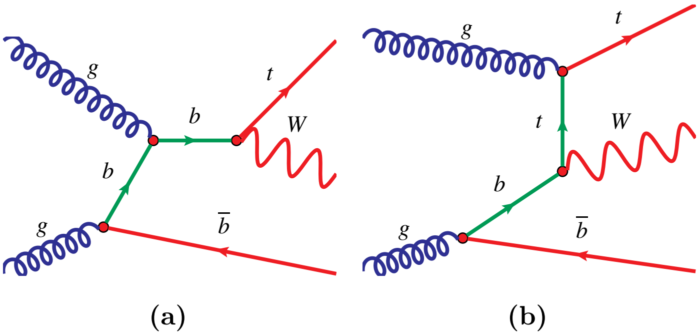

Figure 1.4: tW production at leading order in the 4FS with the b-quark arising from gluon splitting. These diagrams still occur in the 5FS but at a higher order (see §1.5.3 for the leading-order 5FS diagrams).

constant α_s and perform cross-section calculations using a scheme consistent with the number of flavours included in the PDF. The use of the 4FS requires that Feynman diagrams[10] with a b-quark in the initial state be excluded since the proton does not intrinsically contain these quarks. This means that diagrams where the b-quark arises from gluon splitting (as in Fig. 1.4) are the lowest order diagrams contributing to the tW process, for example. An advantage of the 4FS is that detailed modelling of the kinematics of the jet arising from the spectator b-quark is possible unlike in the 5FS where some information is lost when the relevant terms are resummed.

1.3 Monte Carlo event generators

QFT calculations today must be performed at increasingly higher perturbation orders to match the experimental precision that has been achieved. As the number of integrals involved grows exponentially with each order, these calculations in general require the assistance of computer programs using Monte Carlo (MC) methods. MC generators are modular tools that each perform part of the calculations necessary to model the full hadronic interaction. The interaction is intrinsically random in the sense that the specific partons that participate in the interaction and their momenta are distributed according to the PDFs even if the collision is brought about by colliding pure monochromatic beams of protons with known momenta. The MC generators sample these possibilities according to theoretical and empirical descriptions of the underlying physics to produce many random *events* that each represent one possible outcome and history of the interaction. From these events physical observables can be constructed.

Experimentally such MC simulations are indispensable, not only in order to compare experimental observables with these simulated predictions, but also in the execution of the experiment itself. Every part of the experiment from the calibration of detectors to the modelling of backgrounds and the interpretations of results relies on comparisons of observed events to MC simulated events. The inter-dependence of MC simulation and experiment leads to an iterative process by which new experiments constrain modelling and identify inconsistencies in MC generators, and new predictions allow greater experimental precision.

1.3.1 Matrix element generation and parton showering

Matrix element (ME) generators model the hard-scatter interaction, and rely on PDF sets to describe

[10] Feynman diagrams are tools from perturbative calculations where each diagram represents a term in the relevant integrals at a certain order. Fermion propagators are drawn as straight lines, gluons as curly lines, and other bosons as wavy lines. Incoming particles are at the left edge of the diagram, and outgoing particles at the right edge. All internal particles are virtual, and their interactions are represented by vertices where the propagators meet.

the structure of the interacting particles. Parton shower (PS) generators then model soft and collinear QCD radiation resulting in additional outgoing particles, implementing the equivalent of the resummation of large logarithms[11]. These two stages can be simulated by the same MC generator software, but are often handled by separate generators interfaced via a standardised format[12] describing the relevant parameters and the event record containing the generated particles and their ancestry. To meet improving experimental precision, MC generators have implemented increasingly sophisticated and accurate methods. As higher-order corrections to calculations were incorporated, the problem emerged of over-counting contributions to additional radiation in the process of interfacing next-to-leading order (NLO) ME calculations to the approximate NLO corrections in the PS computations. Merging schemes developed to resolve this problem in general depend on the initial state particles and on several trade-offs in their approaches.

The first generator to successfully resolve the over-counting issue was MC@NLO [144]. Its approach is to compute the approximate NLO corrections for a given PS generator and subtract them at the level of the ME computation. One consequence of this procedure is that the weights assigned to particular events by MC@NLO can be negative. While cross-sections calculated with these negative weights end up positive, this situation presents some technical difficulties. The Positive Weight Hardest Emission Generator (POWHEG) method was developed as an alternative approach that accounts for the over-counting while guaranteeing positive event weights [143]. Modelling of top quark production at hadron colliders typically uses MC generators based on the MC@NLO and POWHEG techniques. For background processes where additional hard QCD radiation is particularly important, event generators based on colour-aware multi-leg schemes such as COMIX within the HERWIG framework are preferred [149].

Parton showering models QCD emissions in the soft and collinear limits. The main PS generators used in top quark modelling are HERWIG [119] and PYTHIA [198], with the differences in approaches largely originating from the ordering in which emissions are generated. HERWIG begins with the large angle emissions, and by ensuring angular ordering is able to preserve coherence and partially account for interference effects in the underlying process. This is particularly useful when modelling the decay of unstable resonances, or when information about angular correlations and spin is required. PYTHIA instead begins with high-virtuality emissions, proceeding in order of hardness. This ordering is convenient for merging with ME computations, but does not maintain coherence. There are additional trade-offs between the approaches that motivate preferences for one or the other depending on the kinematics and regions of phase space of interest.

The matching scheme must be designed to ensure a smooth transition between the description of hard, well-separated partons by the ME scheme, and the description of soft and collimated showers of partons by the PS scheme. The task of matching these two regimes becomes increasingly challenging when considering final states with multiple hard jets, and detectors with large acceptances extending into the forward regions. The matching must fix a scale for the transition between regimes, and ensure that the predictions are not overly sensitive to the precise value of this scale. Several different matching schemes are available, with subtle differences in their treatment of parts of the phase

[11] Resummation of logarithmic enhancements in QCD calculations are specific to an observable, but in PS generators the momenta of the final state particle are predicted from which observables can be constructed. The corrections are therefore of a different nature and it is not straightforward to assess the shower accuracy in the same terms as resummation in calculations.

[12] The dominant MC generator interface is the Les Houches Interface for User Processes [94].

space reflected in small disagreements in the distributions of kinematic observables [157]. While the separation of event generation into ME and PS has served the modelling needs of the community well, alternative schemes that unify the two stages are under active investigation [88].

1.3.2 Additional modelling

Interactions between partons other than those directly involved in the hard-scatter interaction form the *underlying event*. Radiation of hard partons from the initial or final state of the hard-scatter interaction (see gluons f and g in Fig. 1.1) contribute to the underlying event. Additional soft components come from the beam remnant—the spectator partons from the colliding protons that are colour-connected to the hard partons—and semi-hard interactions where further partons interact as a result of the collision. These effects are difficult to model, and are managed by parametrising the effects and comparing predictions to experimental data in parts of the phase space that are particularly sensitive to this modelling.

In order to predict observables accurately, the response of the detectors must be considered. This can be done by parametrising the response and applying it by shifting and smearing observables to account for biases and finite resolution in experiment. Where a more accurate treatment is necessary, the particles in the generated event are propagated through space, including a detailed geometrical model of a specific detector and its materials. The detector response can be simulated down to the digitisation by readout electronics, and the events reconstructed according to the experiment's reconstruction software (see §3.1).

In experimental conditions further complications arise which must be incorporated into accurate MC models. One notable instance is the temporal or spatial superposition of multiple collisions between protons, collectively called *pileup*. Pileup is treated in MC event generation by randomly overlaying simulated events on top of the event of interest to account for the associated impact on physical observables.

1.4 Hadronisation and clustering

Partons generated in the hard scatter do not propagate as bare partons due to confinement (see §1.1.2), instead forming colourless hadrons. This *hadronisation* is a soft, non-perturbative process described by empirical models, and the transition from the parton shower to hadronisation is typically chosen to be of order 1 GeV. Typically the process is modelled using "strings" of colour between the partons that are broken up as hadrons form; the colour flow affects the angular development of the resulting hadron shower. Hadrons formed in the shower are in general highly collimated, as in Fig. 1.5, and are referred to as *jets*. Experimentally jets are observed by their energy deposits in calorimeters and by tracks from any charged hadrons in the shower, and they are typically treated as pseudo-particles. Features of jets ranging from kinematics to substructure in the shower depend on the species of parton that initiated the shower, with jets initiated by hard gluons qualitatively different to jets from light quarks or heavy quarks. In the study of top quarks, the signal process contains a large fraction of jets initiated by b-quarks as distinct from background processes, and so it is of particular interest to distinguish such jets from other flavours.

Jets are defined by a choice of clustering algorithm, with a range of such algorithms available. Earlier cone-based schemes simply defined jets as any radiation within a fixed cone around the hardest

Figure 1.5: Illustration of a jet showing a collimated shower of hadrons with associated radiation. Original author unknown.

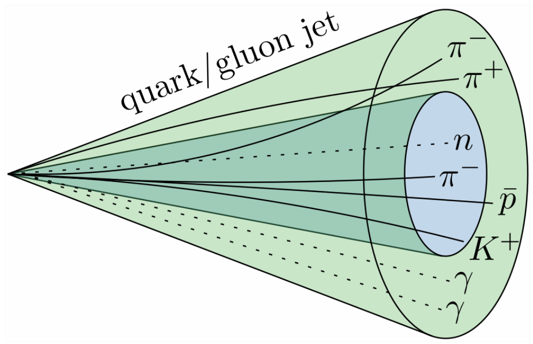

particle, however these suffered from divergences at higher orders of the perturbative expansion due to collinear emissions. Modern clustering algorithms are robust against such technical issues, and are distinguished by features such as their ability to assign hadrons to jets in such a way that the jets correspond as closely as possible to the underlying partons. The prevailing algorithm in use at hadron colliders today is the anti-k_t algorithm [101], a member of the class of sequential recombination algorithms. This class defines the distance $d_{i,j}$ between two entities, and the distance $d_{i,B}$ between an entity and the beam as

$$d_{i,j} \equiv \min\left(k_{t;i}^{2p}, k_{t;j}^{2p}\right) \frac{\Delta_{i,j}^2}{R^2} \qquad \text{and} \qquad d_{i,B} \equiv k_{t;i}^{2p}, \tag{1.1}$$

where $\Delta_{i,j}^2 \equiv (y_i - y_j)^2 + (\phi_i - \phi_j)^2$, $k_{t;i}$, y_i and ϕ_i are the transverse momentum, rapidity, and azimuthal angle of entity i (see §2.2.1 for the coordinate system and definition of rapidity). R is a parameter of the algorithm controlling the clustering radius, and p controls the relationship between the energy and geometric scales. Starting with the list of input entities, all possible $d_{i,j}$ and $d_{i,B}$ terms are computed. If the smallest distance is a distance between two entities, they are combined into a new pseudo-jet entity which is inserted into the list of entities in their place. If the smallest distance is between an entity and the beam, then that entity is removed from the list of entities and is called a jet. The algorithm proceeds by this iterative recombination of entities into jets until all have been clustered. The anti-k_t algorithm refers to the specific choice of $p = -1$ in Eq. (1.1), with $p = 0$ defining the Cambridge-Aachen algorithm [213, 130], and $p = 1$ the k_t algorithm [134].

According to the distance measures used by anti-k_t, hard objects separated by more than $2R$ will be clustered as distinct jets, absorbing any soft particles in their radius. When hard jets overlap, they will be clipped such that the harder object results in a more conical jet shape. When compared to alternative clustering algorithms such as the k_t algorithm, anti-k_t jets are notably more conical. This can be seen in Fig. 1.6 where the same partons are clustered using these two algorithms. The shape of anti-k_t jets bears greater resemblance to early cone-based jet algorithms than other members of the family of sequential recombination algorithms, combining the intuitive shapes of the former with the freedom from singularities that characterises the latter family.

1.5 Top quark production at hadron colliders

After the discovery of the b-quark in 1977 [151] the argument that it should have a weak isospin partner, analogous to the first two quark generations, motivated the search for a new third-generation

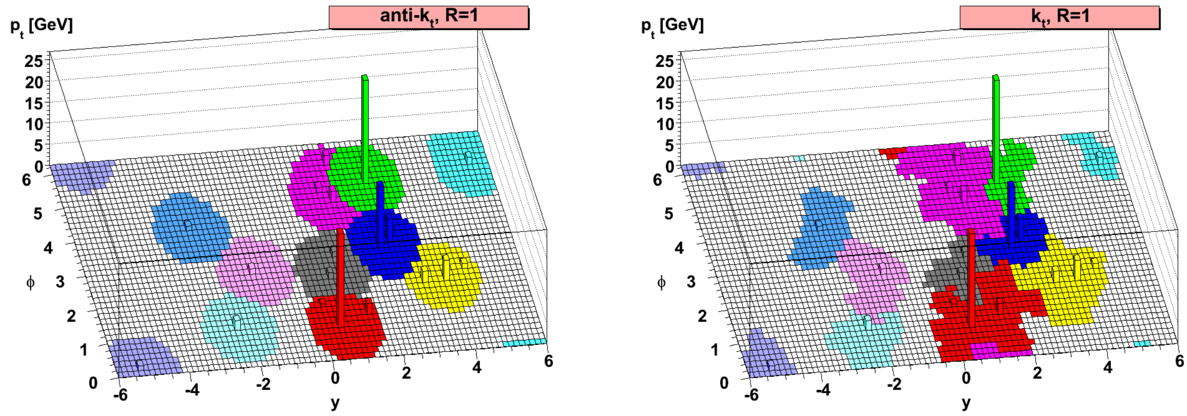

(a) anti-k_t **(b)** k_t

Figure 1.6: Comparison of jet clustering algorithms where in each case the clustering is performed using the same generated partons. The coloured regions identify the areas assigned to each jet, as determined by clustering many artificial particles with negligible momenta. The p_{T} of the particles initiating the jets is shown by the height of the bars. The anti-k_t algorithm produces conical jets, with irregular shapes only where hard particles are found within the radius parameter. In these cases the cone of the harder particle clips the softer cone as seen in (a) where the green jet clips the much softer magenta jet, unlike in (b) where the same magenta jet clips much of the area of the green jet. Taken from Ref. [101].

quark[13]. The b-quark had been discovered by locating a $b\bar{b}$ resonance in $m_{\mu\mu}$, the dimuon mass spectrum, so initial searches for the top quark at lepton colliders in the 1980s sought the reaction $e^- e^+ \to t\bar{t}$. Failure to find evidence of this reaction at lepton colliders hosted by DESY (Deutsches Elektronen-Synchrotron, the German Electron Synchrotron) in Germany and KEK (High Energy Accelerator Research Organisation) in Japan[14] placed a lower limit of around 30 GeV on the mass. The observation of B_d^0-$\bar{B}_d^0$ mixing at DESY in 1987 independently pushed up the lower mass limit [11].

The Super Proton Synchrotron (SPS) at CERN was converted[15] from a synchrotron into the first proton-antiproton collider, S$p\bar{p}$S, hosting the UA1 and UA2 experiments which discovered the W [207, 210] and Z [208, 209] bosons by 1983. The TEVATRON [140, 139] $p\bar{p}$ collider at Fermilab was completed soon after this and began its first run in 1987. The experiments searched for the top quark by collecting samples of singly-produced W bosons and looking for the weak decay $W \to tb$. Failure to observe this process favoured an even higher top mass, in which case the dominant production mechanism would be strong interactions resulting in a top-quark pair which subsequently decays weakly via the process $t \to Wb$. With $p\bar{p}$ collision energies of around 600 GeV at S$p\bar{p}$S and 1.8 TeV at the TEVATRON, direct searches continued to push up the lower bound on the top quark mass. Measurements of the B_d^0-$\bar{B}_d^0$ mixing rate moved the preferred top quark mass above 100 GeV [132], and precise measurements of the W boson mass by UA2 provided the indirect mass measurement $m_t = (160^{+50}_{-60})$ GeV, putting the top quark out of reach of S$p\bar{p}$S. In 1995 the CDF and

[13] The prediction of a third generation of matter before any of its members were experimentally observed was due to Kobayashi and Maskawa [170]. The historical context of their argument is more complicated and interesting than the conventional attribution concedes [201].

[14] For reference LEP, the most powerful electron-positron collider built to date, achieved collision energies of 209 GeV by its retirement in 2000, whereas the $t\bar{t}$ threshold is near 345 GeV.

[15] This conversion from a single-beam synchrotron into a storage ring with high beam luminosities was possible due to the invention of stochastic beam cooling by Simon van der Meer [180].

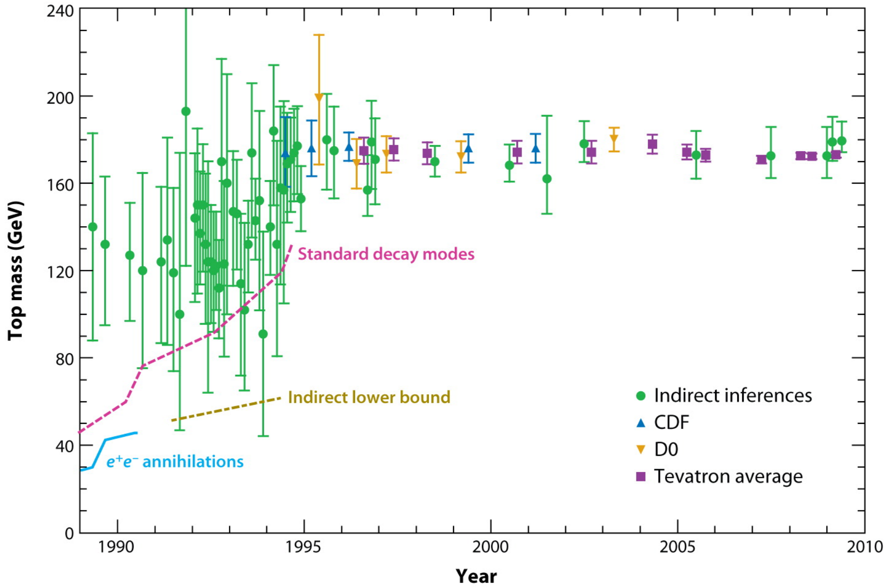

Figure 1.7: History of knowledge about the top quark mass before the LHC. Taken from Ref. [192].

D0 experiments at the TEVATRON discovered the top quark [108, 127] via $t\bar{t}$ production decaying via W bosons to leptons, measuring its mass to be $m_t = (174 \pm 16)\,\text{GeV}$, consistent with global fits to EW precision data available at the time [191, p. 8]. Today's world average value of the top mass is $m_t = (172.9 \pm 0.4)\,\text{GeV}$, combining measurements from the TEVATRON and LHC experiments [205]. The history of m_t constraints and measurements is charted in Fig. 1.7.

Moving from the TEVATRON centre-of-mass energy of $\sqrt{s} = 1.8\,\text{TeV}$ to the LHC design energy of 14 TeV, the $t\bar{t}$ production cross-section increases by several orders of magnitude. Figure 1.8 shows the cross-sections of various processes over these energy ranges; $\sigma_{t\bar{t}}$ indicates the $t\bar{t}$ cross-section which, while lower than the production cross-sections of inclusive QCD jets or the EW bosons W and Z, is larger than the production cross-section of the Higgs boson. Together with the substantially higher beam luminosity of the LHC, this has allowed a large sample of top quarks to be collected for analysis.

1.5.1 Top quark decay

The top quark decays predominantly into a W boson and a down-type quark. Alternative decays have not been directly observed, but upper limits have been placed on their proportions of the total top decay cross-section. The ratio R of the $t \to Wb$ decay to $t \to Wq$ (where q is one of d, s, or b) takes the value $R = 0.957 \pm 0.034$ when combining TEVATRON and LHC measurements [205]. The process $t \to Wb$ therefore accounts for nearly 100 % of top decays, consistent with the $|V_{tb}|$ element of the CKM matrix taking a value close to unity.

The theoretical description of the top quark decay is complicated by its short lifetime. The simplest treatment is to consider the top quark to be stable when it is produced, and then subsequently model the decay. This approximation introduces an error to observables of order Γ_t/m_t where Γ_t is the intrinsic width of the top quark. An improvement to this is the narrow-width approximation used in most calculations of top quark observables in which the top quark is allowed to decay while remaining on-shell. This approximation leads to improved kinematics for the jets and leptons in the decay as well as preserving spin correlations. Finite-width effects that violate the assumptions of the narrow-width approximation are small in observables of single top-quark decays except near

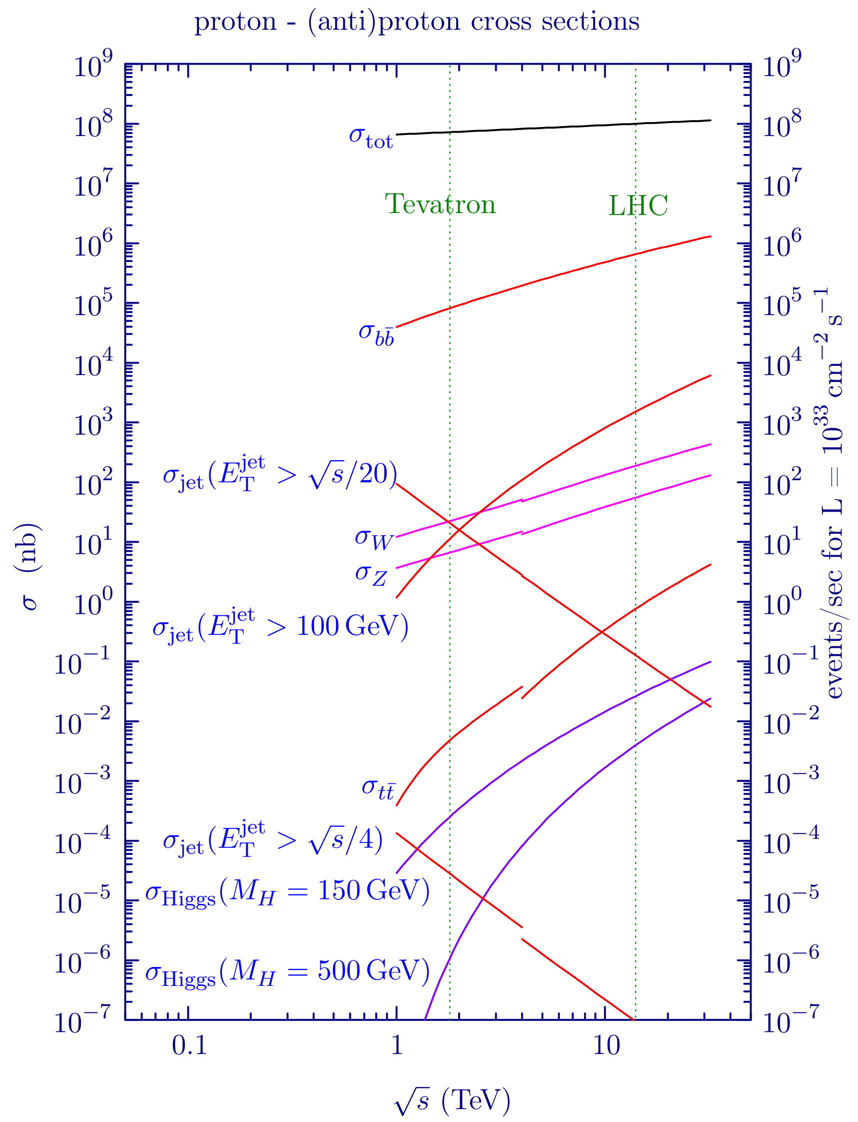

Figure 1.8: Predicted cross-sections of various processes at hadron colliders over a range of collision energies including those of the TEVATRON and the LHC. The right-hand axis shows the event rate corresponding to the cross-sections for a fixed luminosity. The discontinuity in the curves corresponds to a change in the choice of colliding particles from $p\bar{p}$ to pp. Taken from Ref. [105].

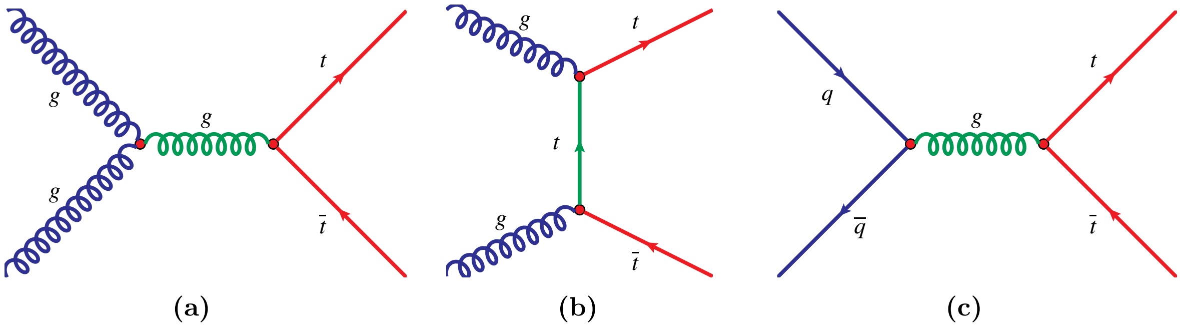

Figure 1.9: Examples of $t\bar{t}$ production at leading order.

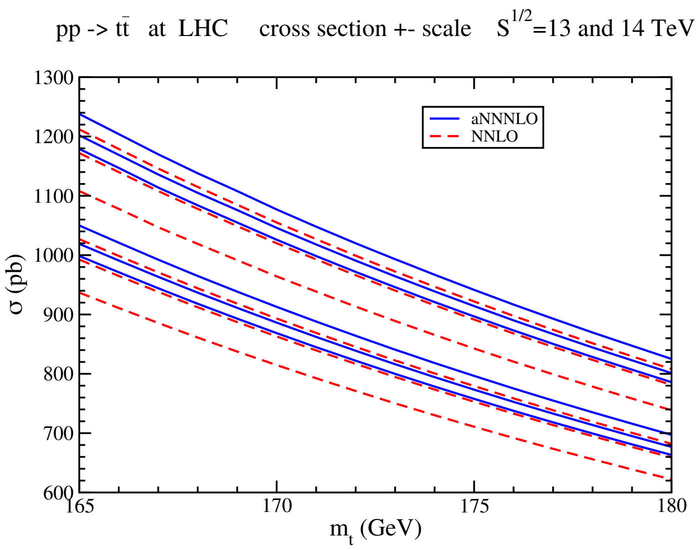

Figure 1.10: Cross-section prediction for $t\bar{t}$ at approximate NNNLO (solid blue) and NNLO (dashed red) as a function of the top quark mass. The top set of lines show the $\sqrt{s} = 14\,\text{TeV}$ predictions while the lower sets show the 13 TeV predictions. Within each set of three lines, the central line corresponds to $\mu = m_t$ while the other lines correspond to the choices $\mu = m_t/2$ and $\mu = 2m_t$. Taken from Ref. [165].

the edges of certain distributions, for example above the invariant mass $m(\ell b) \equiv \sqrt{m_t^2 - m_W^2}$ where nonresonant contributions are large [137, 196].

1.5.2 Top quark pair production

At leading order, $t\bar{t}$ production proceeds according to the strong interaction diagrams in Fig. 1.9, including gluon fusion (gg, Figs. 1.9a and 1.9b) and quark annihilation ($q\bar{q}$, Fig. 1.9c). In $p\bar{p}$ collisions at the TEVATRON, $q\bar{q} \to t\bar{t}$ was the dominant process, but with the large gluon flux in higher-energy pp collisions at the LHC, $gg \to t\bar{t}$ accounts for more than 90 % of $t\bar{t}$ production. In part, this difference is due to the change in the species of colliding particles. In order to produce top quark pairs at rest, an energy of $2m_t$ must be available to the partonic process, from which the threshold constraint $x_i x_j \geq 4m_t^2/s$ follows. The typical parton momentum fraction, x, scales as $1/\sqrt{s}$, and so at the $t\bar{t}$ threshold $x_i x_j$ is smaller at the LHC. Since the gluon PDF is much larger than the quark PDFs at small x (see Fig. 1.2), this also contributes to the difference in dominant production processes [191, p. 10].

Computing the $t\bar{t}$ cross-section requires a choice of normalisation and factorisation scale μ, the value of α_s computed at that scale, and the PDFs. Cross-section predictions at NLO and NNLO [155, 90] in α_s suffered from large uncertainties near the threshold due to high sensitivity to μ. Subsequent work to include logarithmic corrections at next-to-next-to-leading logarithm (NNLL) level [154, 156, 169] greatly reduced this uncertainty. More recently, soft gluon corrections at higher orders have been included to calculate approximate NNNLO cross-sections and kinematic predictions for $t\bar{t}$ [91,

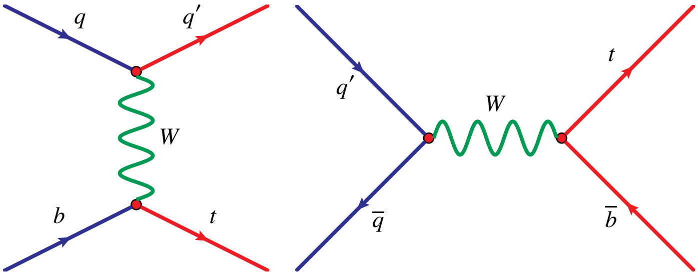

(a) t-channel **(b)** s-channel

Figure 1.11: Examples of t-channel and s-channel production at leading order.

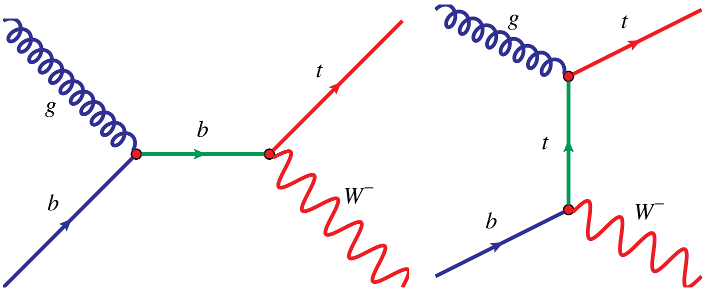

(a) **(b)**

Figure 1.12: Examples of tW production at leading order. The processes are distinguished by the momentum flow; diagram (a) is sensitive to resonant enhancement via the b-quark propagator, for example.

165, 166]. The approximate NNNLO predictions in Fig. 1.10 show a reduced sensitivity to the choice of scale compared to NNLO, and a slight overall enhancement in the cross-section. PDF uncertainties for LHC predictions are smaller than at the TEVATRON since the low-x regime is better constrained by empirical data.

1.5.3 Single top quark production

Whereas top quarks are produced abundantly in pairs via the strong interaction, the production of single top quarks via the weak interaction is a rarer process. There are three canonical single top quark production processes: the t-channel process $qb \rightarrow q't$ where a b-quark radiates a W boson to produce a top quark (Fig. 1.11a); the s-channel process $q\bar{q}' \rightarrow t\bar{b}$ where a pair of quarks annihilate to produce a W boson that then decays into a b-quark and a top quark (Fig. 1.11b); and the tW channel shown in Fig. 1.12 where a top quark is produced in association with a W boson. Of these three processes, the t-channel has the largest cross-section, accounting for about 73 % of the total single top cross-section at the LHC. It was discovered in the final two years of the TEVATRON's data-taking period by the D0 and CDF experiments [126, 107]. The tW processes accounts for about 24 % of the total single top cross-section at the LHC with the s-channel process making up the remaining 2 %. The low cross-sections of the tW and s-channel processes at the TEVATRON meant that they were too rare to be observed there[16].

Predictions for the s-channel and t-channel cross-sections at approximate NNLO are available. At the LHC, these channels exhibit an asymmetry in top and anti-top production, with σ_t measured to be close to twice the value of $\sigma_{\bar{t}}$ in the t-channel. Predictions for the tW process are available at NLO+NNLL and approximate NNNLO accuracy [147, 167]. Evidence for tW production was found at

[16] The tW process was suppressed at the TEVATRON due to the high mass of the final state particles, but with the high collision energies at the LHC the cross-section is correspondingly higher. The s-channel process, on the other hand, is suppressed at the LHC with respect to the TEVATRON since it depends on the density of anti-quarks in the colliding hadrons.

the LHC by the ATLAS and CMS experiments at $\sqrt{s} = 7\,\text{TeV}$ [32, 109], which proceeded to discover the tW process at 8 TeV [52, 114]. Additional rarer single-top processes are predicted by the SM including the production of a single top quark in association with a Z or Higgs boson.

In order to comprehensively search for evidence of new physics, it is necessary to consider both $t\bar{t}$ production, sensitive to the strong coupling of the top quark to gluons, and single-top production, sensitive to the t-W-b vertex. New physics at some scale Λ is expected to modify top quark decays with an effect scaling as some integer power of m_t/Λ and to modify single top quark production with an effect scaling as some integer power of $\sqrt{s}/\Lambda$ [203]. At hadron colliders at the energy frontier where $\sqrt{s} \gg m_t$ this makes the production cross-section a much more sensitive avenue for probing such new physics than the decays.

Certain new physics processes would be expected to affect all channels of single top quark production equally. For example the addition of a fourth generation of heavy quarks would be expected to modify the CKM matrix element V_{tb}, and since the cross-sections of all three canonical single top channels are proportional to $|V_{tb}|^2$ they would be modified in the same direction. Other new physics can modify the single top quark cross-sections independently. An additional heavy particle, for example a new charged boson, would be expected to contribute a resonant enhancement to the s-channel process (which has an intermediate virtual W boson). This would manifest as a larger overall s-channel cross-section and potentially a resonant peak in the mass of the $t\bar{b}$ system, without greatly affecting the other channels. Modifications to the t-W-b vertex would affect all three channels, but since the top quark and W boson both appear in the final state of the tW process, it is insensitive to most other new physics such as flavour changing neutral currents [203]. Study of all single top channels separately therefore offers information about the nature of any new physics that would be impossible to resolve by studying any channel in isolation.

1.5.4 Ambiguity in the tW process

The t-channel and s-channel processes are unambiguous and distinct processes in both the 4FS and the 5FS beyond leading order, however the same is not true of the tW and $t\bar{t}$ processes. The $gg \to tWb$ partonic process (e.g. Fig. 1.13b) and the $gg \to t\bar{t}$ partonic process (e.g. Fig. 1.13a) feature the same outgoing particles when the top decay is considered. These processes interfere when the tWb process features a second top quark propagator that is on-shell. In order to create a fully consistent picture of this interference it is necessary to consider the parent $W^+bW^-\bar{b}$ process in the 4FS. The amplitude of this process is then defined by

$$|\mathcal{A}_{tWb}|^2 = |\mathcal{A}_{1t}|^2 + 2\Re(\mathcal{A}_{1t}\mathcal{A}^*_{2t}) + |\mathcal{A}_{2t}|^2 \tag{1.2}$$

where $\mathcal{A}_{1t}$ is the amplitude of the singly-resonant component (Fig. 1.13b) and $\mathcal{A}_{2t}$ is the amplitude of the doubly-resonant component (Fig. 1.13a) which is the dominant contribution at LHC energies. Note that there is also a non-resonant component (Fig. 1.13c) but its contribution is negligibly small. As a practical matter it is still desirable to be able to isolate the tW part—the singly-resonant part—of this from the $t\bar{t}$ part in order to consider them separately and therefore it is necessary to define a scheme to perform this separation.

Earlier attempts to resolve the interference suffered from various shortcomings, for example a kinematic cut on the mass of the bW system [204] was difficult to apply experimentally due to

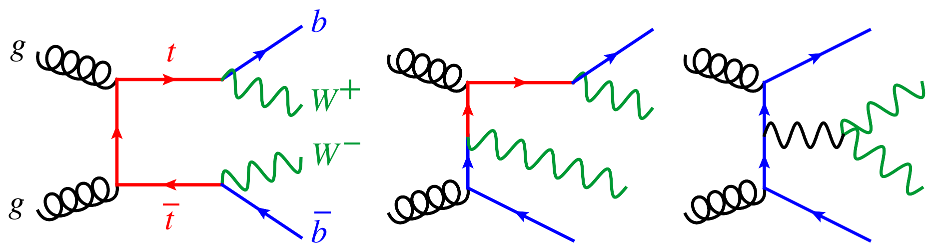

(a) $\mathcal{A}_{2t}$ (b) $\mathcal{A}_{1t}$ (c) $\mathcal{A}_{0t}$

Figure 1.13: Diagrams showing the $W^+ b\, W^- \bar{b}$ process with diagrams exhibiting (a) double top quark resonance, (b) single top quark resonance, and (c) no top quark resonance. Taken from Ref. [129].

the ambiguity in assigning observed particles to the final state partons, while more sophisticated subtraction schemes [164] led to difficulty in incorporating NLO corrections. Schemes in common use today fall into two main families [145]:

- diagram removal (DR) involves discarding some of the terms in Eq. (1.2). DR1 keeps only the purely singly-resonant term $\mathcal{A}_{1t}$ while DR2 keeps also the interference term. This scheme is not gauge invariant and comes with some technical complications related to numerical stability and the fact that the matrix element is not necessarily positive definite.

- diagram subtraction (DS) involves crafting a subtraction term that will cancel the doubly-resonant term $\mathcal{A}_{2t}$ exactly when on the resonant pole while vanishing elsewhere. While this procedure retains gauge invariance, it does so by reshuffling the external momenta to put the internal $\bar{t}$ on-shell. Since this process affects the mass of the $W b$ system but is not a Lorentz transformation, it introduces an arbitrary choice into the design of the subtraction term.

Predictions for the differential cross-section of $pp \to t\, W^- b$ with respect to the $b\, W$ invariant mass are shown in Fig. 1.14 for the first- and second-generation variants of the DR and DS schemes, highlighting their different modelling of the interference region around the top quark mass and in the tail of the mass distribution. The second generation of schemes is thought to more accurately describe the behaviour, whereas the first generation somewhat overestimates the cross-section beyond the top quark mass, and may also overestimate the tail of the $p_{\mathrm{T}}(b)$ distribution [129]. The schemes also evolve differently with the choice of the renormalisation and factorisation scale, with the two second-generation schemes relatively flat across the range of scales relevant to the LHC while the first-generation schemes predict a cross-section that falls slightly as the scale increases.

More recently, an MC generator called `bb4l` [162] has been implemented in the POWHEG framework, providing a full, resonance-aware NLO description of the partonic process $pp \to \ell\nu_\ell\bar{\ell}\bar{\nu}_\ell b\bar{b}$. This description offers many technical advantages over alternative generators, including its exact treatment of spin correlations and off-shell effects in the top decay, incorporation of finite-width effects, and consistent treatment of the entire phase space by using a finite b-quark mass. Implementations of the second-generation DR and DS schemes and the `bb4l` generator were not fully available at the time of much of the work in this thesis and so only the first-generation DR and DS schemes are treated.

The interference effects are relevant primarily in events with two b-quarks in the final state and so by rejecting events with a second b-quark it is possible to greatly suppress not only the doubly-resonant contribution but also the interference term.

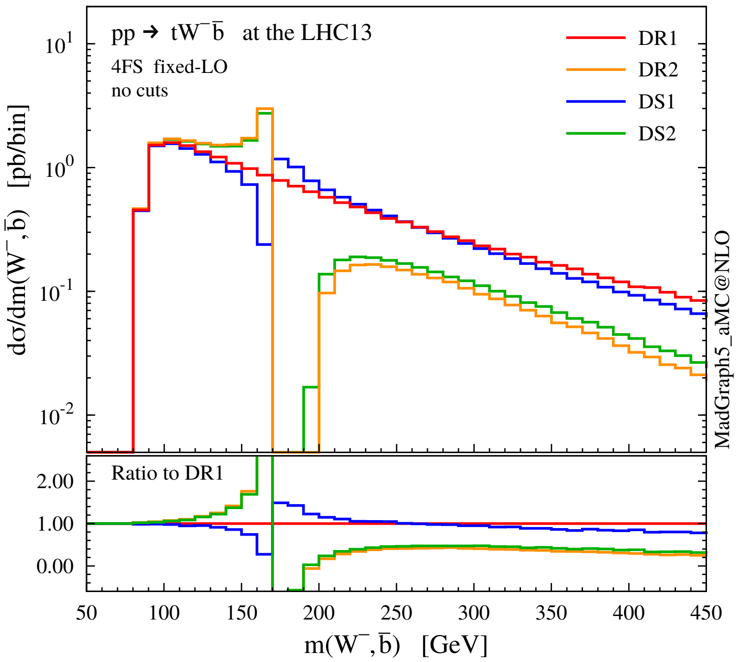

Figure 1.14: Invariant mass of the bW system in the $pp \to tW^-b$ process comparing different DR and DS schemes for isolating the singly-resonant component. The term "LHC13" refers to the LHC with $\sqrt{s} = 13\,\text{TeV}$. Taken from Ref. [129].

Chapter 2

The Large Hadron Collider and the ATLAS Experiment

The LHC [136] at CERN is currently the world's most powerful particle accelerator. It is located in a 27 km tunnel 100 m to 120 m below the ground near Geneva, Switzerland. Of the six experiments hosted at the LHC, two are general purpose particle detectors: the ATLAS experiment [77] and the CMS (Compact Muon Solenoid) [117] experiment. Along with the other more specialised experiments, these contribute to a diverse physics program at the LHC.

2.1 The Large Hadron Collider

The LHC accelerates beams prepared by CERN's accelerator complex. The majority of the operational schedule is devoted to proton-proton (pp) collisions, with some time reserved for heavy ion beams (usually lead ions, with one period more recently devoted to xenon ions). Following the illustration in Fig. 2.1, the proton beams start from a hydrogen gas source; the gas is ionised, and then accelerated by the linear accelerator LINAC2 to 50 MeV. The beams are then accelerated by the Proton Synchrotron BOOSTER to 1.4 GeV, the Proton Synchrotron to 25 GeV and then the SPS to 450 GeV. From the SPS the beams are injected into the two LHC beamlines for acceleration to their final collision energy, currently 6.5 TeV for a beam centre-of-mass collision energy of 13 TeV, and the beams are crossed at the four interaction points (IPS) located in the experimental halls. The accelerator complex supplies beams in a number of configurations for many of CERN's other experiments in addition to those at the LHC.

The first running period of the LHC, called Run 1, included a period of commissioning in 2010, followed by data-taking at $\sqrt{s} = 7$ TeV in 2011 and at 8 TeV in 2012 with a combined integrated luminosity of 25 fb^{-1} collected. After closing for upgrades and maintenance for two years, Run 2 of the LHC began in 2015 at $\sqrt{s} = 13$ TeV. The beam luminosity was progressively increased in the following years, exceeding the design luminosity in order to collect a total integrated luminosity of 140 fb^{-1} by the end of 2018 when Run 2 ended. The LHC is currently shut down for upgrades and maintenance, and is scheduled to come online again in 2021 for its next data-taking run (see §2.1.4).

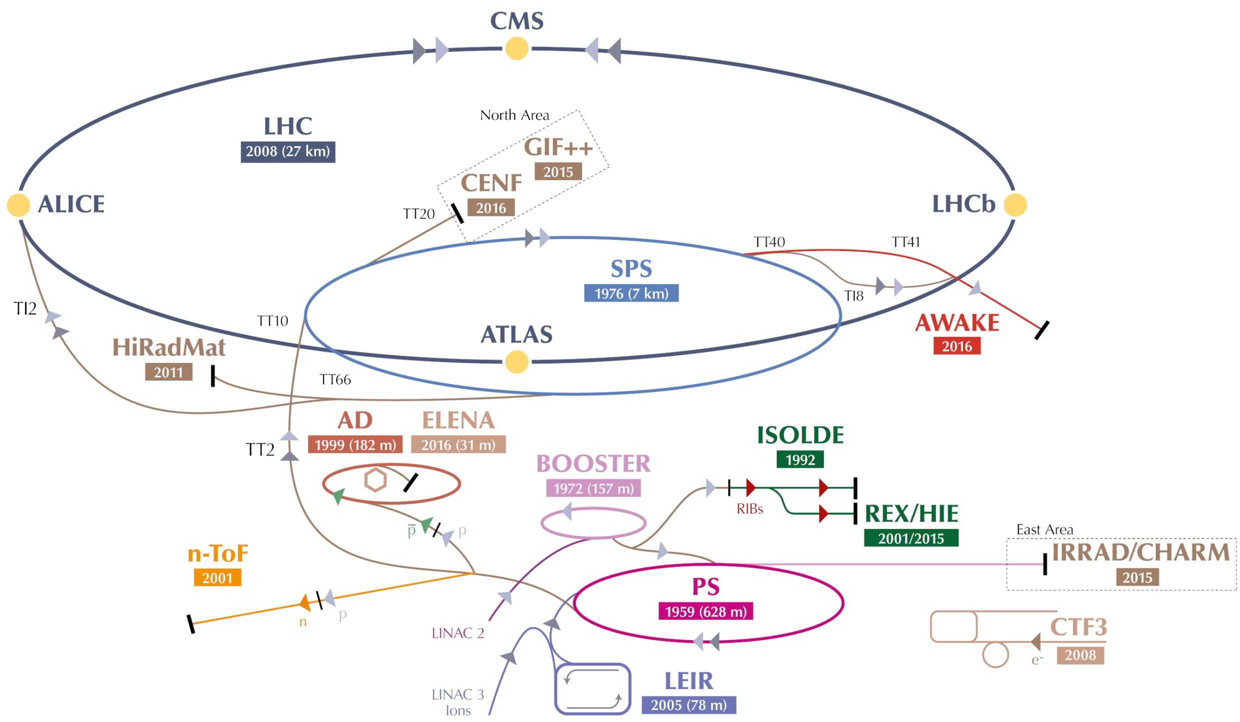

Figure 2.1: The CERN accelerator complex. Taken from Ref. [181].

2.1.1 Design considerations

The particle content of beams for a collider is informed by the physics goals of the design, but also has a number of important engineering consequences. Lepton colliders like LEP (previously hosted in the tunnel now occupied by the LHC) or SuperKEKB [4] (an electron-positron collider in Japan hosting the currently active Belle II experiment) can be tuned to a selected energy since all of the beam particle's momentum is typically transferred into its decay products. This allows for precision measurements of masses of known particles, or tuning to the mass resonance energy of a selected particle to explore its decay spectra. At hadron colliders however, the colliding particles are partons carrying a variable fraction of the beam particle's momentum. Hadron colliders are therefore typically built to explore a range of new physics, particularly at the energy frontier, in order to observe collisions over a wide range of energies.

Synchrotron radiation is a significant challenge at circular lepton colliders, but is less significant at hadron colliders, allowing the latter to operate at much higher beam energies limited by the strength of the magnets and the size of the tunnel. Whereas in fixed-target experiments a beam with low luminosity striking a solid target provides an adequate event rate[1], colliding beams require high luminosities to attain an appreciable rate. Interactions between particles cause beam heating[2], i.e. the development of a transverse momentum component, which is made worse by the need to store the beam for long enough to build up the beam current by injecting bunches of particles.

The electric charge of the beams is also an important consideration since if the beams have

[1] For fixed target experiments though the centre-of-mass energy available to the collision grows only as the square root of the beam energy. At colliders the centre-of-mass energy is the sum of the beam energies.

[2] The stochastic cooling technique of van der Meer mentioned previously—and subsequent developments of the technology—addresses this issue.

opposite charge they can be accelerated in the same magnetic channel, and this was in part behind the choice to collide protons with anti-protons at the TEVATRON. Since the LHC collides protons with protons (and was designed also for an ion acceleration program), it requires two independent magnetic channels and vacuum systems to enable the counter-rotation of the two positive beams.

The magnet system of the LHC consists of 1232 superconducting dipole magnets (with the two beam pipes sharing a common cryogenic system and power supply) to maintain the circular trajectory. The maximum field strength of 8.3 T together with the radius of the tunnel correspond to a maximum possible energy of 7 TeV per beam [184]. A number of other magnets with more complicated field structures are used for focusing and manipulating the beam, including 392 quadrupole magnets.

Beam instabilities arise from a range of sources, from residual electron clouds inside the vacuum to detectors placed very close to the beam and EM interactions between the beams. The LHC is instrumented with monitoring devices that supplement input from the experiments to allow for operation and correction of the beam, and with an interlock system that automatically dumps the beam[3] when its parameters exceed limits established to protect the hardware.

2.1.2 Luminosity and beam parameters

The proton beams injected into the LHC consist of a pattern of up to 2808 bunches of protons, separated by a spacing of 25 ns[4]. Depending on the filling scheme used by the LHC, most bunches are in slots that are paired with a slot in the other beam so that the bunches will collide when the beams are crossed. Some bunch slots may be unpaired or completely unfilled in order to assist with beam background measurements or for operational reasons.

The instantaneous luminosity $\mathcal{L}$ is related to the cross-section σ of a process and the number of events N by

$$N = \sigma \int \mathcal{L}\, \mathrm{d}t.$$

The expected rate of inelastic proton interactions due to a pair of bunches colliding head-on is the product of μ, the average number of inelastic collisions each time the bunches cross, and f_r, the bunch crossing frequency. Given the cross-section for inelastic[5] pp interactions σ_{inel}, it follows that the luminosity of a single colliding bunch pair is given by

$$\mathcal{L}_b = \frac{\mu f_r}{\sigma_{\text{inel}}}. \tag{2.1}$$

Taking the sum over n_b pairs of colliding bunches, the instantaneous luminosity of the beam is given by

$$\mathcal{L} = \sum_{b=1}^{n_b} \mathcal{L}_b = n_b \langle \mathcal{L}_b \rangle_{\text{beam}} = n_b \frac{\langle \mu \rangle_{\text{beam}} f_r}{\sigma_{\text{inel}}}$$

[3]The beam dump system must dissipate the 350 MJ beam energy in the 90 µs full circulation time. This is done by deflecting the beam to sweep out a rough circle of 120 cm diameter on a 7 m deep carbon absorbed block surrounded by concrete and iron shielding [97].

[4]In 2015 from the first stable beams in early June until mid-July the bunch spacing was 50 ns to avoid electron cloud effects. A "scrubbing" procedure was implemented to reduce the electron clouds in the beam pipe, allowing the return to the nominal bunch spacing. While this initial data is usable for physics analysis, it has different beam background characteristics (a smaller spacing means more out-of-time pileup) and so is excluded for simplicity [212].

[5]Elastic or other "glancing" interactions between protons are not typically detected by the main LHC experiments since the products of these interactions have momenta closely aligned to the beam and so do not reach the detectors outside the beam pipe.

Table 2.1: Typical parameters of the LHC beam in each year of operation [46, 95, 5, 212]. The values shown are for pp collisions at the ATLAS interaction point IP1. Note that 2010 was largely devoted to commissioning the accelerators; only a small amount of integrated luminosity was collected for physics analysis.

Parameter		Run 1			Run 2			
	Design	2010	2011	2012	2015	2016	2017	2018
Beam energy [TeV]	7	3.5	3.5	4.0	6.5	6.5	6.5	6.5
Colliding bunch pairs n_b	2808	348	1331	1380	1813	2208	2544	2544
Bunch spacing [ns]	25	150	50	50	25	25	25	25
Protons per bunch [10^{11}]	1.15	1.0	1.3	1.5	1.1	1.1	1.1	1.1
Minimum β^* [m]	0.55	2.0	1.0	0.6	0.8	0.4	0.3	0.25
Emittance at start [µm]		2.2	2.3	2.5	3.5	2.2	2.2	1.9
Peak lumi. [$10^{34}\,\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$]	1.0	0.02	0.35	0.77	0.51	1.4	2.1	2.1
Interactions/crossing $\langle\mu\rangle$		2	9	21	13	25	38	37
Integrated lumi. [fb^{-1}]		0.05	5.5	22.8	4.2	39.7	50.6	66
Days of pp operation			160	201	88	146	140	145
Stable beam time [%]			33	36	35	49	49	49

where here the angle brackets denote averages over bunches in the beam [46]. Note that in general the symbol $\langle\mu\rangle$ denotes the average number of interactions per bunch crossing corresponding to a given dataset, in which context it is an average over time and over multiple fills of the collider. Typical values of the parameters determining luminosity for each year of operation of the LHC are shown in Table 2.1.

In order to maximise the delivered luminosity, given that the crossing frequency and smallest bunch spacing are fixed parameters in the design of the LHC, the parameter to maximise is $\langle\mu\rangle$. In practice this is achieved by concentrating the beams into a smaller interaction volume, and increasing the bunch populations within the operational constraints of the machine. The size of the beam is characterised by the transverse emittance ϵ, and is determined initially by the injectors and subsequently by the ability of the LHC to damp emittance growth throughout a fill. The amplitude function β^* measures the distance from the IP at which the cross-sectional size of the beam is twice that at the interaction region, where smaller values result in higher luminosity, limited by the beam aperture in the focusing region around the IP as shown in Fig. 2.2.

During Run 2, technological advances in the beam injection and optics allowed the LHC to be operated at a peak instantaneous luminosity more than twice the design value of $10^{34}\,\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$. Since this was achieved by increasing $\langle\mu\rangle$, the experiments faced the challenging task of operating with many more interactions in each crossing than the preferred multiplicities, resulting in high-pileup conditions.

2.1.3 Operating conditions and challenges

Experience gained from operating the LHC in Run 1 and in the first part of Run 2 allowed its very successful operation over the bulk of Run 2 with 49 % availability of stable beams, as shown in Table 2.1. Compared to Run 1, the beam optics and structure were more closely aligned to their design configuration. As a fill progresses, the beam luminosity decreases until it becomes preferable

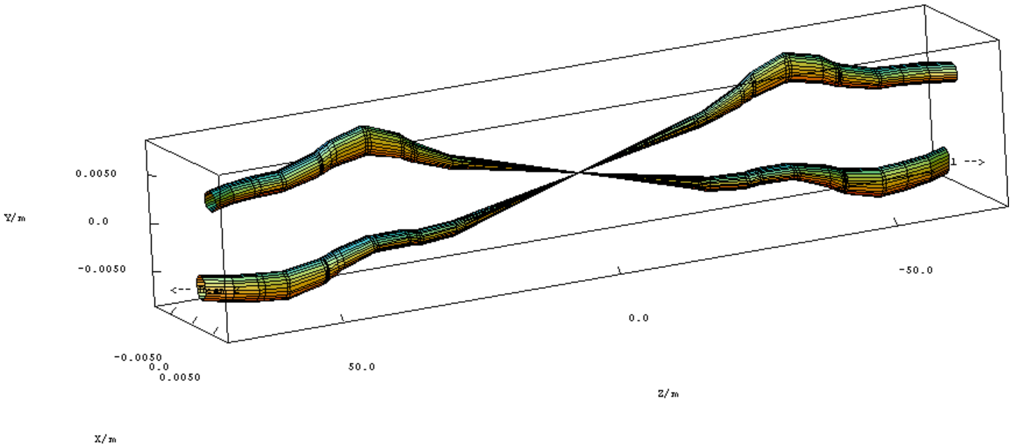

Figure 2.2: Illustration of the 1σ envelopes of the two beams as they are brought together at the ATLAS interaction point at the LHC using parameters from Run 1. Taken from Ref. [163].

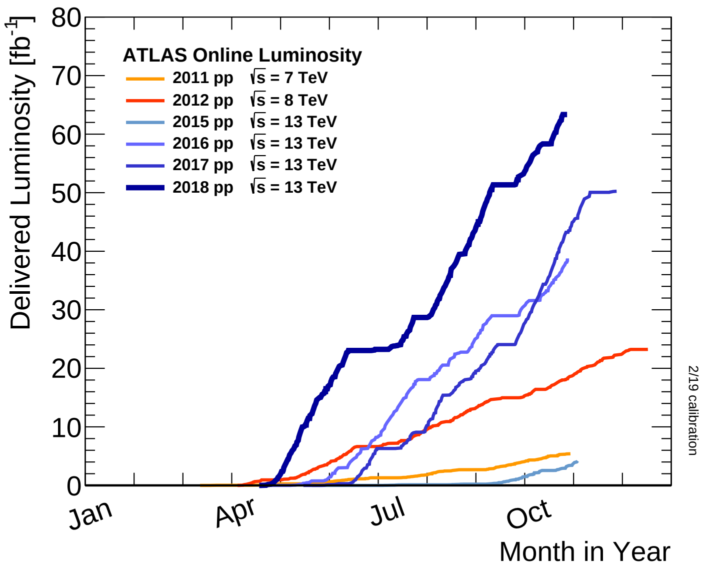

Figure 2.3: Cumulative integrated luminosity of pp collisions delivered by the LHC to ATLAS during each year since 2011. Flat segments correspond to periods of LHC maintenance or operational studies when physics data was not being collected. Taken from Ref. [17].

to dump the beam and incur the overhead of a new injection sequence to fill the LHC from the accelerator complex. The improved stability led to longer fill durations, often exceeding 12 h, and fast turnarounds between fills of as little as 1 h, allowing for more efficient operation. The integrated luminosity delivered by the LHC for each year of operation is shown in Fig. 2.3, where the final three years of Run 2 are clearly distinguished by their rapid accumulation of collision data.

The successful operation of the LHC was threatened by a number of operational issues throughout Run 2 for which various mitigations and novel techniques were developed [212]. A series of problems in the injector system limited each injection to 144 bunches, requiring more injection steps and restricting the beam configuration of the machines earlier in the accelerator chain. A novel injection scheme called batch compression, bunch merging and splitting [128] was implemented in July 2016 within this constraint to increase the beam brightness and reduce transverse emittance. An object was identified lying inside the beam pipe; the beam was steered around it to avoid instabilities until the warming of the cryogenics after Run 2 when the object was extracted and found to be a strip of plastic. Limitations on the minimum β^* of the LHC were overcome by the implementation of a novel achromatic telescopic squeezing [138] optics scheme in 2017, allowing the luminous region at the

IPs to be shrunk. Periodic levelling of the crossing angle at the ATLAS and CMS interaction points during the course of a fill also helped to maximise the delivered luminosity by 3 % to 4 % in 2017. This technique was improved in 2018 with the addition of β^* levelling to further exploit the long fills.

In 2017 the major operation challenge of Run 2 emerged when sudden beam losses in sector 16L2[6] of the LHC resulted in unscheduled beam dumps. The problem was eventually traced to an incident during the replacement of magnets in that sector before Run 2 where gas had entered the beam pipe and subsequently condensed. When this gas was warmed by the beam it would degrade the vacuum and form a plasma that interacted with the beam. An attempt to resolve this by warming the beam screen[7] only worsened the situation. In September 2017 the injection scheme was modified so that groups of 8 filled bunches were followed by 4 empty slots to reduce the heat load on the condensed gas. During the short shut-down of the LHC at the start of 2018, sector 16L2 was warmed and purged of light gases, leaving only a small amount of condensed water vapour. This mitigated the problem in 2018, where the number of related beam dumps was 17 compared to about 70 in 2017. The development of a short conditioning fill configuration allowed the residual effects to be managed so that a return to the nominal bunch configuration was possible.

2.1.4 Future upgrades to the LHC

After a third run of the LHC at its design energy of $\sqrt{s} = 14\,\text{TeV}$, a large dataset will have been accumulated. For measurements that are limited by the available data statistics, the LHC will not offer much opportunity for improvement since halving the statistical uncertainty requires a dataset larger by a factor of 4. In anticipation of this situation, research and development towards an upgrade to the LHC is underway. The High Luminosity LHC (HL-LHC) is scheduled to be installed from 2023 to begin delivering collisions from 2025 at a luminosity greater than that of the LHC by a factor of 5 to 7 [98].

This improvement in luminosity will be achieved through novel accelerator techniques. Many of these techniques have been tested at the LHC, while others are under development, such as “crab cavities” which deflect the head and tail of each bunch immediately before collision to compensate for the crossing angle, thereby maximising the overlap of the colliding bunches [102]. Upgrades throughout the magnets, cryogenic systems, injectors, collimators, and beam instrumentation are planned, with some major upgrades to the CERN accelerator complex already commencing ahead of the third LHC run. The LHC experiments will also undergo major upgrades to cope with the high pileup and associated challenges in the high luminosity environment.

2.2 The ATLAS Detector

The ATLAS detector is installed at IP1 of the LHC and is a general purpose detector designed to facilitate reconstruction of stable particles with sufficient performance to suit the requirements of a diverse physics program. It is depicted in Fig. 2.4 with several of its subdetectors identified. The detector has a cylindrical geometry arranged around the LHC beam axis with almost complete coverage in the azimuthal direction and a large coverage along the beam axis. The innermost layers of

[6] This means the half-cell 16 positions to the left of IP2.

[7] The beam screen is a steel device coated with copper that serves to shield the cold bores of the LHC superconducting magnets from the thermal load due to synchrotron radiation from the circulating beams [10, 194].

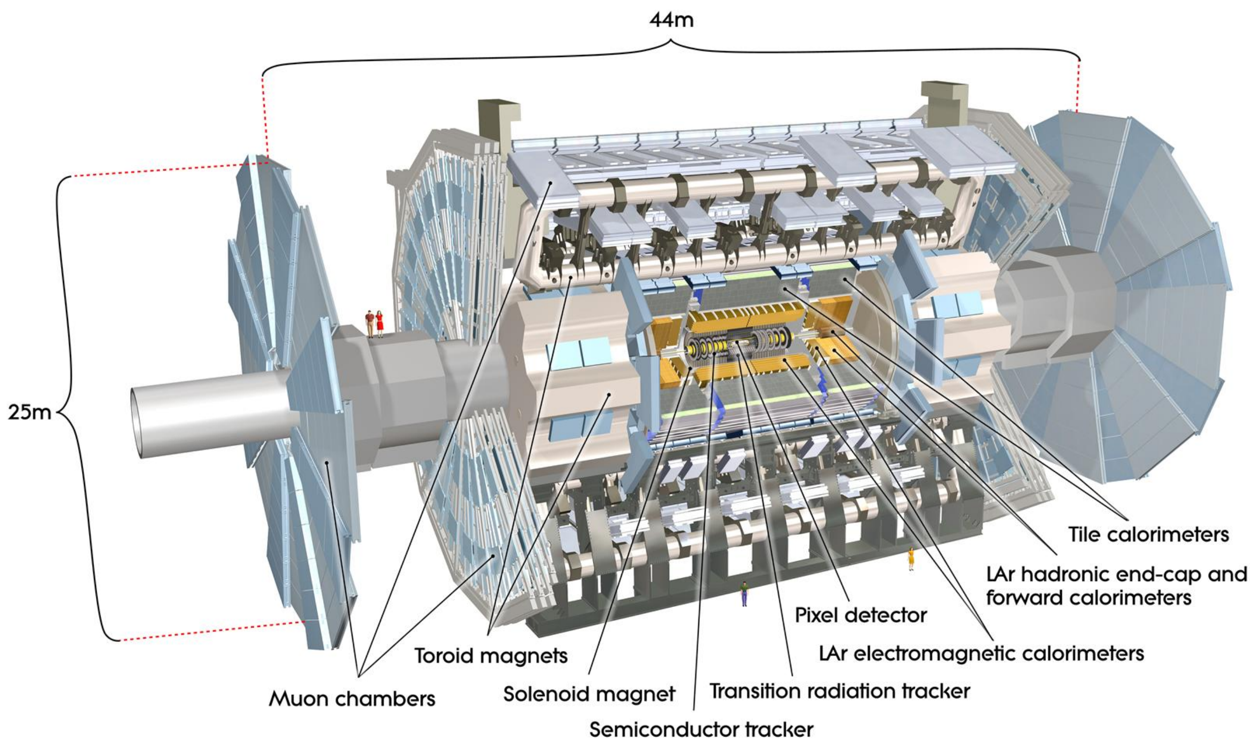

Figure 2.4: The ATLAS detector. Taken from Ref. [187]

the detector track the paths of charged particles in the presence of a magnetic field. The calorimetry system in the next layers measures the energy deposited by particles. The outermost layers of the detector are a dedicated muon system for the identification and measurement of energetic muons which typically penetrate the calorimeters. The large angular acceptance of the detector allows a transverse momentum imbalance to be measured in order to infer the presence of weakly interacting particles such as neutrinos which are not directly detectable within the detector volume.

2.2.1 Coordinate system

The coordinate system used by ATLAS has its origin fixed at the nominal IP inside the beam pipe at the centre of the detector. In practice it is nontrivial to establish the position of the subdetector elements in relation to the origin or to each other since there are large movements associated with thermal expansion and mechanical forces such as coolant flow; good alignment is achieved by an iterative procedure (see §2.6.3). The coordinate choice is right-handed with the z axis pointing along the beam axis, the y axis pointing upwards and the x axis pointing towards the centre of the LHC ring.

A system based on spherical polar coordinates is typically more convenient; these are the radial distance r, azimuthal angle (around the z axis) ϕ and, in place of the polar angle θ the related quantity pseudorapidity defined as $\eta \equiv -\ln\tan(\theta/2)$. In the relativistic limit pseudorapidity approximates rapidity, $y \equiv \frac{1}{2}\ln\left(\frac{E+p_{\mathrm{L}}}{E-p_{\mathrm{L}}}\right)$ where p_{L} is the longitudinal component of momentum. The magnitude of η varies from zero when pointing in the transverse direction, to infinity when pointing along the beam axis, with the ATLAS detector design optimised for good performance in the region $|\eta| < 2.5$.

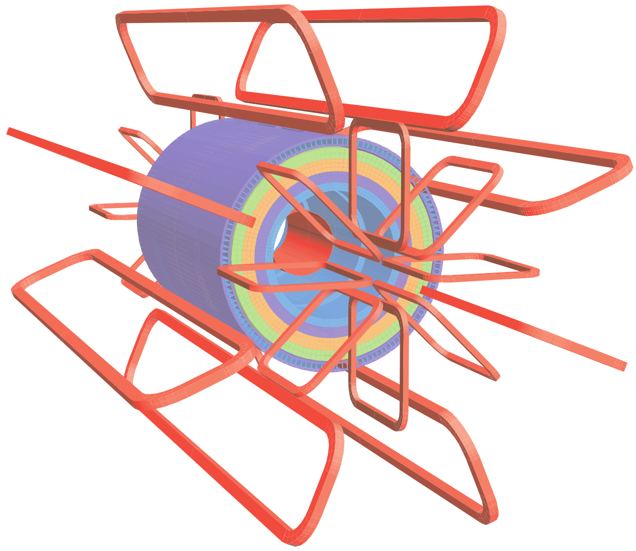

Figure 2.5: Illustration of the bare coils of the ATLAS magnet systems showing their locations and extent in the detector. The red cylinder in the centre is the solenoid magnet. The larger concentric cylinders represent layers of the calorimeter with different EM properties; the magnetised iron in the calorimeter acts as the solenoid return yoke. The 8 barrel toroid loops are arranged around the outside of the two end-cap toroids, which each consist of 8 loops. Taken from Ref. [77].

The angular separation $\Delta R_y \equiv \sqrt{(\Delta y)^2 + (\Delta\phi)^2}$ is Lorentz-invariant under boosts along the z-axis. In practice it is more convenient to use an alternative definition of angular separation in terms of purely angular quantities $\Delta R_\eta \equiv \Delta R \equiv \sqrt{(\Delta\eta)^2 + (\Delta\phi)^2}$, that is with pseudorapidity replacing rapidity. ΔR_η is Lorentz-invariant to longitudinal boosts only in the massless limit. Typically the transverse projections of momentum and energy are of interest, and these are defined as $p_\mathrm{T} \equiv \sqrt{p_x^2 + p_y^2}$ and $E_\mathrm{T} \equiv E \sin\theta$ respectively.

2.2.2 Magnets

There are three superconducting Al/NbTi/Cu magnet systems within the ATLAS detector, responsible for bending charged particles in order to facilitate measurements of momentum. Their design is summarised in Ref. [22] and represents an advance in magnet technology due to the unusually large size and demanding requirements. The components are illustrated in Fig. 2.5.

The central solenoid magnet sits at $r = 1.22\,\mathrm{m}$ to $1.32\,\mathrm{m}$, immersing the Inner Detector (ID) in a $2\,\mathrm{T}$ field. It consists of 1173 turns of superconducting wire over its length of $5\,\mathrm{m}$ carrying a nominal current of $7.73\,\mathrm{kA}$. In order to reduce the amount of material in the active region of the detector, the iron absorber from the barrel Tile Calorimeter (TileCal) is used as the return yoke. Since this is some distance outside the solenoid coil, the field in the tracking volume is not completely uniform, though it has been characterised to high precision.

