Measurement of Associated Z Boson and Heavy-Quark Jet Production in pp Collisions at $\sqrt{s} = 13$ TeV

A thesis submitted by

Alec Drobac

in partial fulfillment of the requirements for the degree of

Doctor of Philosphy

in

Physics

Tufts University

August 2022

Advisor: Pierre-Hugues Beauchemin

Abstract

This thesis presents measurements of the production cross-sections of a Z boson in association with b-jets and with c-jets. The measurements are performed using 139 fb^{-1} of proton-proton collision data collected at $\sqrt{s} = 13$ TeV by the ATLAS experiment at the Large Hadron Collider. The measurement is performed using a novel unfolding technique which simultaneously unfolds the b-jet and c-jet distributions while preserving the correlations between them. Inclusive and differential cross-sections are measured for events containing a Z boson decaying into two electrons or two muons and produced in association with at least one b- or c-jet. Three predictions from Monte Carlo generators, using leading order or next-to-leading order matrix elements and interfaced to parton shower models, are used to test the different shower models and matching/merging schemes. All three predictions overestimate the number of b-jet events while underestimating the number of c-jet events. The results motivate further study, using the analysis techniques developed for this thesis, of the simulation and prediction of heavy-jet production.

Acknowledgements

There are many people without whom this dissertation would not have been possible. I must first acknowledge my advisor, Prof. Pierre-Hugues (Hugo) Beauchemin, whose accessibility, compassion, and enthusiasm for particle physics brought me to this work and to ATLAS. He has been my single greatest source of inspiration for what it means to be a physicist.

I also want to acknowledge my analysis team at ATLAS, with whom I learned and grew as a scientist. In particular, I want to thank Dr. Vincent Croft for his tireless work and teaching on our analysis framework and on unfolding; Dr. Jonathan Bossio and Prof. Federico Sforza, who repeatedly and kindly educated me on all matters ATLAS; and Dr. Semen Turchikhin, whose rapid production of systematic variations was essential to the completion of this thesis.

I have been surrounded at Tufts by a loving and tirelessly supportive community. In particular, I give my profound thanks to: Dr. Katie Mason, Dr. Joshua Mills, and Colette Kaya, my colleagues and friends; Prof. Timothy Atherton, my teaching mentor; the members of the Beauchemin HEP research group, for their inspiration; the entire Tufts physics department, for their joy and companionship; Dr. Miriam Johnston, for her guidance and commiseration on the pursuit of a doctorate; and my parents, family, and friends, for their love.

Finally, I give thanks (as I do so often) that my fiancée, Hannah Johnston, has been by my side throughout this entire journey. Through my greatest triumphs and my darkest moments, she has continually reminded me why I began this work and how far I have come. This thesis is a reflection of not only my dedication to this research but her dedication to me.

Contents

List of Figures

Measurement of Associated Z Boson and Heavy-Quark Jet Production in pp Collisions at $\sqrt{s} = 13$ TeV

Introduction

Particle physics, the study of the fundamental particles and their interactions, attempts to understand the physical universe at its most basic level. The Standard Model (SM) describing all of the non-gravitational interactions between the known particles is perhaps the most successful scientific theory ever developed in terms of predictive power. It has been shown to accurately and precisely describe hundreds of different interactions between particles to astonishing precision: for example, the prediction of the electron's anomalous magnetic moment has been verified to 12 digits, making it the most accurately verified theoretical value in the history of physics.

Particle colliders such as the Large Hadron Collider (LHC), which are used to scatter particles (protons and heavy ions, in the case of the LHC) off of one another, have played a central role in developing and testing the Standard Model. At the extremely high energies of the LHC, the inelastic scattering of the protons, or equivalently the scattering of the quark and gluon constituents of the protons, leads to a varying number of final-state particles. Large detectors such as the ATLAS (A Toroidal LHC ApparatuS) detector identify those final-state particles, and analyses are performed in order to reconstruct from those particles the interaction which took place between the initial quarks and gluons.

Since the scattering of quarks and gluons is governed by quantum mechanics, their production and interactions are probabilistic in nature. There is no way for a collider experiment to explicitly control which final state will be produced in a given collision, or which production mechanism will be responsible for a particular final state, since there can be many interactions leading to the same outcome. Each collision at the LHC is effectively the roll of a thousand dice, each with a thousand sides. Therefore, the key to understanding the particles and their interactions is to roll the dice so many times that the probability is in our favor to see the same outcomes multiple times. Once a given topology of final-state particles has been observed enough times, with enough precision, the observations can be compared to the predicted cross-section, which quantifies the probability of of a given scattering process to occur. Only when the observed and predicted cross-sections are consistent with one another can we conclude whether or not the observed final state is likely due to the underlying interaction predicted by the theory. In this sense, modern experimental particle

physics, which seeks to measure increasingly rare particle interactions, is the science of overcoming extremely small odds through the brute force of repetition.

The interactions which govern the scattering processes we observe in high-energy physics experiments include the strong, weak, and electromagnetic interactions. The non-Abelian nature of the strong interaction, as described by the theory of Quantum Chromodynamics (QCD), leads to the confinement of quarks and gluons in composite particles such as protons and also to the production of jets. These features of QCD lead to messy scattering processes which are difficult to theoretically model; this further distances the final-state particles seen by particle detectors from the particles involved in the initial scattering process. Therefore, scattering processes which can serve as tests of QCD modelling are an especially important piece of a particle physics experiment's physics program.

One such test of QCD is the measurement of Z+jet events. Around one out of every billion pp collisions at the LHC results in the production of a Z boson, the neutral mediator of the weak interaction, and the probability for a Z to be produced in association with hadronic jets is even smaller. These Z+jets events are an essential probe of QCD because the production of the Z forces the event to be at an energy scale at which the theory of perturbative QCD (pQCD) can be used to perform fully analytical calculations of the matrix element, the hard-scattering process; phenomenological corrections, which are described in Chapter 1, are then applied in order to produce a prediction relevant to experimental pp collisions. Of particular interest is the measurement of Z+jets events which include jets originating from b- and c-quarks; such jets are referred to as "heavy-flavor jets" because of the large b- and c-quark masses. Measurements of these Z+heavy-flavor jets (Z+HF) events are important for a number of reasons. The first is that these measurements help deepen our understanding of features of pQCD — for example, the role gluon splitting into massive quarks plays in scattering processes — as well as the proton's structure. The parton distribution functions, which describe the momentum carried by the quarks and gluons inside the proton, must be determined from experiment, and so Z+HF measurements can help determine the b- and c-quark distribution functions. For c-quarks in particular, there is also the potential to detect an intrinsic charm component of the proton, an as-yet unobserved bound state of the proton which includes a $c\bar{c}$ component more massive than the proton itself. Z+HF events also must be measured precisely in order to improve the irreducible background estimation of other measurements. For example, the $H \to b\bar{b}$ channel is the dominant decay mode for the Higgs boson, with a branching ratio of $\approx 58\%$, and Z+HF events are one of the largest backgrounds to the observation of these Higgs decays [1, 2]. Measuring this decay channel is key to understanding the Higgs self coupling as well as the symmetry-breaking mechanism of the SM. Various beyond the Standard Model (BSM) searches for hypothetical particles such as

supersymmetric top squarks [3] and new pseudoscalars [4] will also benefit from a more precise modelling of their Z+HF background, increasing the potential to discover new physics.

Recent measurements of Z+HF events [5, 6] demonstrate discrepancies between theoretical predictions and the data taken at the LHC; however, the precision of the measurements needs to be improved in in order to determine the source of these discrepancies. This thesis seeks to improve the precision of the Z+HF measurement through the use of the full Run 2 ATLAS dataset (139 fb^{-1}) and the development of a novel method of unfolding: by simultaneously correcting the b-jet and c-jet contributions for detector smearing, simultaneous measurements of both distributions can be made while preserving their correlations. These techniques have been developed for the first time in this thesis, and they have the potential to address the tensions seen by the earlier measurements and to impact our understanding of the SM and beyond.

This thesis is organized as follows: Chapter 1 provides a theoretical overview, including brief descriptions of the Standard Model, QCD, and the production of Z bosons. Chapter 2 describes the LHC and the ATLAS experiment, while Chapter 3 describes how particles and jets are identified and reconstructed from ATLAS data. Chapter 4 motivates in detail the study Z+b-jet and Z+c-jet events. Chapter 5 describes the data and simulated samples used in the analysis, and Chapter 6 describes the selections applied to the samples, as well as the measured observables. Chapter 7 describes the procedure used to fit the light-jet background component to the data before background subtraction and unfolding; the unfolding technique developed for this analysis, in order to account for the correlation between b-jets and c-jets, is presented in Chapter 8. Chapter 9 gives an overview of the sources of uncertainty and their relative contributions to the overall measurement uncertainty. Chapter 10 presents the results of the analysis, including the inclusive total cross-section and differential cross-section measurements. Finally, Chapter 11 summarizes what has been learned and suggests potential future studies.

Chapter 1

Theoretical Background

1.1 The Standard Model

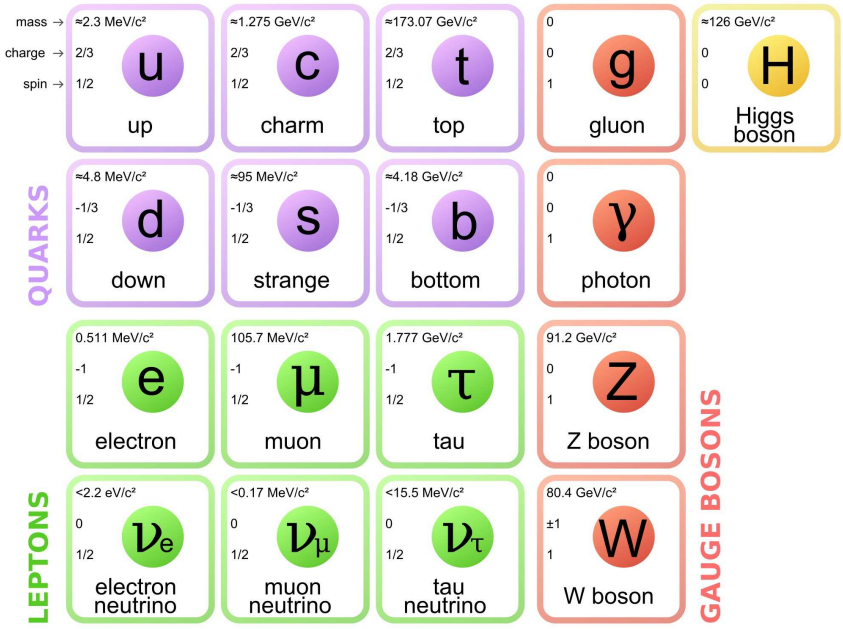

Figure 1.1: The fundamental particles of the Standard Model of particle physics [7].

The Standard Model (SM) of particle physics describes all of the known fundamental particles of nature (Figure 1.1) and their interactions. While there are notable gaps in this theory, not least of which is the total absence of gravity, it has been experimentally verified countless times by decades of particle physics experiments.

The fundamental particles are divided into two distinct groups. The first group is the fermions, with half-integer spin; they are further divided into the quarks, which have the color charge of the strong interaction, and the leptons, which do not. The other group is the gauge bosons, with integer spin, which mediate the electromagnetic, weak, and strong interactions, plus the Higgs

boson. The SM Lagrangian describes how these particles interact with each other. As will be covered in the next Section, the forms of the interaction terms between these particle fields stem from the requirement that the Lagrangian densities $\mathcal{L}$ of these particle fields have local gauge invariance.

1.2 Gauge theories

The analysis discussed in this thesis is primarily a study of QCD. Here I provide a brief overview of the gauge theory and its primary characteristics, though I start from the somewhat simpler theory of Quantum Electrodynamics (QED) to develop much of the mathematics.

1.2.1 Quantum electrodynamics

For the purposes of illustrating how the local gauge invariance of a theory involving vector and fermion fields results in a fundamental interaction between the fields, I will begin by discussing QED. The starting point is a Lagrangian density (I will simply refer to it as the Lagrangian from here) containing the Dirac (spin-$\frac{1}{2}$) fermions particles we have observed in experiments as well as the electromagnetic field strength tensor $F_{\mu\nu} = \partial_\mu A_\nu - \partial_\nu A_\mu$, where A_μ is the photon field:

$$\mathcal{L}_o = \overline{\psi}\left(i\gamma^\mu \partial_\mu - m\right)\psi - \frac{1}{4}F_{\mu\nu}F^{\mu\nu}. \tag{1.1}$$

So far the interactions between these particles have not been specified; this is the "free" Lagrangian, in that the Hamiltonian one could obtain from this Lagrangian would include no potential energy. The goal is to obtain the simplest possible Lagrangian which can describe what we see in experiments. This means that, among other requirements, the Lagrangian needs to be invariant under the Poincaré group of transformations. This group of transformations includes the Lorentz transformation of special relativity as well as space-time translations and rotations; invariance under these transformations ensures conservation of energy, momentum, and angular momentum, respectively. Invariance will also ensure that the theory is stable, in that there is no state with energy lower than the vacuum state; local, in that signals travel no faster than the speed of light; unitary, meaning the predictions of the theory can be interpreted probabilistically; and renormalizable, a feature which will be described later in this Chapter.

The Lagrangian 1.1 is not Poincaré-invariant, however. Under the transformations $U(\Lambda)$, where Λ is the Poincaré transformation in Minkowski space and $U(\Lambda)$ the equivalent transformation acting on vector fields, the photon field A_μ transforms as an invariant vector only up to a "gauge" (phase)

$\alpha(x, \Lambda)$:

$$U(\Lambda)A_\mu(x)U^*(\Lambda) = \Lambda^\nu_\mu A_\nu + \partial_\mu \alpha(x, \Lambda) \tag{1.2}$$

The gauge $\alpha(x, \Lambda)$ has no physical meaning. The situation is analogous to that of classical electromagnetism, where the scalar potential V is determined only up to a constant c, so the substitution $V \to V + c$ changes nothing. The crucial difference, however, is that c has a single value at all locations, and so it represents an example of only *global* gauge invariance; here, $\alpha(x, \Lambda)$ varies at each space-time point x, and so to satisfy Poincaré invariance the theory of QED is required to be *locally* gauge invariant. A local gauge transformation will result in the vector field transforming according to

$$A_\mu \to A_\mu + \partial_\mu \alpha(x, \Lambda) \tag{1.3}$$

and the fermion fields transforming according to

$$\psi(x) \to e^{i\alpha(x)}\psi(x), \tag{1.4}$$

and so the requirement of local gauge invariance boils down to the construction of a Lagrangian which does not change if the substitutions Equation 1.3 and Equation 1.4 are made. The mass term $-m\bar{\psi}\psi$ in Equation 1.1 is already invariant under the substitution in Equation 1.4, but the first term involving the partial derivative $\partial_\mu \psi$ is not. Therefore, in order to make this Lagrangian locally gauge invariant, the partial derivative is replaced with a covariant derivative:

$$\partial_\mu \to D_\mu \equiv \partial_\mu - ieA_\mu. \tag{1.5}$$

Making this substitution, the QED Lagrangian becomes

$$\begin{aligned} \mathcal{L} &= i\overline{\psi}\left(i\gamma^\mu D_\mu - m\right)\psi - \frac{1}{4}F_{\mu\nu}F^{\mu\nu} \\ &= \overline{\psi}\left(i\gamma^\mu \partial_\mu - m\right)\psi + e\overline{\psi}\gamma^\mu \psi A_\mu - \frac{1}{4}F_{\mu\nu}F^{\mu\nu} \\ &= \mathcal{L}_o + \mathcal{L}_{\text{int}}. \end{aligned} \tag{1.6}$$

This Poincaré-invariant Lagrangian now has an interaction term, $e\overline{\psi}\gamma^\mu\psi A_\mu$, describing how charged fermions ψ interact with the photon field A_μ; the interaction is governed by the coupling strength e. In summary, by starting with the most general free Lagrangian and adding terms in order to

satisfy local gauge invariance, a theory of QED has been constructed which fully determines the possible interactions, is physically meaningful, and is extremely consistent with observations [8, 9].

1.2.2 Quantum chromodynamics

Similarly to how the Lagrangian of QED is largely specified using the requirement of local $U(1)$ gauge invariance, the Lagrangian of QCD is constructed by requiring local invariance under the gauge transformations described by the group $SU(3)$. The most general free Lagrangian of QCD looks similar to Equation 1.1:

$$\mathcal{L}_o = \overline{q}_i \left(i\gamma^\mu \partial_\mu - m\right) q_i - \frac{1}{4} G^a_{\mu\nu} G^{\mu\nu}_a. \tag{1.7}$$

There are several key differences, however. Because fermions transform according to the fundamental representation of the gauge group, which for $SU(3)$ is of dimension 3, there are now three quark fermion fields q_1, q_2, q_3 instead of a single fermion field ψ. The three fields correspond to the three "colors" of strong charge: red, green, and blue. In the same manner, the free QCD Lagrangian includes $a = 1, \ldots, 8$ for each generator T_a of the group $SU(3)$ because the vector gauge fields transform according to the adjoint representation of the group, which for $SU(3)$ is of dimension 8. The QCD field strength tensor

$$G^a_{\mu\nu} = \partial_\mu G^a_\nu - \partial_\nu G^a_\mu - g f_{abc} G^b_\mu G^c_\nu \tag{1.8}$$

is therefore formed from eight gluon fields G^a_μ. But there is also a new term, $-g f_{abc} G^b_\mu G^c_\nu$, in the QCD field strength tensor $G^a_{\mu\nu}$ which is not present in $F_{\mu\nu}$. This is because the group $SU(3)$ is non-Abelian, which means that the generators do not commute with one another:

$$[T_a, T_b] = i f_{abc} T_c, \tag{1.9}$$

where the structure constants f_{abc} characterize the commutation relations. QED is a representation of the Abelian $U(1)$ group, meaning the generators of the group commute with one another and the structure constants are all equal to zero. The non-Abelian nature of QCD dramatically changes the nature of its interactions when compared to QED, as we will see multiple times in this Chapter.

As in QED, each of the gluon fields G^a_μ is only Poincaré invariant up to a local gauge:

$$G^a_\mu \to G^a_\mu - \partial \alpha_a(x, \Lambda) - f_{abc} \alpha_b(x, \Lambda) G^c_\mu. \tag{1.10}$$

Once again, the Lagrangian must be adapted so that it is locally gauge invariant. Under a local $SU(3)$ gauge transformation, the fermion fields transform according to

$$q(x) \to e^{i\alpha_a(x)T_a} q(x). \tag{1.11}$$

Therefore, as in QED, the partial derivative in Equation 1.7 must be replaced with a covariant derivative D_μ. The final gauge-invariant QCD Lagrangian is

$$\mathcal{L} = \bar{q}_i \left(i\gamma^\mu \partial_\mu - m\right) q_i - g\left(\bar{q}\gamma^\mu T_a q\right) G^a_\mu - \frac{1}{4} G^a_{\mu\nu} G^{\mu\nu}_a. \tag{1.12}$$

The result is an interaction Lagrangian for QCD which looks quite similar to Equation 1.6: it includes a term $-g\left(\bar{q}\gamma^\mu T_a q\right) G^a_\mu$ describing the interactions between the colored fermions q_i and the colored gluon fields G^a_μ, governed by the QCD coupling strength g. But the extra term in $G^a_{\mu\nu}$, a result to the non-Abelian nature of QCD, means that, unlike the photon of QED, there are also self-interaction terms for the gluon fields; whereas the photon carries no electric charge, the gluons have color. The 3- and 4-gluon vertices these terms described are sketched in Figure 1.2. The importance of these terms will be discussed in the following Sections.

$$\begin{aligned}\mathcal{L}_o &= \bar{q}_i \left(i\gamma^\mu \partial_\mu - m\right) q_i - g\left(\bar{q}\gamma^\mu T_a q\right) G^a_\mu - \frac{1}{4} G^a_{\mu\nu} G^{\mu\nu}_a \\ &= \underbrace{\text{“}\bar{q}q\text{”} \quad + \quad \text{“}(\partial G)^2\text{”} \quad + \quad g\text{“}\bar{q}qG\text{”}}_{\text{QED, QCD Lagrangians}} \quad + \quad \underbrace{g\text{“}G^2\partial G\text{”} \quad + \quad g^2\text{“}G^4\text{”}}_{\text{QCD only}}\end{aligned}$$

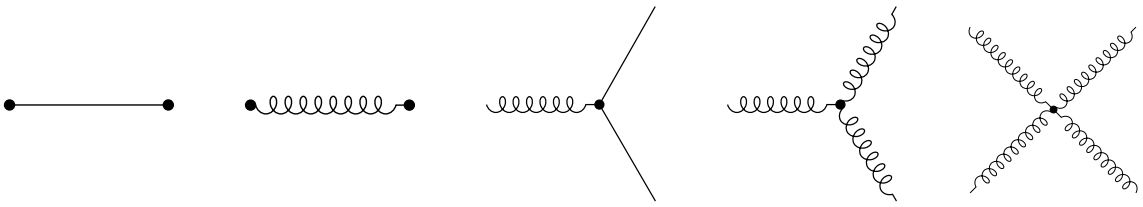

Figure 1.2: Above: QCD Lagrangian terms summarized (using “ ”) according to their field content. Below: pictoral representations of the terms of the QCD Lagrangian. The first three terms appear analogously in the QED Lagrangian as well; the latter two, which involve boson interaction terms, only appear in QCD. Adapted from [8].

1.3 Renormalization

So far, the couplings e and g_s have been used for QED and QCD, respectively, without much thought as to what they represent. In a Feynman diagram, such as the middle diagram of Figure 1.2, they represent the strength of the interaction between the fermion and the boson, called the *bare* coupling. For example, the strength of the interaction between the photon and an electron

is governed by the electric charge e. The value of this bare coupling is not determined by the theory, however, and cannot be measured directly; experimenters such as myself (and perhaps the reader!) only have access to the *effective* coupling, the strength of the interaction due to the sum of all possible diagrams of the given process. Renormalization is the procedure by which the bare quantities are replaced by their experimental counterparts in order for physical theories using these quantities to be representative of what we measure in experiments. However, many of the possible diagrams being included in the renormalized definitions are divergent, as they involve integrals over the four-momenta of the interacting particles, and those momenta can be arbitrarily large. Therefore, the cost of renormalization is to make the couplings "constants" dependent on an arbitrary energy scale at which their values can be determined through experiment. Below is a brief overview of the renormalization of QED and QCD and the ramifications of the procedure in each case.

1.3.1 Quantum electrodynamics

I will start once again with QED, for which the $\mathcal{O}(e^2)$ corrections to the $e^-\gamma e^-$ vertex are shown in Figures 1.3b–1.3e. The latter three diagrams show loop corrections to the electron four-vector current. Integrating these loops over all fermion energies leads to high-energy ("ultraviolet" (UV) in the usual parlance) divergences; however, due to the fact that QED is a locally gauge-invariant theory, the UV divergence of the diagram in Figure 1.3c is exactly cancelled by the similar divergences of the diagrams in 1.3d and 1.3e. This cancellation, known as the Ward identity, means that only 1.3b needs to be estimated in order to renormalize the QED coupling ([10, 11]).

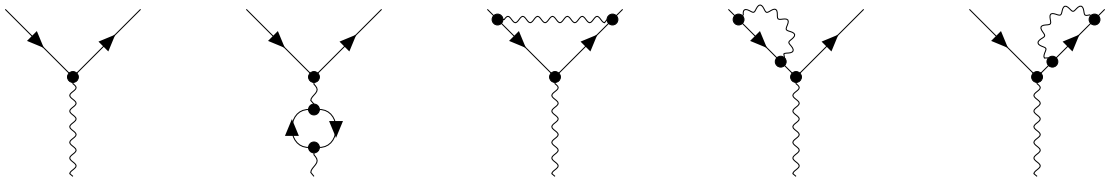

(a) (b) (c) (d) (e)

Figure 1.3: Lowest-order QED vertex diagram (a) and $\mathcal{O}(e^2)$ corrections (b-d). Adapted from [12].

Versions of the diagram in Figure 1.3b with two, three, or infinitely many loops are possible, and all of the versions, which are based on the bare coupling of the theory, must be accounted for in order to determine the effective coupling. To deal with this infinite series of corrections, the lowest-order photon exchange diagram can be replaced with the infinite series of photon loop corrections (Figure 1.4). This cannot be done for free, however; whereas the bare photon propagator is defined as $P_o = e_o^2/q^2$, when the lowest-order photon exchange diagram is redefined as the sum of all the possible loop diagrams, each loop contributes a correction factor $\pi(q^2)$, and so the effective

propagator for the redefined lowest-order diagram goes as

$$P(q^2) = P_o + P_o\pi(q^2)P_o + P_o\pi(q^2)P_o\pi(q^2)P_o + ... = \frac{P_o}{1 - \Pi(q^2)e_o^2} \equiv \frac{e^2(q^2)}{q^2}, \tag{1.13}$$

where $\Pi(q^2) = \pi(q^2)/q^2$. The second term, for example, corresponds to the one-loop diagram, with a photon propagator on either side of the loop (like in Figure 1.3b).

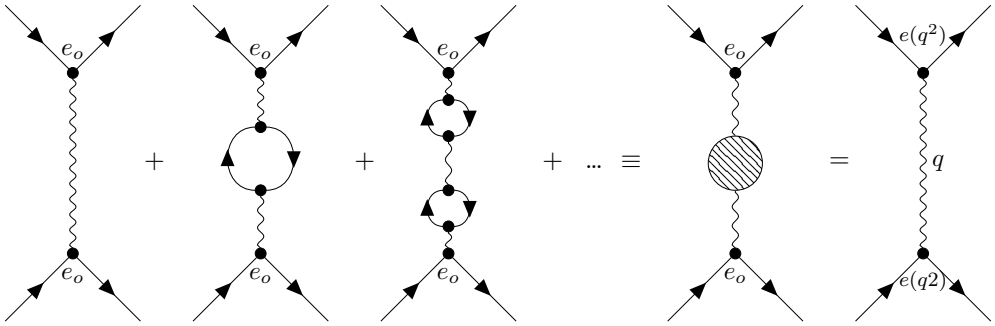

Figure 1.4: Renormalization of QED: the running coupling $e(q^2)$ is the bare charge e_o with the higher order loop corrections absorbed. Adapted from [12].

In summary, Equation 1.13 shows that calculating the infinite series of diagrams in Figure 1.4 results in a photon propagator whic has the same form as the free propagator but replaces the bare coupling e_o with an effective coupling $e(q^2)$. It is worrisome, though, that the effective coupling is defined in terms of an infinite series of infinite coefficients (since q^2 can be arbitrarily large). This is not inherently a problem, since q^2 can simply be chosen to be the scale at which e can be determined experimentally, thus avoiding the issue of have to evaluate $e(q^2 = \infty)$ analytically. However, this coupling is used to calculate particle interaction probabilities, so what happens if experimenters want to make a prediction at some other scale? To get around this, it is first assumed that we know the effective coupling at some reference scale μ^2. Equation 1.13 can then be rearranged to redefine the bare coupling e_o as

$$e_o = \frac{e(\mu^2)}{1 + \pi(\mu^2)e(\mu^2)} \tag{1.14}$$

and this definition of the bare coupling can then be used to define the effective coupling at the scale of interest q^2:

$$e(q^2) = \frac{e(\mu^2)}{1 - [\pi(q^2) - \pi(\mu^2)]\, e(\mu^2)}. \tag{1.15}$$

Now $e(q^2)$ is defined in terms of the finite difference $[\pi(q^2) - \pi(\mu^2)]$ of two divergent quantities. Thus the overall consequence of renormalization is that the effective coupling strength is defined

relative to the coupling strength at some other, arbitrarily chosen, reference scale at which it is determined experimentally; it is said that the coupling strength *runs* with q^2. Once this reference scale is chosen, the renormalized theory completely specifies the coupling strength at all scales through the q^2 dependence of Equation 1.15.

As far as the Lagrangian is concerned, μ^2 is an arbitrary parameter, so no observable $\mathcal{O}$ measured in experiments can directly depend on it (though it will depending on the running of e in reference to that scale). This requirement is expressed by the renormalization group equation

$$\left[\mu^2\frac{\partial}{\partial\mu^2}+\mu^2\frac{\partial e}{\partial\mu^2}\frac{\partial}{\partial e}\right]O=0 \tag{1.16}$$

or, more compactly,

$$\beta(e)=\mu^2\frac{\partial e}{\partial\mu^2}. \tag{1.17}$$

Here β is a perturbative expansion of scale dependence of e. The coefficient of the first term, $\beta_0 e$, corresponds to the one-loop correction and is given by

$$\beta_0=\frac{1}{12\pi}(2n_f-11C_A), \tag{1.18}$$

where n_f is number of degrees of freedom for fermions in the loops of Figure 1.4 and $\delta_{cd}C_A=\sum f_{abc}f_{abd}$ defines the quadratic Casimir operator C_A in terms of the structure constants describing the group (such as in Equation 1.9) [13].

Replacing $e(q^2)$ by $\alpha(q^2)=e^2(q^2)/4\pi$ to make the comparison to QCD in the next subsection visually simpler, the solution to the renormalization group equation looks like

$$\frac{1}{\alpha(\mu^2)}=\frac{1}{\alpha(q^2)}+\beta_0\ln\left(\frac{q^2}{\mu^2}\right)+... \tag{1.19}$$

QED, an Abelian theory, has structure constants all equal to 0, and therefore $C_A=0$; at sufficiently high energies, there are $n_f=3$ fermion families which can participate in the vacuum polarization loops. Taking these values in Equations 1.19 and 1.18, the QED running coupling is obtained:

$$\alpha_{\text{em}}(q^2)=\frac{\alpha(\mu^2)}{\left[1-\frac{1}{3\pi}\alpha(\mu^2)\ln\left(\frac{q^2}{\mu^2}\right)\right]}. \tag{1.20}$$

This coupling strength increases as q^2 increases, as evidenced by the minus sign in the denominator;

simply put, the electromagnetic interaction gets stronger at higher energies. For example, taking $\alpha_{\text{em}} = 1/137$ at the reference scale $\mu = 1$ MeV, $\alpha(M_Z) = 1/129$ at the mass of the Z boson (≈ 100 GeV) [14].

1.3.2 Quantum chromodynamics

The renormalization procedure of QCD is very similar to that of QED, but once again, the non-Abelian nature of QCD is of great significance to the final result. The nonzero structure constants f_{abc} result in $C_A = 3$; combining this with the $n_f = 3$ quark families, $\beta_0 = -9/(4\pi)$, and the running QCD coupling becomes

$$\alpha_s(q^2) = \frac{\alpha(\mu^2)}{\left[1 + \frac{9}{4\pi}\alpha(\mu^2)\ln\left(\frac{q^2}{\mu^2}\right)\right]}. \tag{1.21}$$

In comparison to Equation 1.20, the minus sign in the denominator has become a plus sign, which means that the QCD running coupling strength *decreases* as q^2 increases or, equivalently, as the interaction distance decreases. This phenomenon is known as *asymptotic freedom* [15]. As noted in Figure 1.5, at low energies ($\lesssim 1$ GeV), $\alpha_s \approx 1$. This makes performing a pertubative expansion approximation of QCD at this energy scale impossible: the terms in the perturbative expansion are too large, and so their sum diverges. On the other hand, at energies $\gtrsim 100$ GeV, $\alpha_s \approx 0.1$, which allows for perturbation theory to be applied to QCD; since this is roughly the energy scale of Z boson production, analyses focused on the associated production of this particle are an ideal choice for perturbative QCD (pQCD) investigations.

1.3.3 (Anti)screening

The difference in behavior of the running couplings can also be seen by observing the shielding behavior of the fermions coupled to these fields. In QED, an electron can spontaneously emit a photon, which can decay into an electron-positron pair. Pairs produced in this way are then polarized by the charge of the electron such that the positrons are closer to the electron. Because of this, the electron charge is "screened:" as a probe moves closer to the electron, it penetrates the electron-positron screen and "sees" more of the electron's charge. This means that, in effect, the charge increases as the distance to the particle is decreased.

This screening process occurs in QCD as well, but an additional, and stronger, anti-screening process also takes place, in which e.g. a particle with red color charge is surrounded by other red charges. The net result is that the observed color charge of a quark instead decreases as the distance to the quark is decreased (Figure 1.6), or equivalently as q^2 increases.

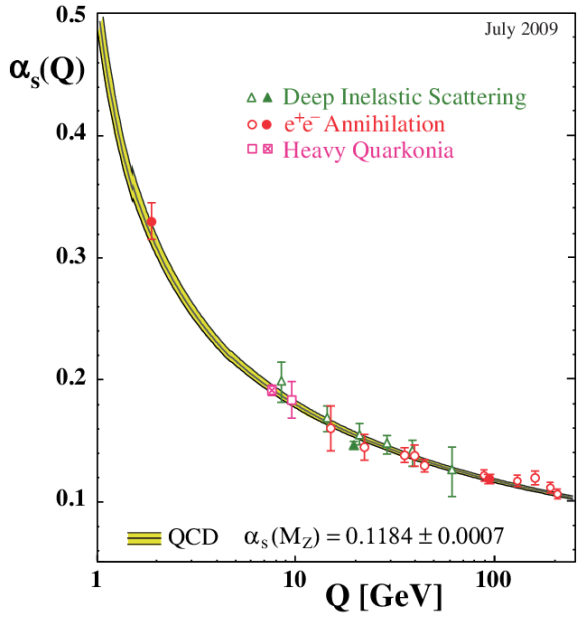

Figure 1.5: Running of the strong coupling constant α_s with energy scale Q. The curve shows QCD predictions for the combined world average value of $\alpha_s(M_Z)$ [16].

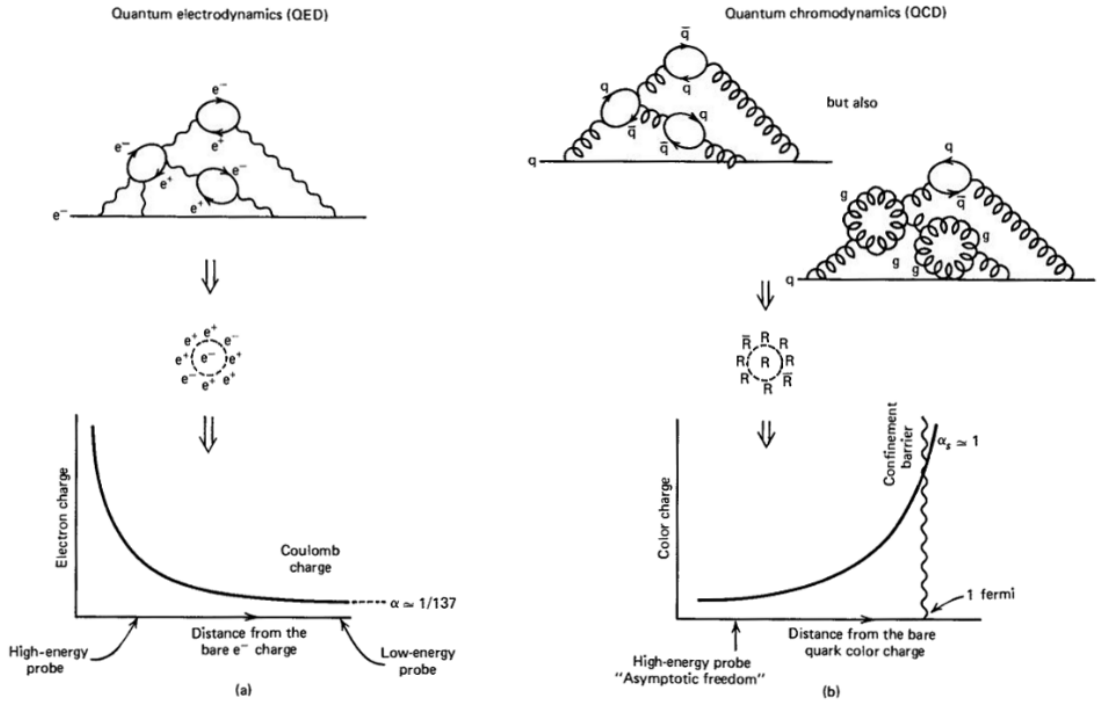

Figure 1.6: Screening of (a) electric and (b) color charge [8].

1.4 QCD cross-section measurements

With renormalized field theories defined in terms of effective couplings, the next question to address is how we, the particle physics experimenters, can use these theories to make predictions. The two primary observable quantities which it would be useful to predict and then measure are the particle resonance width Γ, the inverse of its decay rate, and the cross-section σ, a measure of the probability

of interaction between two particles in a scattering process. The generic cross-section goes as

$$\sigma = \frac{1}{64\pi^2 s} \int |\mathcal{M}_{fi}|^2 d\Omega, \tag{1.22}$$

where s is the square of the center-of-mass energy; the evaluation of the matrix element (ME)

$$\mathcal{M}_{fi} = -\langle f| \mathcal{L}_{\text{int}}(x=0) |i\rangle - \frac{-i^2}{2!} \int d^4x \langle f| T[\mathcal{L}_{\text{int}}(x)\mathcal{L}_{\text{int}}(x=0)] |i\rangle + \dots, \tag{1.23}$$

which describes the probability of the transition $i \to f$, is the central calculation which must be performed. The ME is a perturbative expansion of the interaction Lagrangian $\mathcal{L}_{\text{int}}$, and each higher order term in the expansion corresponds to the addition of another loop in the scattering process, such as in Figure 1.4 (which are accounted for in the integral of Equation 1.23 in a time-ordered manner, hence the $T[...]$ in the second term). At leading order, the ME corresponds to a diagram like the first in Figure 1.4, which includes two vertices; from the expression for $\mathcal{L}_{\text{int}}$ in Equation 1.6, we know that each vertex depends on the coupling constant e, so the ME will depend on $e^2 \propto \alpha$. Therefore, $|\mathcal{M}_{\text{LO}}|^2 \propto \alpha^2$ for a leading order QED diagram. If we perform the perturbative expansion of the ME and consider higher order diagrams, such as the next two diagrams in Figure 1.4 (though these are not the only possibilities), we have to sum the individual amplitudes and then square. At each higher order, another loop is introduced, and so there is another factor of α in the ME of the diagram; taking this into account, we can pull out these factors of α and write the perturbative expansion as

$$\mathcal{M}_{fi} = \alpha M_{\text{LO}} + \alpha^2 \sum_j M_{1,j} + \alpha^3 \sum_k M_{2,k} + \dots, \tag{1.24}$$

where $\mathcal{M}_{\text{LO}} = \alpha M_{\text{LO}}$. Here $\sum_j \mathcal{M}_{1,j}$ sums over the amplitudes of all next-to-leading order diagrams, $\sum_k \mathcal{M}_{2,k}$ sums over the next-to-next-to-leading order diagram amplitudes, etc. $|\mathcal{M}_{fi}|^2$ is then the square of the sum of these amplitudes [12]:

$$\begin{aligned} |\mathcal{M}_{fi}|^2 &= \left(\alpha M_{\text{LO}} + \alpha^2 \sum_j M_{1,j} + \dots\right)\left(\alpha M^*_{\text{LO}} + \alpha^2 \sum_{j'} M^*_{1,j'} + \dots\right) \\ &= \alpha^2 |M_{\text{LO}}|^2 + \alpha^3 \sum_j \left(M_{\text{LO}} M^*_{1,j} + M^*_{\text{LO}} M_{1,j}\right) + \alpha^4 \sum_{j,j'} M_{1,j} M^*_{1,j'} + \dots. \end{aligned} \tag{1.25}$$

Thus the perturbative expansion of the ME is an expansion in terms of the coupling constant α, which, according to our QED Lagrangian 1.6, governs the strength of the particle interactions. The same expansion, but in terms of α_s, is seen in pQCD (Equation 1.12). The interaction Lagrangians are the heart of the Standard Model, and measurements of the cross-section provide

direct information about the underlying physics of the scattering processes.

In this manner, using pQCD, we can make cross-section predictions for processes where i and j are partons (quarks and gluons) interacting via the strong coupling. At the LHC, however, it is not quarks which are collided, but protons *containing* quarks. So we also have to account for the likelihood that a given quark of each proton is involved in the underlying scattering event that we see. The cross-section of a hard-scattering process initiated by two hadrons (protons) a and b can therefore be written as

$$\sigma_{ab\to X} = \sum_{i,j} \int_0^1 dx_a \int_0^1 dx_b \underbrace{f_i(x_a, \mu_F^2) f_j(x_2, \mu_F^2)}_{\text{PDFs}} \otimes \underbrace{\hat{\sigma}_{ij\to X}\big(p_i, p_j, \alpha_s(\mu_F^2), Q^2/\mu_F^2\big)}_{\hat{\sigma}} . \tag{1.26}$$

As conveyed by the convolution operator $\otimes$, the cross-section can be *factorized* into two parts: the short-distance, high-energy, perturbative hard process $\hat{\sigma}$ describing the probability that the scattering of quarks i and j leads to the final state X, and the long-distance, low-energy, non-perturbative parton distribution functions (PDFs) $f_i(x)$ describing the probability to find quark i (j) with momentum fraction x_1 (x_2) in hadron P_1 (P_2). The factorization scale μ_F acts as a resolution parameter by effectively separating these two portions; emissions above this scale are absorbed into $\hat{\sigma}_{ij}$, while emissions below the scale are handled by the PDFs.

The factorization theorem 1.26 not only separates the perturbative and non-perturbative portions of the calculation but says that they are independent; this means that the same estimation of the non-perturbative physics can be used in any other prediction. With the perturbative portion handled by the renormalized theory of QCD, we can now focus on this non-perturbative portion, which includes the PDFs, parton showers, and fragmentation and hadronization [13].

1.4.1 Deep inelastic scattering

Consider elastic scattering processes such as those shown in Figure 1.7. As the energy of an incident particle increases, the cross-section decreases because the target is more and more likely to break apart; this is inelastic scattering. With enough energy we reach the regime of deep inelastic scattering (DIS), where the incident particle now scatters off of the constituents of the initial target particle. If those constituent parts of the original target are point-like — that is, they do not have constituents themselves — we then see that the deep inelastic cross-section is independent of q^2; the inelastic scattering of the probe with the target as a whole can be understood in terms of the elastic scattering of the probe with the constituents of the target.

This lack of q^2 dependence in the scattering cross-section was predicted for the proton by James Bjorken [18]. The inelastic scattering cross-section for an electron (with initial momentum $k = (E, \boldsymbol{k})$ in the laboratory frame) off of a proton (with initial momentum ($p = (M, 0)$) through

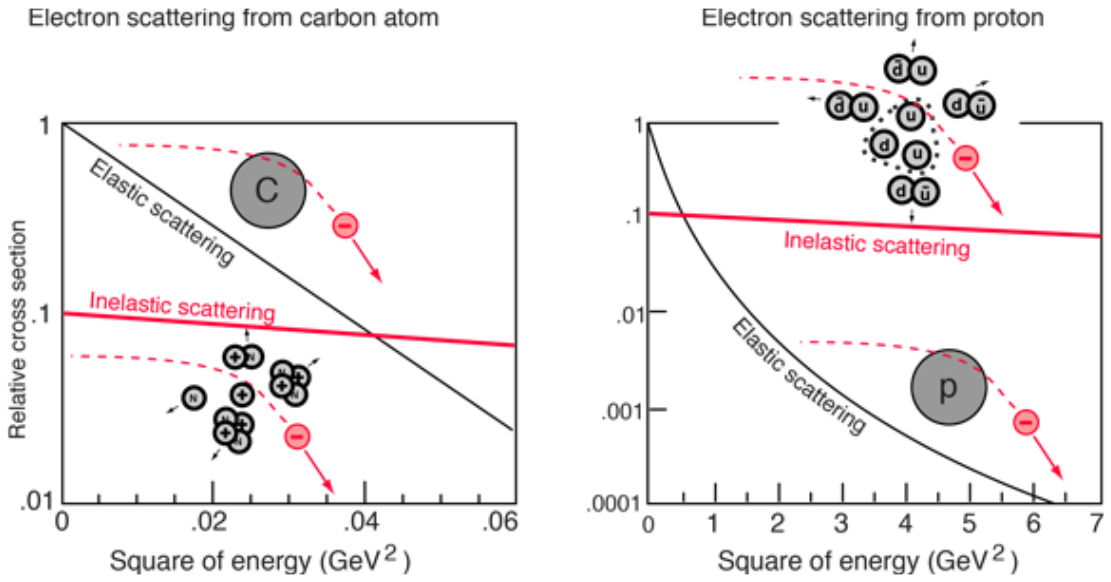

Figure 1.7: Electron elastic and inelastic scattering off of a carbon atom (left) and proton (right) [17].

angle θ, agnostic to the underlying structure of the proton, is

$$\left(\frac{d\sigma}{d\Omega}\right) = \frac{\alpha^2 E^2}{4k^2 \sin^4\left(\frac{\theta}{2}\right)} \left(W_2(\nu, q^2)\cos^2\left(\frac{\theta}{2}\right) + 2W_1(\nu, q^2)\sin^2\left(\frac{\theta}{2}\right)\right), \tag{1.27}$$

where $W_{1,2}$ are proton structure functions of $\nu = (p\cdot q)/M$ and q^2. Bjorken theorized that as q^2 got very large, these structure functions would exhibit scaling behavior, losing their q^2 dependence:

$$\left.\begin{array}{l} Q^2 \to \infty \\ \dfrac{\nu}{Q^2}\text{fixed} \end{array}\right\} \begin{array}{l} MW_1(\nu, Q^2) \to F_1(x) \\ \nu W_2(\nu, Q^2) \to F_2(x). \end{array} \tag{1.28}$$

Here we have made the switch to $Q^2 = -q^2$ — it is equivalent the inverse wavelength squared of the photon, or the resolving power of our DIS microscope — and introduced $x = Q^2/2M\nu$, the Bjorken scaling variable [18]. The above relations were confirmed by a collaboration of researchers at MIT and at the Stanford Linear Accelerator (SLAC), as shown in Figure 1.8; the structure functions exhibit exactly the behavior suggestive of point-like constituents of the proton [19].