The barrel toroid and the two end-cap toroids provide a field of $3.9\,\mathrm{T}$ to $4.1\,\mathrm{T}$ in the muon system, operating at a temperature of $4.7\,\mathrm{K}$ and a nominal current of $20.5\,\mathrm{kA}$. The barrel component is made from eight loops in independent cryostats, with an outer diameter of $20.1\,\mathrm{m}$ and a total cold mass of $370\,\mathrm{t}$. Each end-cap toroid consists of eight loops in a common cryostat with a $10.7\,\mathrm{m}$ outer diameter and a total cold mass of $160\,\mathrm{t}$, and is removable for access to the detector during maintenance periods.

2.3 The ATLAS Inner Detector

The Inner Detector (ID) [19, 20] was designed to provide high granularity tracking of charged particles in the region close to the IP with a p_T resolution of $\sigma_{p_\mathrm{T}}/p_\mathrm{T} = (0.05\,\% \times p_\mathrm{T}\,\mathrm{GeV}) \oplus 1\,\%$ [77] (where

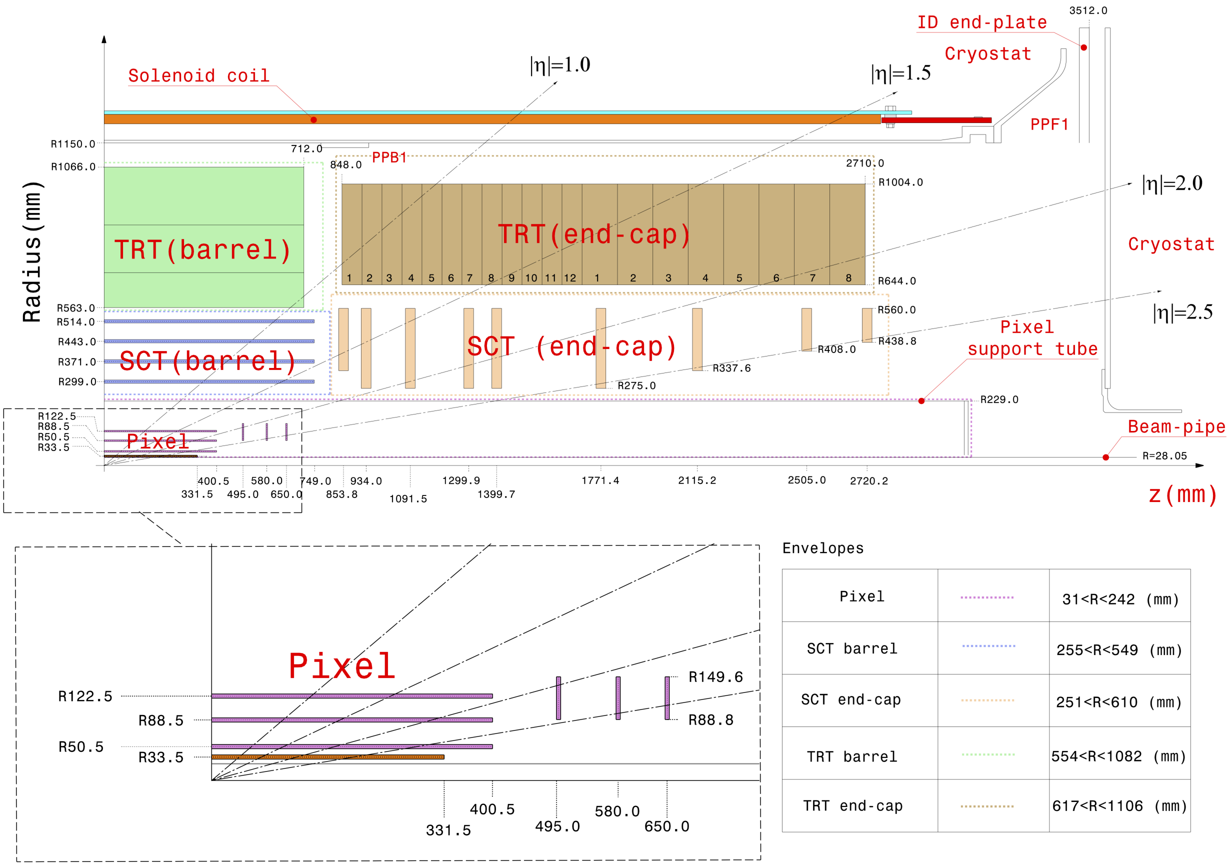

Figure 2.6: One quarter of the r–z cross-section of the ATLAS Inner Detector during Run 2. For the negative z values the detectors are a mirror image of what is shown here, and they are approximately rotationally symmetric around the z axis. Note that r is written as R in the diagram. Taken from Ref. [74].

the sum is in quadrature). It consists of two inner silicon semiconductor detectors surrounded by a gas-filled straw tube tracker, all immersed in the solenoid's 2 T field.

The subdetectors of the ID, as with most of ATLAS, have barrel and end-cap components to accommodate the cylindrical geometry. A quarter cross-section of the ID is shown in Fig. 2.6 with lines of fixed $|\eta|$ for reference. The arrangement of the barrel and end-cap components is such that most particle trajectories within $|\eta| < 2.5$ intersect a consistent number of tracking layers. In the transition between the barrel and end-cap regions there is a loss of performance due to the larger amount of material including services such as cabling, support structures, and cooling, as can be seen in Fig. 2.7.

The ID utilises detector elements with active materials that are ionised by the passage of charged particles but that absorb little of the energy of those particles. The design constraints are therefore to maximise the density of discrete tracking elements while minimising the amount of material between the IP and the calorimetry systems.

Over the course of operations of the LHC until the end of Run 2 the ID has slowly degraded due to radiation damage. Of the 98.65 million channels in the ID, approximately 4 million or 4.1 % are no longer operational. The majority of these failed channels are in the innermost layer[8], the pixel

[8] The 12 million channels of the inner pixel layer, the IBL, were still 99.3 % operational by the end of Run 2 since they were installed more recently than the other pixel layers.

Figure 2.7: Material in the ID in radiation lengths with respect to $|\eta|$. Each detector shows an increase in material in the transition region between the barrel and end-cap, complemented by additional material from services shown in black. Taken from Ref. [19].

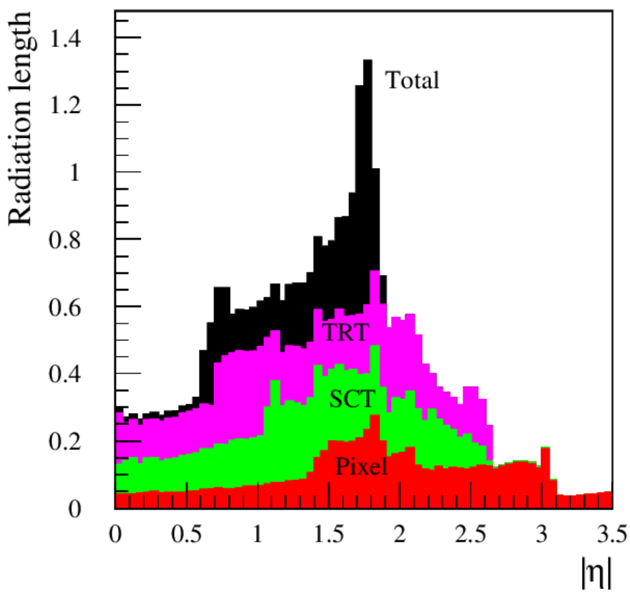

detector, which has the highest density of channels as well as the highest-radiation environment.

2.3.1 The pixel detector

The pixel detector is a silicon semiconductor detector with 92 million pixel channels within a total length of 1.3 m and $r = 3.1$ cm to 24.2 cm. Throughout Run 1 of the LHC it consisted of three concentric barrel layers and three end-cap wheels on each side. In advance of Run 2 of the LHC, an additional barrel layer called the Insertable B-Layer (IBL) was installed [21], shown as the brown layer closest to the beam pipe in the inset of Fig. 2.6. The transverse (r–ϕ) resolution is 10 µm and the longitudinal (z) resolution is 115 µm, allowing for tracking in the dense environment of the LHC at its current luminosity.

In the barrel region of the pixel detector, each detector module is 62.4 mm by 22.4 mm and contains 61 440 pixel elements read out by 16 chips. The modules are mounted on a circular support structure at an angle so that the active elements overlap, providing hermetic coverage. The position of the detector material in the barrel is shown in Fig. 2.8 by reconstructing the points of origin of particles arising from hadronic interactions within the material of the detector. In the end-cap region the detector modules have a similar design and are mounted flat onto supporting ring structures.

The active detector elements are made from silicon wafers with a nominal 150 V depletion voltage that are bump-bonded to chips containing the readout electronics. Electron-hole pairs created by the passage of a charged particle are separated, and the resulting current is amplified and digitised by comparison to a calibrated threshold. This hit information is then read out on a clock cycle according to the decision of the trigger system (see §2.6.1). The pixel detector is operated at a nominal temperature of $-10\,^\circ$C.

2.3.2 Semiconductor Tracker

The Semiconductor Tracker (SCT) is a silicon microstrip tracking detector consisting of four layers with a spatial resolution of 16 µm in r–ϕ and 580 µm in z. In the barrel region the $6.36\times 6.40\,\mathrm{cm}^2$ modules are arranged with an 80 µm pitch, overlapping so as to provide hermetic coverage for $p_\mathrm{T} > 1$ GeV tracks. Since the SCT is a strip detector rather than a pixel detector it is designed to obtain the

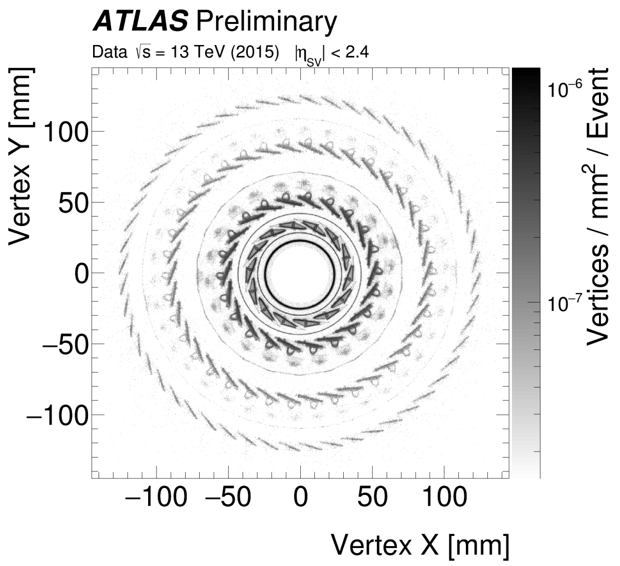

(a) Pixel detector

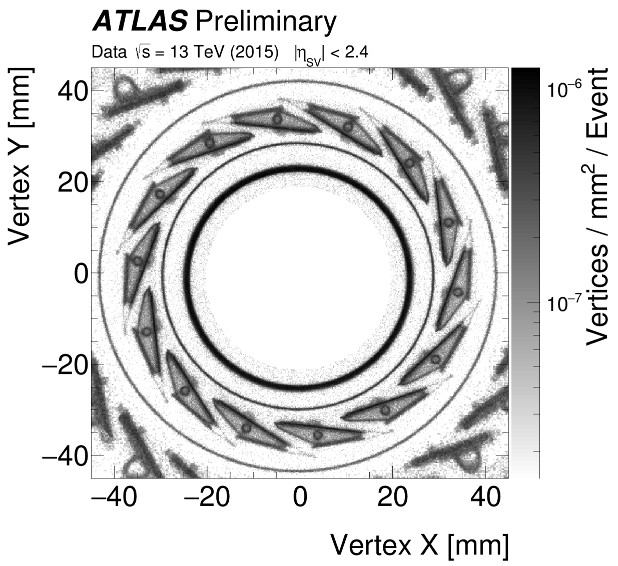

(b) IBL and beam pipe enlargement

Figure 2.8: Reconstructed secondary vertices from hadronic interactions with the detector material measured with a sample of data taken during Run 2. The density of vertices allows the position of the Pixel and IBL detector modules to be seen. The vertex positions were integrated over a large amount of collected data since each pixel layer constitutes a thickness of less than 1.39 % of a radiation length. Taken from Ref. [74].

z measurement via a small-angle stereo mechanism: each module consists of back-to-back pairs of active silicon areas with the strips of one side mounted parallel to the beam axis and the other mounted at an angle of 40 mrad. A similar design is used in the end-caps, including the small-angle stereo. The detection principle and operating conditions of the SCT are similar to the pixel detector. The SCT has 6.3 million readout channels.

2.3.3 Transition Radiation Tracker

The Transition Radiation Tracker (TRT) is a straw tube tracker with modules filled with 4 mm diameter gas-filled straws of up to 150 cm in length. The 50 000 straws in the barrel are each divided into two in the centre to allow for independent readout of each end. There are a further 250 000 straws in the TRT end-caps for a total of 350 000 readout channels. In the barrel region the straws are aligned parallel to the beam axis, while in the end-caps the straws are oriented radially (aligned in the r–ϕ plane). Approximately 36 TRT hits are recorded for each typical track.

Each straw has a 30 µm diameter gold-plated tungsten wire running through the centre and is filled with a gas mixture of xenon, oxygen and CO_2. Under the operating potential of 1600 V, the detector is able to measure the drift time of charged particles within the gas excited by the passage of energetic particles, resolving the coordinates of the hit. Sheets of polypropylene foil are placed between the straws (punched with holes to accommodate the straws), so that particles passing through the volume encounter alternating layers of foil and gas; the resulting variations in the dielectric constant cause the particles to produce transition radiation photons. The effect of this radiator material on the energy loss distribution is illustrated in Fig. 2.9. The transition radiation photons are produced at X-ray frequencies and are absorbed by the xenon gas. Since

Figure 2.9: An illustration of the typical energy loss distribution for energetic electrons in a generic transition radiation detector. Curve (b) shows a detector without gaps in the radiator meaning that the energy loss is due solely to ionisation, whereas curve (a) shows a detector with radiator gaps meaning that energy loss is due to a combination of ionisation and transition radiation. In both cases the amount of material in the radiator is the same. Taken from Ref. [150]

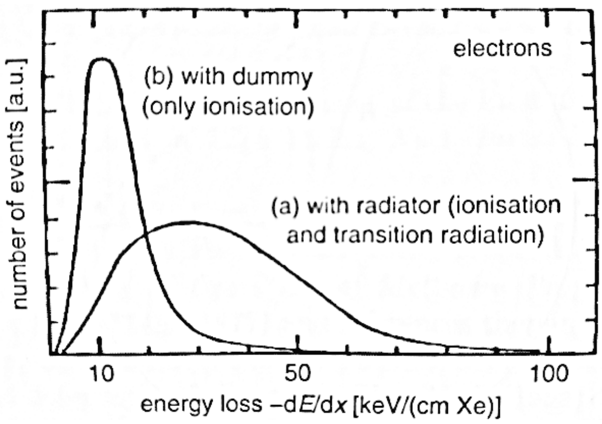

electrons produced directly in the high energy collisions are less massive than other particles passing through the TRT, and are typically energetic, they induce a larger transition radiation effect and a correspondingly larger TRT response. For this reason the TRT is operated using two thresholds with the lower one used for tracking and the higher one used for particle identification.

The ion tail in the gas mixture extends up to 60 µs meaning that the residual response in the detector long outlasts the 25 ns interval between LHC bunch crossings. Coupled with the high rate of irradiation this means that high occupancy is a significant problem for the TRT. The readout electronics must process the signal in order to mitigate the degradation of performance at high occupancy. Despite the lower spatial resolution of the TRT compared to the two silicon detectors its presence improves the overall tracking efficiency due to the additional hit information as shown in Fig. 2.10.

2.4 Calorimetry with ATLAS

The ATLAS calorimeters [16] are responsible for providing accurate measurements of the energy of particles over a range of energies from a few GeV to several TeV. They are particularly important for measuring the energy of jets (see §3.6), which are an important signature of many SM processes including tW production. They are separated into an EM layer designed to stop photons and electrons, inside a hadronic layer which is designed to stop hadrons such as pions. Two types of sampling calorimeter designs are employed: a lead-liquid argon (LAr) calorimeter in the EM layer and in the forward calorimeter with its more challenging radiation environment, and a scintillating tile calorimeter (TileCal) in the hadronic layer. The layout of the calorimeters is shown in Fig. 2.11. The relative resolution of the calorimeters improves with the energy of particles, scaling approximately as $E^{-1/2}$.

2.4.1 Electromagnetic calorimeter

The electromagnetic calorimeter (ECAL) is a lead-LAr detector made with accordion-shaped layers of lead absorbers alternating with active LAr layers. The geometry is intended to keep the depth of absorbing material relatively flat with respect to $|\eta|$.

As energetic electrons and photons enter the ECAL they initiate showers of secondary particles which propagate into the calorimeter volume. For high-energy photons the dominant energy loss mechanism is pair production of electrons and positrons, while for high-energy electrons and positrons

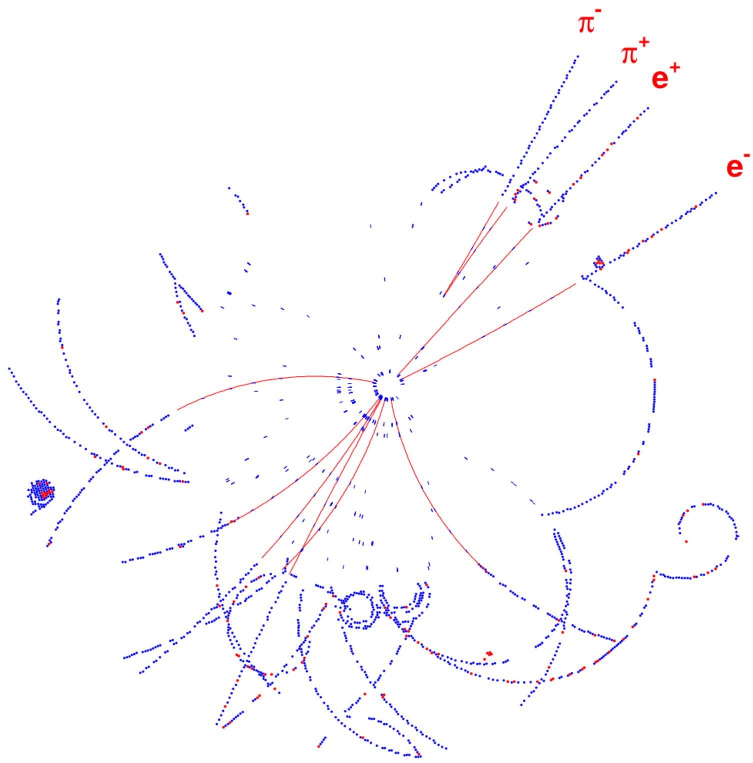

Figure 2.10: Display of an example simulated $B_d^0 \rightarrow J/\psi K_S^0$ event showing the position and granularity of hits (in blue) in the ID. The inner silicon layers provide better hit precision but the larger number of hits from the TRT assists with the track reconstruction. The red points show high-threshold (transition radiation) hits in the TRT, suggesting a higher likelihood that tracks built from these points were made by electrons. The red lines are extrapolations of the tracks between hits. Taken from Ref. [19]

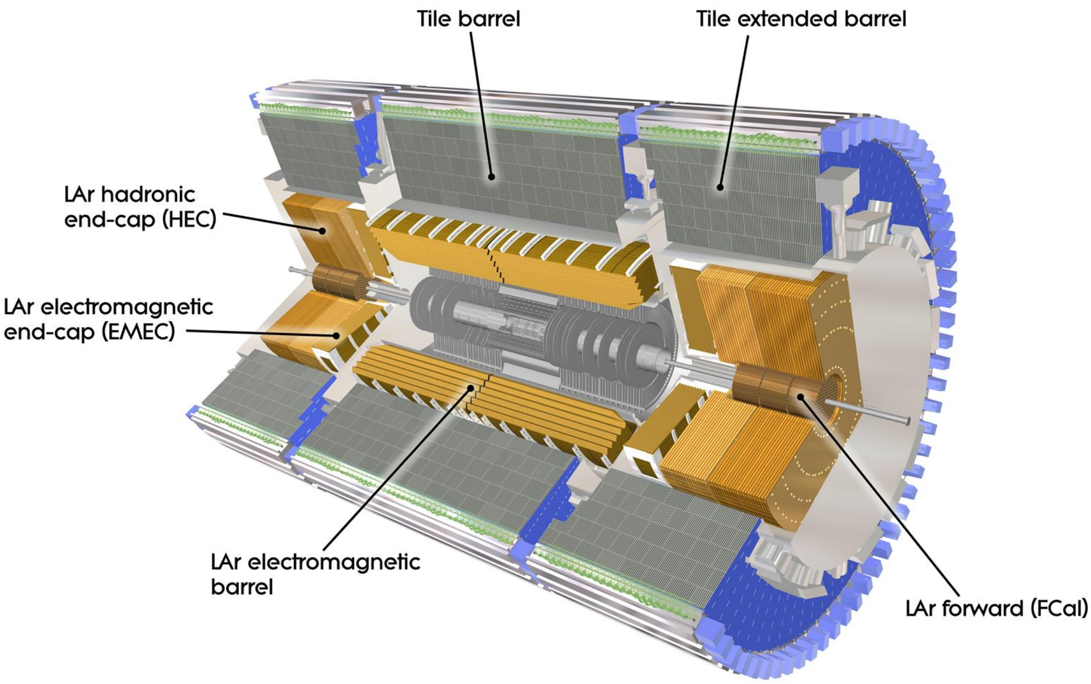

Figure 2.11: The ATLAS calorimeters. The forward LAr calorimeter and the tile barrels are part of the hadronic calorimeter. Taken from Ref. [186]

the main process is bremsstrahlung radiation. As the shower develops through these processes and the energy per particle falls into the MeV range, different energy loss processes begin to dominate. For photons, these are Compton and Rayleigh scattering and the photoelectric effect; below a critical energy, electrons and positrons primarily dissipate their energy through ionisation rather than producing new secondary particles [150, 205]. In order to absorb most of the energy of an incident photon (which is more penetrating than an electron of the same energy) an absorption depth of at least 10 to 15 radiation lengths (X_0) is required, and the ATLAS ECAL is at least 20 X_0 deep for all $|\eta|$ outside the barrel-end-cap crack region as shown in Fig. 2.12a.

The energy deposited in the calorimeters is to first order proportional to the multiplicity of particles in the EM showers, so counting the charged particles reaching an active LAr layer provides an approximation of the energy. The characteristics of the shower assist with particle identification, primarily in the longitudinal and lateral extent of its development. For this reason, and also to allow calorimeter deposits to be matched to precision tracks from the ID, the granularity of the sampling is also important. The energy resolution of the ECAL for a particle with energy E (in GeV) is approximately $10\,\%/\sqrt{E} \oplus 0.4\,\% \oplus 0.3/E$ where the sum is in quadrature. The first term accounts for statistical fluctuations in the shower development and sampling, the constant term covers non-uniformity in the detector and uncertainty in the calibration, and the third term is related to electronic noise in the readout system.

The barrel region covers $|\eta| < 1.475$ and the end-cap covers $1.375 < |\eta| < 3.2$. According to the physics goals of ATLAS the region $|\eta| < 2.5$ is the precision region, and so in this region the ECAL performs three longitudinal samplings whereas in the forward regions only two samplings are performed and at a coarser granularity. In the precision region, the first sampling has a granularity of $\Delta\eta \times \Delta\phi = 0.003 \times 0.1$ to provide precise η location, the second is coarser at 0.025×0.025 and the third is coarser again at 0.05×0.025. This granularity is shown in Fig. 2.13. A presampler LAr layer installed directly inside the ECAL cryostat wall is used to provide an additional measurement in the range $|\eta| < 1.8$ in order to correct for energy lost before entering the calorimeter; this corresponds to the $|\eta|$ range that accounts for the greatest depth of material (in radiation lengths) in front of the calorimeter, as seen in Fig. 2.7.

2.4.2 Hadronic calorimeters

Since in most materials the nuclear interaction length, λ_I, is significantly longer than the radiation length[9], hadrons tend to penetrate the ECAL and interact with material behind it. The ATLAS calorimetry system has a total depth of at least 10 λ_I across its $|\eta|$ coverage as shown in Fig. 2.12b. A large number of inelastic hadronic interactions occur, in addition to EM interactions, to induce a shower consisting primarily of pions (distributed in roughly equal fractions of the three pion species π^0, π^+ and π^-), but also to a much smaller extent neutrons, kaons, and leptons (including both charged leptons and neutrinos). Roughly one third of the secondary particles produced are neutral hadrons which decay into photon pairs that in turn initiate their own EM subshowers. Charged hadrons lose energy through ionisation but also through interactions with nuclei, ejecting nucleons

[9] For elements with $Z > 23$ the radiation length is proportional to $Z^{-2/3}$ while for elements with $Z > 15$ the nuclear interaction length is approximately proportional to $A^{1/3}$. The ratios of these quantities to the density of the material, X_0/ρ and λ_I/ρ, have units of physical length; calorimeter design must consider these together with the availability and cost of materials [205, §34.9].

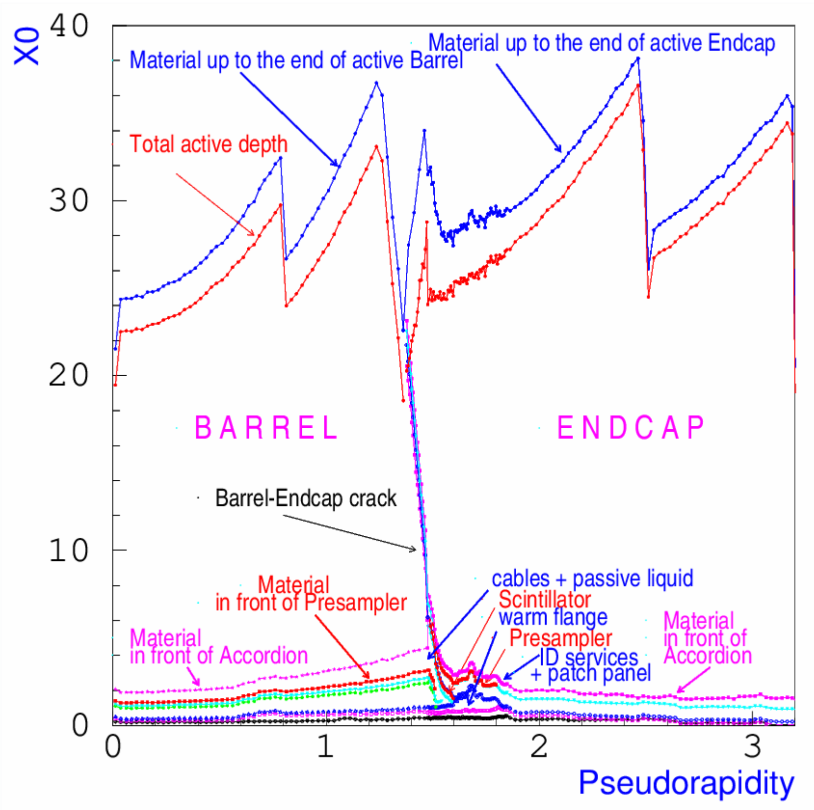

(a) ECAL

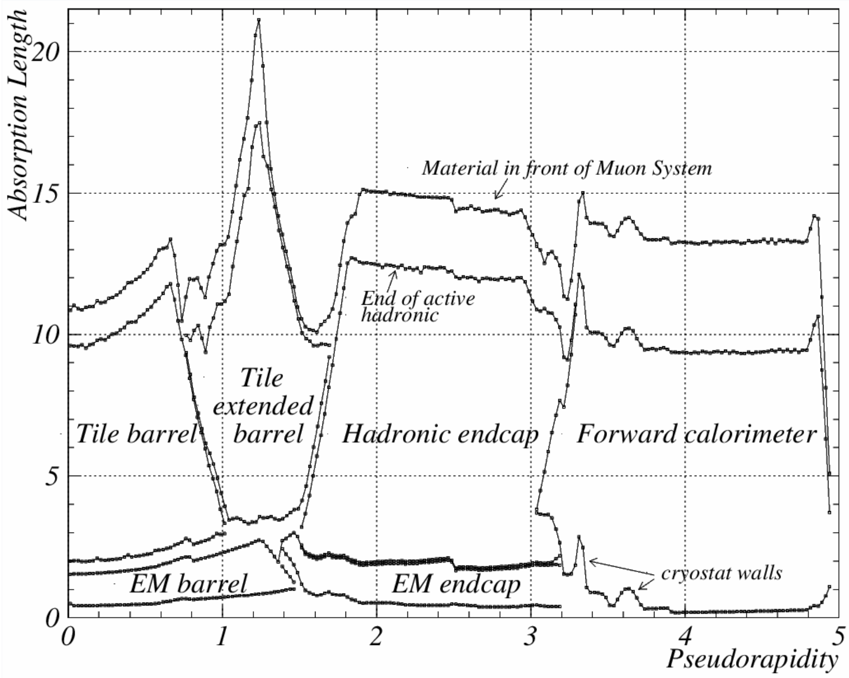

(b) Total calorimetry system

Figure 2.12: Amount of material in the ATLAS calorimeters as a function of $|\eta|$, for (a) the ECAL measured in radiation lengths and (b) the entire calorimetry system measured in interaction lengths. Both figures taken from Ref. [16]

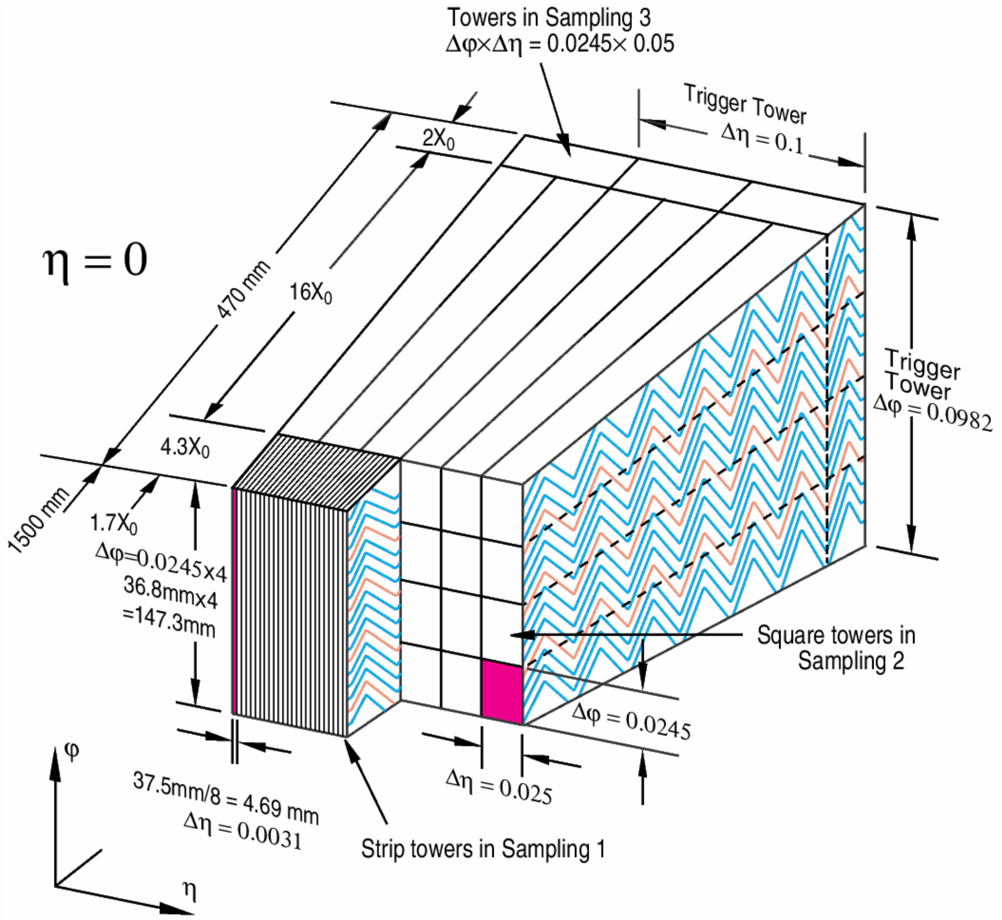

Figure 2.13: Granularity of the ECAL readout system. The first sampling is highly granular in η due to the strip design, whereas the second and third samplings are performed in coarser grids in $\Delta\eta \times \Delta\phi$. The coloured lines illustrate the accordion geometry of the lead absorber and LAr layers. Taken from Ref. [16]

and exciting nuclei which subsequently emit photons. The energy spent to overcome the nuclear binding energy and associated with undetectable nuclear recoil or late capture of neutrons accounts for approximately one third of the total energy in the hadronic shower; this fraction is either undetectable or delayed on the order of several µs. Compared to pure EM showers these hadronic showers develop further laterally as well as longitudinally. They are also highly variable, both because there are fewer hard particles over which fluctuations in the detection efficiency can be averaged out, and because of the invisible fraction. Event-by-event fluctuations in detection efficiency and the EM fraction of the shower combined with sampling fluctuations and leakage degrade the resolution of hadronic calorimeters[10] with respect to EM calorimeters [150, 205].

In the forward region, $1.5 < |\eta| < 4.9$, two hadronic LAr calorimeters are employed based on the same operating principle as the EM LAr calorimeter but with a higher density. The hadronic end-cap calorimeter extends from $1.5 < |\eta| < 3.2$ and uses copper plates of 25 mm to 50 mm thickness as absorbers. The forward calorimeter sits concentrically inside the hadronic end-cap covering the $3.1 < |\eta| < 4.9$ range and consists of three modules of which the first is properly an ECAL with copper absorbers and the latter two are hadronic calorimeters (HCALs) with tungsten absorbers. Taken together, these forward systems facilitate the calorimetry of jets that are boosted with respect to the centre-of-mass frame. They also permit a reduction in the tails of the $E_{\mathrm{T}}^{\mathrm{miss}}$ distribution (see §3.7) by increasing the solid angle coverage of the detector.

The bulk of the HCAL material is in the TileCal, an alternative design of sampling calorimeter in which iron absorbers alternate with polystyrene scintillating tiles that are read out through fibres into photomultiplier tubes (PMTs) to measure three samplings. As shown in Fig. 2.11 the TileCal consists of a barrel component and two extended barrel sections of the same radius; there is also a small intermediate section in the region $0.8 < |\eta| < 1.0$ designed to fill as much of the gap between the other sections as possible while accommodating the necessary services. The TileCal extends from $r = 2.28$ m to 4.25 m and is sufficiently rigid so as to support itself, the other calorimeters, and the ID so that there is no mechanical contact with the muon system which surrounds it. Since the iron mass in the TileCal acts as the solenoid return yoke and is also affected by the toroid, there is a magnetic field in its volume varying greatly in strength across its geometry. This must be accounted for both in the placement of sensitive components and in the simulation of shower development.

The readout granularity is $\Delta\eta \times \Delta\phi \simeq 0.1 \times 0.1$ (reduced to $\sim 0.2 \times 0.1$ in the outermost layer, noting that the η segmentation is only approximately projective), chosen to facilitate handling of shower leakage from the ECAL and to permit some compensation for the nonlinearity in response associated with the invisible energy fraction in hadronic showers. The readout is sufficiently fast to be used to make trigger decisions and to provide a real time estimate of the total luminosity.

2.5 The ATLAS Muon Spectrometer

Energetic muons penetrate the inner layers of the detector losing comparatively little of their energy in the calorimeters. In order to exploit their striking signature ATLAS has a dedicated set of detectors called the Muon Spectrometer (MS) [23] built around the core ID and calorimeters and immersed in the 0.5 T to 2 T magnetic field of the toroid magnet. Precision muon tracking is provided by

[10] The ATLAS hadronic calorimeters were designed to achieve an energy resolution for jets with energy E (in GeV) of $50\,\%/\sqrt{E} \oplus 3\,\%$ in the central region and $100\,\%/\sqrt{E} \oplus 10\,\%$ in the forward region [77].

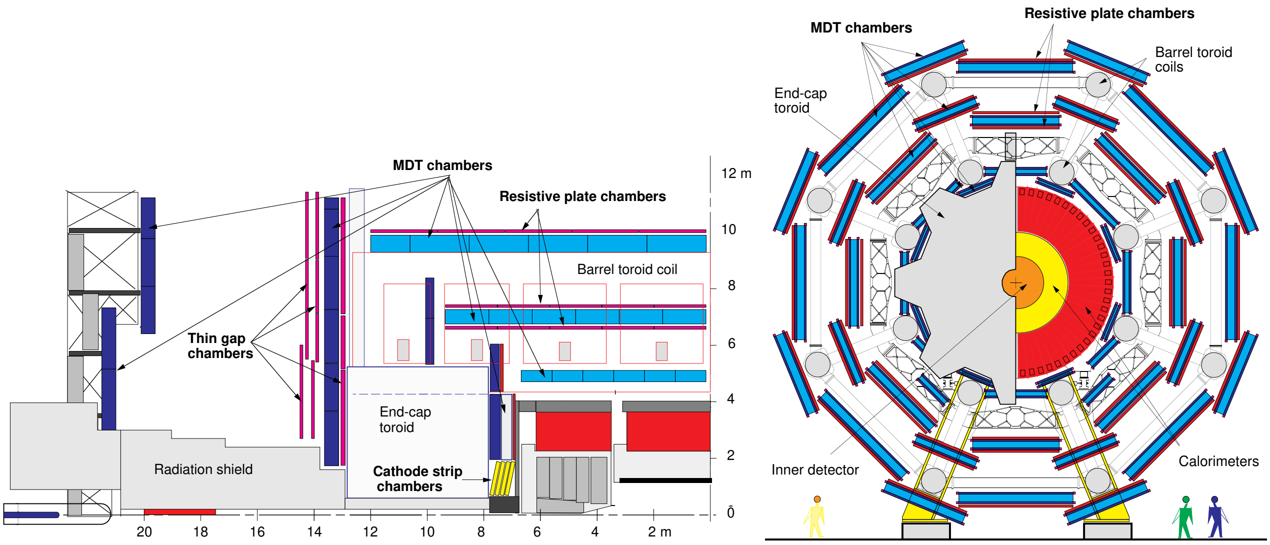

(a) r–z section **(b)** x–y section

Figure 2.14: Schematic of the MS in its Run 1 configuration shown in (a) r–z and (b) x–y sections. In Run 2 additional MDTs and RPCs were added to improve coverage, particularly in the region near the feet of the ATLAS detector. Taken from Ref. [23].

Monitored Drift-Tubes (MDTs) and in the forward region Cathode Strip Chambers (CSCs). Muon trigger information and supplementary tracking is provided in the barrel region by Resistive Plate Chambers (RPCs) and in the end-caps by Thin Gap Chambers (TGCs). The layout of the four muon detectors is shown in Fig. 2.14: three concentric barrel layers at $r = 5\,\text{m}$, $7.5\,\text{m}$, and $10\,\text{m}$ cover the region $|\eta| < 1$ and four end-cap disks at $z = 7\,\text{m}$, $10\,\text{m}$, $14\,\text{m}$, and $21\,\text{m}$ cover the forward range $1 < |\eta| < 2.7$ on each side. Due to the large volume of the MS and the necessary supporting structures and services several variations to the designs of each type of detector are implemented. This is necessary to accommodate geometrical constraints while maintaining adequate overlap of active areas to ensure hermeticity to muon tracks with a p_{T} of 20 GeV across the full pseudorapidity coverage. The momentum measurement of muons is based on the deflection of their trajectories in the large volume of the toroid field rather than on calorimetry. The MS resolution degrades with the particle momentum, scaling approximately as p^{-1}.

2.5.1 Monitored Drift-Tubes

The MDTs are 30 mm diameter[11] aluminium tubes of various lengths filled with a gas mixture consisting primarily of argon with smaller amounts of CH_4 and N_2, and a central 50 µm diameter wire. The tubes in each of the three chambers are stacked in three layers (increased to four for the innermost chamber) on either side of a spacer frame for a total of 20 individual tube layers along a track trajectory. In the barrel the tubes are positioned perpendicular to the beam axis and in the end-caps they are aligned azimuthally in 16 wedges; the tubes are therefore approximately parallel to the magnetic field lines.

Each MDT chamber typically provides 6 to 8 η measurements corresponding to the tubes hit in the bending plane of the muon. When combined with the drift time information, this results

[11] Some of the additional MDT chambers installed in preparation for Run 2 use tubes with a smaller radius in order to improve their performance inside a higher muon flux.

in an effective measurement of a single coordinate to 40 µm precision along with a single angle of 3×10^{-4} rad precision.

2.5.2 Cathode Strip Chambers

The CSCs provide finer granularity tracking in the challenging forward environment at $2 < |\eta| < 2.7$ within a small physical extent. They are multiwire proportional chambers placed in two coaxial rings of 8 chambers in the inner 1 m of the innermost muon end-cap on each side of the detector. They are segmented azimuthally to match the MDTs occupying the remainder of the same end-cap. The CSCs have anode wires spaced at 2.54 mm intervals inside a chamber filled with a gas mixture of CO_2, argon, and CF_4. The avalanche induced by ionisation by passing muons on the anodes creates a charge on the 1 mm to 2 mm strip cathodes on the internal wall of the chambers running perpendicular to the anode wires.

The distribution of charge over a number of cathode strips then allows the determination of the hit coordinate transverse to the cathode strips (i.e. in the bending plane) with a resolution of 60 µm due to charge interpolation between neighbouring strips, and the longitudinal anodes are read out with a 5 mm resolution. In order to reduce the impact of resolution lost due to the inclination of tracks the chambers are mounted with a tilt so that their active surface is normal to rays extended from the IP, as shown in Fig. 2.14a.

2.5.3 Resistive Plate Chambers

The RPCs provide muon trigger information in the muon barrel region $|\eta| < 1.05$[12]. While the spatial resolution (5 mm to 10 mm) is much coarser than that from the MDTs, it has a timing precision of 1 ns, shorter than the LHC bunch crossing interval of 25 ns. They are mounted in three doublet RPC layers: one on each side of the middle barrel MDT layer and the third on one side of the outermost RPC layer.

Each RPC consists of a parallel pair of resistive Bakelite plates held apart by resistive spacers opening a 2 mm gap filled with tetrafluoroethane ($C_2H_2F_4$). The uniform electric field between the plates causes each incident ionisation electron to induce a charge avalanche that propagates to the far plate. The accumulated charge is read out by two orthogonal series of strips to measure coarse hit information and allow an approximate momentum measurement. A coincidence between the first RPC layer and one of the other layers is required in order to produce a trigger signal.

2.5.4 Thin Gap Chambers

The TGCs provide muon trigger information in the muon end-cap region from $1 < |\eta| < 2.4$. They are also able to provide azimuthal measurements out to $|\eta| < 2.7$ to complement the precise η measurements provided by the MDTs. One doublet and two triplet layers of TGCs on the middle MDT layer provide azimuthal and triggering functionality, and two layers of TGCs on the inner MDT layer provide a supplementary azimuthal measurement without the triggering functionality.

The TGCs are multiwire proportional chambers filled with a 55 % CO_2 and 45 % n-pentane (n-C_5H_{12}) gas mixture using anode wires separated by 1.8 mm, larger than the anode-cathode distance of 1.4 mm. Anode wires are bunched into groups of 4 to 12 wires for readout when used as a trigger, with

[12] Additional RPCs were installed on the MDTs added to the barrel-end-cap transition region in advance of Run 2.

the group size varied over the geometry of the detector to meet the spatial resolution requirements at different η. The anode wires are aligned parallel to the MDT tubes, and the cathodes are read out using strips perpendicular to this in order to provide the ϕ measurement.

2.6 Data acquisition and triggers at ATLAS

At the high nominal luminosity of the LHC and its bunch crossing rate of 40 MHz the proton interaction rate is on the order of 1 GHz. The computational resources required to transfer, process and store this amount of data are not technically available. In order to reduce the data-taking rate to a manageable level ATLAS employs a data acquisition and trigger system that makes rapid online decisions about whether or not to read out, reconstruct and store an event for offline analysis. In Run 2 of the LHC ATLAS used a two-stage trigger system with the Level 1 (L1) trigger implemented at low level in hardware and the High Level Trigger (HLT) running in a commodity processor farm.

The ATLAS triggers select events based on a trigger menu informed by the physics goals of the collaboration, typically requiring energetic particles that are common signatures of many SM and BSM processes. In order to avoid a systematic rejection of events that while not predicted are nonetheless interesting, and also for use in calibrations, there are trigger items that accept events randomly to produce a small unbiased sample. The exact composition of the trigger menu is varied to adapt to the luminosity delivered by the LHC so that as much data is recorded as can be permitted by the online operational constraints. Compared to Run 1, the increased beam energies in Run 2 resulted in an event rate 2 to 2.5 times higher for the same luminosity (see §1.5 and Fig. 1.8), so the Run 1 trigger systems would not have coped in the same configuration [64]. Almost all trigger systems were upgraded for Run 2 to accommodate the increased beam energy and the larger peak luminosity.

2.6.1 Level 1 trigger

The L1 trigger must reduce the event rate from 40 MHz to 100 kHz with a latency of at most 2.5 µs. During this window, corresponding to 100 LHC bunch crossings, the detector channels are not read out, and so each channel has a hardware signal pipeline that buffers its response until a trigger accept signal causes it to be read out for further processing. The L1 trigger decisions are made based on information with reduced granularity from the muon and calorimeter systems alone.

The L1 calorimeter triggers use input from both the tile and LAr calorimeters to deduce the presence of high-p_{T} electrons, photons, jets, and tau leptons, or a high total energy or $E_{\mathrm{T}}^{\mathrm{miss}}$ in the event. The calorimeters are read at this stage as projective towers with a granularity of approximately $\Delta\eta \times \Delta\phi = 0.1 \times 0.1$ as shown for the ECAL in Fig. 2.13, meaning that each tower captures approximately 7200 analog input channels, adding them together in hardware. Each tower has configurable p_{T} thresholds, and information from all the towers is used together to determine $E_{\mathrm{T}}^{\mathrm{miss}}$ and the scalar E_{T} sum.

The L1 muon triggers use the RPC and TGC both separately and together in order to identify muon candidates passing configurable p_{T} thresholds. L1 topological triggers were introduced in 2016, combining the muon and calorimeter information to allow for more complicated selections based on quantities such as the angular separation of two candidate objects.

The trigger subsystems pass a summary of the gathered information to a central processor that makes the final L1 trigger decision. To avoid overwhelming the data-taking pipeline the L1 processor imposes certain periods of dead time during which no trigger accept is possible. Dead time can occur due to any part of ATLAS (including the HLT) signalling that it is unable to process new events. There are also two kinds of mandatory dead time: simple dead time blocks the trigger for the five bunch crossings following an L1 accept signal, and complex dead time blocks the trigger when there have been more than a certain number of L1 accept signals in the preceding bunch crossings. The dead time is closely monitored during data-taking to ensure that the detector is successfully reading out data at the expected rate.

Certain desirable trigger items, for example items selecting events containing at least one electron or muon with a relatively low p_{T} threshold, result in an unacceptably high trigger rate. As a compromise such triggers can be used if they are *prescaled*, meaning that all but every nth accept signal will be ignored by the processor. The prescale setting is recorded so that the resulting sample can be correctly interpreted.

2.6.2 High Level Trigger

The HLT system[13] must reduce the 100 kHz L1 rate to a final event rate of approximately 1 kHz. The complete detector readout at full granularity from a bunch crossing event along with the details of the L1 decision are fed to a farm of commodity computers for reconstruction and processing. The algorithms used for the reconstruction are similar to those discussed in Chapter 3 with certain elements replaced by faster approximations according to the latency requirements of the HLT. Due to the availability of detailed physics information the HLT items can be based on correspondingly more complicated decisions. Once an event is accepted by the HLT it is recorded to disk where it is held for further processing.

The prescale factors for certain trigger items are adjusted over the course of a fill in order to keep the HLT farm busy without causing its buffers to fill to the point where it must discard data to keep up. Figure 2.15 shows the rate of each group of triggers over the course of a single fill from July 2016. The changes in prescale factors are visible as steps in the rates, while temporary dips are caused by dead time and small spikes are cause by noise bursts in the detector. The single muon and electron triggers used in the analyses in this thesis are shown in bright red (the top of the stack) and bright green (the third from the top) respectively.

2.6.3 Calibration and processing

As events are recorded, they are reprocessed using the full offline reconstruction algorithms described in Chapter 3. Over the two days following any fill of the LHC the corresponding dataset is placed in a calibration loop. During this period it is incrementally reprocessed and analysed to correct for features such as the relative spatial alignment of parts of ATLAS. The average position and size of the luminous region of the LHC beams is reconstructed and used as a constraint in algorithms that extrapolate tracks to vertices, the points from which they originated.

Figure 2.16 illustrates how a relative misalignment in precision tracking layers can result in a systematic bias in the position of tracks and consequently of vertices. Figure 2.17 shows the evolution in the alignment correction to the horizontal position of the pixel detector over the course of a year, with movement on the scale of 1 µm. Larger movements are observed in the longitudinal direction,

[13] In Run 1 an intermediate software trigger layer existed, but this was merged into a single HLT system for Run 2.

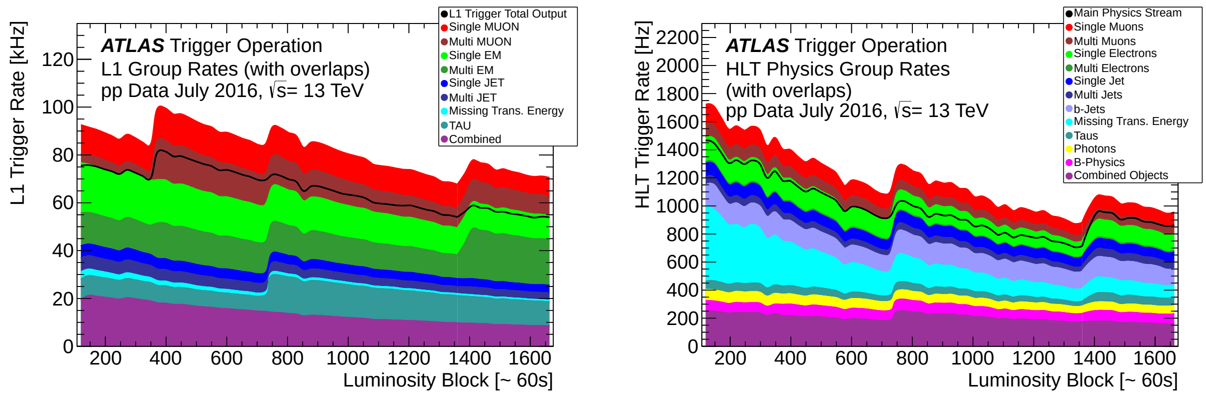

(a) L1 **(b)** HLT

Figure 2.15: Trigger rates over a single fill from July 2016 with a peak instantaneous luminosity of $1.02 \times 10^{34}\,\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$. Each luminosity block corresponds to about 60 s of data-taking. The solid black lines show the total rates, while the stacked histograms show the rates of each type of trigger. An event can be triggered by multiple types of triggers. Taken from Ref. [18].

Figure 2.16: Effect of alignment in the IBL and pixel detector layers. The dashed line represents the track reconstructed using the hits (represented as black circles) in each layer. Due to an (exaggerated) misalignment, the middle pixel layer is actually at the position of the red line, and the hit in that layer should be corrected to be at the position of the red circle. The true track is shown by the solid black line.

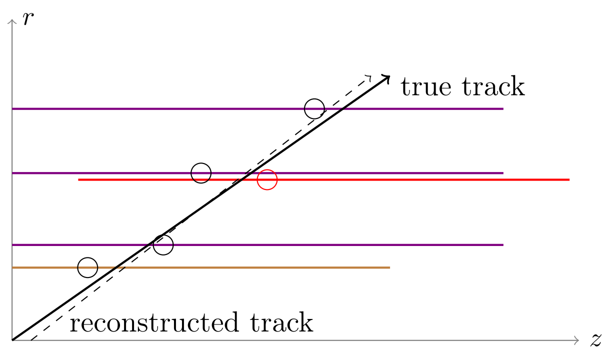

as well as twisting distortions in the IBL due to temperature-dependent mechanical effects, and much larger vertical movements of the pixel detector whenever it is switched on from standby [66]. The complete alignment fit involves tens of thousands of degrees of freedom corresponding to mechanical degrees of freedom in detector elements. It is performed iteratively to accommodate the interdependence of the inputs and outputs to the alignment process.

After several iterations the calibration is finalised and the data are inspected alongside information about the status of the detector to flag events that are not suitable for physics analysis due to features such as noise bursts in electronics. The final dataset is then distributed to data storage facilities across the globe as part of the Worldwide LHC Computing Grid [131, 93] where it can be analysed.

2.7 Luminosity determination with ATLAS

An accurate determination of the instantaneous luminosity is essential for operation of both the LHC and ATLAS, and the uncertainty in the determination of the integrated luminosity enters into all cross-section measurements performed with ATLAS. ATLAS has a number of luminometers, of which the primary detector in Run 2 was LUCID [82], a Cherenkov detector mounted on the beam pipe 17 m from the interaction point on either side. Cherenkov photons are produced in quartz windows inside LUCID and are picked up by PMTs and translated into hits when one of the PMTs produces a response above a configured threshold.

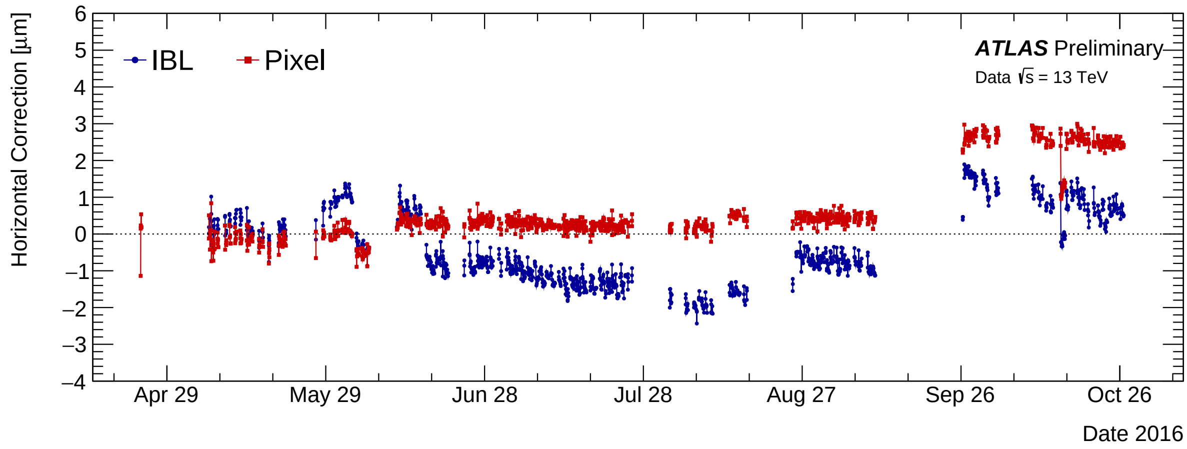

Figure 2.17: The global alignment correction applied to the horizontal (x) position of the IBL and pixel detectors over 2016. This corresponds to the physical movement of the detectors with respect to their baseline alignment, using the SCT barrel as a fixed frame of reference. Note that further alignment corrections are applied at the level of individual IBL and pixel modules. Taken from Ref. [35].

Additional luminosity measurements are performed during operation, including a calorimeter-based counter implemented in the TileCal, a charged track counter in the pixel detector, and other dedicated luminometers such as the Beam Conditions Monitor, a diamond sensor mounted 1.84 m from the interaction point. These luminosity measurements require periodic calibration, for example the TileCal's response is calibrated using a radioactive caesium source that is circulated between the calorimeter cells.

The absolute luminosity calibration is performed by van der Meer (VdM) scans. During VdM scans the LHC sweeps through a number of steps in beam separation in the x and y directions, effectively reducing the luminosity by reducing the beam overlap and so also the size of the luminous region, called the *beam spot*. For the horizontal and vertical convoluted beam size parameters, Σ_x and Σ_y respectively, the bunch luminosity can be written as

$$\mathcal{L}_b = \frac{\mu_{\text{vis}} f_r}{\sigma_{\text{vis}}} = \frac{f_r n_1 n_2}{2\pi \Sigma_x \Sigma_y},$$

where μ_{vis} and σ_{vis} are the visible values of μ and σ_{inel} (i.e. the observable equivalent of Eq. (2.1)), and n_1 and n_2 are the bunch populations. Assuming the beam spot is a three-dimensional Gaussian profile, Σ_x and Σ_y are interpreted as the standard deviation of that profile in the x and y directions. The product $\Sigma_x \Sigma_y$ can be determined by comparing measurements taken at different separation points in a VdM scan [38].

The calibrated luminosity collected in each year of Run 2 is shown in Fig. 2.18 as a function of the average number of interactions per bunch crossing, $\langle\mu\rangle$. The increase in $\langle\mu\rangle$ each year until 2018 is a result of the changes to the LHC injection schemes and fill configuration (see §2.1.3). In 2017, the filling scheme used for part of the year (as a mitigation for the beam loss issues) resulted in a peak instantaneous luminosity at the start of each fill that exceeded the capacity of the detector and the HLT. This was managed by levelling the beam to $\langle\mu\rangle = 59$ at IP1 until the instantaneous luminosity fell below that threshold. This is the reason for the second peak in the 2017 distribution. The small

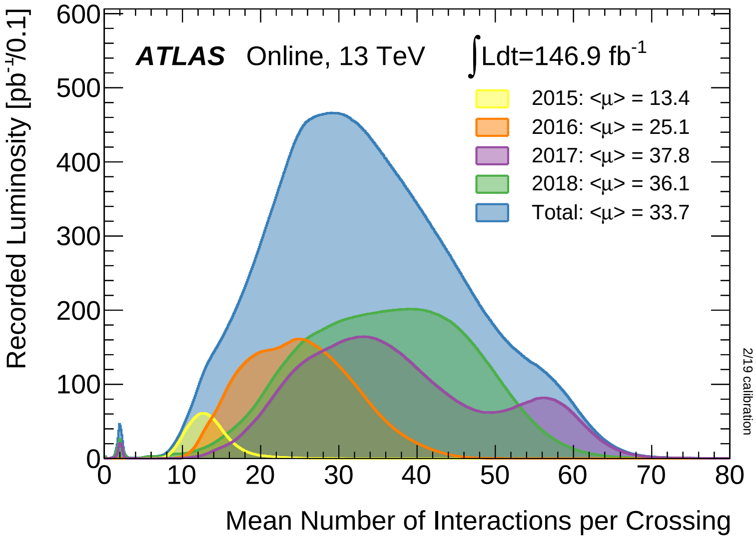

Figure 2.18: Pileup distribution of pp collisions delivered by the LHC to ATLAS during Run 2. Taken from Ref. [17].

peaks in the very low $\langle\mu\rangle$ region for 2017 and 2018 correspond to special runs of the LHC designed to achieve low pileup conditions necessary for some specialised measurements.

Chapter 3

Reconstruction and selection of data

Reconstruction is the process of transforming and interpreting event data for analysis. Reconstruction algorithms build up lists of candidate *physics objects* (electrons, muons, jets, etc.) based on the response of the detector and apply calibrations to correct for various detector effects. Simulated events can also be appropriately processed and digitised so that they can be reconstructed by the same algorithms that are used on observed events. Since the reconstruction algorithms are tuned and validated based on data as they arrive, the specific details of some of the algorithms changes over time as particular datasets are studied and understood.

The aim of the reconstruction algorithms is to be able to identify *prompt* particles (those originating from the hard process as opposed to secondary interactions) with well-understood *efficiency* and *purity* characteristics. Efficiency measures the proportion of a given type of object that is successfully reconstructed, while purity measures the proportion of reconstructed objects that is reconstructed correctly. The algorithms can reconstruct a completely different object that has mimicked the signature of their target object (for example a jet reconstructed as an electron) in which case the candidate is called *fake*. A compromise between high purity and high efficiency must always be reached, according to the requirements of a given analysis, and so to be as broadly applicable as possible these algorithms are calibrated for a range of working points and settings. The performance of the algorithms is checked in MC simulation, but it is also necessary to assess their performance using real data. This is usually done by isolating a sample of events where the property in question can be measured by a second, independent method, for example calibrating jet momentum by using constraints from momentum conservation in events where a jet is balanced against an object with known momentum (see §3.6.1).

This chapter introduces the techniques used to reconstruct the physics objects that will be used for the analyses in Chapters 5 and 6. Study of the tW system requires good reconstruction of its decay products, which include electrons (see §3.4) and muons (§3.5) from leptonic W decays, jets (§3.6) from the top quark decay, and missing transverse momentum (§3.7) from which the presence of neutrinos can be inferred. For each type of object, several reconstruction techniques are used within the ATLAS experiment ranging from general purpose to highly specialised in terms of their intended applications. The discussion in this chapter primarily covers the techniques used directly in this thesis with an emphasis on the relevant settings and calibrations, and on details that are necessary for analysis of the measurement uncertainty.

3.1 Simulation of events

For predictions to be directly comparable with observation either the observation must be corrected to account for effects such as the finite resolution of physical detectors, or the predictions must be corrected to simulate such effects. Even in cases where the former strategy is adopted for the interpretation of results, fully reconstructed simulated events are used extensively for calibration, correction, and validation throughout the reconstruction and analysis framework. This section therefore describes the preparation of simulated events to the level of detector response such that they can be reconstructed using the same algorithms as observed data.

After matrix element generation (see §1.3.1) events are said to be at *parton level*, containing descriptions only of the hard processes. Event generators typically assign MC weights to each event so that a finite set of such events can be used to efficiently sample the full kinematic phase space of the physical process. The parton level events are then processed by a parton shower generator and decayed to stable particles so that the events are said to be at the *particle level*.

The interaction of particles with the ATLAS detector is simulated using the GEANT4 [3] toolkit, which models a comprehensive set of physics processes over energies ranging from a few hundred eV up to the TeV scale. The toolkit is configured using a detailed model of the detector geometry and material and the magnetic field in which it is immersed, and simulates the trajectory and interactions of each type of particle. The response of individual detector elements is then simulated by digitising the hits from GEANT4 into voltages and other readout signals. This takes into account any damaged modules or elements, as well as channel-dependent variations in the responses. The effect of pileup is modelled by simulating additional collisions that are then overlaid onto the event along with samples of electronic noise and ambient background conditions. Note that the accuracy of these descriptions depends on knowing both the condition of the detector and the operation conditions of the LHC (perhaps most importantly the pileup distribution) and yet simulated samples typically need to be prepared in advance of data-taking for efficient utilisation of computational resources. This is handled by preparing dedicated simulations for each specific dataset, and by reweighting prepared simulated events to match the actual conditions.

In terms of computational resources, the most expensive element of this simulation pipeline is the simulation of interactions with the calorimeters. For many purposes it is sufficient to parametrise the response of the calorimeters to produce an approximately correct simulation, using a technique called the ATLAS fast calorimeter simulation (ATLFASTII) [78] which is approximately 20 times faster than the full GEANT4 modelling of the calorimeters.

3.2 Data quality

All analysis of ATLAS data begins by selecting a subset of recorded data that conforms to certain quality requirements. During the calibration and processing of data as described in §2.6.3 data are labelled according to the purpose (e.g. calibration or physics analysis) of the running period in which they are collected, and the reliability of the detector during that period. This is done at the granularity of luminosity blocks, which typically correspond to 1 min of delivered collisions so that parts of a data-taking run affected by transient issues with the detector or data-taking can be isolated without affecting the remainder of the run. Knowledge of operational defects affecting ranges of luminosity blocks is encoded into a database from which *good run lists* are prepared. For some purposes it is preferable to maximise the available sample and tolerate or compensate for even severe defects. The general-purpose good run lists for physics analysis, however, require that all parts of the detector are operating well, and these are the lists used to select datasets in this thesis. Occasionally datasets are reprocessed using updated information that can compensate for certain issues and increase the total integrated luminosity usable from the same running periods.

Further quality requirements are imposed at the level of individual events. Any event affected by a noise burst or data corruption in the calorimeters or SCT in particular are excluded, as well as any with other rare issues such as incomplete information in the events immediately following the online recovery of a subdetector. Individual physics objects in events can be flagged as poorly measured based on details of specific reconstruction algorithms or the presence of certain types of backgrounds in the event [26]; analyses can then choose to exclude events containing such objects.

3.3 Common objects used in reconstruction

The response of the tracking detectors is used to reconstruct *track* and *vertex* objects, and the response of the calorimeters is used to construct topo-clusters, as described in §3.3.3. These common objects are then used in other algorithms to reconstruct physics objects for analysis.

3.3.1 Tracks

Tracks are reconstructed using the ID tracking layers, with extensions in the MS handled separately. First the two-dimensional hit information from the silicon detectors must be transformed into three-dimensional positions called space points. For pixel detector hits this is straightforward, but for SCT hits the small-angle stereo must be exploited together with an additional constraint from the position of the beam spot (see §2.7). In the pixel detector, clustered hits are resolved into individual hits using a neural network [14, 63].

The trajectory of charged particles is modelled as a helix, which is true to the extent that the magnetic field is uniform. Sets of three space points are used to seed track candidates, and then a Kalman filter method [146] is used to follow the road of possible trajectories allowed by the seed points. The method works by predicting the position of the hit on the next surface (moving outwards from the IP) and updating its fit parameters as it includes successive hits. This approach allows outliers to be identified by their large contribution to the χ^2 of the fit and thereby excluded. The track candidates are then checked more carefully using detailed information about the material of the detector, and ambiguities are resolved by preferring higher-quality track candidates. A search is performed for compatible extensions of the track into the TRT volume using its lower precision hits, but these extensions do not modify the reconstructed precision tracks. The full algorithm is described in Ref. [120].

The *transverse impact parameter* d_0 is defined to be the distance to the *perigee*—the closest approach of a track to the beam axis—in the transverse plane. The *longitudinal impact parameter* $z_0 \sin\theta$ is the distance to the perigee in the longitudinal plane in which the magnetic field is approximately parallel (near the origin) to the z-axis. These definitions are illustrated in Fig. 3.1. By convention d_0 is assigned a positive sign if the track has positive angular momentum around the origin in the x–y plane. The impact parameter resolution varies as a function of both η and p_{T}, with

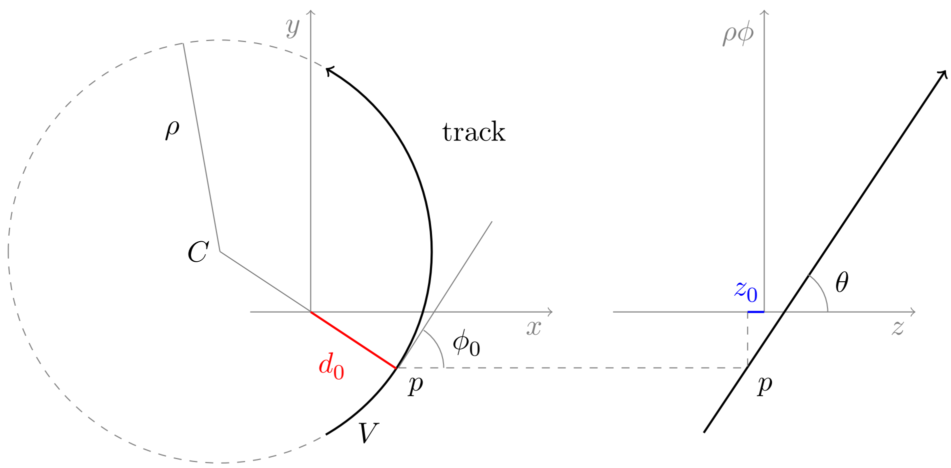

Figure 3.1: Definition of the track impact parameters. In the transverse plane (left), the particle track is parametrised as a helix with centre C and radius ρ. It originates from the vertex, V. The closest approach of the track to the z-axis (i.e. the beam axis) is the perigee, p. The distance from the coordinate origin to p is the transverse impact parameter d_0. In the plane longitudinal to the track (right), the z coordinate of the perigee is z_0, and the track forms an angle θ with the z-axis. In the case illustrated, $d_0 > 0$ and the particle has negative electric charge.

the best resolution for high-p_T tracks in the central region of the detector. For d_0 the resolution ranges from about 10 µm to 140 µm and the z_0 resolution ranges from about 90 µm to 400 µm as measured in a sample of dijet events [36], with an improvement of up to 40 % in both parameters compared to Run 1 in soft tracks. The impact parameter significance is the impact parameter scaled by its resolution.

The track reconstruction efficiency is relatively flat with respect to track p_T, and falls from about 85 % in the central region to about 60 % in the forward region for the tightest track selection shown in Fig. 3.2. The tight track selection is distinguished from the loose selection by the requirement that tracks have at least 9 hits in the silicon tracking layers of which one is in the innermost two layers, and that they have no holes (tracking surfaces with no hits between others with hits) in the pixel layers. The rate of fake tracks—reconstructed tracks that do not correspond to charged particle trajectories—is estimated to be negligible for tight tracks, while rising slowly as a function of pileup in loose tracks beyond about $\mu = 15$ [28].

3.3.2 Vertices

Once a set of tracks is reconstructed in an event they are extrapolated towards the IP and a number of vertices are reconstructed using a χ^2 fit, in order to provide information about the ancestry of particles associated with those tracks. Vertices inside the beam spot compatible with proton collisions are called *primary vertices*, while vertices that are displaced from the beam spot are called *secondary vertices* and are generally the result of secondary interactions in the material of the detector or of the decay of a short-lived particle that itself originated in a different vertex. While each event has many primary vertices the majority of these are due to pileup and typically only one primary vertex will correspond to a hard-scatter process that would have triggered the event. The set of tracks contributing to such a *hard-scatter primary vertex* will have a significantly higher $\sum {p_T}^2$ value than for pileup vertices, and so this quantity is used to identify the hard-scatter vertex. Details of the

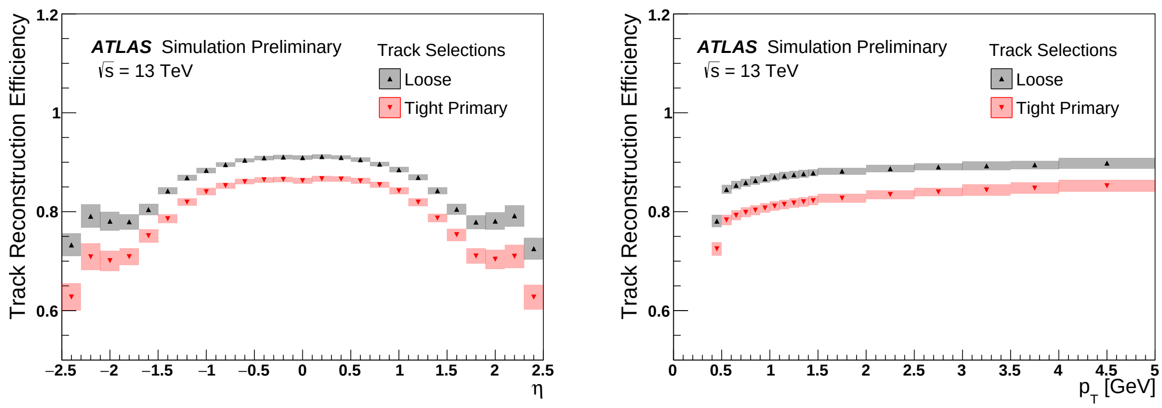

(a) (b)

Figure 3.2: Efficiency of track reconstruction in simulated minimum bias events (those collected by a trigger item that is designed to capture inelastic collisions with very loose requirements, to avoid bias) for two track selections as a function of (a) true η and (b) true track p_{T}. The shaded regions show the total systematic uncertainty. Taken from Ref. [28].

algorithm are described in Ref. [188].

Vertex reconstruction is essentially fully efficient when more than 4 to 5 tracks are used in the fit. The vertex position resolution is determined from a sample of data using a split vertex method. This method breaks tracks from reconstructed vertices into two subsets with approximately equal $\sum {p_{\mathrm{T}}}^2$ (by ranking tracks by p_{T} and assigning them alternately to each subset) and then fits a new vertex separately for each subset. The distance between the two daughter vertices is then a measure of the intrinsic resolution and is used to construct a scale factor parametrised by the track multiplicity that is used to correct the resolution estimated by the standard vertex fitting procedure in other events. In Run 2 the vertex resolution ranges from about 0.15 mm to 0.02 mm in x and y, and 0.3 mm to 0.04 mm in z as the track multiplicity increases [81]. Ahead of the HL-LHC upgrade, ATLAS plans to install a High Granularity Timing Detector [76] that will improve the time resolution in order to distinguish tracks that are physically close but temporally separated so that good vertex reconstruction performance can be achieved in very high pileup environments.

The vertex fit uses the beam spot centroid as a constraint, but this position is determined from the location of vertices. This cyclic dependency is resolved iteratively as discussed in §2.6.3. Accurate determination of the beam spot parameters relies on knowledge of the true vertex resolution, in the form of a vertex resolution scale factor, k. Although k is known to vary as a function of the track multiplicity, and such a parameterisation is used in the vertex fit, the beam spot fit uses only a single global parameter, floated in the fit. In order to reduce the small bias in the beam spot determination incurred as a result of this approximation, only primary vertices are used in the beam spot fit, as these in general have higher track multiplicity than pileup vertices. In higher-luminosity conditions it is more common for multiple physical vertices to be reconstructed as a single vertex, distorting the track multiplicity distribution. Effects such as these, and the desire to improve the statistics available to the fit by including non-primary vertices, have lead to efforts to change the implementation of the vertex resolution scale factor in the beam spot fit[1].

[1]The author qualified for the ATLAS experiment on the basis of work towards this goal [202].

3.3.3 Topological clusters

The formation of *topological cluster*s (topo-clusters) [79] is designed to exploit the full three-dimensional response of the calorimeters to identify correlated signals in nearby cells while being robust against electronic and pileup noise in the event. The initial clustering is done at the EM scale, meaning that no corrections are performed to compensate for the invisible fraction of the hadronic showers (see §2.4.2) and the differences in detector response to EM and non-EM processes. For each cell the significance $\varsigma_{\text{cell}}^{\text{EM}}$ is defined to be the ratio of the cell's signal to the average expected noise, and three thresholds S, N, and P are defined for the magnitude of this significance such that $S > N > P$. First all cells with $|\varsigma_{\text{cell}}^{\text{EM}}| > S$ are ranked by significance. Starting from the most significant, these seeds are grown by adding neighbours that pass the P threshold, and if any of these neighbouring cells additionally pass the N threshold their neighbours are also added. If any neighbouring cell added in this way passes the N threshold and is part of another cluster then that cluster is merged with the current one. This growth is repeated iteratively, so that the N parameter determines the growth characteristics of the clustering, the P parameter controls the envelope around the cluster, and the S parameter affects the multiplicity of the cluster seeds.

The clustering algorithm described above works using the magnitude of $\varsigma_{\text{cell}}^{\text{EM}}$ because negative cell signals are possible due to noise. Clusters with overall negative energy are generally due to out-of-time pileup and so can be used to estimate the local pileup conditions, while including negative cell energies in otherwise positive energy clusters cancels some of the positive noise fluctuations, avoiding bias that would be entailed by only incorporating the positive noise contributions. Only the positive energy clusters are used for reconstructing physics objects.

Because jets may be found in proximity to other jets, a splitting condition is applied to the clusters to avoid greedily merging all such jets. Local maxima, defined to be cells with an energy larger than 500 MeV and without any neighbouring cells more energetic than themselves, are used as the cores of the split clusters, and the cells of the original large cluster are divided between the local maxima. In cases where a cell neighbours more than one sub-cluster it is shared between the two highest energy sub-clusters where its contributions are weighted by its distance to the respective centres.

The η and ϕ of cluster barycentres are calculated by taking the sum of the η and ϕ coordinates of the contributing cells (with respect to the origin of the ATLAS coordinate system) respectively, weighted by the energy signal in the cells and by their geometric distance to the centre. The total cluster energy $E_{\text{clus}}^{\text{EM}}$ is similarly defined to be the weighted sum of cell signals. Using these quantities all clusters with $E_{\text{clus}}^{\text{EM}} > 0$ are treated as massless pseudo-particle objects as input to various other reconstruction algorithms. A set of moments is calculated, characterising the location, extent, and orientation of various features of the cluster. Hadronic showers, being typically wider in shape, are reconstructed as multiple topo-clusters more frequently than EM showers.

Clusters are calibrated in such a way that no assumption is made about the type of particle to which they are associated. Because the calorimeter response is different for EM and hadronic showers, topo-clusters are first assigned a likelihood to be of EM origin. This is based on comparing simulations of charged and neutral pions in the calorimeter, and uses information about the depth and average density of the cluster. A local hadronic correction is then applied to compensate for the invisible component of hadronic showers, weighted by the EM likelihood so that the calibration is smooth across different clusters. A final calibration step corrects for cells that are missed from the

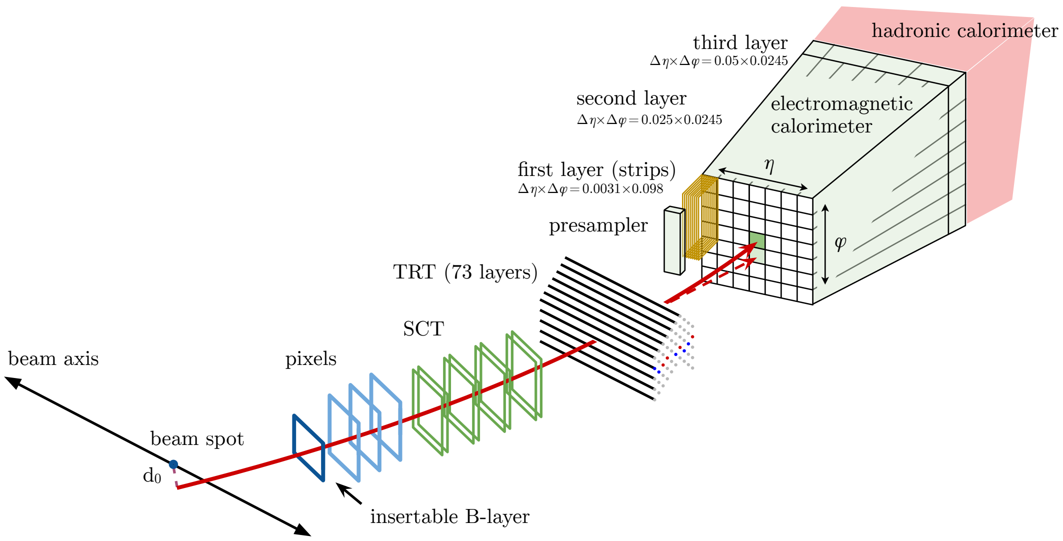

Figure 3.3: Illustration of the trajectory of an electron (in solid red) through the tracking detectors and the ECAL. The dashed line shows the trajectory of a photon emitted by the electron in the tracking system. Taken from Ref. [30].

cluster and deposits in inactive material in the detector.

3.4 Electron reconstruction

Electrons pass through the tracking layers of the detector and deposit typically all of their energy in the ECAL as shown in the schematic Fig. 3.3. Electrons are therefore reconstructed [30] by matching tracks to energy clusters in the ECAL. Through bremsstrahlung emission at any point along its trajectory, a primary electron can lose a significant fraction of its energy to photons. These photons may themselves pair produce new electrons such that the primary electron will often be observed as a collimated set of electrons and photons. Photons from other sources may also pair produce electrons, possibly before interacting with the detector, a process referred to as *photon conversion*. Therefore care must be taken to distinguish such electrons from primary electrons, and also to account for the energy of primary electrons carried by nearby photons.

First deposits in the ECAL are clustered. Using a granularity of $\Delta\eta \times \Delta\phi = 0.025 \times 0.025$ the energy of the EM cells is summed into projective towers. A sliding window algorithm [172] then identifies cluster candidates containing at least 2.5 GeV of energy by searching the entire ECAL and discarding duplicate clusters when candidates are found to overlap. Above $E_{\mathrm{T}} = 15\,\mathrm{GeV}$[2] this algorithm is over 99 % efficient at identifying clusters, as determined from checks in simulated samples. Note that these clusters are not related to the topo-clusters of the previous section; they are formed from the ECAL response calibrated at the EM scale only.

Track candidates reconstructed as described in §3.3.1 that have at least 4 silicon hits and coincide approximately with ECAL deposits are re-fitted using an alternative procedure called the Gaussian-sum filter [37] to account for energy loss through interactions with the detector material. These electron track candidates are then matched to the cluster seeds and, provided that the tracks have hits in the

[2] Since electrons are very light compared to their typical momenta in hard collisions, the quantities p_{T} and E_{T} are often used interchangeably. See §2.2.1 for the definitions of these quantities.

pixel layers and do not originate from secondary vertices, they are accepted as electron candidates. Finally the clusters are expanded to include the energy of nearby cells. This process results in a 97 % to 99 % efficiency to reconstruct a good electron candidate from true electrons above $p_{\mathrm{T}} = 15\,\mathrm{GeV}$[3], as measured in simulation and in a sample of $Z \to e^+e^-$ data [30].

3.4.1 Identification

From the assembled electron candidates, an electron likelihood is then calculated to describe the probability that the candidate corresponds to a true prompt electron. The alternative hypotheses include that the candidate is a true non-prompt electron (either from photon conversion or other secondary processes such as hadronic decays) or a fake electron reconstructed from another object such as a jet. The likelihood is based on features including the leakage of energy into the HCAL, characteristics of the EM shower including width and the fraction of energy deposited in each layer, hit multiplicity, impact parameter, and the compatibility of the match between track and cluster. The likelihood discriminant is then compared to different thresholds to create three working points for electron identification: `Loose`, `Medium`, and `Tight`, corresponding to an efficiency to reconstruct an $E_{\mathrm{T}} = 40\,\mathrm{GeV}$ electron of 93 %, 88 %, and 80 % respectively. Samples of $J/\psi \to e^+e^-$ events are used to estimate the likelihood down to as low as $E_{\mathrm{T}} = 4.5\,\mathrm{GeV}$. The `Tight` working point is 55 % efficient at $E_{\mathrm{T}} = 4.5\,\mathrm{GeV}$, increasing to 90 % efficient at $E_{\mathrm{T}} = 100\,\mathrm{GeV}$ [30].

3.4.2 Isolation

Prompt electrons typically have very little activity in the cone immediately around them whereas non-prompt electrons tend to be associated with the production of other objects, motivating the need to place requirements on the isolation of electron candidates. When objects are highly boosted with respect to the collision centre of mass, or in signal processes where multiple particles are produced with correlated momenta, prompt electrons are not necessary completely isolated. A range of information from tracking and calorimetry is combined to construct more sophisticated isolation discriminants.

Calorimeter-based isolation measures the energy due to additional activity in the vicinity of the electron candidate. In Run 2 this is done by summing the energy of all positive-energy topo-clusters with a barycentre lying within a cone of size $\Delta R_\eta = 0.2$ of the electron candidate[4]. The energy within a fixed rectangular region in the core is subtracted to remove the electron's own contribution. A calibrated E_{T}-dependent correction is applied to account for energy due to the electron that leaked outside the fixed-size core. An additional correction for pileup contributions based on the average energy density of the event is applied.

Track-based isolation measures the additional activity by taking the sum of the p_{T} of good quality tracks in a cone around the electron candidate, subtracting the contributions from any tracks that, when extrapolated into the second layer of the ECAL, lie within $\Delta\eta \times \Delta\phi = 0.05 \times 0.1$ of the electron cluster position. The radius used for the track-based quantity can either be fixed or inversely

[3]Tracks candidates with p_{T} as low as 400 MeV are considered but track fitting does not become efficient until about 10 GeV at which point it is 98 % efficient.

[4]Previously the energy response of the calorimeter cells was used, but the properties of topo-clusters with regard to noise and pileup significantly improve the performance. These topo-clusters are used at the EM scale, without applying the final hadronic calibration.

proportional to the electron p_{T}.

Requirements are then placed on the calorimeter-based and track-based estimates of additional activity, calibrated to achieve an isolation efficiency independent of η but dependent on the E_{T} of the electron (called the *gradient isolation* algorithm). The looser variant of the gradient isolation algorithm is about 95 % efficient at low E_{T} whereas the tight variant is 90 % efficient. At $E_{\mathrm{T}} = 60\,\mathrm{GeV}$ both variants are about 99 % efficient [30]. Alternative isolation criteria are defined that exhibit flat efficiency with respect to the electron candidate p_{T}.

3.4.3 Performance

It is possible for an electron to be reconstructed with the wrong EM charge sign, predominantly due to the assignment of the incorrect track to the electron when there are three closely aligned tracks due to an electron and positron pair produced from bremsstrahlung emission from the primary electron. This predominantly occurs in the region $1.5 < |\eta| < 2.2$ due to the large amount of inactive material in the detector. Charge mis-identification can also occur when tracks are very forward or correspond to very high momentum particles, since in these cases the track's curvature in the magnetic field is very small. The rate of charge mis-identification ranges from less than 0.1 % in central electrons up to about 3 % in electrons with $|\eta| = 2.5$ [30].

The combined efficiency for electron reconstruction, identification, and isolation is shown in Fig. 3.4, demonstrating the good performance over the precision η range and the improvement of efficiency with transverse energy. The "Tight + Gradient" working point in the figure corresponds to the main electron selection used in the analyses presented in this thesis. When selecting a sample of electrons for analysis, it is typical to exclude electrons with low E_{T} since the electrons reconstruction efficiency falls steeply there, and also electrons near $1.4 < |\eta| < 1.5$ where the efficiency dips due to the transition between the barrel and end-cap regions of the detector.

The energy of electrons is calibrated [29] using data and the energy resolution is measured. A multivariate regression technique developed using simulated samples is used to correct the energy of ECAL clusters to account for energy lost to material before the calorimeters, energy leaked into the hadronic calorimeter, and energy leaked to neighbouring calorimeter cells. Relative calibrations of each ECAL layer are then determined using a $Z \to \mu^+\mu^-$ sample since muons do not typically lose much energy through interactions with passive material in front of the calorimeter. Corrections for the ambient energy in the event due to pileup are calibrated as a function of instantaneous luminosity and the position of colliding bunches in the LHC bunch train to correct for localised variations in the pileup background. Further corrections are applied to account for localised variations in the calorimeter response due to the module geometry, voltage irregularities, and gaps in the detector coverage. Finally the absolute energy scale is determined using a $Z \to e^+e^-$ sample by comparing the observed energy scale to that found in simulated samples. A fully calibrated sample of data is compared to MC simulation in Fig. 3.5. For an electron with $p_{\mathrm{T}} = 10\,\mathrm{GeV}$ the typical uncertainty in energy is in the range 0.3 % to 0.8 %.

The transverse momentum of the electron is constructed as the ratio $p_{\mathrm{T}}(e) = E/\cosh\eta$ where E is the calibrated cluster energy discussed above, and η is the pseudorapidity measured precisely by the track reconstruction. In this way the strengths of the tracking and calorimetry systems are combined whereas in principle the momentum can be redundantly measured by both systems.

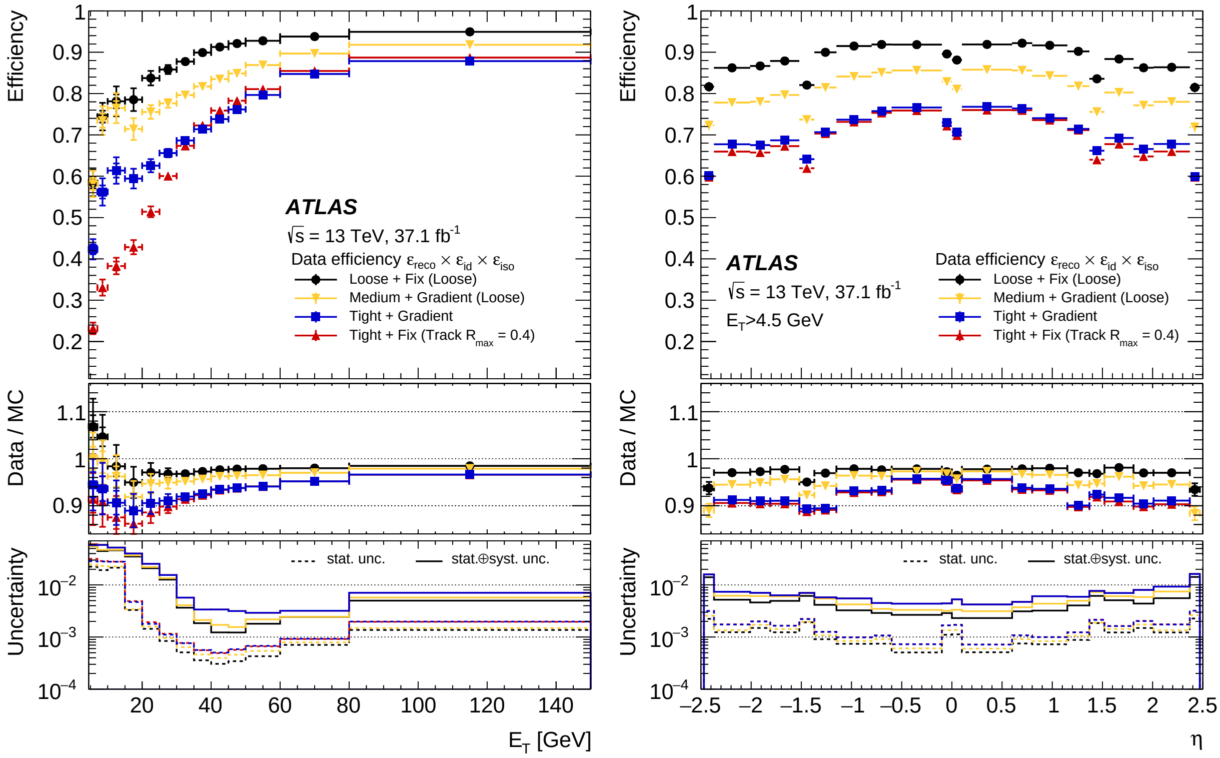

(a) **(b)**

Figure 3.4: Efficiency of several electron reconstruction working points with respect to (a) E_{T} and (b) η, as measured in data. The lower panels show the ratio to the efficiency seen in simulated MC events along with the relative uncertainty. Taken from Ref. [30].

Figure 3.5: Distribution of the invariant mass of two electrons in a sample of $Z \to e^+e^-$ candidates comparing calibrated data and MC simulation. The MC sample is normalised to the same area as the data sample. The error band in the bottom panel shows the uncertainty arising from the calibration. Taken from Ref. [29].

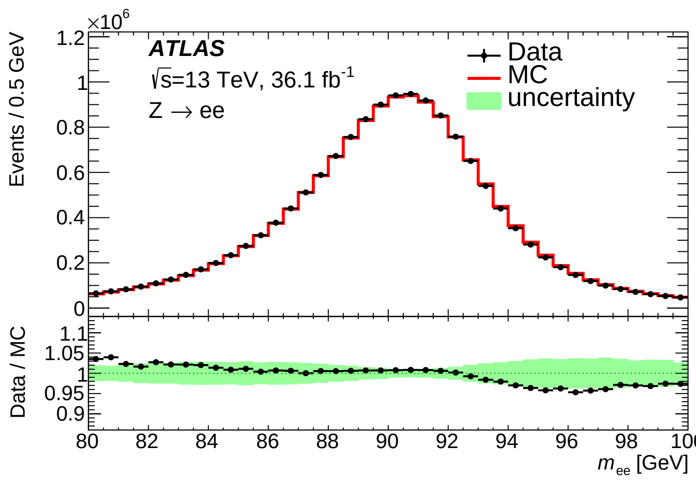

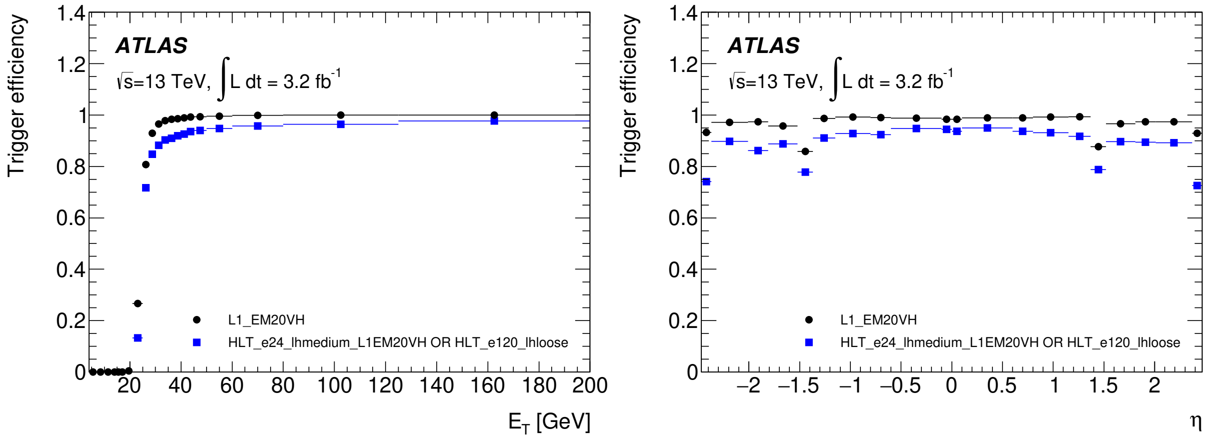

(a)

(b)

Figure 3.6: Trigger efficiencies for the logical OR of `HLT_e24_lhmedium_L1EM20VH` and `HLT_e120_lhloose` with respect to the electron (a) E_{T} and (b) η. The quantities on the horizontal axes are those of the electrons when fully reconstructed with the offline algorithms. These measurements were performed using the 2015 dataset and show the lowest unprescaled single lepton triggers used for that dataset. Taken from Ref. [64].

3.4.4 Single electron triggers

The trigger menu is designed so that primary physics triggers such as the single lepton triggers relevant to this thesis can be operated without prescaling. With respect to the configuration in Run 1, the single electron trigger thresholds for 2015 were raised from 18 GeV to 20 GeV at L1, but the HLT threshold was maintained at 24 GeV (although with the trigger electron isolation requirement now dropped for 2015) [13]. From 2016 the LHC operated at a higher luminosity requiring that the thresholds be slightly raised to cope with the higher trigger rate [80].

The electron HLT triggers proceed by clustering energy deposits in calorimeter cells within the $\Delta\eta \times \Delta\phi = 0.4 \times 0.4$ regions identified by the L1 trigger seed, placing requirements on the total energy and the shape of the deposit, and then requiring an approximate match to a high-p_{T} track reconstructed by a fast tracking algorithm. Candidate electrons passing these requirements are then reconstructed with more detailed algorithms that perform a multivariate likelihood based identification with loose, medium, and tight working points. The lowest unprescaled electron trigger in 2015 was `HLT_e24_lhmedium_L1EM20VH`[5]. From 2016 the lowest unprescaled trigger was `HLT_e24_lhtight_nod0_ivarloose`, which imposes a more stringent electron likelihood requirement in the HLT and a loose isolation requirement. These required an L1 seed with $p_{\mathrm{T}} > 20$ GeV subject to a veto on energy deposits in the HCAL to avoid triggering on charged hadrons. Higher-p_{T} electrons can leak energy into the HCAL and fail this additional requirement, so the primary triggers are intended to be complemented by a higher-p_{T} single electron trigger without this veto.

The performance of the electron triggers was studied using a sample of $Z \to e^+e^-$ events selected by requiring two electrons with an invariant mass consistent with the Z boson. In each event the electron which did not trigger the event is checked against the trigger conditions and the fraction of

[5] Trigger names follow a conventional structure. They begin with the level (`L1` or `HLT`) followed by the particle multiplicity, type, and p_{T} threshold (e.g. `L1_2MU4` means at least 2 muons with $p_{\mathrm{T}} > 4$ GeV at L1) where an uppercase particle name like `MU` indicated an L1 object and a lowercase name like `mu` indicated a HLT object. Various other requirements are indicated by further suffixes.

such electrons passing the requirements is plotted in Fig. 3.6.

3.5 Muon reconstruction

Hits in the MS are relatively strong indicators of muons as most other charged particles are stopped in earlier layers of the detector. Muon candidates are therefore initially reconstructed [56] using only the MS before reconciling these objects with tracks reconstructed in the ID.

In the MS MDT track segments are built in the bending plane using a technique based on Hough transforms [160], and CSC track segments are separately built using an IP constraint. A combinatorial search is then used to match segments in different MS layers to build track candidates. A global χ^2 fit is then used to remove hits that are in poor agreement with the track candidate to which they were assigned.

The majority of muon candidates are formed by matching MS track candidates to ID tracks and re-fitting globally to improve the track parameter estimates, with several further types of muon candidates formed to recover efficiency. If an ID track when extrapolated outwards is found to be consistent with a single MS track segment, it forms a candidate muon on the assumption that the muon's trajectory intersected with only a single layer of chambers. If an extrapolated ID track is consistent with a calorimeter energy deposit with certain restrictions then it also forms a muon candidate in order to recover muon efficiency in regions where the MS acceptance is reduced due to detector services. Good MS candidate tracks that are consistent with an inwards extrapolation to the IP are also used as candidates since the MS covers a larger $|\eta|$ range than the ID. Ambiguities between these classes of muon candidates are resolved such that each candidate is assigned to exactly one category, and these are separately studied and calibrated.

3.5.1 Identification and isolation

The main sources of non-prompt muons are pion and kaon decays in the ID, and to a large extent these can be suppressed by requiring consistency in the momentum measured by the ID and the MS and by requiring a smooth track. Muon identification is done by placing requirements on the ID hit multiplicity, the disagreement between the p_{T} and the q/p (charge divided by momentum) of the muon candidates in the ID and MS, the track fit χ^2, and the location of the muon candidate, taking into account the type of candidate. Different requirements are used to define progressively more restrictive `Loose`, `Medium`, and `Tight` muon identification working points. Unlike the electron identification working points, which correspond to thresholds of a likelihood, the muon working points differ in the type of muon candidates they permit as well as specific quality requirements. The `Medium` working point used in this thesis selects predominantly muon candidates that are constructed from both ID and MS tracks, with only about 0.5 % of candidates tracked by a single system. A loose requirement is placed on the compatibility of the q/p values measured by the ID and the MS in order to reject hadronic background.

Muons are identified down to $p_{\mathrm{T}} = 4\,\mathrm{GeV}$ though with reduced efficiency in the lower p_{T} candidates and in the forward regions outside of the acceptance of the ID. For muons with $p_{\mathrm{T}} > 20\,\mathrm{GeV}$ at the `Medium` working point the identification efficiency as measured in a simulated $t\bar{t}$ sample is 96.1 % with a 0.17 % efficiency to identify non-prompt muons from hadronic decays. The full reconstruction efficiency was checked using $Z \to \mu^+\mu^-$ and $J/\psi \to \mu^+\mu^-$ samples as shown in Fig. 3.7, showing

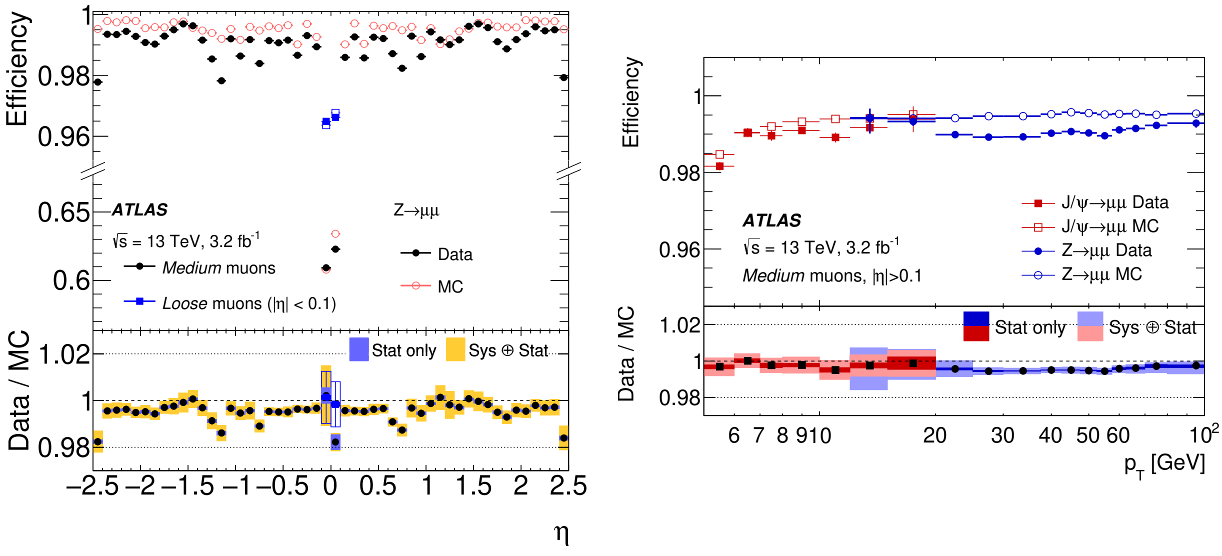

(a) **(b)**

Figure 3.7: Efficiency of the Medium muon reconstruction working point with respect to (a) η in events with $p_{\mathrm{T}} > 10\,\mathrm{GeV}$ and (b) p_{T} in events with $0.1 < |\eta| < 2.5$, as measured in 2015 data. The lower panels show the ratio to the efficiency seen in simulated MC events. Taken from Ref. [56].

good performance across the full precision η range outside of a crack region near $\eta = 0$, and across a large p_{T} range [56].

The prompt muons produced in the decay of particles such as W bosons will tend to be isolated, so a measure of the activity around the candidate is useful to distinguish prompt muons from non-prompt muons; this isolation measure is constructed using a track-based quantity and a calorimeter-based quantity. The track-based quantity is the scalar p_{T} sum of all tracks with $p_{\mathrm{T}} > 1\,\mathrm{GeV}$ excluding the muon track itself that fall inside a cone of size $\Delta R_\eta = \min(0.3, 10\,\mathrm{GeV}/p_{\mathrm{T}}(\mu))$ where $p_{\mathrm{T}}(\mu)$ is the p_{T} of the muon candidate. The calorimeter-based quantity is defined analogously to the equivalent quantity for electrons described in §3.4, taking the sum of topo-clusters around the muon deposit and correcting for ambient energy in the event. Different isolation working points are then defined including the gradient isolation criterion used in this thesis for which the performance is relatively flat with respect to η.

3.5.2 Momentum calibration

The momentum of the reconstructed muons in MC simulation is calibrated to account for imperfections in the modelling of the momentum resolution and scale. The calibration is determined separately in several η–ϕ regions corresponding to parts of the detector with similar muon momentum resolution. Simulated $Z \to \mu^+\mu^-$ and $J/\psi \to \mu^+\mu^-$ samples are fit to real data, using isolation requirements in the former case and a data-driven estimate of the continuum background in the latter case. These fits determine corrections to the ID and MS momentum response, and the final corrected momentum of the muon candidate is taken to be the weighted average of the corrected measurements in the two detectors. The effect of the momentum correction is shown in Fig. 3.8 where the correction is seen to result in a broadening and shifting of the simulated resonance peaks so that they agree with the observed peak shapes within uncertainty.

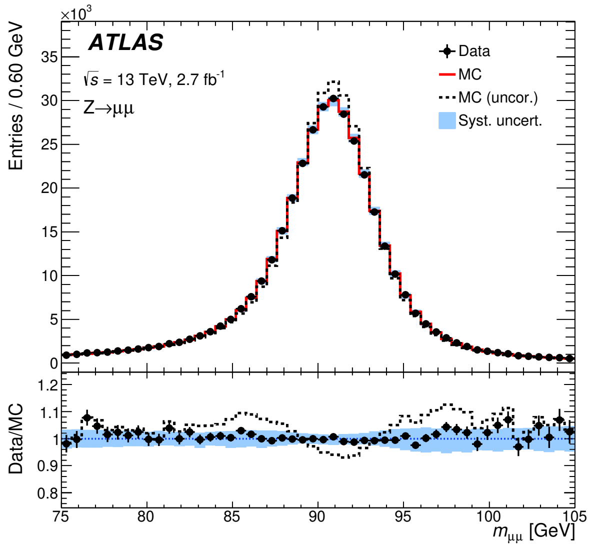

(a) $Z \to \mu^+\mu^-$

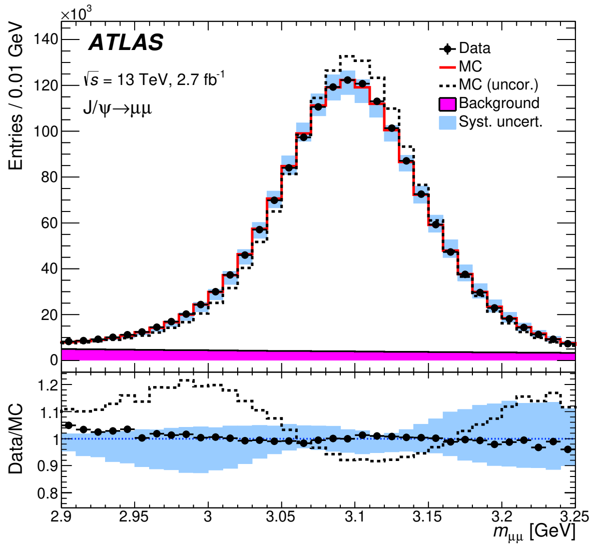

(b) $J/\psi \to \mu^+\mu^-$

Figure 3.8: Distribution of the invariant mass of two muons in a sample of (a) $Z \to \mu^+\mu^-$ candidates and (b) $J/\psi \to \mu^+\mu^-$ candidates, comparing calibrated data and MC simulation with the uncorrected simulation. The error band in the bottom panel shows the uncertainty arising from the calibration. Taken from Ref. [56].

3.5.3 Single muon triggers

The single muon L1 trigger threshold for 2015 was kept at 15 GeV, as in Run 1, while the HLT threshold was lowered from 24 GeV in Run 1 to 20 GeV in 2015 due to the more sophisticated HLT reconstruction algorithms available in Run 2 [64]. The muon HLT triggers take seed information from the dedicated TGC and RPC detectors and supplement it with precision information from the other muon detectors and from the ID tracking. The lowest unprescaled single muon trigger in 2015 was `HLT_mu20_iloose_L1MU15` which imposed loose isolation requirements, excluding muon candidates where there are tracks in a neighbouring $\Delta R_\eta = 0.2$ cone with a total p_T exceeding 12 % of the p_T of the muon itself. To recover efficiency for high-p_T muons this trigger was intended to be complemented by an additional high-p_T muon trigger without the isolation requirement.

From 2016 the LHC operated at a higher luminosity requiring that the thresholds be slightly raised to cope with the higher trigger rate. The lowest unprescaled single muon trigger in 2016 was `HLT_mu26_ivarmedium`, which requires a muon of at least 26 GeV, and with tighter isolation requirements than the corresponding 2015 trigger.

Similarly to what was done with the single electron triggers, the muon trigger efficiencies were studied in a $Z \to \mu^+\mu^-$ sample with the resulting efficiency plotted in Fig. 3.9, showing a relatively sharp turn-on with a plateau after about 20 GeV. The efficiency of the L1 trigger in the barrel region is lower than that in the end-caps, and varies with η as a result of the layout of the MS [64].

3.6 Jet reconstruction

Collimated jets of particles arising from the hadronisation of strongly-interacting partons in the hard scatter are treated as pseudo-particles. Jets at ATLAS are reconstructed with the anti-k_t clustering

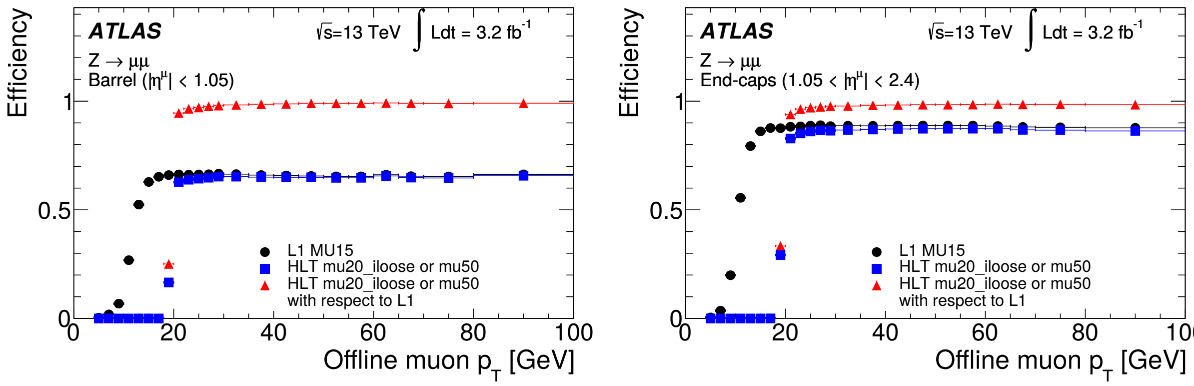

(a) barrel region **(b)** end-cap regions

Figure 3.9: Trigger efficiencies for the logical OR of `HLT_mu20_iloose_L1MU15` and `HLT_mu50` with respect to the offline muon p_{T} in (a) the barrel region $|\eta| < 1.05$ and (b) the end-cap regions $1.05 < |\eta| < 2.4$. These measurements were performed using the 2015 dataset and show the lowest unprescaled single lepton triggers used for that dataset. Taken from Ref. [64].

algorithm discussed in §1.4 using topo-clusters as inputs[6] and $R = 0.4$. Any jets reconstructed with a p_{T} below 7 GeV are discarded. In order to calibrate and study jet reconstruction there is a need to define some equivalent object in MC simulation with known kinematics; this is done by applying the anti-k_t algorithm with $R = 0.4$ to objects in the generator record corresponding to stable[7] particles other than muons, neutrinos, or particles from pileup.

Jets are assigned a dimensionless area according to a *ghost-matching* procedure [100]. This procedure measures the extent of the jet in the clustering algorithms, although they have no well-defined physical area, by inserting four-vectors with infinitesimal p_{T} uniformly across the solid angle of the event prior to re-clustering. The fraction of these “ghost” objects that are clustered into a jet then defines its dimensionless jet area. ID tracks associated with the hard-scatter primary vertex and with a $p_{\mathrm{T}} > 500$ MeV are matched to jets using this same ghost-matching procedure: their momentum is rescaled to a vanishingly small value and then the jets are re-clustered. MS-only track segments are similarly ghost-matched to jets for the purpose of accounting for particles leaked through the HCAL into the MS. By associating ID tracks to a jet it can be assigned to the corresponding primary vertex. Its four-momentum is then recalculated using the vertex as its origin instead of the origin of the ATLAS coordinate system used in the definition of the topo-cluster kinematics.

3.6.1 Jet calibration

For analyses such as those presented in this thesis, the uncertainty due to the calibration of jet energy can contribute a large fraction of the total measurement uncertainty; the calibration of jets is therefore of particular importance. The calibration is performed in a series of steps combining information from the calorimeters with the matched tracks and reference measurement [43].

[6] Other types of jets are defined using different methods of clustering energy deposits in the calorimeter, for example particle flow jets [44] that subtract deposits due to charged hadrons by using momenta measured in the ID, these are not used for the jets in this thesis.

[7] Stable here refers to particles with a lifetime τ such that they propagate a distance $c\tau > 10$ mm.

Pileup correction

The pileup activity in each event is estimated by re-clustering topo-clusters of positive energy in the central region $|\eta| < 0.2$ using the k_t algorithm (see §1.4) with $R = 0.4$, which is more sensitive to soft radiation than the anti-k_t algorithm. The median p_T density (the jet p_T normalised by the jet area), ρ, of these k_t jets is taken to be the pileup p_T density of the event. The density ρ is estimated from the central region of the detector and, in more forward regions, exhibits a dependence on pileup due to the non-uniformity of pileup sensitivity across the detector. Using MC simulated samples, fits are performed in different $|\eta|$ regions to obtain coefficients $\alpha(|\eta|)$ and $\beta(|\eta|)$ for correction terms proportional respectively to N_PV (the number of reconstructed primary vertices, sensitive to in-time pileup) and to μ (instantaneous pileup measured by ATLAS, sensitive to out-of-time pileup). The pileup-corrected momentum ${p_\mathrm{T}}^\mathrm{corr}$ of an anti-k_t jet with momentum ${p_\mathrm{T}}^\mathrm{reco}$ and area A is then given by [43]

$${p_\mathrm{T}}^\mathrm{corr} = {p_\mathrm{T}}^\mathrm{reco} - \rho \times A - \alpha(|\eta|) \times (N_\mathrm{PV} - 1) - \beta(|\eta|) \times \mu.$$

Absolute jet energy scale calibration

The absolute jet energy scale (JES) is determined using an MC simulated sample. After correcting for pileup, isolated reconstructed jets are matched to truth jets within a cone of $\Delta R_\eta = 0.3$ to determine their truth-level energy, E^truth. In several bins of E^truth and η (calculated from the centre of the detector rather than the vertex) Gaussian fits are performed to the core of the distribution of $E^\mathrm{EM}/E^\mathrm{truth}$ where E^EM is the reconstructed jet energy at the EM scale. For each bin the averaged jet response is defined to be the mean of the fitted Gaussian, and this response is observed to be lowest in transition regions between detector technologies as expected. In each bin of η, a fit is then performed to determine the coefficients of a polynomial in $\ln(E^\mathrm{EM})$ parametrising the averaged jet response. Finally this polynomial is inverted to obtain a JES calibration as a function of η. Jets that evolve near a transition region between detector technologies can have a bias in their η calculation due to the different jet response on each side of the jet. A correction to η (applied only to the p_T and η components of the four-momentum) is determined using a similar calibration procedure based on the truth-level η.

Sequential calibration

At this point in the calibration the jet energy is independent of p_T and η but still depends on the composition of jets, in particular whether a jet was initiated by a quark or by a gluon, since this affects the shape and energy distribution of the jet. For each variable correlated with the calorimeter response, the jet four-momentum is corrected as a function of the truth-level p_T and detector η using a similar numerical inversion technique to that used for the initial JES calibration. This calibration is applied sequentially to correct for the fraction of jet energy deposited in the first TileCal layer, the fraction deposited in the third LAr layer, the number of tracks matched to the jet, the average p_T-weighted distance between the jet axis and the matched tracks, and the number of MS tracks matched to the jet. As each stage of this sequential calibration the corrections are normalised to ensure that the average energy response is unaffected. The overall effect is to improve the jet energy resolution (JER), the resolution with which the energy of jets is measured after correcting for the JES.

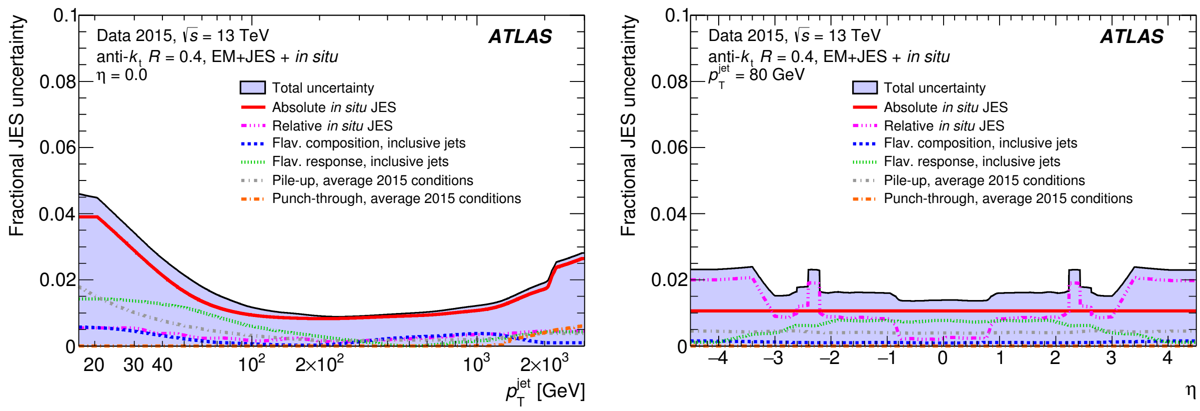

(a) (b)

Figure 3.10: Total uncertainty on the JES calibration as a function of (a) jet p_{T} at $\eta = 0$ and (b) η at $p_{\mathrm{T}} = 80\,\mathrm{GeV}$. These measurements were performed using the 2015 dataset and assume a flavour composition of the jets based on dijet events. Taken from Ref. [64].

In-situ calibration

For real data (but not MC simulation) a final *in-situ* calibration is performed to account for the residual differences between data and simulation due to imperfections in the descriptions of the detector and the details of the physical modelling of jets and the event. The first part of this calibration exploits the momentum balance of events with a dijet topology, that is where there are two isolated jets that are back-to-back in ϕ. A sample of such events where one jet falls into the central region of the detector ($|\eta| < 0.8$) and the other falls into the forward region ($0.8 < |\eta| < 4.5$) is used to inter-calibrate the energy response across these regions. The next part of the calibration applies to the central region and uses samples of events where a single jet with p_{T} less than about 1 TeV is produced in association with a Z boson or a photon. Since the boson energy response is independently calibrated (via $Z \rightarrow e^+e^-$ and $Z \rightarrow \mu^+\mu^-$ decays) it can be used as a reference for the jet energy by assuming that the momenta of the two objects balance, which is true in the case where there are no additional jets initiated by partons radiated by the boson. The last part of the calibration uses events with at least three jets where the high-p_{T} jet is balanced against the fully-calibrated lower-p_{T} jets in order to extend the calibration up to higher momenta jets.

Uncertainty

The overall uncertainty on the JES calibration starts from around 4.5 % at $p_{\mathrm{T}} = 20\,\mathrm{GeV}$ and decreases to 1 % at 200 GeV since the energy response is better understood for harder jets. The uncertainty again increases slowly at the higher p_{T} reaches of the calibration due to the high-p_{T} extension method discussed above. The uncertainty is relatively flat with respect to η except for an increase in the transition regions of the detector and a higher uncertainty in the forward regions where the calorimeter granularity is reduced. These trends are shown in Fig. 3.10 assuming events with a dijet topology. Note however that knowledge of the flavour composition (i.e. the fraction initiated by gluons rather than light quarks) of jets in a sample is one component of this uncertainty, and so in a given analysis where the flavour composition is known more accurately based on modelling of the relevant physical processes, it is in principle possible to reduce the uncertainty from this source.

3.6.2 Rejection of bad jets

In addition to collimated streams of hadrons from hard partons, any other energy deposited in the calorimeters can potentially be reconstructed as a jet. Some of these deposits are due to non-collision sources such as noise bursts in calorimeter cells, ambient background in the experimental hall or from cosmic rays, unintended interactions between the beams and gas around the IP, or a diffuse halo of muons and pions formed through long-ranged beam-beam interactions. Jets reconstructed from calorimeter noise do not exhibit the characteristic pulse shape of real jets in the calorimeter, and generate both positive and negative contributions to the energy response. Jet quality discriminants are constructed by comparing the expected to observed pulse shape, and measuring the total negative energy response in cells. Non-collision sources also result in unusual distributions of energy across the segments of the calorimeters, so further discriminants are constructed from energy ratios between different calorimeter segments, and between the fraction of charged tracks and the largest energy deposit in the matched cluster. These quantities were each found to be successful at distinguishing fake jets reconstructed from noise by using a sample of real dijet events in which a momentum imbalance is a good indicator of the presence of a fake jet. A set of requirements on the jet quality discriminants is used to identify fake jets while remaining 99.5 % efficient at selecting real jets with $p_{\mathrm{T}} > 20\,\mathrm{GeV}$ from pp collisions [71, 26].

Jets associated with pileup are also a major challenge at hadron colliders. While the average energy contribution due to pileup is subtracted as part of interpreting the calorimeter response, residual local fluctuations due to pileup activity are reconstructed as jets. A discriminant called the jet-vertex-tagger (JVT) was developed during Run 1 to identify jets from pileup activity, combining two quantities in a likelihood. The first quantity is a ratio of the p_{T} sum of tracks associated with the jet that originate from the hard-scatter primary vertex to the p_{T} sum of tracks associated with the jet that originate from any other vertex, corrected to remove pileup dependency in the quantity[8]. The second quantity is the ratio of the p_{T} sum of tracks associated to the jet which originate from the hard-scatter primary vertex to the fully calibrated jet p_{T}. The JVT was studied in samples of real and simulated $Z \to \mu^+\mu^-$ events with associated jets and found to perform well and with no observable pileup dependence [62]. For the JVT working point relevant to this thesis, 98 % of pileup jets are rejected while maintaining a 92 % efficiency to select non-pileup jets. The number of jets with $p_{\mathrm{T}} > 20\,\mathrm{GeV}$ before and after the JVT requirement is shown as a function of $\langle\mu\rangle$ in Fig. 3.11, showing the successful removal of pileup jets.

The usual process for removing bad jets is to begin by applying a JVT requirement to all jets with $p_{\mathrm{T}} > 20\,\mathrm{GeV}$, and then remove any jets that overlap with other hard objects in the event such as charged leptons. If any of the remaining jets are identified as fake jets by the discriminants discussed above, the entire event is excluded from the selection.

3.6.3 *b*-jet tagging

In top quark analyses it is essential to be able to identify jets that correspond to the b-quarks produced in the hard process; in the leading order description this distinguishes tW from $t\bar{t}$ and other background processes. It is also necessary to accurately know the efficiency of any algorithm

[8] Earlier in Run 1 a quantity called jet vertex fraction (JVF) [65] was used in place of the JVT. It was defined similarly to this first quantity but without the correction to remove pileup dependency. Because of this, JVF requirements became less efficient at selecting non-pileup jets as the pileup activity in the event increased.

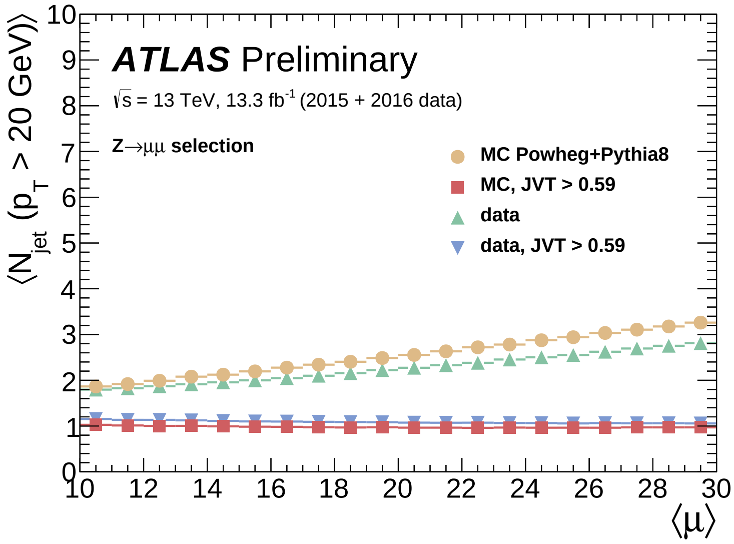

Figure 3.11: Number of jets with $p_T > 20\,\text{GeV}$ before and after applying JVT requirements in both MC simulated and real $Z \to \mu^+\mu^-$ events collected in 2015 and 2016. Before the requirements are applied the number of jets increases with increased pileup activity, measured by the average number of interactions per bunch crossing $\langle\mu\rangle$, whereas after the JVT requirements the jet multiplicity is almost flat with respect to additional pileup activity. Taken from Ref. [45].

used to identify these b-jets, as well as the probability to mistakenly b-tag a jet initiated by a c quark (the *c-jet tagging efficiency*) or a light parton (u, d, s, or g; collectively called the *mistag rate*). The b-tagging algorithms exploit the fact that hadrons containing b quarks can have lifetimes of approximately 1.5 ps meaning that energetic hadrons can travel on the order of a few mm before decaying and hence it is often possible to associate b-jets with a secondary vertex. The b-tagging algorithms therefore rely primarily on the ID tracking performance, and are sensitive to the resolution of the primary vertex which is used as a reference to calculate the impact parameters (see §3.3.1).

Compared to b-tagging in Run 1 [60], in Run 2 ATLAS benefited from the improved tracking resolution associated with the addition of the IBL (see §2.3.1). This improvement is shown in Fig. 3.12 where the light jet and c-jet rejection factors both substantially improve for a given b-tagging efficiency. The initial Run 2 algorithm used in 2015 [27] was improved for 2016 using an improved training method [59], and again for 2017–18 using improved pixel simulation and an expanded set of algorithms [58]. This section describes b-tagging as was used for the analysis of the 2016 dataset.

Three algorithms are used to identify b-jets, and their outputs are combined to form a single discriminant called `MV2`. The first selects tracks with at least 1 GeV and at least 7 silicon hits with $|d_0| < 1\,\text{mm}$ and $|z_0 \sin\theta| < 1.5\,\text{mm}$. Based on information from a reference sample, a b-tag likelihood discriminant is calculated, taking into account detailed information about the track quality from the tracking algorithm. The second algorithm also starts with a set of good-quality tracks, but attempts to find a subset that can be reconstructed as a secondary vertex with good fit quality that is not consistent with a secondary vertex originating from other known sources of secondary interactions or long-lived particles. The third algorithm is based on the `JetFitter` approach described in Ref. [33, p. 445], attempting to reconstruct vertices inside the jet corresponding to the decays of b- and c-hadrons.

The outputs of these three algorithms are combined using a boosted decision tree (BDT) (see §4.4.1) which is also provided with information about the jet kinematics and is trained on a simulated $t\bar{t}$ sample to distinguish b-jets from other jets. The resulting discriminant is called `MV2c10`. The performance was measured in a statistically independent $t\bar{t}$ sample and four working points (thresholds on the discriminant) were defined corresponding to 60 %, 70 %, 77 %, and 85 % b-tagging efficiencies. The working points represent a trade-off between b-tagging efficiency and the mistag rate; at the

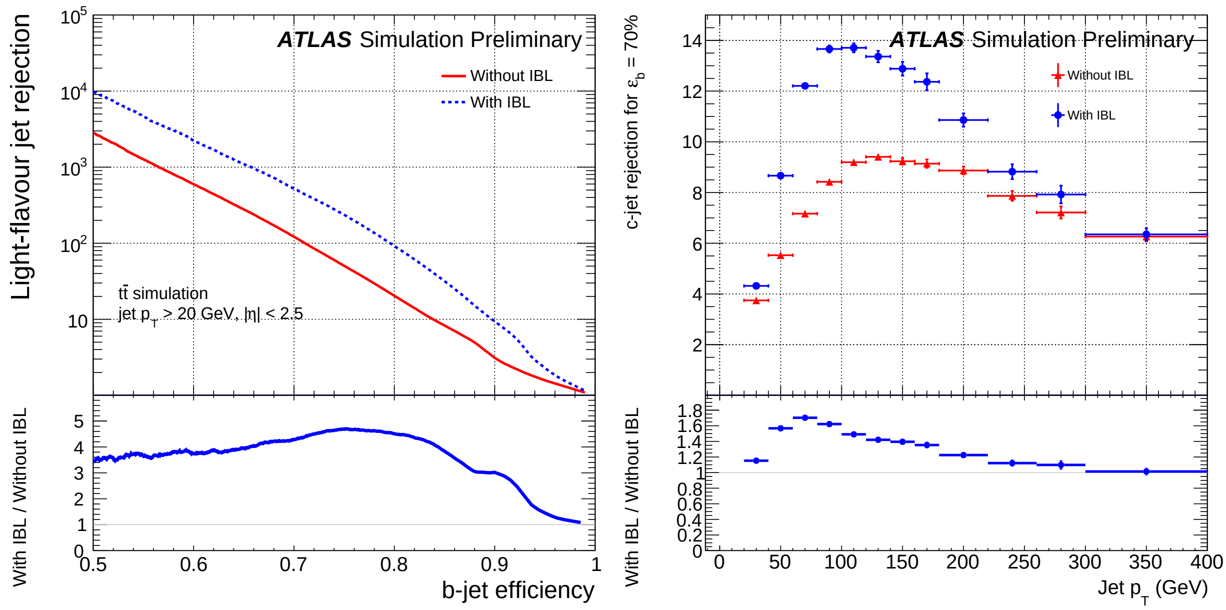

(a) **(b)**

Figure 3.12: Comparisons of b-tagging performance with and without the IBL in a simulated sample, expressed as (a) light jet rejection as a function of b-tagging efficiency, and (b) c-jet rejection as a function of jet p_{T}. In both cases better performance corresponds to a larger area under the curve. Taken from Ref. [67].

77 % working point used in this thesis the c-jet tagging efficiency is about 22 % and the mistag rate is less than 1 % [25]. The working point with 85 % b-tagging efficiency is also more efficient at tagging c-jets by a factor of two, and so would result in a higher rate of misidentified b-jets. The efficiency of the algorithm at the 77 % efficiency working point is shown in Fig. 3.13 as a function of the jet p_{T} and $|\eta|$ as measured using a sample of $t\bar{t}$ events collected in 2015 and 2016.

A sample of real dijet events with a topology consistent with the presence of 2 b-jets was used to check the discriminant [54]. In each event, a b-jet was selected using the 85 % working point, and the other jet was labelled as a probe jet and the MV2c10 discriminant for these probe jets was then compared to simulation. Good agreement was demonstrated in the b-jet-like region, and the small residual disagreement was used to calibrate the discriminant in simulated events. The largest sources of systematic uncertainty in the calibration are due to MC modelling and JES determination.

3.7 Missing transverse momentum reconstruction

Neutrinos are not expected to interact with the ATLAS detector due to their extremely low interaction cross-sections, and so they cannot be directly reconstructed; they are the only particles in the SM that are expected to be invisible in this way. While the boost along the beam axis is not known as it depends on the fraction of the proton momenta carried by the partons interacting in the hard scatter, the transverse momentum of the collision is expected to be zero. Reconstructing an imbalance in the transverse momentum of visible particles in the event then allows a missing transverse momentum vector to be reconstructed. Other objects that escape the acceptance of the detector by propagating in crack or very forward regions (including inside the beam pipe), or noise and misidentified objects can also contribute to the momentum balance of the event. The missing transverse momentum

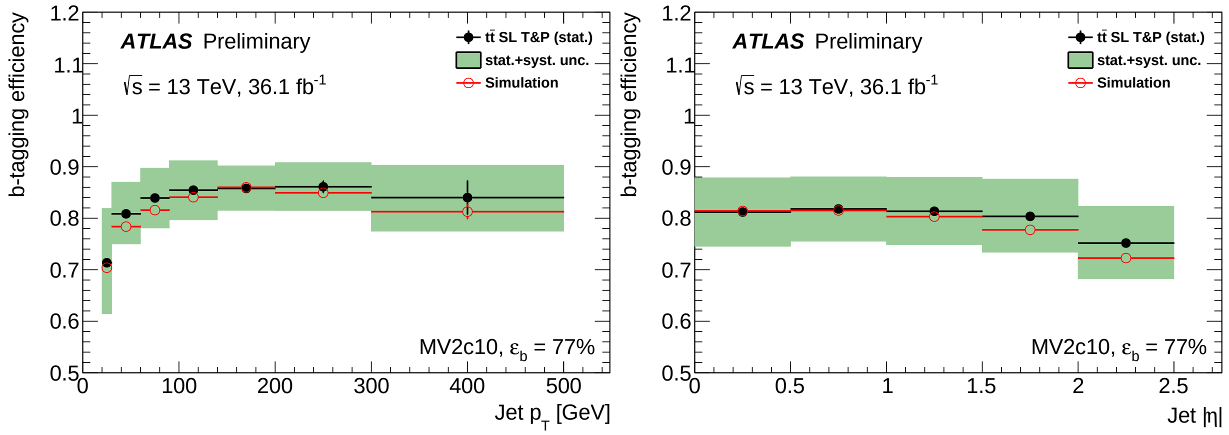

(a) **(b)**

Figure 3.13: Efficiency of the MV2c10 b-tagging algorithm at the 77 % working point as a function of (a) jet p_{T} and (b) jet $|\eta|$. This calibration was performed using a sample of $t\bar{t}$ events decaying to a single charged lepton and associated jets. Taken from Ref. [25].

vector $\vec{E}_{\mathrm{T}}^{\mathrm{miss}}$ therefore cannot be directly interpreted as the total momentum of neutrinos in the event without correcting for these possibilities. As each leptonic decay of a W boson produces a neutrino, it is expected that events analysed in this thesis will have at least two neutrinos contributing to their $\vec{E}_{\mathrm{T}}^{\mathrm{miss}}$.

The momentum balance is calculated separately for hard objects arising from the hard-scatter primary vertex and for soft objects arising from the hard scatter but not the hard objects themselves, for example arising from activity in the underlying event. The x component of the missing momentum is then calculated as

$$E_x^{\mathrm{miss}} = 0 - \sum_{i \in \{\mathrm{hard}\}} p_{x,i} - \sum_{j \in \{\mathrm{soft}\}} p_{x,j}, \tag{3.1}$$

with the y component calculated similarly, where $\{\mathrm{hard}\}$ and $\{\mathrm{soft}\}$ are respectively the sets of hard and soft objects associated with the hard scatter. The missing transverse momentum vector is then defined to be $\vec{E}_{\mathrm{T}}^{\mathrm{miss}} \equiv (E_x^{\mathrm{miss}}, E_y^{\mathrm{miss}})$, and its magnitude is denoted $E_{\mathrm{T}}^{\mathrm{miss}} \equiv |\vec{E}_{\mathrm{T}}^{\mathrm{miss}}|$.

The specific hard objects included in the calculation depend on the selection made by the analysis, and must be chosen after removing overlapping objects (for example when an electron is additionally reconstructed as a jet the jet should be excluded). The hard objects are fully calibrated before performing the $E_{\mathrm{T}}^{\mathrm{miss}}$ reconstruction. The soft term includes only the momentum contributions from good quality soft ID tracks associated to the hard-scatter primary vertex. Previously topo-clusters not associated with a hard object were also used in the soft term calculation, but this results in an unacceptably high residual dependence on pileup activity.

The performance of the $E_{\mathrm{T}}^{\mathrm{miss}}$ reconstruction was checked in 2015 data by studying $W \to e\nu_e$ and $W \to \mu\nu_\mu$ events where the momentum in the hard scatter can be estimated by using the W boson mass as a constraint. Figure 3.14a shows that the deviation from linearity in the $E_{\mathrm{T}}^{\mathrm{miss}}$ response decreases at higher momentum imbalances where the response is dominated by the hard objects. The simulated $t\bar{t}$ sample shows a less linear response than the $W \to \ell\nu_\ell$ samples due to the larger event-by-event fluctuations in pileup activity on account of the higher jet multiplicity [61].