1.4.2 Parton distribution functions

The evidence of point-like partons leads to the parton model of the proton, first proposed by Richard Feynman, in which the total inelastic cross-section is expressed as a sum of elastic collisions with the partons (Figure 1.9). As discussed at the beginning of this Section, the essence and validity of this model is what allows for the factorization of the pp cross-section. In the parton model, we

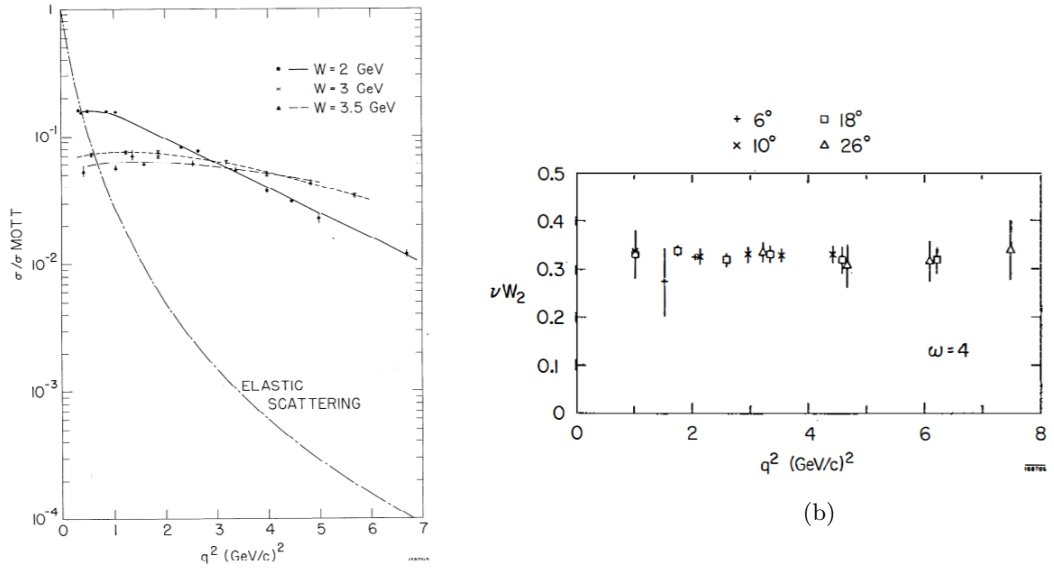

(a) (b)

Figure 1.8: (a) Early (1969) data suggesting that the proton structure functions are nearly independent of q^2 [20]. (b) Friedman, Kendall, and Taylor won the 1990 Nobel Prize for their work evidencing point-like constituents of the proton [19].

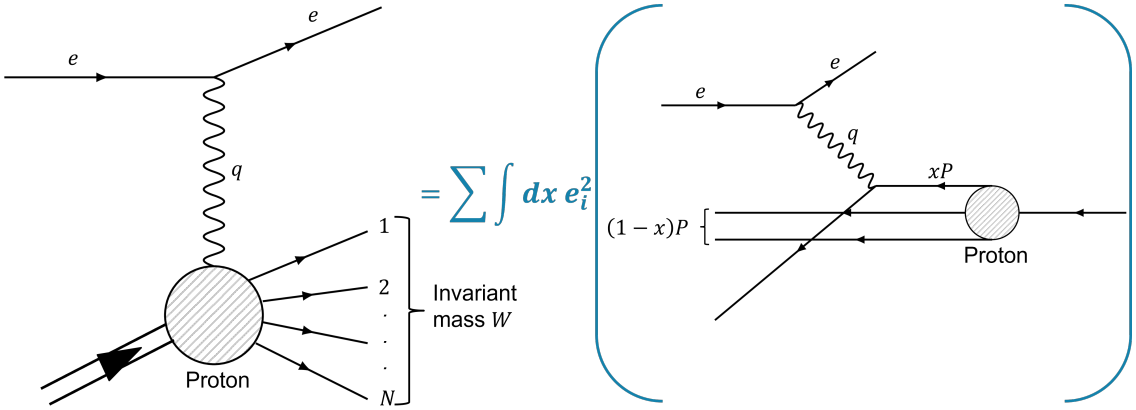

Figure 1.9: A schematic of the parton model. Adapted from [8].

assume that the quarks in the proton are effectively free. Furthermore, we assume that the energy of the proton $E \gg m_p$; this means that we can not only neglect m_p, but we can also neglect the component of the momentum of the struck quark transverse to the direction of the proton. This means that the momemtum of the protons is the sum of the momenta of the quarks. The structure functions $F_{1,2}$ can then be written, using the Callan-Gross relation

$$2xF_1(x) = F_2(x) = \sum_i e_i^2 x f_i(x), \tag{1.29}$$

in terms of x, which we interpret as the momentum fraction carried by a given parton, and the PDFs $f_i(x)$, which we interpret as describing the probability that the struck parton i carries a fraction x of the proton's total momentum. The Callan-Gross relation is a result of the spin-$\frac{1}{2}$

nature of the quarks and is extremely well supported by experimental data.

As stated earlier, the PDFs are process-independent and only need to be obtained once. This is crucial because the PDFs are a necessary input to the calculation of *any* cross-section involving hadrons in the initial state. As long as the PDFs are known at the same scale as the factorization scale of a given process, they can be used for a two-hadron process like Equation 1.26 or for the *ep* scattering process we've been discussing (Figure 1.10) [13]. However, the PDFs are also non-perturbative, which means that they cannot be calculated using perturbation theory. Whereas the hard scattering process $\hat{\sigma}$ treats the initial state particles as free, elastically-scattering quarks, which puts the calculation in the asymptotically-free regime of perturbative QCD, the PDFs describe all of the quarks and their momenta in reference to the bound state of the proton. This includes not only the three valence quarks but also the sea quarks produced via gluon radiation, which are by definition not asyptotically free, putting the PDFs beyond the scope of perturbative QCD.

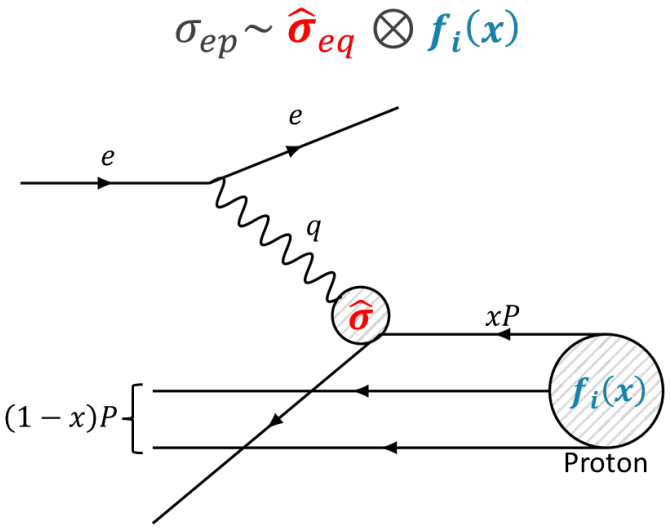

Figure 1.10: The parton-model description of a hard-scattering process, factorized into the perturbative PDFs $f_i(x)$ and the non-perturbative hard process $\hat{\sigma}$.

This raises the question of how the PDFs account for the gluons which are present within the proton. It has been experimentally demonstrated that the quarks in the proton only account for around half of the proton's momentum; this means that the other half must be carried by the gluons, which the PDFs so far do not mention. But accounting for the gluons conflicts with the assumption of Bjorken scaling: as Q^2 increases, more gluon emission is resolved by the probe, and so the probe particle sees the quarks as having lower momentum fraction x (see Figure 1.11). Therefore, when gluon brehmsstrahlung is also accounted for, the Bjorken scaling of the PDFs is violated. This behavior was missed by the earlier scaling experiments, such as those in Figure 1.8, because they did not probe high enough in Q^2 or low enough in x. Fortunately, in the next Section, we will see how we can not only incorporate the lack of Bjorken scaling into the parton model but also use a PDF determined at one energy scale for a prediction at another.

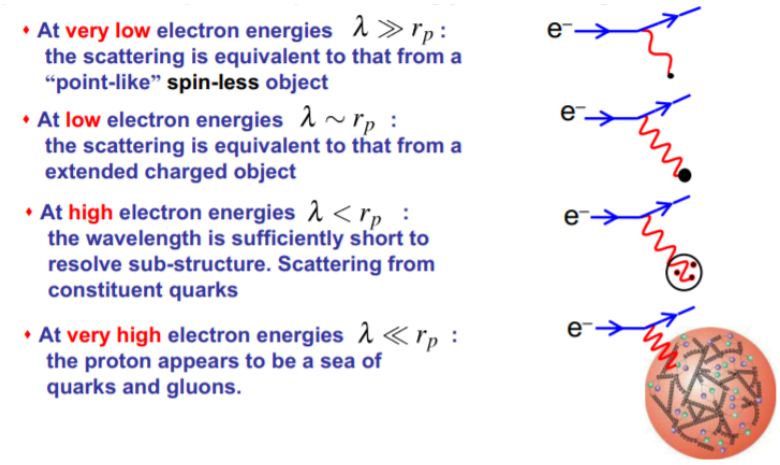

Figure 1.11: A visualization depicting how the increase in momentum transfer Q^2 leads to increased resolution of the proton, and thus the violation of Bjorken scaling [12].

1.4.3 PDF evolution and determination

To account for the possibility that a parton a of a given momentum fraction x may have come from a parent parton b with some larger momentum fraction y, which could either have been a quark or gluon, we use the Dokshitzer-Gribov-Lipatov-Altarelli-Parisi (DGLAP) equation [21–23]:

$$\frac{d}{d\log Q^2} f_a(x, Q^2) = \frac{\alpha_s}{2\pi} \sum_{b\in\{q,g\}} \int_x^1 \frac{dy}{y} P_{ba}\left(\frac{x}{y}\right) f_b(y, Q^2). \tag{1.30}$$

The Altarelli-Parisi splitting functions P_{ba} describe the probability that parton b splits into parton a, and these splitting functions are used to determine the dependence of the structure functions on Q^2. The DGLAP equations cannot determine the PDFs, but if the PDFs are already known at some reference energy scale Q_o^2 they can then be determined at any other scale Q^2 using the DGLAP equation. The evolution of the structure function F_2 with Q^2 is shown in Figure 1.12. This is somewhat analogous to the renormalization group equation 1.19, but for the factorization of the cross-section.

To obtain the PDFs at the original reference scale, it is assumed that the PDFs are smooth functions of x which depend on a finite number of parameters. There are several parameterizations using different functional forms, including the Martin-Stirling-Thorne-Watt (MSTW) and Coordinated Theoretical Experimental Project on QCD (CTEQ) collaborations [25, 26]. Experimental data on a range of fixed-target and collider processes with well-know cross-sections are used to determine the values of the parameters necessary to specify the PDFs; a sample of PDFs obtained using the MSTW parameterization is given in Figure 1.13.

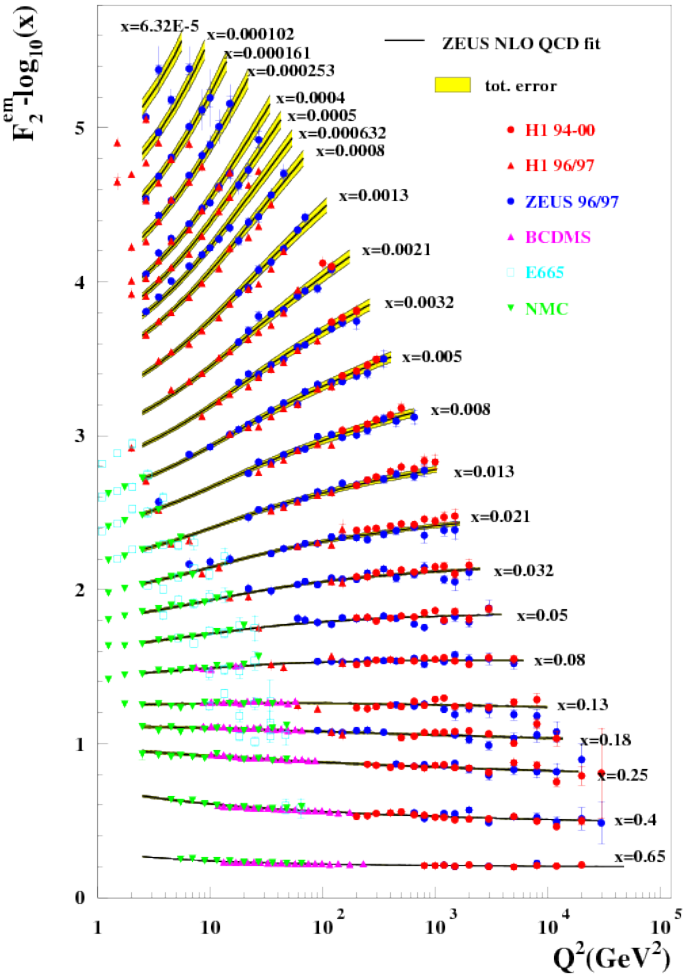

Figure 1.12: Deviations of the F_2 structure function from Bjorken scaling. Data from HERA and fixed-target experiments are compared to the ZEUS NLO fit [24].

1.5 Parton showers

The cross-section in Equation 1.26 is an inclusive cross-section which represents the probability of the process $ab \to X$. Here, X can be a final state accompanied by any number of extraneous particles from initial- and final-state radiation; for example, if we're interested in Z bosons in the final state, the inclusive cross-section will include any number of accompanying jets. But if we want to study the kinematics of the leading jet recoiling against the Z in the final state, we run into a problem. We can use the DGLAP equation to determine the PDFs at different energy scales corresponding to different levels of resolved radiation or lack thereof. But we need something analogous for the hard process scattering cross-section calculation, because it is by definition of Equation 1.26 an inclusive total cross-section calculation; the amount of initial and final-state radiation is not specified. This is where Monte Carlo event generators [28] come in. The

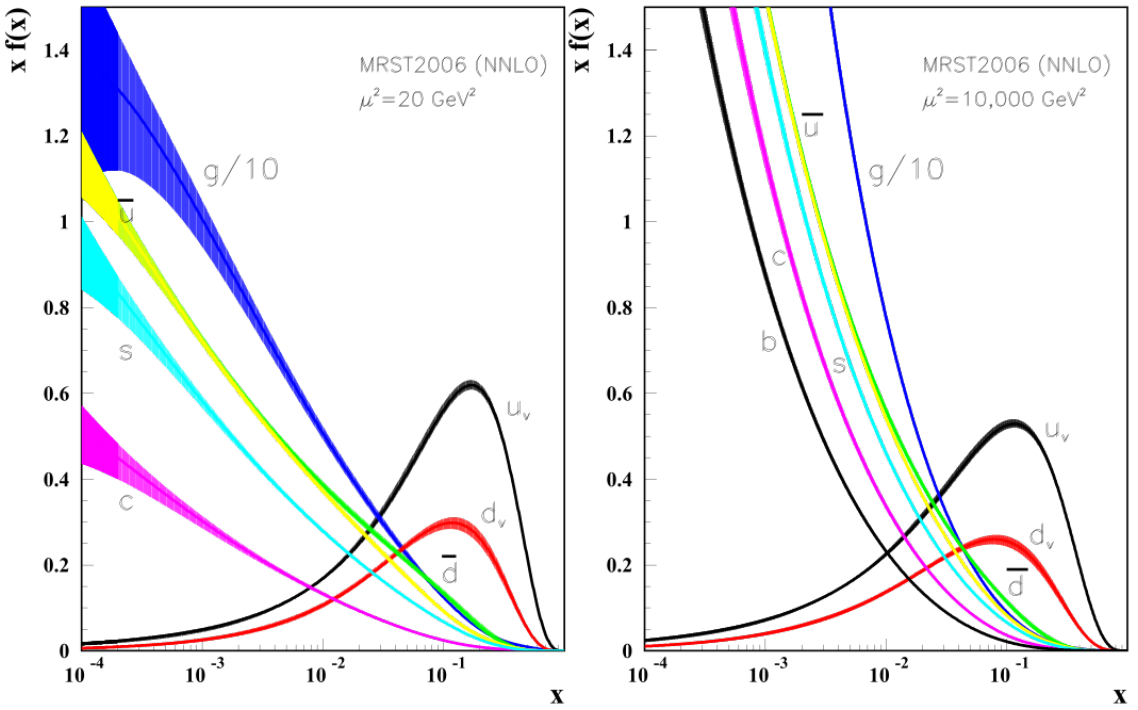

Figure 1.13: Distributions of the parton distribution functions times x using the NNLO MSTW parameterization at $Q^2 = 10\,\mathrm{GeV}^2$ (left) and $Q^2 = 10000\,\mathrm{GeV}^2$ (right) [27].

generators simulate a *parton shower* as described by the Sudakov form factors [29]

$$\Delta_a(t_0, t) \equiv \exp\left[-\frac{\alpha_S}{2\pi}\sum_{b\in q,g}\int_{t_0}^{t}\frac{dt'}{t'}\int_x^1 \frac{dy}{y}\hat{P}_{ab}\left(\frac{x}{y}\right)\right], \tag{1.31}$$

which give the probability that a parton i evolves from time t_0 to time t without emission of a parton j, governed once again by the splitting functions $\hat{P}_{ji}$[1] [13]. These provide a DGLAP-like expression which is convoluted to the scattering process prediction:

$$\frac{d}{d\log Q^2}\log\frac{f_a(x,t)}{\Delta_a(t_c)} = \frac{\alpha_s}{2\pi}\sum_{b\in q,g}\int_x^1\frac{dy}{y}P_{ba}\left(\frac{x}{y}\right)\frac{f_b(y,t)}{f_a(x,t)}. \tag{1.32}$$

The event generators solve this expression in order to account for unresolved splittings. In the final state, highly virtual partons lose energy to splittings until the cutoff scale $Q_c = \frac{1}{t_c}$ is reached, at which point all of the outgoing partons have to be packaged together into hadrons, as described in the following Section. In the initial state, this algorithm is played backwards: the initial-state parton gains virtuality until it reaches the momentum fraction of the PDF.

These parton shower models allow for predictions which account for the various soft and collinear emissions which are not accounted for in fixed-order ME calculation but are in the PDFs

[1] The splitting functions here are unregularized, in that an infrared cutoff is used: emissions with z above the cutoff are deemed unresolvable, so these Sudakov form factors give the probability of evolving without *resolvable* emission.

through DGLAP evolution.[2] However, they do not model hard partons well: the splitting functions correspond to the soft and collinear limit of the full splitting process, and so the parton shower models are not built to adequately model high-energy or large-angle emissions. Therefore, higher-order ME calculations with hard splitting explicitly included are required. But the parton shower models do still model the hard splittings — just not well — and so there is the potential for double-counting these hard splittings with both the parton shower and ME calculations. Because of this, the hard splittings have to be removed from the parton shower modelling with a matching (using a LO ME calculation) or merging (using an NLO or higher-order ME calculation) procedure; the specifics of such procedures are beyond the scope of this thesis.

1.6 Hadronization and underlying events

The next non-perturbative step of the pp collision is the formation of hadrons from the partons available after the parton shower is performed. As the quarks and gluons in the shower separate in the center-of-mass frame, the energy scale keeps dropping until they *hadronize*,[3] i.e. combine to form colorless hadrons. The end result is *jets* of hadrons moving roughly in the direction of the original q and $\bar{q}$. Since this transition is non-perturbative, it must be modelled using experimental data.

Fortunately, we can use a very similar formalism to what was used to describe the PDFs. Recall that the parton distribution functions $f_i(x)$ describe the probability that a parton i carries momentum fraction x of the proton's overall momentum. We can define the hadronization functions $D_q^h(z)$ and $D_{\bar{q}}^h(z)$, which describe the probability that a hadron h is found in the jet of the original quark or antiquark, respectively, carrying a fraction z of its momentum. Just like the PDFs, these functions are subject to momentum conservation, i.e.

$$\sum_i \int_0^1 dx\, x f_i(x) = \sum_h \int_0^1 dz\, z D_q^h(z) = 1, \tag{1.33}$$

so that the sum of the energies of the hadrons is equal to the energy of the parent quarks. The hadronization functions are also, like the PDFs, subject to scaling violations due to gluon emission and can be evolved with a DGLAP evolution equation.

The hadronization of the partons must be modelled based on the observation that the parton momenta and flavors carry over to the resulting hadrons (so-called parton-hadron duality) [13, 27]. There are multiple phenomenological models to describe this carry-over. The two most commonly-

[2] Analytical calculations are possible as well, but most theoretical predictions for the LHC proceed use the generator-based predictions described in this Section.

[3] The term *fragmentation* is usually used in conjunction with hadronization, often interchangeably; for example, the functions $D_q^h(Z)$ which I define in this Section are usually referred to as the "fragmentation functions." To avoid confusion, however, I am going to stick with "hadronization," as I find that term more intuitive.

used models are: the string hadronization model, in which separating quarks lose energy to the color field, which collapses into a string and then breaks through $q\bar{q}$ production to form hadrons (used by the Pythia event generator); and the cluster hadronization model, in which color-singlet clusters of partons form through non-perturbative $g \to q\bar{q}$ splitting. A comparison of these models is shown in Figure 1.14; a further discussion of the Monte Carlo generators used in this analysis can be found in Chapter 5.

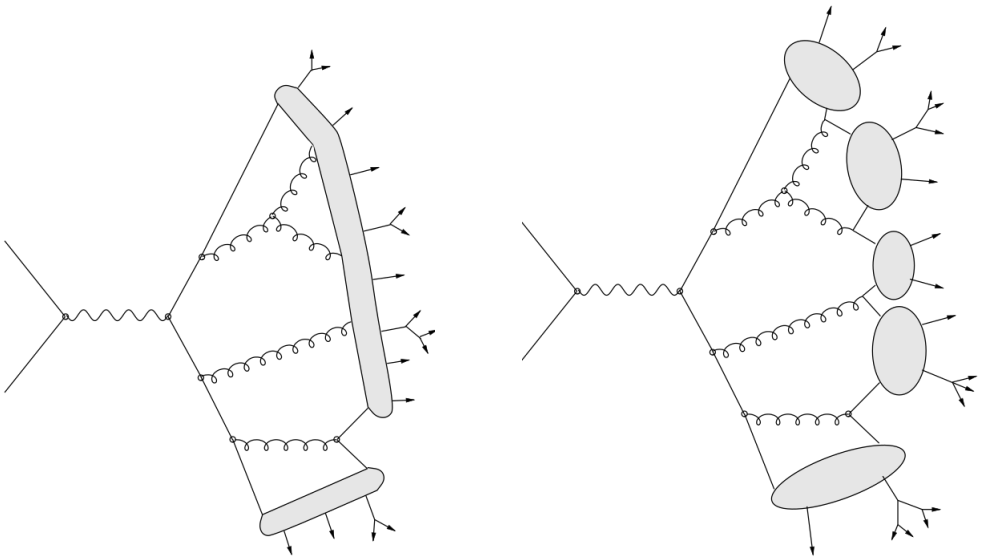

(a) String hadronization model (b) Cluster hadronization models

Figure 1.14: Two different models of parton shower hadronization [13].

The last non-perturbative portion of the pp collisions are the potential interactions of the partons not involved in the main scattering process. The "underlying event" can involve other, separate collisions as well as gluon strings between the main scattering partons and the spectator quarks and gluons. All of this, plus everything described in this chapter, make up the "minimum bias" events recorded by the loosest triggers of the detector's trigger system. These non-perturbative portions of the pp collisions, when combined with the pQCD predictions for the main scattering interaction, provide the full picture of a collision event at the LHC (Fig 1.15). In the next Chapter, an overview of the LHC itself and the detectors used to observe these events is provided.

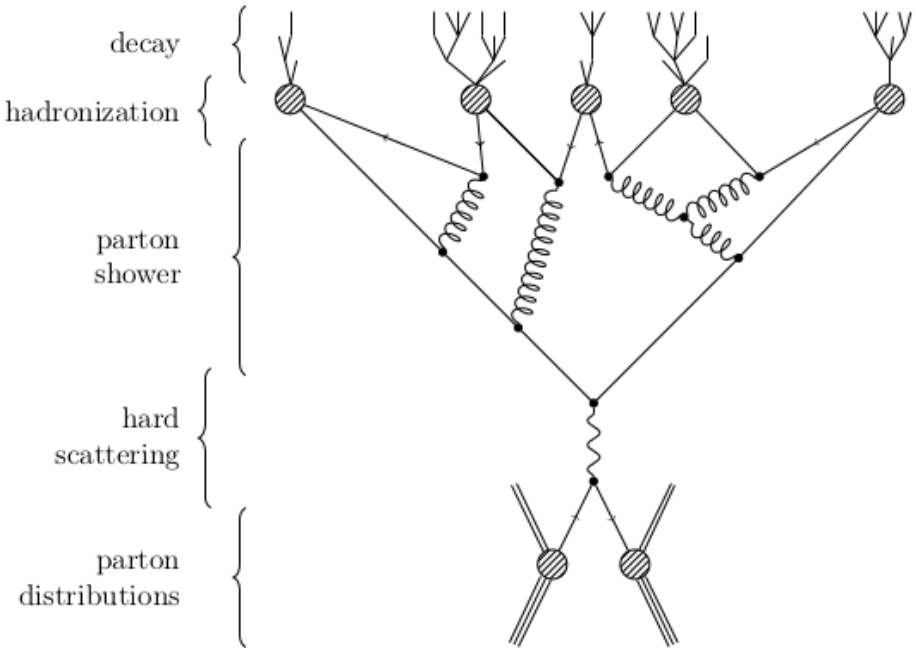

Figure 1.15: A breakdown of the various perturbative and non-perturbative portions of the pp cross-section [30].

Chapter 2

The Large Hadron Collider and the ATLAS detector

2.1 The Large Hadron Collider

The Large Hadron Collider (LHC), commissioned in 2008 as the final stage of the CERN accelerator complex (Figure 2.1), is located underneath the border of France and Switzerland. With its 27 km circumference, it is the largest particle collider in the world, and in fact the largest machine ever built. Although the LHC does spend a few weeks every year colliding heavy ions such as lead (at center-of-mass energies up to $\sqrt{s} = 5$ TeV), it is primarily a proton-proton collider. The LHC has had two runs of operation: during Run 1, from 2011 to 2012, the equivalent of 36 fb^{-1} of data was collected with proton-proton (pp) collisions reaching a $\sqrt{s} = 7$ to 8 TeV; during Run 2, from 2015 to 2018, 139 fb^{-1} of pp data was collected with $\sqrt{s} = 13$ TeV.

The LHC's two counter-circulating beams are made to collide at four interaction points, at which the four main detectors are placed: A Toroidal LHC ApparatuS (ATLAS) [32] and the Compact Muon Solenoid (CMS) [33], which are the two general-purpose detectors; the Large Hadron Collider beauty (LHCb) [34] detector, which focuses on studying the b-quark to investigate matter-antimatter asymmetry; and A Large Ion Collider Experiment (ALICE) [35], which studies heavy-ion collisions.

2.1.1 Proton synchrotrons

The LHC is designed as a charged hadron synchrotron. In a synchrotron, particles are accelerated using radio frequency (RF) cavities, which produce an oscillating electric potential. The oscillation of this potential means that the accelerating force the charged particles receive as they pass through

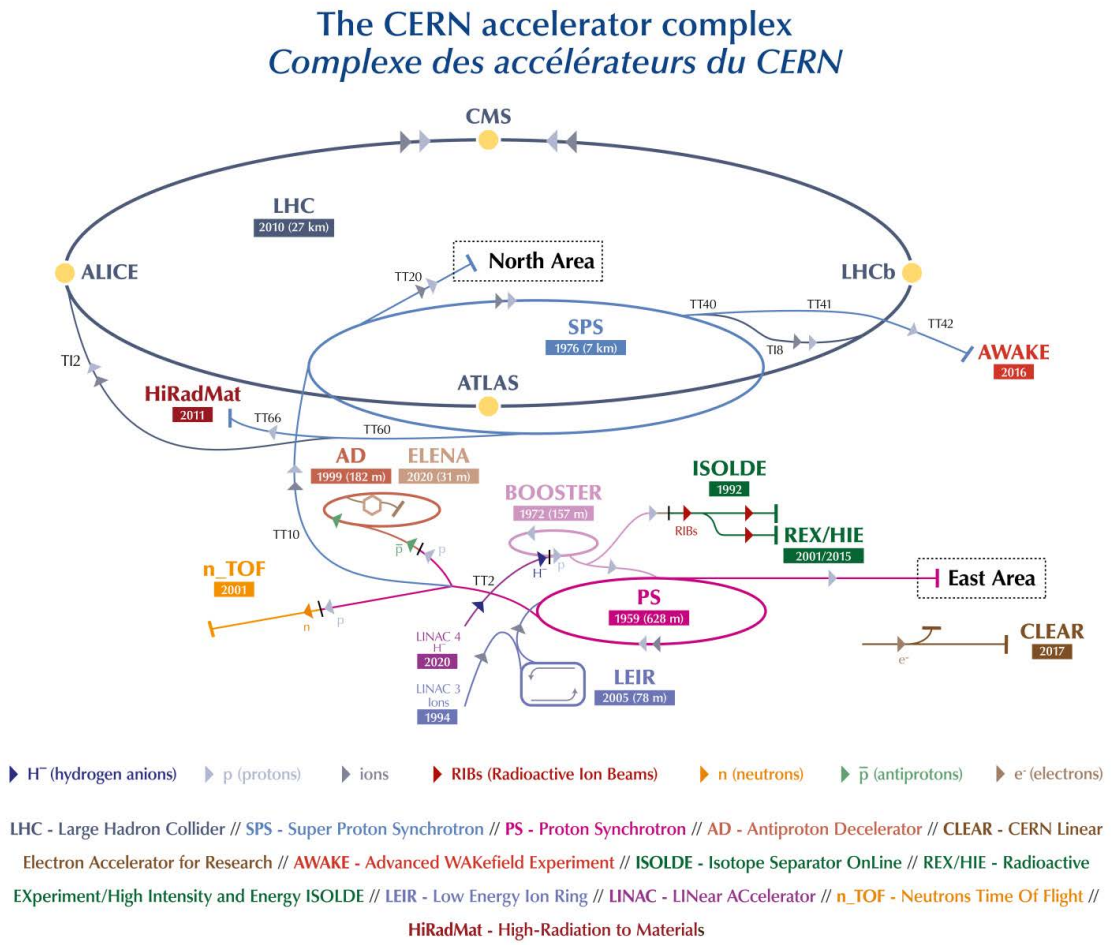

Figure 2.1: The CERN accelerator complex, culminating in the Large Hadron Collider (LHC) [31].

the RF cavities depends on the time of their arrival: particles which arrive at the cavity too "early" see a smaller potential and are boosted less, and those which arrive "late" see a larger potential and receive a larger boost (see Figure 2.2). In this way, the particle density peaks around the phase of the average field strength. Dipole magnets are used to bend the paths of the particles into a ring, and quadrupole magnets near the interaction points focus the particle in the plane transverse to the beam pipes to maximize the probability of collisions at these points. Higher-multipole magnets are used to correct imperfections in the magnetic field.

As the particles move faster, the frequency of the RF cavity potential must increase to boost the particles as they pass. To keep the radius of the particles' paths constant, the magnetic field strength of the dipole magnets must also increase. Thus, as the name "*synchro*tron" suggests, the frequency of the RF potential and the strength of the dipole magnetic field are synchronized to ensure that the particles continue to accelerate while staying along the accelerator's central beamline.

At the LHC, in order to collide protons, two separate beams of protons circulate in opposite

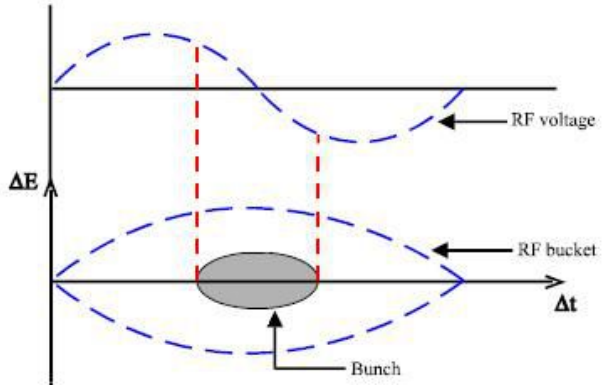

Figure 2.2: Diagram illustrating how the varying RF voltage applies accelerating potentials [36]. Protons arriving "late," with too little energy to match the RF frequency (left red dashed line), experience a larger accelerating potential, while protons arriving "early," with too much energy (right red dashed line), experiencing a smaller accelerating potential.

directions. Because the two beams circulate particles of the same charge, the orientation of the RF cavities and of the dipole magnetic field of each beam is opposite the other beam. To accelerate the protons, each beam is equipped with 8 RF cavities, which have the frequency of their potentials increased to a maximum of 400 MHz; the peak value of the potential is 2 MV. In concert with the increasing frequency of the RF cavities and speed of the particles, a total of 1232 superconducting dipole magnets, each producing a magnetic field of a maximum of 8.3 T, curve particles moving along the beamlines. Finally, 392 quadrupole magnets, with a nominal (maximum) field gradient of 223 T/m (241 T/m), focus each beam as it approaches the interaction points. The dipole and quadrupole magnets are superconducting, which requires them to be held at a temperature of 1.9 K; this is done using a liquid helium distribution system, which sends refridgerated liquid helium in a pipe parallel to the LHC pipe [32, 36, 37].

2.1.2 Energy

The protons accelerated by the LHC start out as negative hydrogen ions, which are initially accelerated by Linear Accelerator 4 (LINAC 4).[1] The hydrogen ions are stripped of their electrons during injection into the Proton Synchrotron Booster (PSB, "Booster" in Figure 2.1), which accelerates then the protons from 160 MeV to 1.4 GeV. The protons are then injected into the Proton Synchrotron (PS); at 25 GeV, they enter the Super Proton Synchrotron (SPS); and they are finally injected into the LHC at 450 GeV. The first collisions which were achieved at the LHC, in November 2009, occured at a center-of-mass energy of $\sqrt{s} = 900$ GeV; in March 2010, $\sqrt{s} = 7$ TeV was reached, with 8 TeV becoming possible in 2012. For the LHC's Run 2, the operational period between April 2015 and the end of 2018, the protons were accelerated to 6.5 TeV, resulting in $\sqrt{s} = 13$ TeV collisions; the data from this Run are what will be discussed in this thesis [37].

[1] LINAC 4 replaced Run 2's LINAC 2 during the 2019-2020 long shutdown.

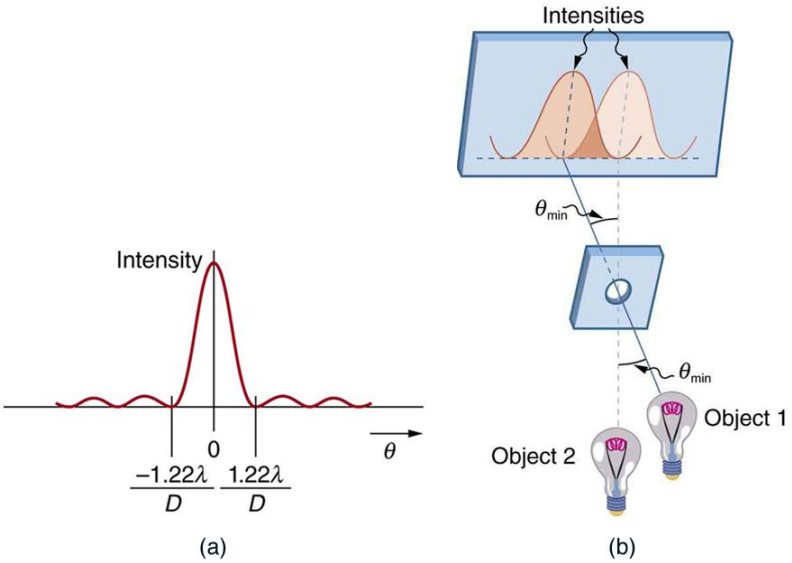

Figure 2.3: (a) Diffraction pattern of a circular aperture. (b) Limit of resolvability of two images: the center of the diffraction pattern of one image is over the minimum of the diffraction pattern of the other [38].

The reason to invest so heavily in the final center-of-mass energy of the proton collisions is that this energy determines the resolution at which the LHC can view the underlying structure and interactions within the proton. de Broglie's inverse relationship between particle momentum and wavelength,

$$\lambda_{\mathrm{dB}} = \frac{\hbar}{p}, \tag{2.1}$$

tells us that increasing the momentum of a particle decreases its wavelength. Moreover, the Rayleigh criterion

$$\sin\left(\theta_{\mathrm{R}}\right) = \frac{\lambda}{D} \tag{2.2}$$

tells us that for a given aperture of diameter D, the resolution θ_R is proportional to the wavelength; this means that a smaller wavelength allows for points separated by a smaller angular distance to be distinguished (see Figure 2.3). Combining these two statements tells us that the faster the protons collide together, the smaller the objects which can be resolved by the detectors (resolution $\sin\left(\theta_{\mathrm{R}}\right) \sim 1/\text{momentum } p$). In this sense, the LHC and its detectors are simply a (very large) proton microscope, with the protons acting as the light producing visible images and the detector acting as the viewing lens.

As a further consequence of the increased energy, rare processes involving heavy particles or very energetic final states are more easily resolvable: because of the Heisenberg uncertainty principle,

more energetic particles decay more quickly, or equivalently over a shorter distance traveled, but the high energies of the probing protons at the LHC can resolve these decays. As a result, as shown in Figure 2.4, higher energies result in larger cross-sections for massive particles: for example, the $p\bar{p}$ collisions at the Tevatron at 1.8 TeV led to the discovery of the top quark, and the LHC quickly exceeded the Tevatron's lifetime top quark production by virtue of its higher center-of-mass energy.

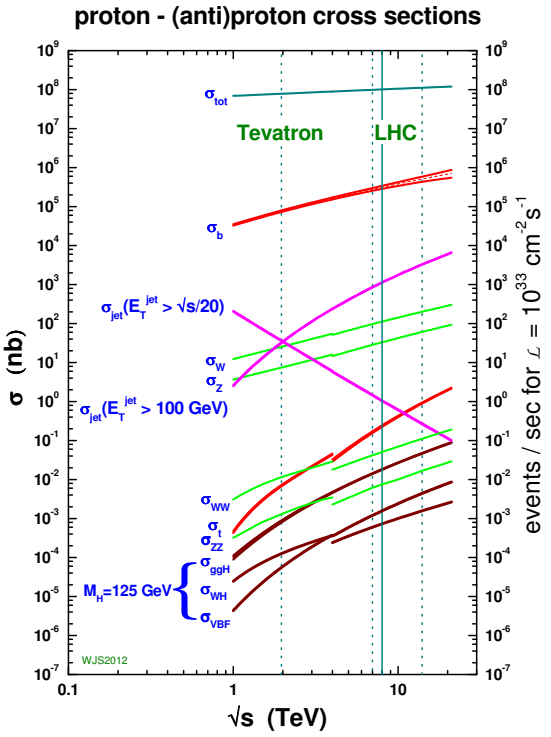

Figure 2.4: Production cross-sections for several processes as a function of center-of-mass energy [39].

2.1.3 Luminosity

To draw any meaningful conclusions about what the detectors are seeing, no matter how energetic the collisions are, requires studying many of these collisions to achieve statistically-significant results. The rate of collisions is closely tied to the luminosity, the number of protons passing through a given area of the detector per unit time. (This is why the LHC is not a proton-*anti*proton collider: it is much more difficult to produce a large number of antiprotons, and so the luminosity of such a collider would be far smaller).

The primary means of increasing the luminosity is keeping the protons tightly packed in bunches so that the number of particles per unit area is high. Bunches of around 1.2×10^{11} protons are formed at injection into the PS, which the RF cavities of the LHC then focus; in the LHC's nominal

operation mode, 2808 such bunches are present in each beam of the LHC. As mentioned above, the oscillating potentials produced by the RF cavities of the LHC squeeze the particles together longitudinally (along the beamline). Any defects in the positions of particles transverse to beam axis are corrected using stochastic cooling: particles straying from the bunch have their positions measured and a corresponding "kick" is given to the bunch, and over many revolutions these corrections normalize the particle spread and contain the bunch. In the LHC, these techniques produce bunches of approximately 10^{11} particles each. The bunches are spaced such that they arrive at the ATLAS detector every 25 ns, corresponding to a bunch-crossing frequency of 40 MHz. A smaller bunch spacing would increase the rate of bunch crossings and the luminosity, but it would also decrease the time available to the detectors to process signals; because of this, the spacing is not decreased further. These bunches are further laterally condensed using the quadrupole magnets, whose field lines direct the protons towards the center of the beam pipe.[2] While the size of the bunches is not constant as the bunches travel through the LHC, it reaches its minimum at the detectors, where powerful insertion magnets at each opening of the detector squeeze the bunches down to 20 μm across.

The instantaneous luminosity is defined as

$$\mathcal{L} = \frac{f N_b^2}{4\pi\sigma_x\sigma_y}, \tag{2.3}$$

where f is the frequency of bunch crossings, N is the number of protons per bunch, and $\sigma_{x,y}$ are the size of a bunch at the interaction point in both directions transverse to the beam direction, about 16 microns for both beams ($\sigma_{x,y}$ can be though of as the size D of the aperture in Equation 2.2 of the LHC "microscope") [36]. Given the values quoted above, we can compute an instantaneous luminosity of 10^{34} cm^{-2}s^{-1}, or 10^{34} bunch crossings per centimeter squared per second. This is not the number of proton collisions; on average, we can expect that the $2 * 10^{11}$ protons involved in a single bunch crossing will result in 40 interactions to study. Further still, only some of these interactions are of interest to a given researcher at a given time, so the actual number of events of interest,

$$N = \sigma L, \tag{2.4}$$

where $L = \int \mathcal{L} dt$, depends on the particular cross-section σ of the process that is to be studied. Increasing the luminosity necessarily increases the expected number of events of any process, producing more data to analyze.

[2] The planned High-Luminosity LHC program requires, among other upgrades, 12 quadrupole magnets with stronger magnetic fields to create even tighter bunches.

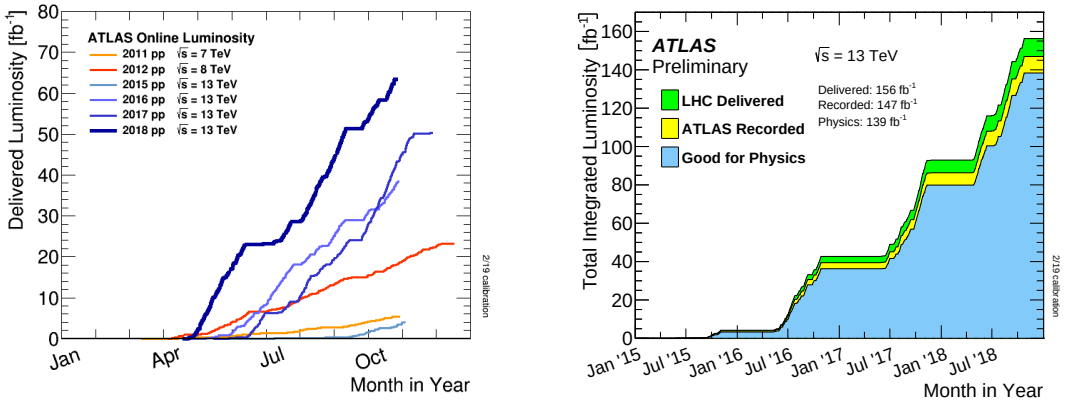

(a) (b)

Figure 2.5: (a) Cumulative luminosity of high-energy proton collisions delivered by the LHC to ATLAS across Runs 1 and 2. (b) Cumulative luminosity delivered by the LHC, during stable beam conditions at $\sqrt{s} = 13$ TeV, to ATLAS (green), recorded by ATLAS (yellow), and certified to be good quality data (blue) from 2015 to 2018 [40].

The integrated luminosity $\int \mathcal{L}\, dt$ provides a measure of the total data collection of the LHC. It is usually expressed in inverse femtobarn, where 1 femtobarn $= 10^{-39}$ cm^2 is the unit used to express cross-sections. The ATLAS experiment, described in Section 2.2, has received a total integrated luminosity of 189.3 fb^{-1} across Runs 1 and 2 (Figure 2.5a). During Run 2 specifically, 156 fb^{-1} was delivered, 139 fb^{-1} of which was recorded by ATLAS during stable beam conditions (Figure 2.5b) [41]. With that tremendous amount of data, ATLAS and the other detectors of the LHC have continued to experimentally verify the success of the Standard Model: as shown in Figure 2.6, ATLAS has verified SM-predicted cross-sections over more than 10 orders of magnitude, making the Standard Model one of the most successful scientific theories ever developed. In addition, ATLAS has also placed severe constraints on a wide array of beyond the SM predictions; see Figure 2.7.

2.2 The ATLAS detector

As mentioned above, ATLAS is one of the two general-purpose detectors of the Large Hadron Collider. As the measurements in this thesis are performed on ATLAS data, I will focus on describing the structure of the ATLAS detector here.

2.2.1 Coordinate systems

The ATLAS detector (Figure 2.8, left) is a cylindrical toroid with forward-backward symmetry with respect to the interaction point where the protons bunches cross. It is 44 m long and 25 m in diameter, with a total weight of approximately 7000 tons. The toroidal shape allows for an almost complete 4π coverage of the area surrounding the interaction point.

To discuss particle positions and momenta in this space, we use a right-handed Cartesian

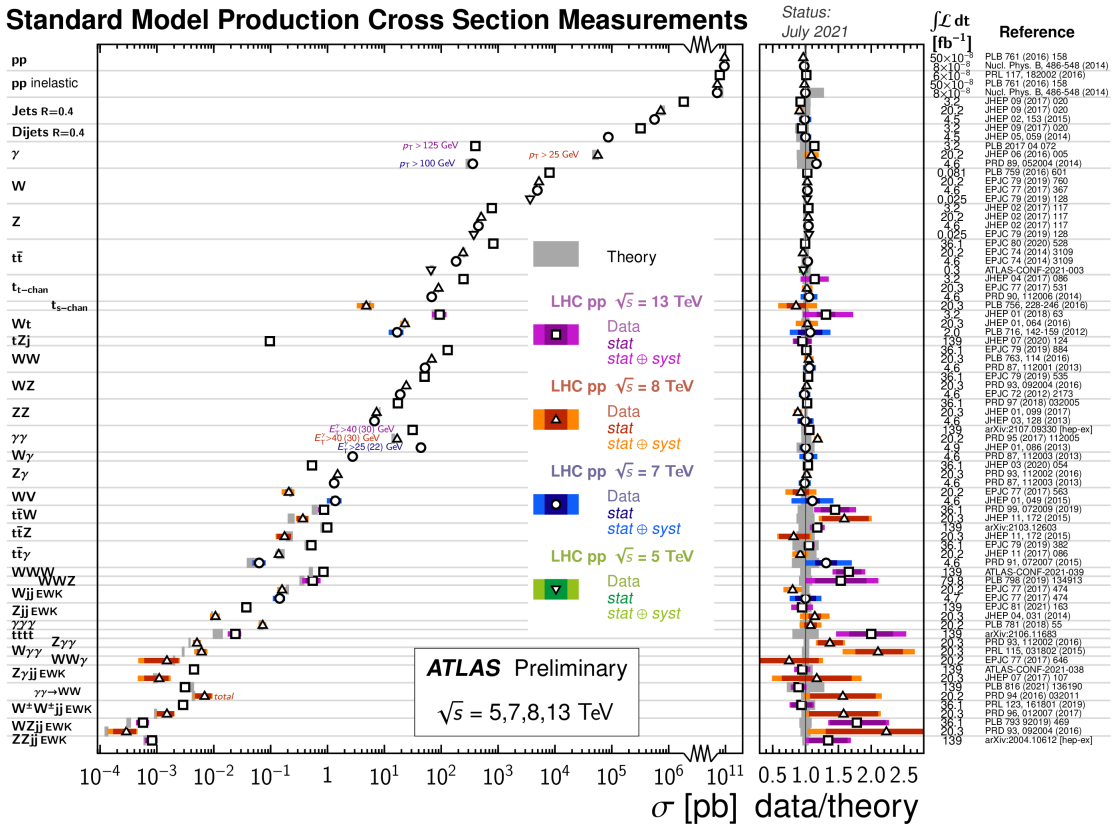

Figure 2.6: Standard Model measurements performed by the ATLAS collaboration using proton-proton collisions with center-of-mass energies of 5, 7, 8, an 13 TeV [42].

coordinate system where the z-axis lies along the beamline, the y-axis points upwards, and the x-axis points towards the center of the LHC ring. A polar coordinate system can also be used, with r extending from the interaction point perpendicular to the z-axis and the ϕ angle, rotating about the z-axis, measuring inclination from the x-axis. A θ angle, which instead measures inclination from the z-axis, is also often used. The coordinate systems are shown on the rightmost panels in Figure 2.8.

Collisions occur in the center-of-mass reference frame of the protons, but not necessarily of the partons which actually participate in the main scattering processes; therefore, the z-momenta of the incoming particles are not necessarily balanced. It is more convenient to perform calculations in the parton center-of-mass frame, where these momenta are balanced, and to then Lorentz boost into the laboratory frame. Doing so requires Lorentz-invariant observable quantities. Two such commonly-used quantities are the transverse momentum

$$p_{\mathrm{T}} = \sqrt{p_x^2 + p_y^2} \tag{2.5}$$

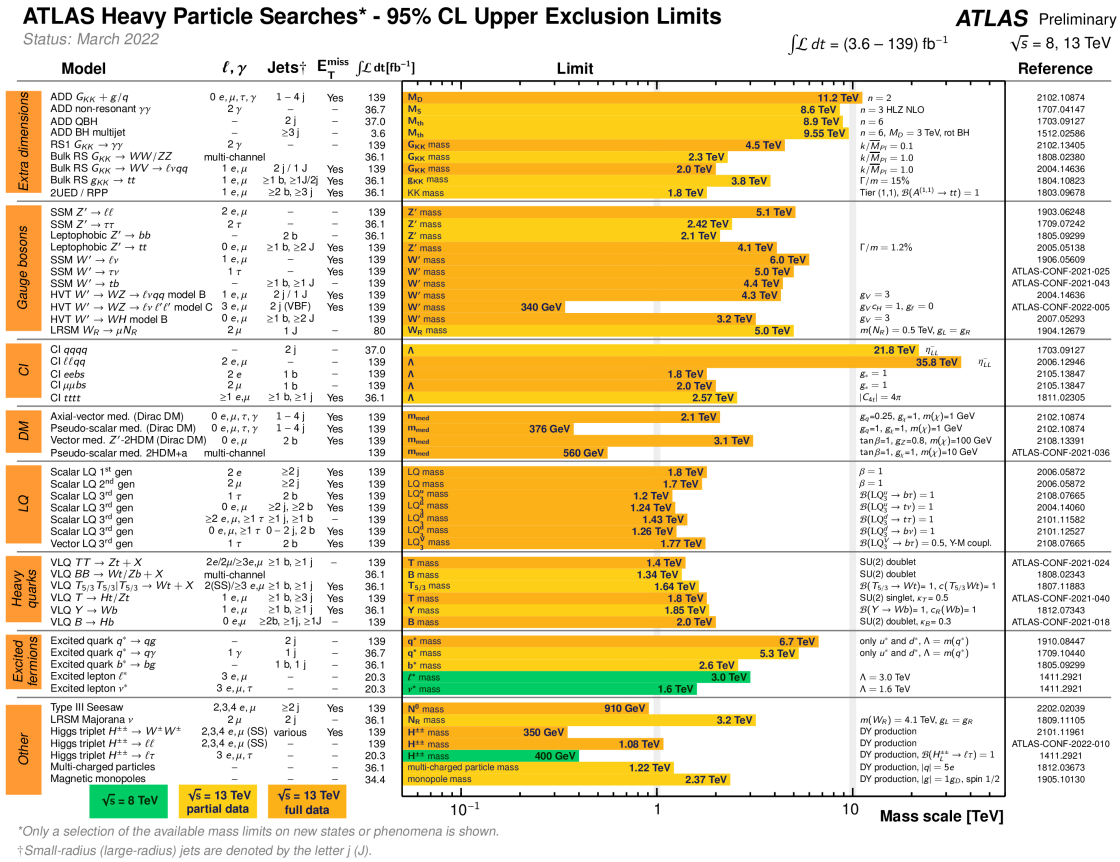

Figure 2.7: Ranges of new particle masses excluded by ATLAS searches at the 95% confidence level [43].

and the pseudorapidity

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

which is used instead of θ because, for highly-relativistic massive particles, it is a good approximation of the Lorentz-invariant rapidity $y = \frac{1}{2}\ln\left[(E + p_z)\,/\,(E - p_z)\right]$. As η and ϕ (which is also Lorentz-invariant) are most commonly used to describe the trajectory of a particle through the ATLAS detector, another useful quantity is the angular distance ΔR between objects:

$$\Delta R = \sqrt{\Delta\eta^2 + \Delta\phi^2}. \tag{2.7}$$

The ATLAS detector is composed of concentric subdetectors which rely on different technologies and particle interaction principles to collect as much precise kinematic and identification information as possible of the particles coming from the collisions. They are described in the following subsections, moving from the center of the detector outwards.

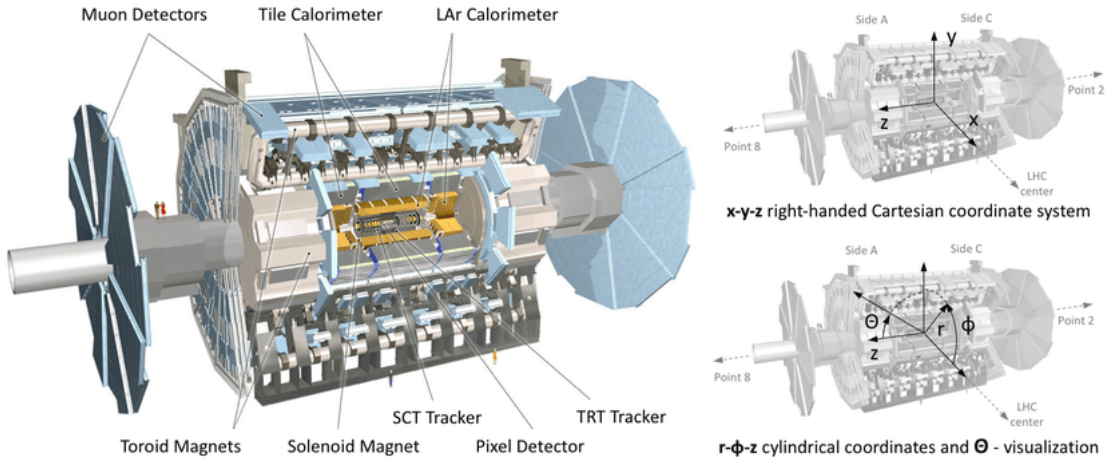

Figure 2.8: Cut-away view of the ATLAS detector (left), with right-handed Cartesian coordinate system (right, top) and polar coordinate system (right, bottom) illustrated [44].

2.2.2 Inner detector

Particles exiting a proton-proton collision at the interaction point first encounter the beryllium beam pipe [45]; the vast majority do not interact with the pipe, as it is meant to be an inactive part of the detector. The first active part of the detector most particles interact with is therefore the **Inner Detector (ID)** [46], whose purpose is to measure the momenta of charged particles. A cylindrical detector covering the range $\eta \leq 2.5$, it consists of three subcomponents (Figure 2.9). The pixel detector is the first and highest-granularity detector and is made up of three concentric barrel layers, and two endcaps of three layers each, of rectangular silicon "pixels". When charged particles pass through the pixels, ions are freed and read out by 80.36 million electronic readout channels. The pixel detector has an intrinsic accuracy of 10 μm in the radial ($R-\phi$) plane and 115 μm in the axial (z, in the barrel region) and radial (R, in the endcap regions) planes. In Run 2, the Insertable B-Layer (IBL) [47], made up of additional pixel sensors (12.04 million additional readout channels), was added between the pixel detector and the beam line to improve the resolution of the charged particle origins, which are especially important for discerning b-hadrons (hence "B-layer"). Beyond the pixel detector are the 4 layers of the Semiconductor Tracker (SCT) [48, 49], which is made up of strips of silicon. The detection principle is the same as that of the pixel detector, but the strips sacrifice some resolution to reduce the overall cost of the detector: the SCT has an intrinsic accuracy of 17 μm in the radial ($R-\phi$) plane and 580 μm in the axial and radial planes. Finally, the Transition Radiation Tracker (TRT) [50, 51] surrounds the SCT with long, thin drift tubes filled with ionizable gas. A charged particle of $\eta < 2$ typically passes through 35-40 tubes. Charged particles passing through the tubes free ions which drift towards an anode wire in the center of each tube, resulting in a readout current; the energy deposited by the drifting ions help distinguish electrons from radiative photons, which result in much larger energy depositions.

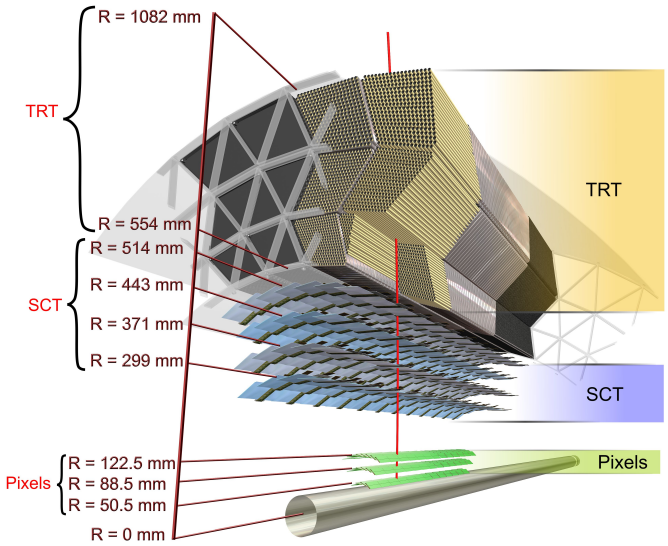

Figure 2.9: Schematic view of the ATLAS inner detector (ID). The ID consists of silicon pixels and strips on the interior and drift tubes in the outermost layers [52].

Enclosing the ID is the solenoid magnet [53], which produces a 2T magnetic field. Charged particles have their trajectories deflected by this magnetic field proportional to their transverse momentum p_{T}. This p_{T} can therefore be determined by using the spatial hits in the inner detector to reconstruct charged particle tracks and then measuring the impact parameter of the tracks. For tracks with $p_{\mathrm{T}} > 30$ GeV, the ID was measured at commissioning to have a transverse impact parameter resolution of 22.1 ± 0.9 μm and a relative momentum resolution of $\sigma_p/p = (4.83 \pm 0.16) \times 10^{-4}$ GeV$^{-1} \times p_{\mathrm{T}}$. The uncertainty in this resolution,

$$\sigma(1/p_{\mathrm{T}}) = A_{p_{\mathrm{T}}} \oplus \frac{B_{p_{\mathrm{T}}}}{p_{\mathrm{T}}\sqrt{\sin\theta}}, \tag{2.8}$$

is dominated by multiple scattering error ($B_{p_{\mathrm{T}}}$) at low p_{T} and by the intrinsic p_{T} resolution and residual misalignment errors ($A_{p_{\mathrm{T}}}$) at high p_{T} [54].

2.2.3 Calorimeters

Beyond the solenoid and its magnetic field are the **Electromagnetic and Hadronic Calorimeters** [55, 56], in that order. The purpose of the calorimeters, which are shown in Figure 2.10, is to measure the energy of incoming particles. In a calorimeter, incoming particles first interact with a dense absorption (passive) material. This leads to electromagnetic or hadronic showers of particles, which cause ionization or light production, respectively, in the active components of the calorimeters. The resulting current or light signal, proportional to the energy of the detected particles, is then read out by the connected electronics. The differing active materials used in

the electromagnetic and hadronic calorimeters (which are specified below) result in different electromagnetic radiation lengths X_o and nuclear interaction lengths λ_I for the two calorimeters.[3] As λ_I is generally much larger than X_o for dense materials, hadrons typically pass through the electromagnetic calorimeter before showering in the hadronic calorimeter.

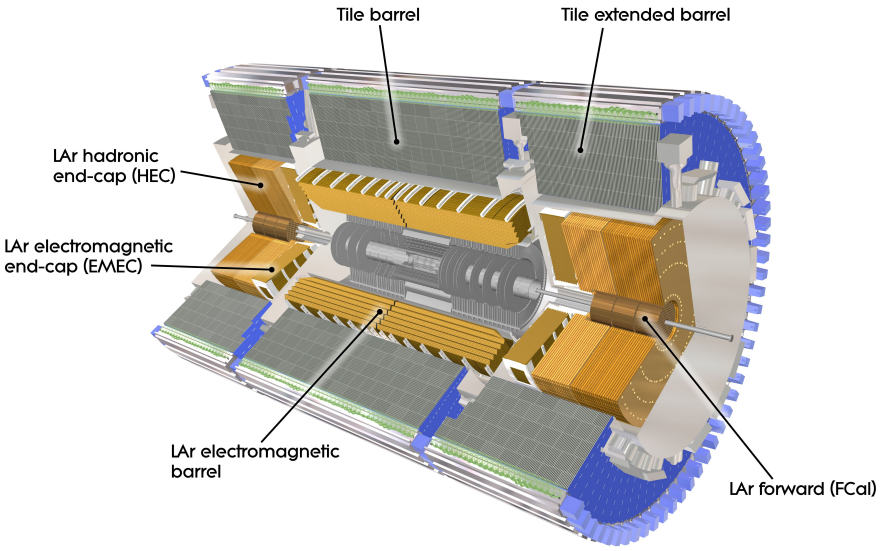

Figure 2.10: Cut-away schematic of the ATLAS electromagnetic and hadronic calorimeters [58].

The electromagnetic calorimeter (ECal) uses an accordion geometry of lead absorptive components and liquid argon (LAr) active components. This geometry of overlapping layers allows for full ϕ coverage, without any cracks, and also for fast readout from the front and back of the electrodes. The calorimeter is divided into two regions: two half-barrels, separated by a small gap at $z = 0$, cover $|\eta| < 1.475$, while two endcaps perpendicular to the beamline cover $1.327 < |\eta| < 3.2$. Energy is determined most precisely in the region $|\eta| < 2.5$, where the calorimeter has three lead-LAr layers; the region $2.5 < |\eta| < 3.2$ is covered by two layers. Figure 2.11 shows a section of a module in the barrel region. The central (second) layer is the longest, so most often the center of the electromagnetic shower, and the majority of the deposited energy, is in this layer. By contrast, the third layer has half the resolution in η, while the first layer has 20 times the resolution in η but a quarter the resolution in ϕ. The electron likelihood identification, described in Section 3.4.2, uses the information obtained from modules such as this.

The energy resolution [59] of the ECal can be parameterized as

$$\frac{\sigma(E)}{E} = \frac{a}{\sqrt{E}} \oplus \frac{b}{E} \oplus c, \tag{2.9}$$

[3] X_o is the mean distance over which a relativistic electron will lose all but $1/e$ of its energy to brehmsstrahlung, while λ_I is the mean distance needed for the flux of relativistic primary hadrons to reduce to a fraction $1/e$ [32, 57].

where a is the stochastic term, b is the electronic noise term, and c is a constant term including detector instabilities and miscalibration. The resolution is designed to be

$$\frac{\sigma(E)}{E} = \frac{10\%}{\sqrt{E}} \oplus \frac{170 \text{ MeV}}{E} \oplus 0.7\%. \tag{2.10}$$

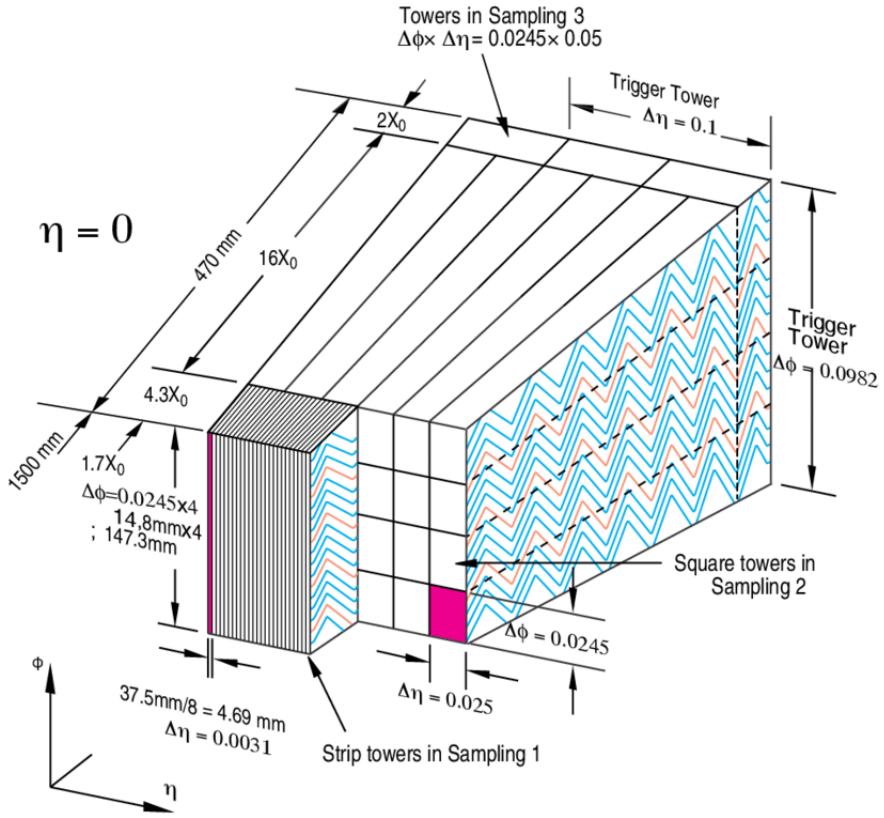

Figure 2.11: Sketch of a barrel module of the electromagnetic calorimeter, with the granularity of the cells in each of the layers given [32].

The hadronic calorimeter (HCal) has three different sections (see Figure 2.10). The tile calorimeter, consisting of 3 layers of steel (absorptive) and scintillator (active) plates, covers the region $|\eta| < 1.7$. The endcap copper-LAr calorimeters (HEC) cover the region $1.327 < |\eta| < 3.2$. Finally, the forward calorimeter (FCal) covers the region $3.1 < |\eta| < 4.9$ and consists of three layers: FCal1 is a copper-LAr calorimeter like the HEC, while Fcal2 and Fcal3 are tungsten-LAr hadronic calorimeters. The tile calorimeter's energy resolution, parameterized as in Equation 2.9 but with the negligible electronic noise term dropped, is measured to be

$$\frac{\sigma(E)}{E} = \frac{52.9\%}{\sqrt{E}} \oplus 5.7\%.$$

Similarly, the design energy resolutions of the HEC and FCal are $\frac{50\%}{\sqrt{E}} \oplus 3\%$ and $\frac{100\%}{\sqrt{E}} \oplus 10\%$, respectively.

2.2.4 Muon spectrometer

The outermost layer of the ATLAS detector is the **Muon Spectrometer (MS)** [60], designed to measure the momenta of the muons which, due to their higher mass, most often make it through the calorimeters while losing minimal energy through interactions with atomic nuclei. The trajectories of the particles in the MS are bent by superconducting toroidal magnets, which are shown in Figure 2.12: the large barrel toroid covers the region $|\eta| < 1.4$, while two endcap toroids cover $1.6 < |\eta| < 2.7$.

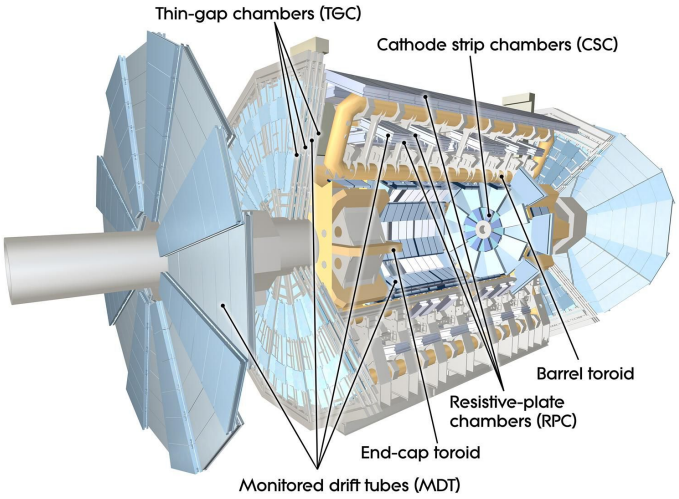

Figure 2.12: Cut-away schematic of the muon spectrometer [61].

Monitored drift tubes (MDTs) and cathode strip chambers (CSCs) are used for measurement and tracking of the particle trajectories. In the region $|\eta| < 2.0$, three layers of the MDTs are used, while in the region $2.0 < |\eta| < 2.7$ one layers of CSCs is used together with two MDT layers (see Figure 2.13). The MDTs work similarly to the TRT of the inner detector, in that charged particles cause ionized electrons to drift towards an anode wire within the chambers. The CSCs are multi-wire proportional chambers, where instead of a single wire concentric to a cylindrical chamber several wires lie between two cathode strips. The spatial resolution of the MDTs during Run 2 was measured to be $81.7 \pm 2.2\mu$m on average, while the momentum resolution is $\approx 3\%$ for $10 < p_{\mathrm{T}} \lesssim 100$ GeV, slowly rising to $\approx 10\%$ for $p_{\mathrm{T}} = 1$ TeV [62, 63]. To reach this precision, each wire's position must be known to within 80 μm; to account for the detector's deformation due to temperature and magnetic field effects, an optical base alignment system is used to monitor the position of the MS modules in real time.

Muons are a common component of signal processes of interest to the ATLAS collaboration, including the $Z \to \mu\mu$ events used for both the trigger efficiency and heavy-flavor jets studies discussed in this thesis. Therefore, the muon trigger system, which determines whether to store events containing muon, is particularly important. The fast muon triggering in the MS is performed

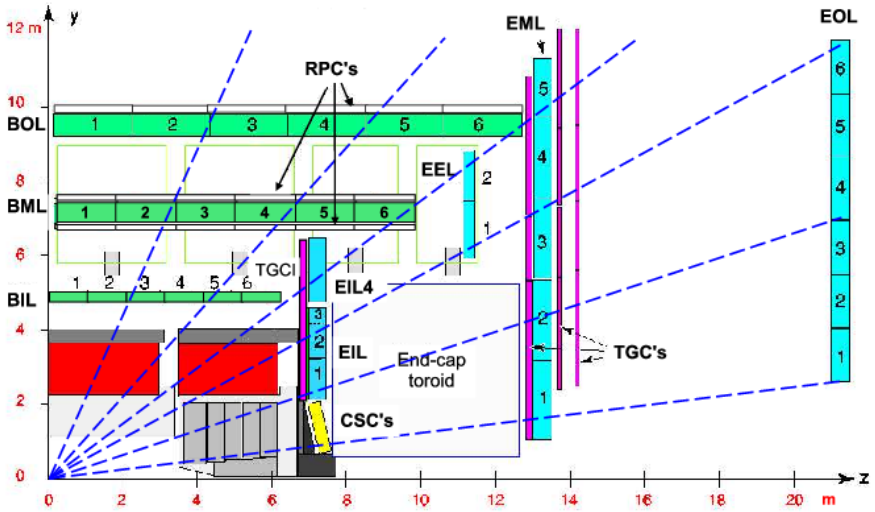

Figure 2.13: Profile schematic of the muon spectrometer. Monitored drift tubes are shown in blue and green [64].

by resistive plate chambers (RPCs) in the barrel region ($|\eta| < 1.05$) and thin-gap chambers (TGCs) in the endcap regions ($10.5 < |\eta| < 2.4$). While the TGCs operate in much the same way as the CSCs, being multiwire proportional chambers as well, the RPCs are parallel electrode-plate detectors, using the electric field between the plates to cause avalanches along the ionization tracks left by the passing particles.

In summary of the ATLAS subdetectors, Table 2.1 gives an overview of their η coverage as well as their p_{T} and E_{T} resolution.

Detector component	**Required resolution**	**$\lvert\eta\rvert$ coverage**	
		Measurement	Trigger
Tracking	$\sigma_{p_{\mathrm{T}}}/p_{\mathrm{T}} = 0.05\%\ p_{\mathrm{T}} \oplus 1\%$	± 2.5	
EM calorimetry	$\sigma_E/E = 10\%/\sqrt{E} \oplus 0.7\%$	± 3.2	± 2.5
Hadronic calorimetry (jets)			
- barrel and endcap	$\sigma_E/E = 50\%/\sqrt{E} \oplus 3\%$	± 3.2	± 3.2
- forward	$\sigma_E/E = 100\%/\sqrt{E} \oplus 10\%$	$3.1 < \lvert\eta\rvert < 4.9$	$3.1 < \lvert\eta\rvert < 4.9$
Muon spectrometer	$\sigma_{p_{\mathrm{T}}}/p_{\mathrm{T}} = 10\%$ at $p_{\mathrm{T}} = 1$ TeV	± 2.7	± 2.4

Table 2.1: Expected coverage and resolution of the ATLAS subdetector systems [32].

2.2.5 Trigger system and data acquisition

During Run 2, proton bunches crossed in the ATLAS detector every 25 ns. Given that the average bunch crossing generates around 1 MB of data, this would mean 40 TB/s of data generated, which is impossible to transmit out of the detector given the readout technology employed. Even if it was possible to read out the information at that rate, there would be the problem of storing more

than two petabytes of new information every hour.

However, not all of the information generated by these bunch crossings is of interest to particle physicists. In fact, most of it is superfluous; if the protons merely glance off of one another, for example, then the collision is not energetic enough to generate final states which contain more massive particles. What matters are the events which contain various useful signatures. These can include events of physics interest, such as those involving a Higgs boson, or events used to perform calibrations or efficiency measurements, such as the $Z \to \mu^+\mu^-$ events used to determine the muon trigger efficiencies (see Section 3.3.5). The ATLAS trigger system [65, 66] is therefore responsible for quickly making decisions on whether or not a given event is one worth storing for further analysis.

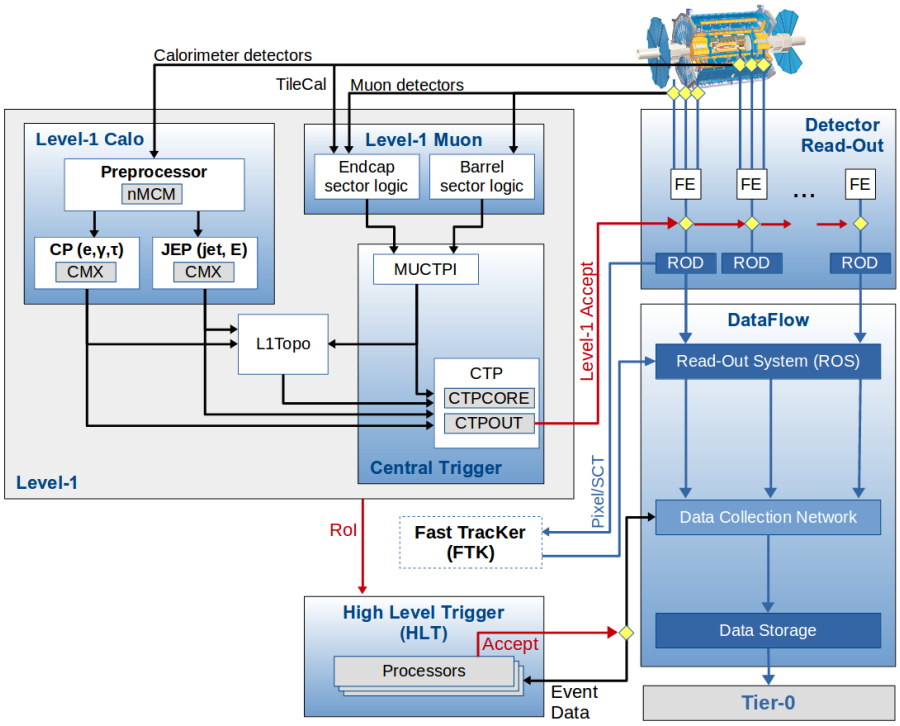

Figure 2.14: Schematic of the ATLAS trigger and data acquisition system for Run 2 [67].

Figure 2.14 gives an overview of the ATLAS trigger and data acquisition system in use during Run 2. Outflowing particles passing through the calorimeters and the muon detectors leave energy deposits which are converted into electical signals. These signals are sent to the hardware-based first-level (L1) trigger [68] while they are simultaneously buffered in the front-end (FE) detector read-outs. The calorimeter (L1Calo) and muon (L1Muon) portions of the L1 pass inputs to the Central Trigger Processor (CTP) which is responsible for all L1 trigger decisions. In addition, prior to Run 2, a topological trigger processor (L1Topo) was added to allow for further selections based on the combined kinematic information from the calorimeter and muon detectors. These selections can reduce backgrounds for the various triggers, saving bandwidth which can then be allocated to

other processes of interest to expand their kinematic range. Overall, the L1 trigger has 2.5 μs to make a decision and reduces the event rate from 40 MHz to 100 kHz.

The events passing the L1 acceptance are passed to the Read-Out System (ROS), where they are again buffered before being processed by the software-based High-Level Trigger (HLT) [69]. The HLT receives Region-of-Interest (RoI) information from the L1 trigger to be used for reconstruction (similar to that done by the offline software) by the HLT trigger algorithms. The processing time allotted to the HLT is 200 ms, and the HLT reduces the event rate further to 1 kHz, corresponding to 1.2 GB/s. Events in the ROS accepted by the HLT are then stored in CERN's Tier-0 computing center for offline reconstruction.

Event selection is controlled through the use of the trigger menu, which configures the list of active trigger chains. The trigger chains, constructed using a L1 trigger and several HLT triggers, are designed to select specific physics signatures (leptons, jets, missing transverse energy, etc.). These trigger chains may or may not have a prescale factor applied; if a prescale factor of value n is applied, a given event has a probability $1/n$ of being selected. These prescale factors are applied to a L1 or to an HLT trigger in a chain and are used to optimize the total resource usage among the physics priorities; for example, during Run 2, the main goal of the trigger menu design was to maintain the unprescaled single-electron and single-muon trigger thresholds around 25 GeV, to ensure the collection of most leptonic W and Z decays.

The data selected for storage by the ATLAS trigger system, in the form of electronic signals, are not yet ready to be used in physics analyses; those signals first have to be translated into physics objects and their characteristics. This process of "reconstruction" is the subject of the next Chapter.

Chapter 3

Offline Reconstruction and Identification

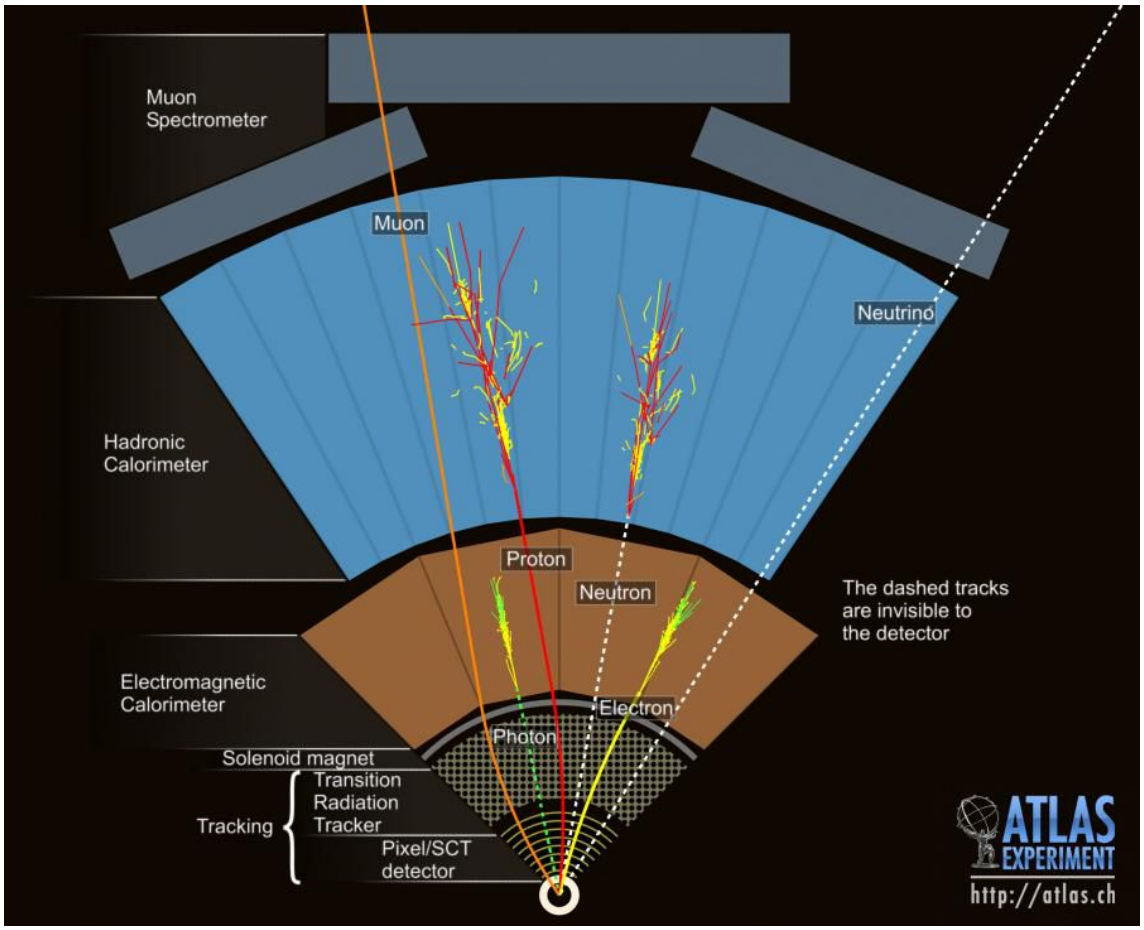

Figure 3.1: An illustration of different types of particles passing through the various layers of the ATLAS detector [70].

The final output of a pp collision in the ATLAS detector, from an experimental perspective, is not the vast array of particles we are interested in, but rather a series of electronic signals

registered by the various components of the detector (Figure 3.1). These signals are then used to *reconstruct* physics objects and their kinematic variables. For example, Figure 3.1 shows that electrons can be reconstructed from energy deposits in the electromagnetic calorimeter matched to tracks left in the inner detector, but neutrons have to be reconstructed solely using energy deposits in the hadronic calorimeter. Neutrinos, which pass completely undetected through ATLAS, have to be reconstructed using missing transverse energy (momentum): the transverse momenta of the outgoing reconstructed particles have to sum to zero in the center-of-mass frame of the collision.

Performance groups within ATLAS work on providing instructions, recommendations, and software for identifying specific objects. The recommendations are designed to strike a balance between identifying as large a percentage of those objects which were present in the event as possible (efficiency) while ensuring that as few background or misidentified objects as can be managed are included (purity). These reconstructions then have to be calibrated to account for detector effects as well as the resolution of the detector: the MC simulated data has to be smeared to match the resolution seen in the measured data. Finally, uncertainties which will propagate through the analyses using these reconstructed objects have to be determined [71].

3.1 Tracks

The momenta and signs of the electric charge of particles can be inferred from the curvature of their trajectories within the magnetic field produced by the solenoid and the toroid. These trajectories are reconstructed as tracks [72, 73] in the ID. Tracks are assembled from three-dimensional measurements of energy deposits in the silicon pixel and SCT, referred to as space-points. Seeds are formed from three space-points in the pixel and in the first layer of the SCT. With a combinatorial Kalman filtering technique [74], additional hits in the outer layers of the SCT are associated with the seeds to form track candidates. Track candidates having $p_{\mathrm{T}} > 400$ MeV and $|\eta| < 2.5$ are then fed to the ambiguity solver, which scores the candidates based on how likely they are to correspond to the correct trajectory of a charged particle. Clusters associated to track candidates increase the scores of those candidates, while holes — measurements on a detector surface expected given the trajectory projection, but not observed — lower their score. Scores are also penalized using the χ^2 of the track fit, determined using a Newton-Raphson method: the χ^2 is calculated using the track-hit residuals $r_i = e_i(\boldsymbol{\tau}) - m_i$, where m_i is the position of the i^{th} hit and $e_i(\boldsymbol{\tau})$ is the position of intersection of the fitted track on the surface on which the i^{th} hit was measured (see Figure 3.2). The track is parameterized using $\boldsymbol{\tau} = (d_0, z_0, \phi_0, \theta_0, q/p)$, where d_0 and z_0 are the transverse and longitudinal impact parameters, respectively, ϕ_0 and θ_0 are polar angles, and q/p is the charge-to-momentum ratio of the particle which formed the track.

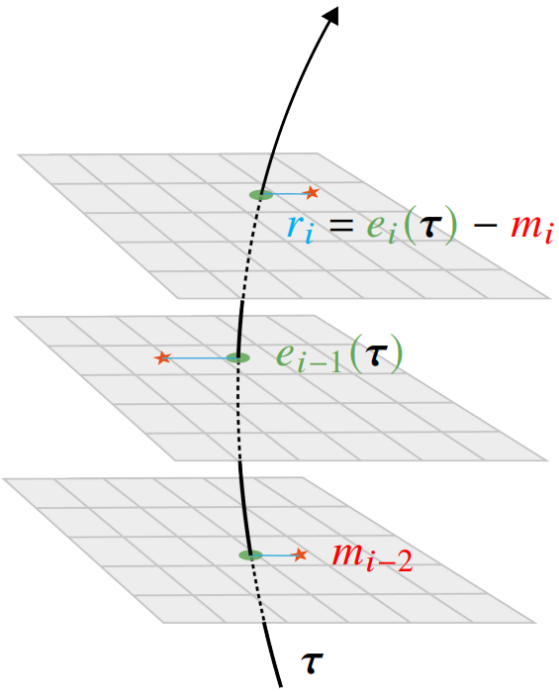

Figure 3.2: Schematic representation of a charged particle passing through planes of the inner detector [75].