The performance of the soft term was checked using a sample of $Z \to \mu^+\mu^-$ events for which

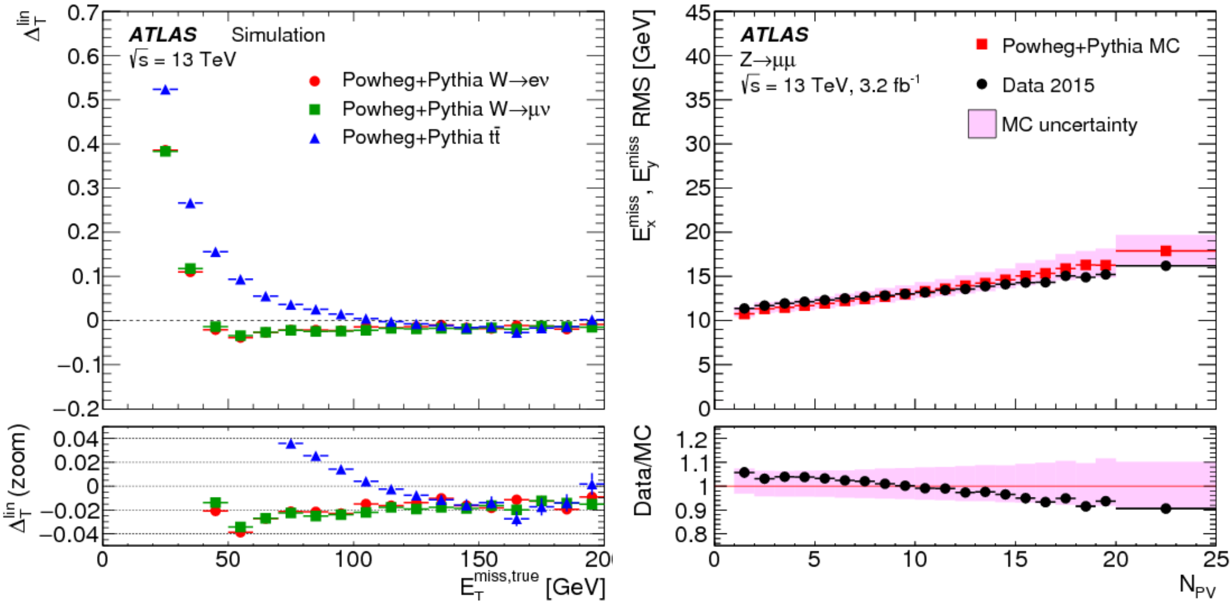

(a) nonlinearity **(b)** resolution

Figure 3.14: Performance of the $E_{\mathrm{T}}^{\mathrm{miss}}$ reconstruction in data collected in 2015. In (a) the deviation from linearity of the $E_{\mathrm{T}}^{\mathrm{miss}}$ response, $\Delta_{\mathrm{T}}^{\mathrm{lin}}$, is shown for several simulated samples, and in (b) the resolution of the reconstructed $E_{\mathrm{T}}^{\mathrm{miss}}$ is measured as a function of the number of primary vertices. Taken from Ref. [61].

there should be no genuine $E_{\mathrm{T}}^{\mathrm{miss}}$ since there are no neutrinos from the hard process. Even in the absence of invisible particles, there is an observation bias towards nonzero values of $E_{\mathrm{T}}^{\mathrm{miss}}$ because the momentum of each object contributing to the calculation is measured with a finite resolution, making exact balancing statistically unlikely. The resolution of the $E_{\mathrm{T}}^{\mathrm{miss}}$ was determined using simulated samples and splitting the reconstructed $E_{\mathrm{T}}^{\mathrm{miss}}$ into projections onto components parallel and perpendicular to the true Z momentum. A measure of the combined x and y resolution is shown in Fig. 3.14b as a function of the number of primary vertices in a $Z \to \mu^+\mu^-$ sample with no true $E_{\mathrm{T}}^{\mathrm{miss}}$, demonstrating a small residual dependence on pileup.

Chapter 4

Analysis strategy

This chapter outlines common methodological elements of the measurements presented in Chapters 5 and 6. The event selection is described in §§4.1-4.3, followed by a description of the classification of signal events in §4.4, and a discussion of the treatment of measurement uncertainty in §§4.5-4.7.

4.1 The dilepton sample

Since the decay of top quarks via the process $t \to Wb$ occurs with a branching fraction of close to unity, the possible decays of the tW system can be enumerated by considering the possible decays of the two W bosons. Each of the W bosons can decay either leptonically into a charged lepton and its corresponding neutrino, or hadronically into jets. Therefore at first order we can expect the decay products of the tW system to include two energetic leptons (*dilepton*), one energetic lepton associated with jets (*lepton+jets*, also called *single lepton*), or many jets without an energetic lepton (*fully hadronic*). The relative rate of these three classes of decay channels is shown in Fig. 4.1.

Placing a requirement on the jet and lepton multiplicity when selecting events allows for the three decay channels to be separated into disjoint samples. Each of these samples has a different profile of background contributions according to the other SM processes resulting in the same event contents, and on the performance of reconstruction algorithms for the different events. Choosing a sample with which to study tW production then amounts to finding a compromise between the branching fraction of the corresponding decay channel and the signal to background ratio. The fully

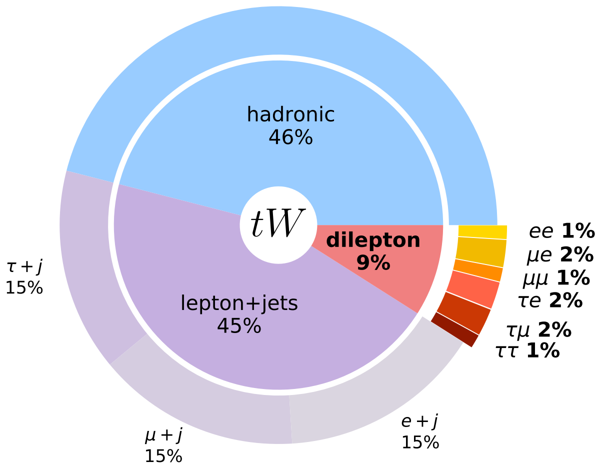

Figure 4.1: Branching fraction of a pair of W bosons arising from the decay of tW or $t\bar{t}$. Decays are classified by the number of charged leptons.

hadronic sample accounts for 46 % of the signal events, but is contaminated with contributions from the many SM multi-jet processes that dominate at hadron colliders. On the theoretical side there are additional challenges in modelling the signal and backgrounds since high jet multiplicities require fixed-order calculations at correspondingly higher orders, and predictions are more sensitive to the matching of additional radiation between ME and PS generators. Experimentally, the details of jet reconstruction and good performance in reconstructing jets in the forward regions of the detector become more important to identify signal events. The single lepton sample accounts for almost the same fraction of the signal as the fully hadronic, but has a smaller though still challenging background profile. The dilepton sample accounts for less than 9 % of the total signal, but has the advantage that comparatively fewer SM processes result in a pair of energetic leptons, and these leptons are well reconstructed by ATLAS. This allows measurements to be viable with smaller samples than would be necessary for the other decay channels. While in the longer term ATLAS will measure tW production in all three of the decay channels, the measurements in this thesis use the dilepton sample.

As shown in Fig. 4.1 over half of the dilepton fraction is accounted for by channels where one of the leptons is a tau lepton. The massive tau leptons have a short lifetime and decay with a branching ratio of about 17 % per flavour into an electron or muon with a pair of neutrinos, with the remaining decays dominated by hadronic processes. Hadronic tau decays occur against a large background of processes that produce jets, and are a challenging signal to extract. In this thesis, events containing hadronic tau decays are excluded from the dilepton sample definition, while events with leptonic tau decays are retained. Unless otherwise stated, for the remainder of this thesis the word *lepton* is used as a shorthand for electrons and muons only.

For each charged lepton in the decay products of the tW system, there will be a corresponding ν_e or ν_μ neutrino (with an additional ν_τ in the case of a decay via a tau lepton). To reconstruct the tW system kinematically from its decay products there must be at least as many kinematic constraints from measurements as there are unknown quantities. With a single neutrino in the event this condition is satisfied, but with two or more neutrinos escaping the detector and a measurement only of the sum of their transverse momenta, it is not possible to fully reconstruct the W bosons. There is an additional ambiguity in assigning the observed leptons to either the top quark or the original W boson. Existing techniques to resolve these ambiguities in the reconstruction of under-constrained systems, for example by assigning a likelihood to each possible assignment of decay products to mother particles, have not yet been applied to tW cross-section measurements at ATLAS.

4.1.1 Composition of the dilepton sample

Within the dilepton sample subdivisions are made based on the multiplicity of selected jets and b-tagged jets to assist in separating the contributions from different sources of dilepton events. These regions are labelled according to the scheme njmb where n is the total jet multiplicity and m is the number of the jets that are b-tagged (such that $n \geq m$). To first order, based on the Feynman diagrams in §1.5.2 and §1.5.3 and Fig. 4.2 the `1j1b` region would contain most of the tW events, the `2j2b` region would consist almost entirely of $t\bar{t}$ background, the `2j1b` region would contain $t\bar{t}$ events where one b-jet was not tagged, and the zero tag regions `1j0b` and `2j0b` would contain small tW and $t\bar{t}$ contributions where b-jets were missed. In reality detector acceptance, reconstruction inefficiencies, and higher-order diagrams complicate the picture as will be shown in the following chapters. Nonetheless the contributions to each region are sufficiently different that this subdivision

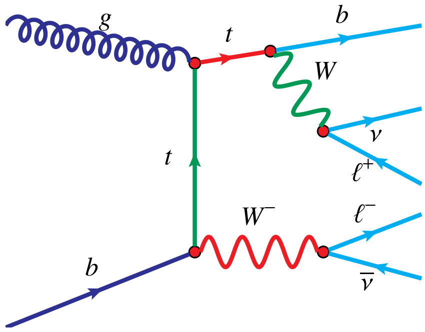

Figure 4.2: Example of tW production at leading order with the subsequent decay to leptons. The b-quark will be reconstructed as a b-tagged jet, and the presence of the two neutrinos will be inferred from the transverse momentum imbalance in the event.

proves to be useful.

The $t\bar{t}$ production cross-section is an order of magnitude larger than the tW cross-section, so it will dominate the final selection. In addition to the tW and $t\bar{t}$ processes, a number of other SM processes are predicted to contribute to the selected events. The production of ZZ, WZ, or WW—collectively called *diboson* production—can result in similar event content to the signal process when the bosons decay leptonically. Leptonic decays of the Z boson result in a pair of oppositely-charged same-flavour[1] leptons with an invariant mass consistent with the Z mass. The $Z + j$ processes, in which a single Z boson is produced in association with jets, can also result in a pair of leptons with jets. The $W + j$ process can contribute slightly when, for example, one of the jets is mistakenly reconstructed as an electron. Processes where $t\bar{t}$ or tW are produced in association with other particles, for example $t\bar{t}W$, can resemble tW events, however these have small cross-sections with respect to tW.

4.2 Modelling signal and background processes

The signal and background processes are modelled using MC simulation and theoretical cross-section calculations to obtain predictions of the contribution of each process to the selected sample, both in terms of an overall yield and the detailed kinematics and properties of these events. The MC samples are produced centrally within the ATLAS Collaboration in consultation with the authors of the MC generator software and other experiments such as CMS to agree on appropriate settings and techniques. In this way the predictions are validated across a variety of regions in phase space as they are shared between a number of different measurements and experiments.

After the MC samples are prepared they are reweighted at the event level and at the object level to improve the agreement with data in a calibration process discussed in §4.6. One notable example of this is the reweighting to correct for pileup conditions. Because the MC samples to be used in association with a particular period of data-taking need to be produced before the data have been collected, the actual $\langle\mu\rangle$ distribution is not known. Instead a best guess of the distribution is used to simulate events and once the actual distribution is known the MC sample is reweighted to match it.

Pileup is modelled in all samples by generating soft QCD events with PYTHIA 8.186 [199] using a set of tuned parameters developed by ATLAS called the A2 [75] tune, and the MSTW2008LO [179] PDF set. These events are then randomly overlaid on top of the hard processes in accordance with the expected $\langle\mu\rangle$ distribution, such that the pileup response is included in the detector simulation

[1] The decay $Z \to \tau^+\tau^-$ where one tau lepton subsequently decays into an electron and the other to a muon is possible. Due to the small branching fraction of a pair of tau leptons to this specific combination of decay products, this is a smaller contribution to the selected events compared to $Z \to e^+e^-$ and $Z \to \mu^+\mu^-$.

Table 4.1: MC generator configurations for simulating the nominal samples for each process.

Process	Matrix element	Parton showering
tW (DR)	POWHEG-BOX v1, CT10	PYTHIA6, CTEQ6L1, P2012
$t\bar{t}$	POWHEG-BOX v2, CT10	PYTHIA6, CTEQ6L1, P2012
$W + j$	POWHEG-BOX v2, CT10	PYTHIA8, CTEQ6L1, AZNLO
$Z + j$	SHERPA, COMIX, OPENLOOPS, CT10/NNPDF	
diboson	SHERPA, COMIX, OPENLOOPS, CT10	
pileup	PYTHIA8, MSTW, A2	

and reconstruction. This is important to be able to accurately model the energy resolution of the calorimeters, and the calculation of whole-event quantities like $E_{\mathrm{T}}^{\mathrm{miss}}$.

To ensure consistency, all MC simulation and fixed-order calculations use the same value of 172.5 GeV for the mass of the top quark, and the value 0.1080 for the branching fraction of the W boson to each lepton flavour. The properties of bottom and charm hadron decays were simulated using EVTGEN 1.2.0 [173] except in the case of samples generated with SHERPA which implements these calculations internally. The MC generators discussed in this section are summarised in Table 4.1.

4.2.1 tW signal and $t\bar{t}$ background

The matrix element calculation was performed with POWHEG-BOX [185, 143, 6], using v1 for tW [193] and v2 for $t\bar{t}$ [103] with the CT10 [171] PDF set. The tW process was defined by using the DR scheme to remove contributions from doubly-resonant diagrams (see §1.5.4). Spin correlations were preserved in the calculation. The parton shower, fragmentation, and underlying event were simulated using PYTHIA 6.428 [198] with the CTEQ6L1 [190] PDF set and the Perugia 2012 (P2012) [200] tune.

The 13 TeV tW theoretical cross-section prediction was taken from Ref. [168] to be

$$\sigma_{\text{theory}}(tW) = 71.7 \ \pm 1.8\,(\text{scale}) \ \pm 3.4\,(\text{PDF})\ \text{pb}$$

where the first uncertainty models the sensitivity of the calculation to the choice of renormalisation and factorisation scales and the second covers the uncertainty of the MSTW PDF sets. The nominal value of the factorisation and renormalisation scales was set to m_t and they were independently varied up and down by a factor of 2 to estimate the uncertainty. The calculation was performed at NLO incorporating NNLL soft-gluon corrections (see §1.5.3).

The 13 TeV $t\bar{t}$ cross-section calculation was performed using `Top++ 2.0` [123] to NNLO accuracy including resummation to NNLL (see §1.5.2), and set to

$$\sigma_{\text{theory}}(t\bar{t}) = 831.8 \ {}^{+19.8}_{-29.2}\,(\text{scale}) \ \pm 35.1\,(\text{PDF} + \alpha_s)\ \text{pb}.$$

The first uncertainty covers the factorisation and renormalisation scale choices and the second was determined by following the PDF4LHC prescription, using the 68 % confidence level MSTW NNLO, CT10 NNLO, and NNPDF 2.3 5F FFN [86] sets and varying the choice of α_s. The factorisation and renormalisation scales were set to $\sqrt{m_t^2 + {p_{\mathrm{T}}}^2}$ where p_{T} is the transverse momentum of the top quark before radiation [39]. In the POWHEG-BOX formalism, additional radiation is split into singular and non-singular parts using a parameter called h_{damp}, which affects the modelling of high-p_{T} radiation; this term was set to m_t.

4.2.2 Prompt bosonic backgrounds

$Z + j$ MC samples were generated with SHERPA 2.1.1 [148] for the measurement in Chapter 5 and SHERPA 2.2.1 [148] for Chapter 6 [55]. For the matrix element generation, diagrams with up to 2 additional final state partons at NLO and up to 4 additional partons at leading order (LO) were calculated using COMIX [149] and OPENLOOPS [104]. The parton showering was done with SHERPA [195] using the ME+PS@NLO [158] prescription for matching additional radiation to the ME generator predictions. The CT10 PDF set was used in the total cross-section analysis, while the NNPDF NNLO PDF set was used in the differential analysis. In both cases a generator-level cutoff $M_{\ell\ell} > 40$ GeV was applied. The NNLO cross-section prediction was used for normalisation.

Diboson MC samples include events with 4 charged leptons, 3 charged leptons and a neutrino, or 2 charged leptons and 2 neutrinos in the final state. These were simulated using SHERPA 2.1.1 for both analyses, including all diagrams with 4 electroweak vertices. For final states with up to one additional parton, NLO corrections were used, while for other final states with up to 3 additional partons COMIX and OPENLOOPS LO calculations were used. The parton showering was done with SHERPA using the ME+PS@NLO prescription for matching. The CT10 PDF set was used, and the NLO cross-section provided by the generator was used for normalisation.

For the measurement in Chapter 5 the very small contribution from $W + j$ events was also considered. These MC samples were generated using POWHEG-BOX v2 and the CT10 PDF set. Parton showering was done using PYTHIA8 and the CTEQ6L1 PDF set with the AZNLO [49] tune [55].

4.2.3 Non-prompt and fake background

Events where an additional non-prompt lepton produced in secondary interactions is selected, or where a lepton is reconstructed where there was not a true lepton in the final state of the hard process also contribute to the selected samples. These contributions depend on the details of the detector response and the reconstruction algorithms and are typically not well modelled by MC samples. Where these events are expected to contribute significantly, for example when measuring tW production in the single lepton decay channel, their contribution must be estimated using data-driven techniques. In the analyses presented here, the selection requirements exclude most of these events and, given the very small size of the residual contribution, the estimated contribution from MC samples is used directly. For each of the MC samples described in the previous sections, any events where a selected lepton is not a prompt lepton from a W decay[2] are excluded from the background sample and instead placed in a separate non-prompt background sample. The composition of this non-prompt sample is shown in Fig. 4.3. The largest source of MC non-prompt events is the $t\bar{t}$ sample in all regions, with small contributions from the $Z + j$ and tW samples in the `1j1b` region, and smaller contributions from the tW and $W + j$ samples in the `2j1b` region.

4.3 Suppressing minor backgrounds

Whereas tW and $t\bar{t}$ events are typically similar in terms of event content and kinematic distributions, the other sources of dilepton events are fundamentally different physical processes and tend to differ

[2] Simulated MC events contain records of the true ancestry of each particle in the event, recorded during the MC generator calculations, allowing the promptness of selected particles to be directly checked.

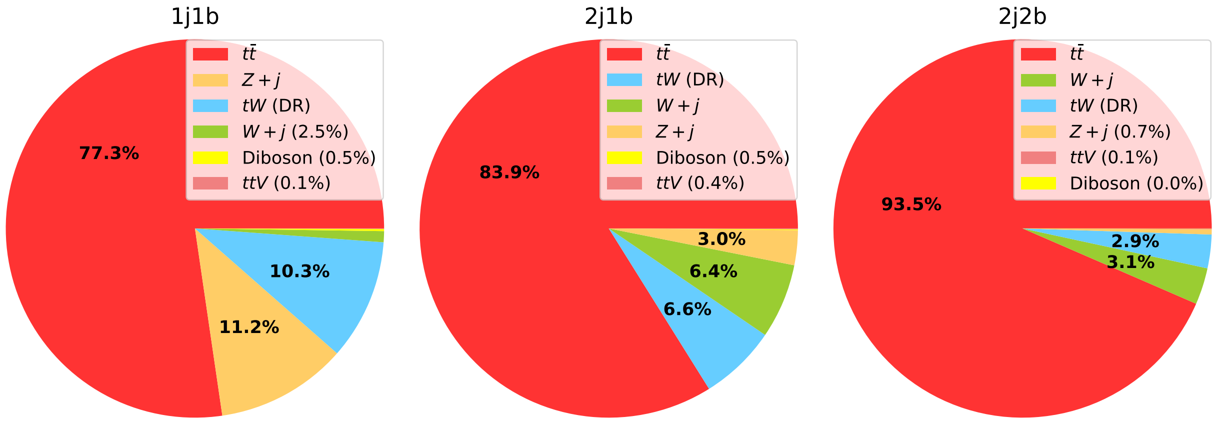

Figure 4.3: Composition of the MC non-prompt sample in each region, broken down according to the original MC sample from which the events were collected.

in several observables from the top processes. A number of requirements are placed on the lepton invariant mass $M_{\ell\ell}$ and the missing transverse momentum $E_{\mathrm{T}}^{\mathrm{miss}}$ in order to reduce the contributions from $Z+j$, diboson and sources of fake or non-prompt leptons. The composition of these backgrounds is significantly different in events where the two leptons are the same flavour (ee or $\mu\mu$) compared to events where the leptons have different flavour ($e\mu$) so the requirements are chosen differently for these two classes of events.

In the case of $Z+j$ events the two leptons are expected to have the same flavour and an invariant mass consistent with the Z boson mass, so any same-flavour events with $81\,\mathrm{GeV} < M_{\ell\ell} < 101\,\mathrm{GeV}$ are excluded. Signal events are expected to have $E_{\mathrm{T}}^{\mathrm{miss}}$ from the invisible neutrino, so same-flavour events with $E_{\mathrm{T}}^{\mathrm{miss}} < 40\,\mathrm{GeV}$ are excluded, and opposite-flavour events with $E_{\mathrm{T}}^{\mathrm{miss}} < 50\,\mathrm{GeV}$ are excluded with the threshold relaxed to 20 GeV for events with $M_{\ell\ell} > 80\,\mathrm{GeV}$. To exclude a large fraction of the remaining minor background contributions in same-flavour events an additional pair of requirements is imposed: events must have $M_{\ell\ell} > 40\,\mathrm{GeV}$ and must satisfy $4 \times E_{\mathrm{T}}^{\mathrm{miss}} > 5 \times M_{\ell\ell}$ when $M_{\ell\ell} < 101\,\mathrm{GeV}$ and $2 \times M_{\ell\ell} + E_{\mathrm{T}}^{\mathrm{miss}} > 300\,\mathrm{GeV}$ otherwise. The effect of these requirements is shown in Fig. 4.4; the $Z+j$ and diboson background are almost entirely eliminated along with a substantial fraction of the fake and non-prompt events.

The details of these requirements were chosen to optimise the overall significance of the cross-section measurement in Chapter 5. This was done by scanning over a number of threshold values for each requirement and plotting the signal efficiency, $t\bar{t}$ background rejection, and the ratio of the number of signal events to the square root of the number of background events. These configuration that achieved the best simultaneous maximisation of these measures was chosen by eye. The requirements were retained for the measurement in Chapter 6 after checking that small variations did not significantly affect the result.

4.4 Classifying tW events

Due to the similarity in final states of the tW and $t\bar{t}$ processes it is challenging to classify a given event as produced by one or the other. In order to exploit as much information about the event as possible it is necessary to consider multivariate techniques to identify features that characterise the

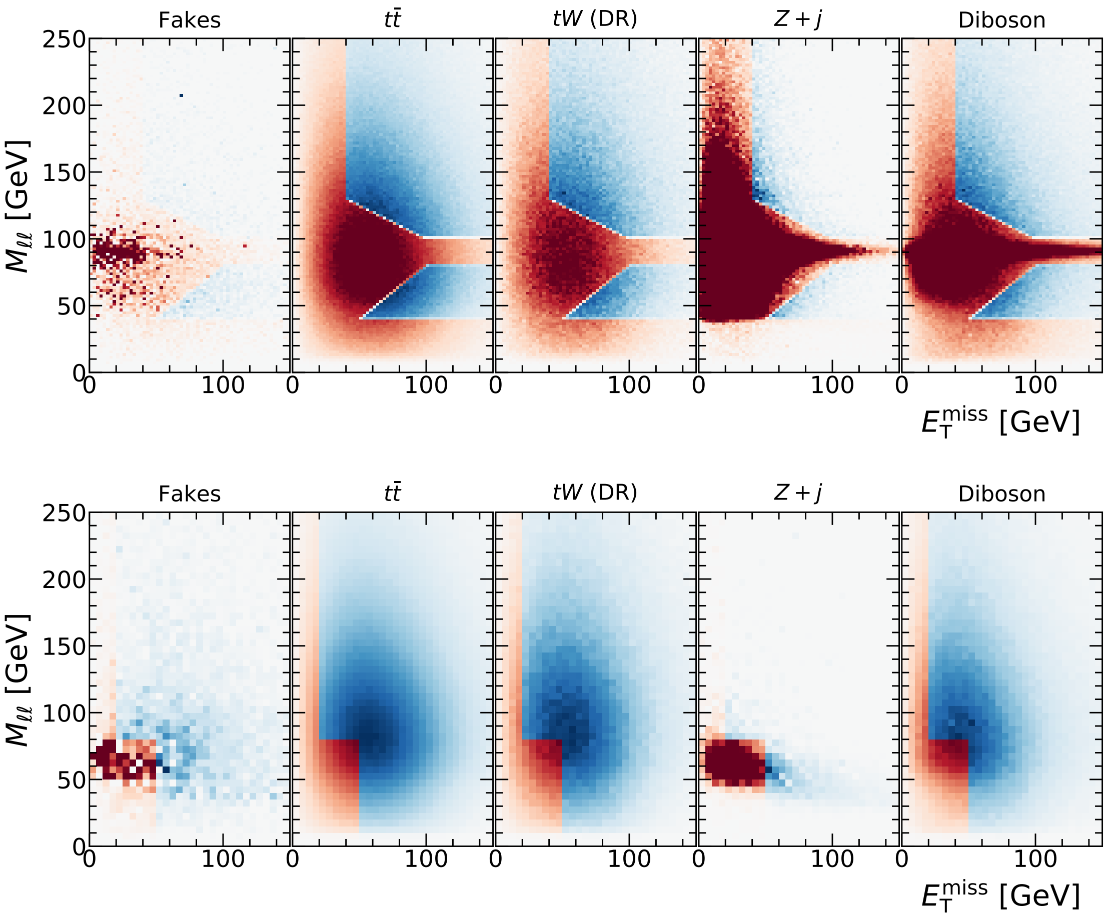

Figure 4.4: The $M_{\ell\ell}$ and $E_{\mathrm{T}}^{\mathrm{miss}}$ requirements applied to each MC sample separately for e^+e^- and $\mu^+\mu^-$ events (top) and $e^{\pm}\mu^{\mp}$ events (bottom). The bins in red are excluded by the requirements, while those in blue are retained. Most of the $Z+j$ and fake events are effectively excluded while the losses in the signal sample are acceptable. Darker areas indicate parts of the parameter space with more events, but the units are arbitrary.

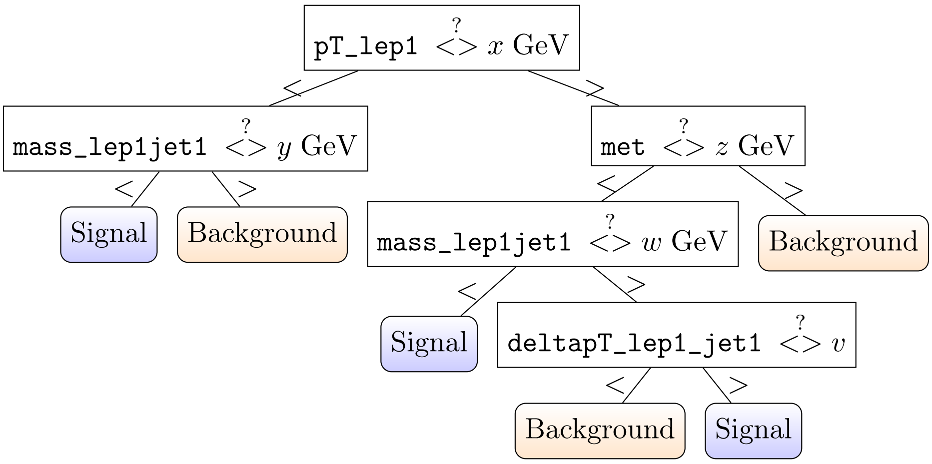

Figure 4.5: Illustration of a single decision tree showing how binary cuts on certain variables can partition events into a number of leaf nodes. The leaf nodes are labelled 'Signal' or 'Background' based on the relative abundance of signal events accumulating at each node during training.

two processes. The specific multivariate tools used in this analysis are boosted decision trees (BDTs) which are comparatively simple with respect to models produced by neural networks while retaining reasonable performance in terms of their classifying power.

4.4.1 Boosted decision trees

A decision tree is a collection of Boolean predicates that filters data events into one of several leaf nodes, each of which is assigned to one of two classes. An example of one such tree is shown in the schematic in Fig. 4.5. The unfilled rectangular nodes in the tree correspond to comparisons of a particular observable of the event with some threshold value, where the outcome of the comparison determines the next node selected. Eventually each event ends up in a leaf node labelled either "Signal" or "Background" which determines the class to which the classifier has assigned the event. *Regression trees* are a generalisation of these classification trees where each node is assigned a value from signal-like to background-like rather than a discrete label.

An alternative way to visualise the tree is to consider the multidimensional phase space spanned by all inputs to the tree (i.e. the set of observables involved in the binary decisions); the boundary between the signal and background classes in this phase space is encoded by the structure of the tree. One such visualisation is illustrated in Fig. 4.6 for a binary classifier using two features. The level of detail in the decision boundary is limited by the complexity and topology of the decision tree. If allowed to incorporate arbitrary complexity then this boundary will be able to divide the phase space perfectly such that there is no contamination of background points in the signal class and vice versa. Such a boundary is highly sensitive to the specific data used for developing the tree and does not necessarily generalise to unseen data.

Rather than using solitary decision trees, the measurements in this thesis employ collections of decision trees called BDTs which classify events based on a weighted average of the classifications provided by their constituent decision trees. In this analysis the BDT implementation in software was based on the software package TMVA [159].

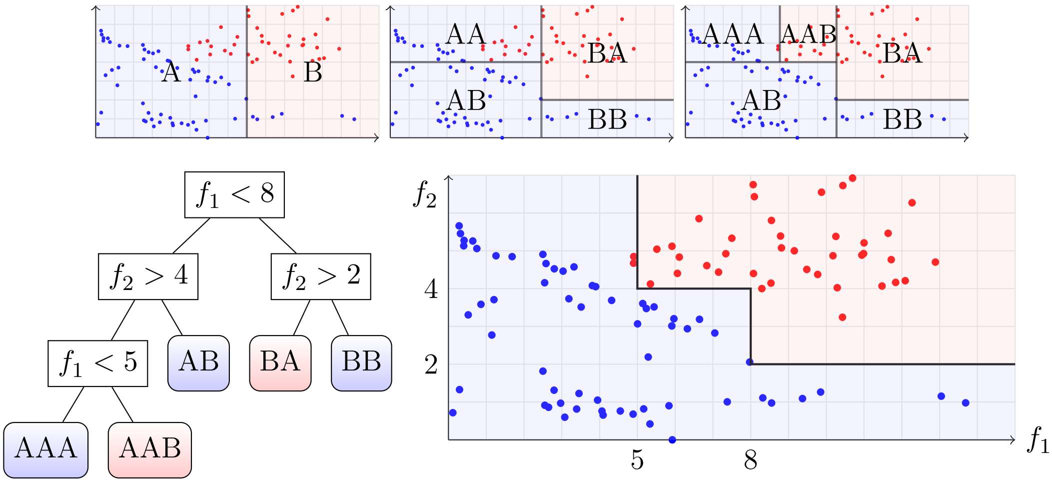

Figure 4.6: Data drawn from two populations (represented by the dots of different colours) are plotted with respect to two features f_1 and f_2. The upper plots show the iterations of a tree growth algorithm. In the first step the phase space is divided into two regions A and B; in the second step each of these regions is further divided into two, and in the final step region AA is again subdivided. The regions are coloured to match the population with the greatest representation in that region. The tree on the lower left encodes the same sequence of binary cuts. The lower right plot shows the final classifier and its decision boundary.

4.4.2 Training and tuning BDTs

Training a BDT involves growing an optimal ensemble of trees using a sample of events, the *training sample*, that are labelled, meaning that each event is known to be either signal or background (because it was produced in an MC simulation). TMVA implements a number of techniques to stabilise the training process against statistical fluctuations and to improve the performance of the final classifier. Tree growth proceeds iteratively, with each iteration introducing a new binary subdivision of an existing region, as illustrated in the upper part of Fig. 4.6. If the number of training points in any of the new regions is below a certain threshold then that region becomes a leaf node, otherwise it is further divided in later iterations. Once the tree growth stops, each leaf node is labelled as either signal or background according to the dominant population of training events contained in the corresponding region of phase space, or assigned a value related to their relative abundance in the case of regression trees.

The training procedure outlined above leaves a number of tuning parameters unspecified. The optimal choice of parameters is the one that affords the best separation of signal from background events without sacrificing the ability of the classifier to generalise beyond the training sample to unseen events. The optimisation scheme and validation will be discussed in §4.4.4. The following is a description of the BDT model parameters and their effects, using the context and nomenclature of TMVA specifically.

Tree topology and growth

The depth of a decision tree is determined by the maximum number of Boolean decisions needed to describe the path of any event from the tree's root to a leaf node; for example the tree in Fig. 4.5 has a depth of 4. A deeper tree can capture information about the interrelation of a larger number of variables. In practice when using boosting the best outcomes are achieved by forcing the individual trees to be relatively shallow, with a depth no greater than about 5. The termination of the growth is controlled by choosing a threshold number of events below which nodes will no longer be subdivided. It should be chosen so as to avoid very small nodes since these are overly sensitive to fluctuations in training data. In TMVA these parameters correspond to the settings `MaxDepth` and `MinNodeSize` (which is expressed as a percentage of the total number of training events).

TMVA also provides parameters which primarily tune the trade-off between computational resources demanded by the training procedure and the performance of the resulting BDT. One such parameter controls the granularity into which the interval for each event variable is divided when attempting to find the optimal cut threshold during tree growth. In the illustration in Fig. 4.6 this corresponds to the granularity of the grid along which the boundary is constrained; a finer grid allows a better boundary to be found but requires more possible boundaries to be considered at each iteration.

Gradient boosting algorithm

Boosting assembles a collection of decision trees into an ensemble that performs much better than the individual trees. In contrast to classifiers such as random forests [153] which are built from the consensus of many large, independent decision trees, boosted classifiers are built iteratively using weak classifiers. A large number of small decision trees are built into an ensemble by minimising at each iteration a differentiable loss function that quantifies the overall error of the classifier. In this thesis TMVA's implementation of gradient boosting is used, which is based on the approach of Ref. [142].

An event can be described as a point $\vec{x} \equiv (x_1, x_2, \ldots, x_d) \in X$ where each x_k is the value of one of d physical observables. An event can be labelled using the set of classes $Y \equiv \{-1, +1\}$ whose elements represent background and signal respectively. The training sample is a set $\{\vec{x}_i, y_i\}_1^N$ where each simulated event $\vec{x}_i$ has a known label $y_i \in Y$ since it is known by construction whether the event came from a signal or a background sample. Regression trees and BDTs are effectively maps $F : X \to [-1, +1]$ that assign each event a continuous response on the interval from background-like to signal-like. Regression trees can be considered as parametrised classifiers defined by a set of parameters encoding the sequence and location of the subdivisions in the parameter space X along with the value of each leaf node. The optimal BDT is the classifier that combines a collection of these regression trees chosen in such a way that the expected value of the binomial log-likelihood loss function

$$L(F(\vec{x}), y) = \ln\left(1 + e^{-2F(\vec{x})y}\right)$$

is minimised over the training dataset. The goal of boosting is to find a BDT $\hat{F}$ that approximates this optimal classifier.

The BDT training is iterative, starting with an initial classifier F_0 made from a single regression tree that attempts to minimise the average loss over the N training events. For each subsequent iteration $m \geq 1$, an additional regression tree $f(\vec{x})$ is added to the classifier so that the updated BDT response is then $F_m(\vec{x}) = F_{m-1}(\vec{x}) + f(\vec{x})$. The optimal choice of $f(x)$ would be equal to the difference between F_m and the true label. Using this observation, the gradient of the loss function is evaluated using the classifier from the previous iteration, $\bar{y}_i \equiv -\partial L(F_{m-1}(\vec{x}_i), y_i)/\partial F_{m-1}(\vec{x}_i)$, and the parameters of $f(\vec{x})$ are then chosen to minimise $\sum_{i=1}^{N} (\bar{y}_i - f(\vec{x}))^2$. This effectively trains the

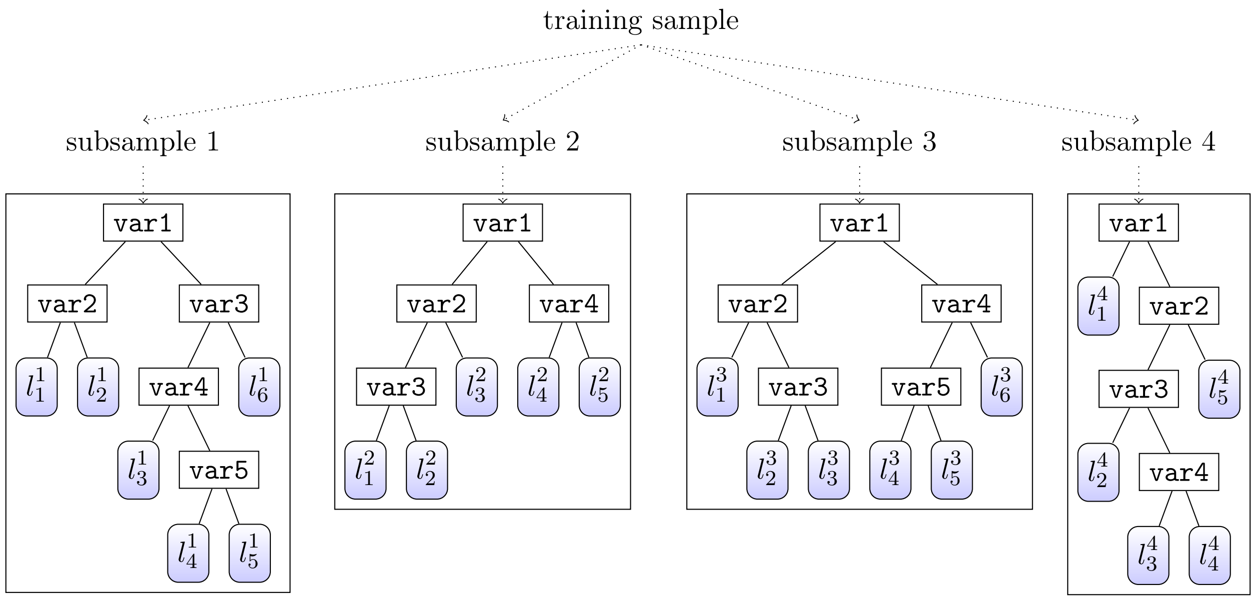

Figure 4.7: Illustration of a BDT consisting of four individual trees. Each tree was trained using a subsample (produced by the bagging and boosting procedure) of the original training sample. The value l_i^j assigned to a given event by the jth tree is combined with the responses from the other trees to arrive at the overall classifier response.

new tree on the residuals of the previous classifier as measured by the loss function, following its steepest gradient. Instead of adding $f(x)$ at each iteration, a global learning rate (or *shrinkage*) parameter $0 < \nu < 1$ can be used to reduce the rate of descent by instead adding $\nu f(x)$. Smaller choices of the shrinkage parameter mean that the BDT follows the gradient of the loss function more closely during training, but requires a larger number of trees to produce a good result.

Terminating at the Mth iteration, the final BDT response $F_M(\vec{x})$ is then a function that continuously classifies a given event as signal-like or background-like, minimising the loss over the training data. Keeping M relatively small can help to avoid over-fitting the training data, but this approach is not in general as effective as reducing the shrinkage [142]. The response of the BDT constructed in this procedure is effectively a weighted sum of the responses of its constituent trees, with the weights absorbed into the values assigned to each leaf node of the trees.

Minimising statistical effects

Bagging is a procedure for stabilising the BDT by averaging out statistical fluctuations that can affect the structure and performance of the trees. The training sample is treated as a probability distribution and training events are chosen for each iteration randomly from the sample (with replacement) to create subsamples each with a fixed fraction of the total number of events in the full training sample. This means that each decision tree in the BDT's ensemble is trained on a different subset of training data as illustrated in Fig. 4.7.

4.4.3 Quantifying the similarity of distributions

The Kolmogorov-Smirnov (KS) test is a useful statistical test for comparing the compatibility of two samples [189, pp. 623-28]. Cumulative distributions $F_N(x)$ are constructed for each sample such that

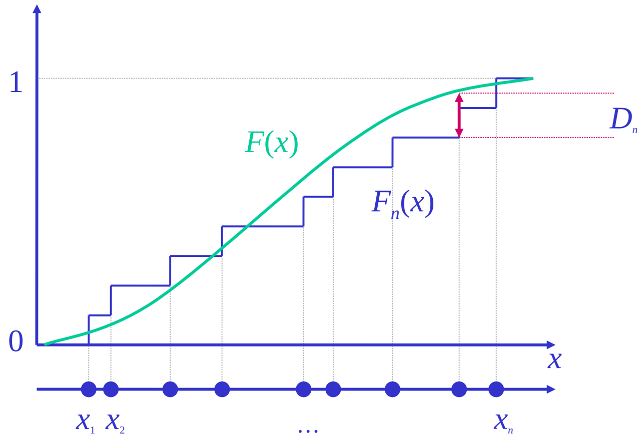

Figure 4.8: The Kolmogorov-Smirnov test statistic, D. The input data (x_i) are converted to a cumulative distribution function F_N which is then compared to a second cumulative distribution. In this illustration this is a known probability distribution function (pdf) $F(x)$ but it can instead be a second sample, as in the applications in this thesis. D is the largest vertical difference between the two curves. Courtesy of Luca Lista.

the value of $F_N(x)$ is the fraction of the data that falls to the left of x,

$$F_N(x) \equiv \frac{1}{N} \sum_{i=1}^{N} I_{(-\infty,x]}(X_i)$$

where N is the size of the sample and $I_{(-\infty,a]}(z)$ is equal to unity when $z \leq a$ and zero otherwise. The two-sample KS test statistic D is then the largest vertical difference between two such cumulative distribution functions,

$$D \equiv \max_{-\infty<x<\infty} |F_{1,N_1}(x) - F_{2,N_2}(x)|.$$

The construction of this test statistic is illustrated in Fig. 4.8.

A statistical test is then constructed with the null hypothesis H_0 being that the two samples are drawn from the same distribution. Under H_0, $\lim_{N\to\infty} D = 0$ up to some technical details. The significance level of an observed value D is approximated by

$$P(K \leq x) = 1 - 2\sum_{i=1}^{\infty} (-1)^{i-1} e^{-i^2x^2}, \qquad \text{where} \quad K = \frac{N_1 N_2}{N_1 + N_2} D,$$

and N_1 and N_2 are the number of data points in the samples. The accuracy of this approximation improves asymptotically as the sample sizes increase. An observed significance of α excludes H_0 at the $(1-\alpha)$ confidence level. Note that this is only true in general for continuous distributions; for binned data the test can still be performed, however the significance will be overestimated. If the binning is fine compared to features of the sample distributions, this overestimation is not severe. The KS test statistic is sensitive both to translation and to differences in shape.

When classifying events based on observable features, it is useful to have a measure of the power of each feature at separating signal events from background events. The *separation power*, S, is used to quantify differences in the distribution of that feature in the signal sample compared to the same feature in the background sample. For two quantities x and y identically binned in N bins with widths b_i, this is defined by

$$S^2 = \frac{1}{2} \sum_{i=1}^{N} \frac{(x_i - y_i)^2}{x_i + y_i} b_i.$$

The separation power is sensitive to differences in any individual bin of the distributions.

4.4.4 Assessing classifier performance and characteristics

The ideal classifier efficiently separates unseen signal events from background events. Overtraining occurs when the classifier has embedded detailed information about the specific events in the training sample in order to achieve good (or even perfect) classification of events in a manner that does not generalise to unseen events. To guard against this, one part of the initial sample of simulated events is reserved to be the testing sample, while the remainder is used for the training sample of §4.4.2. After the BDT is trained, distributions of its response are prepared for the training sample and for the testing sample which has remained unseen by the BDT until this point. A KS test is performed to determine whether the two BDT responses are consistent with having been drawn from the same distribution.

More sophisticated cross-validation techniques are available, such as dividing the training data into several subsets and training multiple BDTs on different combinations of these subsets. These techniques allow for a less biased assessment of overtraining, however they rely on having a large enough training set such that the individual subsets are not overly sensitive to statistical effects.

To evaluate the ability of the BDT to separate signal from background, the rate of signal acceptance is plotted as a function of background rejection to form the so-called receiver operating characteristic (ROC) curve. The area under this curve ranges from 0.5 for a classifier that performs no better than random chance, up to unity for a perfect classifier. In Fig. 4.9 the performance of two BDTs is compared: one is an optimised BDT making use of several features, the other is a toy BDT using only the single feature with the greatest separation power. The single-feature BDT, with strictly less information available to it, produces a ROC curve that encloses an area of 0.64, while the ROC curve of the optimised BDT encloses a slightly larger area of 0.69. The ROC curve of the simplistic BDT is closer to the worst-performance diagonal line, meaning that for any given signal efficiency it will provide poorer background rejection, and so it has worse classification performance. The two ROC curves are so similar in this example because the single most-separating feature explains the majority of the separation power of the optimal classifier in the `1j1b` region as will be shown in Chapter 5.

It is also useful to consider the separation power of the BDT response to simulated samples of signal and background events, and compare it to the separation power of the individual observables that are used by the classifier. This also provides a crude indication of the relative contribution of each observable to the classifier, although this measure does not take into account the actual structure of the BDT.

4.4.5 Features for classification

Reconstructed events contain observables specific to the event itself, such as $E_{\mathrm{T}}^{\mathrm{miss}}$, along with observables related to each selected object in the event. In principle all available information could be taken without further processing and used as input features to train a BDT. This would not make efficient use of the training dataset however, both because better separation sometimes requires combinations of features that are not easily reproduced in the structure of BDTs, and because many of the features are typically correlated, introducing redundancies in the parameter space that the BDT explores. This section introduces all features used in this thesis.

The kinematic observables are constructed using the momentum four-vectors of selected leptons and jets in the event and a four-vector representing the missing momentum constructed by using

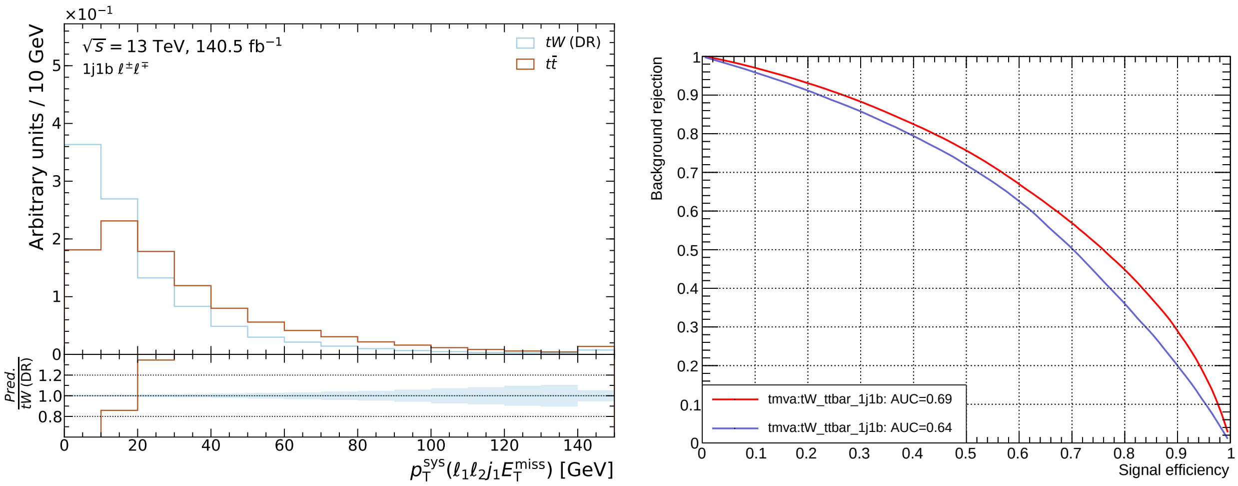

(a) best feature

(b) ROC curve

Figure 4.9: Illustration of the performance of a BDT in the `1j1b` region. In (a) the shape of the single most separating feature in the region is shown for the signal tW MC sample (blue) and background $t\bar{t}$ MC sample (brown). In (b) two ROC curves are shown: in blue is the performance of a toy BDT that is trained using only the single feature from (a), and in red is the performance of an optimised BDT similar to what is used in Chapter 5 combining information from several observables.

$\vec{E}_{\mathrm{T}}^{\mathrm{miss}}$ and treating the mass and the z component as zero. The basic kinematic quantities p_{T}, η, ϕ, E, and m for each four-vector are used. The quantity ΣE_{T} is defined as the scalar sum of the transverse momenta of the objects that contribute to the $E_{\mathrm{T}}^{\mathrm{miss}}$ calculation. For two systems of objects, s_1 and s_2, formed by taking the vectorial sum of a number of individual four-vectors, the following observables are defined:

- $\Delta R(s_1, s_2)$, the η–ϕ separation,
- $\Delta p_{\mathrm{T}}(s_1, s_2)$, the difference in the scalar p_{T} values of the systems, and
- $\Delta\phi(s_1, s_2)$, the ϕ separation.

For a set of four-vectors $o_1 \ldots o_n$, the following observables are also defined:

- $p_{\mathrm{T}}(o_1 \ldots o_n)$, the p_{T} of the system formed by taking the vectorial sum of the four-vectors,
- $H_{\mathrm{T}}(o_1 \ldots o_n)$, the scalar sum of the transverse momenta,
- $\sigma(p_{\mathrm{T}})(o_1 \ldots o_n)$, the ratio of $p_{\mathrm{T}}(o_1 \ldots o_n)$ to $(H_{\mathrm{T}}(o_1 \ldots o_n) + \Sigma E_{\mathrm{T}})$,
- $m(o_1 \ldots o_n)$, the invariant mass of the system,
- $E(o_1 \ldots o_n)$, the energy of the system,
- $E/m(o_1 \ldots o_n)$, the ratio of energy to mass,
- $C(o_1 \ldots o_n)$, the centrality defined by the ratio of the scalar sum of p_{T} to the sum of energy.

The transverse mass, $m_{\mathrm{T}}(o_1, o_2) \equiv [2p_{\mathrm{T}}(o_1)p_{\mathrm{T}}(o_2)(1 - \cos\Delta\phi(o_1, o_2))]^{1/2}$, is defined for a pair of objects o_1 and o_2. Objects are written using the symbol ℓ_i for leptons and j_i for jets where the subscript indicates the ranking according to p_{T} such that ℓ_1 is the leading lepton and ℓ_2 is the sub-leading lepton in the event. Where it is unambiguous the subscripts are dropped. In selection regions where all jets are b-tagged, the symbol b is instead used for the jets.

4.4.6 Optimising BDT parameters

The most important BDT parameters to optimise are the shrinkage, the number of trees, and the maximum depth of individual trees. These must be jointly optimised, since for example reducing the shrinkage will require an increase in the number of trees as discussed in §4.4.2. The optimal configuration depends also on the selection of input features, which itself depends on the configuration of the BDT. This joint optimisation problem is handled by an iterative procedure as follows.

Firstly the full list of observables must be pruned to remove correlated observables. The full list is ranked by separation power as evaluated using the training dataset, and the covariance matrix is computed. If any pair of observables has a linear correlation greater than 80 %, the observable with the lower separation power is removed. This removes the most highly correlated features even if they are traditionally preferred observables[3], since the information that allows them to separate signal from background is retained in the remaining quantities.

Next a series of BDTs are trained starting with the decorrelated list of input features and successively removing the lowest-ranked features. The list that produces the BDT with the highest separation power is retained. If multiple lists produce BDTs with similar separation powers, the list with the fewest features is preferred. Fixing the list of features, the configuration of the BDT is varied by scanning over a grid of possible values for each parameter. The configuration that produces the BDT with the highest separation power is retained. Throughout this procedure, any BDT for which a KS test comparing its response to the signal sample to its response in the background sample is found to have a significance of less than 0.1 is excluded from consideration on the basis that it is likely overtrained.

The process then repeats starting with the new BDT configuration and the full decorrelated list of features. After several iterations the most highly-ranked input features and the BDT configuration typically stabilise. At this point the optimisation is stopped and the resulting BDT is validated by comparing its response in the testing and training datasets to identify any differences that are not detected by the KS test, and verifying that the response is smooth across its range.

4.5 Treatment of uncertainty

The two sets of measurements presented in this thesis differ in their treatment of uncertainties, as will be discussed in Chapters 5 and 6. Fundamentally, both measurements must infer the composition of observed data in terms of the physical processes responsible for the events. While tools such as the BDT discussed above assist in separating signal from background, there is an irreducible contamination of background events in any selection, and in fact there are expected to be more $t\bar{t}$ events than tW events. Information from MC simulation and theoretical calculations is used to establish the expected contributions from each process under consideration, and this is subsequently used as part of the measurement procedure.

To estimate the uncertainty of the measured cross-sections, it is therefore necessary to consider a wide range of sources related to the experimental measurement of observables, to the theoretical modelling of the physical processes, and to any differences between MC simulated samples and the observed performance of the detector. §4.6 provides an overview of the calibration of relevant parts

[3] For example $E_{\mathrm{T}}^{\mathrm{miss}}$ can be pruned from the input features if a highly-correlated observable such as the p_{T} sum of $\vec{E}_{\mathrm{T}}^{\mathrm{miss}}$ and another object is found to have higher separation power.

of the detector and reconstruction algorithms, and the corresponding sources of uncertainty. §4.7 discusses methods for estimating the impact of modelling uncertainty on the measurements.

4.6 Experimental sources of uncertainty

Sources of uncertainty related to calibration are for the most part handled by breaking the uncertainties into approximately uncorrelated components, and for each component producing simulated samples corresponding to 1σ fluctuations of the uncertainty in the positive and negative directions. This is implemented either as a pair of weight vectors used to reweight the nominal MC sample, or as a pair of independent MC samples produced with different settings. The pair of templates is checked to ensure it brackets the nominal sample, and is then used to describe the effect of the systematic source on both the shape of observables and the normalisation of the sample.

Due to the large number of additional events that need to be simulated for systematic variations implemented as alternative samples, these are usually reconstructed using AtlfastII to save computational resources. To avoid conflating the performance of the fast simulation with each systematic source, AtlfastII versions of the nominal samples are also produced to use in comparisons with these variations such that samples with different reconstruction schemes are never compared when estimating systematic uncertainty.

In many cases the effect of a particular source of uncertainty can either be broken into many components (corresponding to independent calibrations in different parts of the detector, for example) or combined into a single component by conservatively assuming that the components are uncorrelated. In a profile likelihood fit (see §5.3) each of the components can be constrained by other information in the sample and correlations with other sources of uncertainty, and so there is an advantage in pursuing the finest available granularity. On the other hand using too many components reduces the stability of the fit due to the correspondingly large number of nuisance parameters. If the measurement is not particularly sensitive to a given source of uncertainty there will be no appreciable reduction in the overall uncertainty through profiling in any case, so it is in general preferable to opt for a smaller number of components of uncertainty in such cases.

4.6.1 Lepton efficiency and calibration

After the calibration of electrons and muons discussed in §3.4 and §3.5 there is a residual disagreement between the efficiency of lepton identification, isolation, reconstruction, and trigger requirements between MC simulated samples and the observed data. This disagreement is quantified using $Z \to \ell^+\ell^-$ and $J/\psi \to \ell^+\ell^-$ samples and converted into scale factors that are used to reweight the MC events to achieve similar efficiencies to those observed in data. The scale factors have an associated statistical uncertainty due to the finite size of the samples used in the comparison. The lepton scale factors are determined in bins of η and ϕ and it is possible to consider various schemes of correlating the contributions of each bin to model the effect on the analysis. Because the lepton scale factors do not contribute substantially to the overall uncertainty of the measurements in this thesis, they are treated as overall effects. The templates for electron identification, isolation, reconstruction, and trigger efficiencies are produced by varying the scale factors according to this statistical uncertainty to produce one pair of weight vectors for each scale factor. For the muon scale factors, in addition to the effect of finite statistics in the determination of the scale factors, the systematic uncertainty

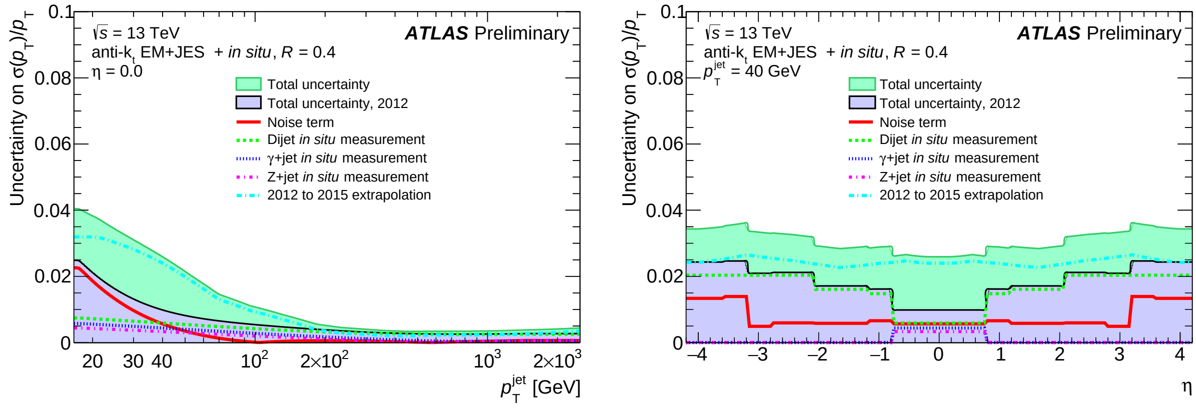

(a) **(b)**

Figure 4.10: Total uncertainty on the JER calibration as a function of (a) jet p_{T} at $\eta = 0$ and (b) η at $p_{\mathrm{T}} = 40\,\mathrm{GeV}$. The uncertainty of the 2012 calibration is shown by the black curve and blue filled region. The additional uncertainty added to conservatively cover the extrapolation to 2015 conditions is shown by the dashed cyan line, and the total uncertainty of the extrapolated calibration is shown by the green filled region. Taken from Ref. [40].

of the method is separately estimated, so that two pairs of weight vectors are produced for each of the identification of low- and high-p_{T} muons, isolation, track-to-vertex association, and trigger requirement scale factors.

The resolution of leptons in MC samples is smeared to match that observed in data, and the lepton energy scale is likewise shifted to match the observed one. The electron energy scale and resolution calibration is performed independently in different parts of the detector and can be considered separately, but here is combined into one pair of systematic variations for fluctuations in the scale and another for the smearing. For muons the ID track smearing, MS track smearing, and the muon momentum scale are each accounted for with systematic variations.

4.6.2 Jet calibration

The JER (see §3.6.1) calibration is performed according to the procedure in Ref. [40]; since this was prepared early in Run 2 of the LHC, it was based on an extrapolation from the 2012 calibration from the end of Run 1 after performing additional checks and adding a conservative uncertainty to cover the extrapolation. The JER in MC samples is smeared to match that in data. For each comparison of different samples (between the 2012 sample and either the 2015 sample or a reference MC sample) during the calibration, the agreement of the resolution is measured by the *smearing factor* defined to be the signed square root of the difference between squares of the resolutions. The total JER uncertainty is estimated by adding these smearing factors. Figure 4.10 shows the JER uncertainty of the original 2012 calibration with the extrapolated calibration overlaid.

The two measurements in this thesis use slightly different calibrations and uncertainties for the JES but the overall approach is the same: they are based on preliminary calibrations using the 2015 dataset, without any extrapolation from earlier calibrations but not yet benefiting from the improvements of the final calibration of Ref. [43]. With respect to that final calibration, those used here have a slightly larger JES uncertainty in high-p_{T} jets and in forward regions of the detector. In particular, an issue related to jet cleaning in the 2015 and 2016 datasets has since been identified that was not known at the time the present results were published. An error in associating topo-clusters to electrons and photons in the transition region of the detector between the barrel and end-cap regions was found to result in the double counting of some energy deposits and a subsequent increase in the tails of the $E_{\mathrm{T}}^{\mathrm{miss}}$ distribution. This is not expected to have any impact on measurements here since hard electrons in that region of the detector are excluded by the event selection requirements, and so do not contribute to the $E_{\mathrm{T}}^{\mathrm{miss}}$ calculation.

A set of almost 80 sources of systematic uncertainty are estimated to describe the total JES uncertainty. About 65 of these are related to the *in-situ* calibration using $Z/\gamma + j$ samples and multi-jet balancing, and these are reduced to a set of 6 orthogonal components by diagonalising their correlation matrix[4]. The other 12 sources cover pileup, the η inter-calibration, leakage into the MS, high-p_{T} jets inaccessible to the *in-situ* scheme, and differences in jet response and flavour composition. In the case of unknown flavour composition of jets in an analysis, it is assumed that half of the jets are initiated by quarks and the other half by gluons, with a 100 % uncertainty on the composition. In total, 18 JES uncertainty components are used in the analyses in addition to the JER term, and these are implemented as pairs of systematic variation samples where their effects are individually fluctuated.

4.6.3 $E_{\mathrm{T}}^{\mathrm{miss}}$ calibration

As discussed in §3.7, the $E_{\mathrm{T}}^{\mathrm{miss}}$ calculation involves a term related to the hard objects in the event, and a separate term involving soft objects. Since the hard term is calculated directly from the leptons and jets selected in the analysis, the uncertainty of this part of the $E_{\mathrm{T}}^{\mathrm{miss}}$ is accounted for by the corresponding uncertainties on the lepton and jet energy scales by recomputing the $E_{\mathrm{T}}^{\mathrm{miss}}$ quantity in the systematic variations.

The treatment of the soft term uncertainty is based on a calibration using $Z \to \mu^+\mu^-$ events in 2015 data for which no true missing momentum associated with invisible particles is expected [61]. The deviations of the reconstructed $E_{\mathrm{T}}^{\mathrm{miss}}$ soft term from zero in this sample is projected onto vectors parallel and perpendicular to the hard object axis. Averaging over bins of the phase space and comparing to MC simulation, these are interpreted as the longitudinal and transverse $E_{\mathrm{T}}^{\mathrm{miss}}$ soft term resolution. A systematic variation is compared for each of the two components by smearing the resolution. For the longitudinal projection, an additional pair of systematic variations is produced by scaling its magnitude up or down, and this variation is taken to be correlated with the smearing variations.

4.6.4 b-jet tagging efficiency

The efficiency of b-tagging in MC simulation is corrected to match that observed in data. Uncertainties on these scale factors are estimated by studying samples of data with enriched contributions from b-jets, c-jets, or light jets, summarised as three pairs of weight vectors fluctuating the b-tagging response. The comparison with b-jet enriched samples was done in 2015 data, but the other two

[4] The reduction is prepared by the jet performance community within ATLAS and is regularly updated as the calibration improves. The precise number of reduced components changes according to details of the calibration. While there is negligible impact to the overall uncertainty as a result of this process, the reduced components are no longer easily traced back to the original components which have physical interpretations, and as such cannot be used when combining measurements with other experiments.

comparisons were based on the Run 1 calibration [60] and covered by an additional extrapolation uncertainty to account for changes in the detector geometry and operating conditions. For the second analysis presented here the extrapolation from Run 1 was replaced by a new calibration using only 2015 data and simulation. The updated calibration confirmed after the fact that the extrapolation uncertainty was chosen appropriately and covered differences in the conditions. Two additional components to the uncertainty are considered to account for the extrapolation of the scale factors up to higher-p_{T} jets; two pairs of weight vectors implement the effect of this uncertainty.

4.6.5 Luminosity

The luminosity is determined by ATLAS using the VdM scan technique discussed in §2.7. The calibration scale is derived for each year of data-taking, and the total uncertainty is typically about 2 % with the current set of luminometers and method used by ATLAS. Since this is the uncertainty on a single measurement of the total luminosity it is not implemented as a pair of templates and only affects the overall normalisation.

4.7 Sources of modelling uncertainty

Some sources of modelling uncertainty can be treated similarly to the experimental uncertainties by comparing the effect of extreme choices or by propagating the uncertainties from a corresponding calculation. Others such as choices of calculation schemes exist only as a discrete number of alternative implementations, and so are estimated by comparing the predictions of the alternatives.

4.7.1 Monte Carlo modelling for tW and $t\bar{t}$

To estimate the uncertainty of the matrix element modelling in tW and $t\bar{t}$, MC samples were prepared using MG5_aMC@NLO [7] instead of POWHEG-BOX (see §4.2.1). The showering was performed using HERWIG++ 2.7.1 [119] for these samples. For both tW and $t\bar{t}$, another set of samples was generated to estimate the uncertainty from the parton showering. These samples used POWHEG-BOX for the matrix element generation and HERWIG++ for the parton showering. In all cases UE-EE-5 [24] is used to model the underlying event. The matrix element uncertainty is estimated by comparing the samples with POWHEG-BOX or MG5_aMC@NLO, interfaced to HERWIG++ in both cases. The uncertainty due to the parton showering (including the hadronisation and underlying event models) is estimated by comparing the samples with POWHEG-BOX interfaced to PYTHIA6 or HERWIG++. These comparisons are designed to cover differences in the kinematic predictions resulting from the different possible choices of the trade-offs in calculations discussed in §1.3.1.

Additional QCD radiation by the initial or final state particles in the hard scatter is modelled using pairs of alternative samples with the radiation set lower or higher than the nominal configuration. In the POWHEG generators there are three tuning parameters that control the radiation: the h_{damp} parameter and the choice of factorisation scale (μ_F) and renormalisation scale (μ_R); in addition there is a parameter in the P2012 tune related to the choice of α_s in the showering. Each of these parameters is simultaneously varied to produce the samples with varied radiation[5] [39]. The sample

[5] An improved generator setup was developed within ATLAS after the work completed in this thesis. In addition to different generator tunes the way that h_{damp} is varied in these samples is changed and the alternative choice of independently varying the tuning parameters is under investigation.

Table 4.2: MC generators and settings used for the alternative samples, which are compared to the nominal samples (indicated with an asterisk) to estimate systematic uncertainties. The final column identifies whether the detector simulation was the full GEANT4 simulation or the ATLFASTII fast simulation.

Process		Matrix element	Parton showering	Sim.
tW (DS)		POWHEG-BOX v1, CT10	PYTHIA6, CTEQ6L1, P2012	full
tW (DR)	*	POWHEG-BOX v1, CT10	PYTHIA6, CTEQ6L1, P2012	full
		POWHEG-BOX v2, CT10	HERWIG++, UE-EE-5	fast
		POWHEG-BOX, low radiation	PYTHIA6, P2012	fast
		POWHEG-BOX, high radiation	PYTHIA6, P2012	fast
		MG5_aMC@NLO	HERWIG++, UE-EE-5	fast
$t\bar{t}$	*	POWHEG-BOX v2, CT10	PYTHIA6, CTEQ6L1, P2012	full
		POWHEG-BOX v2, CT10	HERWIG++, UE-EE-5	fast
		POWHEG-BOX, low radiation	PYTHIA6, CTEQ6L1, P2012	fast
		POWHEG-BOX, high radiation	PYTHIA6, CTEQ6L1, P2012	fast
		MG5_aMC@NLO	HERWIG++, UE-EE-5	fast

with low radiation has h_{damp} set to m_t and both μ_R and μ_F increased by a factor of 2, while the high radiation sample has $h_{\text{damp}} = 2m_t$ and both μ_R and μ_F set to half of the nominal choice. Half of the difference between the two variations is then taken to be the uncertainty for the modelling of additional radiation.

An ideal estimate of the uncertainty of the DR scheme used to resolve the ambiguity between tW and $t\bar{t}$ (see §1.5.4) would account for the size of the interference between the processes in the part of the phase space corresponding to the event selection criteria. At the time of the measurements in this thesis, the best approximation to this was a comparison between the nominal sample and a sample that was prepared identically except using the DS scheme. In principle this does not guarantee that the uncertainty is well-estimated because neither scheme is fully theoretically consistent, however there is empirical evidence that the difference between DR and DS predictions brackets the predictions of the full resonant-aware calculation (see §7.2.2) [69].

The alternative MC samples discussed in this section are summarised in Table 4.2.

4.7.2 Modelling of other backgrounds

Since the contributions of the other sources of background events after applying selection requirements are small, the uncertainties in their overall normalisations are estimated based on experience gained with the relevant MC samples in ATLAS, and by a comparison between alternative generators in the context of the tW analysis. In particular for determining the normalisation, a global 5 % uncertainty associated with the NNLO $Z + j$ cross-section and a 6 % uncertainty on the NLO diboson cross-sections were found to be suitable. Comparisons between samples prepared using the SHERPA or POWHEG generators show differences in the predicted diboson contributions that vary in magnitude from 0.5 % in the `2j2b` region to 25 % in the `2j1b` region; based on these considerations, the normalisation of diboson events is assigned a 25 % uncertainty across all regions. The normalisation of $Z + j$ events with one b-jet is assigned a 50 % uncertainty, increased to 100 % for $Z + j$ events with two b-jets. The contribution of events with a fake or non-prompt lepton selected is assigned a 100 %

uncertainty. These values are chosen to overestimate the true uncertainty since the small number of these background events remaining after the analysis selection mean that even a large uncertainty has very little impact on the total uncertainty of the cross-section measurement compared to more important sources associated with $t\bar{t}$ and tW modelling.

4.7.3 Parton distribution function modelling

The PDF4LHC15 [99] combined PDF set is used to estimate the uncertainty due to the PDF modelling. This PDF set provides a set of 30 symmetric eigenvectors[6] to summarise the sources of uncertainty in the model. The tW and $t\bar{t}$ samples produced using MG5_aMC@NLO are reweighted according to the central value and these eigenvectors to construct a set of positive and negative variation samples covering the PDF uncertainty. The full difference of the central value compared to the nominal samples using the CT10 PDF set, and between the nominal PDF4LHC sample and each of its variations are each taken as a source of uncertainty for the measurements. The difference between the alternative PDF sets is larger than any of the internal uncertainties of the PDF4LHC set.

[6]The PDF4LHC15 prescription recommends using the full set of 100 replicas, however since the overall impact of the PDF modelling on the measurement is small, the reduced symmetric set can be used without issue.

Chapter 5

Total cross-section

The material presented in this chapter is based on work by the author published as Ref. [50].

This chapter presents the measurement of the total cross-section of tW production using a sample of $\sqrt{s} = 13\,\text{TeV}$ pp collision data with an integrated luminosity of $3.2\,\text{fb}^{-1}$ collected in 2015. The dilepton sample is divided into regions based on the number of jets and b-tagged jets present. BDTs are trained to discriminate between tW and $t\bar{t}$ events in regions with exactly one b-tagged jet and zero or one additional jet. The binned response of these two BDTs in an MC simulated sample, along with the total predicted event yield in the region with exactly 2 b-tagged jets, are fit to the observed event yields using a maximum likelihood fit. The measured cross-section is found to be consistent with theoretical predictions for $\sqrt{s} = 13\,\text{TeV}$.

5.1 Samples used in the analysis

As discussed in §2.1.3 (see Fig. 2.3), the 2015 dataset is the smallest dataset collected in Run 2 as much of that year was dedicated to commissioning the LHC at the new collision energy of 13 TeV. It accounts for only 2 % of the integrated luminosity of the full Run 2 dataset collected at ATLAS. A small subset of the data was collected using the 50 ns LHC bunch spacing configuration that was used in Run 1, but for simplicity this is excluded so that all data analysed here uses the Run 2 nominal bunch spacing of 25 ns. The average number of collisions per bunch crossing was $\langle \mu \rangle = 13.4$, compared to the Run 2 average of $\langle \mu \rangle = 33.7$, due in part to the lower instantaneous luminosity.

The nominal MC simulated samples were prepared according to the discussion in §4.2, and alternative samples for evaluating systematics were prepared following §4.6 and §4.7. The nominal samples were produced using the full detector simulation with GEANT4 while some of the systematic variation samples were produced using ATLFASTII.

Calorimeter issues such as LAr noise bursts are avoided by requiring events to pass the jet cleaning requirements discussed in §3.6.2. The lowest unprescaled single lepton triggers (with their associated complementary high-p_{T} triggers as discussed in §3.4.4 and §3.5.3) are then used to select the samples of real and simulated data. For electrons this was the logical OR of the `HLT_e24_lhmedium_L1EM20VH`, `HLT_e60_lhmedium`, and `HLT_e120_lhloose` triggers, and for muons it was the logical OR of `HLT_mu20_iloose_L1MU15` and `HLT_mu50`. Events must pass the single electron or single muon trigger requirement.

5.1.1 Object definitions

The definition of physics objects is based on the reconstruction algorithms discussed in Chapter 3. Electrons are required to pass the `Tight` identification likelihood while muons must pass the `Medium` identification working point (see §3.4.1 and §3.5.1). Both types of leptons must additionally pass the gradient isolation algorithm (see §3.4.2) and pass requirements on the impact parameters of their ID tracks to ensure that they are compatible with the beam spot position. For electrons the requirements are $|d_0|/\sigma_{d_0} < 5$ and $|z_0 \sin\theta| < 0.5\,\text{mm}$ and for muons they are $|d_0|/\sigma_{d_0} < 3$ and $|z_0 \sin\theta| < 0.5\,\text{mm}$ (see §3.3.1). Electrons must fall within the central region $|\eta| < 2.47$, with those in the transition region between the barrel and end-cap $1.37 < |\eta| < 1.52$ excluded. Muons must fall within $|\eta| < 2.5$. Jets are reconstructed from topo-clusters using the anti-k_t algorithm with $R = 0.4$, and must lie in the central region $|\eta| < 2.5$ with a p_{T} of at least 25 GeV. Jets with $p_{\text{T}} < 60\,\text{GeV}$ and $|\eta| < 2.4$ must satisfy a JVT requirement (see §3.6.2) to exclude jets reconstructed from noise and pileup activity. b-tagging is performed using the 77 % efficiency working point. The $E_{\text{T}}^{\text{miss}}$ is calculated using the selected leptons and jets in the event, with the addition of the soft term calculated from other tracks originating from the hard-scatter primary vertex.

Since the reconstruction of each type of object is independent, it is necessary to account for cases where the same response in the detector is reconstructed as multiple different objects. If an electron and a muon share a common ID track this is likely due to a hard photon emission by a muon (final-state radiation or bremsstrahlung) and so the electron candidate is ignored. Any jets within a cone $\Delta R < 0.2$ of an electron are rejected since they are probably reconstructed from energy deposited by the electron. After this rejection, any electron candidates inside a hollow cone $0.2 < \Delta R < 0.4$ around any remaining jet are rejected on the basis that they are more likely to be constituents of a real hadronic jet rather than prompt electrons. There is a remaining ambiguity between muons and jets due to the chance that muons are accompanied by an energetic photon emitted either through bremsstrahlung or final-state radiation that is reconstructed as a jet. To resolve this any jet with fewer than 3 ID tracks within a cone $\Delta R < 0.4$ of a muon candidate is rejected, and then any muon candidate within a cone $\Delta R < 0.4$ of a remaining jet is rejected.

5.1.2 Event selection

Each event is required to contain exactly two leptons with $p_{\text{T}} > 20\,\text{GeV}$ and opposite electrical charge, at least one of which has $p_{\text{T}} > 25\,\text{GeV}$. An event is rejected if it has any additional lepton with $p_{\text{T}} > 20\,\text{GeV}$. To ensure that one of the two selected leptons was responsible for triggering the event, one of the leptons is required to be matched to the trigger object; in the case of an electron trigger there must be a selected electron within a cone size of $\Delta R = 0.07$ around the trigger object, and in the case of a muon trigger there must be a muon in a cone of $\Delta R = 0.1$ around the trigger object. Contributions from the $Z + j$, $W + j$, and diboson background processes are suppressed using the flavour-dependent requirements on $E_{\text{T}}^{\text{miss}}$ and $M_{\ell\ell}$ presented in §4.3. The requirements are summarised in Table 5.1.

For MC simulated events information about the ancestry of particles is used to identify events where at least one of the selected leptons is not in fact associated with a prompt lepton from the hard-scatter interaction. These events are separated into a sample labelled as fake events, and contribute to the expected event yields.

Table 5.1: Summary of the event selection requirements, including common requirements and the additional requirements that depend on the flavour of the leptons in the event.

Common	At least one jet with $p_T > 25\,\text{GeV}$, $\lvert\eta\rvert < 2.5$	
	Exactly two leptons of opposite charge with $p_T > 20\,\text{GeV}$,	
	$\lvert\eta\rvert < 2.5$ for muons and $\lvert\eta\rvert < 2.47$ excluding $1.37 < \lvert\eta\rvert < 1.52$ for electrons	
	At least one lepton with $p_T > 25\,\text{GeV}$, veto if third lepton with $p_T > 20\,\text{GeV}$	
	At least one lepton matched to the trigger object	
Different flavour	$E_T^{\text{miss}} > 20\,\text{GeV}$	
	$E_T^{\text{miss}} > 50\,\text{GeV}$,	if $M_{\ell\ell} < 80\,\text{GeV}$
Same flavour	$E_T^{\text{miss}} > 40\,\text{GeV}$	
	$M_{\ell\ell} > 40\,\text{GeV}$	
	veto,	if $81\,\text{GeV} < M_{\ell\ell} < 101\,\text{GeV}$
	$4E_T^{\text{miss}} > 5M_{\ell\ell}$,	if $M_{\ell\ell} < 101\,\text{GeV}$
	$2M_{\ell\ell} + E_T^{\text{miss}} > 300\,\text{GeV}$,	if $M_{\ell\ell} \geq 101\,\text{GeV}$

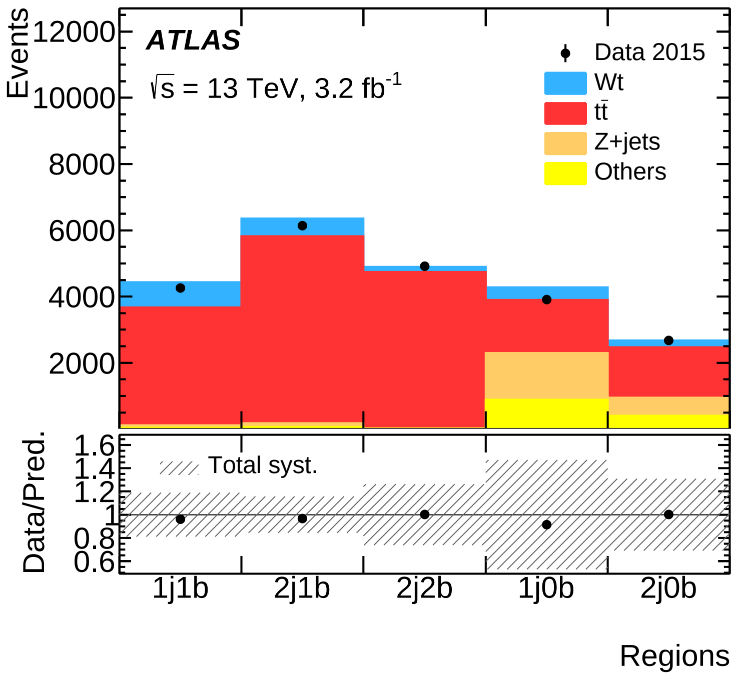

Figure 5.1: Expected number of events for each signal and background process, with the number of observed events superimposed. The expected yields are normalised according to their theoretical cross-sections and the luminosity observed in data. Each bin corresponds to a sub-selection based on jet multiplicity: the first number is the total jet multiplicity and the second is the number of b-jets in the events. The bottom panel shows the ratio of observed to predicted event yield with the total pre-fit systematic and statistical uncertainty shown by the hashed area.

The events remaining after making the selection described in this section are shown in Fig. 5.1. In categories of events where there is at least one b-tagged jet, the contributions of the tW and $t\bar{t}$ processes dominate with only very small contributions from other processes. The `2j2b` region contains only a small number of signal events, and is a high-purity $t\bar{t}$ sample. Categories without b-tagged jets contain large contributions from the prompt bosonic background processes and are used to check their modelling, but are not used directly in the measurement of the tW cross-section. Note that the systematic uncertainty represented by the hashed area in Fig. 5.1 is very large compared to the statistical uncertainty on the observed yields, which is so small that the error bars on the points are not visible. The uncertainty bands are also much larger than the disagreement between the observed and predicted yields in each region. This is largely due to the fact that the total systematic uncertainty here combines many correlated sources of uncertainty without accounting for those correlations, resulting in an over-estimation. The process of fitting the predicted to the observed data will allow the correlations between sources of uncertainty to be explored to provide a better estimate of the total uncertainty.

5.2 BDT performance

A separate BDT is trained for each of the `1j1b` and `2j1b` regions. In the `2j2b` region, the requirement that events include 2 b-jets excludes all but a small number of tW events drawn from a sparsely populated region of the phase space. As a result it is particularly expensive in terms of computational resources to obtain a `2j2b` tW sample large enough to allow for the reliable training and testing of a BDT. Trial BDTs trained with the available MC samples displayed evidence of overtraining that suggested they would not generalise to the observed data. The fact that this region contains only a small contribution from the signal sample also means that it has a smaller contribution to the overall significance of the signal. Attempts to subdivide the region using bins of kinematic quantities such as the p_{T} of selected objects did not result in significant improvements in the measurement, and experiments where the region was excluded from the measurement altogether revealed that its primary contribution is to control the total normalisation of the $t\bar{t}$ process rather than more detailed information. For these reasons no BDT is used in the `2j2b` region.

The two BDTs are each separately optimised according to the method discussed in §4.4.6. The signal events used are a combination of the tW samples simulated using the DR and the DS definitions of the process since these samples exhibit largely similar features for the purpose of distinguishing events from $t\bar{t}$. The nominal $t\bar{t}$ MC sample is used for the background events, such that the trees are trained only to distinguish tW from $t\bar{t}$ events. For both BDTs, the shrinkage parameter was set to 0.05. For the `1j1b` BDT, 600 trees with a maximum depth of 2 were used, while for the `2j1b` BDT 700 trees with maximum depth 3 were used. To improve the performance of the `2j1b` BDT, given the smaller number of signal events in that region, a finer granularity was used when searching the parameter space for subdivisions during tree growth, and a larger fraction of the total signal events was required to reach each leaf node.

The final lists of inputs to the optimised BDTs are shown in Table 5.2 along with the separation power of each (see §4.4.3). The optimised `2j1b` BDT uses a longer list of features, but each individual feature provides less separation power compared to the `1j1b` BDT where the single observable $p_{\mathrm{T}}(\ell\ell b E_{\mathrm{T}}^{\mathrm{miss}})$ accounts for the majority of the separation power. Distributions of the 6 most

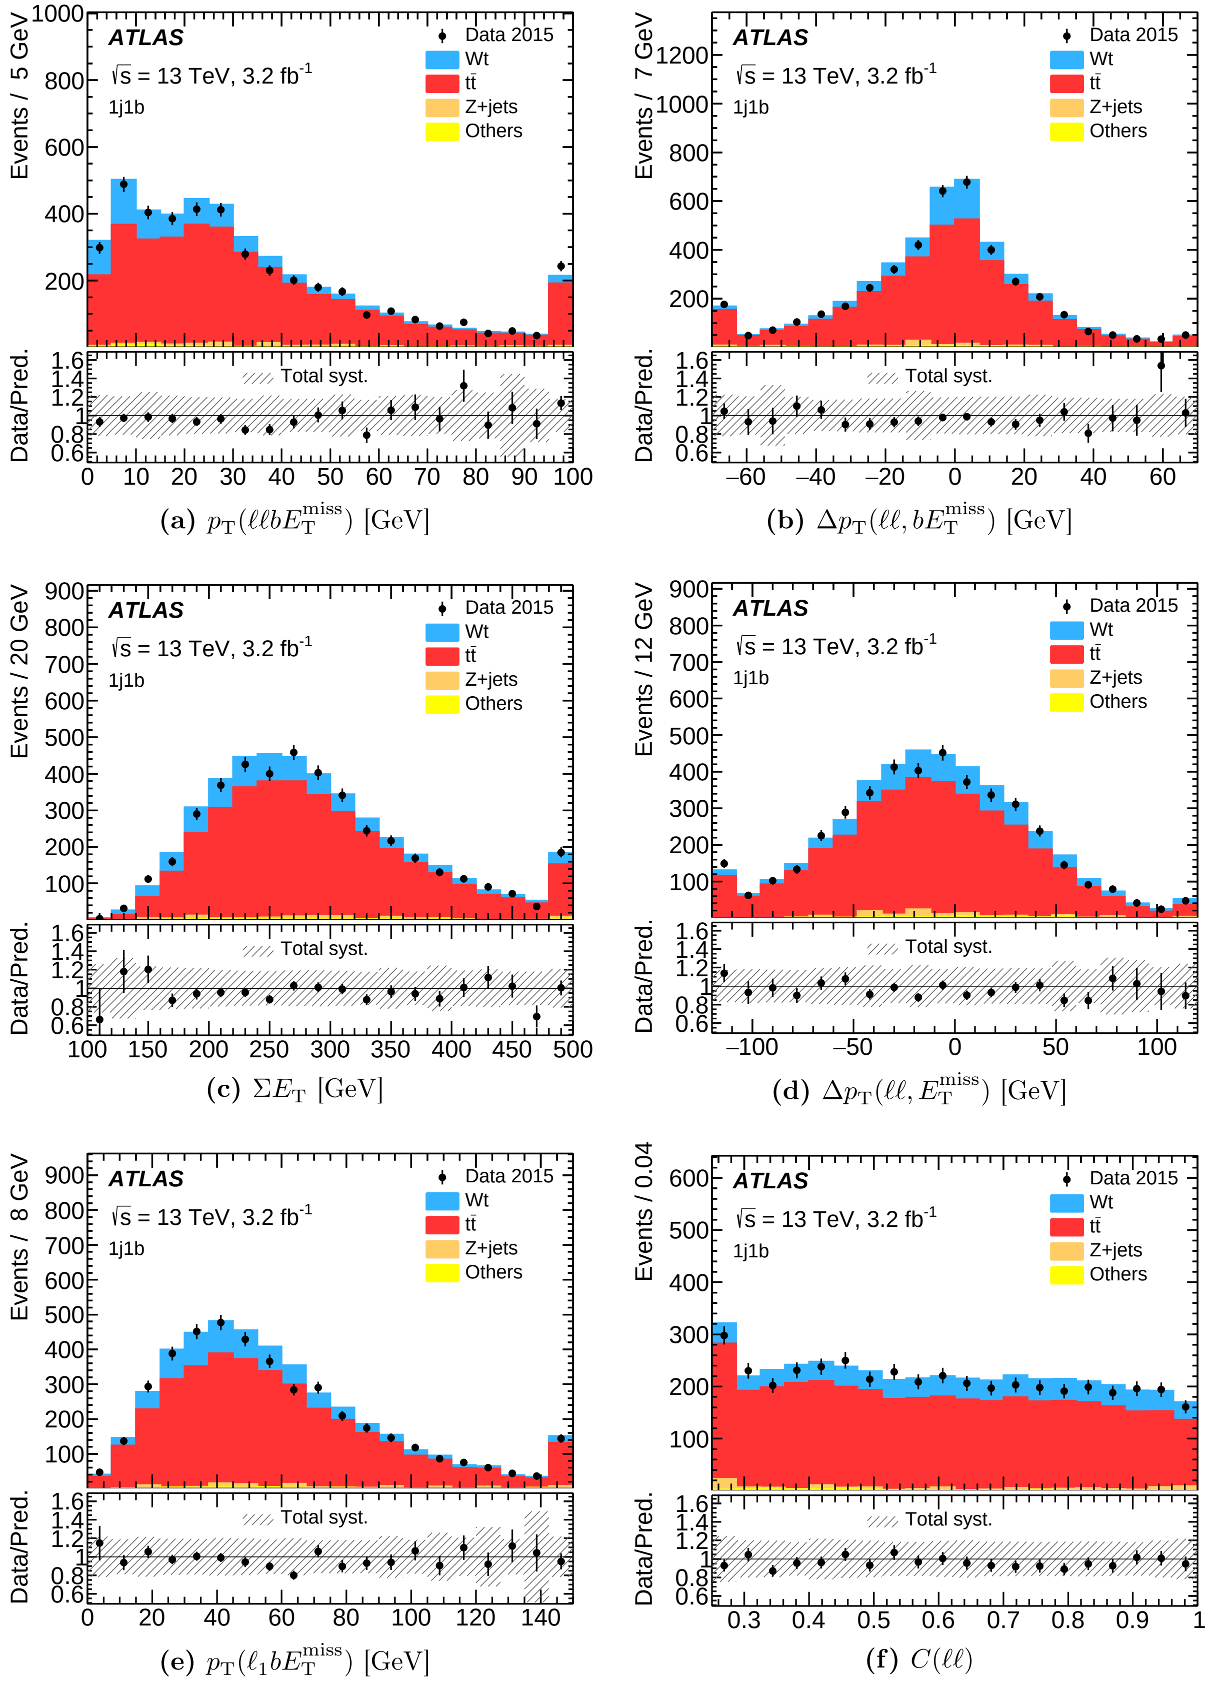

(a) $p_{\mathrm{T}}(\ell\ell b E_{\mathrm{T}}^{\mathrm{miss}})$ [GeV]
(b) $\Delta p_{\mathrm{T}}(\ell\ell, b E_{\mathrm{T}}^{\mathrm{miss}})$ [GeV]
(c) ΣE_{T} [GeV]
(d) $\Delta p_{\mathrm{T}}(\ell\ell, E_{\mathrm{T}}^{\mathrm{miss}})$ [GeV]
(e) $p_{\mathrm{T}}(\ell_1 b E_{\mathrm{T}}^{\mathrm{miss}})$ [GeV]
(f) $C(\ell\ell)$

Figure 5.2: Distributions of the 6 most separating observables used as inputs to the BDT optimised for the `1j1b` region, showing the expected contribution from each process with the observed distributions overlaid and their ratio plotted in the lower panel. The first and last bins include events that respectively underflow or overflow each histogram.

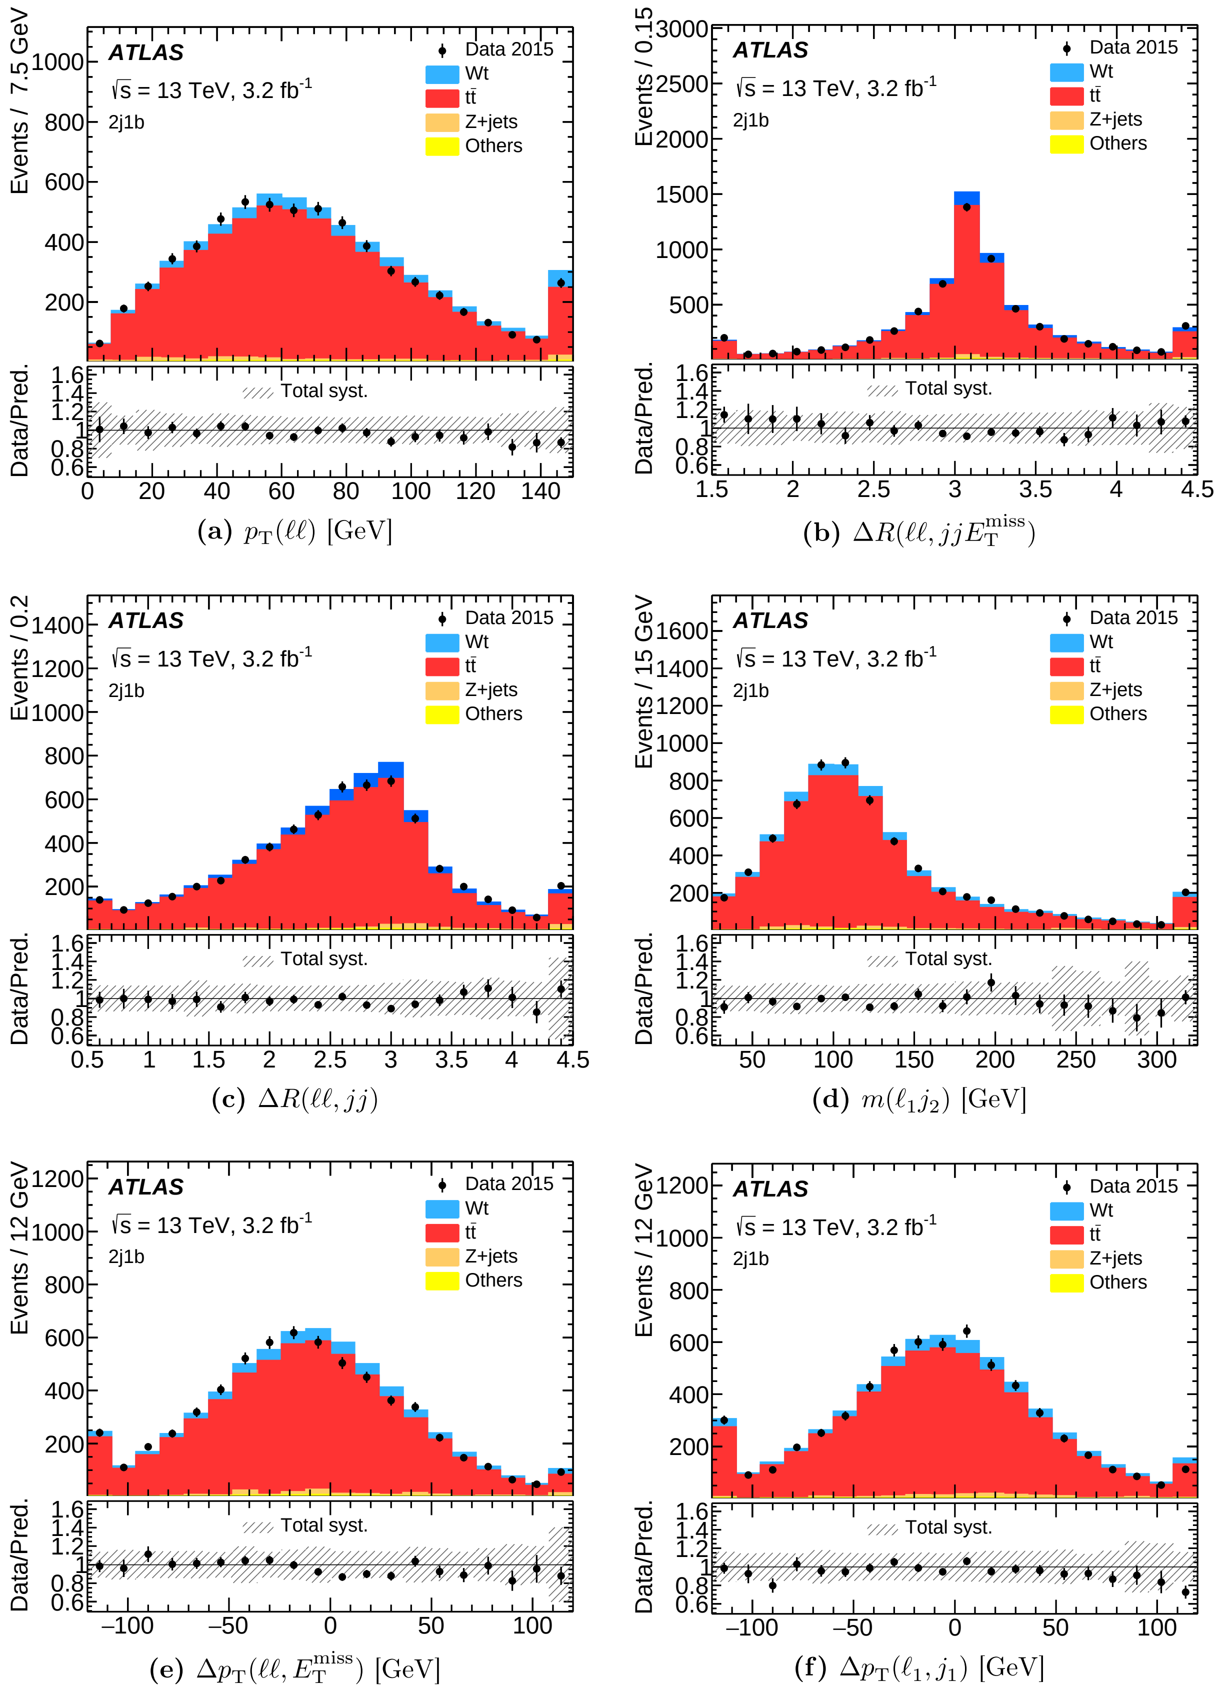

Figure 5.3: Distributions of the 6 most separating observables used as inputs to the BDT optimised for the `2j1b` region, showing the expected contribution from each process with the observed distributions overlaid and their ratio plotted in the lower panel. The first and last bins include events that respectively underflow or overflow each histogram.

Table 5.2: Features used in the optimised BDTs for the `1j1b` and `2j1b` regions ranked according to their separation power S. See §4.4.5 for the definitions of the observables. The overall separation power of each BDT is shown in the final row.

`1j1b`		`2j1b`	
Observable	S $[10^{-2}]$	Observable	S $[10^{-2}]$
$p_{\mathrm{T}}(\ell\ell b E_{\mathrm{T}}^{\mathrm{miss}})$	5.3	$p_{\mathrm{T}}(\ell\ell)$	1.7
$\Delta p_{\mathrm{T}}(\ell\ell, b E_{\mathrm{T}}^{\mathrm{miss}})$	2.9	$\Delta R(\ell\ell, jj E_{\mathrm{T}}^{\mathrm{miss}})$	1.7
ΣE_{T}	2.7	$\Delta R(\ell\ell, jj)$	1.5
$\Delta p_{\mathrm{T}}(\ell\ell, E_{\mathrm{T}}^{\mathrm{miss}})$	1.2	$m(\ell_1 j_2)$	1.4
$p_{\mathrm{T}}(\ell_1 b E_{\mathrm{T}}^{\mathrm{miss}})$	0.9	$\Delta p_{\mathrm{T}}(\ell\ell, E_{\mathrm{T}}^{\mathrm{miss}})$	1.4
$C(\ell\ell)$	0.9	$\Delta p_{\mathrm{T}}(\ell_1, j_1)$	1.4
$\Delta p_{\mathrm{T}}(\ell_1, E_{\mathrm{T}}^{\mathrm{miss}})$	0.8	$m(\ell_1 j_1)$	1.3
		$p_{\mathrm{T}}(\ell_1)$	1.3
		$\sigma(p_{\mathrm{T}})(\ell\ell j_1 E_{\mathrm{T}}^{\mathrm{miss}})$	1.2
		$\Delta R(\ell_1, j_1)$	1.2
		$p_{\mathrm{T}}(j_2)$	0.9
		$\sigma(p_{\mathrm{T}})(\ell\ell jj E_{\mathrm{T}}^{\mathrm{miss}})$	0.9
		$m(\ell_2 jj)$	0.3
		$m(\ell_2 j_1)$	0.3
		$m(\ell_2 j_2)$	0.1
BDT response	8.6	BDT response	10.9

separating observables in the `1j1b` BDT are shown in Fig. 5.2 and the top 6 observables for the `2j1b` region are shown in Fig. 5.3. Note that by convention in this thesis where the tails of distributions extend past the range of a histogram, those tails are incorporated into the bin at the nearest edge of the histogram. MC predictions are normalised according to their theoretical cross-sections and the luminosity of the sample of observed data. In all cases the modelling of the observables is shown to be in good agreement with the observed distributions, with the exception of a small shift in the position of the peak of the $\Delta p_{\mathrm{T}}(\ell\ell, E_{\mathrm{T}}^{\mathrm{miss}})$ distribution in the `2j1b` region (Fig. 5.3e). The size of this disagreement is small, and it occurs in the region around zero whereas the best separation of signal by that observable occurs in the tails of the distribution (see Fig. 5.5e). Since the separation power of the observable is also small, its effect on the BDT response is expected to be negligible.

The shapes of the 6 most important variables for the `1j1b` region are shown in Fig. 5.4 in the MC samples. For the tW process the system of the two charged leptons tends to have a higher centrality than that for the $t\bar{t}$ process, and a p_{T} more similar to that of the system composed of the $E_{\mathrm{T}}^{\mathrm{miss}}$ and the leading jet. $t\bar{t}$ events tend to have a higher ΣE_{T} and harder p_{T} for systems involving the $E_{\mathrm{T}}^{\mathrm{miss}}$ and the leading lepton and jet. The shapes of the 6 most important variables for the `2j1b` region are shown in Fig. 5.5; similar differences in the kinematic distributions can be seen, although there are no variables with as large a separation power as the top three in the `1j1b` region. Note that while these figures and the separation power demonstrate that kinematic differences do exist between the signal and background samples, this does not necessarily translate directly into a ranking of the relative importance of each feature to the BDT classifier response.

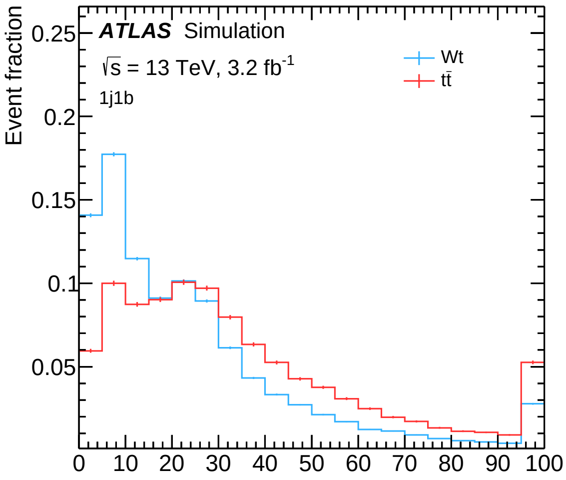

(a) $p_{\mathrm{T}}(\ell\ell b E_{\mathrm{T}}^{\mathrm{miss}})$ [GeV]

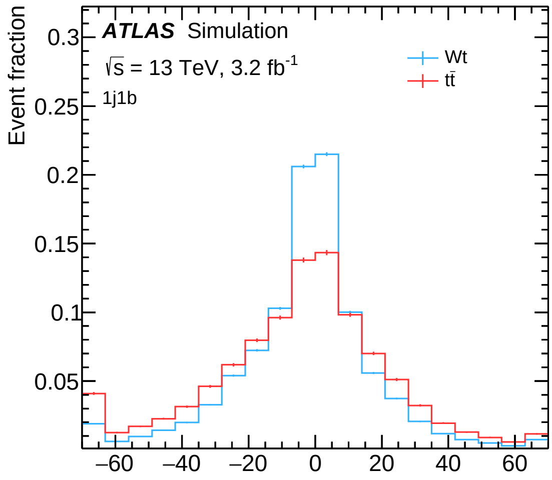

(b) $\Delta p_{\mathrm{T}}(\ell\ell, b E_{\mathrm{T}}^{\mathrm{miss}})$ [GeV]

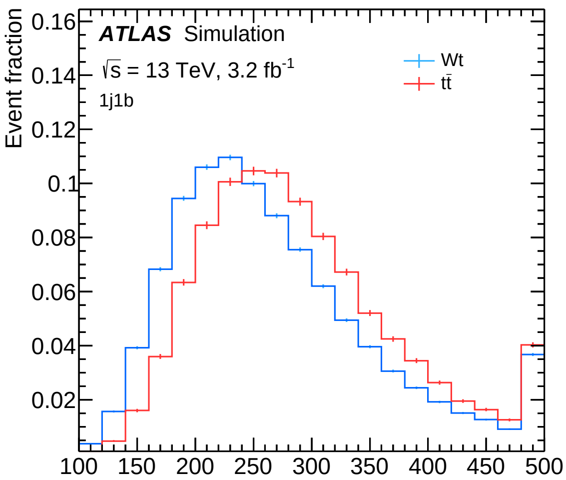

(c) ΣE_{T} [GeV]

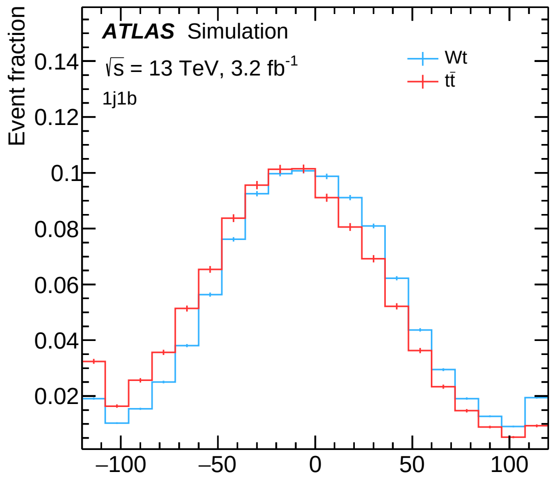

(d) $\Delta p_{\mathrm{T}}(\ell\ell, E_{\mathrm{T}}^{\mathrm{miss}})$ [GeV]

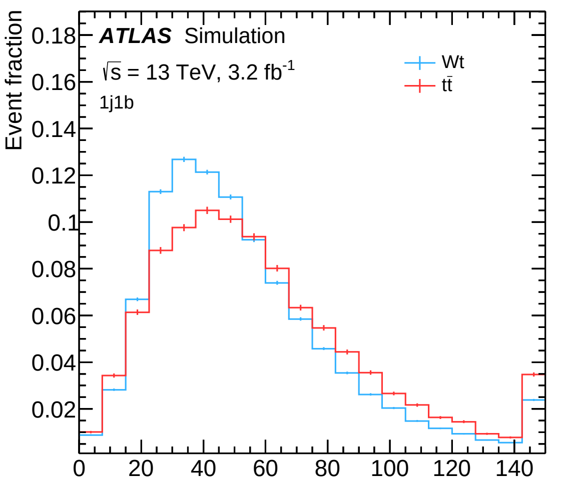

(e) $p_{\mathrm{T}}(\ell_1 b E_{\mathrm{T}}^{\mathrm{miss}})$ [GeV]

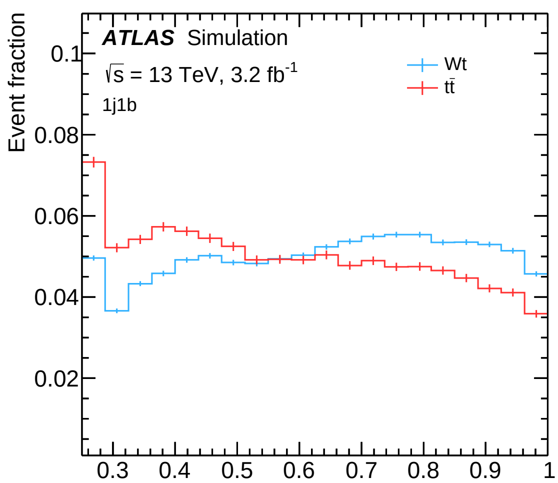

(f) $C(\ell\ell)$

Figure 5.4: Distributions of the 6 most separating observables used as inputs to the BDT optimised for the `1j1b` region, showing the expected contribution from each process normalised to unity. The first and last bins include events that respectively underflow or overflow each histogram.

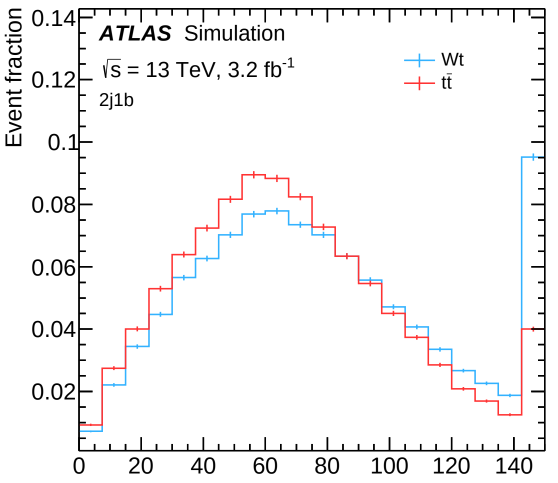

(a) $p_{\mathrm{T}}(\ell\ell)$ [GeV]

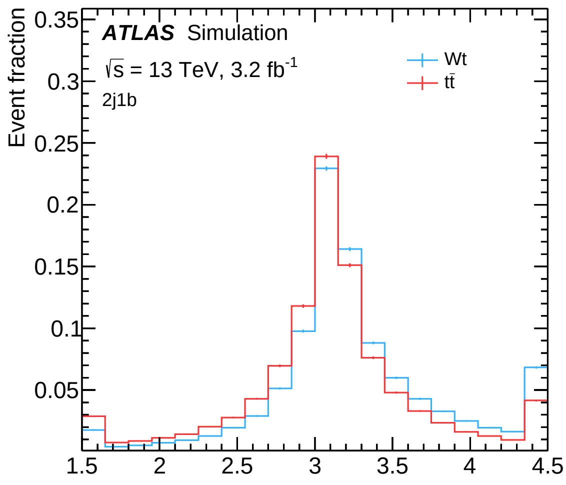

(b) $\Delta R(\ell\ell, jjE_{\mathrm{T}}^{\mathrm{miss}})$

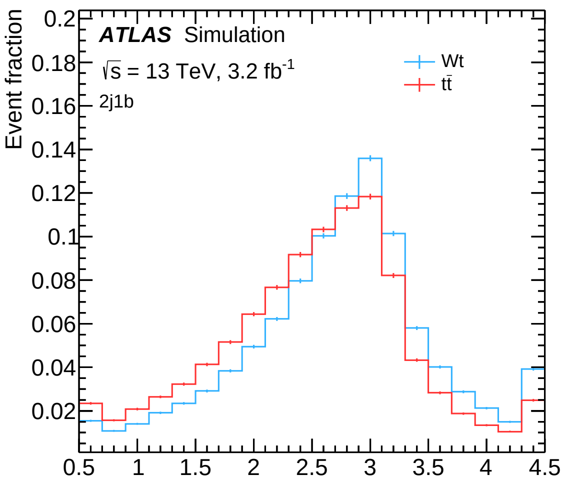

(c) $\Delta R(\ell\ell, jj)$

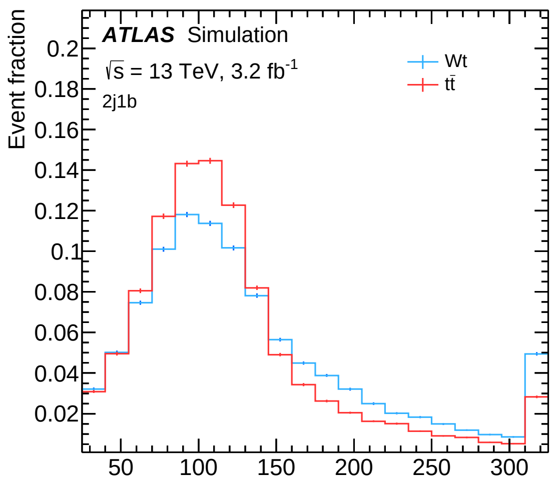

(d) $m(\ell_1 j_2)$ [GeV]

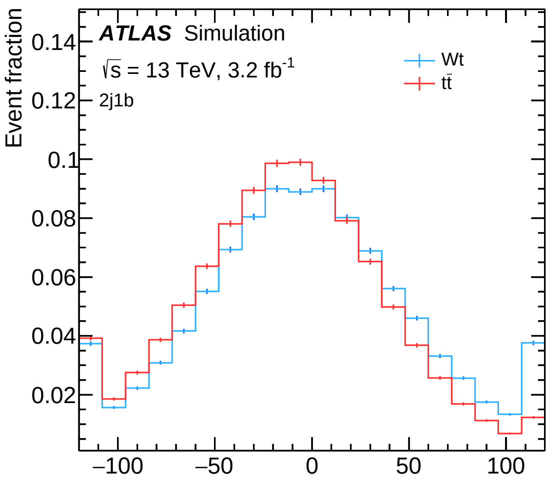

(e) $\Delta p_{\mathrm{T}}(\ell\ell, E_{\mathrm{T}}^{\mathrm{miss}})$ [GeV]

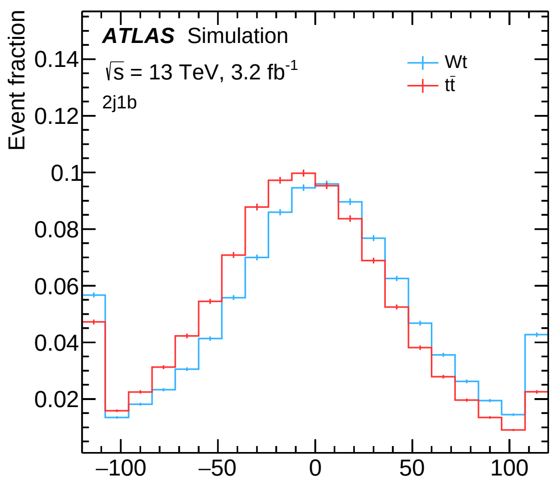

(f) $\Delta p_{\mathrm{T}}(\ell_1, j_1)$ [GeV]

Figure 5.5: Distributions of the 6 most separating observables used as inputs to the BDT optimised for the `2j1b` region, showing the expected contribution from each process normalised to unity. The first and last bins include events that respectively underflow or overflow each histogram.

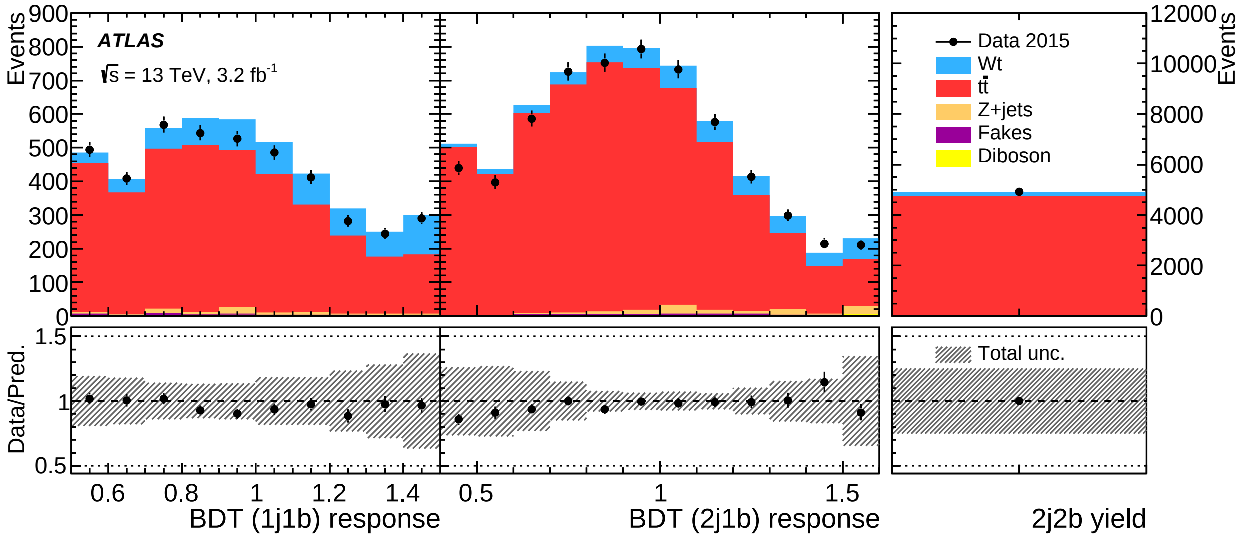

Figure 5.6: Bins used in the fit, shown before the fit is performed. The hashed region in the ratio panel corresponds to the total uncertainty. This combines the statistical uncertainty with the contributions from each individual source of systematic uncertainty assuming no correlations. The uncertainty of the background normalisations are unknown prior to fitting, so are set to zero.

5.3 Maximum likelihood fit

Information from the simulated samples is used together in a binned profile likelihood fit using the observed data to constrain sources of systematic uncertainty and background contributions, and from the results of this fit the cross-section measurement is extracted.

The fit begins by calculating the expected yields of each process in several bins of the response of each BDT, and in the entire 2j2b region as shown in Fig. 5.6, such that the fit is in effect performed over a single one-dimensional histogram. The likelihood function at the heart of the fit must encode Poisson terms that account for the probability of observing a set of yields in the experiment, along with additional pdfs describing parameters that account for fluctuations away from the central values of each source of systematic uncertainty. It contains one parameter of interest related to the signal cross-section along with a number of nuisance parameters related to the cross-sections of the background processes and to the preferred values of the sources of systematic uncertainty.

The fit is performed using the HistFitter [83] software tool, which uses Minuit to perform the minimisation, and is based on HistFactory, RooStats, and RooFit [211].

5.3.1 Sources of uncertainty

The sources of uncertainty are modelled according to the discussion in §4.6 and §4.7. These include experimental sources, modelling of the tW and $t\bar{t}$ processes, normalisation of the smaller background processes, and uncertainty due to the choice of PDF. When considering the sources of uncertainty relevant to all three regions used in the fit and all processes modelled, there are around 300 individual components. Many of these components are negligible when compared to the dominant components, and so these are omitted from the treatment in the fit discussed below, reducing the total number to around 50 parameters related to the systematic uncertainties[1]. The uncertainty on the integrated luminosity is $\pm 2.1\,\%$ as determined using a VdM scan performed in August 2015 following the methodology of Ref. [38].

5.3.2 Implementation of the fit

Rather than working directly with cross-sections, the normalisation of the different processes is handled in the fit by using strength terms denoted by μ, defined as the ratio of a cross-section to its predicted value. The signal strength, μ_{sig}, is the parameter of interest to be determined in the fit. An observed μ_{sig} greater than unity would correspond to the observation of an excess of tW events with respect to the SM prediction, while a μ_{sig} of zero would indicate the complete absence of signal events. The strengths of the background processes are grouped into a vector $\vec{\mu_p}$.

For each source of systematic uncertainty, and for each signal or background process, a pair of histograms is prepared corresponding to upward and downward fluctuations of the effect of that source. These *templates* then contain information about the way in which a change in, for example, the choice of an energy scale in modelling the signal, or a calibration factor in one part of the reconstruction, can impact on the shape and normalisation of the distribution[2]. Each set of three templates (up, down, and nominal[3]) is then interpolated into a continuous function describing its effect on the nominal yields of the relevant process using a single continuous parameter θ_j. To reduce the impact of statistical fluctuations in the simulated samples used to construct the templates, a smoothing algorithm[4] is applied to the ratio of the variations to the nominal in each set of templates where the algorithm was found not to greatly affect the shape of the distribution[5]. When θ_j is zero this corresponds to the nominal value (i.e. no effect on the nominal yields), when it is equal to -1 it corresponds to the downward template, and when it is equal to 1 it corresponds to the upward template. In general the pdf of each of the θ_j parameters is assumed to be Gaussian with unit width. The set of all of these systematic fluctuation terms is written as a vector $\vec{\theta}$.

Note that while this procedure makes intuitive sense for describing an underlying continuous choice, such as that of a modelling parameter, some sources of uncertainty are inherently discrete. For the PS modelling uncertainty for the signal sample, for example, the set of available histograms consists of only the two histograms corresponding to the two different MC PS generators available. The nominal histogram is treated as the baseline template, and the other histogram is converted into a pair of symmetric variations by taking the upward fluctuation to be equal to half of the difference from the nominal histogram, and the downwards template to be identical but in the opposite direction. The corresponding θ_j can then be interpreted as the preference of the data for

[1] In addition to stabilising the fit, this reduction in sources helps to avoid over-counting spurious effects on shape uncertainty due to the limited statistics of the variation samples used in the comparisons. Several approaches to pruning nuisance terms are currently in use, but for this analysis it was done by hand, verifying that the conclusions from the fit were unaffected by removing terms.

[2] The shape and normalisation components of each source are fully correlated and controlled by a single nuisance parameter. Sources that are found to be sufficiently flat with respect to the nominal distribution are treated as normalisation-only sources to simplify the fit.

[3] The nominal sample uses either the full calorimeter simulation or AtlfastII as discussed in §4.6.

[4] The "353QH twice" smoother of Ref. [141] is used, as implemented in the Root [96] framework.

[5] This restriction is intended to avoid losing information in the fine structure of certain variations that might be physically relevant. The size of the smoothing effect is measured using a two-sample KS test to check the similarity of each smoothed and original distribution. Guided by this test, distributions whose shape is expected to contribute substantially to the overall uncertainty are excluded from the smoothing step.

the alternative generator, despite the underlying choice being binary.

The fit procedure interprets the observed data as a sample from an underlying pdf that reflects the probability of observing a given set of yields assuming the SM is true [121]. It proceeds by maximising the likelihood function L, which is the pdf defined by

$$L\left(\vec{n},\vec{\theta_0}\;\middle|\;\mu_{\text{sig}},\vec{\mu_p},\vec{\theta}\right) := \prod_{i=1}^{N} P\left(n_i\;\middle|\;\lambda_i(\mu_{\text{sig}},\vec{\mu_p},\vec{\theta})\right) \times C_{\text{syst}}\left(\vec{\theta_0},\vec{\theta}\right).$$

The first term is a product of Poisson distributions P of the number of events n_i observed in each bin from the set of N bins. The likelihoods λ_i encode the expected yields in each bin for a given set of strength terms and systematic fluctuations. The second term, C_{syst}, is a probability function describing how the systematic fluctuations vary about their central values $\vec{\theta_0}$, here taken to be

$$C_{\text{syst}}\left(\vec{\theta_0},\vec{\theta}\right) = \prod_{j=1}^{M} G\left(\theta_{0,j}-\theta_j\right),$$

where M is the number of systematic sources, G is the standard normal distribution, and typically $\vec{\theta_0} = 0$ [83]. Performing the fit is equivalent to finding the set of values for the unknown parameters μ_{sig}, $\vec{\mu_p}$, and $\vec{\theta}$ that maximises L. In practice this is done by minimising the value of $-2\ln L$. During the fit, the covariance matrix is calculated, facilitating the calculation of the combined uncertainty on μ_{sig} taking into account the correlations between sources of systematic uncertainty as determined by the data. The confidence contours yielded by this procedure have the correct coverage under the assumption that the pdfs are Gaussian [175, p. 92].

5.3.3 Validation of the analysis

In order to avoid biasing the results of the analysis, the real dataset was not used in the fit until all important choices of the analysis had been finalised. It is useful to estimate the sensitivity of the measurement by studying the "average" significance obtained by fits to each of an ensemble of random datasets (e.g. fluctuating the bin contents of the simulated data according to Poisson statistics and their systematic uncertainty). This would be computationally expensive in practice, so it is preferable to replace this ensemble of pseudo-datasets with a single representative dataset. Using the simulated dataset as this representative results in a statistically well-behaved estimate of the measurement sensitivity using the asymptotics formalism of Ref. [122], and in this context the dataset is called the *Asimov dataset*.

Using the Asimov dataset in the fit, the signal strength is always measured to be unity, and all nuisance parameters take on their nominal values. This Asimov fit allows the covariance matrix of the fit parameters to be calculated, along with the overall expected uncertainty and the relative impact of each source. By studying the expected behaviour of the fit it is possible to assess the sensitivity and behaviour of the fit without introducing bias from the observed yields. It was determined that the design of the analysis was sufficiently sensitive to permit a significant measurement to be performed with real data. Checking many configurations of the fit, for example removing certain sources or regions, the detailed behaviour of the fit was checked so that any large or asymmetrical constraints were understood. Choices related to the correlation of sources, the smoothing of shape systematics, and the removal of shape or normalisation contributions from particular sources were fixed at this

Table 5.3: Expected event yields and the breakdown of the contribution from each process. The yields are calculated using the same integrated luminosity as the observed dataset. The uncertainties shown are the total uncertainty before the fit, assuming that the sources are uncorrelated.

	1j1b	2j1b	2j2b
Observed events	4254	6138	4912
Expected total events	4430 ±780	6350 ±810	4900 ±1200
Expected tW events	760 ±140	526 ± 87	164 ± 45
Expected $t\bar{t}$ events	3550 ±700	5640 ±820	4700 ±1200
Expected $Z+j$ events	74 ± 37	118 ± 59	7.5 ± 7.4
Expected fake events	25 ± 25	37 ± 36	14 ± 14
Expected diboson events	23.7± 6.0	24.9± 6.3	0.91± 0.23

Table 5.4: Fitted event yields and the breakdown of the contribution from each process. The uncertainties shown are the total uncertainty after the fit, taking into account the correlations between sources.

	1j1b	2j1b	2j2b
Observed events	4254	6138	4912
Fitted events	4257	6139	4908
Fitted tW events	910 ±210	640 ±160	210 ± 82
Fitted $t\bar{t}$ events	3230 ±210	5340 ±160	4670 ±110
Fitted $Z+j$ events	69 ± 35	87 ± 46	7.6 ± 7.5
Fitted fake events	30 ± 26	40 ± 38	15 ± 14
Fitted diboson events	23.5± 6.0	24.8± 6.2	0.91± 0.23

stage.

5.4 Results

By performing the maximum likelihood fit to the observed data, the signal strength and nuisance parameters are determined. The expected event yield in each region is shown in Table 5.3 and the final fitted event yields are shown in Table 5.4. In each region the tW yield is higher than expected and the $t\bar{t}$ yield is lower than expected, but the yields are within the uncertainty of the predictions. Compared to an expected signal strength of $\mu_{\text{sig}} = 1.00^{+0.33}_{-0.29}$, the observed signal strength is $\mu_{\text{sig}} = 1.31^{+0.41}_{-0.34}$.

The bin-by-bin fitted yields and uncertainties are shown in Fig. 5.7. With respect to the pre-fit distributions shown above in Fig. 5.6, the post-fit uncertainties are substantially reduced due to accounting for the correlation between components of the uncertainty and due to constraints imposed by the observed data. The fitted distribution agrees well with the observed data, with small residuals randomly distributed around zero. The background processes other than $t\bar{t}$ contribute few events to the selected regions.

The behaviour of the most important nuisance parameters in the fit is summarised in Fig. 5.8. To produce this figure four fits are performed for each nuisance parameter in addition to the nominal fit.

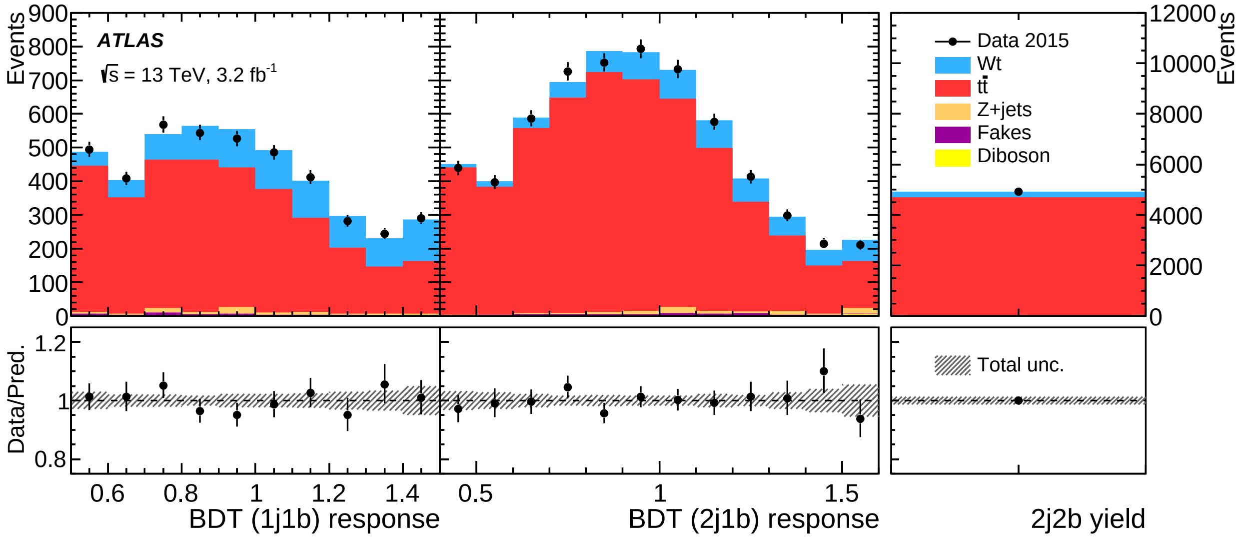

Figure 5.7: The post-fit distributions for each region. The hashed region in the ratio panel corresponds to the total uncertainty. This combines the statistical uncertainty with the contributions from each individual source of systematic uncertainty taking into account all correlations between them.

In each case the parameter is altered while fixing all other nuisance parameters to the values they took in the nominal fit[6], and the impact of the alteration on the newly-fitted signal strength is observed. To obtain the pre-fit impact for a given nuisance parameter, it is first fixed to the -1σ extreme of its pre-fit uncertainty and then to its $+1\sigma$ extreme. The green band extends from the smaller to the larger value of $\Delta\mu$, the difference of μ_{sig} from its value in the nominal fit in each case. The hatched purple band is constructed in the same manner but by using the $\pm1\sigma$ extremes of the post-fit uncertainty instead. The black points show the pull, $(\hat{\theta} - \theta_0)/\Delta\theta$, which is the difference between the post-fit $(\hat{\theta})$ and pre-fit (θ_0) values of the nuisance parameter divided by the total pre-fit uncertainty. The error bars around these points shows the post-fit uncertainty of the nuisance parameter.

The overall contributions to the uncertainty are summarised in Table 5.5, with related nuisance parameters combined into categories of uncertainty. The uncertainty estimates are obtained in the same way as the post-fit impact on μ_{sig} in Fig. 5.8. Since the estimates are made by considering each source in isolation, the resulting values do not combine to the total obtained from the nominal fit when summed in quadrature.

5.5 Interpretation

The SM prediction for the tW production cross-section at NLO+NNLL ([168]; see §4.2.1) is

$$\sigma_{\text{theory}}(tW) = 71.7 \ \pm 1.8\,(\text{scale}) \ \pm 3.4\,(\text{PDF}) \ \text{pb}.$$

[6]Note that this procedure differs from the convention used in similar figures in other ATLAS analyses where the other nuisance parameters are allowed to float in the fit rather than being fixed to their nominal values.

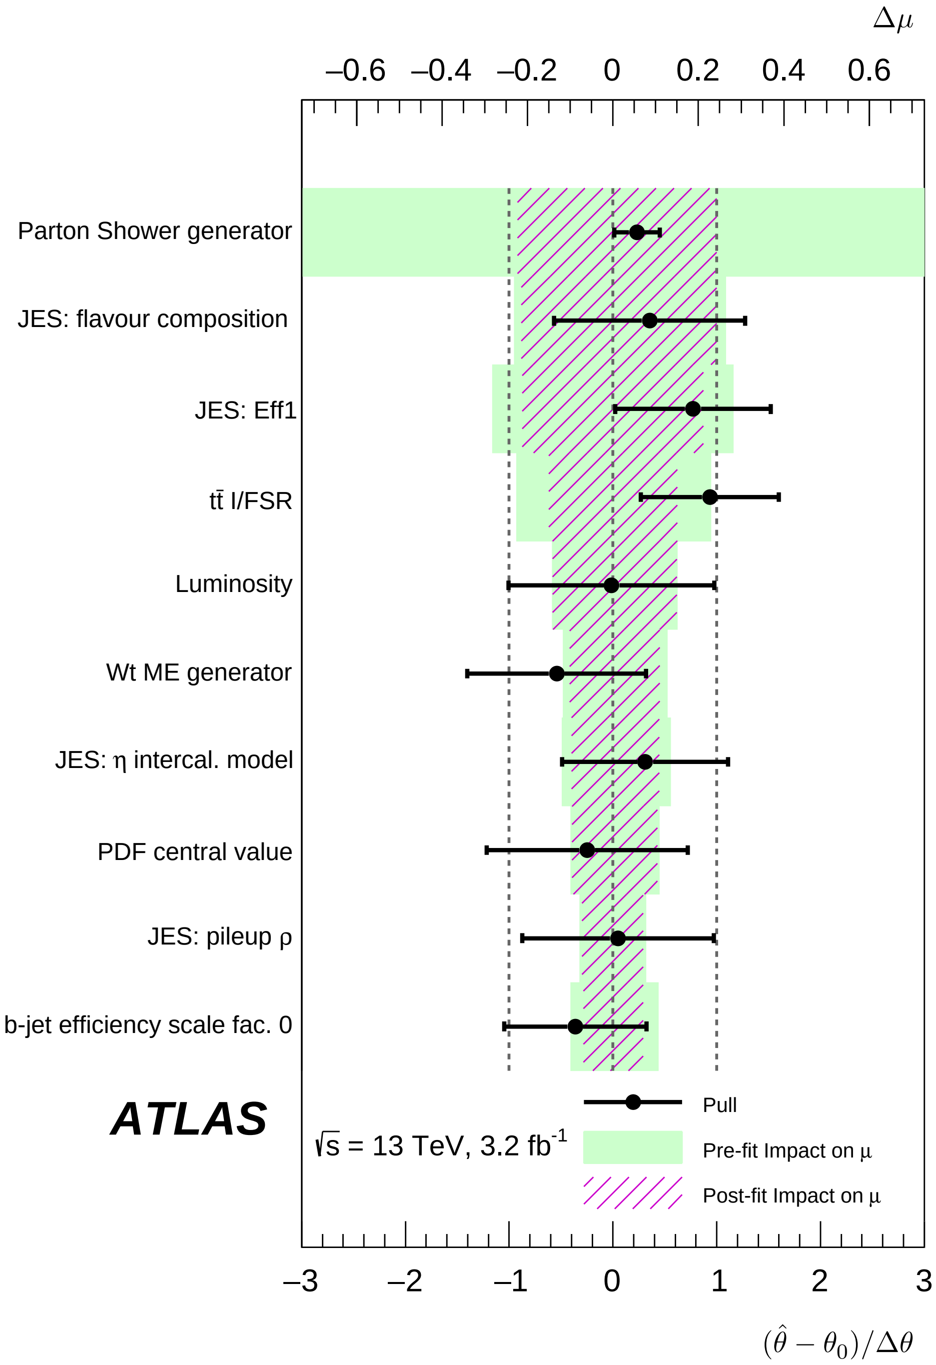

Figure 5.8: The top nuisance parameters in the fit ranked by their post-fit impact on the signal strength. The black points correspond to the lower axis scale and show the pull of each parameter. The green bands and the purple hatched bands correspond to the upper axis scale and show the impact of each parameter on the signal strength before and after the fit respectively.

Source	$\Delta\sigma_{tW}/\sigma_{tW}$ [%]
Jet energy scale	21
Jet energy resolution	8.6
$E_{\mathrm{T}}^{\mathrm{miss}}$ soft terms	5.3
b-tagging	4.3
Luminosity	2.3
Lepton efficiency, energy scale, resolution	1.3
NLO matrix element generator	18
Parton shower and hadronisation	7.1
Initial-/final-state radiation	6.4
Diagram removal/subtraction	5.3
Parton distribution function	2.7
Non-$t\bar{t}$ background normalisation	3.7
Total systematic uncertainty	30
Data statistics	10
Total uncertainty	31

Table 5.5: Summary of contributions to the total uncertainty. Entries in the table summarise more than one individual nuisance parameter, grouping them into related categories. The top section of the table shows experimental sources of uncertainty, the middle section shows uncertainty related to the modelling of the signal and background processes, and the final section shows statistical uncertainty and the totals. Note that due to the method by which the contributions are estimated, they do not sum (in quadrature) to the total quoted systematic uncertainty.

Multiplying the central value by the fitted signal strength, the cross-section is measured to be

$$\sigma_{tW} = 94 \ \pm 10\,(\text{stat.}) \ {}^{+28}_{-22}\,(\text{syst.}) \ \pm 2\,(\text{lumi.}) \ \text{pb},$$

where the uncertainty is quoted separately for the data statistics, total systematic, and luminosity respectively. To assess the significance of this measurement, a statistical test is performed based on the q_0 test statistic for the discovery of a positive signal of Ref. [122]. The background-only null hypothesis ($\mu_{\text{sig}} = 0$) is rejected at a significance of 4.5σ. Performing the same test with the Asimov dataset yields an expected significance of 3.9σ. That the significance of the measurement is higher than expected corresponds to the observation of more signal events than predicted.

The uncertainty in the measurement is dominated by systematic sources, though the data statistics available with the early Run 2 sample used in the analysis still result in a substantial uncertainty. Of the systematic sources, the largest single contribution is from the JES. This is consistent with the fact that discrimination between tW and $t\bar{t}$ is sensitive to the kinematics of jets in the selected events, particularly in events with two jets. Note also that the JES calibration used conservative estimates of uncertainty to cover the extrapolation of parts of the calibration from Run 1 measurements (see §4.6.2). The next largest single contribution is from the modelling of tW and $t\bar{t}$ with the NLO ME generator. This is due to differences in the predictions of MG5_aMC@NLO and Powheg-Box of kinematic quantities to which the analysis is sensitive.