Vertices are then defined as points in space where two tracks, or their fitted projections in the ID, meet. Through an iterative process, tracks are weighted according to their compatibility with a seed vertex position, and a fit is made to update the vertex position. At the last iteration, tracks with insufficient weights are discarded, to be used with another vertex. Tracks are associated to vertices in this way until either no tracks remain or there are no more vertices to associate to the remaining tracks. Vertices with at least two tracks of $p_{\mathrm{T}} > 400$ MeV are considered as primary vertex candidates, and among these the vertex with the largest sum of the squared transverse momenta $\sum_i (p_{\mathrm{T},i})^2$ of its associated tracks is selected and associated with the hard scatter.

3.2 Topological clusters

Electromagnetic and hadronic shower energy depositions in the calorimeters are reconstructed as topological calorimeter clusters (topo-clusters) [76]. The clusters are groups of calorimeter cells which received energy deposits with a sufficiently high signal-to-noise ratio (see Figure 3.3), where the noise is the root-mean-square of the energy distribution measured in events from random bunch crossings, i.e. the average energy expected to be deposited in a calorimeter cell [77]. A seed cell with energy 4σ above the noise level is used to start a topo-cluster, and neighboring cells are included in the cluster. If a neighboring cell has energy 2σ above the noise level, it is treated as a secondary seed cell and its neighbors are included. If no more secondary seeds are present,

all immediately surrounding cells with energy above a very low 0σ threshold are included in the cluster. Finally, if a cluster has two or more local maxima ($E_{\text{cell}}^{\text{EM}} > 500$ MeV), it is split between the signal peaks into two smaller clusters. Cells which were a part of the original cluster are then assigned to and shared between the two highest-energy clusters after the splitting. The results of this splitting procedure can be seen in Figure 3.3(c), where several clusters effectively overlap on their boundaries.

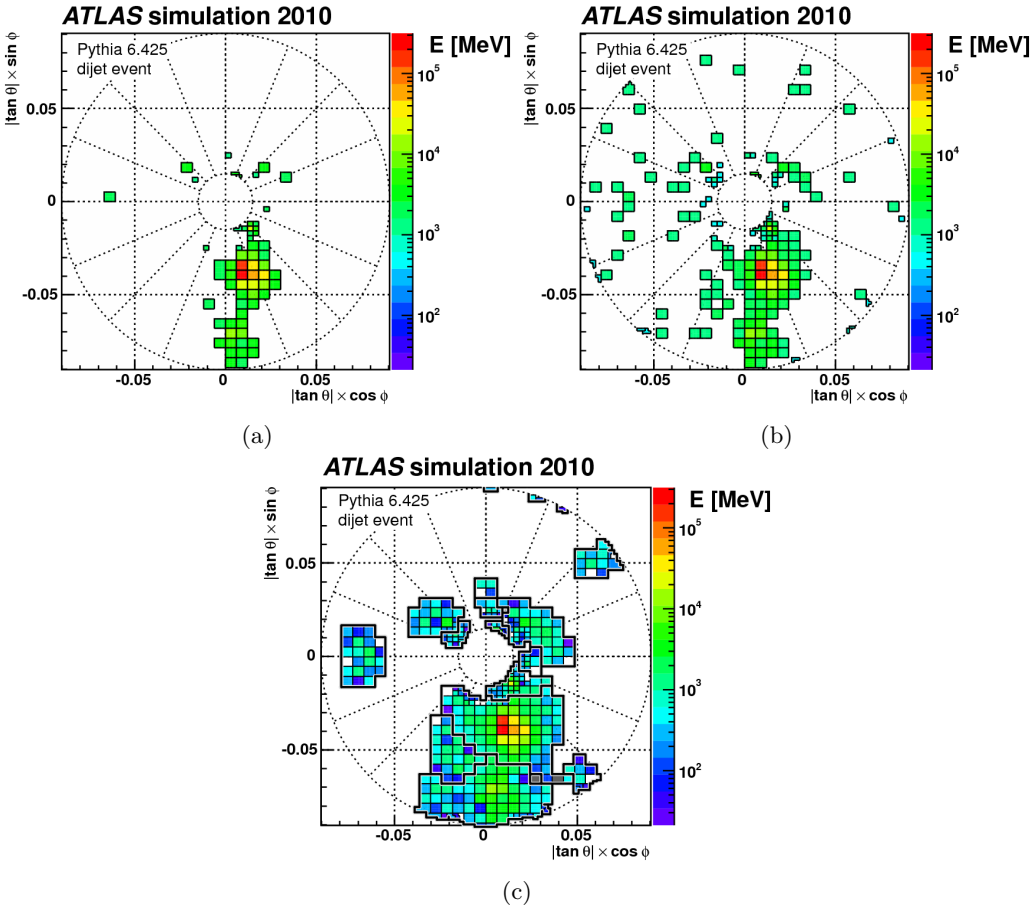

Figure 3.3: Stages of the topo-cluster algorithm in the first module (FCAL0) of the FCAL calorimeter for a simulated dijet event with at least one jet entering this calorimeter [76]. Shown in (a) are cells with $E_{cell} > 4\sigma$ noise, used to seed the topo-clusters; in (b), cells with $E_{cell} > 2\sigma$ noise, used to grow the cluster; and in (c) boundary cells, with the outlines of the clusters added.

The energy of the topo-cluster, the sum of the energies deposited in the constituent cells, is calibrated at the electromagnetic scale. However, the ATLAS hadronic calorimeter is non-compensating, meaning that the response (i.e. how much deposited energy is converted into signal) of the calorimeter to electromagnetic showers is not the same as its response to hadronic showers; though the discrepancy varies with the energy of the incident particle, on average the hadronic response is lower. To recover the hadronic energy loss, Local Hadronic Calibration [78] is used to correct the topo-cluster energies before they are used to reconstruct hadronic jets. In this calibration, using the GEANT4 simulation of the ATLAS detector [79], clusters are identified as

either electromagnetic or hadronic and then are corrected accordingly: both types receive differing corrections for the energy lost outside of the active calorimeter volumes and discarded due to the noise threshold, and the hadronic clusters specifically receive a correction for invisible and escaped energy.

3.3 Muon reconstruction

Muon reconstruction [80, 81] is performed independently in the ID and MS. In the ID, muon tracks are reconstructed as described in Section 3.1; the ID makes no distinctions between the tracks of different charged particles. In the MS, a Hough transform [82] is used to search for hits in the MDTs which align to form a candidate trajectory within the bending plane of the detector. This is done in each of the three layers of MDTs, resulting in MDT segments in each layer (see Figure 3.4). An algorithm then takes the segments from the middle layers, where more hits are available, and attempts to fit them with segments from the inner and outer layers; this is also where the muon sagitta is defined, which is listed in Chapter 9 as a source of experimental uncertainty because it is proportional to the muon momentum, a key kinematic variable. A track candidate is formed if more than two segments are fitted together, aside from in the barrel-endcap transition region where a signal high-quality segment is sufficient. For each track candidate, a global χ^2 fit is performed on the associated hits; tracks passing the selection criteria for this fit are accepted.

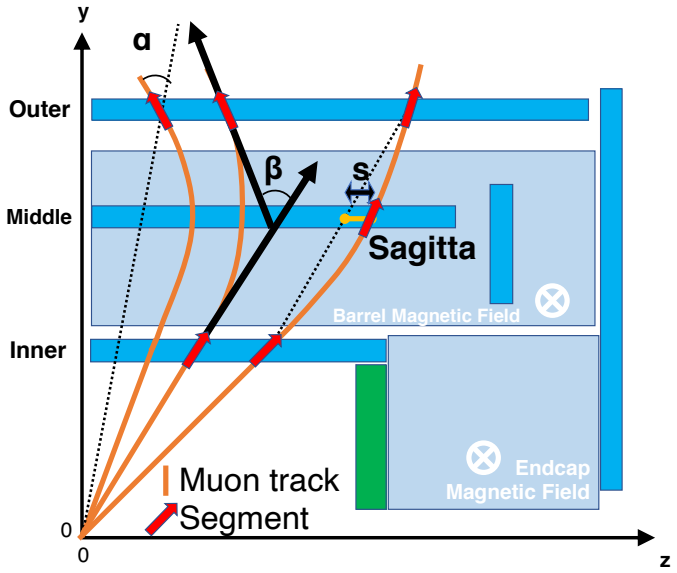

Figure 3.4: Schematic diagram of muons passing through the three MDT layers (blue) of the MS. The muon sagitta is defined as the distance from the position of the segment in the middle MDT station to a virtual straight muon track drawn between the segments in the inner and the outer MDT stations [83].

Reconstructed muons are categorized according to which subdetectors are used in the recon-

struction. Combined (CB) muons are reconstructed by matching MS tracks to ID tracks and performing a combined global refit using hits from both subdetectors. The tracks are associated to the beamline in the transverse plane and to the primary vertex in the longitudinal plane. If an MS track cannot be matched to an ID track, it instead is extrapolated to the beamline and reconstructed as an MS-extrapolated (ME) muon; these are used to extend the coverage of muon reconstruction from the $|\eta| < 2.5$ range of the ID to the full $|\eta| < 2.7$ range of the MS. In a complementary manner, the reconstruction can begin with an ID track, and if the inside-out algorithm can match it to three loosely-aligned hits in the MS, then an inside-out (IO) muon is reconstructed. If instead the ID track can be associated to a full MS track segment, it is reconstructed as a segment-tagged (ST) muon, while if it can be matched with an energy deposit in the calorimeter it is reconstructed as a calorimeter-tagged (CT) muon. The majority of reconstructed muons are CB muons, so the others will not be focused on in this thesis.

3.3.1 Identification

Muon identification is performed by making quality selections to pick out muons with well-reconstructed momenta and to suppress backgrounds (mainly pion and kaon decays). These quality selections are adjusted to create five different muon identification quality working points: Medium, Loose, Tight, Low-p_{T}, and High-p_{T}. The different quality working points, which require varying numbers of hits, detector layers containing hits, and fit significance, are used in different physics analyses depending on the purity, efficiency, and momentum resolution needed. The Medium identification criteria provides a balance of efficiency and purity suitable for most ATLAS analyses while also keeping the reconstruction and background rejection uncertainties low. However, a Higgs analysis might choose the Loose working point to maximize the reconstruction efficiency, while another analysis might opt for the Tight working point in order to maximize the purity of the muon sample, at the cost of a lower reconstruction efficiency.

This analysis uses exclusively Medium muons. They are required to have three or more hits in at least two MDT layers, except in the region $|\eta| < 0.1$, where one MDT with three hits but no more than one MDT with no hits is required. The reconstruction efficiency for all Medium muons with $p_{\mathrm{T}} > 5$ GeV and $|\eta| < 2.5$ was 97% during Run 2.

3.3.2 Isolation

Muons are also categorized according to their relative isolation from hadronic activity in the ID and calorimeters, which serves as a useful parameter for background rejection. The isolation is defined using two variables: the track-based isolation variable $p_{\mathrm{T}}^{\mathrm{varcone30}}$ sums the transverse momenta of tracks within $\Delta R = \min(10\,\mathrm{GeV}/p_{\mathrm{T}}^{\mu}, 0.3)$ of the muon candidate, while the calorimeter-based

isolation variable $E_{\mathrm{T}}^{\mathrm{topocone20}}$ does the same for the transverse energies of topo-clusters within $\Delta R = 0.2$. Ratios of these variables to the muon transverse momentum p_{T}^{μ} are then used to create the selection criteria for the seven different isolation working points. Like the quality working points, these are optimized for different analyses. In this analysis, the `FCTight` working point is used: it requires the p_{T} sum within a variable-radius cone around the combined track in the ID system to be smaller than 0.06 times the muon p_{T}.

3.3.3 Reconstruction efficiency

To obtain the muon reconstruction efficiency within the acceptance of the ID ($|\eta| < 2.5$), a tag-and-probe method is used. In this method, one leg of a $J/\psi \to \mu\mu$ or $Z \to \mu\mu$ decay event is deemed the "tag." The tag leg needs to have fired the lowest unprescaled muon trigger, to trigger online readout of the event, and it must pass stringent identification criteria in order to ensure a pure sample of probe muons. The other leg of the decay is then the "probe" and is independently reconstructed; it is used to measure the reconstruction efficiency. If both muons fulfill the selection criteria of the tag muon candidate, they are considered as tag candidates in turn to avoid bias.

$Z \to \mu\mu$ events are used to obtain the reconstruction efficiency for muons with $p_{\mathrm{T}} > 15$ GeV. The selection criteria on the tag muon include: satisfying the Tight isolation and Medium quality working points; $p_{\mathrm{T}} > 27$ GeV and $|\eta| < 2.5$; requirements on the significance of the transverse impact parameter d_0 ($|d_0|/\sigma(d_0)$, where $\sigma(d_0)$ is the total uncertainty on d_0) and the longitudinal impact parameter z_0; and association with the primary vertex. The probe muon must satisfy differing requirements depending on the type of muon and the particular efficiency measurement being performed.

The reconstruction efficiency is then computed as the number N_P^X of probe muons P that are reconstructed and identified according to X criterion, divided by the total number of probes N_P^{All}:

$$\epsilon(X|P) = \frac{N_P^X}{N_P^{\mathrm{All}}}. \tag{3.1}$$

An efficiency "scale factor," defined as the ratio

$$\mathrm{SF} = \frac{\epsilon^{\mathrm{Data}}(\mathrm{X})}{\epsilon^{\mathrm{MC}}(\mathrm{X})}, \tag{3.2}$$

measures the agreement between the measured efficiency $\epsilon^{\mathrm{Data}}(\mathrm{X})$ and the simulated efficiency $\epsilon^{\mathrm{MC}}(\mathrm{X})$, which is determined using the same events, phase space, and methods. This SF is then used to correct simulations for real detector behavior. SFs are produced in 2D $\eta - \phi$ maps, as the

reconstruction efficiency has been found to be flat in p_{T}.

For the $\sqrt{s} = 13$ TeV data recorded in Run 2, the reconstruction efficiency for Loose and Medium muons was measured to be $> 98\%$ in the majority of the detector phase space ($0.1 < |\eta| < 2.5$ and $p_{\mathrm{T}} > 5$ GeV), while the efficiency for Tight muons was $> 95\%$. In both cases, the agreement between data and simulation was excellent ($\approx 0.5\%$ difference for Loose and Medium muons, $\approx 1\%$ for Tight muons).

3.3.4 Momentum calibration

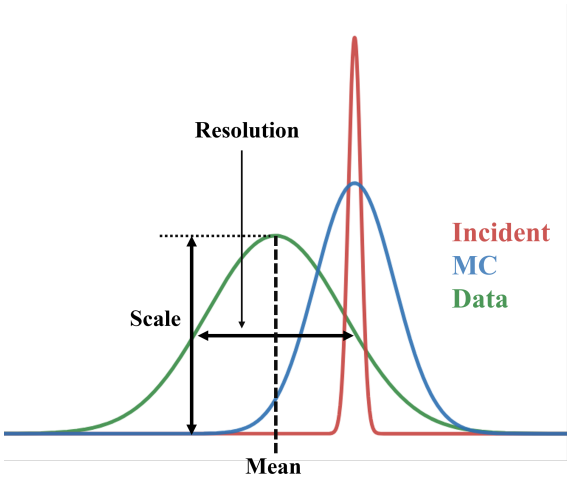

Figure 3.5: A schematic to illustrate the procedure of momentum or energy calibration.

In this and future Sections, I will be discussing calibration procedures for reconstructed objects, so I will first provide a quick overview of the goal of these procedures.

For a given quantity, say the muon transverse momentum p_{T}, a distribution, such as the red one in Figure 3.5, can be obtained using either a controlled test experiment in the ATLAS detector or a simulation. This distribution can be characterized by its resolution (width), scale (height), and mean value. Given the control over the inputs in this scenario, the resolution will be quite small. Using data collected by the ATLAS detector, this distribution cannot be reconstructed with perfect accuracy; the data (green) will result in a measured p_{T} distribution with a much larger resolution as well as a different scale and mean. The first step of the muon transverse momentum calibration is therefore to calibrate the muon transverse momentum data so that its scale and mean align with the test data. In addition, in order for MC to data comparisons to be viable, the MC (blue) p_{T} distribution needs to have a resolution comparable to the measured data. Therefore, the MC distribution is smeared, enlarging its resolution to match the resolution of the data. The resolution effects intrinsic to the detector are only dealt with through unfolding, which is discussed in Chapter 8.

In the ATLAS muon momentum calibration procedure, fine corrections to the simulated muon

momenta reconstructed in the ID and MS subdetectors are applied. The corrections to the simulated muon momenta are defined using parameters extracted from data using a maximum-likelihood fit comparing the $J/\psi \to \mu\mu$ and $Z \to \mu^+\mu^-$ invariant mass distributions; they are then applied in (η, ϕ) regions with relatively homogeneous detector technology and performance. To correct the momentum scale, the parameters extracted from data account for inaccuracies in the description of the magnetic field and the size of the detector in the plane perpendicular to the field, both of which affect the momentum scale. Likewise, to correct the muon momentum resolution, the parameters account for inaccuracies in describing energy losses in the calorimeters and other materials between the ID and MS; intrinsic resolution limitations due to the spatial resolution of hits; and residual misalignment of the MS.

3.3.5 Trigger efficiency

The efficiency of the muon trigger system [84] is also evaluated using a tag-and-probe method. Because I have personally been involved in muon trigger efficiency scale factor production throughout my doctorate, including the development of a centralized production framework and tutorial, I will describe these efficiency measurements in perhaps more detail than is strictly necessary for this thesis.

To select a clean di-muon sample of Z boson decay events, a pair of oppositely-charged muons originating from the same interaction vertex, using transverse and longitudinal impact parameter selections, is required. The invariant mass of the pair is required to be consistent with the Z-boson mass, $|m_Z - m_{\mu\mu}| < 10$ GeV. If one of the two muons has $p_{\mathrm{T}} > 28$ GeV and satisfies the *Loose* isolation requirement, it is considered as a tag candidate. The tag candidate must further have an angular distance of $\Delta R < 0.1$ from an object that fired the lowest unprescaled muon trigger (`HLT_mu20_iloose_L1MU15` during 2015 and `HLT_mu26_ivarmedium` during 2016-2018). These strict requirements on the tag, as mentioned earlier, are for the purpose of ensuring as little bias as possible in the selection of the probe candidates; looser requirements on the tag could lead to the preferential selection of probes with particular kinematics, etc.

The other muon is then taken as a probe candidate. As opposed to the tag candidate, the probe candidate is subject to selections defined by the identification working point and trigger for which the efficiency is desired. The identification working point selections include the number of hits in the MDT chambers and the compatibility of tracks reconstruced in different subdetectors; see Section 3.3.1 for the Medium woking point as an example. To ensure it is matched to the trigger, the probe must lie within $\Delta R < 0.1$ of an object which caused the trigger to fire. The trigger efficiency is then defined as the fraction of probe candidates associated to at least one trigger muon object.

Trigger efficiencies are determined for L1 (Figure 3.6) and HLT (Figure 3.7) triggers; the HLT efficiencies shown are for the L1-HLT chain, as the HLT efficiency relative to the L1 efficiency is close to unity. The slight HLT inefficiency in the turn-on curve region is due to the fact that no isolation requirements are imposed on the offline-selected muons; past the turn-on curves, the efficiency plateaus extend to $p_{\mathrm{T}} \approx 500$ GeV and beyond but are measured using $t\bar{t}$ and W+jet events at $p_{\mathrm{T}} \gtrsim 100$ GeV for better statistics. The efficiencies are lower in the barrel region because of reduced geometrical coverage, the result of support structures and the need for service access to the ID and calorimeters. The drop in L1 efficiency in the endcap region for 2017 and 2018 was due to gas leaks in the RPC chambers. Figure 3.7 also shows the trigger efficiency scale factors in the bottom panel, which account for the differences in L1-HLT trigger chain efficiency between data and MC. The discrepancy is more noticeable in the barrel region due to a different MC RPC efficiency. It can be seen that these scale factors do not have a strong dependence on p_{T}, and so they are delivered in two-dimensional (η, ϕ) histograms; examples are shown in Figure 3.8.

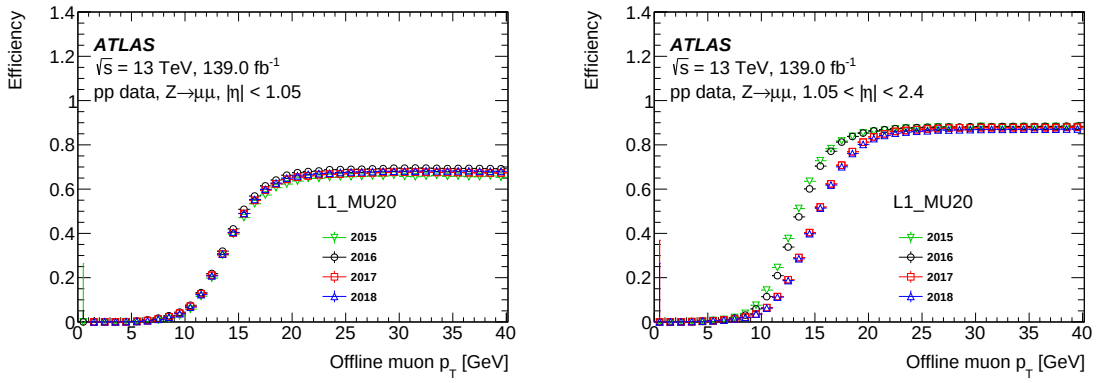

Figure 3.6: Efficiency of the single-leg trigger `L1_mu20`, in the barrel (left) and endcap (right) regions, as a function of muon p_{T} in 2015-2018 [84].

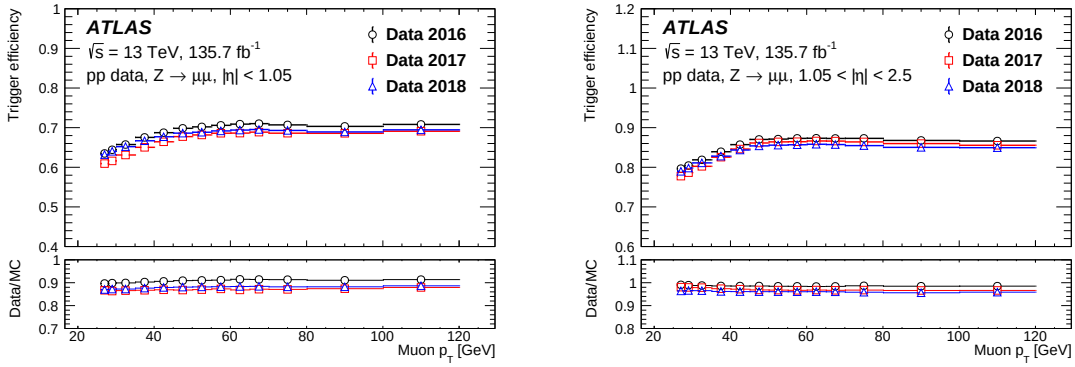

Figure 3.7: Efficiency and scale factor for passing either the `HLT_mu26_ivarmedium` or the `HLT_mu50_ivarmedium` trigger, in the barrel (left) and endcap (right) regions, as a function of muon p_{T} in 2016-2018.

The sources of systematic uncertainty considered in the trigger efficiency measurements are

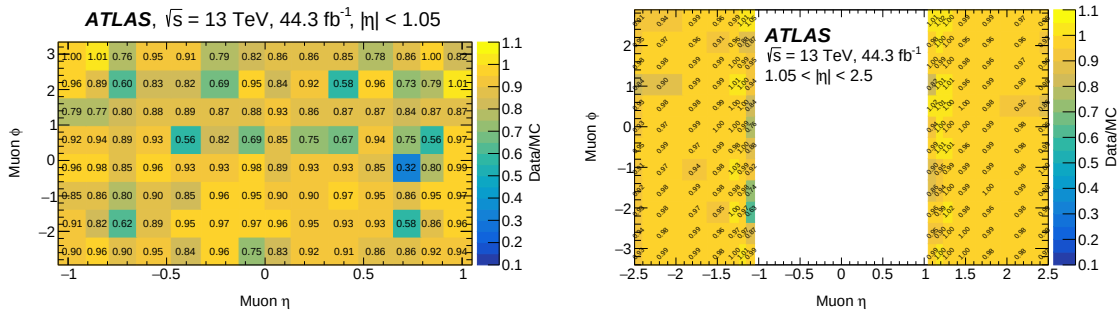

Figure 3.8: 2017 trigger efficiency scale factors for passing either the `HLT_mu26_ivarmedium` or the `HLT_mu50_ivarmedium` trigger, in the barrel (left) and endcap (right) regions, shown in two-dimensional bins of η and ϕ.

listed below. In each case, the systematic uncertainty is obtained by taking the difference in efficiency determined with and without the systematical variations described; the total systematic uncertainty is then the sum in quadrature of these differences.

- **Pile-up dependence:** In order to evaluate the uncertainty related to pile-up interactions, a cut on the number of reconstructed vertices is imposed, the trigger efficiency above and below this cut is computed, and each is compared to the nominal. The cut is made at 11 reconstructed vertices for 2015 and 2016, and at 19 reconstructed vertices for 2017 except for Period K 2017, a high-pile-up period, for which the cut is set at 25 reconstructed vertices. This procedure was performed separately for data and MC.

- **Correlation between the tag and probe muons:** Muon pairs from the $Z \to \mu^+\mu^-$ decay tend to be back-to-back in ϕ. Since the barrel and endcaps have 16-fold and 12-fold symmetry respectively, this means that if the tag muon ends up in a highly-efficient region of the detector, the probe muon is likely to as well. To evaluate potential bias associated with this effect, the trigger efficiency is calculated using an extra selection cut, $\Delta\phi(\text{tag, probe}) < \pi - 0.1$. Selecting only tag and probe muons that satisfy the condition provides a sample enriched in muons that are not back-to-back.

- **Probe selection criteria:** The effects of various probe selection cuts on trigger efficiency are investigated. These include:

  - *Charge:* Since charge affects the behaviour of muons in the magnetic field present in the MS, it could conceivably affect muon trigger efficiency. The impact of this systematic was evaluated by comparing the trigger efficiency using only all-positive or all-negative probe muons; each sample was then compared to the nominal.

  - *Impact parameters:* To investigate the effect of the probe impact parameter selections on the efficiency, the efficiency is calculated without impact parameter cuts.

  - *Isolation:* An isolation cut is added to the probe muon selection in order to check the effect on the trigger efficiency of requiring isolated probes.
  - *Transverse momentum:* The efficiency scale factors are only delivered in bins of (η, ϕ), so their dependence on p_{T} is monitored by splitting probe muons into two groups according to a p_{T} cut; each group is then compared to the nominal. The value of the cut depends on the p_{T} threshold of the trigger of interest: for trigger thresholds of 26 GeV the cut is 40 GeV, while for thresholds of 50 GeV a cut at 70 GeV is used. Trigger efficiency is then calculated using only the subset of probe muons with p_{T} above or below that cut; each subset is then compared to the nominal.

- **Background contribution:** To evaluate the background-related uncertainty, the efficiency measurements are repeated by enlarging the Z-mass window by ± 5 GeV.

3.4 Electron reconstruction

Much of the procedure for reconstructing electrons is similar to that used to reconstruct muons, including the assembly of hits in the ID into tracks; the use of calorimeter- and track-based isolation working points; the creation of several identification working points for use in different types of physics analyses; and the application of momentum calibrations. However, the electron reconstruction faces the additional challenge of abundant brehmsstrahlung energy loss and pair production. EM energy deposity in the electromagnetic calorimeter (ECal) also have to be identified as coming from electrons or photons. Notable differences and additions to the identification and reconstruction processes are given below.

3.4.1 Reconstruction

Electrons are also reconstructed [85, 86] from clusters of ECal energy deposits. In 2015-2016, a sliding-window algorithm [87] was used. First, the $\eta - \phi$ space of the ECal is split into a 200×256 grid of elements ("towers") of size $\Delta\eta \times \Delta\phi = 0.025 \times 0.025$ (the granularity of the second layer of the ECal). For each tower, the energy depositions in the three layers (and the presampler, for $|\eta| < 1.8$; see Figure 3.9) are summed together. Electromagnetic energy cluster candidates are then seeded from localized energy deposits using the sliding-window algorithm: the energy of a 3×5 section of towers is summed, and if it exceeds 2.5 GeV it is kept as a candidate, and the window moves over one tower in η or ϕ, repeating the process and removing the smaller of cluster candidates in close proximity. In this manner, the reconstruction efficiency of clusters with $E_{\mathrm{T}} > 15$ GeV is greater than 99%.

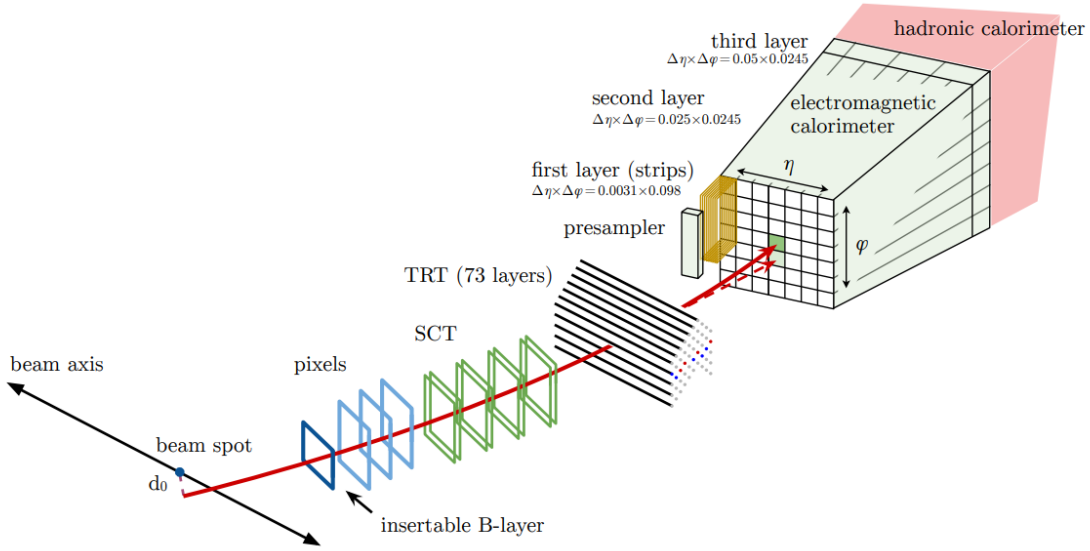

Figure 3.9: A schematic illustration of an electron passing through the ATLAS detector (red solid line). The red dashed line represents a photon produced by the electron interacting with the tracking material [85].

In 2017-2018, superclusters were used instead of the sliding-window algorithm to reconstruct electrons. EM topo-clusters, sorted by E_{T} and having $E_{\mathrm{T}} \geq 1$ GeV, are sequentially used as seed candidates. For electrons and photons, nearby (within the 3×5 cell window) "satellite" topo-clusters are added to the seed candidate; for electrons only, additional satellite clusters within a 5×12 cell window are added if their best-matched track is the same as the best-matched track of the seed cluster (see Figure 3.10). The electron reconstruction efficiency using the supercluster method plateaus at $\geq 95\%$ for $E_{\mathrm{T}} \geq 15$GeV; this is less than the efficiency of the prior sliding-window method, but the dynamically-sized superclusters better capture the energy of brehmsstrahlung photons or electrons from converted photons.

ECal clusters are then matched matched to ID tracks to distinguish electrons from photons. The tracks are refitted using a Gaussian Sum Filter (GSF) [89], which is based on the Kalman filtering technique mentioned in Section 3.1, to account for energy loss due to brehmsstrahlung. The seed-cluster matching algorithm works by checking to see if a given electron track candidate falls within the EM cluster region-of-interest. Clusters matched to ID tracks, signalling an origin from an electron produced in the beam interaction region, are classified as electrons, while clusters which cannot be matched to tracks are classified as unconverted (pair-producing in the ECal) photons.

Particle charge can be reconstructed from the curvature of their tracks in the ID. For electrons, this identification is complicated by the presence of multiple tracks due to pair production as the result of brehmsstrahlung. A single instance of pair production will lead to three tracks (the tracks of the electron-positron pair plus that of the original electron), where two of the tracks will have the

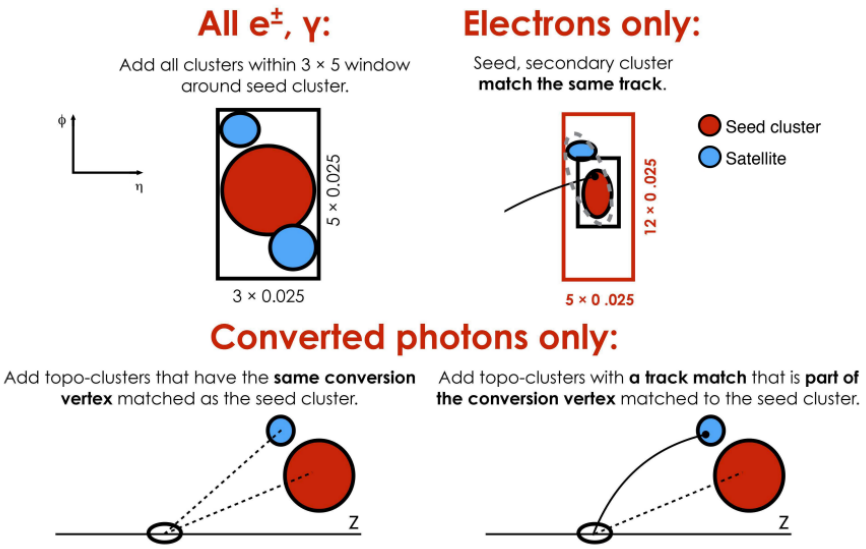

Figure 3.10: Diagram of the superclustering algorithm for electrons and photons, showing seed clusters in red and satellite clusters in blue [88].

correct charge. To reduce the misidentification rate, a boosted decision tree trained on a simulated single-electron sample is used to apply an additional selection criterion.

The reconstruction efficiency of electrons with "good-quality" tracks consisting of at least seven silicon hits and at least one pixel hit was between 97% and 99% for $E_{\mathrm{T}} > 15$ GeV for 2015-2016 data; $E_{\mathrm{T}} = 15$ GeV, the reconstruction efficiency was determined solely from simulation.

3.4.2 Identification

Whereas the muon identification is performed through the use of quality requirements in the form of cuts, electrons are identified in the online selection through the use of a likelihood (LH). The inputs to the electron LH include tracking and calorimeter quantities i to distinguish prompt electrons from hadronic jets, converted photons, and electrons from heavy-flavor hadron decays. Signal (prompt electron) and background (non-prompt electrons, as well as jets and electrons from photon conversions) probability density functions (pdfs) $P_{S,i}(x_i)$ and $P_{B,i}(x_i)$, respectively, are then determined for quantity i at value x_i. Using simulation prior to 2017 and subsequently using 2015-2016 data [88], the pdfs are determined by smoothing the distributions of the quantities with the adaptive kernel density estimator in the TMVA toolkit [90]. An overview of the different pdfs used in Run 2 (2015-2016) is given in Table 1 of Ref. [85]. The values of the pdfs are then used to construct the likelihood function for each event:

$$L_{S(B)}(\mathbf{x}) = \prod_{i=1}^{n} P_{S(B),i}(x_i). \tag{3.3}$$

From these LHs, a discriminant d_L is determined as

$$d_L = \frac{L_S}{L_S + L_B}. \tag{3.4}$$

The identification working points for the electron reconstruction (VeryLoose, Loose, Medium, and Tight) are then defined based on a value of the inverse sigmoid-transformed discriminant

$$d'_L = -\tau^{-1}\ln(d_L^{-1} - 1). \tag{3.5}$$

The inverse sigmoid is used because the discriminant is often peaked at 0 (background) and 1 (signal), and so applying the inverse sigmoid zooms in on the more central peaks. This prescription is taken from the TMVA toolkit with $\tau = 15$.

Some advantages to the LH method of identification for electrons over the cut-based method used for the muons include the ability to use discriminating quantities with similar distributions without loss of efficiency and the avoidance of failing electrons because they do not satisfy the selection criterion of a single quantity. In this analysis, Tight electrons are used, which are identified at an efficiency of 80%.

3.4.3 Isolation

Similarly to muons, track-based and calorimeter-based variables are constructed to measure the relative isolation of electron candidates. In this analysis, the `FCTight_FixedRad` working point was used: it requires the sum of the calorimeter energy deposits and the p_{T} of the ID tracks within a variable-width (decreasing as a function of the p_{T}) cone of the electron to be less than 0.6 times the electron's transverse energy.

3.4.4 Momentum calibration

After the energies of the electrons and photons are calculated from the reconstructed clusters to determine electron candidates, electron energy calibration [91] is performed. Several steps are taken to correct the electron energy scale in data before calibrating it with respect to smulation. First, an accurate estimate of the reconstructed cluster energies are determined in simulation and data by correcting for energy losses in materials upstream of the calorimeter, using cells in the immediate (η, ϕ) vicinity of the cluster and beyond the liquid argon (LAr) calorimeter, with a multivariate regression algorithm. Then, based on studies in simuation of energy deposits in the ECal, corrections are applied to data to adjust the relative scale of the first and second ECal layers; muon energy deposits from $Z \to \mu\mu$ events are used in the studies, as they are relatively insensitive

to the passive material in front of the ECal. Next, the energy shift due to pile-up is corrected in data by subtracting the average expected shift cell-by-cell and also by applying a correction to each ECal layer accounting for the average number of interactions per bunch crossing as a function of eta. Finally, relative non-uniformities in the amounts of passive material (due to, for example, gaps in the absorbers in the calorimeters) and in the applied high voltage are corrected in data.

After all of these corrections are applied, the electron energy scale in data is calibrated in regions of η. Simultaneously, the electron energy resolution calibration is applied to simulation (also in regions of η). Both the energy scale and energy resolution calibration terms are determined by comparing invariant mass distributions in data and simulation. The calibration is validated by studying $J/\psi \to ee$ samples. Systematic uncertainties, including those related to pile-up, layer calibration, and the material in front of the ECal, are determined using MC variations.

3.5 Jets

One of the most common features of proton-proton interactions at the LHC are collimated sprays of hadrons produced by the fragmentation of energetic quarks and gluons; these sprays are referred to as jets. Jets are a fundamental piece of many measurements and searches, and for those searches the uncertainty in the jet energy often is the largest experimental uncertainty. Therefore, the careful and precise reconstruction, calibration, and selection of jets is extremely important.

3.5.1 Reconstruction

Jets are reconstructed from topo-clusters using the anti-k_t algorithm [92] implemented in the FastJet software package [93]. The algorithm defines the distance d_{ij} between two objects i and j, with the objects being either topo-clusters or previously-grouped topo-clusters called "pseudojets," and the distance d_i between object i and the beam as follows:

$$\begin{aligned} d_{ij} &= \min\left(\frac{1}{p_{T,i}^2}, \frac{1}{p_{T,j}^2}\right)\frac{\Delta R_{ij}^2}{R^2}, \\ d_i &= \frac{1}{p_{T,i}^2}. \end{aligned} \tag{3.6}$$

If d_{ij} is smaller than d_i, objects i and j are combined into a pseudojet; otherwise, object i is called a jet and removed from subsequent iterations of the clustering algorithm. The distances d_{ij} and d_i are redefined and objects are combined or removed until no objects are left.

There are many different jet-clustering algorithms available, such as the k_t algorithm, which simply uses $p_{T,i}^2$ instead of $1/p_{T,i}^2$. The advantages of the anti-k_t algorithm are that all but the softest jets result in circular clustering cross-sections, as seen in the bottom right of Figure 3.11,

and that the algorithm is infrared and collinearly safe: neither soft particles nor collinear splittings change the jet definition (see Figure 3.12).

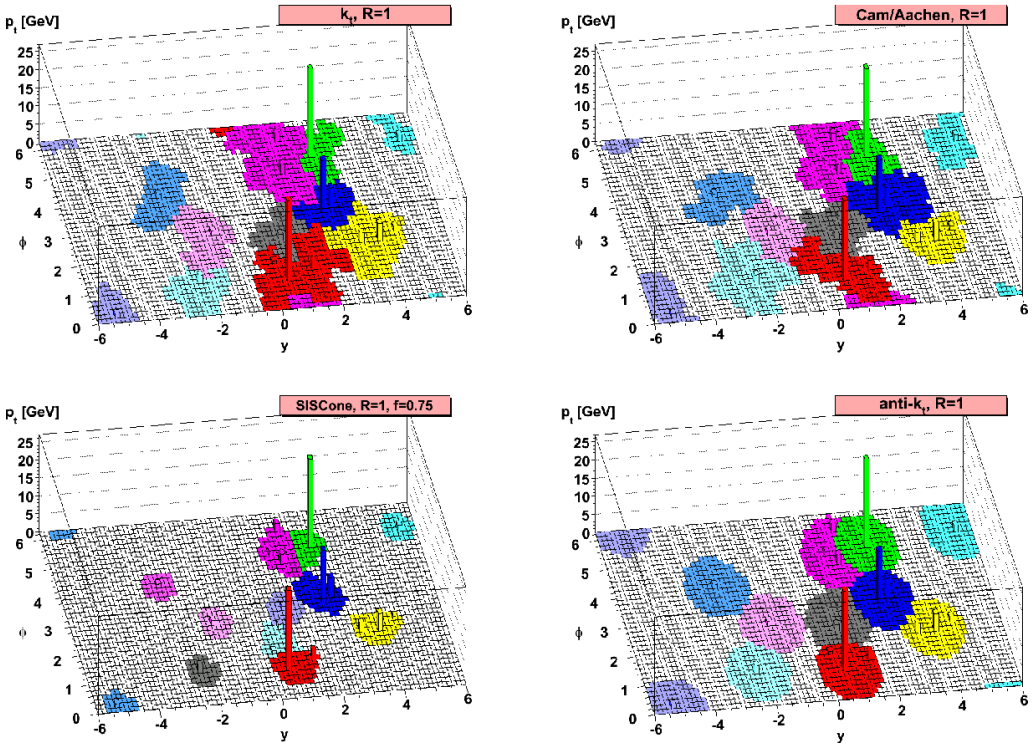

Figure 3.11: A sample event, along with soft "ghosts" artificially inserted, with the jets clustered using different algorithms. The anti-k_t algorithm is shown in the bottom right. Each color denotes a separate reconstructed jet [92].

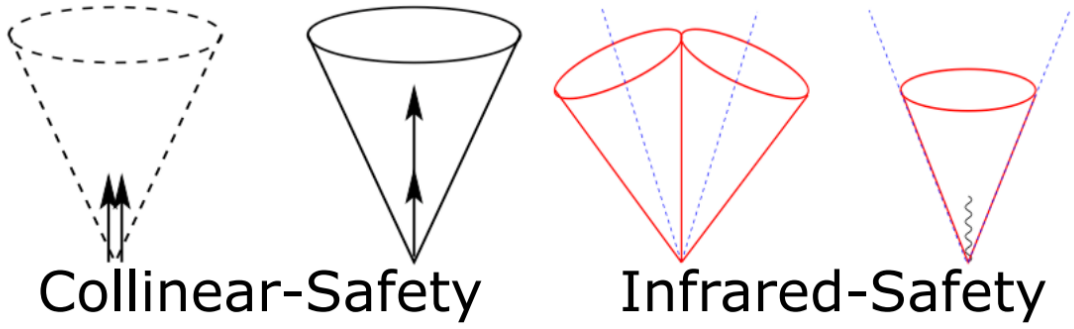

Figure 3.12: Schematic demonstrating the problems which collinear and infrared safety avoid: a collinear splitting should not cause a jet to be not reconstructed (left), nor should a soft emission change the jet definition (right) [94].