Three of the top four nuisance parameters are pulled from zero in the fit as can be seen in Fig. 5.8. These indicate a preference for the alternative PS generator and setting for the additional radiation in $t\bar{t}$ ("$t\bar{t}$ I/FSR"), and one component of the jet energy scale uncertainty ("JES: Eff1") is also slightly pulled. The information about these nuisance parameters learned from the data is only applicable to the part of the phase space determined by the event selection and object definitions of the analysis and including the structure of the BDT discriminant; any preference for alternative modelling choices does not necessarily generalise to events with additional jets or alternative decays of the tW and $t\bar{t}$ systems. In order to assess the implications of these preferences it would be necessary to measure the

Figure 5.9: Summary of measurements of the tW production cross-section with ATLAS at different centre-of-mass energies. The measurement discussed in this chapter is the red square. Taken from Ref. [73].

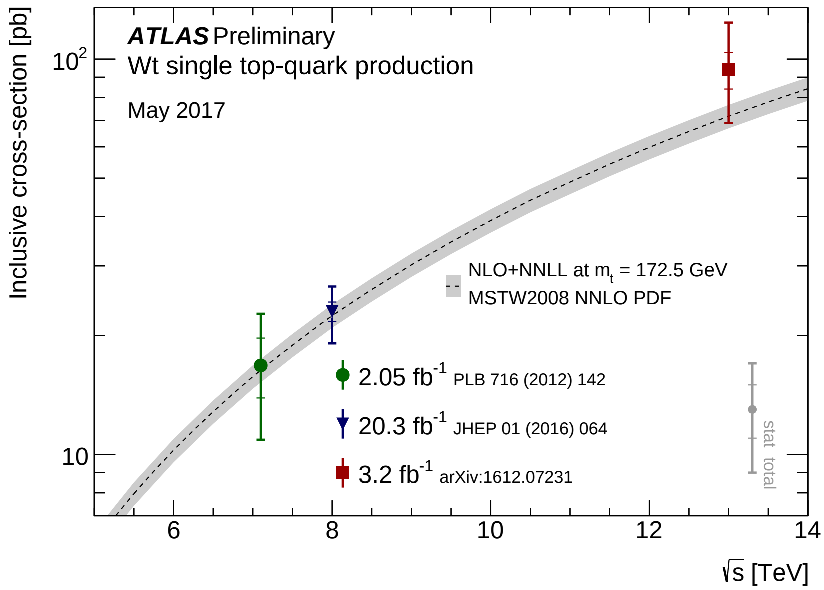

fiducial cross-section (i.e. a direct measurement of the cross-section in a restricted part of the phase space as opposed to an inclusive cross-section) and extrapolate it to the full phase space while taking into account the effect of uncertainties on the acceptance and efficiency corrections related to the extrapolation.

Some nuisance parameters have post-fit uncertainties significantly smaller than their pre-fit uncertainties. This is particularly evident in the case of the PS modelling uncertainty which has a pre-fit impact of more than 60 % on μ_{sig}. In this case the large reduction in uncertainty of the nuisance parameter is due to information provided by the $t\bar{t}$-rich 2j2b region in the fit. This region constrains the overall $t\bar{t}$ normalisation and the relative proportion of tW to $t\bar{t}$ and so constrains this particular modelling term. Uncertainty on the overall $t\bar{t}$ normalisation also explains the asymmetry observed in the impacts of some nuisance parameters.

The measurement is shown in Fig. 5.9 alongside the theoretical calculation of the tW production cross-section and the two previous measurements of the same quantity by ATLAS at lower centre-of-mass energies [32, 52]. The measurement is shown in Fig. 5.10 alongside similar measurements performed by the CMS Collaboration [109, 114, 111], and measurements by ATLAS and CMS of single top quark production in the t-channel and s-channel. Note that the CMS measurement of σ_{tW} at 13 TeV using 35.9 fb^{-1} of data is below the NLO+NNLL prediction, providing further evidence that the slight excess observed in the present analysis is likely as a result of a statistical fluctuation in the observed dataset. The overall picture is of consistency with the SM, with no significant deviations observed from the predicted cross-section.

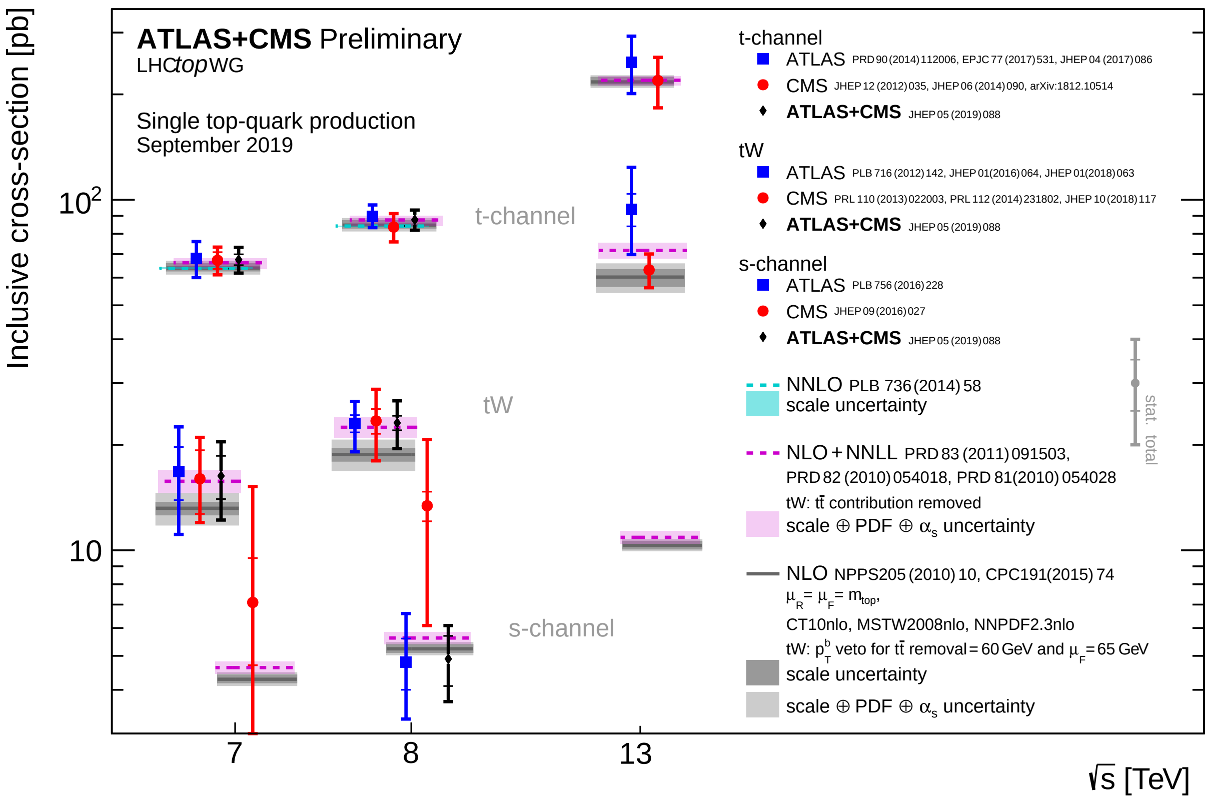

Figure 5.10: Summary of measurements of single top quark production cross-sections at the LHC with ATLAS and CMS at different centre-of-mass energies. The measurement discussed in this chapter is the blue square in the middle row at 13 TeV. Taken from Ref. [73].

Chapter 6

Differential cross-sections

The material presented in this chapter is based on work by the author published as Ref. [47].

The pp collision data collected with ATLAS in 2016 had an integrated luminosity an order of magnitude higher than in the previous year. Whereas it was already possible to measure the total tW production cross-section with the 2015 dataset, with this larger dataset it became possible to measure the tW production cross-section as a function of various quantities. Doing so provides information about the kinematics of tW production that can be compared to predictions from MC simulations to check their agreement with data for the first time.

The analysis presented in this chapter performs normalised differential cross-section measurements of tW production with respect to 6 kinematic quantities in events with two charged leptons and exactly one jet which is b-tagged. The differential cross-sections are measured as a function of:

- $E(b)$, the energy of the b-jet;
- $m(\ell_1 b)$, the invariant mass of the system comprised of the highest-p_{T} charged lepton and the b-jet;
- $m(\ell_2 b)$, the invariant mass of the system comprised of the sub-leading lepton and the b-jet;
- $E(\ell\ell b)$, the energy of the system of the two leptons and the b-jet;
- $m(\ell\ell b)$, the invariant mass of the same system; and
- $m_{\mathrm{T}}(\ell\ell\nu\nu b)$, the transverse mass of the system containing both charged leptons, the b, and the missing momentum due to the two neutrinos. Whereas the definition of transverse mass given in §4.4.5 applied only to a pair of objects, this quantity is defined as

$$m_{\mathrm{T}}(\ell\ell\nu\nu b) \equiv \Big[\big(p_{\mathrm{T}}(\ell_1) + p_{\mathrm{T}}(\ell_2) + p_{\mathrm{T}}(\nu\nu) + p_{\mathrm{T}}(b)\big)^2 - \big(p_x(\ell_1) + p_x(\ell_2) + p_x(\nu\nu) + p_x(b)\big)^2 - \big(p_y(\ell_1) + p_y(\ell_2) + p_y(\nu\nu) + p_y(b)\big)^2 \Big]^{1/2},$$

  where p_x and p_y are the components of momentum in the x and y directions. The object $\nu\nu$ represents the vector sum of the neutrino four-vectors at particle level (see §6.2.2), while at reconstruction level (see §6.2.1) the $E_{\mathrm{T}}^{\mathrm{miss}}$ four-vector is used instead.

These quantities are related to the kinematics of the t and W boson, but are not direct proxies due to the ambiguity in assigning decay products to the particles from the hard-scatter interaction in the kinematically under-constrained dilepton topology.

Table 6.1: Single lepton triggers used with each dataset. In each period at least one of the triggers is required to have passed in order for the event to be selected. The trigger names follow a consistent scheme: `HLT_eXX` or `HLT_muXX` indicate that an electron or muon respectively were required to be found in the HLT with at least the p_{T} indicated, in GeV; `lhloose`, `lhmedium`, and `lhtight` specify the lepton identification working point; `iloose`, `ivarloose`, and `ivarmedium` specify the isolation working point; `L1EM20VH` and `L1MU15` identify the L1 triggers used as input to the HLT decision; `nod0` indicates relaxed requirements on the track impact parameter.

	2015	2016
Electron	HLT_e24_lhmedium_L1EM20VH	HLT_e26_lhtight_nod0_ivarloose
	HLT_e60_lhmedium	HLT_e60_lhmedium_nod0
	HLT_e120_lhloose	HLT_e140_lhloose_nod0
Muon	HLT_mu20_iloose_L1MU15	HLT_mu26_ivarmedium
	HLT_mu50	HLT_mu50

Events that are used in the measurement must contain exactly one jet, and that jet must be b-tagged; this corresponds to using only the `1j1b` region in the terminology introduced in Chapter 5. With respect to the previous chapter, there are minor differences in the event selection criteria. The data are corrected to account for detector effects such as resolution by unfolding to particle level (see §6.2.2), allowing for direct comparison with several predictions from simulated samples.

6.1 Data and simulated samples used in the analysis

The ATLAS $\sqrt{s} = 13$ TeV pp collision dataset collected in 2015 and 2016 has an integrated luminosity of 36.1 fb^{-1} accounting for 25 % of the entire Run 2 dataset. The data were collected with an LHC bunch spacing of 25 ns; the small dataset collected in 2015 with a bunch spacing of 50 ns is not used. The 2016 dataset has a pileup of $\langle\mu\rangle = 25.1$, approximately twice that of the 2015 dataset due to changes in the LHC beam configuration to increase the beam luminosity, but lower than the Run 2 average of $\langle\mu\rangle = 33.7$.

The set of triggers used for collecting samples of events with energetic leptons was modified for the 2016 data-taking period, as shown in Table 6.1. The most significant change was an increase in the minimum p_{T} thresholds in the HLT trigger from 24 GeV to 26 GeV for electrons and from 20 GeV to 26 GeV for muons. This was necessary to reduce the rate with which events were collected in light of the higher luminosity environment later in Run 2. The post-2015 triggers also differ in their requirements on the identification and isolation thresholds for the leptons (see §3.4.4 and §3.5.3).

MC simulated samples were prepared according to the considerations discussed in §4.2, §4.6 and §4.7, with the nominal samples produced using the full detector simulation with GEANT4 and some of the systematic variation samples produced using ATLFASTII.

6.2 Object definitions

Objects defined in terms of the response of the detector are called *reconstruction-level* objects, and are convenient to work with experimentally due to their correspondence with detector observables. In general these definitions depend on the detector, but also on the specific analysis choices and the particular calibrations used by the analysis. Observables derived from these objects, such as their

p_T, are then not directly comparable with predictions for the kinematic distributions of particles within the QFT. For the cross-section measurement in the previous chapter, this was overcome by transforming the outgoing particles in the MC generator records into reconstruction-level objects by first simulating their interaction with the detector and then reconstructing the simulated response using the same reconstruction process as for the observed data. The resulting MC reconstruction-level objects are then directly comparable with their experimental counterparts provided that the detector response was accurately modelled and simulated.

The goal of the present measurement is to produce distributions that can be directly compared to theoretical predictions now and in the future. If the measurement were performed at reconstruction level then to test any theoretical predictions it would be necessary to use detailed knowledge of the ATLAS detector and its calibrated response as it was in 2016. This approach is clearly not tractable. Instead the effects of the detector are removed by unfolding (see §6.5) to reinterpret the reconstruction-level observables as *particle-level* observables. This then forms a well-defined interface between theory and experiment with which comparisons can be made. The unfolding technique requires that observables for MC events be calculated both at reconstruction level and at particle level.

6.2.1 Reconstruction-level objects

The reconstruction-level physics objects used in the measurement are based on the reconstruction algorithms discussed in Chapter 3. They differ only slightly from those in the analysis of Chapter 5 due to incremental improvements in the calibration and reconstruction performance as a result of greater experience with the detector in Run 2. The leptons are required to be detected in the central region of the detector, and must pass identification, isolation, and track quality requirements. Electrons must satisfy $|\eta| < 2.47$, with the transition region between the barrel and endcap ECAL at $1.37 < |\eta| < 1.52$ excluded; for muons the requirement is $|\eta| < 2.5$. Jets are again required to satisfy $p_T > 25\,\text{GeV}$, and the $77\,\%$ efficiency working point for b-tagging is used. All identified jets with $p_T > 20\,\text{GeV}$ and selected leptons are used together with a soft term incorporating other unmatched tracks originating from the primary vertex to calculate the E_T^{miss}.

The procedure to ensure that the response of the same physical particle is not reconstructed as multiple physics objects is slightly different from the one described in §5.1.1. The differences do not have a large effect on the selection, and are motivated by a move towards a standard procedure across all similar ATLAS analyses. Any muon candidate that includes a calorimeter energy deposit and shares a common ID track with an electron candidate is removed. For any remaining cases where a muon and electron candidate share a track, the electron candidate is removed. Next, jets within a cone of $\Delta R = 0.2$ of an electron are considered, but only the jet closest to the electron is removed. Any electron candidate within $\Delta R < 0.4$ of a remaining jet is removed. Next, jets with fewer than three tracks that are within $\Delta R < 0.2$ of a remaining muon are removed, followed by the removal of any muon candidates within $\Delta R < 0.4$ of a remaining jet.

6.2.2 Particle-level objects

Particle level quantities are defined in terms of the stable particles of MC generator event records using standard definitions where possible to form an interface between theory and experiment [70].

In this context stability of a particle is taken to mean having a lifetime τ such that $c\tau > 10\,\text{mm}$.

Prompt leptons are defined to be those that do not arise from the decay of hadrons; in particular they must include a Z or W boson in their ancestry. To account for the associated radiation, these bare leptons are dressed by incorporating the four-momenta of all photons found within a cone[1] of $\Delta R_y < 0.1$ that do not originate from a hadronic decay. All particles other than the bare leptons and the photons used to dress them are then clustered using the anti-k_t algorithm with $R = 0.4$ (see §1.4) to form the particle-level jets. These jets are then b-tagged using a ghost matching technique similar to that described in §3.6. The jet clustering is performed again, but with b-hadrons modified so that their p_T is reduced to negligible values. Any of the original jets that are associated with a new jet that contains a ghost b-hadron with an original $p_\text{T} > 5\,\text{GeV}$ is considered to be b-tagged.

In place of E_T^miss, the sum of the momenta of the neutrinos in the leptonic decays of W bosons is used. This is motivated by the fact that E_T^miss is an experimental proxy for the neutrino momentum given that ATLAS does not detect neutrinos, but the specifics of the actual E_T^miss calculation are an experimental detail.

6.3 Event selection

The definition of particle-level events is chosen to be as similar as possible to the event selection requirements at reconstruction level. In both cases, each event must contain exactly two charged leptons with $p_\text{T} > 20\,\text{GeV}$ and opposite electrical charge. At least one of these charged leptons must have $p_\text{T} > 27\,\text{GeV}$. For the reconstruction-level selection, at least one of the leptons must also be associated with the trigger object from the single lepton trigger. This is established by checking that an electron lies within a cone of $\Delta R_y = 0.07$ of the trigger object, or in the case of a muon that it lies within a cone $\Delta R_y = 0.1$ of the trigger object.

At reconstruction level, contributions from $Z + j$, $W + j$, and diboson background processes are reduced by a set of requirements on the E_T^miss and $M_{\ell\ell}$ as per the discussion in §4.3. No equivalent requirements are imposed at particle level. About one third of the nominal tW MC sample passes the dilepton requirements, about one quarter of the total also pass the backgrounds suppression cuts, and about 6 % of the total sample also passes the jet multiplicity requirements. Taking the same sample at particle level, about half of the events pass the dilepton requirements, and about one quarter of the total also pass the jet multiplicity requirements.

Fake and non-prompt leptons are identified in MC simulated samples by examining the production history of the selected leptons. Any selected lepton that does not originate from the decay of a W or Z boson (possibly via a τ lepton) is labelled as fake and its event is placed in the fake sample. The fake sample consists of all fake events regardless of whether they were originally found in $t\bar{t}$ or tW samples. Due to the other selection requirements there is a very small rate of fake events in the selection.

Events that are used in the measurement must contain exactly one jet, and that jet must be b-tagged. Other regions, including the two-jet regions `2j1b` and `2j2b` which were used in the previous measurement, are here used only to validate modelling of background processes and are not part of the cross-section measurements. The event yields in the regions checked are shown in Fig. 6.1.

[1] Since the mass of the particles is known in MC simulations at particle level the angular distance uses rapidity rather than pseudorapidity.

Figure 6.1: Expected number of events for each signal and background process, with the number of observed events superimposed. The expected yields are normalised to their theoretical cross-sections for the luminosity observed in data. The bins correspond to the total number of jets and the number of b-jets in the events. The bottom panel shows the ratio of observed to predicted event yield with the total systematic uncertainty shown by the hashed area.

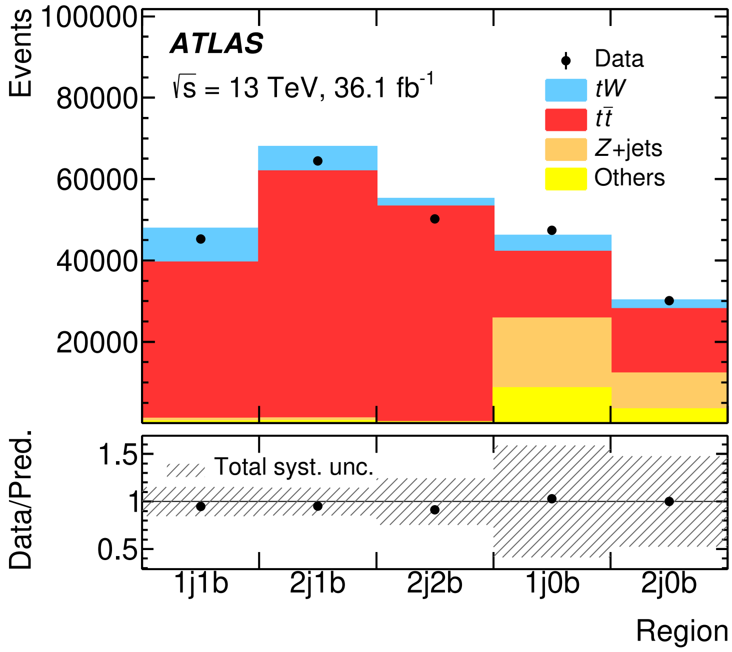

Table 6.2: Summary of the event selection requirements at particle level defining the fiducial selection.

Leptons	Exactly 2 electrons or muons with $p_T > 20\,\mathrm{GeV}$ and $\lvert\eta\rvert < 2.5$, and at least one must satisfy $p_T > 27\,\mathrm{GeV}$.
Jets	Exactly 1 jet with $p_T > 25\,\mathrm{GeV}$ and $\lvert\eta\rvert < 2.5$, and the jet is b-tagged using a ghost matching technique (see §3.6).
E_T^{miss} and $M_{\ell\ell}$	No requirements.

By requiring only a single b-jet, the regions of phase space where the interference between $t\bar{t}$ and tW is greatest are excluded, and therefore the measurement is expected to be less sensitive to the modelling of this interference. The cost of this choice is that many signal events are discarded with the other regions, reducing the statistics available for the measurement. The particle-level event selection requirements defining the fiducial region are summarised in Table 6.2.

The distributions of the six observables (see §6.5) for which differential cross-sections will be measured are shown in Fig. 6.2. In each case the ratio between the observed data and the prediction is close to unity with respect to the total uncertainty. In Fig. 6.2e, there is a slight slope in the ratio panel suggesting that the prediction for $m_T(\ell\ell E_T^{\mathrm{miss}} b)$ is slightly harder than observed.

6.4 Extracting a *tW*-enriched sample

Whereas in the previous chapter the distribution of the responses of a BDT classifiers were used as input to a maximum likelihood fit, the present measurement requires only that a reasonably pure tW sample is available for the cross-section determination. Instead of using the continuous classification score varying between signal-like and background-like, we therefore select a threshold value and discard all events on the background-like side of this threshold.

A BDT classifier is trained and optimised for the `1j1b` region according to the techniques discussed in §4.4. In order to increase the size of the simulated signal sample used for training, the tW DS

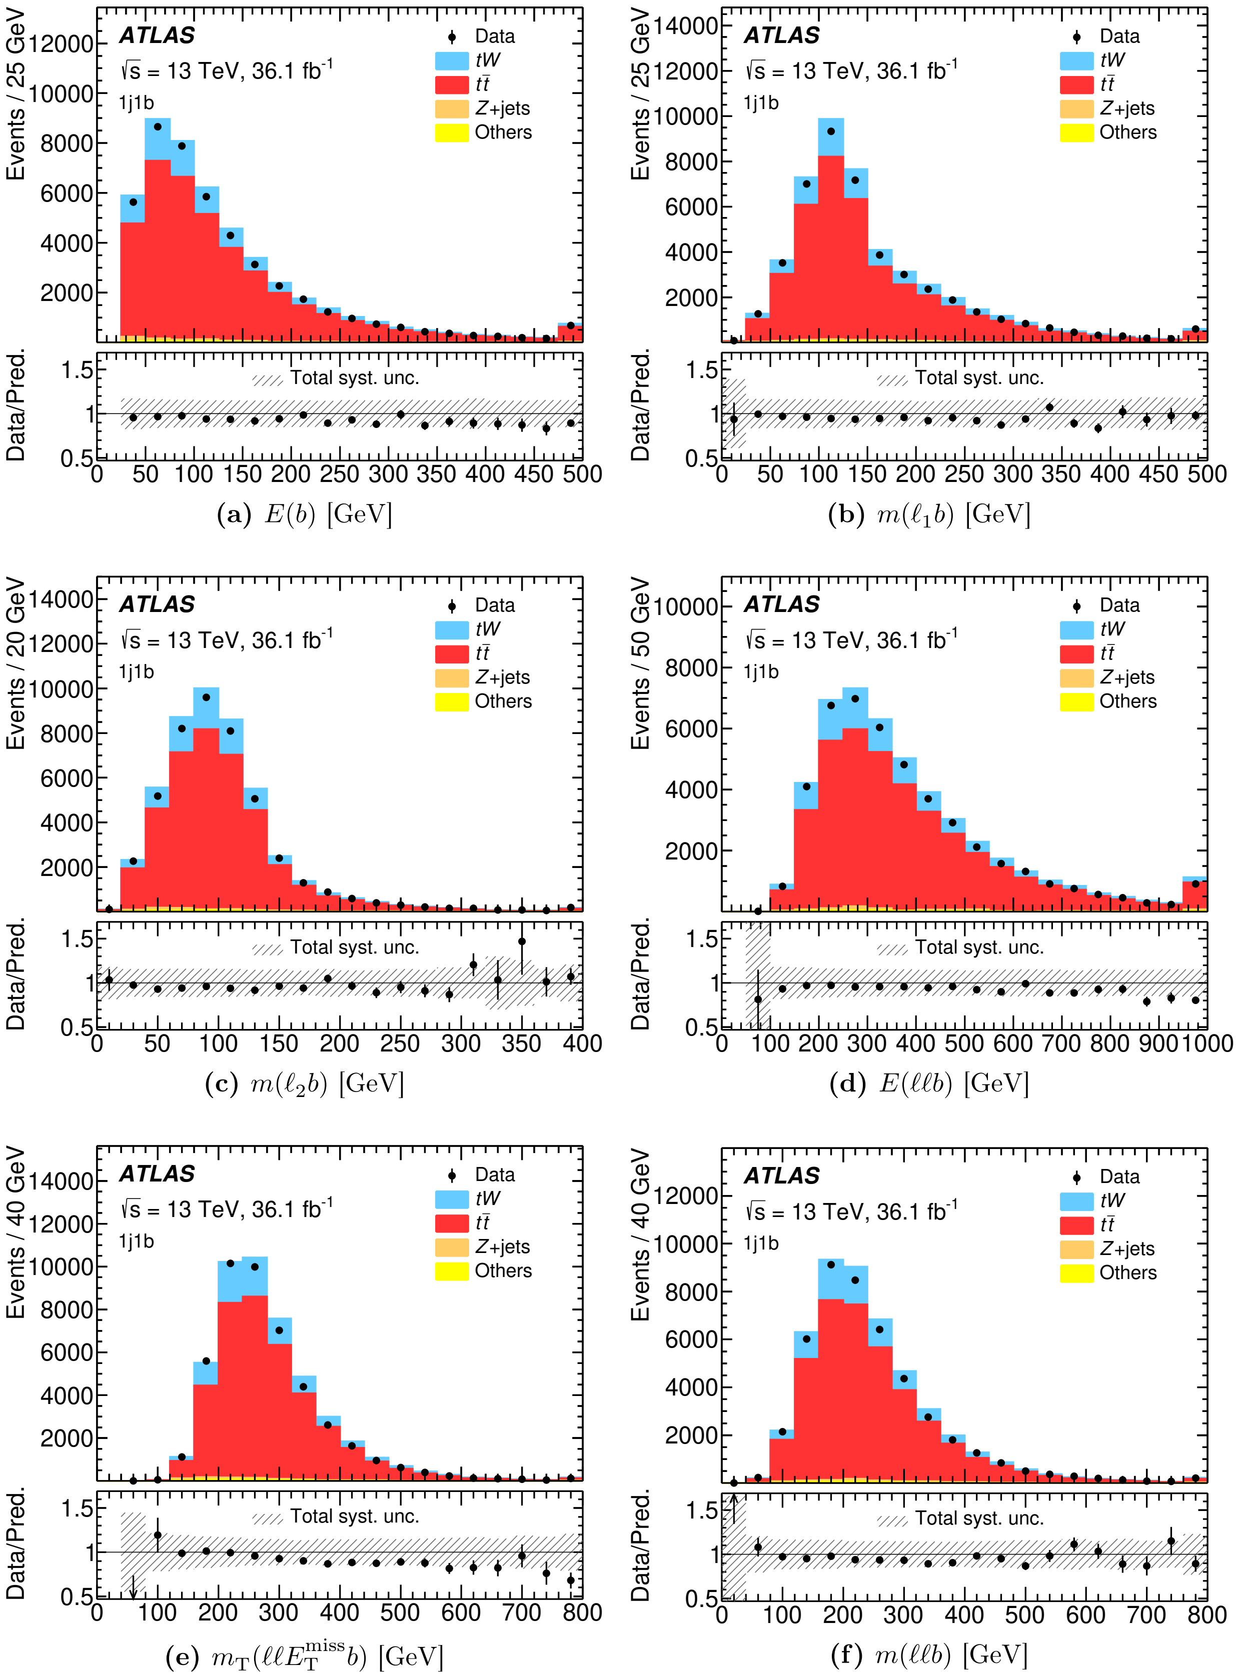

(a) $E(b)$ [GeV]

(b) $m(\ell_1 b)$ [GeV]

(c) $m(\ell_2 b)$ [GeV]

(d) $E(\ell\ell b)$ [GeV]

(e) $m_{\mathrm{T}}(\ell\ell E_{\mathrm{T}}^{\mathrm{miss}} b)$ [GeV]

(f) $m(\ell\ell b)$ [GeV]

Figure 6.2: Distributions of the 6 variables to be unfolded in the `1j1b` region, showing the expected contribution from each process with the observed distributions overlaid and their ratio plotted in the lower panel. The first and last bins include events that respectively underflow or overflow each histogram.

Table 6.3: Features used in the optimised BDTs for the `1j1b` region ranked according to their separation power S (see §4.4.3). The overall separation power of the BDT is shown in the final row.

Observable	S $[10^{-2}]$
$p_{\mathrm{T}}(\ell\ell b E_{\mathrm{T}}^{\mathrm{miss}})$	4.1
$\Delta p_{\mathrm{T}}(\ell\ell b, E_{\mathrm{T}}^{\mathrm{miss}})$	2.5
ΣE_{T}	2.3
$\eta(\ell\ell b E_{\mathrm{T}}^{\mathrm{miss}})$	1.3
$\Delta p_{\mathrm{T}}(\ell\ell, E_{\mathrm{T}}^{\mathrm{miss}})$	1.1
$p_{\mathrm{T}}(\ell\ell b)$	1.0
$C(\ell\ell)$	0.9
$m(\ell_2 b)$	0.2
$m(\ell_1 b)$	0.1
BDT response	8.1

MC sample is combined with the nominal tW sample. The input features provided to the training procedure are the same quantities defined in §4.4.5, with the objects being the highest-p_{T} lepton ℓ_1, the sub-leading lepton ℓ_2, the b-jet b, and $E_{\mathrm{T}}^{\mathrm{miss}}$ (or the combined neutrino 4-vector $\nu\bar{\nu}$ at particle level). The final input features used in the optimised classifier are listed in Table 6.3 along with their separation powers. The feature that allows the greatest separation between tW and $t\bar{t}$ is the p_{T} of the entire system of selected objects. Note that while the selection and ranking of features in similar to the corresponding list for the `1j1b` classifier in the previous chapter (see Table 5.2), some changes are to be expected. This is because of the way in which the BDT optimisation procedure considers correlation between features, meaning that a set of variables might be substituted for another that includes similar kinematic information. The distributions of the selected input features are shown in Fig. 6.3, demonstrating good agreement between the predicted and observed distributions except for the low tail of the ΣE_{T} distribution in Fig. 6.3c, which is covered by the estimated uncertainty.

The response of the final classifier is shown in Fig. 6.4. The $t\bar{t}$ background events account for a progressively smaller fraction of the content of each bin moving from left to right. The BDT response in the observed data is found to be consistent with its response to the simulated MC samples.

Finally, in order to extract a signal-enriched sample, a threshold of 0.3 is set on the BDT response, with events below this threshold discarded. The threshold was determined by testing a range of possible choices and optimising for the lowest total uncertainty on the cross-sections over all bins in the final result. The total systematic uncertainty varies as a complicated function of the BDT response threshold, with some bins in some of the distributions increasing while others decrease due to the dependence of some sources of systematic uncertainties on the parts of the phase space occupied by the selected objects. The total statistical uncertainty reduces as the requirement becomes more permissive due to the higher signal efficiency, for example its value in each bin for a trial threshold of 0.1 is about 60 % of its value with a threshold of 0.4. The chosen value of 0.3 corresponds to the lowest average total uncertainty across all bins of all distributions in the final result. The effects of the BDT requirement on the yield of each process and the observed data are shown in Table 6.4. Whereas about three quarters of the signal events are lost by the requirement, over 90 % of the $t\bar{t}$ events are also discarded, improving the signal-to-background ratio from about 17 % to about 35 %. Importantly, this background reduction is associated with a reduction in the relative importance of the systematic uncertainty due to the background modelling, which is in general large.

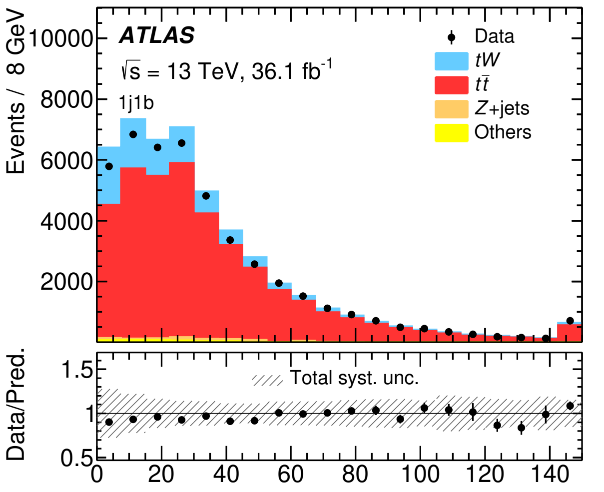

(a) $p_{\mathrm{T}}(\ell\ell b E_{\mathrm{T}}^{\mathrm{miss}})$ [GeV]

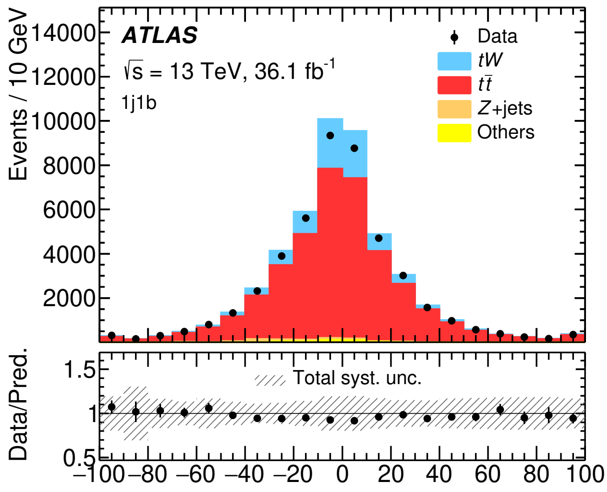

(b) $\Delta p_{\mathrm{T}}(\ell\ell b, E_{\mathrm{T}}^{\mathrm{miss}})$ [GeV]

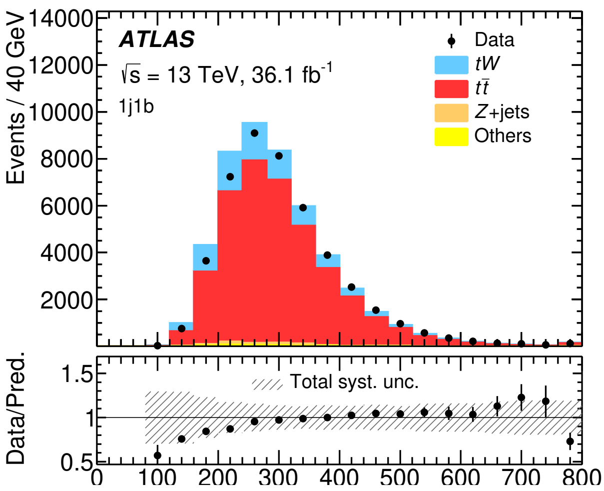

(c) ΣE_{T} [GeV]

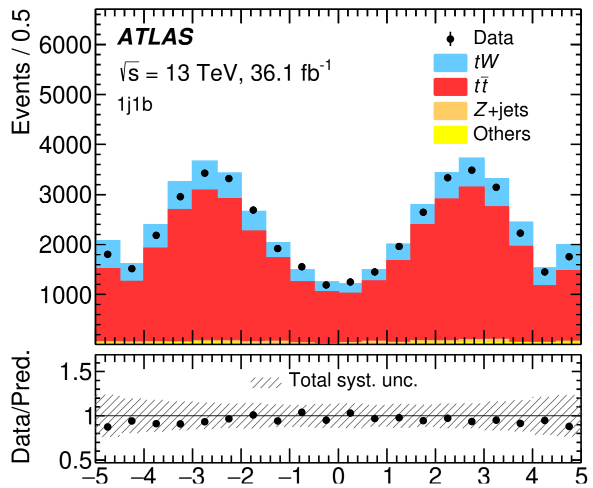

(d) $\eta(\ell\ell b E_{\mathrm{T}}^{\mathrm{miss}})$

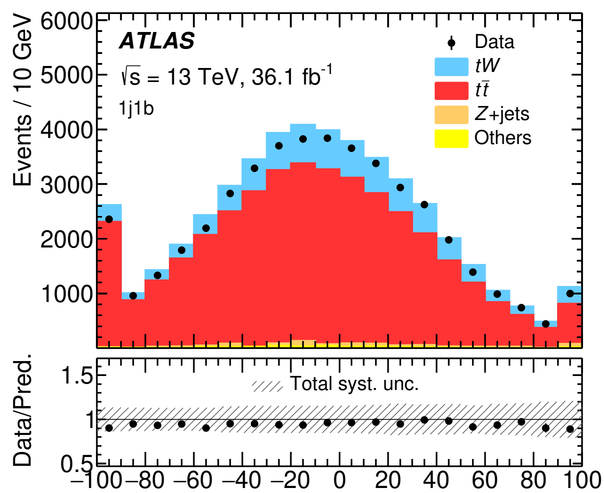

(e) $\Delta p_{\mathrm{T}}(\ell\ell, E_{\mathrm{T}}^{\mathrm{miss}})$ [GeV]

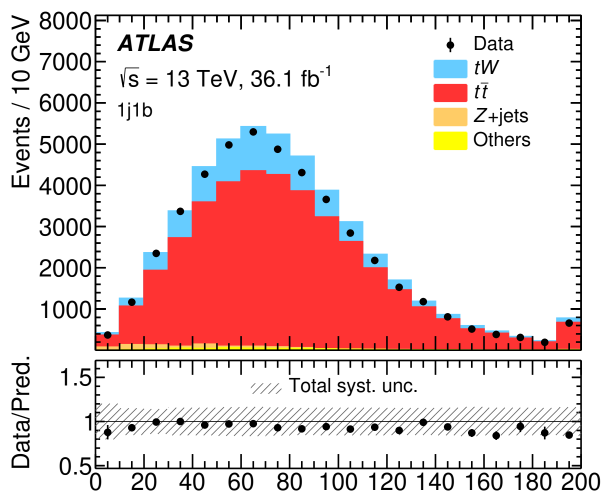

(f) $p_{\mathrm{T}}(\ell\ell b)$ [GeV]

Figure 6.3: Distributions of the 7 most important quantities used as input to the BDT, showing the expected contribution from each process with the observed distributions overlaid and their ratio plotted in the lower panel. The first and last bins include events that respectively underflow or overflow each histogram. The eighth and ninth quantities, $m(\ell_2 b)$ and $m(\ell_1 b)$, are shown above in Fig. 6.2c and Fig. 6.2b respectively.

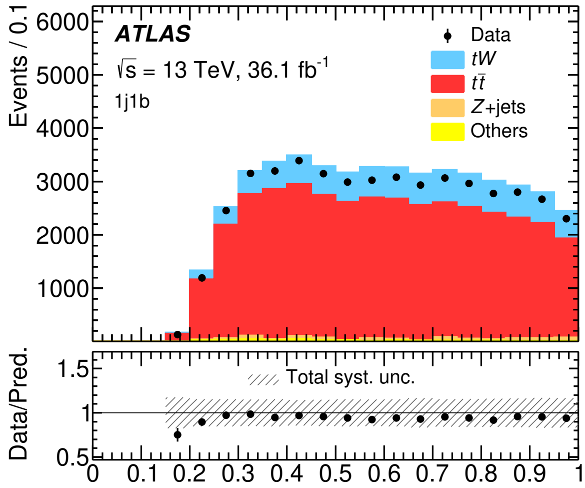

(g) $C(\ell\ell)$

Figure 6.3: (Continued.)

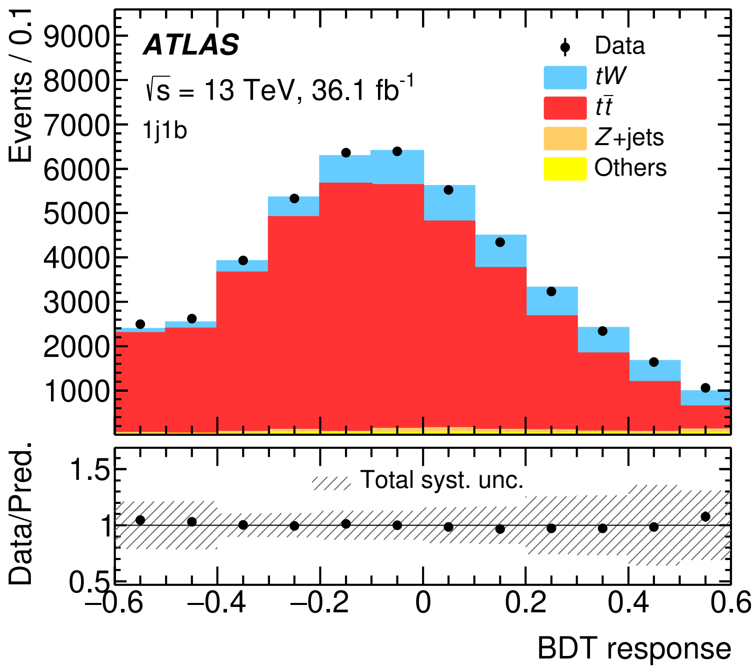

Figure 6.4: Response of the optimised BDT classifier, comparing the response to the observed dataset compared to the simulated MC samples normalised to the expected cross-section. The hashed region on the ratio panel shows the combined contribution of sources of systematic uncertainty. The first and last bins include the underflow and overflow respectively.

Table 6.4: Event yields for the `1j1b` region, showing the observed yields alongside the predicted yields which are decomposed into the contributing signal and background processes. The central column shows the yields before the BDT requirement while the rightmost column shows the yields after imposing the requirement.

Process	Events	Events BDT response > 0.3
tW	8300 ± 1400	1970 ± 560
$t\bar{t}$	38 400 ± 6600	3400 ±1300
$Z + j$	620 ± 310	159 ± 80
Diboson	230 ± 58	81 ± 20
Fakes	220 ± 220	19 ± 19

Figure 6.5 shows the distribution of the quantities that will be used in the cross-section measurement, with the BDT requirement imposed. Compared to the distributions before the requirement (see Fig. 6.2), the higher signal purity is evident, while the agreement between prediction and observation remains generally good. In the tails of the distributions of $m(\ell_2 b)$ and $m(\ell\ell b)$ (Figs. 6.5c and 6.5f) the agreement is noticeably poorer, however with the binning used in these figures the occupancy of the tails is very small and the uncertainty is correspondingly large. For the cross-section measurement the binning is optimised to reduce these statistical effects. In selecting the quantities for the cross-section measurements, it is important to avoid bias introduced by requirements placed on the BDT response. Any observables that correlate highly with inputs that contribute significantly to the BDT response are therefore unsuitable as candidates for unfolding to particle level. After considering this restriction the only two observables chosen for unfolding that are also used as input to the BDT are $m(\ell_1 b)$ and $m(\ell_2 b)$ which are the least significant inputs and were found not to be subject to the shaping effects that exclude other quantities.

6.5 Correcting for detector effects

The process of correcting the observed distributions to estimate the corresponding true distributions, moving from reconstruction-level quantities to particle-level quantities, is called *unfolding*[2] [125, p. 203]. Experimentally measured quantities have features smeared by resolution effects, in addition to residual measurement effects after applying calibration procedures. The corrections applied by unfolding therefore typically result in the narrowing of peaks and stretching or shifting of certain parts of distributions. Since the data are binned, these effects can all be described by the bin-to-bin migration of events. If such migrations are well-understood, a series of correction factors can be written down that allow a particle-level distribution to be estimated from a reconstruction-level one.

Naturally the necessary correction factors are unknown since they could otherwise be removed by further calibration, so the task of unfolding is to estimate them using all available knowledge of the detector's response. When using MC simulated samples, the true underlying distributions are known, so, by comparing the reconstruction- and particle-level information in those samples, it is possible to estimate the migration of events between bins when moving from the particle-level quantities to the reconstruction-level quantities. These bin-to-bin migrations can be described by a response matrix, M, that encodes the response of the detector. The elements M_{ij} of M are interpreted as the probability that an event measured in bin i of the reconstruction-level distribution has a true value for the particle-level quantity that falls in bin j of that distribution. The off-diagonal terms of M therefore have the effect of smearing since they correspond to events moving from a bin to its neighbours.

Given the vector $\mathbf{s}$ describing the number of signal events in each bin of the particle-level quantity, and $\mathbf{b}$ describing the number of background events, the expected number of events in each reconstruction-level bin, $\mathbf{n}$, is given by the matrix relation

$$\mathbf{n} = M\mathbf{s} + \mathbf{b}.$$

To unfold, an estimator $\hat{\mathbf{s}}$ must be found for the expected number of signal events. If the reconstruction-

[2] In other contexts this process is also called deconvolution or unsmearing.

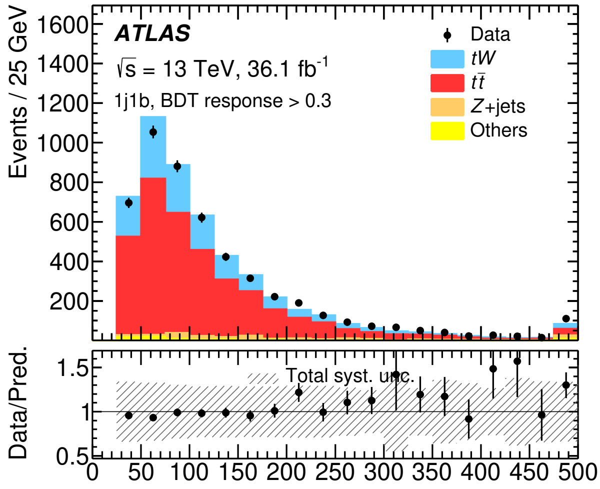

(a) $E(b)$ [GeV]

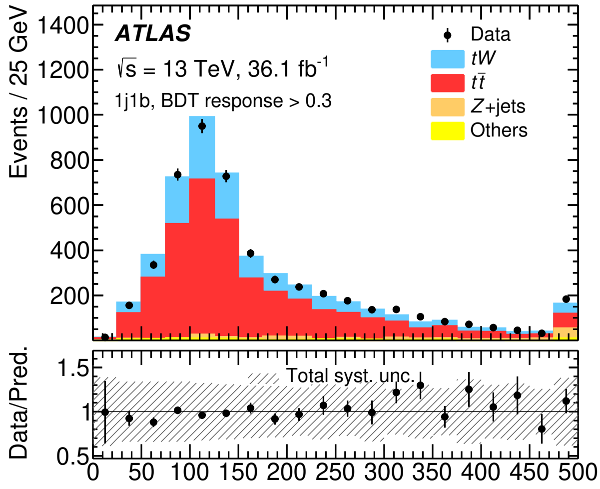

(b) $m(\ell_1 b)$ [GeV]

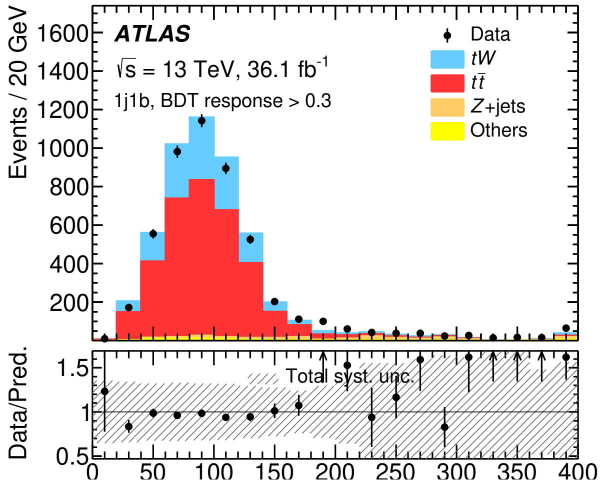

(c) $m(\ell_2 b)$ [GeV]

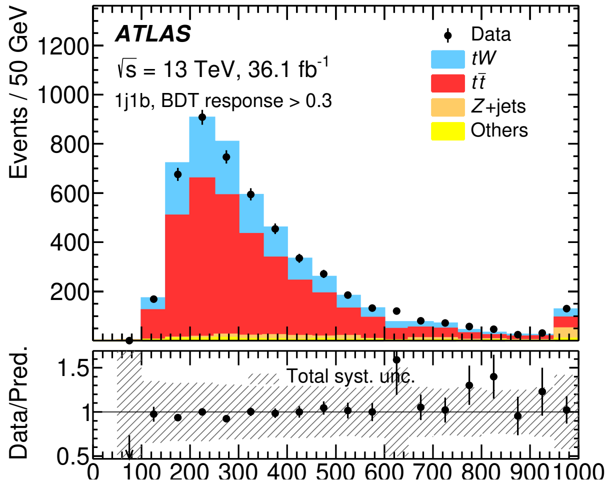

(d) $E(\ell\ell b)$ [GeV]

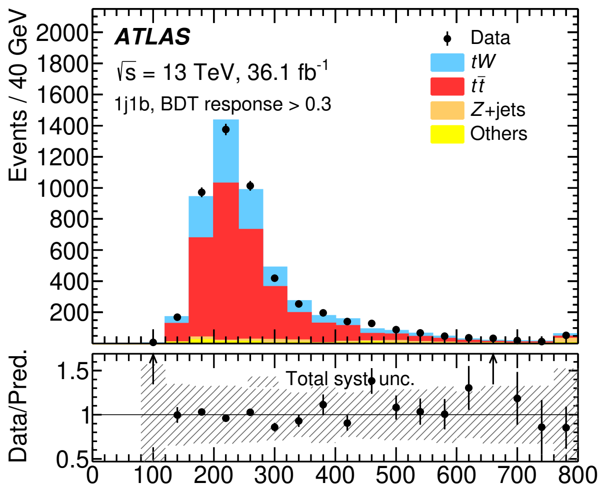

(e) $m_{\mathrm{T}}(\ell\ell E_{\mathrm{T}}^{\mathrm{miss}} b)$ [GeV]

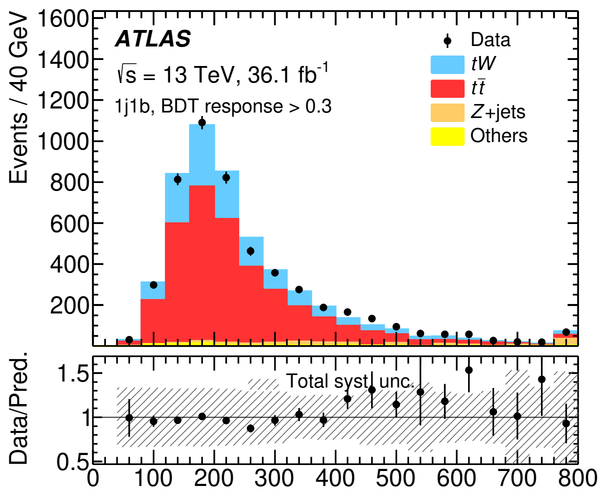

(f) $m(\ell\ell b)$ [GeV]

Figure 6.5: Distributions of the 6 variables to be unfolded after applying the BDT requirement, showing the expected contribution from each process with the observed distributions overlaid and their ratio plotted in the lower panel. The first and last bins include events that respectively underflow or overflow each histogram.

and particle-level quantities have the same number of bins, one such estimator can be constructed by inverting the response matrix:

$$\hat{\mathbf{s}} = M^{-1}(\hat{\mathbf{n}} - \mathbf{b}),$$

where $\hat{\mathbf{n}}$ is an estimator of the total number of events in the reconstruction-level bins. The obvious choice for $\hat{\mathbf{n}}$ is the yields observed experimentally, and this indeed results in an unbiased estimator $\hat{\mathbf{n}}$ that corresponds to the solution obtained by maximising a likelihood function constructed by treating the yield in each bin as an independent Poisson variable.

In practice unfolding is not typically implemented as a matrix inversion because of issues related to singularity and the constraints it imposes on the relationship between the binning of the reconstruction- and particle-level quantities. The matrix inversion approach often leads to pathological situations such as that described in Ref. [121, p. 160] where the estimator $\hat{\mathbf{s}}$ exhibits very large uncertainties and correlations of approximately -1 between neighbouring elements. Such behaviour arises from the fact that the experiment is only performed once, and so the yields are not precise estimators of $\mathbf{n}$. Statistical fluctuations in the yields are amplified by the application of M^{-1} because they are details of the fine structure of the distribution, and the response matrix is deliberately sensitive to such fine structure which is necessary to perform the unfolding. In general more sophisticated techniques are necessary for a practical unfolding procedure.

6.5.1 Iterative Bayesian unfolding

The unfolding approach adopted here follows the iterative Bayesian method of Ref. [124]. In the terminology of this method, the observation of an event in a particular reconstruction-level bin is called an *effect*, and is due to an underlying *cause* which is that the event belonged to a particular true particle-level bin (including cases where the cause is due to a background process). The probability of a cause C_i having been responsible for an observed effect E_j is then given by an expression derived [125, p. 205] from Bayes' theorem:

$$P\left(C_i \,|\, E_i\right) = \frac{P\left(E_j \,\middle|\, C_i\right) P_\circ(C_i)}{\sum_{l=1}^{n_C} P\left(E_j \,\middle|\, C_l\right) P_\circ(C_l)}. \tag{6.1}$$

The $P_\circ(C_i)$ terms are the appropriately-normalised priors for each of the n_C causes and can be taken as uniform. The $P\left(E_j \,\middle|\, C_i\right)$ terms are the probability of observing an effect E_j given a cause C_i, and these correspond to the M_{ji} elements of the response matrix discussed above. Each effect must be attributable to at least one cause, but a given cause does not necessarily produce any effect.

The collection of N_{obs} observations results in the observed yields $\vec{n}(E) \equiv \left\{n(E_j)\right\}_{j=1}^{n_E}$ for the n_E possible effects. The best estimate for the number of events in a true particle-level bin is

$$\hat{n}(C_i) = \frac{1}{\epsilon_i} \sum_{j=1}^{n_E} n(E_j) P\left(C_i \,\middle|\, E_j\right)$$

where $\epsilon_i \equiv \sum_{j=1}^{n_E} P\left(E_j \,\middle|\, C_i\right)$ is the efficiency of finding cause C_i for any possible effect. From these unfolded yields, the probabilities $\hat{P}(C_i) \equiv P\left(C_i \,|\, \vec{n}(E)\right)$ can be calculated by taking the ratio of each $\hat{n}(C_i)$ to the total number of unfolded events. These probabilities can be considered to be an updated prior given the data. The method is then to repeat the process, substituting $\hat{P}(C_i)$ where $P_\circ(C_i)$

appears in Eq. (6.1). Each repetition is a new iteration of the method, though repeatedly updating the prior using the same observations means that the technique is not strictly Bayesian in character[3].

The iterative Bayesian unfolding technique is well-suited to the particle physics application. The formulation results in a well-defined relationship between the covariance matrix between bins after unfolding and the equivalent matrix before unfolding, allowing careful treatment of sources of uncertainty. Background processes can be incorporated simply by adding new causes. The only free parameter in the technique is the number of iterations to be performed. In the limit of a very large number of iterations, neglecting pathological cases, the method recovers the true distribution. In this limit however the method is subject to the same deficiency as the matrix inversion approach whereby statistical fluctuations are amplified and the unfolded distributions are useless. In practice, the results of the unfolding are typically good after as few as three iterations, and by halting early the sensitivity to statistical fluctuations in the observations is reduced[4].

6.5.2 Unfolding implementation

The iterative Bayesian approach is implemented in the RooUnfold [2] software package. The above discussion has implicitly assumed that for each MC simulated event both a set of reconstruction-level observables and a set of particle-level observables exist. In reality due to the event selection requirements it is possible for an event to be accepted at one level and rejected at the other. This could occur for example if a particle-level electron with $p_{\mathrm{T}} = 20.2\,\mathrm{GeV}$ is reconstructed as an electron with $p_{\mathrm{T}} = 19.9\,\mathrm{GeV}$ due to the failure to reconstruct and include a small part of its EM shower in the calorimeter, thereby causing the event to fail the reconstruction-level selection. In order to account for such events, two correction factors are calculated for each bin. C_j^{oof} is the out-of-fiducial correction for reconstruction-level bin j, and is defined to be ratio of the number of events in that bin that are accepted at both reconstruction and particle level ($N_j^{\mathrm{p+r}}$) to the number of events in that bin that are only accepted at reconstruction level. C_i^{eff} is the efficiency to reconstruct a fiducial event for particle-level bin i, and is defined to be the ratio of $N_i^{\mathrm{p+r}}$ to the number of events in that bin that are only accepted at particle-level.

The complete unfolding procedure is then applied as

$$N_i^{\mathrm{ufd}} = \frac{1}{C_i^{\mathrm{eff}}} \sum_j M_{ij}^{-1} C_j^{\mathrm{oof}} \left(N_j^{\mathrm{obs}} - B_j \right) \tag{6.2}$$

where N_i^{ufd} is the number of signal events in bin i after unfolding, N_j^{obs} is the number of observed events in bin j, and B_j is the expected number of background events in bin j. Note that the unfolding matrix M_{ij}^{-1} does not involve actual matrix inversion and is rather intended to represent the application of the method of §6.5.1. Note also that the expected contribution from background events is subtracted before unfolding.

The boundaries of the bins used for each reconstruction- and particle-level quantity when unfolding are not constrained by the method. Prudent choice of the boundaries can however assist in reducing the overall uncertainty and further mitigating the amplification of statistical fluctuations. An ideal binning should be larger than the intrinsic resolution of the detector (and therefore of the smearing

[3]For this reason the technique is sometimes called Bayesian-inspired unfolding or simply D'Agostini unfolding.

[4]This feature of the method responsible for reducing the dependence on random fluctuations is called *regularisation*. Most practical unfolding techniques offer some mechanism for regularisation to avoid the problems discussed earlier.

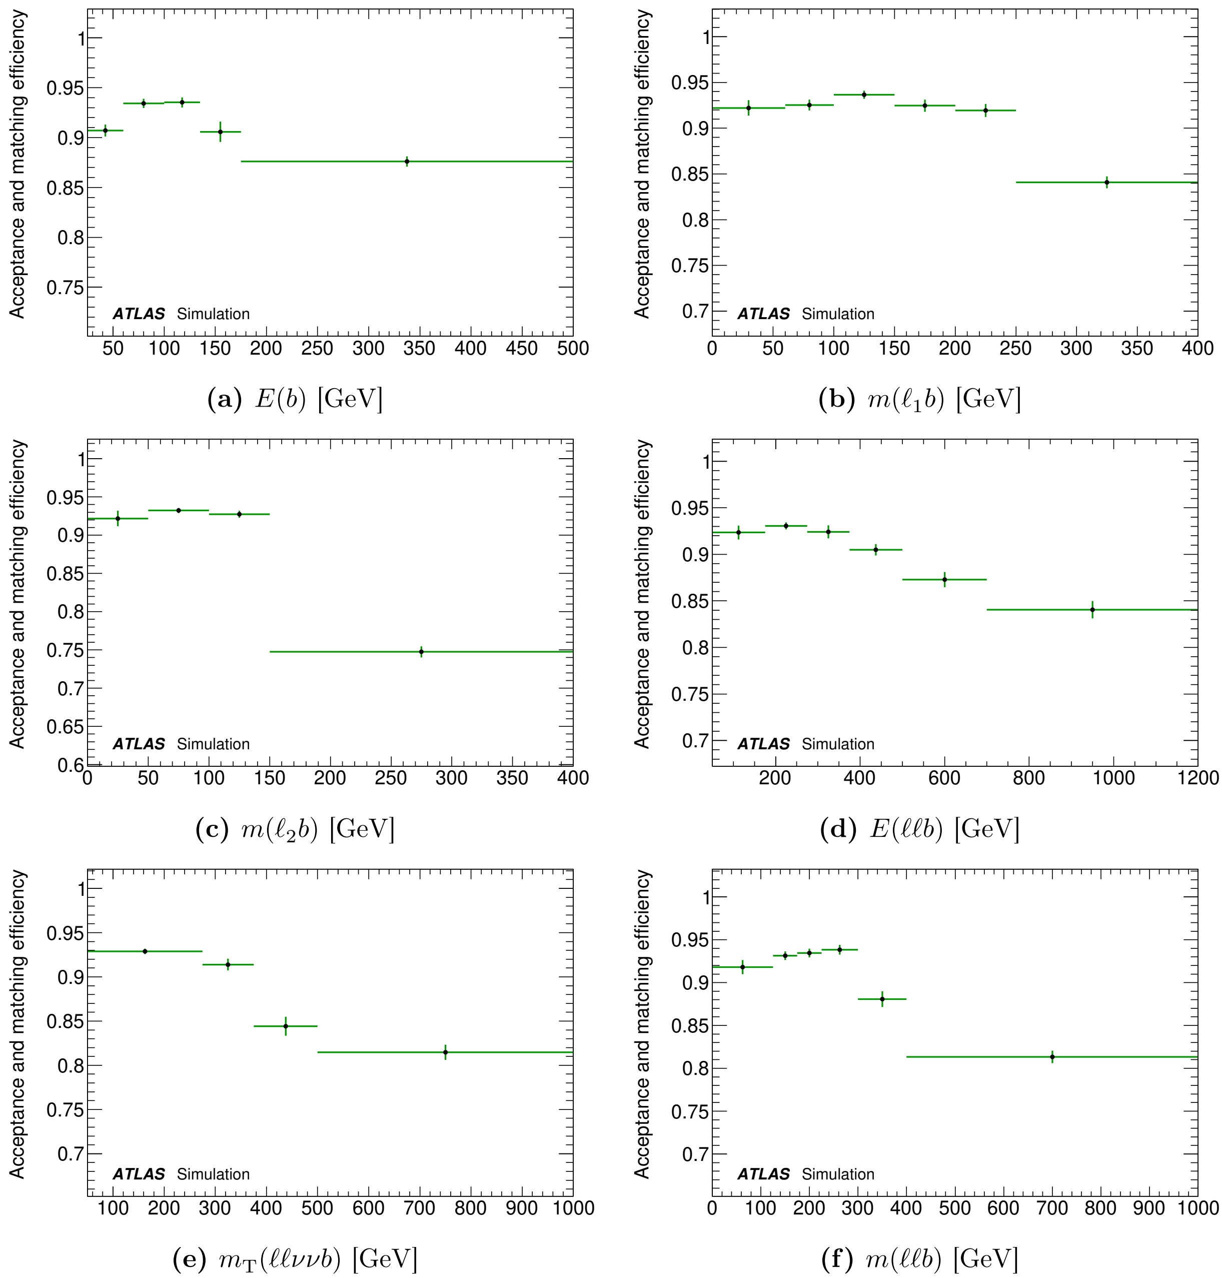

(a) $E(b)$ [GeV] **(b)** $m(\ell_1 b)$ [GeV]

(c) $m(\ell_2 b)$ [GeV] **(d)** $E(\ell\ell b)$ [GeV]

(e) $m_{\mathrm{T}}(\ell\ell\nu\nu b)$ [GeV] **(f)** $m(\ell\ell b)$ [GeV]

Figure 6.6: Out-of-fiducial corrections, C_j^{oof}, for each variable that is unfolded. The value in each bin is calculated using the MC simulated samples and constructing the ratio of the number of events that are accepted at both particle and reconstruction level to the number that are accepted at reconstruction level only. The error bars correspond to the statistical uncertainty arising from the finite size of the simulated samples used for this procedure.

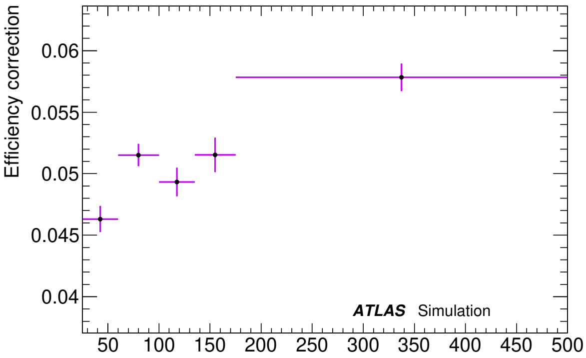

(a) $E(b)$ [GeV]

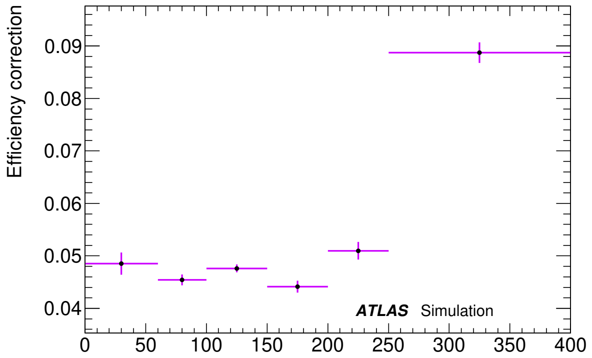

(b) $m(\ell_1 b)$ [GeV]

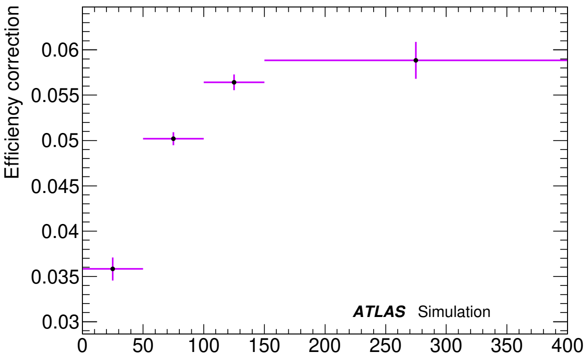

(c) $m(\ell_2 b)$ [GeV]

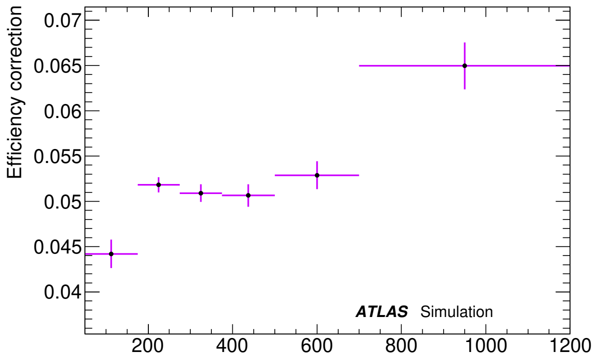

(d) $E(\ell\ell b)$ [GeV]

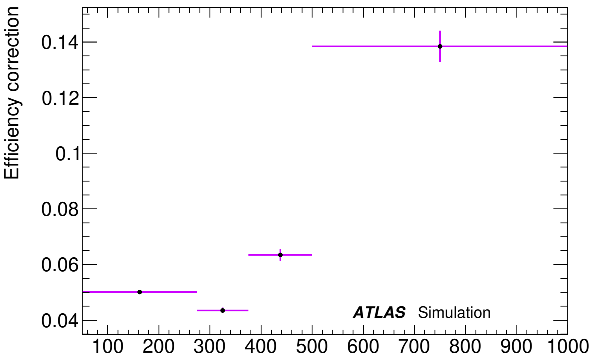

(e) $m_{\mathrm{T}}(\ell\ell\nu\nu b)$ [GeV]

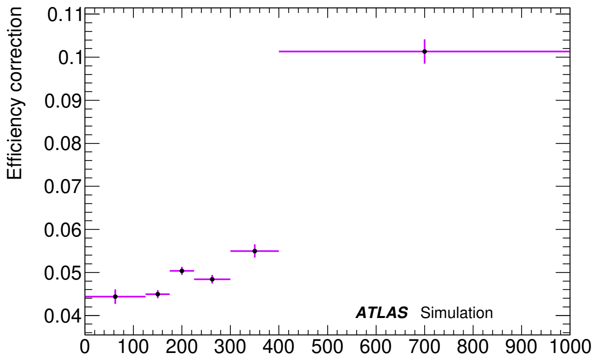

(f) $m(\ell\ell b)$ [GeV]

Figure 6.7: Efficiency corrections, C_i^{eff}, for each variable that is unfolded. The value in each bin is calculated using the MC simulated samples and constructing the ratio of the number of events that are accepted at particle level only to the number that are accepted at both particle and reconstruction level. The error bars correspond to the statistical uncertainty arising from the finite size of the simulated samples used for this procedure.