The jets reconstructed from electromagnetic topo-clusters (even uncalibrated ones, since jets are calibrated directly, as described in the next Section) are referred to as *EMTopo* jets. They are defined at the electromagnetic energy scale $E^{\text{jet}}_{\text{meas}}$ at which the calorimeter correctly measures electromagnetic shower energy depositions. EMtopo jets were the standard collections used in ATLAS analyses until the end of Run 2. The jets which are used in this analysis, however, are reconstructed using the "particle flow" approach, in which information from the tracker is

combined with the topo-clustering process described above [95]. In EMTopo jets, the energy deposited by charged particles in the calorimeter is used to determine the jet four-momentum. *PFlow* jets instead use the momenta of the tracks which are associated with the topo-clusters of charge particles, in addition to the topo-clusters of neutral particles. One of the main advantages of this approach is that it improves the overall resolution of low-p_{T} jets (see Section 3.5.3), since for those jets the momentum resolution of the tracker is significantly better than the energy resolution of the calorimeter (see Table 2.1). Additionally, the association of tracks to vertices can determine whether particles originate from pile-up vertices and therefore should be rejected. Just as for EMTopo jets, the anti-k_t jet clustering algorithm is used to construct PFlow jets.

3.5.2 Jet energy scale calibration

Simulation is used to calibrate the energies of reconstructed jets, on average, to the energies of their constituent stable particles; this is referred to as the jet energy scale (JES) calibration [77, 96]. This calibration accounts for detector effects which affect the jet energy measurements, including the non-compensation of the calorimeter for hadronic energy deposits, dead material, and leakage of particles beyond the calorimeters. In the simplest calibration scheme, the JES calibration is performed by applying the calibration function $\mathcal{F}_{\mathrm{calib}}$ to the energies[1] of jets reconstructed at the EM scale, $E^{\mathrm{jet}}_{\mathrm{meas}}$:

$$E^{\mathrm{jet}}_{\mathrm{calib}} = E^{\mathrm{jet}}_{\mathrm{meas}} / \mathcal{F}_{\mathrm{calib}}(E^{\mathrm{jet}}_{\mathrm{meas}}), \text{ with } E^{\mathrm{jet}}_{\mathrm{meas}} = E^{\mathrm{jet}}_{\mathrm{EM}} - \mathcal{O}(N_{\mathrm{PV}}) \tag{3.7}$$

This calibration scheme consists of three sequential corrections. First, to account for the extra energy due to additional pp interactions, a pile-up correction is made through an offset applied to the jet transverse energy:

$$E_{\mathrm{T}}^{\mathrm{corrected}} = E_{\mathrm{T}}^{\mathrm{uncorrected}} - \mathcal{O}(\eta, N_{\mathrm{PV}}, \tau_{\mathrm{bunch}}). \tag{3.8}$$

The offset is derived from minimum bias data and is a function of η, the number of primary vertices N_{PV}, and the bunch spacing τ_{bunch}. Next, a correction is applied to the jet four-momentum such that the direction of the topo-clusters points back to the primary vertex instead of the geometrical center of the ATLAS detector, which is initially used as a reference to determine the direction of the jets and their constituents. Finally, the energy of the jets is calibrated by creating a jet response calibration function $\mathcal{F}_{\mathrm{calib}}$. For individual bins of the truth jet energy $E^{\mathrm{jet}}_{\mathrm{truth}}$ and η (determined before the origin correction above), the averaged jet calorimeter energy $\left\langle E^{\mathrm{jet}}_{\mathrm{truth}} \right\rangle$ and the averaged

[1] Calibrations are applied to jet energy E rather than E_{T} because the calorimeters explicitly measure E, and the response curves can therefore be compared directly to the expected energy.

jet energy response $\left\langle \mathcal{R}^{\text{jet}}_{\text{EM}} \right\rangle$ are determined. The calibration function is then obtained, for a given value of η, from a fit on these averaged values for each $E^{\text{jet}}_{\text{truth}}$ bin, and is applied to the measured jet energy to complete the JES calibration:

$$E^{\text{jet}}_{\text{JES}} = \frac{E^{\text{jet}}_{\text{EM}}}{\mathcal{F}_{\text{calib}}(E^{\text{jet}}_{\text{EM}})|\eta}.$$

3.5.3 Jet energy resolution correction

Studies which measure jets or SM particles which decay into jets, as well as BSM searches relying on a precise measurement of the missing transverse energy (MET), require a thorough understanding of the jet energy resolution (JER) of ATLAS. The dependence of the relative JER on the jet transverse momentum p_{T} can be parameterized as

$$\frac{\sigma(p_{\text{T}})}{p_{\text{T}}} = \frac{N}{p_{\text{T}}} \oplus \frac{S}{\sqrt{p_{\text{T}}}} + C, \tag{3.9}$$

where the noise term N is due to detector front-end electronics noise and pileup, the stochastic term S accounts for statistical fluctuations in the amount of energy deposited, and the constant term C includes effects which are constant in p_{T}, such as non-uniformities in the calorimeter response [96]. N dominates below ≈ 30 GeV, while C dominates above ≈ 400 GeV. The relative energy resolution is measured using a dijet balance method (see [96]) and varies from $(24 \pm 1.5)\%$ at $p_{\text{T}} = 20$ GeV to $(6 \pm 0.5)\%$ and 300 GeV (see Figure 3.13); as mentioned earlier, this is a substantial improvement in energy resolution at low p_{T} over that of EMTopo jets (35%; see Figure 3.14).

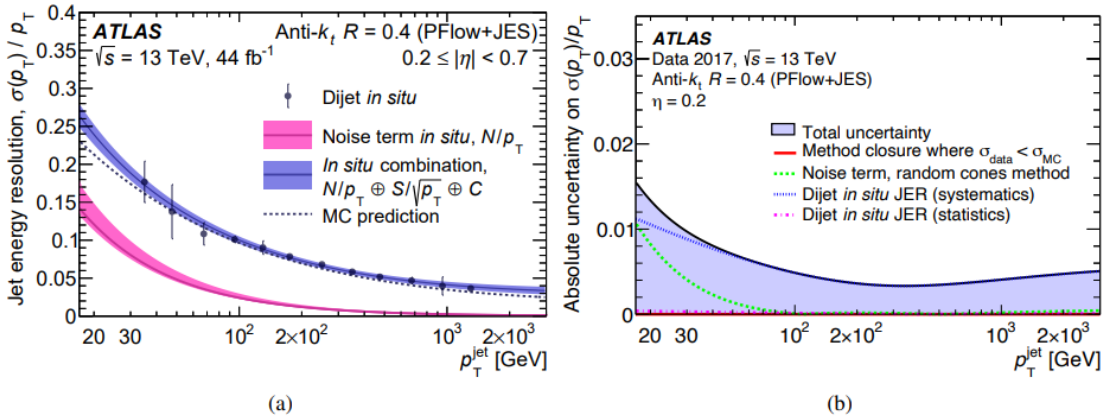

Figure 3.13: (a) Relative jet energy resolution for JES-calibrated PFlow jets as a function of $p^{\text{jet}}_{\text{T}}$, with the noise term ($N$ in Equation 3.9) shown in pink. Error bars indicate the total (statistical and systematic) uncertainties on the determination of the relative resolution using dijet events. (b) Absolute uncertainty on the relative jet energy resolution as a function of jet p_{T}. The fit to the resolution as a function of $p^{\text{jet}}_{\text{T}}$ shows a resolution improvement for PFlow jets over EMTopo jets at low p_{T} [96].

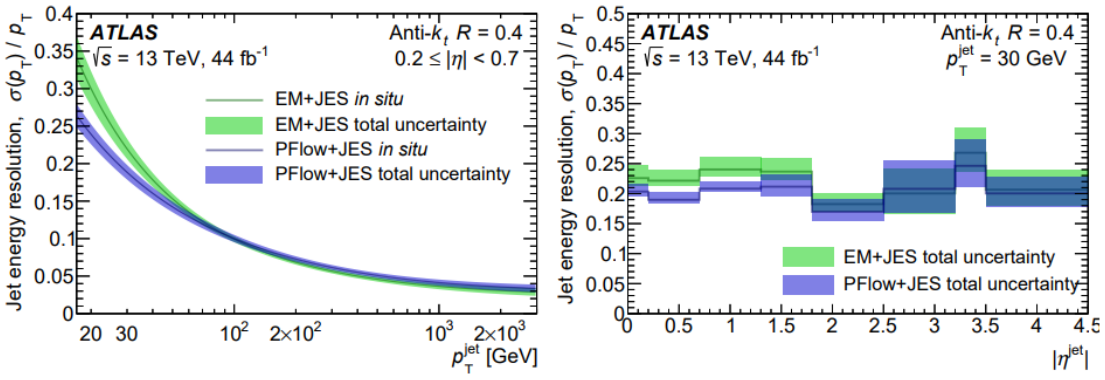

Figure 3.14: Relative jet energy resolution for JES-calibrated PFlow (blue) and EMTopo (green) jets as a function of (a) $p_{\mathrm{T}}^{\mathrm{jet}}$ and (b) $|\eta^{\mathrm{jet}}|$. The fit to the resolution as a function of $p_{\mathrm{T}}^{\mathrm{jet}}$ shows a resolution improvement for PFlow jets over EMTopo jets at low p_{T} [96].

With the JER in data known, the resolution in MC is made to match it as much as possible. This is ensured by smearing MC jets using a Gaussian function in regions of jet p_{T} where the data JER is larger than the resolution in MC. Uncertainties in this smearing, as discussed in Chapter 9, are then propagated through the analysis.

3.5.4 Jet quality selection

Jet selection is applied through cuts to reject jets coming from background events, which include beam-gas and beam-halo events and cosmic ray muons. For example, a jet produced by the emission of a large brehmsstrahlung photon in the calorimeter by a cosmic ray muon can be effectively rejected by cutting on the jet time t_{jet}, as the jet from the brehmsstrahlung photon will not be in-time with the beam collision. Two jet quality selections, Loose and Medium, are defined. The Loose working point, whose efficiency is above 99%, is the recommendation for the majority of analyses and is used in this thesis.

3.5.5 b-jet tagging

In order to study events involving b-jets and/or c-jets, it is necessary to distinguish those events from each other and from events involving light-jets. Fortunately, the "heaviness" of the b-quark results in distinguishing features of their jets as opposed to jets originating from light quarks or gluons. The semi-leptonic width $\Gamma_{\mathrm{sl}}(b \to q)$ for a b-quark to decay into another quark can be expressed as [97]

$$\Gamma_{\mathrm{sl}}(b \to q) = |V_{qb}|^2 \frac{G_F^2 m_Q^5}{192\pi^3} I(\epsilon), \tag{3.10}$$

where V_{qb} are elements of the Cabibbo–Kobayashi–Maskawa quark-mixing matrix [27], $G_F = g^2/8M_W^2$, $\epsilon = \frac{m_q}{m_b}$, and

$$I(\epsilon) = 1 - 8\epsilon^2 + \epsilon^6 - \epsilon^8 - 24\epsilon \ln \epsilon. \tag{3.11}$$

In large part because the V_{qb} are quite small ($V_{cb} = (42.2 \pm 0.8) \times 10^{-3}$, $V_{ub} = (3.94 \pm 0.36) \times 10^{-3}$), the b-quark has a small semi-leptonic decay width, which corresponds to a long lifetime. The b-quark can therefore travel through the beam pipe before decaying in the inner detector. The result is that heavy-quark jets can be distinguished by the measurement of the impact parameter from the secondary vertex, as shown in Fig. 3.15.

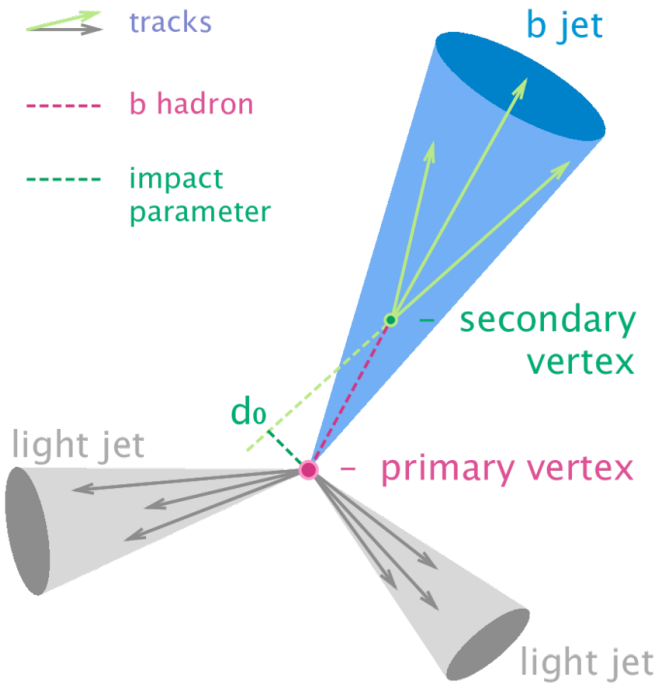

Figure 3.15: A diagram of a three-jet event where one jet originates from a b-quark. The large mass of the b-quark results in a longer lifetime and therefore a displaced vertex for its decay products [98].

There are currently several low-level tagging algorithms in use at ATLAS which work to identify these features. The products of these algorithms are discriminants whose values are based on the likelihood that a given jet is a b-jet or not. These discriminants are then combined using the DL1 deep neural network [99] to produce a more comprehensive flavor tagging algorithm. This analysis makes use of a particular version of this neural network, DL1r, which outputs three nodes for each jet: p_b, p_c, and p_{light}, representing the probability that the jet is a b-jet, c-jet, or light-jet,

respectively. These nodes are used to construct the final discriminant

$$D_{\text{DL1}} = \ln\left(\frac{p_b}{f_c \cdot p_c + (1 - f_c) \cdot p_{\text{light}}}\right). \tag{3.12}$$

f_c is the effective c-jet fraction in the background training sample, which can be tuned to enhance discrimination of b- vs. c-jets or b- vs. light-jets; the value for the DL1r algorithm sets $f_c = 0.018$.

Using the DL1r output scores, an operating working point (OP) can be defined by a cut on the score which corresponds to a particular b-jet tagging efficiencies. The efficiency is defined as

$$\varepsilon^b = \frac{\text{N}^b_{\text{pass}}(D_{\text{DL1r}} > T_f)}{\text{N}^b_{\text{total}}}, \tag{3.13}$$

where $\text{N}^b_{\text{Total}}$ is the total number of b-jets and N^b_{pass} is the number of b-jets with a DL1r score larger than a threshold T_f. T_f is the cut value for the OP and is determined using a simulated $t\bar{t}$ sample. OP are defined for b-jet tagging efficiencies of 60%, 70%, 77%, and 85%. Finally, these efficiency OP can be used as bin edges to form "pseudo-continuous" b-tagging weight bins. The bins correspond to b-jet tagging efficiencies of [0, 60]%, [60, 70]%, [70, 77]%, [77, 85]%, and [85, 100]% [100]. The distribution of jets falling into one of these bins will be used extensively in future Chapters as an observable quantity.

This concludes the discussion on object reconstruction from ATLAS data. The next Chapter will discuss why, in particular, this analysis is focused on heavy-flavor jets produced in association with Z bosons.

Chapter 4

Z Boson Associated Production with Heavy-Flavor Jets

During Run 2, the LHC produced around 10 billion events which included a Z boson, the neutral mediator of the weak interaction, in the final state. Around 120 million of these Z bosons have been observed in ATLAS data through the Drell-Yan decay process, in which the Z decays into two leptons of the same flavor (Figure 4.1a). In addition, the Z is often produced in association with jets originating from quarks or gluons, against which the Z recoils (e.g. Figure 4.1b). Those jets are categorized as light-jets if they originate from a gluon or an up, down, or strange quark and as heavy-flavor (HF) jets if they originate from beauty or charm quarks.[1] HF jets are the focus of this analysis, and in this Chapter I will provide an overview of the study of both flavors.

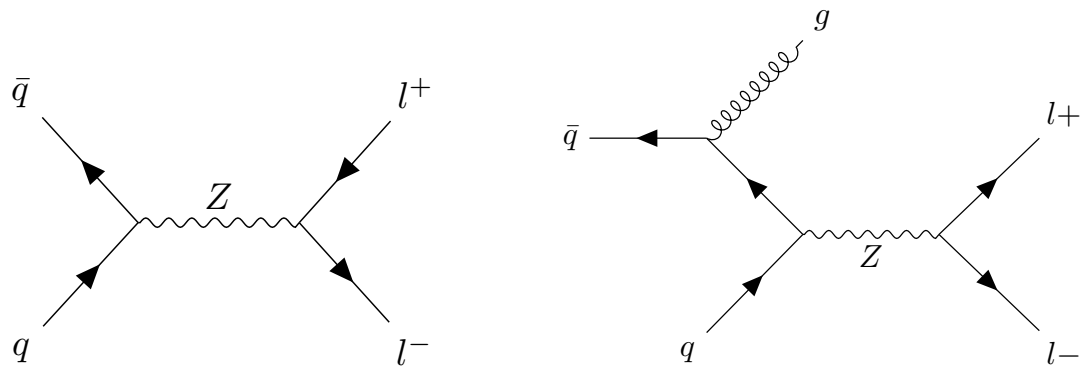

(a) (b)

Figure 4.1: (a) Drell-Yan production mode of a Z production. (b) Drell-Yan Z production in association with a jet from initial-state radiation (b).

[1] Top quarks decay weakly, most often into a b-quark and a W boson, before hadronizing due to their very large mass.

4.1 Z+b-jets

There are two main reasons why Z+b-jets events are particularly interesting. The first is that the modelling of Z+jet events is performed using approximations which work well for light-jets but which begin to break down for HF jets; this is due to the larger masses of the b- and c-quarks, and is described in the following subsections. This tension results in significant systematic uncertainty in the modelling of Z+HF jets events, which can only be alleviated by better understanding the QCD describing these events. The second reason is that the largest branching fraction of the Higgs at its rest mass $m_H = 125$ GeV is to two b-quarks (Figure 4.2), and so Z+b-jets events are a significant background to SM Higgs measurements such as Figure 4.3(b). The same is true for a slew of BSM searches for dark matter or supersymmetry, such as the pseudoscalar search shown in Figure 4.3(c). The modelling of Z+b-jets events is often a leading systematic uncertainty for these studies, limiting their ability to further our understanding of the Higgs or potentially discover new physics.

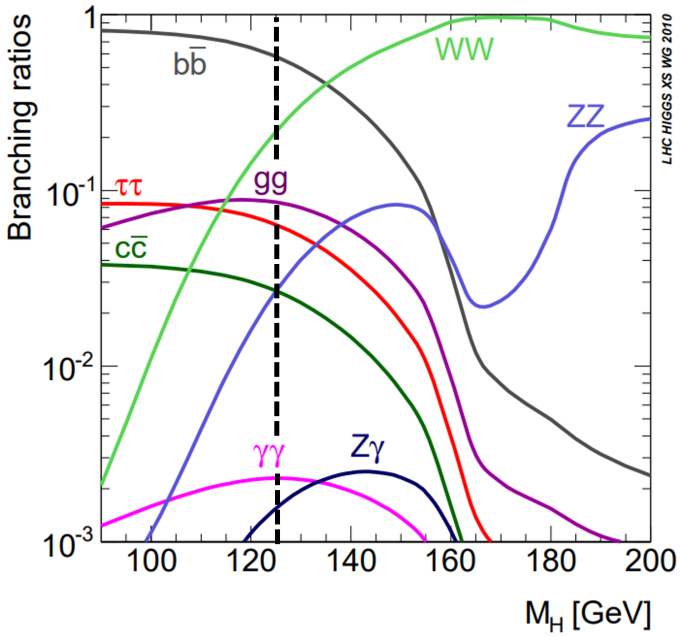

Figure 4.2: Branching ratios for the Higgs boson as a function of the mass of the Higgs, with $m_H = 125$ GeV denoted by the dotted line [101].

Compounding the fact that Z+HF jet events are a large background to other measurements and searches is the fact that current measurements of Z+HF jets events are imprecise. This thesis aims to perform a precision measurement of these events in order to provide experimental guidance on the improvement of Z+HF jets modelling. This will further our understanding of QCD, allow for a better understanding of the Higgs bosons, and increase the sensitivity of future analyses to BSM physics.

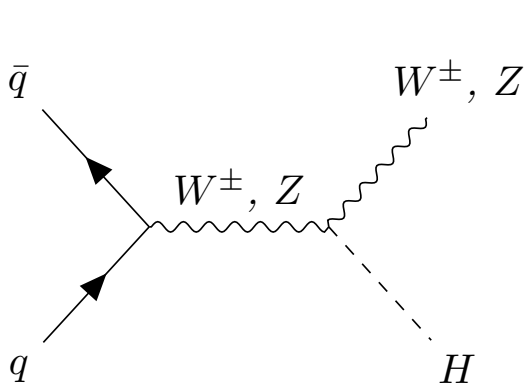

(a)

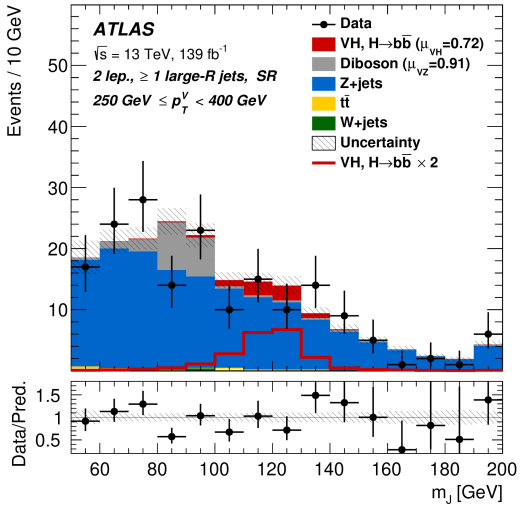

(b)

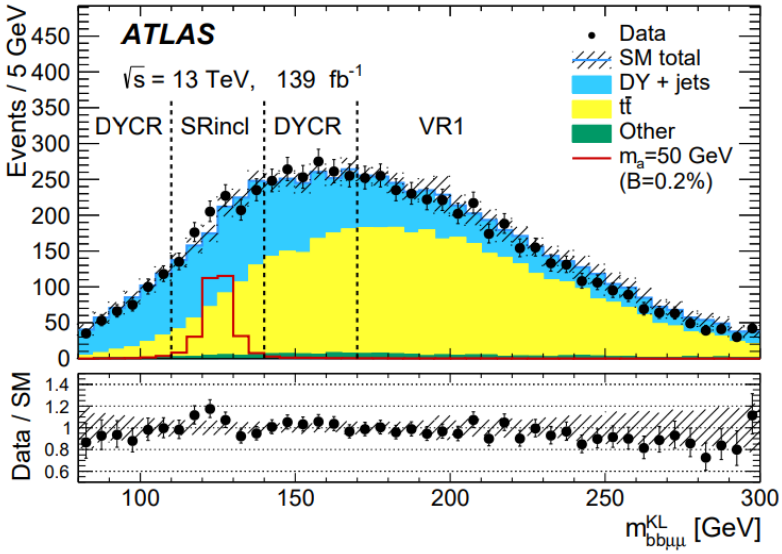

(c)

Figure 4.3: (a) Diagram for production of Higgs boson in association with a vector boson. (b) An example event count for a VH measurement, with Z+HF-jet background in dark blue (right) [102]. (c) An example event count of a search for a pseudoscalar a decaying into a $bb\mu\mu$ final state, with the Z+jet background given in light blue ("DY+jets") [4].

4.1.1 Four-flavor number scheme

One of the reasons why Z+b-jet events are not precisely understood is that they are sensitive to the differences in predictions made using two different calculation schemes, representing different points of view on the presence of b-quarks in the proton. One point of view is to assume that, because the b-quark is much heavier than the proton, it cannot be an initial-state quark in the matrix element of the calculation, and therefore the only way to have b-quarks in the final state is through an explicit gluon splitting in the high-Q^2 matrix element calculation. This gluon splitting can take place in either the initial state (Figure 4.4a) or the final state (Figure 4.4b) of the process. Because only the four lightest quark are included in the PDFs, this is referred to as the "four-flavor number scheme" (4FNS) [103].

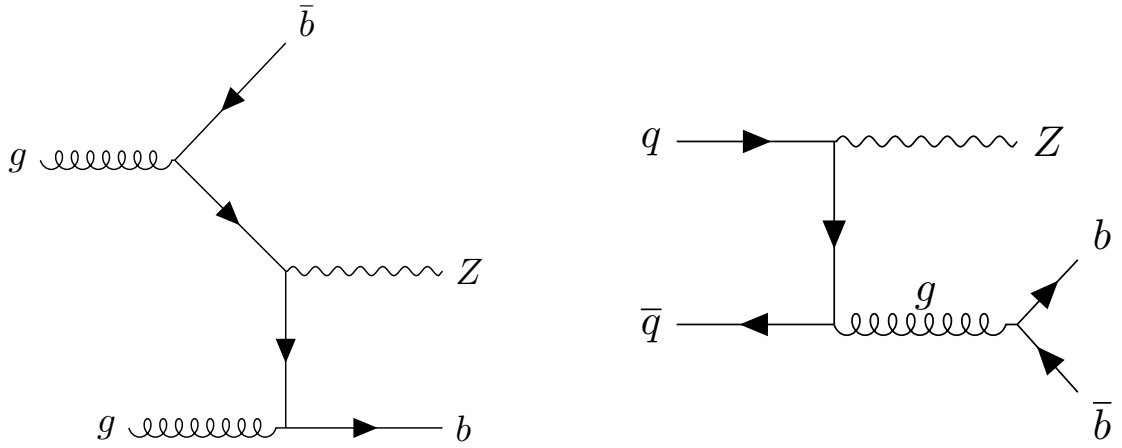

Figure 4.4: Four-flavor number scheme (4FNS) diagrams with the b-jets coming from a gluon splitting in the initial state (a) or the final state (b).

Because the b-quark mass sets a minimum energy scale for the gluon splitting, the matrix element calculation of this prediction, at a fixed order in perturbation theory, will involve terms which are proportional to

$$\alpha_s^n \log^{2n}\left(\frac{E^2}{m_b^2}\right), \tag{4.1}$$

where E is the energy scale of the final-state particles. As the energy scale of the hard scattering process increases, the log term becomes more and more significant. This means that, while $\alpha_s(m_Z) \approx 0.1$ and is usable in perturbation theory, $\alpha_s^n \log^{2n}(E^2/m_b^2)$ will become too large for perturbation theory to be valid. In short, at high energy, the gluon splitting into two b-quarks will make perturbation theory unusable (or at least converge more slowly, thereby requiring higher-order terms to obtain the same level of precision as at a lower energy scale). To fix this, an explicit "resummation" of the log terms must be performed. This is a computationally intensive

procedure which has to be done specifically for the phase space being measured. The final-state resummation can be done with parton showers but, as mentioned earlier, the parton showers make a soft approximation of the splitting functions, and the mass of the b-quark makes this a questionable approximation. Even if an NLO calculation is performed to include the first radiation, the initial-state radiation is still not resummed, and the perturbative expansion will be unusable or very imprecise at $n = 2$ and beyond.

Despite these problems, the 4FNS does have its advantages. It allows for explicit calculations of final states involving two b-jets, and it also accounts for the mass of the b-quark in the calculation, which is a relevant scale for a large portion of the phase space accessible at the LHC. Both of these are in contrast to the calculation scheme described in the next Section.

4.1.2 Five-flavor number scheme

The problems inherent to the 4FNS are of little importance if the general scale of the process of interest is low. If the scale is significantly higher than m_b, however, the 4FNS predictions are spoiled by the large logarithms in the perturbative expansion and cannot be easily resummed. In such a situation, an alternate calculation scheme, in which the b-quark is treated as massless but included in the initial quark state of the matrix element through a b-quark PDF, can be used. This is the "five-flavor number scheme" (5FNS) (Figure 4.5). In this scheme, the incident b-quark line in the matrix element calculation is automatically resummed to all orders because the gluon splitting is obtained via DGLAP evolution of the PDF, rather than putting it into the matrix element explicitly by hand. There are a few tradeoffs, however. The first is that, by setting $m_b = 0$, predictions in regions of phase space where m_b cannot be neglected are expected to be poor. In addition, NLO calculations have to avoid double-counting the initial-state gluon splitting when it enters the matrix element calculation as well. Finally, whereas final states including two b-jets from the initial gluon splitting are observable in the 4FNS (see Figure 4.4a), such final states are ignored in the 5FNS.

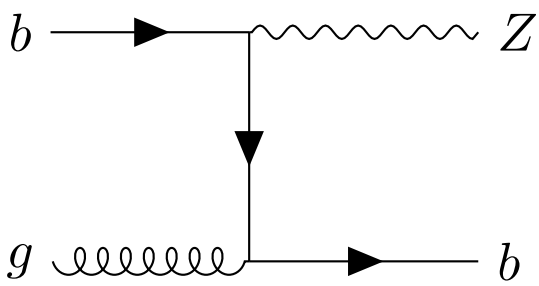

Figure 4.5: Z boson production in association with b-quarks in the 5FNS with the b-quark from the initial state (b).

The two b-jet final states of the 4FNS and the one b-jet final states of the 5FNS leave room

for measurement distinction. However, if the two b-jets of the 4FNS diagram in Figure 4.4b are collinear, or if one of the jets is too soft to leave a reconstructable signal in the detector, the event will appear to have only one one b-jet and therefore be mistaken as the 5FNS diagram of Figure 4.5. Predictions of Z+$b(b)$ events, compared to inclusive Z+jet predictions, are thus complicated significantly by the presence of the b-quark, forcing a choice in approach that must then be studied experimentally. Therefore, it is extremely important to perform measurements which could further our understanding of future Z+b-jet predictions.

It should be noted that the 5FNS production of a W instead of a Z is also possible, but in the W process a c-quark in the initial state is flavor-changed in the charged-current process into a b-quark, and that change of flavor is Cabbibo-suppressed. A measurement of W+b-jet events would therefore have to rely on 4FNS predictions. In addition, such a measurement is inherently less precise, due to the final state of one lepton plus missing energy being harder to distinguish from background signatures. It is therefore important to first understand Z+b-jet events and then to apply what is learned to W+b-jet events.

4.2 Z+c-jets

The c-quark, like the b-quark, is significantly more massive than the total mass of the proton's valence quarks, and so Z+c-jet predictions suffer from many of the same issues which plague Z+b-jet predictions. However, m_c is about three times smaller than m_b, and so the kinematic regions in which m_c can be neglected, and those in which resummation is necessary, differ from the similar regions with respect to m_b. Therefore, comparing c-jet and b-jet measurements in different kinematic regions can provide a more detailed picture of which prediction schemes are appropriate in those regions. For example, a recent ATLAS paper [104] measuring $\gamma + b$ and $\gamma + c$ events at $\sqrt{s} = 8$ TeV found that, for $\gamma + b$, only the 5FNS predictions agree with the data at higher energies ($125 < E_T^\gamma \lesssim 200$ GeV), while for $\gamma + c$ both the 4FNS and 5FNS predictions agree with the data within uncertainties; the authors posit that this is is due to the fact that the scale at which gluon splitting becomes the dominant production mode (over the Compton process $qg \to q\gamma$) is lower for $\gamma + b$ than for $\gamma + c$, and so the need for the resummable 5FNS prediction is needed for $\gamma + b$ at lower energies than for $\gamma + c$.

In addition, Z+c-jet measurements have the potential to address the longstanding question of whether there is an intrinsic heavy-quark component of the proton, as theorized by Brodsky et al. [105] several decades ago. As opposed to the extrinsic quark content, which comes from the QCD brehmsstrahlung and pair-production processes and which is only present at high-Q^2, intrinsic quark content refers to possible bound states of the proton beyond its typical three-quark

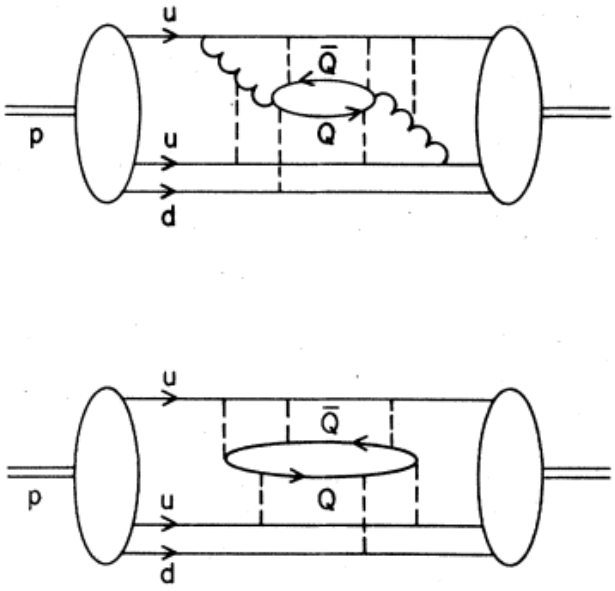

Figure 4.6: Diagrams which give rise to $|uudq\bar{q}\rangle$ bound states of the proton. Curly and dashed lines represent transverse and longitudinal gluons, respectively [106].

$|uud\rangle$ bound state. These other states, including $|uudc\bar{c}\rangle$ and $|uudb\bar{b}\rangle$, among others, contribute to the overall Fock state of the proton, and while they are rare they are expected to lead to a number of phenomenological consequences [106]. The states form due to the gluon exchange of the valence up and down quarks, as shown in Figure 4.6. As the charm quark ($m_c = 1.3$ GeV) is lighter than the bottom quark ($m_b = 4.2$ GeV), the effects of intrinsic charm bound states will be more easily seen through experiment; for example, an enhancement of the charm PDF is expected at $x \geq 0.1$ [107]. Explicitly measuring Z+c processes could therefore not only help understand the acceptability ranges of 4FNS and 5FNS predictions but also help constrain intrinsic charm models, which motivates the inclusion of Z+c-jet measurements in this thesis.

Chapter 5

Data and Simulated Monte Carlo Samples

5.1 ATLAS dataset description

This analysis uses data recorded by ATLAS during Run 2, from 2015 to 2018, at a center-of-mass energy of $\sqrt{s} = 13$ TeV. Only events contained in luminosity blocks (∼1-2 minutes of data taking, depending on the year) listed in the Good Run Lists (GRLs) [108] are used. The GRLs are XML files which contain information for each luminosity block specifying if the LHC was operating under stable beam conditions, if each subdetector was operational, if certain triggers were not being used, or if any number of problems were present. The events included in the GRLs used in this analysis[1] represent a total of 139 fb^{-1} of data taken during Run 2, as shown in Figure 2.5b.

Pileup reweighting, following the recommendations of the Jet/EtMiss working group [109], is performed to account for pp collisions occurring in addition to the collision of interest. In-time pileup, due to the additional pp collisions taking place in the same bunch crossing as the collision of interest, is shown for each of the data-taking years of Run 2 in Figure 5.1. Out-of-time pileup, due to collisions taking place in the bunch crossings proceeding and succeeding the bunch crossing containing the collision of interest, is also relevant for when detectors are sensitive to electronic signals integrated over longer than 25 ns, the spacing between the bunch crossings [110]. To account for these effects, simulated minimum-bias events are overlaid on every MC generator event; these account for the difference in data vs. MC of the average number of interactions per bunch crossing.

[1] 2015: `data15_13TeV.periodAllYear_DetStatus-v89-pro21-02_Unknown_PHYS_-StandardGRL_All_Good_25ns.xml`
2016: `data16_13TeV.periodAllYear_DetStatus-v89-pro21-01_DQDefects-00-02-04_PHYS_` — `StandardGRL_All_Good_25ns.xml`
2017: `data17_13TeV.periodAllYear_DetStatus-v99-pro22-01_Unknown_PHYS_` — `StandardGRL_All_Good_25ns_Triggerno17e33prim.xml`
2018: `data18_13TeV.periodAllYear_DetStatus-v102-pro22-04_Unknown_PHYS_` — `StandardGRL_All_Good_25ns_Triggerno17e33prim.xml`

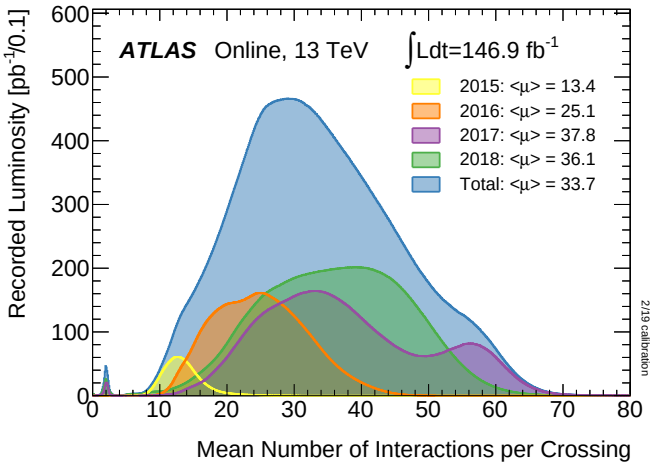

Figure 5.1: Luminosity as a function of the mean number of interactions per bunch crossing $<\mu>$ for 2015-2018 pp collision data taken at $\sqrt{s} = 13$ TeV [40].

5.2 Simulated data description

MC event generators model proton collisions to produce simulated events. Different MC campaigns correspond to the different years of data-taking: campaign MC16a corresponds to 2015-2016 data, campaign MC16d to 2017 data, and campaign MC16e to 2018 data. The different MC campaigns account for the different pileup distributions shown in Figure 5.1. MC samples were produced to estimate both the signal and background components of this analysis; the samples are described below and summarized in Table 5.1.

5.2.1 Signal samples

Three simulated signal samples of Z boson production in association with jets ($Z \to ll$ in Table 5.1) are used in this analysis. The first is produced by Sherpa [111] using the ATLAS Sherpa v.2.2.11 configuration [112]. This configuration performs explicit matrix element (ME) calculations of 0-2 partons at next-to-leading order (NLO) plus 1-3 more partons at leading order (LO). The ME are calcuated with the Comix [113] and OpenLoops [114] libraries. The PDF set used is the NNPDF 3.0 NNLO set [115], which is calculated at next-to-next-to-leading order (NNLO) in QCD with $\alpha_s = 0.018$. This means that the hard scattering processes ($\hat{\sigma}$ in Equation 1.26) are calculated at NNLO, and then are used to extract the PDFs from fits to experimental data providing σ; the data come from a number of different experiments, at different colliders, and so different NNLO calculations are performed with respect to the different experiments and their final states. Higher jet multiplicities are obtained using the Sherpa parton shower [116]. This parton shower model

is based on a Catani-Seymour dipole factorization, which separates the overall shower correction into individual m-parton loop corrections for each $(m+1)$-parton exclusive final state. The model is also is transverse momentum- (k_{T})-ordered; this is beneficial because it relates the cut-off for the parton shower to the minimum transvese momentum necessary to resolve partons [117, 118], and it simplifies the procedure to merge the parton shower correction to the ME calculation. Here the merging is performed using the MEPS@NLO prescription [119, 120]; the matching is governed by a set of tuned parameters developed by the Sherpa authors. Sherpa uses the cluster fragmentation model (see Section 1.6) for the hadronization of shower particles. In this sample, massless b and c quarks are included at ME level while massive b and c quarks are included in the parton shower; this makes this sample a 5FNS prediction.

A second $Z \to ll$ sample is generated using MadGraph5_aMC@NLO v2.2.2 [121–123], which includes 0-4 partons at LO in the ME calculation. The NNPDF 3.0 NNLO PDF set is used with $\alpha_s = 0.018$. The ME calculation is interfaced to Pythia 8.186 [124], which models the parton shower, hadronization, and underlying event (this interface is denoted MG5_aMC+Py8 in Table 5.1). The Pythia parton shower is a k_{T}-ordered dipole shower similar to Sherpa, but the Lund string model (see Section 1.6) is used instead for hadronization. The A14 tune [125] of Pythia, which uses Run 1 ATLAS data to optimize the multi-parton interaction parameters in Pythia 8, is used, along with the NNPDF 2.3 LO PDF set [126]. Overlap between the ME and parton shower emissions is removed at LO using the CKKW-L merging procedure [127, 128]. This sample is also based on the 5FNS prediction scheme.

Finally, an additional 5FNS theoretical prediction is generated using MadGraph5_aMC@NLO v2.6.5 [112], which includes up to three additional final-state partons at NLO. The NNPDF 3.1 NNLO PDF set [129, 130], supplemented with the LUXqed photon PDFs, is used at $\alpha_s = 0.118$. Showering is performed using Pythia 8.240 with the A14 tune and the NNPDF 2.3 LO PDF set. As opposed to the previous MadGraph sample, this sample uses the FxFx NLO matching scheme [131].

For the remainder of this document, these three signal samples will be referred to simply as the Sherpa, MadGraph, and FxFx samples, respectively. The Sherpa sample is used to unfold the data (Chapter 8); the MadGraph sample is used in conjunction with the Sherpa sample to determine modelling uncertainties (Chapter 9); and all three samples will provide theoretical predictions which will be compared to the unfolded measurements (Chapter 10).

5.2.2 Background samples

Several MC samples were generated to estimate the background components to the signal data, and are subtracted from the the data prior to unfolding. The $t\bar{t}$ background was produced us-

Process	Generator	Order of MC pQCD
$Z \to ll\,(l = e, \mu)$	Sherpa v.2.2.11	0-2p NLO 3-4-5p LO
$Z \to ll\,(l = e, \mu)$	MG5_aMC+Py8 (CKKW-L merging)	0-4p LO
$Z \to ll\,(l = e, \mu)$	MG5_aMC+Py8 (FxFx merging)	0-3p NLO
$t\bar{t}$	Powheg+Py8	NLO
VH	Powheg+Py8	NLO
Single top ($s/t/tW$-channel)	Powheg+Py8	NLO
$VV \to \ell\ell/\ell\nu/\nu\nu + q\bar{q}$	Sherpa v.2.2.1	NLO
$W \to l\nu$	Sherpa v.2.2.1	NLO
$Z \to \tau\tau$	Sherpa v.2.2.1	NLO

Table 5.1: Summary of the MC generators used to produce the various background processes. "Order of pQCD" refers to the order in QCD of the matrix element calculation of the MC simulation. Np denotes the number of real parton emissions.

ing the PowhegBox v2 [132, 133] generator at NLO with the NNPDF 3.0 NNLO PDF set. This sample was interfaced to Pythia 8.230 using the A14 tune and the NNPDF 2.3 LO PDF set (Powheg+Py8 in Table 5.1). The sample was then normalized to the NNLO cross-section, including the resummation of next-to-next-to-leading logarithmic (NNLL) soft gluon terms calculated with Top++2.0 [134]. Powheg+Py8 with the NNPDF 3.0 NNLO PDF set was also used to simulate the production of vector bosons W, Z in association with a Higgs boson (VH).

Single top quark production, in association with W bosons (tW) as well as in the s- and t-channels, was also simulated using PowhegBox v2 interfaced with Pythia 8.230 at NLO using the NNPDF 3.0 NNLO PDF set. To remove overlap with the $t\bar{t}$ sample, the diagram-removal [135] procedure was used.

Semi-leptonic diboson (VV) final states, where one W or Z boson decays leptonically and the other decays hadronically, were simulated using Sherpa v.2.2.1 [136], with NLO MEs for up to one parton emission and LO MEs for up to three parton emissions calculated using the NNPDF 3.0 NNLO PDF set. The ME calculations are matched to the Sherpa parton shower using the MEPS@NLO prescription. Sherpa v.2.2.1 and the NNPDF 3.0 NNLO PDF set were also used to generate the V+jets ($W \to l\nu$ and $Z \to \tau\tau$) background samples.

The minimum-bias events used for pileup reweighting included the single-, double-, and non-diffractive pp processes of Pythia 8.186 using the A3 tune [137]. For all MadGraph and PowhegBox samples (see below), the EvtGen v1.6.0 program was used to simulate the decay of b and c hadrons. Finally, all simulated events were processed with the GEANT4 simulator of the ATLAS detector [79] and reconstructed with the same software and algorithms as were used to reconstruct the data.

Chapter 6

Event Selection

This Chapter outlines the selections used in this analysis to define the final state in which the measurement is performed. The purpose of the selections is to balance background removal (signal purity) with signal efficiency such that the residual background and efficiency uncertainties are minimized. Included in this Chapter are estimates of the background after the selections are applied, including comparisons between MC and data.