Observable	Number of iterations
$E(b)$	15
$m(\ell_1 b)$	7
$m(\ell_2 b)$	5
$E(\ell\ell b)$	5
$m_{\mathrm{T}}(\ell\ell\nu\nu b)$	7
$m(\ell\ell b)$	5

Table 6.5: The number of iterations of the iterative Bayesian unfolding method applied to each observable.

effect), should minimise the amount of bin-to-bin migration in the unfolding, and should minimise the statistical uncertainty on each reconstruction-level bin. To this end we adopt bin boundaries that result in no more than about 40 % of the unfolding matrix being off-diagonal, and no more than about 20 % statistical uncertainty in each bin. Bins at the tails of distributions are widened to incorporate enough events that the statistical uncertainty is small, and they extend far enough that no underflow or overflow bin is required. Other bins are constrained to have widths that are a small multiple of the average resolution of the observable, as determined by the width of fitted Gaussian distributions in simulated samples. Where possible the same configuration is used for reconstruction and particle level.

Using the optimised binning configuration, the correction factors C_j^{oof} and $1/C_i^{\mathrm{eff}}$ are shown in Figs. 6.6 and 6.7 respectively. The out-of-fiducial corrections are greater than 0.9 in most bins, corresponding to a small probability of reconstructing events that are not accepted at particle level. The efficiency corrections are less than 0.06 in most bins, corresponding to a small probability that events accepted at particle level are not accepted at reconstruction level. In both cases this is helped by the event selection criteria having been chosen similarly at each level, and chosen such that reconstruction performance is good for the accepted events.

In order to measure the bias of the unfolding procedure, tests are performed where half of the simulated sample is reserved while the other half is used to prepare a test M_{ij}^{-1}. Once this test unfolding matrix is constructed, it is used to unfold the reserved half of the simulated sample and the resulting unfolded distributions are compared to the known particle-level distributions. Repeating this process for several random partitions of the simulated sample, the average statistical uncertainty and the average bias per bin (the relative difference between the unfolded quantities and the known values) are estimated as a function of the number of iterations as illustrated in Fig. 6.8. The statistical uncertainty rises with successive iterations in accordance with the arguments presented earlier. The magnitude of the bias initially increases with successive iterations before arriving at a plateau, suggesting that subsequent iterations offer no further improvement in removing detector effects. The optimal number of iterations is picked by balancing the requirements that the uncertainty be as small as possible while ensuring that the plateau in bias has been reached. The number of iterations for each observable is shown in Table 6.5.

The unfolding matrix for each observable is shown in Fig. 6.9, with the values normalised for ease of comparison. The unfolding matrix for $E(b)$ (Fig. 6.9a) has a larger off-diagonal component than the other observables which slows the convergence of iteration by the above criteria, requiring a larger number of iterations compared to the other observables. This is reflected in the gentle slopes of Fig. 6.8a compared to the other observables which each exhibit a turn-on curve followed by a plateau as in Fig. 6.8b.

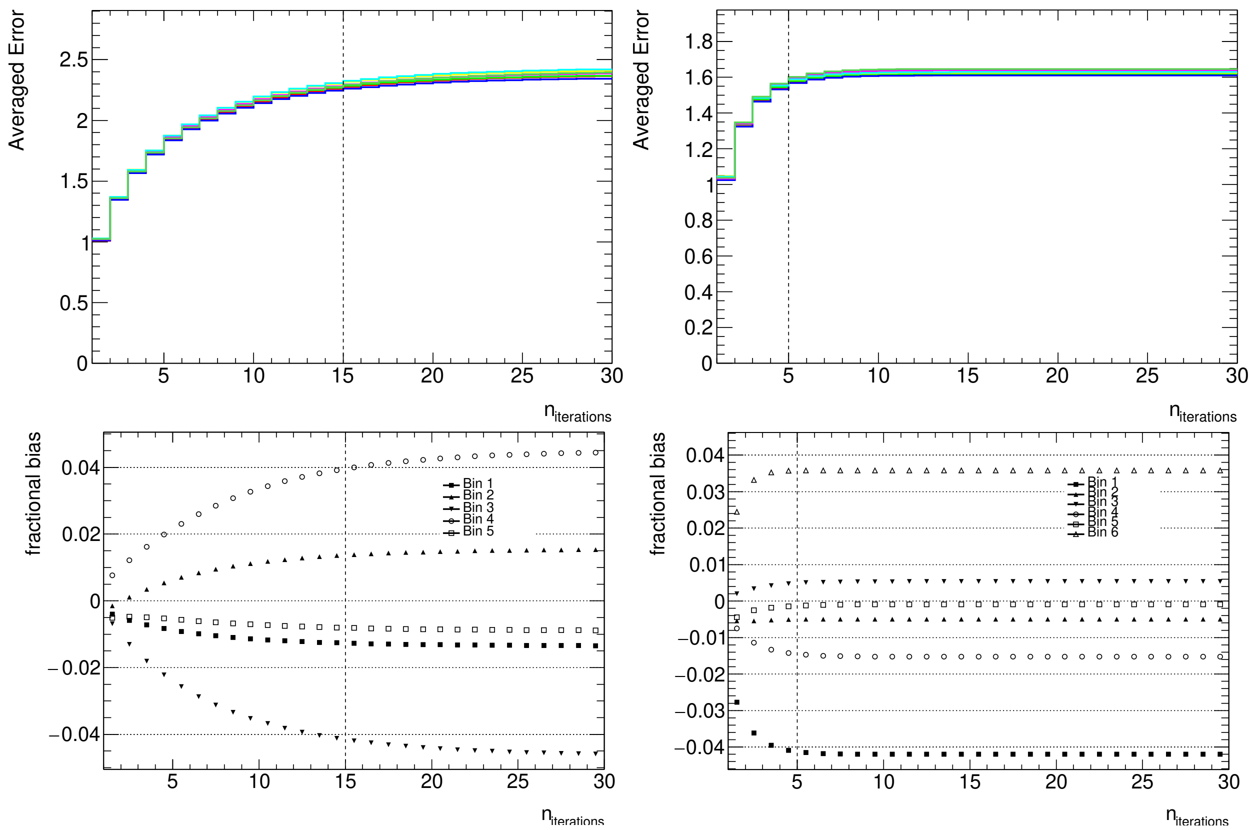

(a) $E(b)$ **(b)** $E(\ell\ell b)$

Figure 6.8: Optimisation of the number of unfolding iterations for 2 of the 6 observables. The top row shows the average statistical uncertainty across all bins of the quantity as a function of the number of iterations. The 8 different coloured lines (almost superimposed) correspond to different random partitions of the data, randomly selected from the 1000 such partitions tested. The bottom row shows the bias in each individual bin of the quantity as a function of the number of iterations.

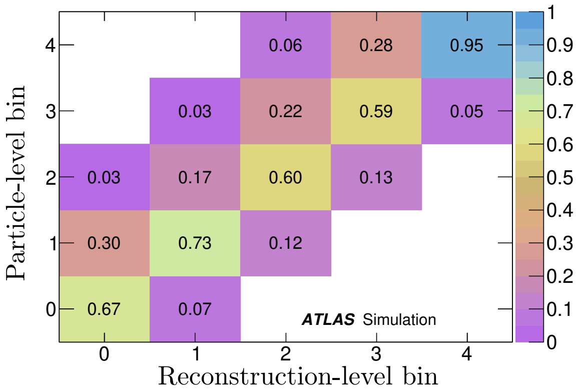

(a) $E(b)$ [GeV]

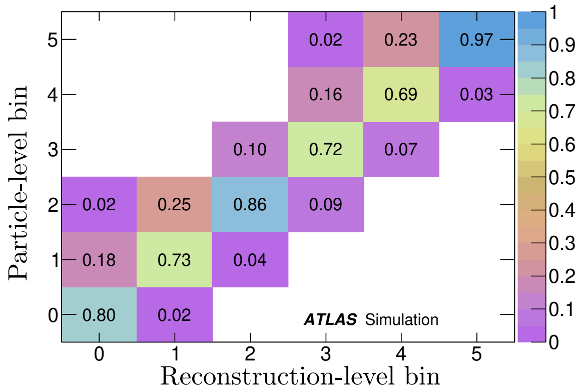

(b) $m(\ell_1 b)$ [GeV]

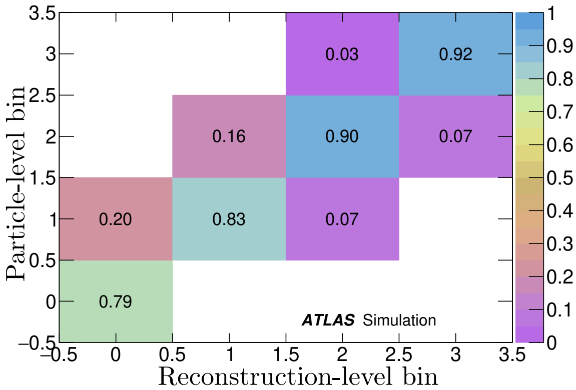

(c) $m(\ell_2 b)$ [GeV]

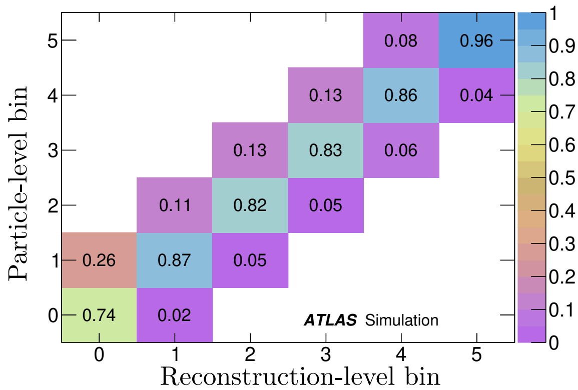

(d) $E(\ell\ell b)$ [GeV]

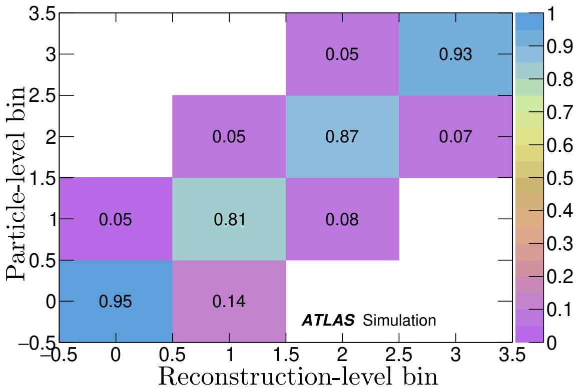

(e) $m_{\mathrm{T}}(\ell\ell\nu\nu b)$ [GeV]

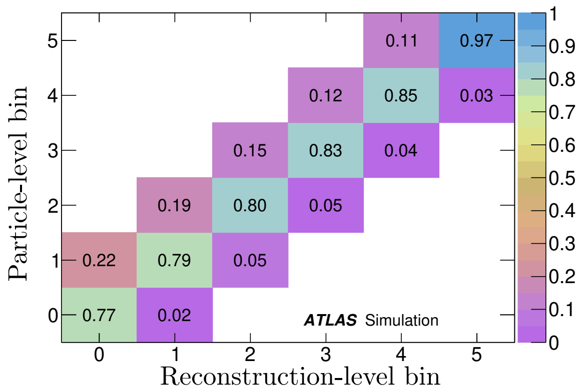

(f) $m(\ell\ell b)$ [GeV]

Figure 6.9: Unfolding matrices, M_{ij}^{-1}, for each variable that is unfolded. The elements of the matrix are calculated by the iterative Bayesian unfolding method using MC simulated samples. Each matrix is normalised so that the sum of each of its columns is unity.

6.5.3 Cross-section determination

The unfolded signal yields, N_i^{ufd}, are calculated using the procedure of Eq. (6.2). For a quantity X, the differential cross-section is then obtained using the width Δ_i of particle-level bin i (for quantity X) and the integrated luminosity of the data, L, as

$$\frac{\mathrm{d}\sigma_i}{\mathrm{d}X} = \frac{N_i^{\mathrm{ufd}}}{L\Delta_i}.$$

This quantity is the absolute cross-section differential in X. The total fiducial cross-section σ_{fid} can be recovered as the sum over the bins of any differential cross-section as

$$\sigma_{\mathrm{fid}} = \sum_i \left(\frac{\mathrm{d}\sigma_i}{\mathrm{d}X} \cdot \Delta_i \right) = \sum_i \frac{N_i^{\mathrm{ufd}}}{L}.$$

The normalised cross-section $\frac{1}{\sigma_{\mathrm{fid}}}\frac{\mathrm{d}\sigma_i}{\mathrm{d}X}$ then describes the shape of the quantity X, allowing for cancellations in systematic effects that simultaneously affect the total cross-section and the bin-by-bin absolute cross-sections, namely those that affect the normalisation.

6.6 Sources of uncertainty

The sources of experimental and modelling systematic uncertainty are described according to the discussion in §4.6 and §4.7. Whereas in the previous chapter initial estimates of uncertainty were improved by a maximum likelihood fit which also calculated the correlations between sources, here no such fit is performed. The uncertainty on the integrated luminosity is $\pm 2.1\,\%$ as determined using a VdM scan performed in August 2015 following the methodology of Ref. [38].

In this measurement the $t\bar{t}$ cross-section is not simultaneously measured, but its value must be known in order to find the expected number of background events in each bin when subtracting the backgrounds. Its nominal value is fixed to the SM prediction, and an uncertainty of 5.5 % is assigned in order to cover theoretical uncertainty on the scale and PDF. Similarly the smaller background processes are normalised according to their SM predictions. A 50 % uncertainty is assigned to the normalisation of $Z + j$ and a 25 % uncertainty is assigned to the normalisation of the diboson process. These uncertainties were chosen conservatively to fully cover differences in modelling of the processes by different MC generators [48].

An additional uncertainty is included to account for the effect of the finite size of the simulated samples on the construction of the unfolding matrix. This is estimated with a bootstrap method, meaning that an ensemble of pseudo-datasets is created by sampling from the full MC simulated dataset, with replacement. For each such pseudo-dataset the unfolding process is followed and new values for C_j^{oof}, C_i^{eff}, and M_{ij}^{-1} are calculated, and these are then used to unfold the full nominal MC sample. Once this is done for all of the pseudo-datasets, there is a distribution of unfolded values for every bin. The width of this distribution is then taken to be that bin's uncertainty due to limited MC statistics. The size of this uncertainty is typically 1 % to 3 %.

To account for uncertainty in the unfolding procedure itself, a series of tests are performed. In each test a Gaussian or linear function is used to reweight the nominal MC sample, and then the resulting sample is unfolded. If the unfolding procedure performed ideally, the injected function

Table 6.6: Samples used in the propagation of different types of systematic uncertainties. tW is the signal sample and $t\bar{t}$ is the background sample. A subscript of 0 indicates the nominal sample, whereas subscripts of + or − indicate the upwards and downwards systematic variation respectively. In each case the reconstruction-level sample is unfolded and compared to the truth information in the particle-level sample.

Type of systematic	Reconstruction-level sample		Particle-level sample
Signal modelling	$tW_\pm$	−	$tW_\pm$
Background modelling	$tW_0 + (t\bar{t}_\pm - t\bar{t}_0)$	−	tW_0
Experimental	$tW_\pm + (t\bar{t}_\pm - t\bar{t}_0)$	−	tW_0
PDF modelling	$tW_\pm + (t\bar{t}_\pm - t\bar{t}_0)$	−	$tW_\pm$

should be recovered perfectly in its effect on the particle-level distribution. Wherever a bias is observed in these tests, the magnitude of the bias is included as a systematic source referred to as unfolding non-closure. In many bins of the unfolded variables no significant bias is observed and so this effect is omitted. The effect is largest in bins covering the extremes of the distributions, where it can reach the 10 % level.

6.6.1 Propagation of uncertainty

The estimated uncertainties on the reconstruction-level quantities are propagated through the unfolding procedure to determine the uncertainties on the measured particle-level quantities. A choice is made as to which uncertainties should be treated as fully correlated or fully uncorrelated between the signal and $t\bar{t}$ background samples, for uncertainties which affect both the signal and background samples. Experimental sources are chosen to be fully correlated between samples since they are related to detector effects that are essentially independent of the physics process. The ME generator uncertainty is chosen to be uncorrelated, whereas the PS generator uncertainty is chosen to be correlated. The uncertainty related to additional radiation and the scale choice are chosen to be uncorrelated. These choices are based on investigations in Ref. [50] and were checked to be no less conservative than alternative choices.

For each source of uncertainty, the propagation is done by choosing reconstruction-level samples that incorporate the relevant variation, unfolding them, and comparing to the known truth information at particle level. For example, in the case of the parton shower modelling for tW, the reconstruction-level predictions from the signal sample prepared with the alternative generator are unfolded. The resulting yield (which is now at particle level since it has been unfolded) is then compared to the "true" particle-level predictions from the same alternative sample. The difference between these predictions is then taken as the propagated uncertainty. The choice of samples for the purpose of these comparisons depends on the type of uncertainty since experimental uncertainties affect only the unfolded side of the comparison whereas modelling uncertainties also affect the true particle-level predictions. In the comparisons the nominal unfolding matrix is fixed and only the input spectra are modified, with the exception of those cases where the AtlfastII sample is used for the nominal sample, in which case an unfolding matrix prepared using that sample is used instead. In this way differences in the detector simulation are never incorporated into other comparisons.

For signal modelling sources (tW ME and PS generators, additional radiation, and the subtraction scheme for the interference term), the reconstruction-level variation samples are unfolded and compared to the truth information at particle-level in the same variation. For sources that affect the background modelling ($t\bar{t}$ ME and PS generators, additional radiation, and background normalisations), the difference between the predictions for the background contribution in the nominal and variation background samples are added to the nominal signal prediction at reconstruction level. This is then unfolded and compared to the nominal signal prediction at particle-level. For experimental sources the same procedure is followed except that the signal variation samples are used at reconstruction-level since these sources affect both the signal and the background. For PDF modelling uncertainties a technical limitation prevents the same procedure from being followed since the POWHEG generator does not currently support PDF reweighting, and so the MG5_aMC@NLO sample is used for the nominal and the variations. The varied PDF modelling is then unfolded and compared to the equivalent signal predictions at particle level. The choice of samples is summarised in Table 6.6.

Figure 6.10 summarises the relative contribution of each category of uncertainties to each bin of the 6 unfolded distributions at particle level. The largest contributions are due to statistical fluctuations related to the small data sample collected in each bin, the signal and background modelling, and the JES calibration. The differential cross-sections in this chapter are normalised to the fiducial cross-section, meaning that the measurement in each bin is divided by the total (unbinned) cross-section. Sources of uncertainty that contribute approximately evenly to each bin of a given distribution will therefore be cancelled by the corresponding term in the denominator of this ratio. In particular the $t\bar{t}$ PS modelling which is large when considering the absolute cross-section has a relatively small effect on the normalised differential cross-sections.

6.7 Results and interpretation

Figure 6.11 shows the normalised differential cross-sections for each of the 6 observables, unfolded to particle-level. The shapes predicted by the DS interference scheme in POWHEG are compared to the nominal DR scheme. Alternative predictions obtained by using the MG5_aMC@NLO generator or by using HERWIG++ for parton showering are also compared to the observed data. Of the cross-sections presented, $E(b)$ most directly corresponds to the tW production kinematics since the single b-jet is, at first order, unambiguously associated with the b-quark in the decay of the top quark. Its cross-section in Fig. 6.11a shows slight evidence of a harder observed spectrum compared to the predictions. The quantities $m(\ell_1 b)$ and $m(\ell_2 b)$ are sensitive to the correlations in the angular distribution of the decay products arising from spin correlations in the tW production. There is very little difference between the DR and DS interference schemes, as expected as a result of the choice of an event selection that minimises the effects of interference between tW and $t\bar{t}$.

The data and bin-by-bin uncertainty are tabulated in Table 6.7. The total uncertainty in each bin is large with respect to the size of differences between the alternative predictions, and in most cases covers any disagreement between the data and the predictions. Statistical and systematic sources of uncertainty contribute similarly to the total uncertainty in many bins.

In order to quantify any preference for alternative predictions for the normalised differential cross-sections, a χ^2 test is performed based on the methodology of Ref. [51]. The χ^2 values are given by

$$\chi^2 = \mathbf{v}^\mathsf{T}\mathbf{C}^{-1}\mathbf{v} \tag{6.3}$$

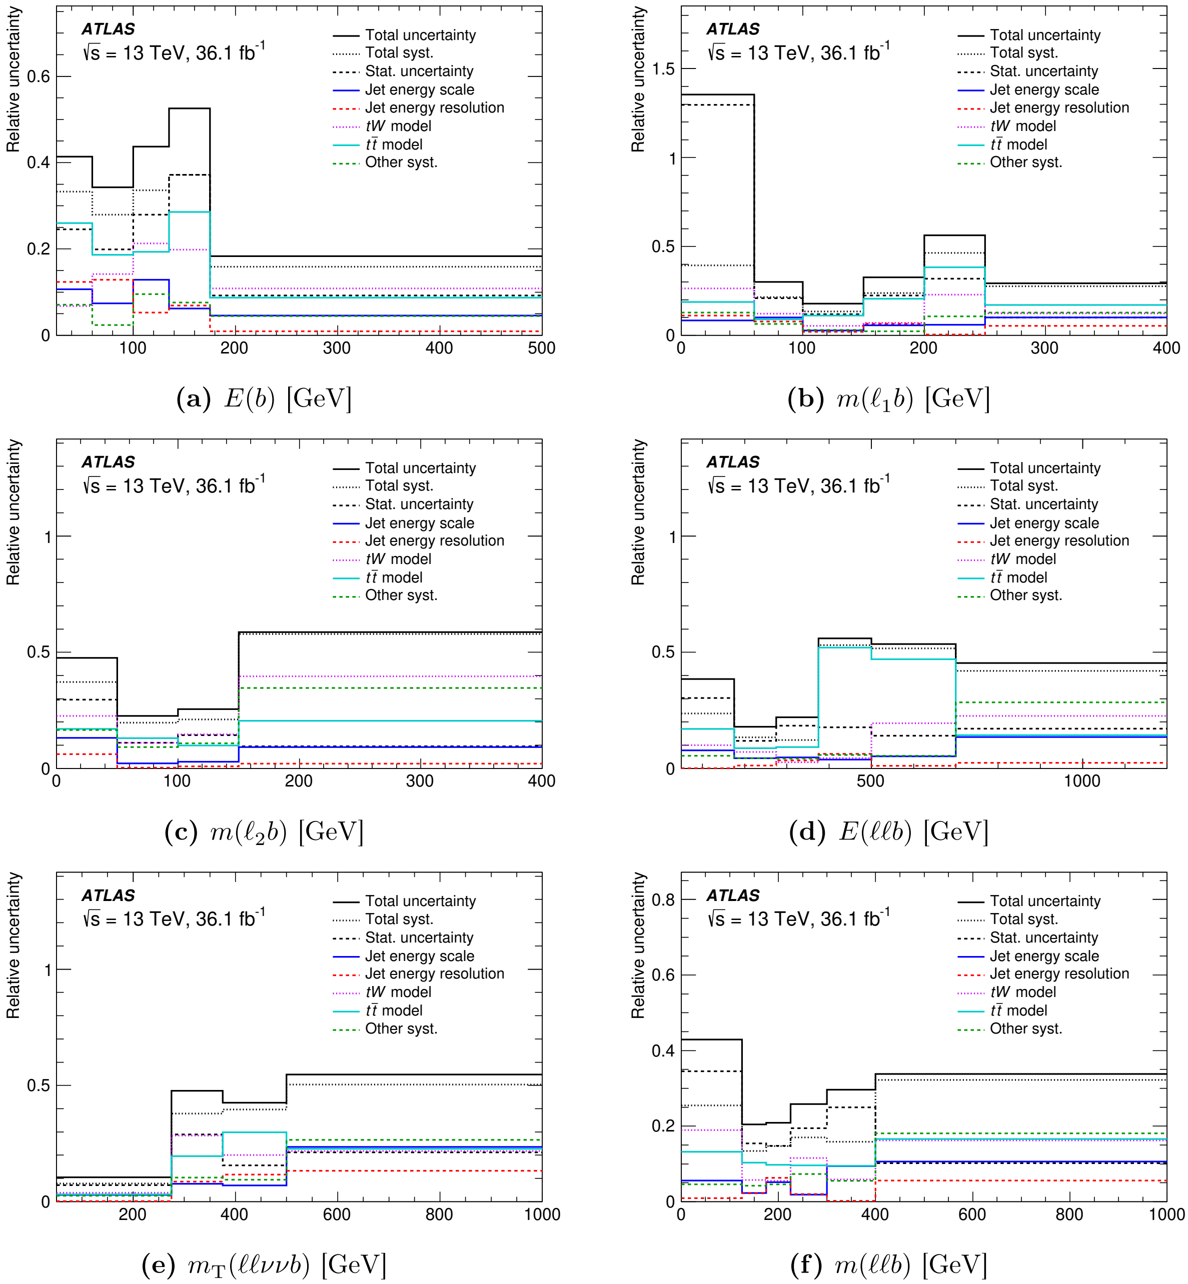

(a) $E(b)$ [GeV] **(b)** $m(\ell_1 b)$ [GeV]

(c) $m(\ell_2 b)$ [GeV] **(d)** $E(\ell\ell b)$ [GeV]

(e) $m_{\mathrm{T}}(\ell\ell\nu\nu b)$ [GeV] **(f)** $m(\ell\ell b)$ [GeV]

Figure 6.10: Summary of the relative uncertainty on the normalised differential cross-sections due to different categories of sources.

Table 6.7: Normalised differential cross-sections of each observable. Statistical, systematic, and total uncertainty are quoted as percentages for each bin.

$E(b)$ bin [GeV]	[25, 60]	[60, 100]	[100, 135]	[135, 175]	[175, 500]	
$(1/\sigma)\ \mathrm{d}\sigma/\mathrm{d}x$ [MeV^{-1}]	4.38	6.13	4.74	2.52	1.03	
Stat. uncertainty	25	20	28	37	9.3	
Total syst. uncertainty	33	28	34	37	16	
Total uncertainty	41	34	44	53	18	
$m(\ell_1 b)$ bin [GeV]	[0, 60]	[60, 100]	[100, 150]	[150, 200]	[200, 250]	[250, 400]
$(1/\sigma)\ \mathrm{d}\sigma/\mathrm{d}x$ [MeV^{-1}]	0.191	4.28	8.06	3.33	1.53	1.14
Stat. uncertainty	130	21	12	22	32	10
Total syst. uncertainty	39	22	13	24	46	28
Total uncertainty	140	30	18	33	56	29
$m(\ell_2 b)$ bin [GeV]	[0, 50]	[50, 100]	[100, 150]	[150, 400]		
$(1/\sigma)\ \mathrm{d}\sigma/\mathrm{d}x$ [MeV^{-1}]	1.84	8.45	5.31	0.879		
Stat. uncertainty	30	11	14	9.6		
Total syst. uncertainty	37	20	21	58		
Total uncertainty	48	23	25	59		
$E(\ell\ell b)$ bin [GeV]	[50, 175]	[175, 275]	[275, 375]	[375, 500]	[500, 700]	[700, 1200]
$(1/\sigma)\ \mathrm{d}\sigma/\mathrm{d}x$ [MeV^{-1}]	0.597	3.22	1.85	1.35	0.832	0.167
Stat. uncertainty	30	12	18	18	14	17
Total syst. uncertainty	24	13	12	53	52	42
Total uncertainty	38	18	22	56	53	45
$m_{\mathrm{T}}(\ell\ell\nu\nu b)$ bin [GeV]	[50, 275]	[275, 375]	[375, 500]	[500, 1000]		
$(1/\sigma)\ \mathrm{d}\sigma/\mathrm{d}x$ [MeV^{-1}]	3.3	1.23	0.856	0.0551		
Stat. uncertainty	7.1	29	16	21		
Total syst. uncertainty	7.8	38	40	50		
Total uncertainty	11	48	43	55		
$m(\ell\ell b)$ bin [GeV]	[0, 125]	[125, 175]	[175, 225]	[225, 300]	[300, 400]	[400, 1000]
$(1/\sigma)\ \mathrm{d}\sigma/\mathrm{d}x$ [MeV^{-1}]	0.51	5.33	5.38	2.42	0.949	0.208
Stat. uncertainty	35	15	15	19	25	10
Total syst. uncertainty	25	13	15	17	16	32
Total uncertainty	43	20	21	26	30	34

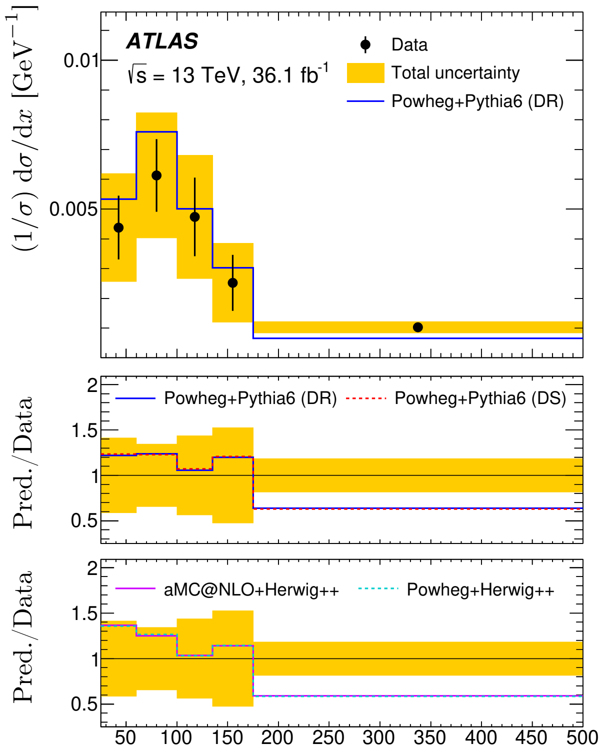

(a) $E(b)$ [GeV]

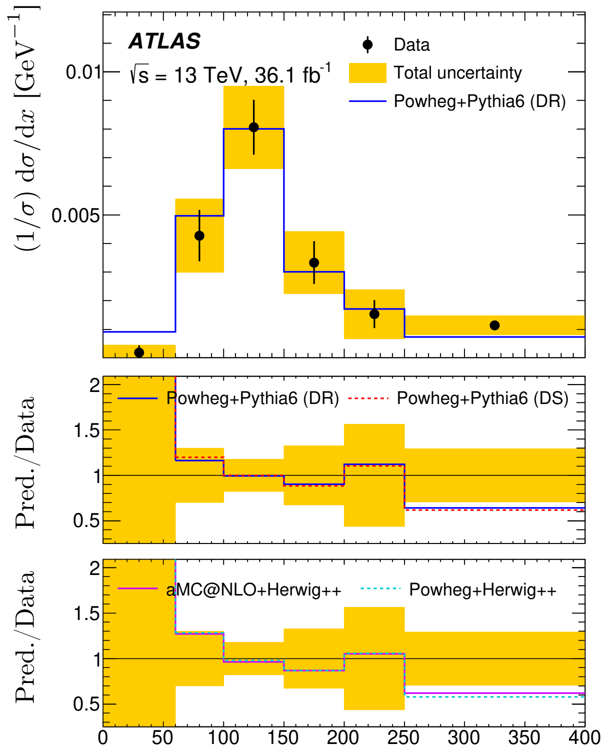

(b) $m(\ell_1 b)$ [GeV]

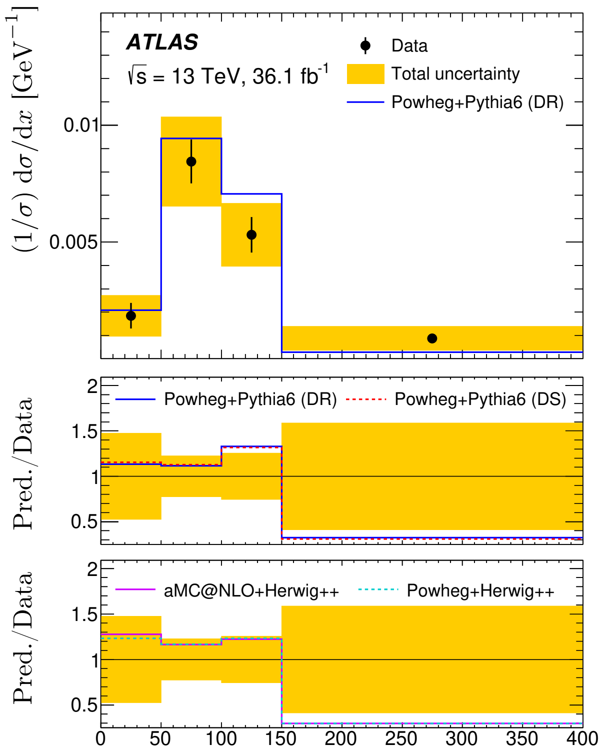

(c) $m(\ell_2 b)$ [GeV]

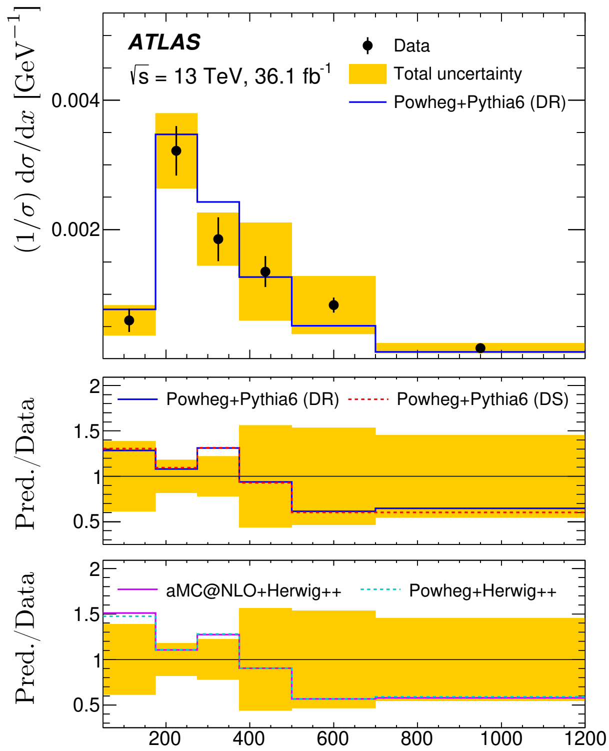

(d) $E(\ell\ell b)$ [GeV]

Figure 6.11: Normalised differential cross-sections of the 6 observables. The black dots show the observed cross-section with statistical error bars. The solid blue line shows the predicted cross-section using the nominal MC generators. The two lower panels compare predictions from alternative generators with the observed cross-section. The yellow shaded region shows the total uncertainty.

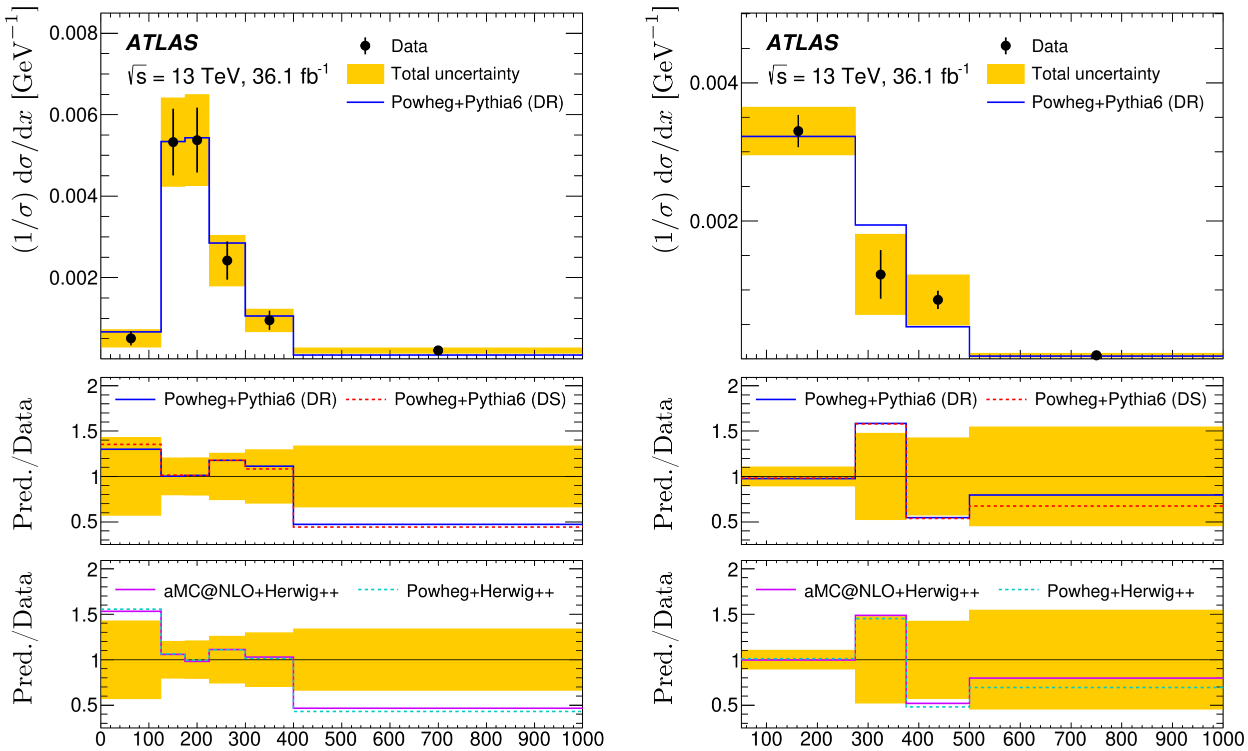

(e) $m_{\mathrm{T}}(\ell\ell\nu\nu b)$ [GeV] **(f)** $m(\ell\ell b)$ [GeV]

Figure 6.11: (Continued.)

where **v** is a vector containing the difference between the observed and predicted cross-sections in each bin of an observable and **C** is the covariance matrix. **C** is computed by combining a covariance matrix intended to describe the correlations in experimental uncertainties with a second matrix describing the other sources.

The first matrix is computed by creating an ensemble of pseudo-datasets by first varying the yield in each bin independently according to a Poisson distribution determined by the statistical uncertainty, followed by a fluctuation according to each experimental source of uncertainty. For this second step, a Gaussian distribution with mean zero and unit width is independently sampled for each experimental source of systematic uncertainty, and is multiplied by the difference between that source's systematic variation and the nominal yield to determine the contribution from each source. By independently fluctuating each experimental source and the statistical variation for each of 10 000 pseudo-datasets, information about the bin-to-bin correlations is preserved. Each pseudo-dataset is unfolded using the nominal unfolding matrix to produce 10 000 cross-section values in each bin. The covariance matrix is then determined using the average variations in these cross-section values[5].

The second covariance matrix is computed for the modelling and unfolding uncertainties. The correlation is set to zero for the MC statistics and non-closure sources since they were independently determined for each bin. All other sources are considered to be fully correlated between bins. These choices for bin-to-bin correlation are then used to transform the diagonal covariance matrix calculated during unfolding [2] into a full covariance matrix. The sum of the resulting matrix and the experimental covariance matrix is then used as the final covariance matrix **C**.

[5]The terms of the covariance matrix are given by $\mathrm{Cov}(i,j) = \left\langle (\mathrm{d}\sigma_i - \langle \mathrm{d}\sigma_i \rangle) \cdot (\mathrm{d}\sigma_j - \langle \mathrm{d}\sigma_j \rangle) \right\rangle_N$ where i and j index the bins in the distribution and the averaging is done over the ensemble of $N = 10000$ pseudo-datasets.

Table 6.8: χ^2 values and corresponding p-values characterising the agreement between the observed normalised differential cross-sections and the predictions using the nominal and alternative MC samples.

Observable	$E(b)$		$m(\ell_1 b)$		$m(\ell_2 b)$		$E(\ell\ell b)$		$m_{\mathrm{T}}(\ell\ell\nu\nu b)$		$m(\ell\ell b)$	
Degrees of freedom	4		5		3		5		3		5	
Prediction	χ^2	p	χ^2	p	χ^2	p	χ^2	p	χ^2	p	χ^2	p
POWHEG+PYTHIA6												
diagram removal	4.8	0.31	5.7	0.34	2.6	0.45	8.1	0.15	2.0	0.56	4.0	0.55
diagram subtraction	5.0	0.29	6.1	0.30	2.6	0.46	9.1	0.11	2.4	0.49	4.4	0.50
high radiation	4.8	0.30	5.3	0.38	2.5	0.48	7.9	0.16	1.9	0.60	3.7	0.60
low radiation	5.0	0.29	5.8	0.33	2.6	0.45	8.4	0.14	2.1	0.56	4.0	0.55
aMC@NLO+HERWIG++	5.6	0.23	5.4	0.37	2.4	0.49	8.7	0.12	1.8	0.61	3.6	0.61
POWHEG+HERWIG++	6.2	0.18	8.1	0.15	2.3	0.52	11.0	0.05	2.0	0.57	5.2	0.40

The χ^2 values obtained using Eq. (6.3) are shown in Table 6.8 along with the corresponding p-values. None of the p-values are low enough to suggest any evidence that a particular generator is disfavoured by the data. This is consistent with the fact that the total uncertainty is of comparable size to any disagreement with the data.

The normalised differential cross-sections of tW production presented here are the first such measurements that have been made. The cross-sections are found to be in good agreement with a range of predictions. The observed cross-sections are published and available for use in the preparation and refinement of MC generators. With the substantially larger dataset collected by the ATLAS experiment over the full Run 2 of the LHC a reduction in the statistical uncertainty, and an improved understanding of the detector in these conditions, there is good potential to achieve a reduction in the total uncertainty in a similar measurement.

Chapter 7

Understanding the tW process

The measurements presented in this thesis are the first measurements of tW production cross-sections with the ATLAS experiment at $\sqrt{s} = 13\,\text{TeV}$. They fit into a broader set of studies advancing our understanding of single top quark production at the LHC. This chapter outlines some recent results concerning tW production and discusses the prospects for updated measurements in the future.

7.1 The state of single top cross-section measurements

In Run 1 of the LHC, tW production was discovered by the ATLAS and CMS collaborations in the dilepton decay channel [52, 114]. Both collaborations also measured single top production in the t-channel using the single lepton channel [34, 110]. Due to the larger cross-section of t-channel production, and the ability to fully reconstruct the system due to the presence of only a single invisible neutrino at first order, the collaborations were able to measure differential cross-sections with respect to the kinematics of the top quark and with precision comparable to theoretical predictions. Measurements of polarisation observables in the t-channel placed constraints on broad classes of new physics affecting the structure of the t-W-b vertex [68, 15]. Both collaborations also measured s-channel production [31, 115], although due to large backgrounds and small production cross-section these measurements did not achieve the same level of precision. A recent cross-collaboration effort to combine these results [12] has yielded LHC combination measurements for single top cross-sections at $\sqrt{s} = 7\,\text{TeV}$ and $8\,\text{TeV}$, shown in Fig. 7.1. This result included a detailed study of the sources of uncertainty in the measurements so that the correlation of individual sources between the experiments was accounted for in the combination.

Since the cross-section is sensitive to the t-W-b coupling, the measurements can be interpreted as a measurement of this coupling. The presence of anomalous left-handed vector couplings would modify the coupling independently of the value of the CKM matrix element V_{tb}, so conventionally a form factor f_{LV} is introduced with its SM value set to unity. The expression $|f_{\text{LV}}V_{tb}|^2$ is then equal to the ratio of the measured single-top cross-section to the theoretical prediction[1] for a given production process. By evaluating this expression using the combined cross-section measurements for each single top process, as shown in Fig. 7.2, a value of $|f_{\text{LV}}V_{tb}| = 1.02 \pm 0.04$ is measured. The largest components of the uncertainty are due to the theoretical cross-section and modelling uncertainties in the measurements. This result is independent of any assumptions about the number of quark

[1] V_{tb} is taken to be unity when calculating the theoretical cross-section prediction.

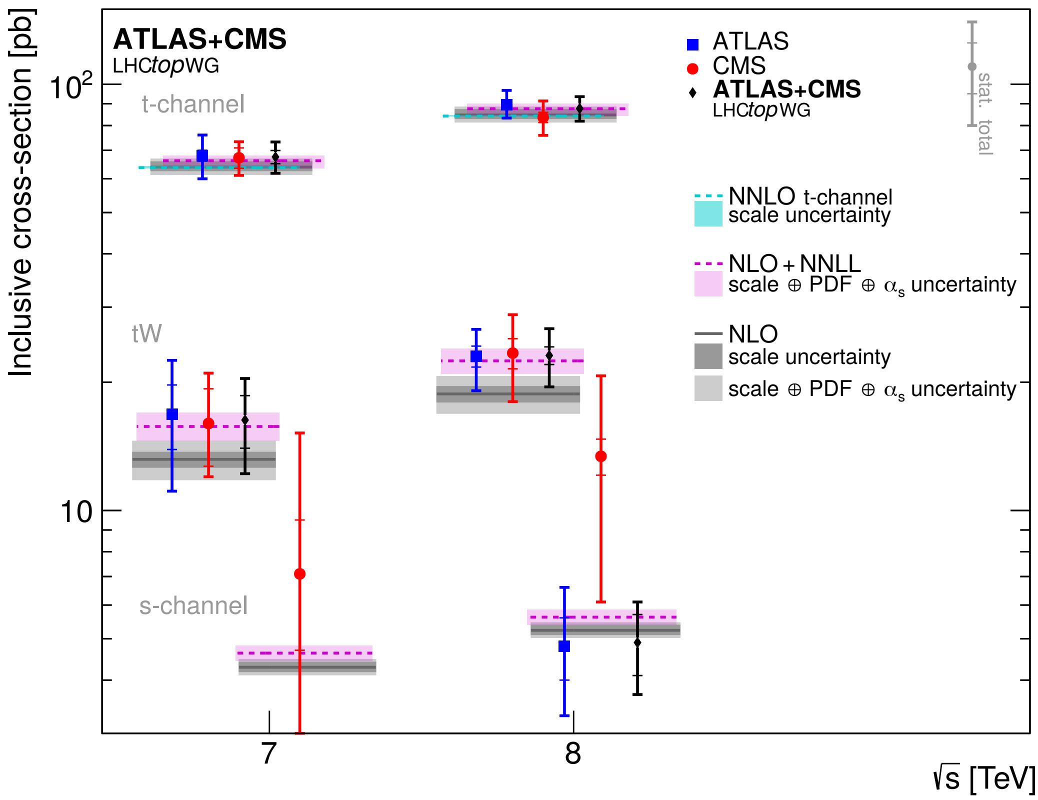

Figure 7.1: Combination of single top cross-section measurements by the Atlas and Cms experiments at Run 1 of the Lhc. The theoretical predictions are shown along with their associated uncertainties. Taken from Ref. [12].

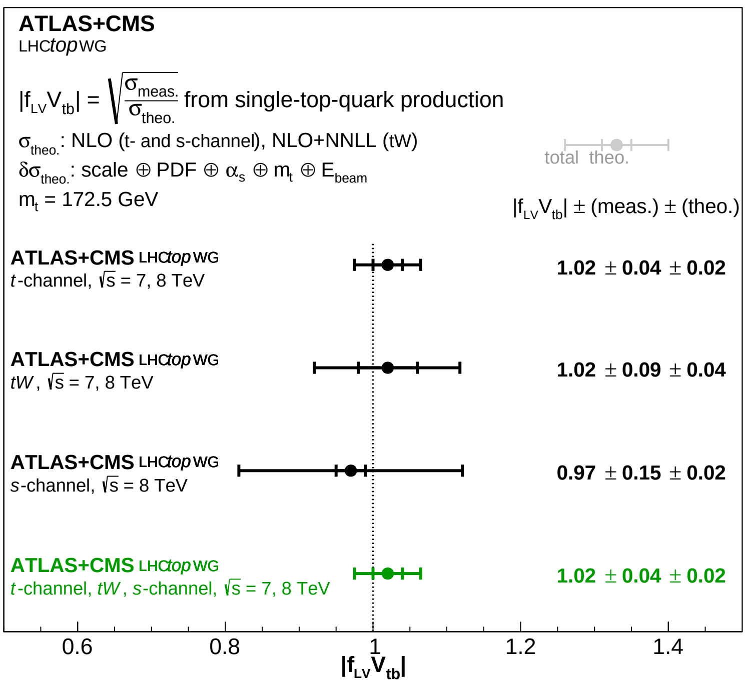

Figure 7.2: Combination of single top cross-section measurements by the Atlas and Cms experiments at Run 1 of the Lhc. Taken from Ref. [12].

generations or the unitarity of the CKM matrix, and so provides a powerful test for new physics that affects the t-W-b coupling.

The early Run 2 measurement of the tW production cross-section with ATLAS was presented in Chapter 5. In §5.5 the result was shown alongside the first Run 2 measurement of tW production by the CMS collaboration using 35.9 fb^{-1} of pp collisions [111]. The CMS measurement also considered the decay of the tW system to two oppositely-charged leptons, but with the additional restriction that the two leptons must have different flavour. The event selection and object definitions were chosen similarly to those in this thesis, although of course the details of event reconstruction with the CMS detector are different. The analysis is broadly similar to the one presented in this thesis, with a BDT constructed from kinematic observables and its response used in a binned maximum likelihood to extract the tW cross-section. One important difference is that jets with $20\,\text{GeV} < p_{\text{T}} < 30\,\text{GeV}$ are treated as separate objects to jets with higher p_{T} when preparing input features for the BDT, and the p_{T} and multiplicity of these softer jets both contribute to the performance of the BDT in the `1j1b` region. The composition of the total uncertainty on the CMS cross-section measurement differs from that of the ATLAS measurement due to differences in the performance of the detectors and differences in the methodology of the uncertainty analysis.

Both the ATLAS and CMS experiments will prepare updated measurements of the tW production cross-section using the full Run 2 dataset which will then be used to produce an updated combined LHC measurement. In addition to measuring tW, t-channel, and s-channel production, the collaborations are pursuing ever-rarer processes including single top quark production in association with a Z boson [53, 113] or Higgs boson [116].

7.2 Prospects for future measurements of tW production

Using the full ATLAS Run 2 dataset, larger than the dataset used in Chapter 6 by a factor of 4, the uncertainty on the cross-section measurements due to limited data statistics is naively expected to be reduced by a factor of $\sqrt{4} = 2$. For the differential cross-section measurements where the data statistics currently dominate the uncertainty in some bins, this improvement would already make an updated measurement worthwhile. For the total cross-section measurement of Chapter 5, however, the total systematic uncertainty is larger than the statistical uncertainty by a factor of 3, limiting the potential of a large dataset to improve the measurement. It is therefore essential to reduce the systematic uncertainty in order to achieve greater precision in the future. The largest contributions to this uncertainty are related to jet reconstruction and signal and background modelling, so reductions in these components will have the largest impact on the total uncertainty.

7.2.1 Calibration improvements

As discussed in §4.6.2, the jet energy calibration used in this thesis was based on an extrapolation of the Run 1 calibration with a conservatively large uncertainty term added to cover uncertainty in the extrapolation procedure. Updated calibrations using the 2015 to 2017 dataset [41, 42] have been prepared without this extrapolation and the associated additional uncertainty, demonstrating a reduced uncertainty particularly in the low-p_{T} range where the extrapolation uncertainty was most pronounced.

For full Run 2 measurements, ATLAS is moving towards a new jet reconstruction technique called particle flow [44], based on similar techniques to those used by CMS and other experiments. Whereas in this thesis jets are reconstructed from topo-clusters built from the calorimeter response, particle flow instead extrapolates ID tracks into the calorimeter and uses the response of both the tracking system and the calorimeters to measure the jet energy. This offers a number of advantages, particularly concerning low-p_{T} particles. The momentum resolution of the tracking system at low p_{T} is significantly better than the energy resolution of the calorimeter for low-energy deposits. The tracking system also offers much better angular resolution, and provides information about whether tracks are likely to have originated at the same vertex as each other. This additional information allows for soft particles which have been swept out of the jet while traversing the magnetic field of the ID to be correctly attributed to the jet, whereas this is not possible using only information from the calorimeters. Particle flow jets are also less affected by pileup jets due to their incorporation of tracking information.

In addition to the improvements from the particle flow jet reconstruction technique, the uncertainty associated with the jet energy can be reduced by exploiting information about the specific samples being analysed. The jet energy calibration is sensitive to whether jets originate from gluons or quarks, and to account for this the uncertainty includes terms related to how precisely the quark-to-gluon ratio is known, as well as terms related to uncertainty in the response of the calorimeter to gluon jets. The default calibration assumes that 50 % of the jets are initiated by gluons, with a 100 % uncertainty on that value. By observing the true fraction in the simulated MC samples used in the tW cross-section measurements, we find that for both tW and $t\bar{t}$ events, the majority of jets are initiated by gluons in the parts of the phase space probed by the analysis. Quark-initiated jets on the other hand account for a larger fraction of jets in the forward direction and at higher p_{T}. By using this information, the associated component of the jet energy uncertainty can be reduced, but at the expense of increasing the importance of the uncertainty related to the parton-dependent response to the jet in the calorimeters.

Improvements in b-tagging algorithms [58] are also expected to assist with updated tW measurements since the presence of an additional b-jet can distinguish $t\bar{t}$ from tW events. It is possible that incorporating a looser b-tagging algorithm, or even directly using a b-tagging likelihood as a discriminating variable could yield improvements in background separation. The calibration of such an approach and its implications on the overall uncertainty of the measurement are not obvious *a priori*.

7.2.2 Modelling the interference between tWb and $t\bar{t}$

A challenge presented by the `2j2b` region is that it captures the final state in which the $t\bar{t}$ and tW processes interfere (see §1.5.4), making the analysis more sensitive to the modelling of this interference. This must be confronted not only by studies of tW production, but also by any searches for new physics for which tW production is considered a background. The ATLAS collaboration studied a sample of events selected to emphasise the interference effects—in contrast to the event selections in this thesis which sought to minimise them—to evaluate the alternative models for resolving the ambiguity in the processes [69]. The study selected dilepton events with exactly 2 b-tagged jets, with vetoes applied to events where the invariant mass of the leptons was near the Z mass or where any

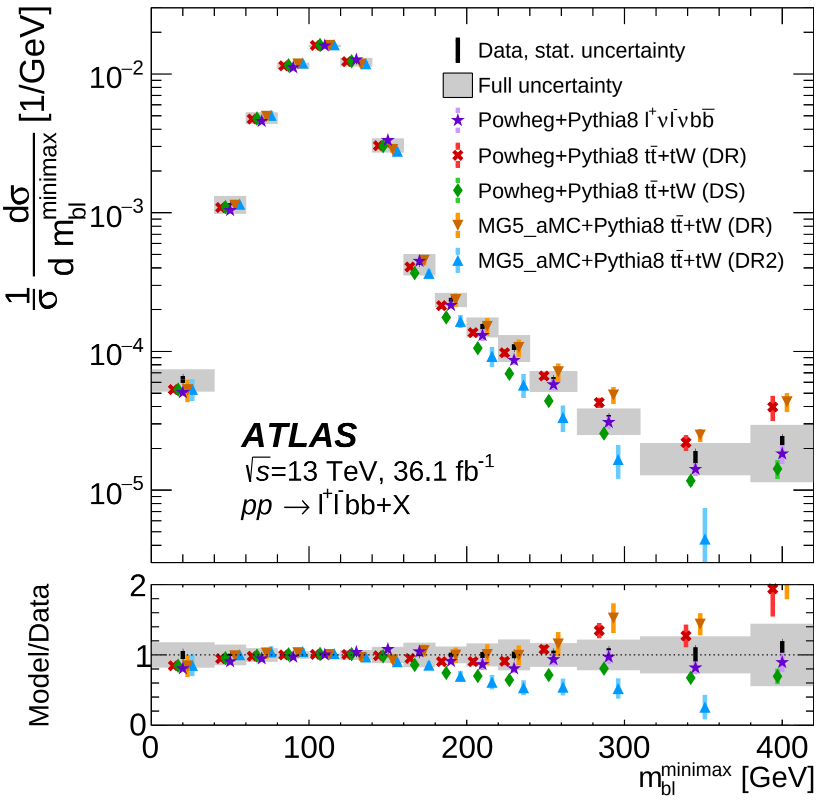

Figure 7.3: The differential cross-section of $tW + t\bar{t}$ with respect to $m_{b\ell}^{\mathrm{minimax}}$. The cross-section is compared to several predictions, with the lower panel displaying the ratio of these predictions to the observed cross-section. The `bb4l` prediction is represented by the purple points, and the DR scheme used as the nominal sample in this thesis is shown by the red points. Taken from Ref. [69].

additional jet passed a looser b-tagging algorithm. The invariant mass variable

$$m_{b\ell}^{\mathrm{minimax}} \equiv \min\{\ \max[m(b_1\ell_1), m(b_2\ell_2)],\ \ \max[m(b_1\ell_2), m(b_2\ell_1)]\ \}$$

was defined to be sensitive to the different kinematics of the doubly- and singly-resonant processes. The doubly-resonant process is suppressed at leading order above the kinematic endpoint $\sqrt{m_t^2 - m_W^2}$, whereas the singly-resonant process is not. The differential cross-section of $tW + t\bar{t}$ with respect to $m_{b\ell}^{\mathrm{minimax}}$ was measured in 36.1 fb^{-1} of pp data and compared to several MC predictions, as shown in Fig. 7.3. The prediction by the resonance-aware `bb4l` generator was found to give the best description of the observed cross-section in the tail of the distribution. The DR and DS predictions respectively overestimate and underestimate the tail, and the newer DR2 prediction significantly underestimates the tail.

This first test of predictions in the interference-dominated regime of $W^+bW^-\bar{b}$ production demonstrates that there is a need for further theoretical work to improve the modelling situation. Experimentally, there is more work to be done in comparing the full range of available predictions at greater precision, and to apply the results of this investigation to improve analysis techniques and uncertainty modelling across measurements and searches sensitive to this regime.

7.2.3 Improvements to the analysis technique

An improvement in the rejection of $t\bar{t}$ and other backgrounds from the sample used to extract tW cross-sections will result in a corresponding reduction in uncertainties related to the modelling of those background processes. The $Z + j$ and diboson background processes contribute a substantially smaller number of events to the final sample, but any of these events that are classified by the BDT as

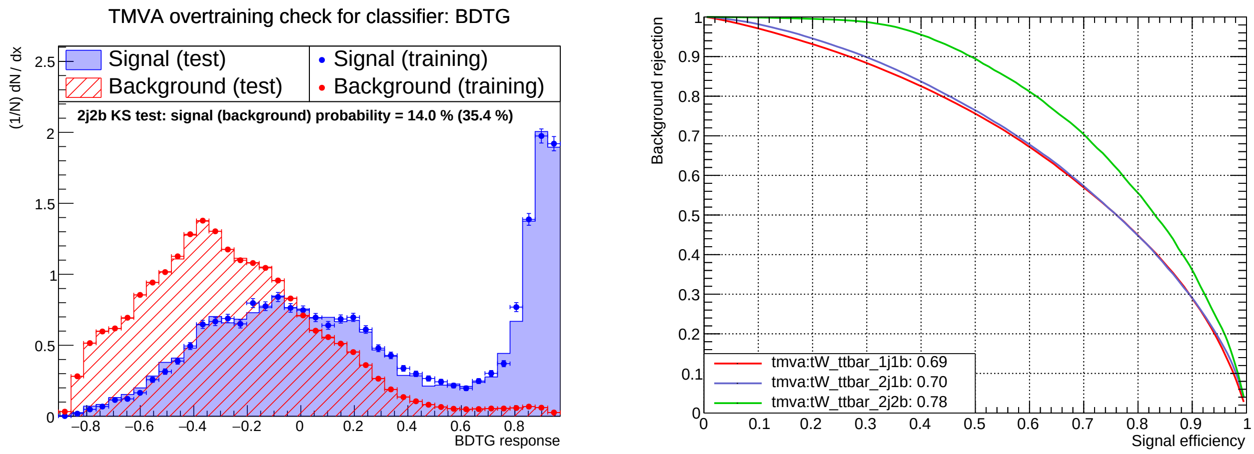

(a) BDT response **(b)** ROC curve

Figure 7.4: Performance of a BDT trained in the 2j2b region to discriminate between tW and $t\bar{t}$ events. The response of the BDT is shown in (a) evaluated using the nominal tW sample (blue) and the nominal $t\bar{t}$ sample (red). The points show the response in the training sample used to prepare the BDT, and the hashed/filled regions show the response in events that were not used in the training to check for overtraining. The curves are normalised such that they have equal area. The ROC curve for the BDTs is shown in (b) in green, compared to those of the 1j1b and 2j1b regions.

signal-like will accumulate in the bins of the BDT response distribution which have the highest ratio of tW to $t\bar{t}$ events. This situation can cause modelling uncertainty on these small backgrounds to become large as they affect the most sensitive parts of the fit. The measurements in this thesis avoided this issue by optimising background rejection cuts, however in a larger sample this approach will be less effective. Removing or relaxing the flavour-dependent requirements placed on $M_{\ell\ell}$ and $E_{\mathrm{T}}^{\mathrm{miss}}$ would further align the reconstruction- and particle-level selections for a differential cross-section measurement, reducing some of the acceptance and efficiency corrections. If the contamination of smaller backgrounds becomes larger as a result of a loosening in the background requirements, it may become necessary to develop more sophisticated approaches to estimate their contributions, for example data-driven methods that extrapolate measurements from background-rich control regions.

The measurement in Chapter 5 did not use a BDT to discriminate between tW and $t\bar{t}$ in the 2j2b region, but using MC samples prepared since then it has become possible to successfully train a BDT in this region. The response and performance of one such classifier is shown in Fig. 7.4. Unlike the other regions for which classifiers were prepared, there is a clear peak of signal events at the high range of the BDT response with almost no background contamination. This allows a very pure sample of tW events to be isolated, albeit in a region which has very few tW events since it requires the radiation of an addition b-quark with respect to the leading-order process. The presence of this signal peak is reflected in the ROC curve, which encloses a substantially larger area of 0.78 compared to the area of 0.70 underneath the ROC curves for the single b-tag regions. This good performance of the classifier suggests an improvement that will be reflected in the full likelihood fit, but it also suggests that there is a distinct population of tW events that are qualitatively different to others in the 2j2b region. Further work to understand the nature of this difference would be enlightening.

In parallel with incorporating more detailed b-tagging information, there are potential gains in using softer jets in the analysis as demonstrated by CMS. Low-p_{T} jets are more challenging to

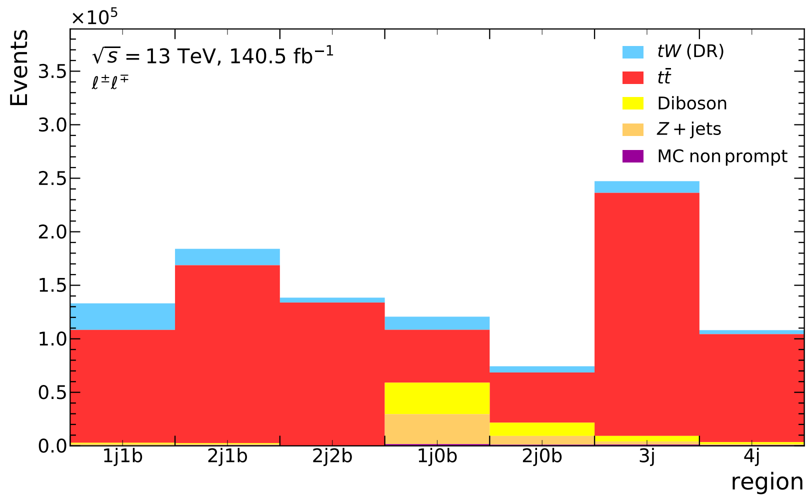

Figure 7.5: Predictions for the yields of processes contributing to the dilepton selection, normalised using their predicted cross-sections and the luminosity of the full Run 2 dataset. In addition to the regions previously introduced, two new regions, `3j` and `4j` are included. These are defined by the presence of 3 or 4 selected jets, with no b-tagging requirement.

reconstruct, but the particle flow technique assists with reducing the uncertainty associated with the energy reconstruction in exactly this regime. The multiplicity or detailed kinematics of any soft jets in the event could then be used to improve the discrimination between tW and $t\bar{t}$ events. This could be done without modifying the definition of jets as objects with $p_{\mathrm{T}} > 25\,\mathrm{GeV}$ used to define the regions of the analysis. It could also be useful to consider events with 3 such jets present, without any restriction on the number of b-tagged jets. This `3j` region has a large contribution of $t\bar{t}$ events, but also includes tW events with additional radiation as shown in Fig. 7.5. Due to the lack of a b-tagging requirement, there are in fact more tW events in the `3j` region than in the `2j2b` region, although there is also a higher contamination from backgrounds other than $t\bar{t}$.

7.2.4 Simultaneous measurements of dilepton processes

In Run 1, the ATLAS collaboration measured the cross-sections of the $t\bar{t}$, W^+W^-, and $Z/\gamma^* \to \tau\tau$ processes simultaneously in a sample of dilepton events [72]. This measurement was based on a global fit technique developed at the CDF experiment at the TEVATRON [106]. By treating all three of these processes—which are the dominant contributors to final states with two oppositely-charged high-p_{T} leptons—as signal processes, the background to the measurement consisted of SM processes with smaller cross-sections such as tW production, alongside events with fake and non-prompt leptons. The signal processes were separated by considering the jet multiplicity and $E_{\mathrm{T}}^{\mathrm{miss}}$, since they typically occupy separate regimes of this two-dimensional distribution. Their cross-sections were then extracted simultaneously with a binned likelihood fit to observed data using this distribution.

Because the cross-sections were measured simultaneously, the correlations between them could be explored in the fit, allowing for the contours of the likelihood to be extracted. The slice of these contours for the $Z/\gamma^* \to \tau\tau$ and $t\bar{t}$ cross-sections are shown in Fig. 7.6, where the comparisons to predictions using PDF sets at fixed order in the QCD expansion show that all NLO predictions and some NNLO predictions fail to describe the data.

With the substantially larger datasets collected in Run 2, an update to this study could also treat tW as a signal process, and incorporate improved b-tagging information to further separate the signal processes. The global fit technique allows for a consistent treatment of the signal processes in the same phase space. For this reason it is sensitive to new physics that might not be evident using cross-section measurements obtained individually.