6.1 Event pre-selection

As mentioned in Chapter 5, only data events which are included in the ATLAS Good Run Lists are included in this measurement. They are also subject to the LooseBad jet cleaning requirement ("Looser" in [138]), which removes jets formed by beam-induced background particles, primarily muons, interacting with the calorimeters. In addition, each event has to pass an unprescaled trigger chain according to the data period and lepton flavor to be analyzed. The triggers are listed in Table 6.1, with the chain being the logical OR of the triggers in a given cell. Triggers with higher p_{T} thresholds generally have looser isolation requirements so that highly-boosted Z decays, in which the two leptons are almost collinear, are not removed.

Years	Muon Channel	Eletron Channel
2015	HLT_mu20_iloose_L1MU15 HLT_mu50	HLT_e24_lhmedium_L1EM20VH HLT_e60_lhmedium HLT_e120_lhloose
2016 - 2018	HLT_mu26_ivarmedium HLT_mu50	HLT_e26_lhtight_nod0_ivarloose HLT_e60_lhmedium_nod0 HLT_e140_lhloose_nod0 HLT_e300_etcut

Table 6.1: The lowest unprescaled triggers for muon and electron channels, which form the trigger chains for the respective years and lepton flavors.

6.2 Object selections

The selections applied to the different physics objects used in this analysis — muons, electrons, jets, and missing transverse momentum — are described in the following subsections.

6.2.1 Muons

Objects reconstructed in the region $|\eta| < 2.5$ as CB or ME muons are identified using tracks matched between the MS and ID, as described in Section 3.3.1. The tracks are associated to the primary vertex using constraints on the transverse ($|d_0|/\sigma(d_0) < 3.0$) and longitudinal ($|z_0 \sin(\theta)| < 0.5$) impact parameters. The muon candidates are identified using the Medium quality working point (see Chapter 3.3.1) and must pass the `FCTight` isolation working point. The p_{T} of the muon candidates must be larger than 27 GeV.

6.2.2 Electrons

Electron candidates are reconstructed in the region $\eta < 2.47$, excluding the crack region $1.37 < |\eta| < 1.52$ between the barrel and endcap ECAL cryostats, from energy deposits in the ECAL matched to tracks in the ID. The tracks are subject to impact parameter ($|d_0|/\sigma(d_0) < 3.0$ and $|z_0 \sin(\theta)| < 0.5$) constraints and are identified using the Tight quality working point (see Chapter 3.4.2). The electron cancidates must pass the `FCTight_FixedRad` isolation working point and have $p_{\mathrm{T}} > 27$ GeV.

Leptons passing the above selection criteria, which are summarized in Table 6.2, are referred to as "good leptons." Only good leptons are considered in the measurements of this analysis.

Lepton Definition	**Muon Channel Criteria**	**Electron Channel Criteria**	**Applied at Truth Level?**
Impact Parameter	$z_0 \sin\theta < 0.5$ $\lvert d_0\rvert/\sigma(d_0) < 3$	$z_0 \sin\theta < 0.5$ $\lvert d_0\rvert/\sigma(d_0) < 5$	
Identification	Medium	Tight	
Isolation	`FCTight`	`FCTight_FixedRad`	
Kinematics	$p_{\mathrm{T}} > 27$ GeV, $\lvert\eta\rvert < 2.5$	$p_{\mathrm{T}} > 27$ GeV, $\lvert\eta\rvert < 2.47$ (ignore crack electrons)	Yes

Table 6.2: Criteria used to define and select good leptons. The rightmost column notes which of these detector-level selections are applied to the particle-level ("truth") objects to produce the theoretical predictions.

6.2.3 Jets

Jets reconstructed from noise-suppressed topo-clusters in the calorimeter are required to have $p_{\mathrm{T}} > 20$ GeV and $|\eta| < 2.5$. To reduce the effects of pileup, jets are required to pass the Tight

working point of the Jet Vertex Tagger (JVT) [109, 139], which is 96% efficient. This requirement ensures that a significant fraction of the tracks have origins which are compatible with the primary vertex. Finally, since leptons and jets are reconstructed independently, a single object can be reconstructed as both a lepton and a jet; therefore, an overlap removal procedure [140] is performed. Muons are discarded if $\Delta R(\text{jet}, \text{muon}) < 0.4$ and the number of tracks associated to the jet is larger than 2, removing ambiguity between muons and jets containing semi-leptonic decays. Meanwhile, electrons are discarded if $0.2 < \Delta R(\text{jet}, \text{electron}) < 0.4$ are removed, while jets are discarded if $\Delta R(\text{jet}, \text{electron}) < 0.2$; the former removal handles electron reconstruction bias as a result of the nearby hadronic activity, while the latter removal ensures electrons aren't duplicated as jets. Only jets passing all of the above jets selections, as well as passing the requirement that $\Delta R(\text{jet}, \text{any good lepton}) < 0.4$, are considered in this analysis.

Jets which pass all of the above selection criteria are referred to as "good jets." Only good jets are considered in this analysis.

6.2.4 *b*-jets

To be considered a *b*-jet, a jet must first pass all of the selections described in the previous subsection. They are then considered "*b*-tagged," i.e. to have originated from a *b*-quark, if they pass the DL1r operational working point (described in Section 3.5.5) corresponding to an 85% *b*-jet tagging efficiency. Good jets which are also *b*-tagged according to this definition (thus passing all of the selections summarized in Table 6.3) are referred to as "good *b*-jets." Only good *b*-jets are measured in this analysis.

Jet Definition	**Muon Channel Criteria**	**Electron Channel Criteria**	**Applied at Truth Level?**
Kinematics	$p_{\text{T}} > 20$ GeV and $\lvert y \rvert < 2.5$		Yes
Overlap Removal	Discard jet if $\Delta R(\text{jet}, e) < 0.2$ Discard e if $0.2 < \Delta R(\text{jet}, e) < 0.4$	Discard μ if $\Delta R(\text{jet}, \mu) < 0.4$	Yes
	$\Delta R(\text{jet}, \text{any signal lepton}) < 0.4$		Yes
Jet Vertex Tagger	Tight		
b-tagging	DL1r 85% working point		Truth b or c

Table 6.3: Criteria used to define and select the good (*b*-)jets which are measured in this analysis. The rightmost column notes which of these detector-level selections are applied to the particle-level ("truth") objects to produce the theoretical predictions. At truth level, *b*-jet and *c*-jet events are selected using their truth flavor information rather than the detector-level *b*-tagging discriminant working point.

6.3 Z final state selection

Events containing a Z boson candidate are selected by requiring that the event contain exactly two good leptons of the same flavor and opposite charge. At least one of the leptons must be matched to the trigger chain used to preselect the events (Table 6.1); this is ensured by requiring that the lepton be within $\Delta R < 0.4$ of object which fired the trigger chain and have a p_{T} greater than at least one of the triggers in the chain. Moreover, the invariant mass of the dilepton pair has to be within $76 < m_{ll} < 106$ GeV, which is centered on m_Z. In order to reject the large top background, events with $p_T^Z < 150$ GeV are discarded if the missing transverse energy (MET) > 60 GeV; this selection is taken from a similar analysis [5] performed using 2015-2016 data, in which the selection reduced the $t\bar{t}$ background by about 55% while reducing the signal by only about 5%. Finally, the Z+HF jet final state is selected by requiring at least one good b-jet.

Table 6.4 summarizes these final-state selections. These and the above object selections were checked using numerous cutflow comparisons: for each analysis framework being used by different members of the analysis team, the number of events after each selection, in data and MC, was obtained and compared to ensure consistency across the frameworks.

Selection	**Criteria**	**Applied at Truth Level?**
Leptons	Exactly 2 good leptons of same flavor, opposite charge	Yes
Trigger matching	≥ 1 good lepton trigger matched to preselection chain	
Invariant Mass	$76 < m_Z < 106$ GeV	Yes
Missing Transverse Energy (MET)	Discard event if: MET > 60 GeV and $p_{\mathrm{T}}^Z < 150$ GeV	
Jets	≥ 1 good b-jet	Yes

Table 6.4: Selections used in this analysis, which are applied to the objects as defined in the above sections.

6.4 Efficiency correction scale factors

The many selections described above are necessary to obtain a clean sample of Z+HF jets events, but with each selection comes the possibility that a signal event will not be measured because it does not pass the selection. To account for each of these selection inefficiencies, efficiency corrections must be applied to data. However, the process of unfolding detector resolution effects (see Chapter 8) also includes such an efficiency correction, calculated from MC in the phase space of the measurement. Therefore, to account for efficiency differences between data and MC, efficiency

scale factors (SFs) of the form

$$SF^x = \frac{\varepsilon^x_{\text{Data}}}{\varepsilon^x_{\text{MC}}} \tag{6.1}$$

are applied to the simulated events for each selection effect x. Efficiency SFs are produced by the various Combined Performance groups of ATLAS by estimating the data and MC efficiencies in a particular phase space. and account for the following effects:

- **Single lepton triggers:** Data and MC efficiencies are measured for each trigger in Table 6.1, for each isolation and identification working point, and are used to created a trigger SF in bins of $\eta - \phi$ ($p_{\text{T}} - \eta$) for muons (electrons). For a given trigger with efficiency $\varepsilon^{l_i,\text{Trig}}$ for the first (largest p_{T}) lepton l_i, the overall trigger efficiency is one minus the product of inefficiencies $(1 - \varepsilon^{l_i,\text{Trig}})$ for each lepton. For this analysis, with two signal leptons, the trigger scale factor is given by

$$SF^{\text{Trig}} = \frac{\varepsilon^{\text{Trig}}_{\text{Data}}}{\varepsilon^{\text{Trig}}_{\text{MC}}} = \frac{1-(1-\varepsilon^{l1,\text{Trig}}_{\text{Data}})(1-\varepsilon^{l2,\text{Trig}}_{\text{Data}})}{1-(1-\varepsilon^{l1,\text{Trig}}_{\text{MC}})(1-\varepsilon^{l2,\text{Trig}}_{\text{MC}})}.$$

  Example muon trigger efficiencies and scale factors can be seen in Figure 3.7, while similar examples for electrons can be seen below in Figure 6.1.

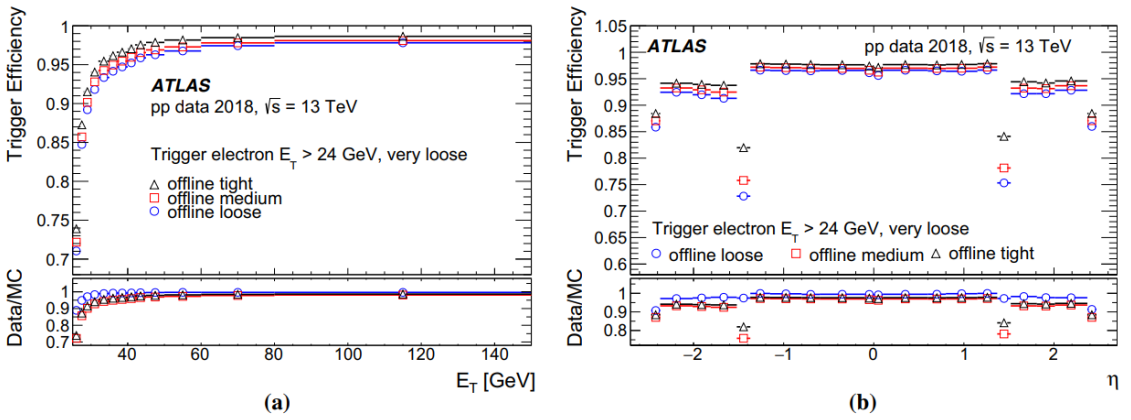

Figure 6.1: Efficiency of the `e24_lhvloose_nod0` trigger for Tight, Medium, and Loose electrons as a function of the offline electron E_{T} (a) and η (b) [86].

- **Lepton selections:** For each muon, the identification working point, track-to-vertex association, and isolation efficiencies are corrected using SFs. Similarly, for each electron, the electron reconstruction, identification, and isolation efficiencies are corrected using SFs. The three respective SFs are multiplied together and applied for each signal lepton. For example,

for the first (largest p_{T}) lepton $l1$, the overall lepton SF is

$$SF^{l1} = SF^{l1,\mathrm{ID}} \times SF^{l1,\mathrm{reco}} \times SF^{l1,\mathrm{iso}}. \tag{6.2}$$

- **JVT:** A JVT SF is determined for each jet in an event, and the product of these SFs is applied to reweight the event, to correct for the efficiency of the JVT Tight requirement.
- ***b*-tagging:** For each jet passing the 85% efficiency working point of the DL1r tagger, a *b*-tagging SF is determined; the product of these SFs is applied to the event to correct for the efficiency of the *b*-tagging selection.

6.5 ≥1 jet MC-data comparisons

The MC samples, once prepared with the above selections, are used to correct the data through the efficiency corrections described above, MC-derived estimates of the larger backgrounds (see Chapter 7), and unfolding (Chapter 8). To make sure that the MC samples are appropriate to use to obtain these corrections (and their uncertainties), it is important to check that a comparison between MC and data does not display unmanageable discrepancies. The comparison is made using observables which have been carefully and precisely measured in other ATLAS/CMS analyses.

Because this is a measurement of SM physics, rather than a search relying on MC predictions to investigate new physics, the MC-to-data agreement does not have to be perfect. Even if the MC is mismodelled and only agrees with the data up to a 20% discrepancy, that 20% discrepancy becomes an uncertainty of only a few percent in the final measurement. This is primarily because a) even if the backgrounds are mismodelled, they are only a small fraction of the total sample (except for the light-jet background discussed in Chapter 7); b) the efficiency scale factors are close to unity in all $\eta - \phi/p_{\mathrm{T}} - \eta$ bins, and so lepton kinematics mismodelling will result in a percent-level uncertainty; and c) the unfolding corrections use MC truth and reco to form the migration matrices, and so in those migration matrices the mismodelling effects mostly cancel out.

Several MC/data comparisons were made in the region of phase space corresponding to the selections of Table 6.4 minus the requirement that the selected jet be *b*-tagged. The observables chosen for these comparisons have been measured numerous times by different experiments and at varying energy scales, etc. The MC-data agreement shown in the below plots is within the range of what we would expect given those previous measurements, which provides some confidence that the MC samples have been prepared adequately for this analysis.

Figure 6.2 shows the p_{T} distribution of the Z boson, which drops off following a power law as the p_{T} increases, as expected from differential cross-section calculations. MadGraph tends to

over-predict the p_{T} distribution observed in data, while SHERPA underpredicts at $\gtrsim 100$ GeV, but both generators are in agreement with the data up to 20% in entire kinematic region of interest. This suggests that, given the uncertainty estimates described above, any potential mismodelling of the Z+jets processes of interest in this analysis will only have a small effect on the precision of the final measurement. Figure 6.3 shows similar features in the predictions of the p_{T} of the leading jet. Leading jet η is shown in Figure 6.4, which is flat as expected. Finally, Figure 6.5 shows the comparison between the observed and expected number of jets reconstructed in a given event. The agreement between SHERPA and the data is better here than for the other observables, with a difference of $\lesssim 5\%$ for most jet multiplicities, while MADGRAPH under-predicts higher jet multiplicities by up to 20%. Overall, the agreement for these observables does not vary between the electron and muon channels, which suggests that there is little sensitivity to any differences in the modelling of different flavors of leptons.

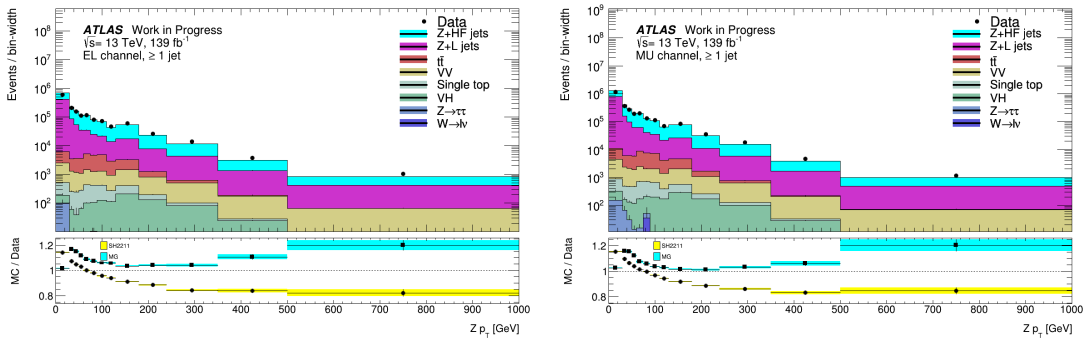

(a) (b)

Figure 6.2: Stacked histograms of the p_{T} of the Z boson, obtained from 2015-2018 data and from SHERPA 2.2.11 MC samples, in the (a) electron and (b) muon channels. The ratio plot shows the MC/data ratio, including the ratio using the MADGRAPH MC samples instead.

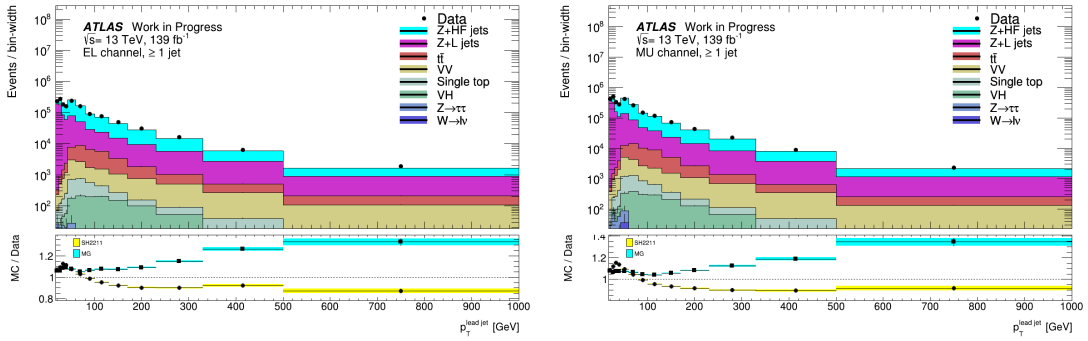

(a) (b)

Figure 6.3: Stacked histograms of the leading jet p_{T}, obtained from 2015-2018 data and from SHERPA 2.2.11 MC samples, in the (a) electron and (b) muon channels. The ratio plot shows the MC/data ratio, including the ratio using the MADGRAPH MC samples instead.

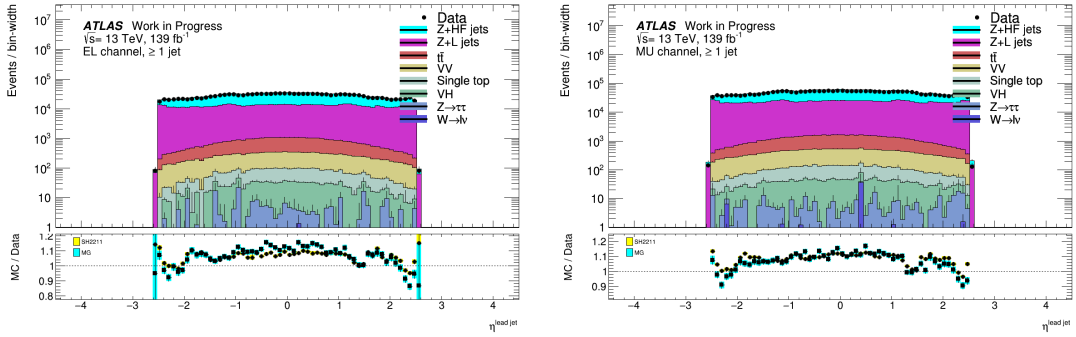

(a) (b)

Figure 6.4: Stacked histograms of the leading jet pseudorapidity η, obtained from 2015-2018 data and from Sherpa 2.2.11 MC samples, in the (a) electron and (b) muon channels. The ratio plot shows the MC/data ratio, including the ratio using the MadGraph MC samples instead.

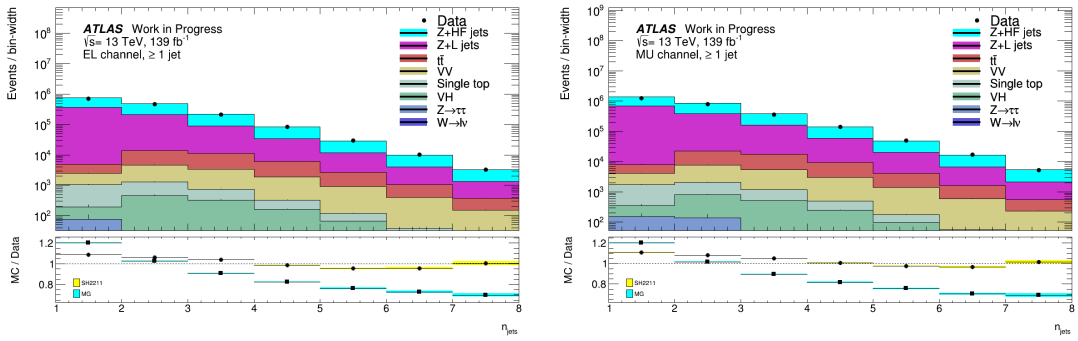

(a) (b)

Figure 6.5: Stacked histograms of the number of reconstructed jets, obtained from 2015-2018 data and from Sherpa 2.2.11 MC samples, in the (a) electron and (b) muon channels. The ratio plot shows the MC/data ratio, including the ratio using the MadGraph MC samples instead.

6.6 ≥1 b-jet signal region

This analysis is performed on events with at least one b-jet, with a b-jet defined as a jet passing all of the jet selections listed in Table 6.4. This detector-level selection will include particle-level b- and c-jets, which will be unfolded simultaneously, as discussed in Chapter 8.

Table 6.5 lists the signal and background percentages in this signal region, in both the electron and muon channels, using Sherpa 2.2.11. The largest background is Z+light-jet events; despite the use of b-tagging to suppress this background, the cross-section is $\approx 50\times$ ($\approx 10\times$) larger than the cross-section for b-jets (c-jets). To avoid suffering from large modelling uncertainties on this background, the light-jet background is fit to the data, as described in Chapter 7. The second largest background is $t\bar{t}$: each top quark decays into a b and a W boson, and the W bosons decay semi-leptonically, and so the final state includes two b-jets and two charged leptons.

Several MC-data comparisons were also performed in this region to validate the MC samples before unfolding, with the Sherpa samples being used in the unfolding procedure. Figure 6.6

Process	**Percentage (%)**		**Weighted event count**	
	Electron	**Muon**	**Electron**	**Muon**
Z+HF jets	55.02	53.29	881524	1517111
Z+Light jets	42.47	44.46	680509	1265683
$t\bar{t}$	1.64	1.45	26225	41314
Diboson	0.65	0.60	10461	16991
Single top ($s/t/tW$-channel)	< 0.5	< 0.5	2349	3664
VH			1122	1715
$Z \to \tau\tau$			154	300
$W \to l\nu$			Negligible	
Total MC Events (Observed Data)			1602327 (1509583)	2846835 (2628258)

Table 6.5: Percentage of signal and background events in the electon and muon channels using Sherpa 2.2.11. The final row gives the total number of events in MC (data).

shows the p_{T} distribution of the leading b-tagged jet: a more rapid departure of the MadGraph prediction from the Sherpa prediction than in Figure 6.3 is seen. The gap is not as large for lepton p_{T}, shown in Figure 6.7. Figure 6.8 shows lepton η, where both Sherpa and MadGraph are in good agreement with the data. The invariant mass m_Z MC-data agreement for both Sherpa and MadGraph is $\approx 5-10\%$ too large; this is essentially an issue of cross-section normalization, as this over-prediction is also seen in the low-p_{T} bins of the leading jet p_{T} distributions. Finally, the MC predictions of $x_F(\text{tagged jet}) \equiv 2p_{\mathrm{T}}\sinh(y)/\sqrt{s}$ differ more significantly in the electron channel than in the muon channel, as shown in Figure 6.10; low statistics in the final bin cause the uncertainty on the predictions to be quite large.

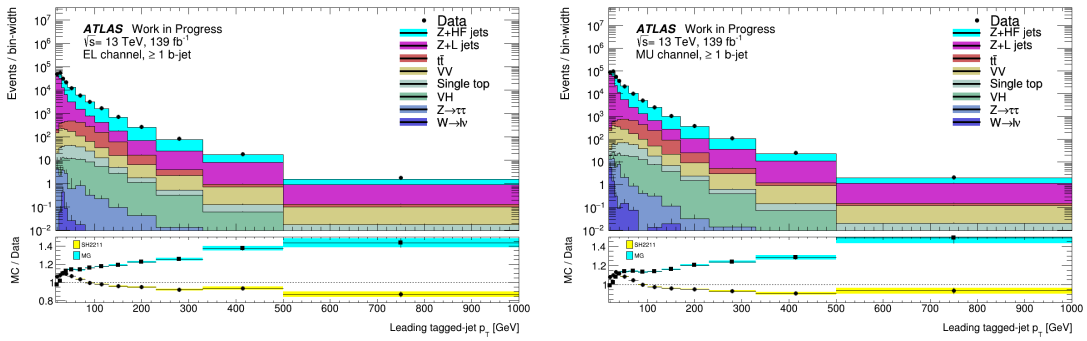

(a) (b)

Figure 6.6: Stacked histograms of the leading b-jet p_{T}, obtained from 2015-2018 data and from Sherpa 2.2.11 MC samples, in the (a) electron and (b) muon channels. The ratio plot shows the MC/data ratio, including the ratio using the MadGraph MC samples instead.

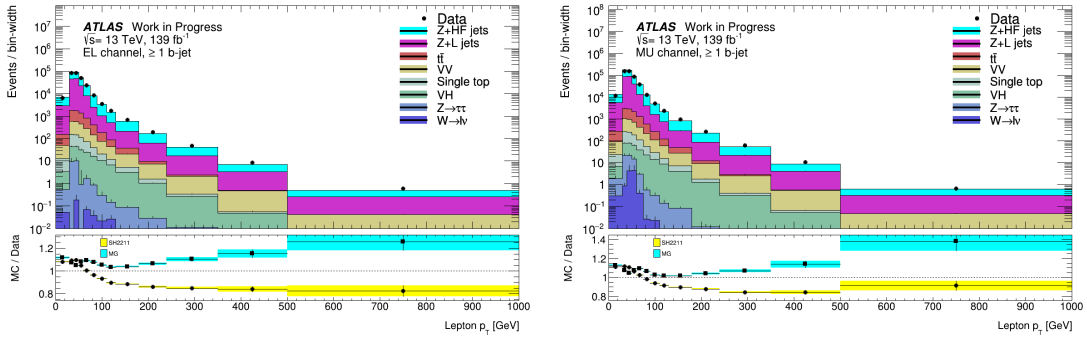

(a) (b)

Figure 6.7: Stacked histograms of the lepton p_{T}, obtained from 2015-2018 data and from Sherpa 2.2.11 MC samples, in the (a) electron and (b) muon channels. The ratio plot shows the MC/data ratio, including the ratio using the MadGraph MC samples instead.

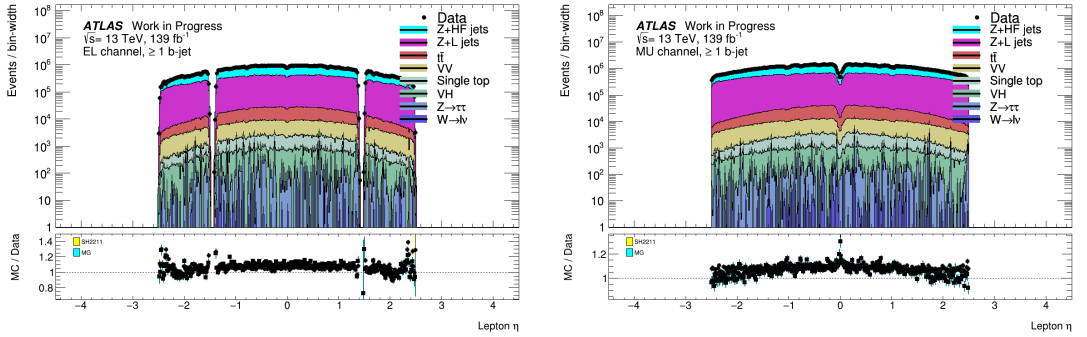

(a) (b)

Figure 6.8: Stacked histograms of the lepton pseudorapidity η, obtained from 2015-2018 data and from Sherpa 2.2.11 MC samples, in the (a) electron and (b) muon channels. The ratio plot shows the MC/data ratio, including the ratio using the MadGraph MC samples instead.

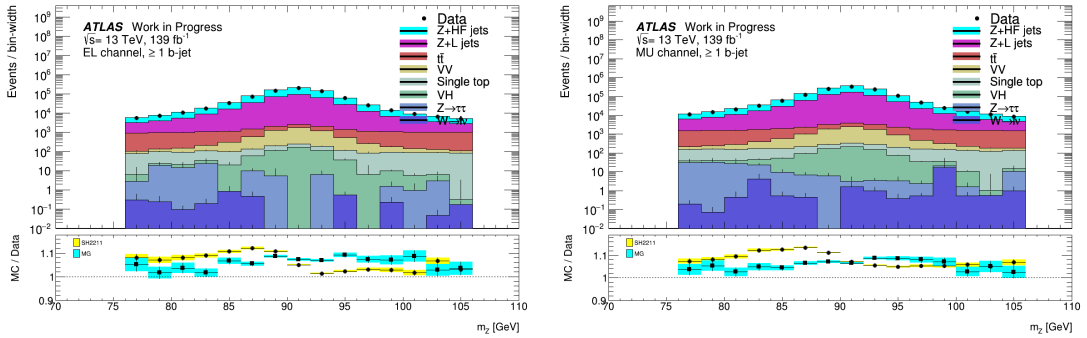

(a) (b)

Figure 6.9: Stacked histograms of the Z boson mass m_Z, obtained from 2015-2018 data and from Sherpa 2.2.11 MC samples, in the (a) electron and (b) muon channels. The ratio plot shows the MC/data ratio, including the ratio using the MadGraph MC samples instead.

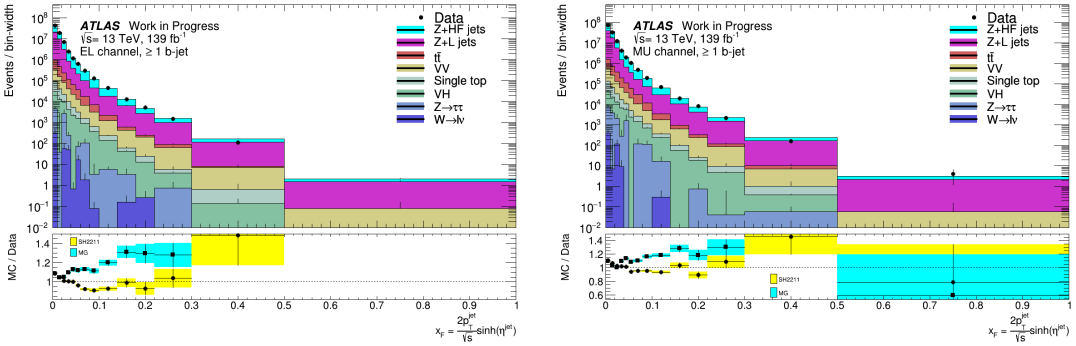

(a) (b)

Figure 6.10: Stacked histograms of x_F(tagged jet), obtained from 2015-2018 data and from SHERPA 2.2.11 MC samples, in the (a) electron and (b) muon channels. The ratio plot shows the MC/data ratio, including the ratio using the MADGRAPH MC samples instead.

6.7 Observables

Four observables were measured in this analysis in the $\geq$1 b-tagged jet signal region. Here is a brief description of why these observables were chosen:

- [**Leading tagged-jet $p_{\mathbf{T}}$**] serves as a test of the modelling of gluon splitting into b-jets, and therefore provides a good measure of the scalability of the 4FNS prediction. As mentioned in Chapter 4, the matrix element calculations in 4FNS predictions are pertubative expansions involving terms of the form $\alpha_s^n \log^{2n}(E^2/m_b^2)$, which are too large for resummation at higher energies. Leading b-jet p_{T} tests a range of final-state energies, and therefore can describe at which energies the 4FNS begins to break down.

- [Z $p_{\mathbf{T}}$] (alternatively written as p_{T}^Z) serves a similar role as an observable to leading b-jet p_{T}, since jet emission recoils against the Z boson, but it is determined via lepton, rather than jet, reconstruction. In general, leptons are the better and more precisely understood reconstructed physics objects, and so p_{T}^Z will have smaller uncertainties and be easier to unfold. However, p_{T}^Z is less sensitive to two b-jet events in which the jet emissions recoil against each other, leaving the Z more or less at rest.

- [ΔR(Z, **leading tagged-jet**)] is more sensitive to the collinearity of the b-jet splitting than the leading b-jet p_{T} (which in turn is more sensitive to the energy scale of the splitting).

- [**Leading tagged-jet x_F**] is included as a means to test the sensitivity of this analysis to intrinsic charm; as noted in [107], at $x_F \geq 0.1$, intrinsic charm would increase the c-quark PDF contribution to processes which include c-jets in the final state.

These four observables are by no means the only observables of physics interest. For example:

- [**Leading** b**-jet** η] could be included for similar reasons as the leading b-jet p_{T}: more energetic b-jets should be more forward (higher in η), and so a degradation of 4FNS predictions is predicted as η increases (though to a lesser extent than for leading b-jet p_{T}, as leading jet η is expected to be fairly flat).

- The [**number of** b**-jets**] can be measured in order to obtain the exclusive b-jet cross-section for a given b-jet multiplicity.

It should be noted in particular that two b-jet observables are not included in this analysis. As will be discussed in Chapter 8, the correlation between b- and c-jets is unfolded by constructing a response matrix in the b-tagging quantile bins; this response matrix would have to be expanded to account for the combined quantiles of the two tagged jets, and this implementation has not been developed. Some examples of two b-jet observables which should be include in future studies include:

- [**Sub-leading** b**-jet** $p_{\mathbf{T}}$], which should contrast the 4FNS and 5FNS predictions because, as mentioned earlier, two b-jet final state diagrams are only calculated in 4FNS predictions; in the 5FNS, a second b-jet can only be introduced via parton showering (which is also possible in the 4FNS); and

- [ΔR**(Leading, sub-leading** b**-jets)**], which would potentially be better handled by 4FNS predictions because a full gluon splitting calculation (with nonzero m_b) is performed without the assumptions of collinearity inherent to parton showering.

The four observables measured in this analysis provide a strong handle on the QCD phenomena which are the focus of this Z+HF measurement, as described in Chapter 5. In the next Chapter, the estimation of the light-jet background is presented.

Chapter 7

Jet Flavor Fit

The largest background by far in the signal region of this analysis, as shown in Table 6.5, is Z+light-jets. This Chapter provides an overview of how this background component, whose distribution is modelled using Monte Carlo (MC) for each observable, is fit to the data to obtain a data-driven estimate of the normalization of this background. The normalized Z+light-jets component is then subtracted from the data before unfolding.

The method used to fit the light flavor fraction is presented in the context of an overall flavor fit, which can be performed to fit the b-, c-, and light-jet flavor fractions of the MC signal sample. Similar flavor fits were used to extract the Z+b-jet components signal in previous ATLAS and CMS measurements. However, as discussed in the next Chapter, this method of separately fitting the b- and c-jet contributions is discarded in favor of a combined unfolding technique. The flavor fit method is therefore only used to estimate the light-jet background contribution to the Z+jets events.

7.1 Flavor fit using PCBT quantile bins

Initially, the light-jet contribution was fit to the data at the same time as the b-jet and c-jet contributions using the distribution of events in the PCBT quantile bins introduced in Section 3.5.5. As discussed there, the PCBT bins describe different b-tagging efficiencies, corresponding to [0, 60]%, [60, 70]%, [70, 77]%, [77, 85]%, and [85, 100]% [100]. This analysis selects b-jets passing the 85% working point, which can end up in any of these quantile bins except the [100, 85]% bin, which is therefore ignored. The PCBT quantile distributions then look like Figure 7.1a, where Z+jets events have been split according to their truth flavor information into b-/c-/light-jets. As expected, the leftmost bin ([77, 85]% b-tagging efficiency) is dominated by light-jets, as the higher b-tagging efficiency results in lower b-jet purity, while the rightmost bin ([0, 60]%) is primarily

b-jets. A maximum likelihood fit is then performed, using RooFit v.6.22.00 [141], where each of the three flavor fractions is allowed a floating normalization while the backgrounds are held fixed. The normalization scale factors (SFs) — the post-fit template normalization divided by the pre-fit normalization — are applied to the pre-fit MC flavor distributions to maximize the agreement with data, resulting in Figure 7.1b.

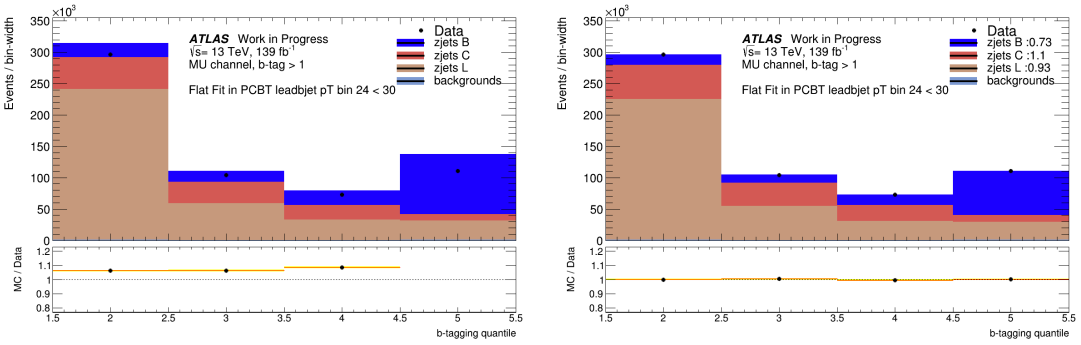

(a) Pre-fit (b) Post-fit

Figure 7.1: PCBT quantile distribution of events with $24 < p_{\mathrm{T}} < 30$ GeV pre-fit (a) and post-fit (b) using Sherpa 2.2.11.

This PCBT fit is performed in each bin of the observable; for example, Figure 7.1 shows the fit being performed in the $24 < p_{\mathrm{T}} < 30$ GeV bin of leading tagged-jet p_{T}. Doing so results in a set of three scale factors for each bin (Figure 7.2). This binwise PCBT fit of the observable results in an excellent MC-data agreement, as shown in Figure 7.3. This is not surprising, given that the bins are fit individually; however, it relies on the PCBT templates, which are discriminated purely using detector effects, and the templates are not the best tool for separating each individual flavor.

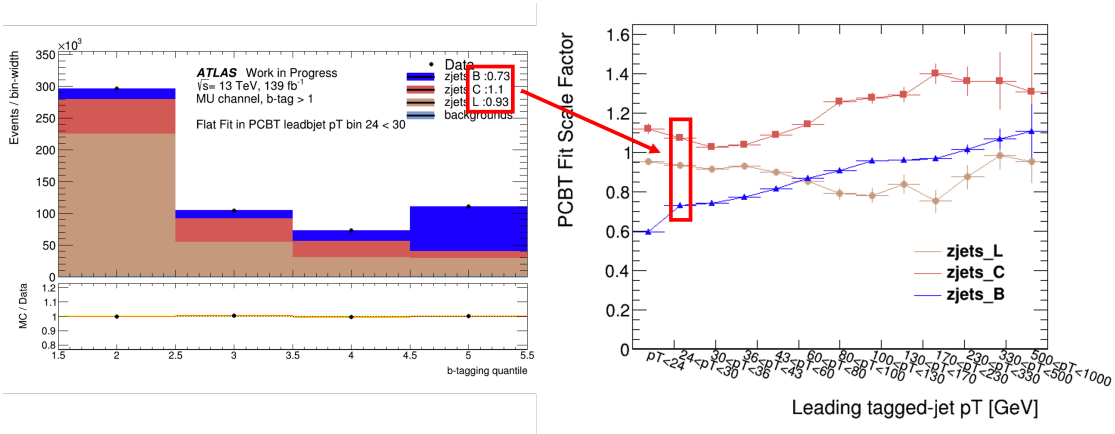

Figure 7.2: An illustration of how the PCBT flavor-fit scale factors (right) are derived for each bin of the observable using the PCBT distribution (left). Shown here are the scale factors resulting from the likelihood fit in in the $24 < p_{\mathrm{T}} < 30$ GeV bin of leading tagged-jet p_{T}, using Sherpa 2.2.11.

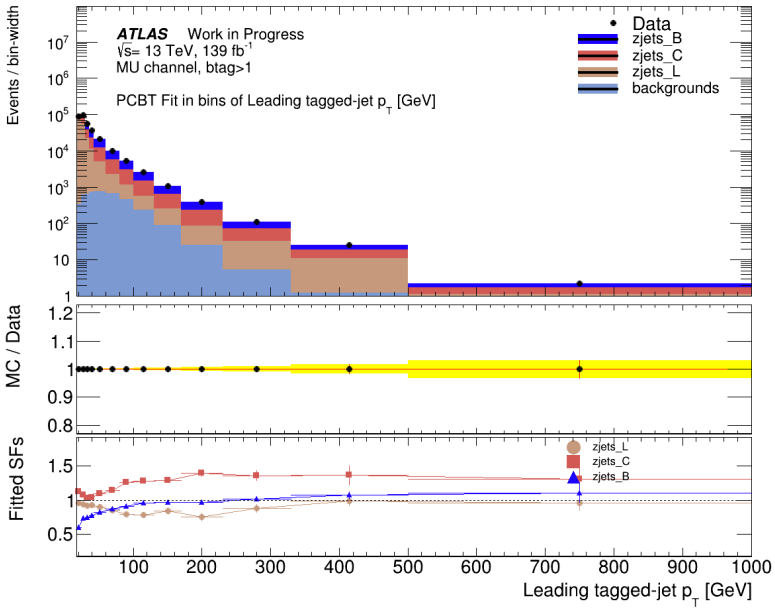

Figure 7.3: Post-PCBT-fit distribution of leading tagged-jet p_{T} in the muon channel, using events generated with Sherpa 2.2.11. The MC/data ratio is shown in the second pad, while the flavor SFs in each bin are plotted in the bottom pad.

7.2 Flavor percentage variance

The flavor fit described above, while resulting in excellent MC-data agreement, also exhibits undesirable variance. As an example, consider Figure 7.4, which shows the jet flavor percentages in each bin of leading tagged-jet p_{T}. Before the PCBT fit is performed, the distributions of jet flavor percentage per bin are relatively smooth, with the c-jet flavor percentage in particular being almost constant for $35 \lesssim p_{\mathrm{T}} \lesssim 300$ GeV. After the PCBT fit, however, the c-jet percentage varies from bin to bin, particularly for $p_{\mathrm{T}} \gtrsim 100$ GeV. The light-jet distribution also gains variance which mirrors the variance in the c-jet distribution: in bins where the c-jet percentage varies upwards from the relative mean, such as the bin of $130 < p_{\mathrm{T}} < 170$ GeV, the light-jet percentage exhibits a comparable downwards variation, and vice-versa. The b-jet distribution, meanwhile, appears relatively unchanged after the PCBT fit.

The reason for the variance introduced by the PCBT fit is the similarity of the c- and light-jet PCBT distributions, as shown in Figures 7.5b and 7.5c. The similarity of the distributions causes there to be no clear singular solution to the maximum likelihood problem — a SF applied to the c-jet distribution results in a similar effect as if applied to the light-jet distribution — and so the solution reached by the maximum likelihood fit is overly sensitive to small differences between the c-jet and light-jet distributions.

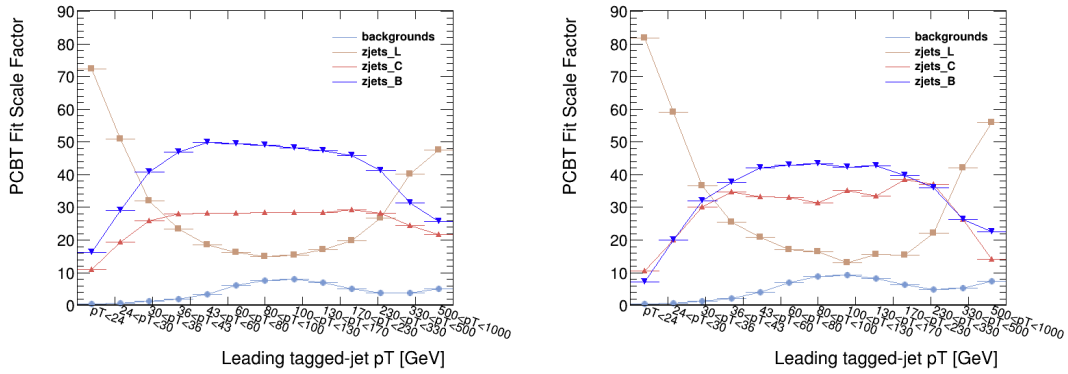

(a) Pre-fit (b) Post-fit

Figure 7.4: Distributions of jet flavor percentages in bins of leading tagged-jet p_{T} (a) pre-PCBT fit and (b) post-fit in the electron channel using MadGraph.

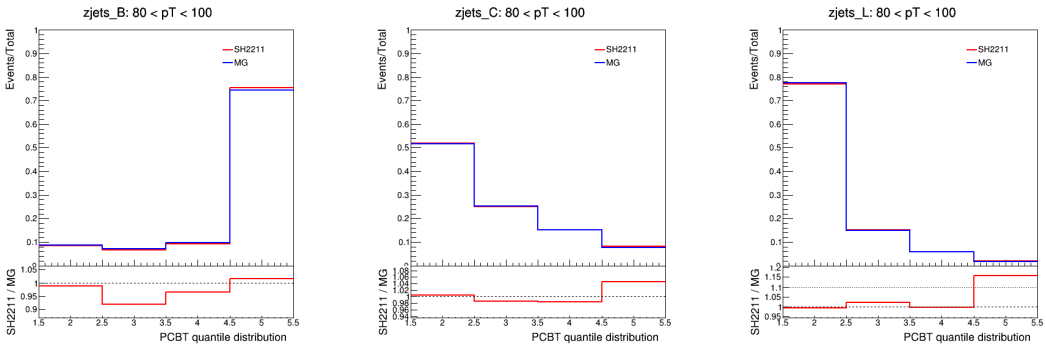

(a) b-jets (b) c-jets (c) Light jets

Figure 7.5: Distributions in the four PCBT quantiles, normalized by the total number of events, of (a) b-jets, (b) c-jets, and (c) light-jets events with $80 <$ leading tagged-jet $p_{\mathrm{T}} < 100$ GeV. Events are generated in the electron channel using MadGraph.

7.3 Control region fit

The method used to circumvent the problem described above is to separately fit the light-jet flavor fraction in a control region. An orthogonal control region is constructed by selecting only events with zero b-jets; as expected, this region is very pure ($\approx 90\%$) in light-jets (see Figure 7.6). A flat fit of the MC light-jet fraction to the data is performed in this control region and the resulting scale factor is applied to each bin of the MC; subsequently, the PCBT fit described earlier in this Section can be performed, but with the the light fraction no longer allowed to fluctuate. This two-step flavor fit results in a MC-data agreement almost as good as using the method described above (Figure 7.7a) while reducing the variance in the smooth flavor percentage distributions (Figure 7.7b).[1]

[1] A binwise fit of the light-jet flavor fraction can be performed instead of a flat fit; however, preliminary studies showed a negligible improvement in the post-fit MC-data agreement is negligible, and the binwise fit takes substantially more time to perform.

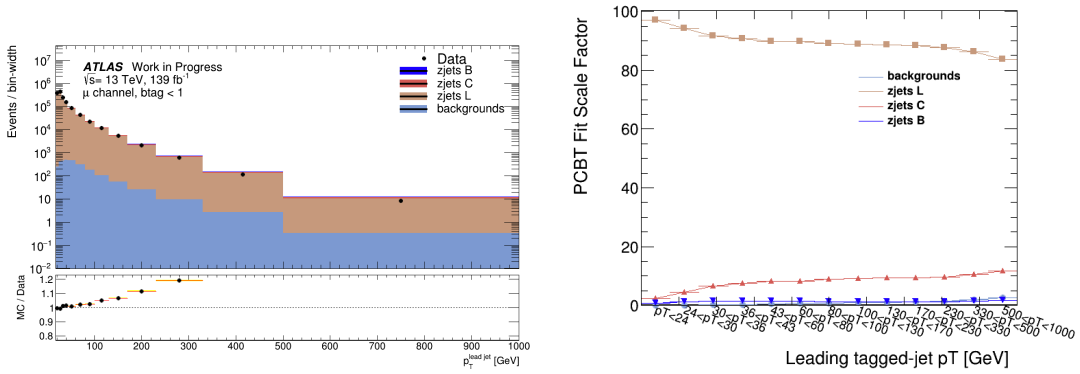

(a) (b)

Figure 7.6: (a) Leading jet p_{T} in the zero b-jet control region, using electron channel events generated using MadGraph. (b) Distributions of jet flavor percentages in bins of leading jet p_{T}.

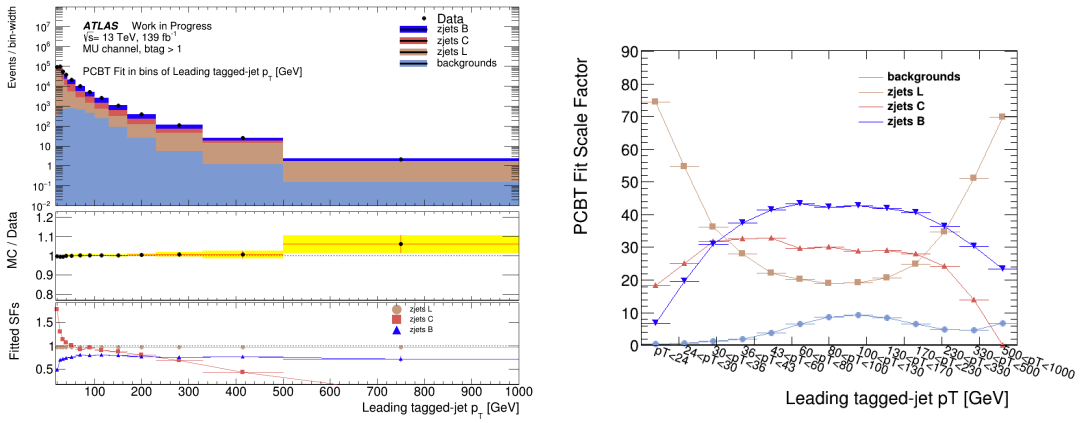

(a) (b)

Figure 7.7: (a) Leading tagged-jet p_{T} after the light-jet fraction fit in the control region and the PCBT fit of the b- and c-jet fractions, using muon-channel events generated using MadGraph. (b) Distributions of flavor percentages in bins of leading tagged-jet p_{T}.

7.4 Unfolding fitted results

The result of the flavor fits performed in this Section are scaled MC b-jet, c-jet, and light-jet contributions which comprise the overall MC signal. From here, the b- and c-jet distributions would each need to be unfolded, the motivation and process for which is described in the next Chapter. As alluded to at the beginning of this Chapter, however, this analysis chooses instead to discard the PCBT fit of the b- and c-jet components in favor of unfolding the b- and c-jet contributions simultaneously. This preserves the correlations between the two contributions and removes any uncertainty which is associated with the PCBT fit. The light-jet fit in the control region, meanwhile, will still be performed, as it provides a more accurate description of the largest background to our signal.

Chapter 8

Unfolding

Inherent to the measurements of Z+HF jet observables are distortions due to the finite resolution of the ATLAS detector. Unfolding, or deconvolution, is the process of correcting the measurements for these distortions. Unfolding must be performed in order to compare experimental data to theoretical predictions, which are not subject to detector resolution effects. Moreover, the results of separate experiments using different detectors will be subject to varying resolution effects; only by unfolding the effects of each detector can these separate results be compared and combined.

In this Chapter, I will present a general overview of the problem of unfolding, unfolding via matrix inversion, and regularized methods of unfolding. I will then discuss how "unrolling" the input b-jet and c-jet contributions prior to unfolding can be used to preserve the correlations between them; this is performed using the PCBT bins in order to simultaneously unfold b-jet and c-jet distributions.

8.1 Overview

Suppose that an experiment has been performed in which a detector was used to take n measurements of some random observable variable x; as an example, consider an experiment using the ATLAS detector to measure the transverse momentum p_{T} of the leading b-tagged jet in each pp collision. The limited resolution of the detector used to perform the measurements means that the measured values of x will fluctuate stochastically from their corresponding true values y. In the ATLAS example, this means that the measured b-tagged jet p_{T} distribution will not match the true b-jet p_{T} distribution. The relationship between the distribution $f(y)$ of the true values y and the measured distribution $g(x)$ is given by a Fredholm integral equation of the first kind:

$$g(x) = \int A(x|y)f(y)dy + b(x), \tag{8.1}$$

where the resolution function $A(x|y)$ describes the effect the the detector and $b(x)$ is the background contribution [142]. The resolution function $A(x|y)$ describes the total effect of the detector on the true distribution:

$$A(x|y) = s(x|y)\varepsilon(y). \tag{8.2}$$

Here the point-spread function $s(x|y)$ describes the probability to observe x given the true value y, while the efficiency $\varepsilon(y)$ describes the probability that an event with true observable value y will actually be measured; the number n of measurements which were made may not equal the number of possible measurements m. Using the ATLAS example, an event which does not fulfill all of the isolation, identification, or reconstruction criteria will not be measured. On the other hand, we may get a fake measured value when there was no event to actually measure; these backgrounds are accounted for with the $b(x)$ term. The problem of unfolding is to solve for $f(y)$ assuming $g(x)$, $A(x|y)$, and $b(x)$ are known.

In an experiment where a finite number of measurements are performed of the observable x, we must discretize Equation 8.1. Let $\boldsymbol{\mu} = (\mu_j, \dots, \mu_M)$,[1] with $\mu_j = \int_{j-1}^{j} f(y)dy$, be the "true histogram," the expected number of events in M bins if x could be measured exactly. The experiment obtains the data histogram $\boldsymbol{n} = (n_i, \dots, n_N)$, with $n_i = \int_{i-1}^{i} g(x)dx$ for N bins [143]. Each of these n_i is itself a (Poisson-distributed) random variable: for each bin i, the expected number of measured events is $\nu_i = E[n_i]$, and so we have the expected histogram $\boldsymbol{\nu} = \nu_i, \dots, \nu_N$. Finally, we can approximate the kernel resolution function $A(x|y)$ of Equation 8.2 with the *response matrix* $\boldsymbol{R}$. Each element R_{ij} of the response matrix gives the conditional probability for an event to be reconstructed from a measurement in bin i of the data histogram given that its true value was in bin j of the true histogram, i.e.

$$R_{ij} = P(\text{observed in bin } i|\text{true value in bin } j). \tag{8.3}$$

The response matrix $\boldsymbol{R}$ of a given observable $\mathcal{O}$ thus describes how strong an effect the detector has on the true distribution of $\mathcal{O}$. A perfectly diagonal response matrix for observable $\mathcal{O}$ describes a detector in which all events in a given truth bin μ_j of $\mathcal{O}$ are reconstructed in the same data bin ν_i of $\mathcal{O}$, in which case only the potential for detector inefficiency and the presence of backgrounds need to be accounted for. On the other hand, a response matrix with many off-diagonal elements describes a detector which tends to smear true values of the same truth bin μ_j of $\mathcal{O}$ into different data bins, i.e. a detector whose resolution prevents the observation of fine structure (sharp variations) in the

[1] I will consistently use bold symbols to differentiate vectors and matrices from variables.

data.

Putting all of this together, the (one-dimensional) discretized form of the Fredholm equation 8.1 is

$$\nu_i = \sum_{j=i}^{M} R_{ij}\mu_j + \beta_i, \tag{8.4}$$

where we continue to account for the possibility of background events β_i in each bin [144]. The task of unfolding in this analysis is therefore to obtain the estimators $\hat{\boldsymbol{\mu}}$ of the true distribution $\boldsymbol{\mu}$ using our data distribution $n_i = \hat{\nu}_i$ and a response matrix $\boldsymbol{R}$ obtained from MC simulation.

8.1.1 Estimator properties

Several concepts are used to evaluate the estimators $\hat{\boldsymbol{\mu}} = (\hat{\mu}_1, \dots, \hat{\mu}_M)$ we can obtain for the true distribution. The first is *consistency*: an estimator is said to be consistent if $\hat{\mu}_j$ converges to μ_j in the limit that infinitely many measurements are taken. This is a basic requirement for a useful estimator; if the estimator is supplied with more data, it should provide a better estimate of the true value. It is also desirable for an estimator to be unbiased, where the *bias* is defined [144] as

$$b_j = E[\hat{\mu}_j] - \mu_j. \tag{8.5}$$

The expectation value $E[\hat{\mu}_j]$ is the expected mean of $\hat{\mu}_j$[2] from infinitely many experiments. An unbiased estimator has $b_j = 0$ for any number of measurements, while an asymptotically unbiased estimator approaches $b_j = 0$ as the number of measurements taken approaches infinity. So whereas a consistent estimator is one which converges to the true value over an infinite number of measurements, an unbiased estimator is one which converges to the true value when a large number of experiments, each of which includes a finite number of measurements, is performed. This means that an unbiased estimator is particularly desirable for comparing the results of different experiments.

Finally, the bias and the *variance* $V[\hat{\mu}_j] \equiv E[\hat{\mu}_j^2] - E[\hat{\mu}_j]^2$ can be used to estimate the inherent error of the estimator by defining the *mean squared error* (MSE) [144]

$$\text{MSE} = V[\hat{\theta}] + b^2. \tag{8.6}$$

These concepts are used to differentiate the different possible sets of estimators which can be

[2]Because the estimator $\hat{\mu}_j$ is a function of the bins of data $\boldsymbol{n} = n_i, \dots, n_N$, which are random variables, $\hat{\mu}_j$ itself is a random variable which can be described by a probability distribution function (often referred to as the "sampling distribution"), its mean, etc.

constructed for a given experiment.

8.2 Inverting the response matrix

The simplest method of unfolding is to invert the response matrix. If we start with the matrix relation $\boldsymbol{\nu} = \boldsymbol{R\mu} + \boldsymbol{\beta}$, we invert to get

$$\boldsymbol{\mu} = \boldsymbol{R}^{-1}(\boldsymbol{\nu} - \boldsymbol{\beta}). \tag{8.7}$$

We are assuming that the measured distribution is estimated with our data ($\hat{\nu}_i = n_i$), which means that the estimators for the true distribution are

$$\hat{\boldsymbol{\mu}} = \boldsymbol{R}^{-1}(\boldsymbol{n} - \boldsymbol{\beta}). \tag{8.8}$$

This is both the least-squares solution and the solution obtained by maximizing the log-likelihood function [144]. In some ways, it is a promising solution: looking at Equation 8.8, we can see that, because we have chosen $\hat{\nu}_i = n_i$,

$$\begin{aligned} E[\hat{\mu}_i] &= \sum_{i=1}^{N} R_{ji}^{-1} E[n_i - \beta_i] = \sum_{i=1}^{N} R_{ji}^{-1} E[\nu_i - \beta_i] \\ &= \mu_i, \end{aligned} \tag{8.9}$$

meaning that the estimators are unbiased. In addition, the variance of the estimators is as small as it can be for an estimator of zero bias [144]. Unfortunately, the use of the data $\boldsymbol{n}$ as the estimators of $\boldsymbol{\nu}$ leads to a problem. The effects of finite detector resolution smear the true distribution; in other words, fine structure, or sharp variations in the data, tend to get flattened out. Unfolding, therefore, attempts to restore this fine structure. The data bins $\boldsymbol{n}$ are random variables, and so they are subject to statistical fluctuations. These statistical fluctuations can potentially introduce artificial fine structure, which unfolding will then exaggerate. The resulting estimators will thus have zero bias but a large variance; because of this, the unfolded results will be imprecise. One natural next step, then, is to find a different set of estimators, which have a smaller variance at the cost of some bias.

8.3 Regularized unfolding

Regularizing the unfolded distribution amounts to encouraging the estimators $\hat{\boldsymbol{\mu}}$ of the true distribution to be smooth, i.e. to suppress the large variance of the matrix-inversion estimators [145].

The downside is that, since the matrix-inversion solution is equivalent to the maximum-likelihood solution, the set of estimators obtained will by definition not be the set which is most likely to describe the true distribution. This can be a worthwhile tradeoff, however, because even if the estimators are less accurate, the decrease in variance can make the estimators more precise, thus decreasing the overall uncertainty in the unfolded results.

The first, and perhaps simplest, approach one can take to decrease the variance of the estimators is to remove some of the fine structure from the inputs. This can be done by, for example, decreasing the number of bins. However, typically the input binning is chosen explicitly to expose finer structure (if present), and so decreasing the number of bins can end up removing too much valuable information. A more systematic approach can therefore be taken to unfold the original inputs but accept a nonzero bias. The following subsections go over a few of these approaches.

8.3.1 Tikhonov regularization

Regularizing the unfolded solution can be constructed as a maximization problem where, instead of maximizing the log-likelihood function $\log L(\boldsymbol{\mu})$, we maximize a linear combination of $\log L(\boldsymbol{\mu})$ and a regularization function $S(\boldsymbol{\mu})$ which will define how we are imposing "smoothness" on the estimators of the true distribution $\boldsymbol{\mu}$ [145]. Maximizing such a linear combination will lead to a set of estimators which differ from the log-likelihood solution. We therefore also need to choose an acceptable degree of difference $\Delta \log L$ from the log-likelihood solution, i.e.

$$\log L(\boldsymbol{\mu}) \geq \log L_{\max} - \Delta \log L. \tag{8.10}$$

The regularized unfolding problem can thus be described as the maximization of

$$\alpha[\log L(\boldsymbol{\mu}) - \log L_{\max} - \Delta \log L] + S(\boldsymbol{\mu}), \tag{8.11}$$

where α is the regularization parameter [144]. This parameter governs how much smoothness we want to impose on the distribution: $\alpha = 0$ means as smooth as possible, though we are completely ignoring the data, while a large value leads to approximately the original inverse solution.

Tikhonov regularization supposes that a good measure of the smoothness of the distribution is the mean value of the square of one of its derivatives. This means that the smoothing function looks something like

$$S[f_{\text{true}}(y)] = -\int \left(\frac{d^k f_{\text{true}}(y)}{dy^k} \right)^2 dy. \tag{8.12}$$

Since we are thinking of $f_{\text{true}}(y)$ as representing a histogram, the integral becomes a sum and the derivative is approximated by using the difference between sucessive bins. For example, choosing $k = 2$ because the second derivative is a measure of the curvature of the distribution, the above expression simplifies to

$$S(\boldsymbol{\mu})_{k=2} = -\sum_{i=1}^{M-1}(-\mu_i + 2\mu_{i+1} - \mu_{i+2})^2. \tag{8.13}$$

Slotting this smoothing function into our minimization expression (Equation 8.11), we see that this method of regularization effectively penalizes the presence of bins very different from one another, i.e. fine structure: if neighboring bins have very different values, then $S(\boldsymbol{\mu})$ will be a large negative term in the maximization problem [144].

8.3.2 Iterative Bayesian unfolding

The Richardson-Lucy method [146, 147], promoted by D'Agostini [148] and often referred to as the iterative Bayesian method, starts with some initial probabilities $\boldsymbol{p}_0 = (p_1, \ldots, p_M)$ for finding an event in a given bin. The initial probabilities are usually set using MC particle-level information; the resulting estimators will then be biased towards this distribution, as will be discussed below. The estimators $\hat{\boldsymbol{\mu}}$ of the true distribution are then

$$\hat{\boldsymbol{\mu}} = n_{tot}\boldsymbol{p_0}, \tag{8.14}$$

where $n_{tot} = \sum_i n_i$ is the total number of events in the measured histogram. The estimators are then updated using the response matrix [149]:

$$\hat{\mu}_i = \frac{1}{\varepsilon_i}\sum_{j=1}^{N} P(\text{true value in bin } i|\text{found in bin } j) = \frac{1}{\varepsilon_i}\sum_{j=1}^{N}\left(\frac{R_{ij}p_i}{\sum_k R_{jk}p_k}\right)n_j. \tag{8.15}$$

These estimators, which include the migration information R_{ij}, will be an improvement over the initial estimators. The improvement can be evaluated using a χ^2 test. If the χ^2 value is too large, the initial probabilities are updated to $\boldsymbol{p} = \hat{\boldsymbol{\mu}}/\mu_{tot}$, where $\hat{\boldsymbol{\mu}}$ is the result of the previous iteration, and Equation 8.15 is evaluated again. If χ^2 is low enough, $\hat{\boldsymbol{\mu}}$ can be taken as the regularized solution.

If this process is iterated a very large number of times, the solution approaches the matrix inversion solution, complete with its large variances; therefore, the question is how many iterations are necessary to perform in order to arrive at a usable solution with a reasonably low χ^2. The number of iterations plays the same role as the regularization parameter α of the Tikhonov regu-

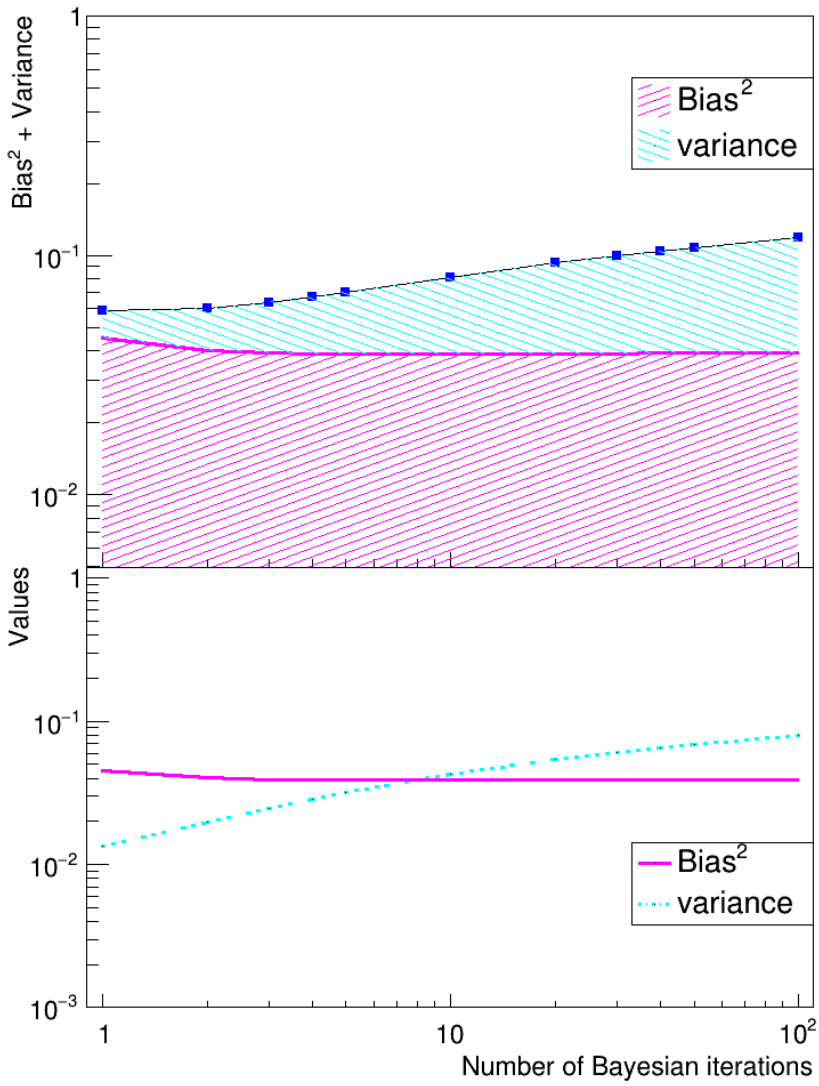

Figure 8.1: Bias squared and variance of unfolding estimators, using the iterative Bayesian method, as a function of the number of iterations. The regularization strength is inversely related to the number of iterations. Unfolding was performed using SHERPA 2.2.11, with MADGRAPH as the target "data."

larization method, in that it determines how much smoothness will be imposed on the unfolding solution: a large number of iterations, since the solution tends towards the matrix inversion solution, is effectively unregularized, while a small number of iterations leads to a highly regularized set of estimators. For example, consider Figure 8.1, which shows the bias squared and the variance as a function of the number of Bayesian iterations performed using the SHERPA sample. For a small number of variations, the bias is at its relative maximum because the initial estimators were formed using MC particle-level information. On the other hand, the variance is at a relative minimum. As the number of iterations increases, the regularization strength decreases; the bias

decreases because the estimators are farther removed from the initial MC particle-level guess, and the variance increases as more fine structure is accentuated.

8.3.3 Singular value decomposition

The method of singular value decomposition (SVD), developed by Höcker and Kartvelishvili [150], is a linear algebraic method of extracting the key features of the response matrix. Since our unfolding equation 8.4 is an expression of the linear transformation of $\hat{\boldsymbol{\mu}}$ into $\hat{\boldsymbol{\nu}}$, we can *decompose* the transformation into separate rotation and stretching transformations. The general form of the SVD is

$$R = USV^T, \tag{8.16}$$

where R is our $M \times N$ response matrix, U and V are $M \times M$ and $N \times N$ orthogonal matrices, respectively, and S is an $M \times N$ diagonal matrix with non-negative diagonal elements called singular values. The U and V matrices represent the rotations of the response matrix, while S applies the stretching or compression. The singular values, then, represent the semi-axes of the response matrix; in 2D, for example, they would represent the semi-axes of an ellipse (see Fig. 8.2). This

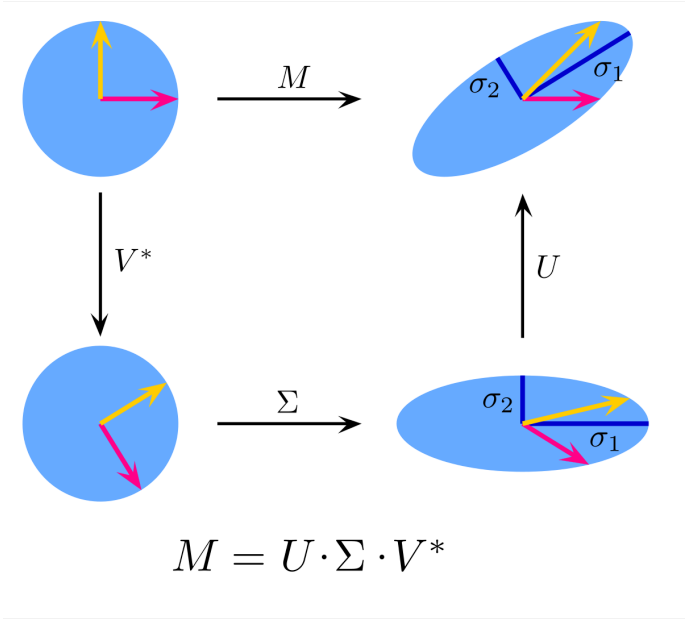

Figure 8.2: A visual depiction of singular value decomposition of a matrix M [151].

decomposition allows for the determination of $\hat{\boldsymbol{\mu}}$ using the singular values — essentially a system of linear equations. Regularization is then applied by adding an extra curvature term and scaling factor to the χ^2 minimization expression, in a manner similar (but not identical) to Tikhonov regularization.

8.3.4 The regularization parameter

All three of the above methods provide some sort of term which governs how much regularization is applied: for the Tikhonov method (and similarly for SVD, although it is not described in detail here), a scaling factor determines how much of a penalty is associated to the curvature terms of the minimization, while for the iterative Bayesian method the regularization term is the number of iterations. But there is no single, correct method for determining the proper regularization strength. Possible regularization strength choices include the minimum of the MSE of Equation 8.6; the point at which the bias squared and variance are equal, as it does for approximately 10 Bayesian iterations in Figure 8.1; and others. In this analysis, the Bayesian unfolding technique was performed using two iterations; this choice was made because, as seen in the Figure 8.1, the bias effectively flattens at around two iterations while the variance continues to rise as more iterations are performed. Performing only two iterations also limits the computing time necessary to perform the unfolding.

8.4 Unfolding inputs

Unfolding is performed in this analysis using RooUnfold [145, 152], which includes implementations of the unfolding methods mentioned above plus several others. The response matrix is formed using events passing both particle-level ("truth") and detector-level (reconstructed or "reco") selections, as outlined in Table 6.4; no reco-truth object matching, such as via a ΔR cut, is applied. The truth distribution is used as the Bayesian prior for the first iteration of the matrix problem; as shown in Figure 8.1, the unfolded result is biased towards this prior, but the bias decreases as the number of iterations increases.

Events which pass the truth selections may fail reconstruction for a number of reasons, including lepton ID and b-tagging inefficiencies, Z boson reconstruction, and small differences in the phase space such as the η "crack" Before the data is unfolded, the MC backgrounds are subtracted, and a scale factor is applied to correct for the presence of fake events. The scale factor is determined by dividing the number of events which passed both the truth and reco selections by the total number of events which passed the reco selections.

8.5 Sequential flavor fit and unfold

For this analysis of Z+b-jet events, one approach for obtaining results from unfolded data would be to do the following: first, perform a flavor fit of the MC to the data to infer the proportionality of signal events (b-jet events) to background events (c-jet and light jet events); second, subtract

the fitted MC background estimates from the data to obtain a signal sample; and finally, using a response matrix determined solely using MC, unfold the background-subtracted data using one of the aforementioned methods. The b-jet and c-jet distributions are unfolded separately to the truth b and truth c distributions, respectively, resulting in two separate sets of estimators.

In this method, a migration matrix is first formed for each heavy flavor with matrix elements M_{ij} = (reco value, true value). This matrix is then normalized so that the bins are of equal width and divided by the true distribution to form the response matrix. Figure 8.3a shows an example migration matrix on the left and the corresponding response matrix on the right. The red dot and arrows provide an illustrative example of a response matrix element: an event in the second bin of p_{T} in truth has a probability of being reconstructed in the first ($\sim 30\%$ chance), second ($\lesssim 70\%$), and third ($\lesssim 10\%$) bin of p_{T}.

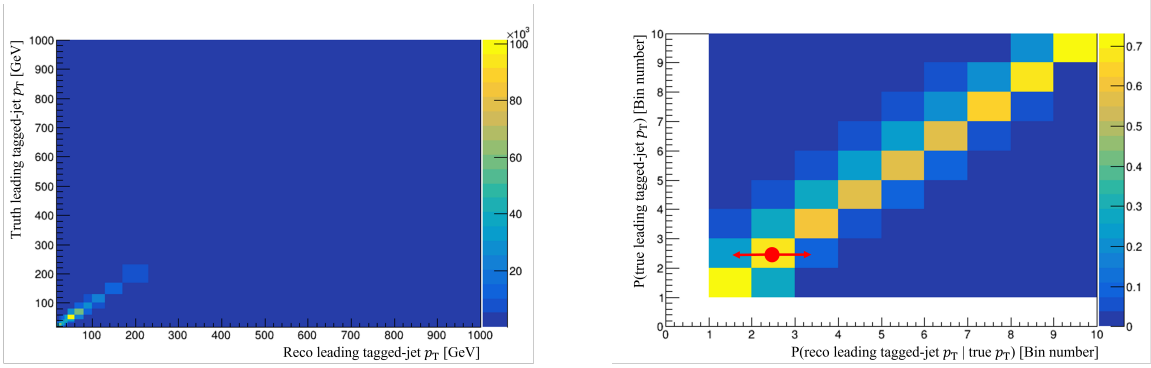

(a) (b)

Figure 8.3: (a) An example migration matrix for the leading heavy-flavor jet p_{T} observable using Sherpa 2.2.11. Rescaling (a) to equal bin widths and dividing by the truth distribution results in the response matrix (b). The red dot and arrows show the probability of migration from the second bin of p_{T} in truth to the first, second, or third bins of p_{T} in reco.

8.6 Unrolled unfolding

Because this analysis is interested in both Z+b-jet events as well as Z+c-jet events, we must considered how the detector's limited resolution plays a role in our ability to distinguish between the two signal components. The above method, which determines separate sets of estimators for b-jet and c-jet distributions, is agnostic to the correlations between the two flavor fractions (and they certainly are (anti-)correlated: this analysis uses a b-tagging discriminant, and so events which are b-tagged are by definition less likely to include c-jets.) These correlations affect the unfolded results but also the uncertainties. For example, to perform a measurement of the Z+b-jet/Z+c-jet ratio, these correlations must be taken into account for the result to have meaning.

To account for these correlations, we unfold the migration between the two flavors before separately plotting b-jet and c-jet distributions. The method used to perfom this simulatenous unfolding

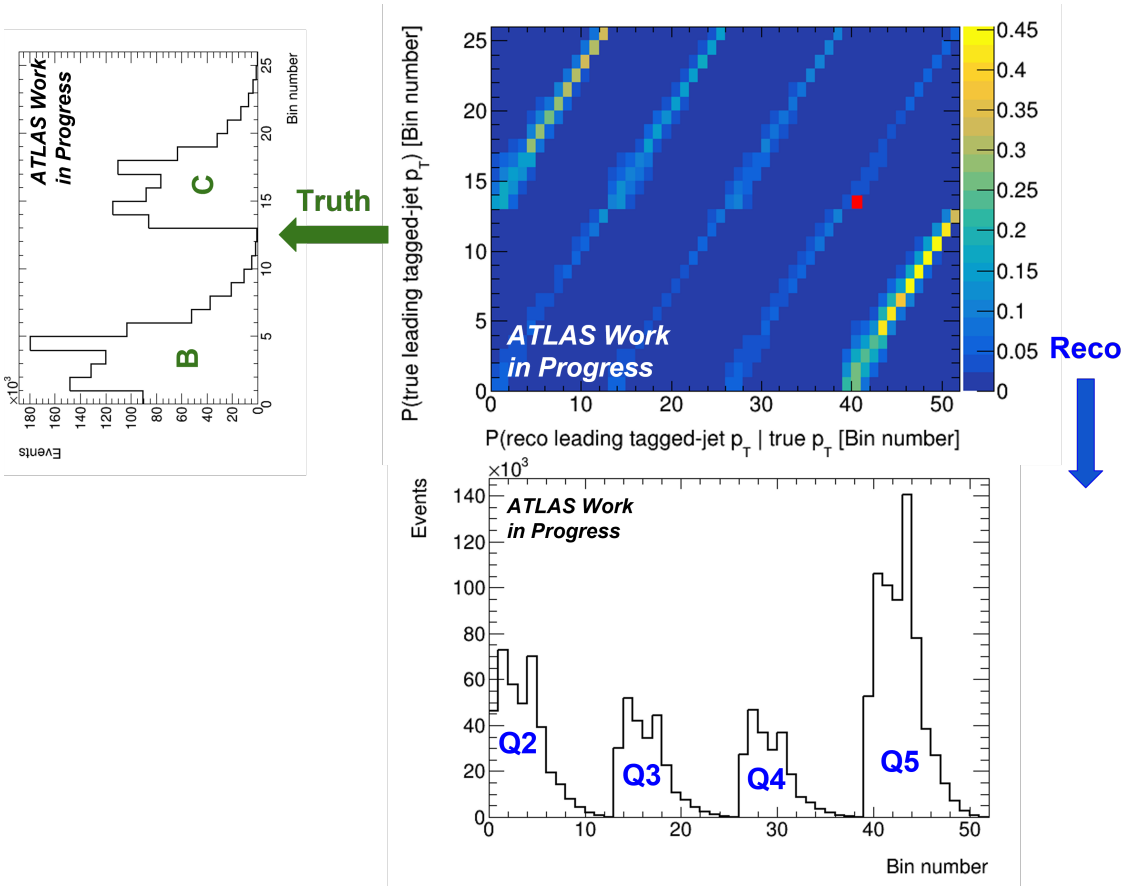

Figure 8.4: A migration probabilities matrix showing the eight possible migrations from particle-level ("Truth") b- and c-jets to each of the four PCBT bins at detector-level ("Reco"), as well as migration between the bins of p_{T}. These distributions were constructed using muon channel events generated by Sherpa 2.2.11.

is to "unroll" our observables into the four PCBT bins and to use the unrolled distributions in the unfolding. This will serve as the reconstructed input to the migration matrix. For the true distribution, we have the b-jet and c-jet distributions separated using the truth information of the originating quarks. These elements are used to construct the migration probabilities matrix shown in Figure 8.4. This matrix demonstrates that there is simultaneous migration from particle-level b- and c-jets to each detector-level PCBT quantile, as shown by the mostly-diagonal octants, and migration between bins of p_{T}, as shown by the off-diagonal elements in each octant. For example, the bin marked by the red dot quantifies the nonzero probability that an event which at particle-level contained a 17-24 GeV c-jet is reconstructed with 24-30 GeV in the highest b-tagging efficiency PCBT quantile.

The regularized unfolding methods discussed in this Chapter, such as Bayesian unfolding, work as well for an $M \times N$ response matrix as they would for a standard square $M \times M$ response matrix. Figures 8.5 and 8.6 shows the closure tests performed for leading b-jet and leading c-jet p_{T}, respectively, in the electron channel for both Sherpa 2.2.11 and MadGraph, demonstrating that the unfolded reconstructed distribution is consistent with the truth distribution.

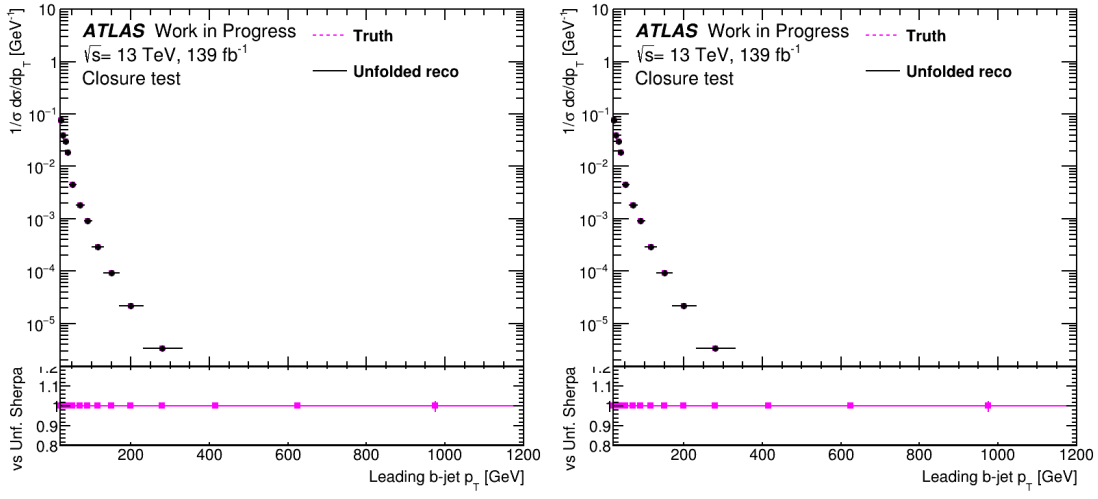

(a) Sherpa 2.2.11 (b) MadGraph

Figure 8.5: Sherpa 2.2.11 (a) and MadGraph (b) closure tests of leading b-jet p_{T} in the muon channel: for both generators, the reconstructed events unfolded using the migration matrix is identical to the true distribution.

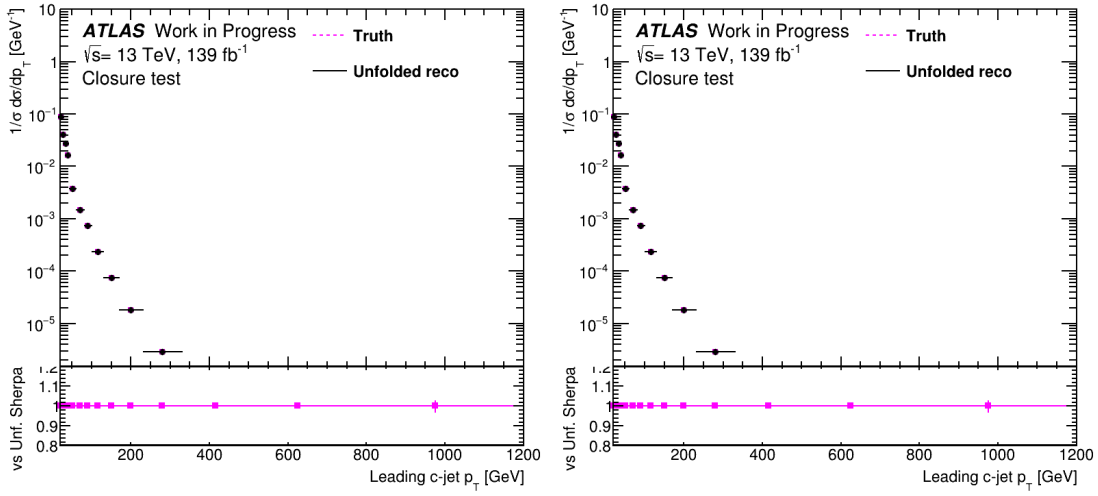

(a) Sherpa 2.2.11 (b) MadGraph

Figure 8.6: (a) Sherpa 2.2.11 and (b) MadGraph closure tests of leading c-jet p_{T} in the muon channel: for both generators, the reconstructed events unfolded using the migration matrix is identical to the true distribution.

An unrolled migration matrix is obtained for each of the observables listed in section 6.7, which allows for the b-jet and c-jet distributions to be obtained separately while still accounting for the correlated detector effects. The results of this unfolding procedure are presented in Chapter 10. Before that, Chapter 9 discusses the handling of the various uncertainties in this measurement.

Chapter 9

Uncertainties

This chapter describes how the total uncertainty in the unfolded results, which are presented in Chapter 10, is estimated. The total uncertainty is determined in each bin of a given observable as the sum in quadrature of the experimental, modelling, and statistical uncertainties. The estimation of each of these uncertainties is described in further detail in the following sections.

9.1 Experimental systematic uncertainties

This section gives an overview of the detector-level uncertainties. The ATLAS combined performance groups provide recommended systematic variations to account for uncertainty in the various MC corrections and calibrations described in Chapter 3. These systematic variations are applied to the MC samples to produce variated distributions, on which the entire measurement procedure — the light-jet component fit, background subtraction, and unfolding — is performed. For each independent systematic variation, the difference between the variated unfolded result and the nominal unfolded result is obtained. The total experimental systematic uncertainty, which is labelled "Experimetal" in the total uncertainty plots below, is then the sum in quadrature of these differences; as the uncertainties are independently obtained, they are treated as uncorrelated. For this analysis, systematic uncertainties were determined using MC16a systematic samples and then applied to the full Run 2 dataset; this was done to save time in the generation of systematic samples and is expected be a valid representation of the total systematic uncertainty due to the overall similarity of the MC productions used. The only significant difference between the MC16 campaigns is the pileup; as discussed in Chapter 5, this is taken into account in the simulated datasets through pileup reweighting, and the impact of pileup on systematic variations is a second-order effect.