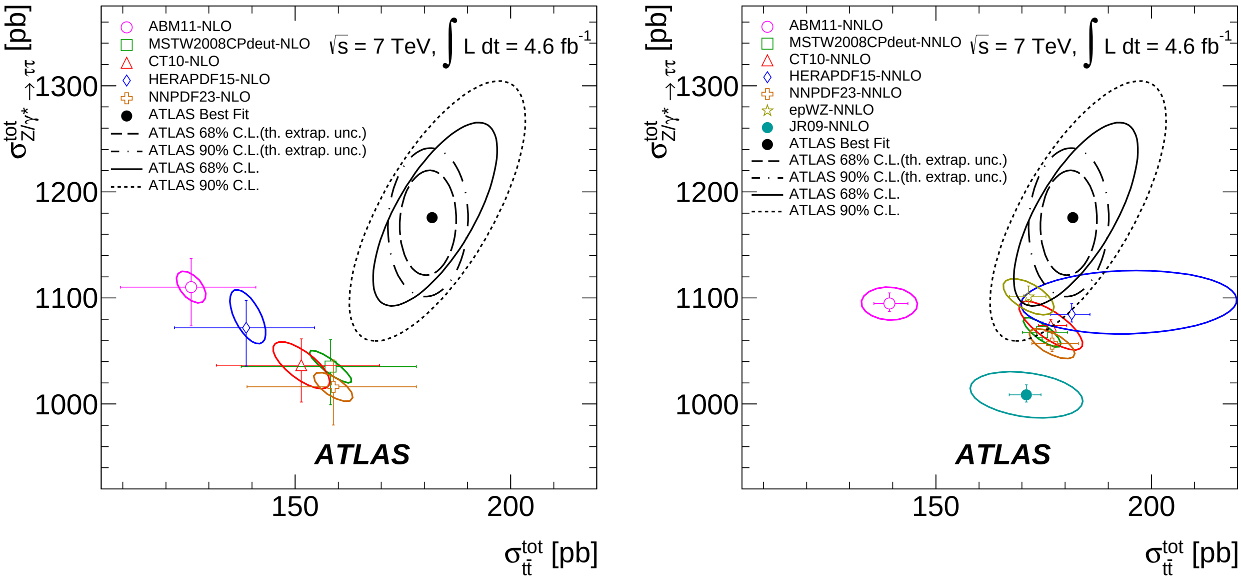

(a) NLO **(b)** NNLO

Figure 7.6: Contours of the likelihood function as a function of $Z/\gamma^* \to \tau\tau$ and $t\bar{t}$ cross-sections. The contours in black were obtained from the observed data, while the coloured points and 68 % confidence level uncertainty contours correspond to predictions made with different PDF sets. Taken from Ref. [72].

7.3 Conclusions

This thesis presented the first cross-section measurements of tW production at $\sqrt{s} = 13\,\text{TeV}$ with the ATLAS experiment. With dilepton events from $3.2\,\text{fb}^{-1}$ of pp collision data collected in 2015, the tW cross-section was measured to be

$$\sigma_{tW} = 94\ \pm\ 10\,(\text{stat.})\ {}^{+28}_{-22}\,(\text{syst.})\ \pm 2\,(\text{lumi.})\ \text{pb},$$

consistent with the theoretical prediction of

$$\sigma_{\text{theory}}(tW) = 71.7\ \pm\ 1.8\,(\text{scale})\ \ \pm 3.4\,(\text{PDF})\ \text{pb}.$$

With dilepton events from $36.1\,\text{fb}^{-1}$ of pp collision data collected in 2015 and 2016, the differential tW cross-section was measured as a function of 6 observables related to the kinematics of the decay products of the tW system. Predictions from several MC generators were found to describe the observables well. These measurements taken together with the broader context of single top studies present a picture of tW and single top quark production that is consistent with the SM. The improvements to modelling of SM processes will assist with searches for rarer SM processes and potential new physics for which single top production constitutes a background.

There are good prospects for improving the precision of tW measurements using the full ATLAS Run 2 dataset of $140\,\text{fb}^{-1}$ by incorporating refined calibrations of the detector's response, improving the analysis techniques using experience gained from these early studies, and benefiting from improvements in the theoretical modelling of the relevant processes. As single top cross-section measurements move into the precision regime, many more specialised measurements are becoming feasible, including differential cross-sections as functions of the kinematics of the tW system itself, exploration of poorly-understood regions of the phase space, and global fits to SM processes.

Acronyms

4FS	four-flavour scheme, p. 12
5FS	five-flavour scheme, p. 12
ATLAS	A Toroidal LHC ApparatuS, p. 25
ATLFASTII	ATLAS fast calorimeter simulation, p. 50
BDT	boosted decision tree, p. 67
BSM	beyond the Standard Model, p. 3
CDF	Collider Detector at Fermilab, p. 17
CERN	European Organisation for Nuclear Research (Conseil européen pour la recherche nucléaire), p. 1
CKM	Cabibbo-Kobayashi-Maskawa, p. 6
CMS	Compact Muon Solenoid, p. 25
CSC	Cathode Strip Chamber, p. 42
DESY	German Electron Synchrotron (Deutsches Elektronen-Synchrotron), p. 17
DR	diagram removal, p. 23
DS	diagram subtraction, p. 23
ECAL	electromagnetic calorimeter, p. 36
EM	electromagnetic, p. 5
EW	electroweak, p. 5
HCAL	hadronic calorimeter, p. 41
HL-LHC	High Luminosity LHC, p. 30
HLT	High Level Trigger, p. 44
IBL	Insertable B-Layer, p. 34
ID	Inner Detector, p. 32
IP	interaction point, p. 25
JER	jet energy resolution, p. 64
JES	jet energy scale, p. 64
JVF	jet vertex fraction, p. 66
JVT	jet-vertex-tagger, p. 66
KS	Kolmogorov-Smirnov, p. 81
L1	Level 1 trigger, p. 44
LAr	liquid argon, p. 36
LEP	Large Electron-Positron Collider, p. 17
LHC	Large Hadron Collider, p. 3
LO	leading order, p. 75
LUCID	Luminosity Cherenkov Integrating Detector, p. 46
MC	Monte Carlo, p. 13
MDT	Monitored Drift-Tube, p. 42
ME	matrix element, p. 13
MS	Muon Spectrometer, p. 41
NLO	next-to-leading order, p. 14
NNLL	next-to-next-to-leading logarithm, p. 20
P2012	Perugia 2012, p. 74

PDF	parton distribution function, p. 10
PMT	photomultiplier tube, p. 41
PS	parton shower, p. 14
QCD	quantum chromodynamics, p. 5
QFT	quantum field theory, p. 3
ROC	receiver operating characteristic, p. 83
RPC	Resistive Plate Chamber, p. 42
SCT	Semiconductor Tracker, p. 34
SM	Standard Model (of particle physics), p. 3
S$p\bar{p}$S	Proton-Antiproton Collider, p. 17
SPS	Super Proton Synchrotron, p. 17
TGC	Thin Gap Chamber, p. 42
TileCal	Tile Calorimeter, p. 32
topo-cluster	topological cluster, p. 54
TRT	Transition Radiation Tracker, p. 35
UA2	Underground Area 2, p. 17
VdM	van der Meer, p. 47

Bibliography

[1] H. Abdolmaleki, A. Khorramian and A. Aleedaneshvar. 'Impact of intrinsic charm on PDFs with EMC and LHC data'. In: *Nuclear and Particle Physics Proceedings* 282-284 (2017). 19th International Conference in Quantum Chromodynamics, pp. 27–31. ISSN: 2405-6014. DOI: 10.1016/j.nuclphysbps.2016.12.006 (cit. on p. 10).

[2] T. Adye. *Unfolding algorithms and tests using RooUnfold*. 2011. arXiv: 1105.1160 (cit. on pp. 122, 134).

[3] S. Agostinelli et al. 'GEANT4: A Simulation toolkit'. In: *Nucl. Instrum. Meth. A* 506 (2003), p. 250. DOI: 10.1016/S0168-9002(03)01368-8 (cit. on p. 50).

[4] K. Akai, K. Furukawa and H. Koiso. 'SuperKEKB collider'. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 907 (2018), pp. 188–199. ISSN: 0168-9002. DOI: 10.1016/j.nima.2018.08.017 (cit. on p. 26).

[5] R. Alemany-Fernandez et al. *Operation and Configuration of the LHC in Run 1*. CERN-ACC-NOTE-2013-0041. 2013. URL: https://cds.cern.ch/record/1631030 (cit. on p. 28).

[6] S. Alioli et al. 'A general framework for implementing NLO calculations in shower Monte Carlo programs: the POWHEG BOX'. In: *JHEP* 06 (2010), p. 043. DOI: 10.1007/JHEP06(2010)043. arXiv: 1002.2581 [hep-ph] (cit. on p. 74).

[7] J. Alwall et al. 'The automated computation of tree-level and next-to-leading order differential cross sections, and their matching to parton shower simulations'. In: *JHEP* 07 (2014), p. 079. DOI: 10.1007/JHEP07(2014)079. arXiv: 1405.0301 [hep-ph] (cit. on p. 89).

[8] U. Amaldi. 'How protons are seen by high energy hadrons'. In: *Soviet Physics Uspekhi* 21.4 (1978), pp. 328–344. DOI: 10.1070/pu1978v021n04abeh005540 (cit. on p. 8).

[9] S. Amendolia et al. 'Measurement of inclusive two-particle rapidity correlations at the ISR'. In: *Physics Letters B* 48.4 (1974), pp. 359–366. ISSN: 0370-2693. DOI: 10.1016/0370-2693(74)90609-1 (cit. on p. 8).

[10] B. Angerth et al. 'The LHC beam screen: specification and design'. In: *4th European Particle Accelerator Conference, London, UK, 27 Jun - 1 Jul 1994*. 1994, pp. 208–212. URL: http://accelconf.web.cern.ch/AccelConf/e94/PDF/EPAC1994_0208.PDF (cit. on p. 30).

[11] ARGUS Collaboration. 'Observation of B^0-$\bar{B}^0$ mixing'. In: *Physics Letters B* 192.1 (1987), pp. 245–252. ISSN: 0370-2693. DOI: 10.1016/0370-2693(87)91177-4 (cit. on p. 17).

[12] ATLAS and CMS Collaborations. 'Combinations of single-top-quark production cross-section measurements and $|f_{\mathrm{LV}} V_{tb}|$ determinations at $\sqrt{s} = 7$ and 8 TeV with the ATLAS and CMS experiments'. In: *JHEP* 05 (2019), p. 088. DOI: 10.1007/JHEP05(2019)088. arXiv: 1902.07158 [hep-ex] (cit. on pp. 136, 137).

[13] ATLAS Collaboration. *2015 start-up trigger menu and initial performance assessment of the ATLAS trigger using Run-2 data*. ATL-DAQ-PUB-2016-001. 2016. URL: https://cds.cern.ch/record/2136007 (cit. on p. 59).

[14] ATLAS Collaboration. 'A neural network clustering algorithm for the ATLAS silicon pixel detector'. In: *JINST* 9 (2014), P09009. DOI: 10.1088/1748-0221/9/09/P09009. arXiv: 1406.7690 [hep-ex] (cit. on p. 51).

[15] ATLAS Collaboration. 'Analysis of the Wtb vertex from the measurement of triple-differential angular decay rates of single top quarks produced in the t-channel at $\sqrt{s} = 8$ TeV with the ATLAS detector'. In: *JHEP* 12 (2017), p. 017. DOI: 10.1007/JHEP12(2017)017. arXiv: 1707.05393 [hep-ex] (cit. on p. 136).

[16] ATLAS Collaboration. *ATLAS calorimeter performance: Technical Design Report*. ATLAS-TDR-1. 1996. URL: https://cds.cern.ch/record/331059 (cit. on pp. 36, 39, 40).

[17] ATLAS Collaboration. *ATLAS Experiment Luminosity Public Results, Run 2*. URL: https://twiki.cern.ch/twiki/bin/view/AtlasPublic/LuminosityPublicResultsRun2 (visited on 07/08/2018) (cit. on pp. 29, 48).

[18] ATLAS Collaboration. *ATLAS Experiment Trigger Operations Public Results*. URL: https://twiki.cern.ch/twiki/bin/view/AtlasPublic/TriggerOperationPublicResults (visited on 03/06/2019) (cit. on p. 46).

[19] ATLAS Collaboration. *ATLAS inner detector: Technical Design Report, 1*. ATLAS-TDR-4. 1997. URL: https://cds.cern.ch/record/331063 (cit. on pp. 32, 34, 37).

[20] ATLAS Collaboration. *ATLAS inner detector: Technical Design Report, 2*. ATLAS-TDR-5. 1997. URL: https://cds.cern.ch/record/331064 (cit. on p. 32).

[21] ATLAS Collaboration. *ATLAS Insertable B-Layer Technical Design Report*. ATLAS-TDR-19. 2010. URL: https://cds.cern.ch/record/1291633 (cit. on p. 34). Addendum: ATLAS-TDR-19-ADD-1. 2012. URL: https://cds.cern.ch/record/1451888.

[22] ATLAS Collaboration. *ATLAS magnet system: Technical Design Report, 1*. ATLAS-TDR-6. 1997. URL: https://cds.cern.ch/record/338080 (cit. on p. 32).

[23] ATLAS Collaboration. *ATLAS muon spectrometer: Technical Design Report*. ATLAS-TDR-10. 1997. URL: https://cds.cern.ch/record/331068 (cit. on pp. 41, 42).

[24] ATLAS Collaboration. *ATLAS Pythia 8 tunes to 7 TeV data*. ATL-PHYS-PUB-2014-021. 2014. URL: https://cds.cern.ch/record/1966419 (cit. on p. 89).

[25] ATLAS Collaboration. *b-tagging efficiency calibration from 2015+2016 data on lepton+jets $t\bar{t}$ events*. 29 June 2017. URL: http://atlas.web.cern.ch/Atlas/GROUPS/PHYSICS/PLOTS/FTAG-2017-003/ (cit. on pp. 68, 69).

[26] ATLAS Collaboration. 'Characterisation and mitigation of beam-induced backgrounds observed in the ATLAS detector during the 2011 proton–proton run'. In: *JINST* 8 (2013), P07004. DOI: 10.1088/1748-0221/8/07/P07004. arXiv: 1303.0223 [hep-ex] (cit. on pp. 51, 66).

[27] ATLAS Collaboration. *Commissioning of the ATLAS b-tagging algorithms using $t\bar{t}$ events in early Run 2 data*. ATL-PHYS-PUB-2015-039. 2015. URL: https://cds.cern.ch/record/2047871 (cit. on p. 67).

[28] ATLAS Collaboration. *Early Inner Detector Tracking Performance in the 2015 Data at $\sqrt{s} = 13$ TeV*. ATL-PHYS-PUB-2015-051. 2015. URL: https://cds.cern.ch/record/2110140 (cit. on pp. 52, 53).

[29] ATLAS Collaboration. 'Electron and photon energy calibration with the ATLAS detector using 2015–2016 LHC proton–proton collision data'. In: *JINST* 14 (2019), P03017. DOI: 10.1088/1748-0221/14/03/P03017. arXiv: 1812.03848 [hep-ex] (cit. on pp. 57, 58).

[30] ATLAS Collaboration. 'Electron reconstruction and identification in the ATLAS experiment using the 2015 and 2016 LHC proton–proton collision data at $\sqrt{s} = 13$ TeV'. In: *Eur. Phys. J.* (2019). DOI: 10.1140/epjc/s10052-019-7140-6. arXiv: 1902.04655 [hep-ex] (cit. on pp. 55–58).

[31] ATLAS Collaboration. 'Evidence for single top-quark production in the s-channel in proton–proton collisions at $\sqrt{s} = 8$ TeV with the ATLAS detector using the Matrix Element Method'. In: *Phys. Lett. B* 756 (2016), p. 228. DOI: 10.1016/j.physletb.2016.03.017. arXiv: 1511.05980 [hep-ex] (cit. on p. 136).

[32] ATLAS Collaboration. 'Evidence for the associated production of a W boson and a top quark in ATLAS at $\sqrt{s} = 7$ TeV'. In: *Phys. Lett. B* 716 (2012), p. 142. DOI: 10.1016/j.physletb.2012.08.011. arXiv: 1205.5764 [hep-ex] (cit. on pp. 22, 108).

[33] ATLAS Collaboration. *Expected Performance of the ATLAS Experiment - Detector, Trigger and Physics*. 2009. arXiv: 0901.0512 [hep-ex] (cit. on p. 67).

[34] ATLAS Collaboration. 'Fiducial, total and differential cross-section measurements of t-channel single top-quark production in pp collisions at 8 TeV using data collected by the ATLAS detector'. In: *Eur. Phys. J. C* 77 (2017), p. 531. DOI: 10.1140/epjc/s10052-017-5061-9. arXiv: 1702.02859 [hep-ex] (cit. on p. 136).

[35] ATLAS Collaboration. *ID alignment trend plots in 2016*. 7 Dec. 2016. URL: https://atlas.web.cern.ch/Atlas/GROUPS/PHYSICS/PLOTS/IDTR-2016-016/ (cit. on p. 47).

[36] ATLAS Collaboration. *Impact Parameter resolutions in di-jet events*. 25 Sept. 2018. URL: https://atlas.web.cern.ch/Atlas/GROUPS/PHYSICS/PLOTS/IDTR-2018-008/ (cit. on p. 52).

[37] ATLAS Collaboration. *Improved electron reconstruction in ATLAS using the Gaussian Sum Filter-based model for bremsstrahlung*. ATLAS-CONF-2012-047. 2012. URL: https://cds.cern.ch/record/1449796 (cit. on p. 55).

[38] ATLAS Collaboration. 'Improved luminosity determination in pp collisions at $\sqrt{s} = 7$ TeV using the ATLAS detector at the LHC'. In: *Eur. Phys. J. C* 73 (2013), p. 2518. DOI: 10.1140/epjc/s10052-013-2518-3. arXiv: 1302.4393 [hep-ex] (cit. on pp. 47, 102, 128).

[39] ATLAS Collaboration. *Improvements in $t\bar{t}$ modelling using NLO+PS Monte Carlo generators for Run 2*. ATL-PHYS-PUB-2018-009. 2018. URL: https://cds.cern.ch/record/2630327 (cit. on pp. 74, 89).

[40] ATLAS Collaboration. *Jet Calibration and Systematic Uncertainties for Jets Reconstructed in the ATLAS Detector at $\sqrt{s} = 13$ TeV*. ATL-PHYS-PUB-2015-015. 2015. URL: https://cds.cern.ch/record/2037613 (cit. on p. 87).

[41] ATLAS Collaboration. *Jet energy resolution in 2017 data and simulation*. 5 Sept. 2018. URL: https://atlas.web.cern.ch/Atlas/GROUPS/PHYSICS/PLOTS/JETM-2018-005/ (cit. on p. 138).

[42] ATLAS Collaboration. *Jet energy scale and uncertainties in 2015-2017 data and simulation*. 27 Nov. 2018. URL: https://atlas.web.cern.ch/Atlas/GROUPS/PHYSICS/PLOTS/JETM-2018-006/ (cit. on p. 138).

[43] ATLAS Collaboration. 'Jet energy scale measurements and their systematic uncertainties in proton–proton collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector'. In: *Phys. Rev. D* 96 (2017), p. 072002. DOI: 10.1103/PhysRevD.96.072002. arXiv: 1703.09665 [hep-ex] (cit. on pp. 63, 64, 87).

[44] ATLAS Collaboration. 'Jet reconstruction and performance using particle flow with the ATLAS Detector'. In: *Eur. Phys. J. C* 77 (2017), p. 466. DOI: 10.1140/epjc/s10052-017-5031-2. arXiv: 1703.10485 [hep-ex] (cit. on pp. 63, 139).

[45] ATLAS Collaboration. *JVT Public Plots for ICHEP 2016*. 1 Aug. 2016. URL: https://atlas.web.cern.ch/Atlas/GROUPS/PHYSICS/PLOTS/JETM-2016-011/ (cit. on p. 67).

[46] ATLAS Collaboration. 'Luminosity determination in pp collisions at $\sqrt{s} = 8$ TeV using the ATLAS detector at the LHC'. In: *Eur. Phys. J. C* 76 (2016), p. 653. DOI: 10.1140/epjc/s10052-016-4466-1. arXiv: 1608.03953 [hep-ex] (cit. on p. 28).

[47] ATLAS Collaboration. 'Measurement of differential cross-sections of a single top quark produced in association with a W boson at $\sqrt{s} = 13$ TeV with ATLAS'. In: *Eur. Phys. J. C* 78 (2018), p. 186. DOI: 10.1140/epjc/s10052-018-5649-8. arXiv: 1712.01602 [hep-ex] (cit. on p. 110).

[48] ATLAS Collaboration. 'Measurement of the $t\bar{t}$ production cross-section using $e\mu$ events with b-tagged jets in pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector'. In: *Phys. Lett. B* 761 (2016), p. 136. DOI: 10.1016/j.physletb.2016.08.019. arXiv: 1606.02699 [hep-ex] (cit. on p. 128).

[49] ATLAS Collaboration. 'Measurement of the Z/γ^* boson transverse momentum distribution in pp collisions at $\sqrt{s} = 7$ TeV with the ATLAS detector'. In: *JHEP* 09 (2014), p. 145. DOI: 10.1007/JHEP09(2014)145. arXiv: 1406.3660 [hep-ex] (cit. on p. 75).

[50] ATLAS Collaboration. 'Measurement of the cross-section for producing a W boson in association with a single top quark in pp collisions at $\sqrt{s} = 13$ TeV with ATLAS'. In: *JHEP* 01 (2018), p. 063. DOI: 10.1007/JHEP01(2018)063. arXiv: 1612.07231 [hep-ex] (cit. on pp. 92, 129).

[51] ATLAS Collaboration. 'Measurement of the differential cross-section of highly boosted top quarks as a function of their transverse momentum in $\sqrt{s} = 8$ TeV proton–proton collisions using the ATLAS detector'. In: *Phys. Rev. D* 93 (2016), p. 032009. DOI: 10.1103/PhysRevD.93.032009. arXiv: 1510.03818 [hep-ex] (cit. on p. 130).

[52] ATLAS Collaboration. 'Measurement of the production cross-section of a single top quark in association with a W boson at 8 TeV with the ATLAS experiment'. In: *JHEP* 01 (2016), p. 064. DOI: 10.1007/JHEP01(2016)064. arXiv: 1510.03752 [hep-ex] (cit. on pp. 22, 108, 136).

[53] ATLAS Collaboration. 'Measurement of the production cross-section of a single top quark in association with a Z boson in proton–proton collisions at 13 TeV with the ATLAS detector'. In: *Phys. Lett. B* 780 (2018), p. 557. DOI: 10.1016/j.physletb.2018.03.023. arXiv: 1710.03659 [hep-ex] (cit. on p. 138).

[54] ATLAS Collaboration. 'Measurements of b-jet tagging efficiency with the ATLAS detector using $t\bar{t}$ events at $\sqrt{s} = 13$ TeV'. In: *JHEP* 08 (2018), p. 089. DOI: 10.1007/JHEP08(2018)089. arXiv: 1805.01845 [hep-ex] (cit. on p. 68).

[55] ATLAS Collaboration. *Monte Carlo Generators for the Production of a W or Z/γ^* Boson in Association with Jets at ATLAS in Run 2*. ATL-PHYS-PUB-2016-003. 2016. URL: https://cds.cern.ch/record/2120133 (cit. on p. 75).

[56] ATLAS Collaboration. 'Muon reconstruction performance of the ATLAS detector in proton–proton collision data at $\sqrt{s} = 13$ TeV'. In: *Eur. Phys. J. C* 76 (2016), p. 292. DOI: 10.1140/epjc/s10052-016-4120-y. arXiv: 1603.05598 [hep-ex] (cit. on pp. 60–62).

[57] ATLAS Collaboration. 'Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC'. In: *Phys. Lett. B* 716 (2012), p. 1. DOI: 10.1016/j.physletb.2012.08.020. arXiv: 1207.7214 [hep-ex] (cit. on p. 7).

[58] ATLAS Collaboration. *Optimisation and performance studies of the ATLAS b-tagging algorithms for the 2017-18 LHC run*. ATL-PHYS-PUB-2017-013. 2017. URL: https://cds.cern.ch/record/2273281 (cit. on pp. 67, 139).

[59] ATLAS Collaboration. *Optimisation of the ATLAS b-tagging performance for the 2016 LHC Run*. ATL-PHYS-PUB-2016-012. 2016. URL: https://cds.cern.ch/record/2160731 (cit. on p. 67).

[60] ATLAS Collaboration. 'Performance of b-jet identification in the ATLAS experiment'. In: *JINST* 11 (2016), P04008. DOI: 10.1088/1748-0221/11/04/P04008. arXiv: 1512.01094 [hep-ex] (cit. on pp. 67, 89).

[61] ATLAS Collaboration. 'Performance of missing transverse momentum reconstruction with the ATLAS detector using proton–proton collisions at $\sqrt{s} = 13\,\text{TeV}$'. In: *Eur. Phys. J. C* 78 (2018), p. 903. DOI: 10.1140/epjc/s10052-018-6288-9. arXiv: 1802.08168 [hep-ex] (cit. on pp. 69, 70, 88).

[62] ATLAS Collaboration. 'Performance of pile-up mitigation techniques for jets in *pp* collisions at $\sqrt{s} = 8\,\text{TeV}$ using the ATLAS detector'. In: *Eur. Phys. J. C* 76 (2016), p. 581. DOI: 10.1140/epjc/s10052-016-4395-z. arXiv: 1510.03823 [hep-ex] (cit. on p. 66).

[63] ATLAS Collaboration. 'Performance of the ATLAS track reconstruction algorithms in dense environments in LHC Run 2'. In: *Eur. Phys. J. C* 77 (2017), p. 673. DOI: 10.1140/epjc/s10052-017-5225-7. arXiv: 1704.07983 [hep-ex] (cit. on p. 51).

[64] ATLAS Collaboration. 'Performance of the ATLAS trigger system in 2015'. In: *Eur. Phys. J. C* 77 (2017), p. 317. DOI: 10.1140/epjc/s10052-017-4852-3. arXiv: 1611.09661 [hep-ex] (cit. on pp. 44, 59, 62, 63, 65).

[65] ATLAS Collaboration. *Pile-up subtraction and suppression for jets in ATLAS*. ATLAS-CONF-2013-083. 2013. URL: https://cds.cern.ch/record/1570994 (cit. on p. 66).

[66] ATLAS Collaboration. *Pixel vertical position evolution*. 13 Mar. 2018. URL: https://atlas.web.cern.ch/Atlas/GROUPS/PHYSICS/PLOTS/IDTR-2018-003/ (cit. on p. 46).

[67] ATLAS Collaboration. *Plots of b-tagging performance before and after the installation of the Insertable B-Layer*. 2017. URL: http://atlas.web.cern.ch/Atlas/GROUPS/PHYSICS/PLOTS/FTAG-2017-005/ (cit. on p. 68).

[68] ATLAS Collaboration. 'Probing the Wtb vertex structure in t-channel single-top-quark production and decay in *pp* collisions at $\sqrt{s} = 8\,\text{TeV}$ with the ATLAS detector'. In: *JHEP* 04 (2017), p. 124. DOI: 10.1007/JHEP04(2017)124. arXiv: 1702.08309 [hep-ex] (cit. on p. 136).

[69] ATLAS Collaboration. 'Probing the Quantum Interference between Singly and Doubly Resonant Top-Quark Production in *pp* Collisions at $\sqrt{s} = 13\,\text{TeV}$ with the ATLAS Detector'. In: *Phys. Rev. Lett.* 121 (2018), p. 152002. DOI: 10.1103/PhysRevLett.121.152002. arXiv: 1806.04667 [hep-ex] (cit. on pp. 90, 139, 140).

[70] ATLAS Collaboration. *Proposal for particle-level object and observable definitions for use in physics measurements at the LHC*. ATL-PHYS-PUB-2015-013. 2015. URL: https://cds.cern.ch/record/2022743 (cit. on p. 112).

[71] ATLAS Collaboration. *Selection of jets produced in* 13 *TeV proton–proton collisions with the ATLAS detector*. ATLAS-CONF-2015-029. 2015. URL: https://cds.cern.ch/record/2037702 (cit. on p. 66).

[72] ATLAS Collaboration. 'Simultaneous measurements of the $t\bar{t}$, W^+W^-, and $Z/\gamma^* \to \tau\tau$ production cross-sections in *pp* collisions at $\sqrt{s} = 7\,\text{TeV}$ with the ATLAS detector'. In: *Phys. Rev. D* 91 (2015), p. 052005. DOI: 10.1103/PhysRevD.91.052005. arXiv: 1407.0573 [hep-ex] (cit. on pp. 142, 143).

[73] ATLAS Collaboration. *Single Top Summary Plots Autumn 2019*. ATL-PHYS-PUB-2019-037. 2019. URL: http://cds.cern.ch/record/2690527 (cit. on pp. 108, 109).

[74] ATLAS Collaboration. *Studies of the ATLAS Inner Detector material using $\sqrt{s} = 13$ TeV pp collision data*. ATL-PHYS-PUB-2015-050. 2015. URL: https://cds.cern.ch/record/2109010 (cit. on pp. 33, 35).

[75] ATLAS Collaboration. *Summary of ATLAS Pythia 8 tunes*. ATL-PHYS-PUB-2012-003. 2012. URL: https://cds.cern.ch/record/1474107 (cit. on p. 73).

[76] ATLAS Collaboration. *Technical Proposal: A High-Granularity Timing Detector for the ATLAS Phase-II Upgrade*. 2018. URL: http://cds.cern.ch/record/2623663 (cit. on p. 53).

[77] ATLAS Collaboration. 'The ATLAS Experiment at the CERN Large Hadron Collider'. In: *JINST* 3 (2008), S08003. DOI: 10.1088/1748-0221/3/08/S08003 (cit. on pp. 25, 32, 41).

[78] ATLAS Collaboration. *The simulation principle and performance of the ATLAS fast calorimeter simulation FastCaloSim*. ATL-PHYS-PUB-2010-013. 2010. URL: https://cds.cern.ch/record/1300517 (cit. on p. 50).

[79] ATLAS Collaboration. 'Topological cell clustering in the ATLAS calorimeters and its performance in LHC Run 1'. In: *Eur. Phys. J. C* 77 (2017), p. 490. DOI: 10.1140/epjc/s10052-017-5004-5. arXiv: 1603.02934 [hep-ex] (cit. on p. 54).

[80] ATLAS Collaboration. *Trigger Menu in 2016*. ATL-DAQ-PUB-2017-001. 2017. URL: https://cds.cern.ch/record/2242069 (cit. on p. 59).

[81] ATLAS Collaboration. *Vertex Reconstruction Performance of the ATLAS Detector at* $\sqrt{s} =$ *13 TeV*. ATL-PHYS-PUB-2015-026. 2015. URL: https://cds.cern.ch/record/2037717 (cit. on p. 53).

[82] G. Avoni et al. 'The new LUCID-2 detector for luminosity measurement and monitoring in ATLAS'. In: *Journal of Instrumentation* 13.07 (2018), P07017–P07017. DOI: 10.1088/1748-0221/13/07/p07017 (cit. on p. 46).

[83] M. Baak et al. 'HistFitter software framework for statistical data analysis'. In: *Eur. Phys. J. C* 75 (2015), p. 153. DOI: 10.1140/epjc/s10052-015-3327-7. arXiv: 1410.1280 [hep-ex] (cit. on pp. 101, 103).

[84] BABAR Collaboration. 'Evidence for an Excess of $\overline{B} \to D^{(*)}\tau^-\overline{\nu}_\tau$ Decays'. In: *Phys. Rev. Lett.* 109 (10 2012), p. 101802. DOI: 10.1103/PhysRevLett.109.101802 (cit. on p. 7).

[85] R. D. Ball et al. 'Parton distributions from high-precision collider data'. In: *Eur. Phys. J.* C77.10 (2017), p. 663. DOI: 10.1140/epjc/s10052-017-5199-5. arXiv: 1706.00428 [hep-ph] (cit. on p. 11).

[86] R. D. Ball et al. 'Parton distributions with LHC data'. In: *Nucl. Phys. B* 867 (2013), p. 244. DOI: 10.1016/j.nuclphysb.2012.10.003. arXiv: 1207.1303 [hep-ph] (cit. on p. 74).

[87] R. D. Ball et al. 'Impact of Heavy Quark Masses on Parton Distributions and LHC Phenomenology'. In: *Nucl. Phys.* B849 (2011), pp. 296–363. DOI: 10.1016/j.nuclphysb.2011.03.021. arXiv: 1101.1300 [hep-ph] (cit. on p. 12).

[88] C. W. Bauer, F. J. Tackmann and J. Thaler. 'GenEvA. I. A New framework for event generation'. In: *JHEP* 12 (2008), p. 010. DOI: 10.1088/1126-6708/2008/12/010. arXiv: 0801.4026 [hep-ph] (cit. on p. 15).

[89] Belle Collaboration. 'Lepton-Flavor-Dependent Angular Analysis of $B \to K^*\ell^+\ell^-$'. In: *Phys. Rev. Lett.* 118.11 (2017), p. 111801. DOI: 10.1103/PhysRevLett.118.111801. arXiv: 1612.05014 [hep-ex] (cit. on p. 7).

[90] M. Beneke, A. Signer and V. Smirnov. 'Top quark production near threshold and the top quark mass'. In: *Physics Letters B* 454.1 (1999), pp. 137–146. ISSN: 0370-2693. DOI: 10.1016/S0370-2693(99)00343-3 (cit. on p. 20).

[91] M. Beneke, Y. Kiyo and K. Schuller. 'NNNLO results on top-quark pair production near threshold'. In: *PoS* RADCOR2007 (2007), p. 051. DOI: 10.22323/1.048.0051. arXiv: 0801.3464 [hep-ph] (cit. on p. 20).

[92] M. Beneke et al. 'Non-resonant and electroweak NNLO correction to the e^+e^- top anti-top threshold'. In: *JHEP* 02 (2018), p. 125. DOI: 10.1007/JHEP02(2018)125. arXiv: 1711.10429 [hep-ph] (cit. on p. 3).

[93] I. Bird et al. *Update of the Computing Models of the WLCG and the LHC Experiments*. Tech. rep. CERN-LHCC-2014-014. LCG-TDR-002. 2014. URL: https://cds.cern.ch/record/1695401 (cit. on p. 46).

[94] E. Boos et al. 'Generic User Process Interface for Event Generators'. In: *Physics at TeV colliders. Proceedings, Euro Summer School, Les Houches, France, May 21-June 1, 2001*. 2001. arXiv: hep-ph/0109068 [hep-ph] (cit. on p. 14).

[95] R. Bruce et al. 'LHC Run 2: Results and challenges'. In: CERN-ACC-2016-0103 (2016), MOAM5P50. 7 p. DOI: 10.18429/JACoW-HB2016-MOAM5P50 (cit. on p. 28).

[96] R. Brun and F. Rademakers. 'ROOT — An object oriented data analysis framework'. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 389.1 (1997). New Computing Techniques in Physics Research V, pp. 81–86. ISSN: 0168-9002. DOI: 10.1016/S0168-9002(97)00048-X (cit. on p. 102).

[97] O. S. Brüning et al. *LHC Design Report*. CERN Yellow Reports: Monographs. Geneva: CERN, 2004. DOI: 10.5170/CERN-2004-003-V-1 (cit. on p. 27).

[98] O. Brüning and L. Rossi. *The High Luminosity Large Hadron Collider*. World Scientific, 2015. DOI: 10.1142/9581 (cit. on p. 30).

[99] J. Butterworth et al. 'PDF4LHC recommendations for LHC Run II'. In: *J. Phys. G* 43 (2016), p. 023001. DOI: 10.1088/0954-3899/43/2/023001. arXiv: 1510.03865 [hep-ph] (cit. on p. 91).

[100] M. Cacciari and G. P. Salam. 'Pileup subtraction using jet areas'. In: *Physics Letters B* 659.1 (2008), pp. 119–126. ISSN: 0370-2693. DOI: 10.1016/j.physletb.2007.09.077 (cit. on p. 63).

[101] M. Cacciari, G. P. Salam and G. Soyez. 'The anti-k_t jet clustering algorithm'. In: *JHEP* 04 (2008), p. 063. DOI: 10.1088/1126-6708/2008/04/063. arXiv: 0802.1189 [hep-ph] (cit. on pp. 16, 17).

[102] R. Calaga. 'Crab Cavities for the High-luminosity LHC'. In: *18th International Conference on RF Superconductivity, Lanzhou, China, 17 - 21 Jul 2017*. 2018. DOI: 10.18429/JACoW-SRF2017-THXA03 (cit. on p. 30).

[103] J. M. Campbell et al. 'Top-Pair Production and Decay at NLO Matched with Parton Showers'. In: *JHEP* 04 (2015), p. 114. DOI: 10.1007/JHEP04(2015)114. arXiv: 1412.1828 [hep-ph] (cit. on p. 74).

[104] F. Cascioli, P. Maierhofer and S. Pozzorini. 'Scattering Amplitudes with Open Loops'. In: *Phys. Rev. Lett.* 108 (2012), p. 111601. DOI: 10.1103/PhysRevLett.108.111601. arXiv: 1111.5206 [hep-ph] (cit. on p. 75).

[105] S. Catani. 'Aspects of QCD, from the Tevatron to the LHC'. In: *1st Les Houches Workshop on Physics at TeV Colliders (PhysTeV 1999) Les Houches, France, June 7-18, 1999*. 2000. arXiv: hep-ph/0005233 [hep-ph] (cit. on p. 19).

[106] CDF Collaboration. 'Cross section measurements of high-p_T dilepton final-state processes using a global fitting method'. In: *Phys. Rev. D* 78 (1 2008), p. 012003. DOI: 10.1103/PhysRevD.78.012003 (cit. on p. 142).

[107] CDF Collaboration. 'Observation of Electroweak Single Top-Quark Production'. In: *Phys. Rev. Lett.* 103 (9 2009), p. 092002. DOI: 10.1103/PhysRevLett.103.092002 (cit. on p. 21).

[108] CDF Collaboration. 'Observation of Top Quark Production in $\overline{p}p$ Collisions with the Collider Detector at Fermilab'. In: *Phys. Rev. Lett.* 74 (14 1995), pp. 2626–2631. DOI: 10.1103/PhysRevLett.74.2626 (cit. on p. 18).

[109] CMS Collaboration. 'Evidence for associated production of a single top quark and W boson in pp collisions at $\sqrt{s} = 7\,\text{TeV}$'. In: *Phys. Rev. Lett.* 110 (2013), p. 022003. DOI: 10.1103/PhysRevLett.110.022003. arXiv: 1209.3489 [hep-ex] (cit. on pp. 22, 108).

[110] CMS Collaboration. 'Measurement of the t-channel single-top-quark production cross section and of the $|V_{tb}|$ CKM matrix element in pp collisions at $\sqrt{s} = 8\,\text{TeV}$'. In: *JHEP* 06 (2014), p. 090. DOI: 10.1007/JHEP06(2014)090. arXiv: 1403.7366 [hep-ex] (cit. on p. 136).

[111] CMS Collaboration. 'Measurement of the production cross section for single top quarks in association with W bosons in proton–proton collisions at $\sqrt{s} = 13\,\text{TeV}$'. In: *JHEP* 10 (2018), p. 117. DOI: 10.1007/JHEP10(2018)117. arXiv: 1805.07399 [hep-ex] (cit. on pp. 108, 138).

[112] CMS Collaboration. 'Observation of a new boson at a mass of 125 GeV with the CMS experiment at the LHC'. In: *Phys. Lett. B* 716 (2012), p. 30. DOI: 10.1016/j.physletb.2012.08.021. arXiv: 1207.7235 [hep-ex] (cit. on p. 7).

[113] CMS Collaboration. 'Observation of single top quark production in association with a Z boson in proton–proton collisions at $\sqrt{s} = 13\,\text{TeV}$'. In: *Phys. Rev. Lett.* 123 (2018), p. 132003. DOI: 10.1103/PhysRevLett.122.132003. arXiv: 1812.05900 [hep-ex] (cit. on p. 138).

[114] CMS Collaboration. 'Observation of the associated production of a single top quark and a W boson in pp collisions at $\sqrt{s} = 8\,\text{TeV}$'. In: *Phys. Rev. Lett.* 112 (2014), p. 231802. DOI: 10.1103/PhysRevLett.112.231802. arXiv: 1401.2942 [hep-ex] (cit. on pp. 22, 108, 136).

[115] CMS Collaboration. 'Search for s channel single top quark production in pp collisions at $\sqrt{s} = 7$ and $8\,\text{TeV}$'. In: *JHEP* 09 (2016), p. 027. DOI: 10.1007/JHEP09(2016)027. arXiv: 1603.02555 [hep-ex] (cit. on p. 136).

[116] CMS Collaboration. 'Search for associated production of a Higgs boson and a single top quark in proton–proton collisions at $\sqrt{s} = 13\,\text{TeV}$'. In: *Phys. Rev. D.* 99 (2018), p. 092005. DOI: 10.1103/PhysRevD.99.092005. arXiv: 1811.09696 [hep-ex] (cit. on p. 138).

[117] CMS Collaboration. 'The CMS experiment at the CERN LHC'. In: *JINST* 3 (2008), S08004. DOI: 10.1088/1748-0221/3/08/S08004 (cit. on p. 25).

[118] J. C. Collins, D. E. Soper and G. Sterman. 'Factorization for short distance hadron-hadron scattering'. In: *Nuclear Physics B* 261 (1985), pp. 104–142. ISSN: 0550-3213. DOI: 10.1016/0550-3213(85)90565-6 (cit. on p. 12).

[119] G. Corcella et al. 'HERWIG 6: an event generator for hadron emission reactions with interfering gluons (including supersymmetric processes)'. In: *JHEP* 01 (2001), p. 010. DOI: 10.1088/1126-6708/2001/01/010. arXiv: hep-ph/0011363 (cit. on pp. 14, 89).

[120] T. Cornelissen et al. 'The new ATLAS track reconstruction (NEWT)'. In: *Journal of Physics: Conference Series*. Vol. 119. 3. IOP Publishing. 2008, p. 032014. DOI: 10.1088/1742-6596/119/3/032014 (cit. on p. 51).

[121] G. Cowan. *Statistical Data Analysis*. Oxford University Press, 1998. ISBN: 0198501560 (cit. on pp. 103, 121).

[122] G. Cowan et al. 'Asymptotic formulae for likelihood-based tests of new physics'. In: *Eur. Phys. J. C* 71 (2011), p. 1554. DOI: 10.1140/epjc/s10052-011-1554-0. arXiv: 1007.1727 [physics.data-an] (cit. on pp. 103, 107). Erratum in: *Eur. Phys. J. C* 73 (2013), p. 2501. DOI: 10.1140/epjc/s10052-013-2501-z.

[123] M. Czakon and A. Mitov. 'Top++: A Program for the Calculation of the Top-Pair Cross-Section at Hadron Colliders'. In: *Comput. Phys. Commun.* 185 (2014), p. 2930. DOI: 10.1016/j.cpc.2014.06.021. arXiv: 1112.5675 [hep-ph] (cit. on p. 74).

[124] G. D'Agostini. 'A multidimensional unfolding method based on Bayes' theorem'. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 362.2 (1995), pp. 487–498. ISSN: 0168-9002. DOI: 10.1016/0168-9002(95)00274-X (cit. on p. 121).

[125] G. D'Agostini. *Bayesian Reasoning in Data Analysis: A Critical Introduction*. Singapore: World Scientific, 2003. ISBN: 9789812383563. DOI: 10.1142/5262 (cit. on pp. 119, 121).

[126] D0 Collaboration. 'Observation of Single Top-Quark Production'. In: *Phys. Rev. Lett.* 103 (9 2009), p. 092001. DOI: 10.1103/PhysRevLett.103.092001 (cit. on p. 21).

[127] D0 Collaboration. 'Search for High Mass Top Quark Production in $p\bar{p}$ Collisions at $\sqrt{s} =$ 1.8 TeV'. In: *Phys. Rev. Lett.* 74 (13 1995), pp. 2422–2426. DOI: 10.1103/PhysRevLett.74.2422 (cit. on p. 18).

[128] H. Damerau et al. *RF Manipulations for Higher Brightness LHC-Type Beams*. CERN-ACC-2013-0210. 2013. URL: https://cds.cern.ch/record/1595719 (cit. on p. 29).

[129] F. Demartin et al. 'tWH associated production at the LHC'. In: *Eur. Phys. J.* C77.1 (2017), p. 34. DOI: 10.1140/epjc/s10052-017-4601-7. arXiv: 1607.05862 [hep-ph] (cit. on pp. 23, 24).

[130] Y. L. Dokshitzer et al. 'Better jet clustering algorithms'. In: *JHEP* 08 (1997), p. 001. DOI: 10.1088/1126-6708/1997/08/001. arXiv: hep-ph/9707323 [hep-ph] (cit. on p. 16).

[131] C. Eck et al. *LHC computing Grid: Technical Design Report. Version 1.06 (20 Jun 2005)*. Technical Design Report LCG. Geneva: CERN, 2005. URL: https://cds.cern.ch/record/840543 (cit. on p. 46).

[132] J. Ellis et al. 'Implications of recent measurements of B meson mixing and $\varepsilon'/\varepsilon K$'. In: *Nuclear Physics B* 304 (1988), pp. 205–235. ISSN: 0550-3213. DOI: 10.1016/0550-3213(88)90625-6 (cit. on p. 17).

[133] R. K. Ellis, W. J. Stirling and B. R. Webber. *QCD and Collider Physics*. Cambridge: Cambridge University Press, 2003. ISBN: 0-521-54589-7 (cit. on p. 10).

[134] S. D. Ellis and D. E. Soper. 'Successive combination jet algorithm for hadron collisions'. In: *Phys. Rev. D* 48 (1993), pp. 3160–3166. DOI: 10.1103/PhysRevD.48.3160. arXiv: hep-ph/9305266 [hep-ph] (cit. on p. 16).

[135] F. Englert and R. Brout. 'Broken Symmetry and the Mass of Gauge Vector Mesons'. In: *Phys. Rev. Lett.* 13 (9 1964), pp. 321–323. DOI: 10.1103/PhysRevLett.13.321 (cit. on p. 6).

[136] L. Evans and P. Bryant. 'LHC Machine'. In: *JINST* 3 (2008), S08001. DOI: 10.1088/1748-0221/3/08/S08001 (cit. on p. 25).

[137] P. Falgari, A. S. Papanastasiou and A. Signer. 'Finite-width effects in unstable-particle production at hadron colliders'. In: *JHEP* 05 (2013), p. 156. DOI: 10.1007/JHEP05(2013)156. arXiv: 1303.5299 [hep-ph] (cit. on p. 20).

[138] S. Fartoukh. 'Achromatic telescopic squeezing scheme and application to the LHC and its luminosity upgrade'. In: *Phys. Rev. ST Accel. Beams* 16 (11 2013), p. 111002. DOI: 10.1103/PhysRevSTAB.16.111002 (cit. on p. 29).

[139] Fermilab. *Proton-Proton Collider Upgrade (New Main Injector, New Tevatron)*. FERMILAB-DESIGN-1988-01. 1988. URL: http://inspirehep.net/record/276291 (cit. on p. 17).

[140] Fermilab. *TEVATRON Phase 1*. FERMILAB-DESIGN-1980-01. 1980. DOI: 10.2172/1345599 (cit. on p. 17).

[141] J. H. Friedman. 'Data analysis techniques for high energy particle physics'. In: SLAC-176 (1974), 96 p. URL: https://cds.cern.ch/record/695770 (cit. on p. 102).

[142] J. H. Friedman. 'Greedy Function Approximation: A Gradient Boosting Machine'. In: *The Annals of Statistics* 29.5 (2001), pp. 1189–1232. ISSN: 00905364. URL: http://www.jstor.org/stable/2699986 (cit. on pp. 80, 81).

[143] S. Frixione, P. Nason and C. Oleari. 'Matching NLO QCD computations with parton shower simulations: the POWHEG method'. In: *JHEP* 11 (2007), p. 070. DOI: 10.1088/1126-6708/2007/11/070. arXiv: 0709.2092 [hep-ph] (cit. on pp. 14, 74).

[144] S. Frixione and B. R. Webber. 'Matching NLO QCD computations and parton shower simulations'. In: *JHEP* 06 (2002), p. 029. arXiv: hep-ph/0204244 (cit. on p. 14).

[145] S. Frixione et al. 'Single-top hadroproduction in association with a W boson'. In: *Journal of High Energy Physics* 2008.07 (2008), pp. 029–029. DOI: 10.1088/1126-6708/2008/07/029 (cit. on p. 23).

[146] R. Frühwirth. 'Application of Kalman filtering to track and vertex fitting'. In: *Nucl. Instrum. Meth. A* 262.2 (1987), pp. 444–450. DOI: 10.1016/0168-9002(87)90887-4 (cit. on p. 51).

[147] A. Giammanco. 'Single top quark production at the LHC'. In: *Reviews in Physics* 1 (2016), pp. 1–12. ISSN: 2405-4283. DOI: 10.1016/j.revip.2015.12.001 (cit. on p. 21).

[148] T. Gleisberg et al. 'Event generation with SHERPA 1.1'. In: *JHEP* 02 (2009), p. 007. DOI: 10.1088/1126-6708/2009/02/007. arXiv: 0811.4622 [hep-ph] (cit. on p. 75).

[149] T. Gleisberg and S. Höche. 'Comix, a new matrix element generator'. In: *JHEP* 12 (2008), p. 039. DOI: 10.1088/1126-6708/2008/12/039. arXiv: 0808.3674 [hep-ph] (cit. on pp. 14, 75).

[150] C. Grupen and B. Shwartz. *Particle Detectors*. 2nd ed. Cambridge Monographs on Particle Physics, Nuclear Physics and Cosmology 26. Cambridge University Press, 2008. ISBN: 9780521840064 (cit. on pp. 36, 38, 41).

[151] S. W. Herb et al. 'Observation of a Dimuon Resonance at 9.5 GeV in 400-GeV Proton-Nucleus Collisions'. In: *Phys. Rev. Lett.* 39 (5 1977), pp. 252–255. DOI: 10.1103/PhysRevLett.39.252 (cit. on p. 16).

[152] P. W. Higgs. 'Broken Symmetries and the Masses of Gauge Bosons'. In: *Phys. Rev. Lett.* 13 (16 1964), pp. 508–509. DOI: 10.1103/PhysRevLett.13.508 (cit. on p. 6).

[153] T. K. H. Ho. 'Random decision forests'. In: *Proceedings of 3rd International Conference on Document Analysis and Recognition*. Vol. 1. 1995, pp. 278–282. DOI: 10.1109/ICDAR.1995.598994 (cit. on p. 80).

[154] A. H. Hoang et al. 'Threshold $t\bar{t}$ cross section at next-to-next-to-leading logarithmic order'. In: *Phys. Rev. D* 65 (1 2001), p. 014014. DOI: 10.1103/PhysRevD.65.014014 (cit. on p. 20).

[155] A. H. Hoang et al. 'Top-Antitop Pair Production Close to Threshold Synopsis of Recent NNLO Results'. In: *EPJ direct* 2.1 (2000), pp. 1–22. ISSN: 1435-3725. DOI: 10.1007/s1010500c0003 (cit. on p. 20).

[156] A. H. Hoang. 'Top pair production at threshold and effective theories'. In: *Acta Phys. Polon.* B34 (2003), pp. 4491–4509. arXiv: hep-ph/0310301 [hep-ph] (cit. on p. 20).

[157] S. Höche et al. 'A critical appraisal of NLO+PS matching methods'. In: *JHEP* 2012.9 (2012), p. 49. DOI: 10.1007/JHEP09(2012)049 (cit. on p. 15).

[158] S. Höche et al. 'QCD matrix elements and truncated showers'. In: *JHEP* 05 (2009), p. 053. DOI: 10.1088/1126-6708/2009/05/053. arXiv: 0903.1219 [hep-ph] (cit. on p. 75).

[159] A. Hoecker et al. 'TMVA: Toolkit for Multivariate Data Analysis'. In: *PoS* ACAT (2007), p. 040. arXiv: physics/0703039 (cit. on p. 78).

[160] J. Illingworth and J. Kittler. 'A survey of the hough transform'. In: *Computer Vision, Graphics, and Image Processing* 44.1 (1988), pp. 87–116. ISSN: 0734-189X. DOI: 10.1016/S0734-189X(88)80033-1 (cit. on p. 60).

[161] M. Ježabek and J. H. Kühn. 'The Top width: Theoretical update'. In: *Phys. Rev. D* 48 (1993), R1910–R1913. DOI: 10.1103/PhysRevD.48.R1910. arXiv: hep-ph/9302295 [hep-ph] (cit. on p. 3). Erratum in: *Phys. Rev. D* 49 (9 1994), p. 4970. DOI: 10.1103/PhysRevD.49.4970.

[162] T. Ježo et al. 'An NLO+PS generator for $t\bar{t}$ and Wt production and decay including non-resonant and interference effects'. In: *Eur. Phys. J.* C76.12 (2016), p. 691. DOI: 10.1140/epjc/s10052-016-4538-2. arXiv: 1607.04538 [hep-ph] (cit. on p. 23).

[163] J. Jowett. *Beam separation and 3D envelopes (1σ) around IP1 for LHC version 6.4 collision optics*. URL: https://lhc-data-exchange.web.cern.ch/lhc-data-exchange/ruggiero.pdf (visited on 18/07/2018) (cit. on p. 29).

[164] B. P. Kersevan and I. Hinchliffe. 'A Consistent prescription for the production involving massive quarks in hadron collisions'. In: *JHEP* 09 (2006), p. 033. DOI: 10.1088/1126-6708/2006/09/033. arXiv: hep-ph/0603068 [hep-ph] (cit. on p. 23).

[165] N. Kidonakis. 'NNNLO soft-gluon corrections for the top-antitop pair production cross section'. In: *Phys. Rev.* D90.1 (2014), p. 014006. DOI: 10.1103/PhysRevD.90.014006. arXiv: 1405.7046 [hep-ph] (cit. on pp. 20, 21).

[166] N. Kidonakis. 'NNNLO soft-gluon corrections for the top-quark p_T and rapidity distributions'. In: *Phys. Rev.* D91.3 (2015), p. 031501. DOI: 10.1103/PhysRevD.91.031501. arXiv: 1411.2633 [hep-ph] (cit. on p. 21).

[167] N. Kidonakis. 'Single-top production in the Standard Model and beyond'. In: *13th Conference on the Intersections of Particle and Nuclear Physics (CIPANP 2018) Palm Springs, California, USA, May 29-June 3, 2018*. 2018. arXiv: 1808.02934 [hep-ph] (cit. on p. 21).

[168] N. Kidonakis. 'Theoretical results for electroweak-boson and single-top production'. In: *Proceedings, 23rd International Workshop on Deep-Inelastic Scattering and Related Subjects (DIS 2015): Dallas, Texas, USA, April 27-May 01, 2015*, p. 170. DOI: 10.22323/1.247.0170. arXiv: 1506.04072 [hep-ph] (cit. on pp. 74, 105).

[169] N. Kidonakis and R. Vogt. 'Theoretical top quark cross section at the Fermilab Tevatron and the CERN LHC'. In: *Phys. Rev. D* 78 (7 2008), p. 074005. DOI: 10.1103/PhysRevD.78.074005 (cit. on p. 20).

[170] M. Kobayashi and T. Maskawa. 'CP-Violation in the Renormalizable Theory of Weak Interaction'. In: *Progress of Theoretical Physics* 49.2 (1973), pp. 652–657. ISSN: 0033-068X. DOI: 10.1143/PTP.49.652 (cit. on pp. 6, 17).

[171] H.-L. Lai et al. 'New parton distributions for collider physics'. In: *Phys. Rev.* D82 (2010), p. 074024. DOI: 10.1103/PhysRevD.82.074024. arXiv: 1007.2241 [hep-ph] (cit. on p. 74).

[172] W. Lampl et al. *Calorimeter Clustering Algorithms: Description and Performance*. ATL-LARG-PUB-2008-002. 2008. URL: https://cds.cern.ch/record/1099735 (cit. on p. 55).

[173] D. J. Lange. 'The EvtGen particle decay simulation package'. In: *Nucl. Instrum. Meth. A* 462 (2001), p. 152. DOI: 10.1016/S0168-9002(01)00089-4 (cit. on p. 74).

[174] LHCb Collaboration. 'Measurement of the Ratio of Branching Fractions $\mathcal{B}(\overline{B}^0 \to D^{*+}\tau^-\overline{\nu}_\tau)$ / $\mathcal{B}(\overline{B}^0 \to D^{*+}\mu^-\overline{\nu}_\mu)$'. In: *Phys. Rev. Lett.* 115 (11 2015), p. 111803. DOI: 10.1103/PhysRevLett.115.111803 (cit. on p. 7).

[175] L. Lista. *Statistical Methods for Data Analysis in Particle Physics*. Lecture Notes in Physics 909. Springer, 2017. ISBN: 9783319628394. DOI: 10.1007/978-3-319-20176-4 (cit. on p. 103).

[176] E. Majorana. 'Teoria simmetrica dell'elettrone e del positrone'. In: *Il Nuovo Cimento (1924-1942)* 14.4 (1937), p. 171. DOI: 10.1007/BF02961314 (cit. on p. 4).

[177] Z. Maki, M. Nakagawa and S. Sakata. 'Remarks on the Unified Model of Elementary Particles'. In: *Progress of Theoretical Physics* 28.5 (1962), pp. 870–880. ISSN: 0033-068X. DOI: 10.1143/PTP.28.870 (cit. on p. 4).

[178] M. L. Mangano and T. J. Stelzer. 'Tools for the simulation of hard hadronic collisions'. In: *Annual Review of Nuclear and Particle Science* 55.1 (2005), pp. 555–588. DOI: 10.1146/annurev.nucl.55.090704.151505 (cit. on p. 9).

[179] A. D. Martin et al. 'Parton distributions for the LHC'. In: *Eur. Phys. J. C* 63 (2009), p. 189. DOI: 10.1140/epjc/s10052-009-1072-5. arXiv: 0901.0002 [hep-ph] (cit. on p. 73).

[180] S. van der Meer. *Stochastic damping of betatron oscillations in the ISR*. 1972. URL: https://cds.cern.ch/record/312939 (cit. on p. 17).

[181] E. Mobs. *The CERN accelerator complex. Complexe des accélérateurs du CERN*. 2016. URL: https://cds.cern.ch/record/2197559 (cit. on p. 26).

[182] T. Morii, C. Lim and S. Mukherjee. *The Physics of the Standard Model and Beyond*. World Scientific, 2004. ISBN: 9789812795601 (cit. on p. 5).

[183] Muon (g-2) Collaboration. 'Measurement of the Negative Muon Anomalous Magnetic Moment to 0.7 ppm'. In: *Phys. Rev. Lett.* 92 (16 2004), p. 161802. DOI: 10.1103/PhysRevLett.92.161802 (cit. on p. 7).

[184] S. Myers. 'The Large Hadron Collider 2008–2013'. In: *International Journal of Modern Physics A* 28.25 (2013). DOI: 10.1142/S0217751X13300354 (cit. on p. 27).

[185] P. Nason. 'A new method for combining NLO QCD with shower Monte Carlo algorithms'. In: *JHEP* 11 (2004), p. 040. DOI: 10.1088/1126-6708/2004/11/040. arXiv: hep-ph/0409146 (cit. on p. 74).

[186] J. Pequenao. *Computer Generated image of the ATLAS calorimeter*. 2008. URL: https://cds.cern.ch/record/1095927 (cit. on p. 37).

[187] J. Pequenao. *Computer generated image of the whole ATLAS detector*. 2008. URL: https://cds.cern.ch/record/1095924 (cit. on p. 31).

[188] G. Piacquadio, K. Prokofiev and A. Wildauer. 'Primary vertex reconstruction in the ATLAS experiment at LHC'. In: *Journal of Physics: Conference Series*. Vol. 119. 3. IOP Publishing. 2008, p. 032033. DOI: 10.1088/1742-6596/119/3/032033 (cit. on p. 53).

[189] W. H. Press et al. *Numerical Recipes in C: The Art of Scientific Computing*. 2nd ed. New York, NY, USA: Cambridge University Press, 1992. ISBN: 0-521-43108-5 (cit. on p. 81).

[190] J. Pumplin et al. 'New generation of parton distributions with uncertainties from global QCD analysis'. In: *JHEP* 07 (2002), p. 012. arXiv: hep-ph/0201195 (cit. on p. 74).

[191] A. Quadt. *Top Quark Physics at Hadron Colliders*. Advances in the Physics of Particles and Nuclei 28, ed. by D. H. Beck, D. Haidt and J. W. Negele. Berlin and Heidelberg: Springer-Verlag, 2007. ISBN: 9783540710608 (cit. on pp. 10, 18, 20). (Orig. pub. in 2006).

[192] C. Quigg. 'Unanswered Questions in the Electroweak Theory'. In: *Annual Review of Nuclear and Particle Science* 59.1 (2009), pp. 505–555. DOI: 10.1146/annurev.nucl.010909.083126 (cit. on p. 18).

[193] E. Re. 'Single-top Wt-channel production matched with parton showers using the POWHEG method'. In: *Eur. Phys. J.* C71 (2011), p. 1547. DOI: 10.1140/epjc/s10052-011-1547-z. arXiv: 1009.2450 [hep-ph] (cit. on p. 74).

[194] A. Romano. 'Electron cloud formation in CERN particle accelerators and its impact on the beam dynamics. Entstehung von Elektronenwolken in den Teilchenbeschleunigern des CERN und ihre Auswirkungen auf die Strahldynamik'. 2018. URL: https://cds.cern.ch/record/2676577 (cit. on p. 30).

[195] S. Schumann and F. Krauss. 'A Parton shower algorithm based on Catani-Seymour dipole factorisation'. In: *JHEP* 03 (2008), p. 038. DOI: 10.1088/1126-6708/2008/03/038. arXiv: 0709.1027 [hep-ph] (cit. on p. 75).

[196] C. Schwinn. 'Theory status of hadronic top-quark pair production'. In: *30th Rencontres de Blois on Particle Physics and Cosmology Blois, France, June 3-8, 2018*. 2018. arXiv: 1810.05096 [hep-ph] (cit. on p. 20).

[197] K. Seidel et al. 'Top quark mass measurements at and above threshold at CLIC'. In: *The European Physical Journal C* 73.8 (2013), p. 2530. ISSN: 1434-6052. DOI: 10.1140/epjc/s10052-013-2530-7. URL: https://doi.org/10.1140/epjc/s10052-013-2530-7 (cit. on p. 3).

[198] T. Sjöstrand, S. Mrenna and P. Z. Skands. 'PYTHIA 6.4 Physics and Manual'. In: *JHEP* 05 (2006), p. 026. DOI: 10.1088/1126-6708/2006/05/026. arXiv: hep-ph/0603175 (cit. on pp. 14, 74).

[199] T. Sjöstrand et al. 'An Introduction to PYTHIA 8.2'. In: *Comput. Phys. Commun.* 191 (2015), p. 159. DOI: 10.1016/j.cpc.2015.01.024. arXiv: 1410.3012 [hep-ph] (cit. on p. 73).

[200] P. Z. Skands. 'Tuning Monte Carlo Generators: The Perugia Tunes'. In: *Phys. Rev. D* 82 (2010), p. 074018. DOI: 10.1103/PhysRevD.82.074018. arXiv: 1005.3457 [hep-ph] (cit. on p. 74).

[201] K. W. Staley. *The Evidence for the Top Quark: Objectivity and Bias in Collaborative Experimentation*. Cambridge University Press, 2004. ISBN: 978-0521827102 (cit. on p. 17).

[202] C. Suster. *Measuring the ATLAS beam spot offline in the run 2 environment*. Restricted to current ATLAS members. 2016. URL: https://cds.cern.ch/record/2196938 (cit. on p. 53).

[203] T. M. P. Tait and C.-P. Yuan. 'Single top quark production as a window to physics beyond the standard model'. In: *Phys. Rev. D* 63 (1 2000), p. 014018. DOI: 10.1103/PhysRevD.63.014018 (cit. on p. 22).

[204] T. M. P. Tait. 'The tW^- mode of single top production'. In: *Phys. Rev.* D61 (1999), p. 034001. DOI: 10.1103/PhysRevD.61.034001. arXiv: hep-ph/9909352 [hep-ph] (cit. on p. 22).

[205] M. Tanabashi, K. Hagiwara, K. Hikasa et al. 'Review of Particle Physics'. In: *Phys. Rev. D* 98 (3 2018), p. 030001. DOI: 10.1103/PhysRevD.98.030001 (cit. on pp. 3–5, 12, 18, 38, 41).

[206] TOTEM Collaboration. 'First measurement of elastic, inelastic and total cross-section at $\sqrt{s}$=13 TeV by TOTEM and overview of cross-section data at LHC energies'. In: *The European Physical Journal C* 79.2 (2019), p. 103. DOI: 10.1140/epjc/s10052-019-6567-0 (cit. on p. 8).

[207] UA1 Collaboration. 'Experimental observation of isolated large transverse energy electrons with associated missing energy at $\sqrt{s}$=540 GeV'. In: *Physics Letters B* 122.1 (1983), pp. 103–116. ISSN: 0370-2693. DOI: 10.1016/0370-2693(83)91177-2 (cit. on p. 17).

[208] UA1 Collaboration. 'Experimental observation of lepton pairs of invariant mass around 95 GeV/c^2 at the CERN SPS collider'. In: *Physics Letters B* 126.5 (1983), pp. 398–410. ISSN: 0370-2693. DOI: 10.1016/0370-2693(83)90188-0 (cit. on p. 17).

[209] UA2 Collaboration. 'Evidence for $Z^0 \rightarrow e^+e^-$ at the CERN $p\bar{p}$ collider'. In: *Physics Letters B* 129.1 (1983), pp. 130–140. ISSN: 0370-2693. DOI: 10.1016/0370-2693(83)90744-X (cit. on p. 17).

[210] UA2 Collaboration. 'Observation of single isolated electrons of high transverse momentum in events with missing transverse energy at the CERN $p\bar{p}$ collider'. In: *Physics Letters B* 122.5 (1983), pp. 476–485. ISSN: 0370-2693. DOI: 10.1016/0370-2693(83)91605-2 (cit. on p. 17).

[211] W. Verkerke and D. Kirkby. *The RooFit toolkit for data modeling*. 2003. arXiv: physics/0306116 [physics.data-an] (cit. on p. 101).

[212] J. Wenninger. *Operation and Configuration of the LHC in Run 2*. CERN-ACC-NOTE-2019-0007. 2019. URL: https://cds.cern.ch/record/2668326 (cit. on pp. 27–29).

[213] M. Wobisch and T. Wengler. 'Hadronization corrections to jet cross-sections in deep inelastic scattering'. In: *Monte Carlo generators for HERA physics. Proceedings, Workshop, Hamburg, Germany, 1998-1999*. 1998, pp. 270–279. arXiv: hep-ph/9907280 [hep-ph] (cit. on p. 16).

Acknowledgements

Over the course of my research I have worked with a great many people, as is to be expected in the context of a large international collaboration, and I am grateful for their assistance.

Within the ATLAS top community and the collaboration more broadly there are too many to thank individually. The AIDA team who introduced me to my first analysis also propelled me towards some of the most supportive communities within the collaboration; in particular I'd like to thank Aldo Saavedra, Tony Limosani, and Kevin Finelli. Thanks to Anthony Morley for showing me the ropes at CERN and introducing me to the operational side of the experiment; I had some exciting times living in the calibration loop. I acknowledge the many conveners and sub-conveners who tirelessly support the analysers. Thanks to my fellow Australians within ATLAS, particularly those I lived and worked alongside at CERN; thank you Millie and Brian. To Rui, Doug, Kevin, Jin, Regina, and Ian: you've each had an impact on my development as a physicist and contributed to my enjoyment of the whole experience; thanks for everything over the years we've worked together.

The particle physics group at Sydney has been a big part of my life for many years, and has always been an easy-going and stimulating research community. Despite the huge diversity of research interests across theory and multiple experiments in a relatively small group, it is still a single group, and I appreciate the breadth of perspectives and ideas I've been exposed to as a result. Thanks to everyone from the group, but especially to Kevin Varvell and Bruce Yabsley who have been my mentors since I joined the field. Kevin has made sure that I've had all possible opportunities and guided me through the task of distilling a project out of the prodigious scope of the experiment. Bruce's perspective and experience have made me a better experimentalist. Thank you both, and thanks also for your assistance in correcting this thesis.

Thanks to my friends and family for their support as always. *Grazie soprattutto a Benedetto; forse sarei riuscito a scrivere una tesi senza di te, ma non sarrebe stata questa tesi e non sarei quello che sono oggi.*