Each uncertainty, unless otherwise specified, is defined using asymmetric $\pm 1\sigma$ "up/down" variations. The uncertainties are listed in full in Tables 9.2 (jet uncertainties) and 9.1 (other uncer-

tainties); an overview of their purpose is as follows [153]:

- **EG**: This group encompasses uncertainties due to imperfect electron calibration. This includes mismodelling of the electron energy resolution (EG_RESOLUTION_ALL) and scale (EG_SCALE).
- **EL_EFF**: These uncertainties address the effects of imprecise electron efficiency measurements. This includes measurements of the reconstruction (EL_EFF_Reco_TOTAL_1NPCOR_PLUS_UNCOR), identification (EL_EFF_ID_TOTAL_1NPCOR_PLUS_UNCOR), and isolation (EL_EFF_Iso_TOTAL_1NPCOR_PLUS_UNCOR) efficiencies which are used to form the efficiency scale factors applied to each signal lepton, as described in Chapter 6. In addition, a systematic is included for the determination of the trigger scale factor (EL_EFF_Trigger_TOTAL_1NPCOR_PLUS_UNCOR).
- **MUON**: This includes uncertainties in the muon measurements, including the smearing of tracks in the ID (MUON_ID) and MS (MUON_MS) and mismodelling of the momentum scale (MUON_SCALE). A charge-dependent scale correction (MUON_SAGITTA_RHO) is included, along with the residual charge dependency after the correction (MUON_SAGITTA_RESBIAS). Finally, both the statistical and systematic uncertainties in the isolation measurement (MUON_ISO_STAT/MUON_ISO_SYST) and the track-to-vertex association (MUON_TTVA_STAT/MUON_TTVA_SYST) are included.
- **MUON_EFF**: The uncertainties on muon efficiency measurements include per-object correction factors (MUON_EFF_STAT/MUON_EFF_SYS) and uncertainties in the trigger efficiencies (MUON_EFF_TrigStatUncertainty/MUON_EFF_TrigSystUncertainty) used to form the applied trigger scale factors. Correction factors are also included specifically for low-p_{T} muons (MUON_EFF_STAT_LOWPT/MUON_EFF_SYS_LOWPT).
- **MET**: The missing transverse energy (MET) uncertainties are uncertainties on the resolution (MET_SoftTrk_ResoPerp) and scale (MET_SoftTrk_Scale) due to the track-based soft-term contributions. These two uncertainties are accounted for using a single (symmetrized) variation as opposed to separate up and down variations. Hard-term uncertainties are covered by the electron, muon, and jet uncertainties and are propagated through the MET calculation.
- **JET**: Uncertainties on the JES calibration are handled with correction factors measured in-situ with respect to well-understood reference objects; these include Zs (relevant for low-p_{T} ($<$ 100 GeV) jets), photons (100 GeV - 1 TeV), and systems of previously-calibrated low-p_{T} jets (1-2 TeV). A correction factor is then obtained as the ratio of the jet p_{T} to the p_{T} of the object the jet is being balanced against. All of these nuisance parameters (NPs) are combined into a smaller set of statistical (JET_EffectiveNP_Statistical), modelling (JET_EffectiveNP_

Plot Category	Source	Systematics
"EGAMMA"	Electron	EG_RESOLUTION_ALL EG_SCALE_AF2 EG_SCALE_ALL
	Electron efficiency	EL_EFF_Reco_TOTAL_1NPCOR_PLUS_UNCOR EL_EFF_ID_TOTAL_1NPCOR_PLUS_UNCOR EL_EFF_Iso_TOTAL_1NPCOR_PLUS_UNCOR EL_EFF_Trigger_TOTAL_1NPCOR_PLUS_UNCOR
"MUON"	Muon	MUON_SCALE MUON_ID MUON_MS MUON_SAGITTA_RESBIAS MUON_SAGITTA_RHO MUON_ISO_STAT MUON_ISO_SYST MUON_TTVA_STAT MUON_TTVA_SYST
	Muon efficiency	MUON_EFF_STAT MUON_EFF_SYS MUON_EFF_TrigStatUncertainty MUON_EFF_TrigSystUncertainty MUON_EFF_STAT_LOWPT MUON_EFF_SYS_LOWPT
"MET"	Missing transverse energy	MET_SoftTrk_ResoPara MET_SoftTrk_ResoPerp MET_SoftTrk_Scale

Table 9.1: Electron, muon, and missing transverse energy uncertainties. All of the systematic uncertainties have separate `__up` ($+1\sigma$) and `__down` (-1σ) variations except for `MET_SoftTrk_ResoPara` and `MET_SoftTrk_ResoPerp`, which are symmetrized around the nominal distributions. The first column categorizes the systematics based on how they are plotted in this Chapter.

`Modelling`), detector (`JET_EffectiveNP_Detector`), and mixed (`JET_EffectiveNP_Mixed`) variations. Additional in-situ calibration are done to balance forward jets against central jets (`JET_EtaIntercalibration`) and very high (1-2 TeV) p_{T} jets against single-particle propagation (`JET_SingleParticle_HighPt`). Other uncertainties accounted for include pileup (`JET_Pileup`), which is relevant for very low-p_{T} jets; punch-through jets which escape the calorimeters (`JET_PunchThrough`); the mismodelling of the jet flavor response (`JET_Flavor_Response`) and composition (`JET_Flavor_Composition`), with b-jets handled individually (`JET_BJETS_Response`); and the mismodelling of the jet vertex tagging (JVT) requirement efficiency (`JET_JvtEfficiency`). An overview of the JES uncertainties as functions of jet p_{T} and η is provided in Figure 9.1.

- `JER`: Uncertainties in the Jet Energy Resolution (JER) smearing correction applied to MC are handled with a set of seven NPs (`JET_JER_DataVsMC_MC16`,`JET_JER_EffectiveNP`).

- `PRW`: The uncertainty in the reweighting of the pileup vertices in simulation to the data is handled with the `PRW_DATASF` variations.

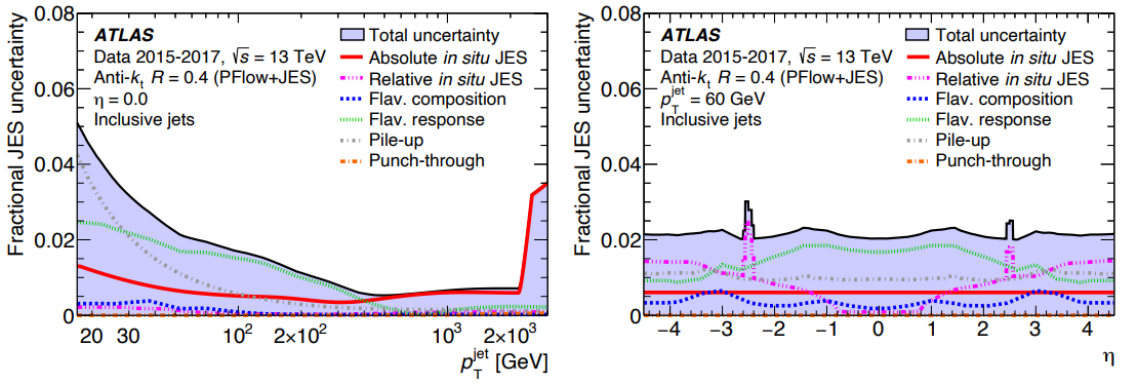

Figure 9.1: Components of the fractional jet energy scale systematic uncertainty as a function of (a) jet p_{T} and (b) jet η at $p_{\mathrm{T}} = 60$ GeV [96].

- `FT_EFF`: The uncertainty in the jet flavor tagging efficiency is handled with a set of nuisance parameters determined using an eigenvector variation method. 44 variations are used for b-jets (`FT_EFF_Eigen_B`) and 19 are used each for c-jets (`FT_EFF_Eigen_C`) and light jets. (`FT_EFF_Eigen_Light`).

9.2 Modelling uncertainties

This Section covers uncertainties related to the theory MC modelling of the signal and background samples. As discussed in Chapter 5 the MadGraph and Sherpa simulated samples use matrix element calculations at different levels of precision, different parton shower models and matching/merging schemes, and different tunes of their free parameters.[1] The difference between the two distributions therefore spans that between the majority of the available MC generators. Including systematics on the signal and background modelling therefore provides an estimate of how strongly the measurement results depend on the theoretical assumptions being made.

Because the Sherpa signal sample is used to form the response matrix, while the background samples are subtracted from the data; the signal and background uncertainties have to be treated separately, as described below.

9.2.1 Signal modelling

Since the MC Z+HF jets distributions are used to form the response matrices with which the data are unfolded, uncertainty in the modelling of those events has to be estimated after the unfolding procedure has been performed. The uncertainty is therefore obtained by unfolding the Sherpa reco distribution using two different response matrices: the first is derived from the Sherpa sample (in

[1] This does not include the PDF dependence, which is expected to be smaller than the difference described above.

Plot Category	Source	Systematics
"JET"	General	JET_EffectiveNP_Detector{1,2} JET_EffectiveNP_Mixed{1,3} JET_EffectiveNP_Modelling{1,4} JET_EffectiveNP_Statistical{1,6}
	Eta intercalibration	JET_EtaIntercalibration_Modelling JET_EtaIntercalibration_NonClosure_2018data JET_EtaIntercalibration_NonClosure_highE JET_EtaIntercalibration_NonClosure_negEta JET_EtaIntercalibration_NonClosure_posEta JET_EtaIntercalibration_TotalStat
	Pileup	JET_Pileup_OffsetMu JET_Pileup_OffsetNPV JET_Pileup_PtTerm JET_Pileup_RhoTopology
	Flavor	JET_Flavor_Composition JET_Flavor_Response JET_BJES_Response
	Punch-through	JET_PunchThrough_MC16 JET_SingleParticle_HighPt
	b-jets	JET_BJETS_Response
	JER	JET_JER_DataVsMC_MC16 JET_JER_EffectiveNP_{1,11} JET_JER_EffectiveNP_12restTerm
"FT"	Flavor tagging efficiency	FT_EFF_Eigen_B_{1,44} FT_EFF_Eigen_C_{1,19} FT_EFF_Eigen_Light_{1,19}
"PRW"	Pileup reweighting	PRW_DataSF

Table 9.2: Jet uncertainties. All of the systematic uncertainties have separate `__up` ($+1\sigma$) and `__down` (-1σ) variations. The first column categorizes the systematics based on how they are plotted in this Chapter.

which case closure is obtained), and the other is derived from the MadGraph sample. The relative difference between the resulting unfolded distributions is taken as an uncertainty on the modelling of the signal samples; this uncertainty is labelled "(Unfolded) Sig. modelling" in the plots below. It should be noted that this modelling uncertainty is larger than the the bias introduced by using two Bayesian iterations: as seen in Figure 8.1, an unfolded solution derive using two iterations suffers from only $\approx 0.1\%$ more bias than a roughly unregularized (100 iterations) solution, and the signal modelling uncertainty described here covers this excess.

9.2.2 Background modelling

The uncertainty in the MC modelling of the background components, which are subtracted from the data and not used in the unfolding, also has to be taken into account. For each of Sherpa and MadGraph the Z+light-jet background component is fit in the control region. The fitted Z+light-jet distribution is then added to the other backgrounds as modelled by the same generator. This results in a post-fit total background distribution for Sherpa and for MadGraph; the relative

difference is taken as the background modelling uncertainty. This uncertainty is labelled "Bkg. modelling" in the plots below.

9.3 Unfolded statistical uncertainty

This covers the statistical uncertainty on the measurement results as propagated through the unfolding procedure. The uncertainty is determined through the toy MC procedure of RooUnfold, in which toy datasets are sampled from the probability density of the input data and unfolded. 10,000 such toy datasets are generated by RooUnfold and unfolded using the nominal MC response matrix to produce a covariance matrix expressing the variances on the unfolded bins. The diagonal elements of this matrix are taken as the per-bin unfolding uncertainties. This uncertainty fully accounts for the original statistical uncertainty in the data as it appears in the unfolded results, as well as any variance introduced by the unfolding procedure itself: unfolding using two Bayesian iterations introduces only $\approx 0.5\%$ more variance over one iteration, as seen in Figure 8.1. This statistical uncertainty is labelled "(Unfolded) Statistical" in the plots below.

9.4 Total and experimental uncertainty results

Presented below are sets of plots for each observable showing the total uncertainty, contribution from each source of uncertainty, and the breakdown of uncertainty components as well as the breakdown of the experimental systematic uncertainties. In general, the background modelling uncertainty is the largest, due to the large light-jet component in the signal region. Because this analysis seeks to measure c-jet as well as b-jet events, the b-tagging discriminant value was chosen to be loose enough so as to not remove too many c-jet events. For example, as opposed the 85% b-tagging efficiency working point chosen for this analysis, the previous ATLAS Z+b-jet analysis [5] used a 70% b-tagging efficiency working point (albeit of a different algorithm); as a result, the light-jet background in the previous analysis was 18%, as compared to the $> 40\%$ background in this analysis. It is therefore imperative to accurately estimate the background modelling uncertainty. While the method presented here is robust and is used in other analyses, future studies should compare this estimation to one derived from systematic variations of the PDFs, matching/merging schemes, etc. used to generate the background MC samples. In addition, an analysis of Z+b-jet/Z+c-jet observable ratios would have much of this modelling uncertainty cancel in the ratio; this is another avenue for future work.

As expected, the largest components of the experimental systematic uncertainty, which tends to be most significant at low-p_{T}, are the jet and flavor-tagging systematic uncertainties. The experimental systematic uncertainties are comparable between the lepton channels aside from

the expected increase in electron-related systematic uncertainty in the electron channel. There are, however, two bins in muon channel leading c-jet x_F distribution with abnormally large jet uncertainties. Given the limited statistics in this and larger x_F bins, this is likely due to a small stochastic fluctuation in one or multiple of the "JET" systematics samples.

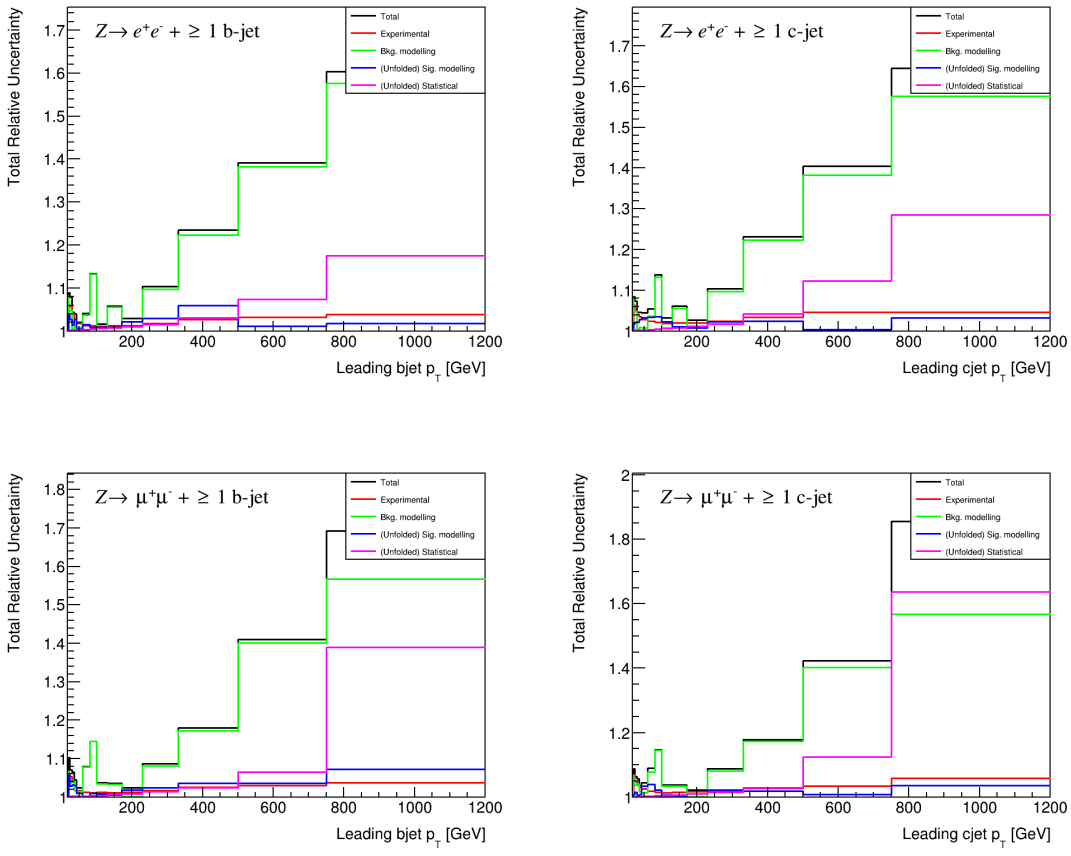

Figure 9.2: Total uncertainty on the $Z \to l^+l^- + \geq 1$ b-jet (left column) and $Z \to l^+l^- + \geq 1$ c-jet (right column) differential cross-sections as functions of the leading heavy-flavor jet p_{T}, in the electron (top row) and muon (bottom row) channels.

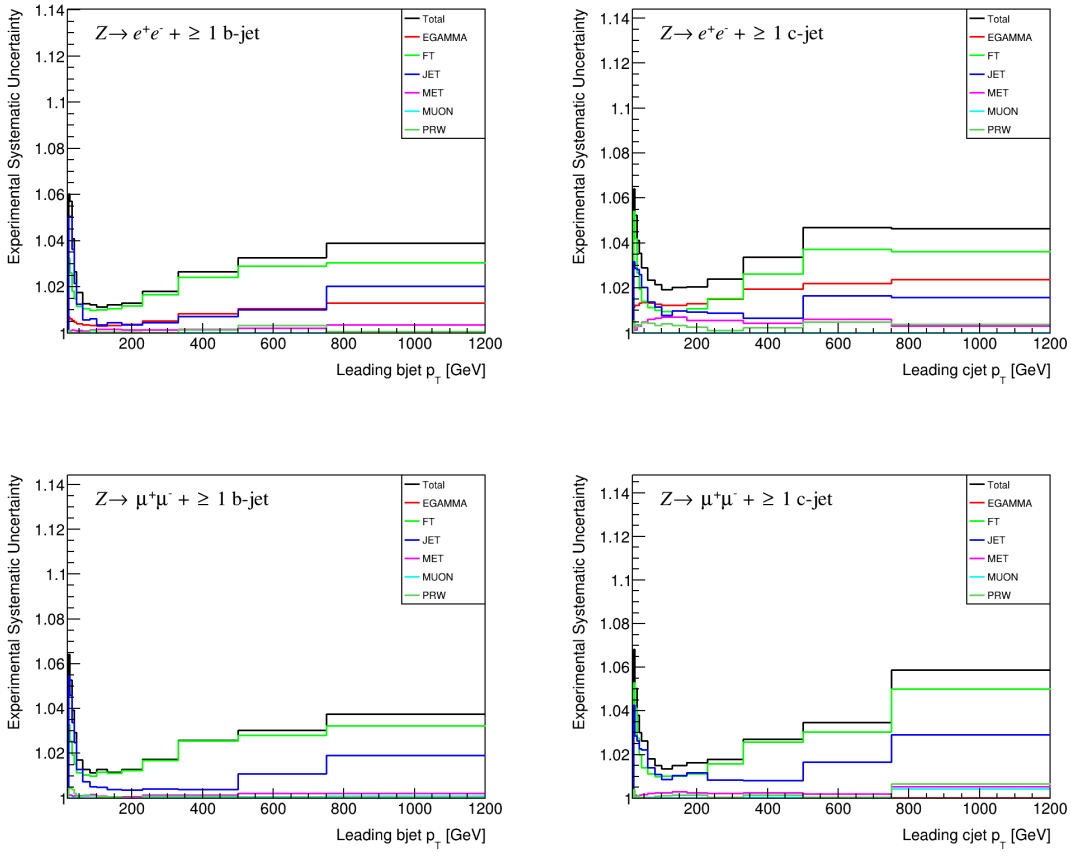

Figure 9.3: Systematic uncertainties on the $Z \to l^+l^- + \geq 1$ b-jet (left column) and $Z \to l^+l^- + \geq 1$ c-jet (right column) differential cross-sections as functions of the leading heavy-flavor jet p_{T}, in the electron (top row) and muon (bottom row) channels.

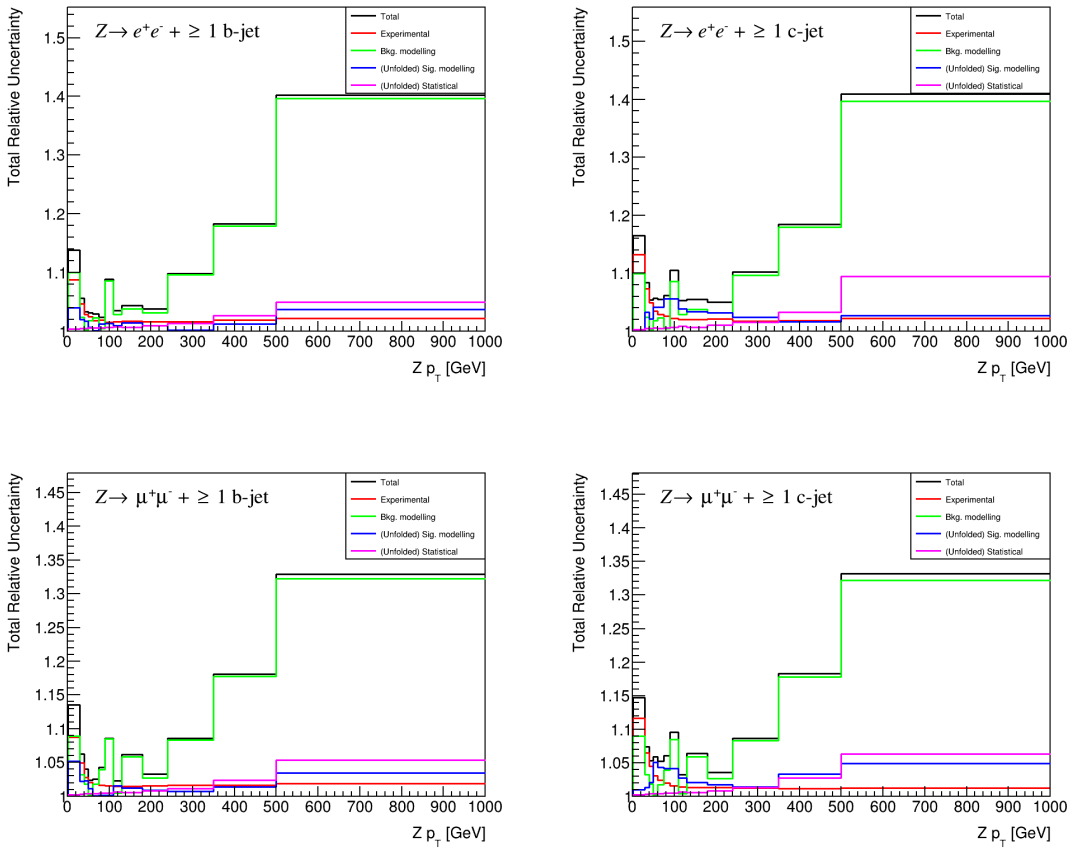

Figure 9.4: Total uncertainty on the $Z \to l^+l^- + \geq 1$ b-jet (left column) and $Z \to l^+l^- + \geq 1$ c-jet (right column) differential cross-sections as functions of p_{T}^{Z}, in the electron (top row) and muon (bottom row) channels.

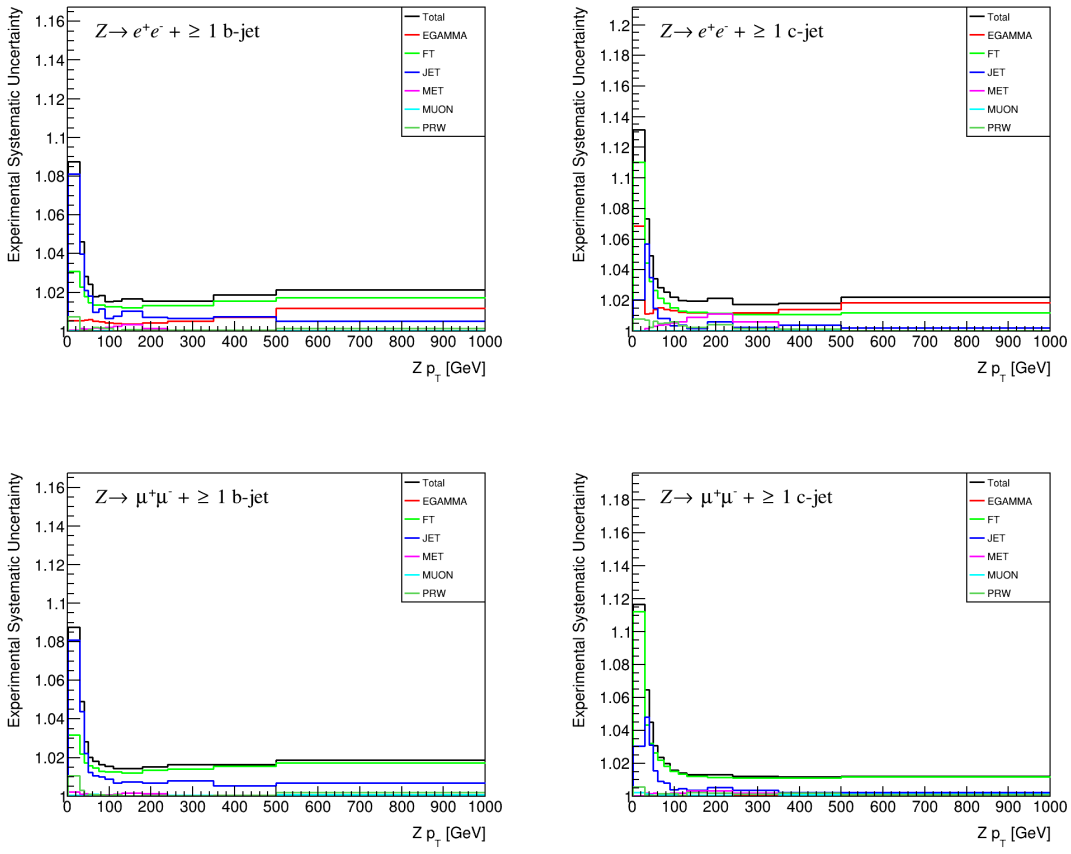

Figure 9.5: Systematic uncertainties on the $Z \to l^+l^- + \geq 1$ b-jet (left column) and $Z \to l^+l^- + \geq 1$ c-jet (right column) differential cross-sections as functions of p_{T}^{Z}, in the electron (top row) and muon (bottom row) channels.

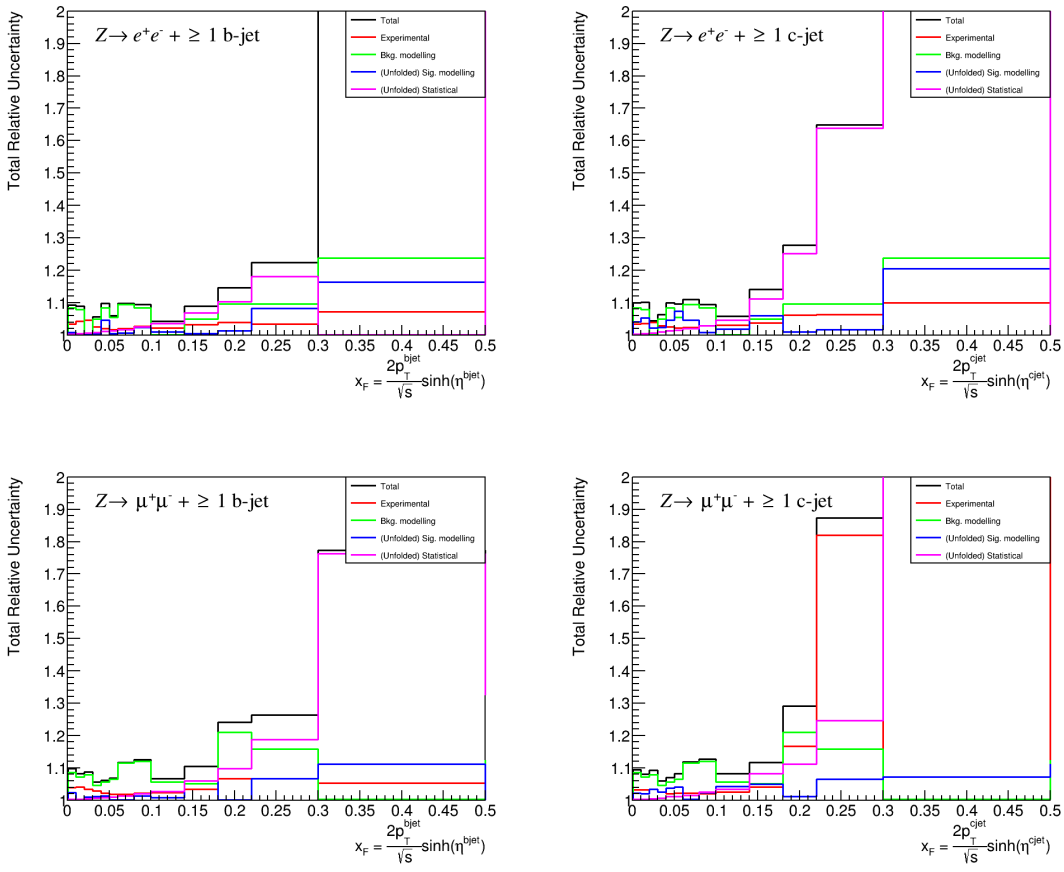

Figure 9.6: Total uncertainty on the $Z \to l^+l^- + \geq 1$ b-jet (left column) and $Z \to l^+l^- + \geq 1$ c-jet (right column) differential cross-sections as functions of the leading heavy-flavor jet x_F, in the electron (top row) and muon (bottom row) channels.

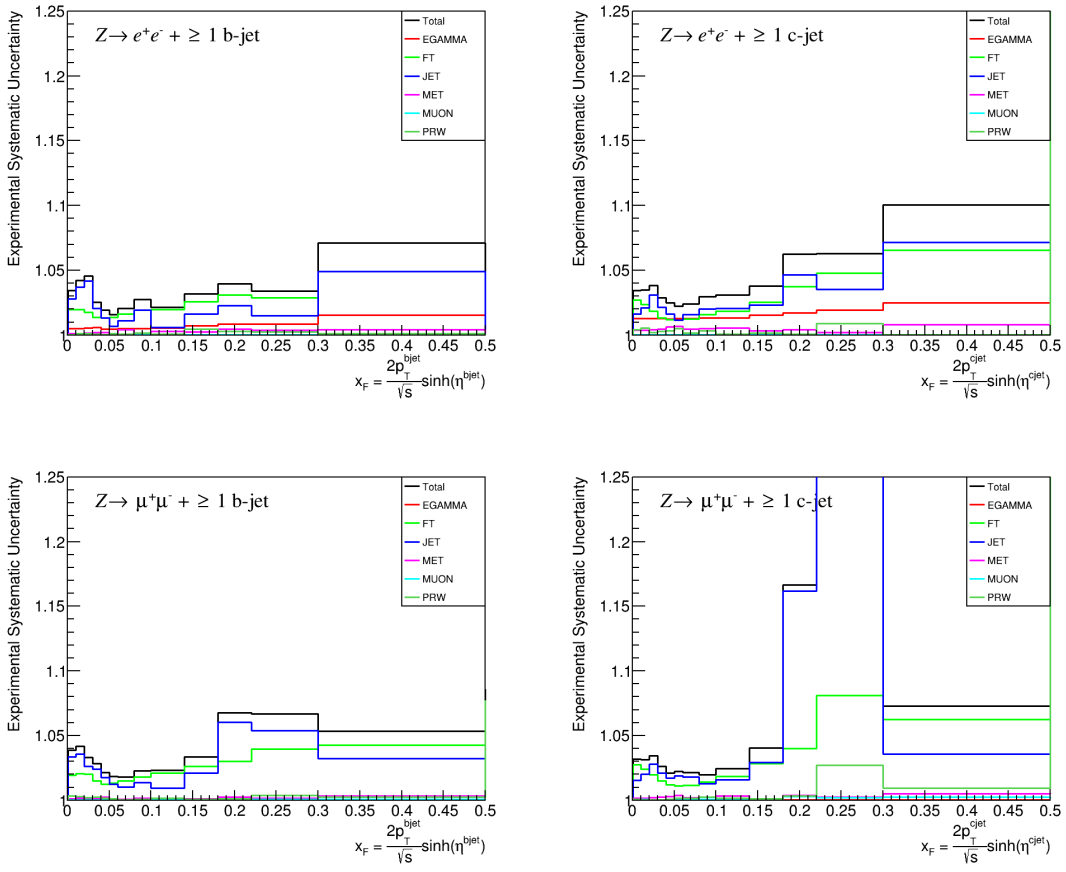

Figure 9.7: Systematic uncertainties on the $Z \to l^+l^- + \geq 1$ b-jet (left column) and $Z \to l^+l^- + \geq 1$ c-jet (right column) differential cross-sections as functions of the leading heavy-flavor jet x_F, in the electron (top row) and muon (bottom row) channels.

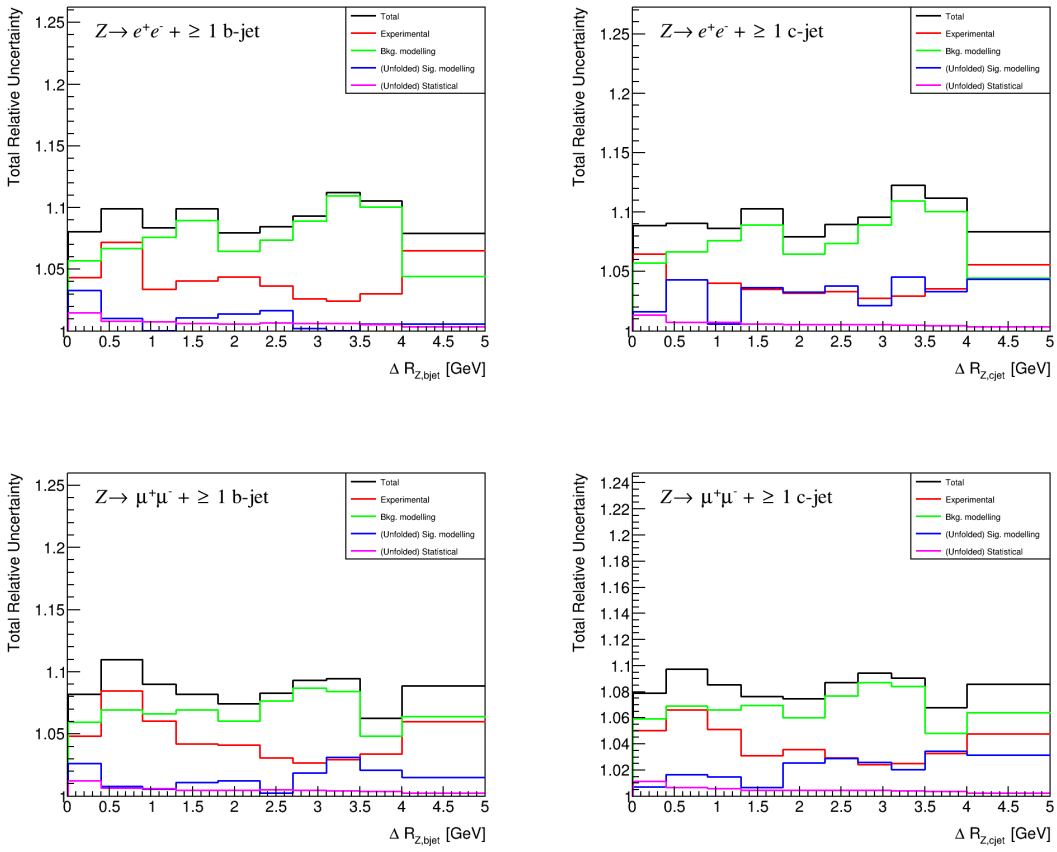

Figure 9.8: Total uncertainty on the $Z \to l^+l^- + \geq 1$ b-jet (left column) and $Z \to l^+l^- + \geq 1$ c-jet (right column) differential cross-sections as functions of ΔR(Z,heavy-flavor jet) in the electron (top row) and muon (bottom row) channels.

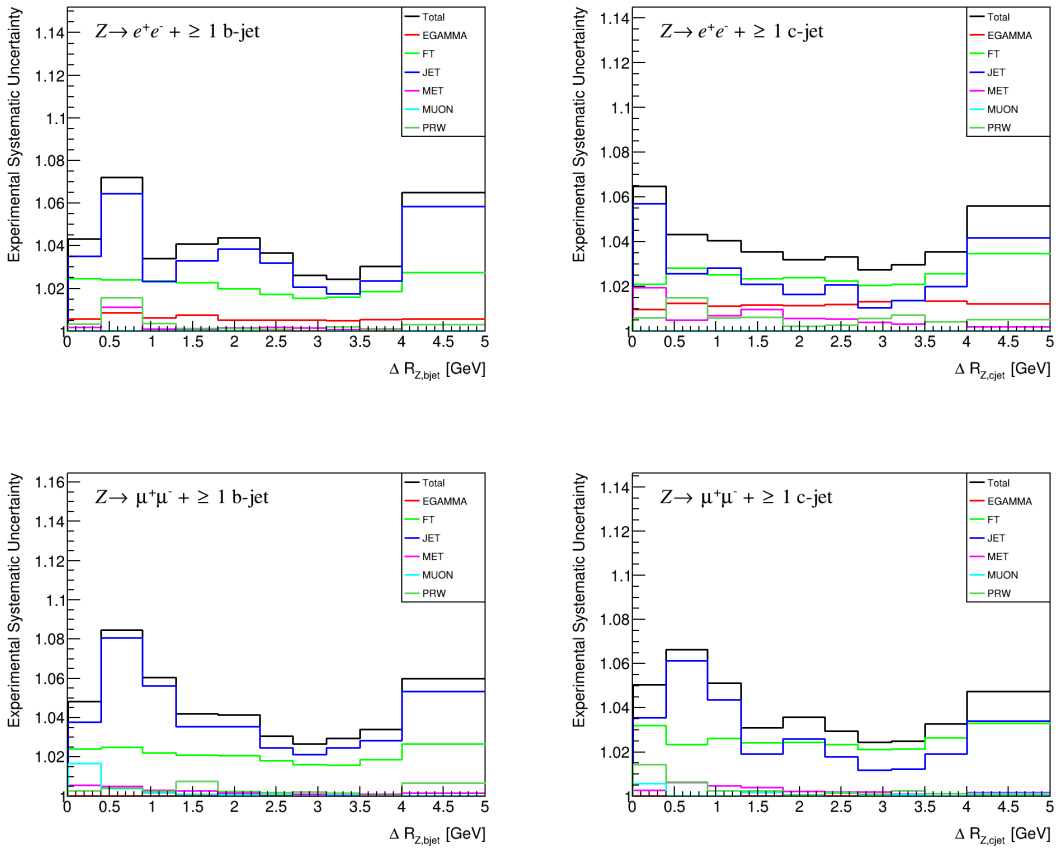

Figure 9.9: Systematic uncertainties on the $Z \to l^+l^- + \geq 1$ b-jet (left column) and $Z \to l^+l^- + \geq 1$ c-jet (right column) differential cross-sections as functions of ΔR(Z,heavy-flavor jet), in the electron (top row) and muon (bottom row) channels.

Chapter 10

Results

10.1 Inclusive cross-sections

Measured and predicted inclusive cross-sections were obtained for $Z \to l^+l^- + \geq 1$ b-jet and $Z \to l^+l^- + \geq 1$ c-jet in both the electron and muon channels; the values are presented in Tables 10.1 and 10.2 for the ≥ 1 b-jet and ≥ 1 c-jet processes, respectively. The measured inclusive cross-sections in each channel are quoted with the total systematic and (unfolded) statistical uncertainties, as described in the previous Chapter, as well as the uncertainty in the luminosity [154], which was determined using the LUCID2 luminosity monitor [155]. The MC-predicted cross-sections are given with the respective MC statistical uncertainties. These results are then combined using the method of best linear unbiased estimators (BLUE, otherwise referred to as the Gauss-Markov or Gauss method, e.g. [156]); the cross-sections are plotted in Figure 10.1.

In both lepton channels, the measured $Z \to l^+l^- + \geq 1$ b-jet is in agreement with the previous ATLAS result [5]. None of the three predicted $Z \to l^+l^- + \geq 1$ b-jet cross-sections are in agreement with the measured cross-section, but the NLO Sherpa (in disagreement with the data at the 2σ level in both lepton channels) and FxFx (1σ) predictions outperform the LO MadGraph ($\geq 4\sigma$) prediction. This suggests that the usage of NLO calculations plays an important role in the accuracy of b-jet final-state predictions; this is reinforced in particular by the contrast between the FxFx and MadGraph samples, which use the same NNPDF set and parton shower model but differ in their complete ME order.

Meanwhile, the $Z \to l^+l^- + \geq 1$ c-jet cross-section is consistently underestimated by all three predictions, with Sherpa and MadGraph in tension with the data at 1.5σ in both lepton channels and FxFx at 2σ. These discrepancies are likely the result of the over-estimation of b-jets predictions causing an under-estimation of c-jets; a dedicated inclusive-flavor Z+jets cross-section measurement recently published by ATLAS [157] shows excellent agreement between data and the

Sherpa 2.2.11 sample used.

$Z \to l^+l^- + \geq 1$ b-jet cross-section		
	σ [pb] $\pm$ syst. $\pm$ stat. $\pm$ lumi.	
	Electron	Muon
Data	$10.798 \pm 0.619 \pm 0.037 \pm 0.184$	$10.651 \pm 0.713 \pm 0.029 \pm 0.181$
	σ [pb] $\pm$ MC stat.	
	Electron	Muon
Sherpa	12.445 ± 0.022	12.557 ± 0.022
MadGraph	13.872 ± 0.059	13.589 ± 0.068
FxFx	11.657 ± 0.048	11.876 ± 0.047

Table 10.1: Inclusive $Z \to l^+l^- + \geq 1$ b-jet cross-section, in the electron and muon channels, as measured in data and predicted by the Sherpa, MadGraph, and FxFx samples.

$Z \to l^+l^- + \geq 1$ c-jet cross-section		
	σ [pb] $\pm$ syst. $\pm$ stat. $\pm$ lumi.	
	Electron	Muon
Data	$21.727 \pm 1.419 \pm 0.063 \pm 0.456$	$21.649 \pm 1.449 \pm 0.049 \pm 0.368$
	σ [pb] $\pm$ MC stat.	
	Electron	Muon
Sherpa	19.051 ± 0.033	19.258 ± 0.032
MadGraph	19.171 ± 0.075	19.216 ± 0.098
FxFx	18.023 ± 0.054	18.229 ± 0.055

Table 10.2: Inclusive $Z \to l^+l^- + \geq 1$ c-jet cross-section, in the electron and muon channels, as measured in data and predicted by the Sherpa, MadGraph, and FxFx samples.

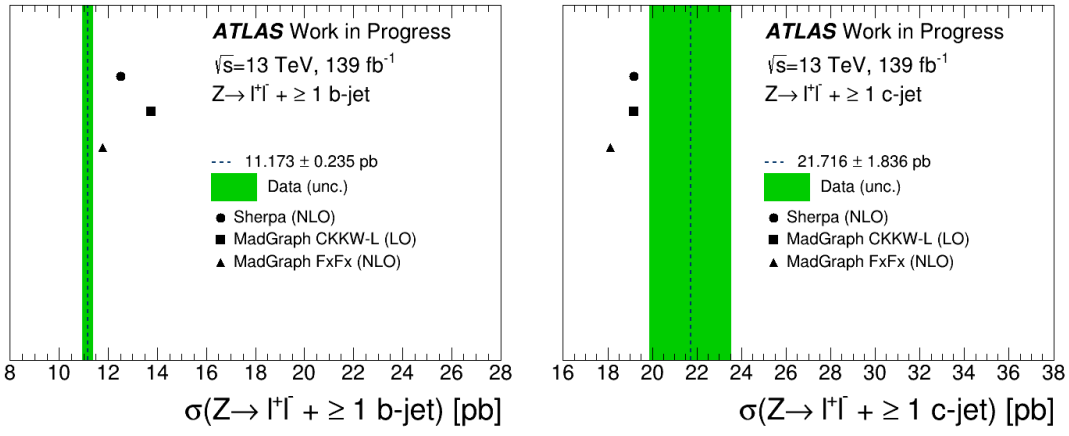

Figure 10.1: Measured and predicted inclusive $Z \to l^+l^- + \geq 1$ b-jet (left) and $Z \to l^+l^- + \geq 1$ c-jet (right) cross-sections. The measured cross-sections are compared to the cross-section predictions of the Sherpa, MadGraph, and FxFx samples. The green bands correspond to the statistical, systematic, and luminosity uncertainties on the measured cross-sections added in quadrature. Only statistical uncertainties are plotted for the theoretical predictions.

10.2 Differential cross-sections

The differential b-jet (left column) and c-jet (right column) distributions for both leptons channels (electron (top row) and muon (bottom row)) for the four observables used in this analysis are presented in Figures 10.2-10.5.

Figure 10.2 shows the differential cross-sections as functions of the leading heavy-flavor jet p_{T}, which are comparable between the two lepton channels. Past ≈ 150 GeV in leading b-jet p_{T}, the MadGraph sample is no longer in agreement with the data within the uncertainty, while the FxFx sample remains so, reinforcing the above conclusion that NLO ME calculations are highly beneficial to the modelling of heavy-flavor jets, particularly at high p_{T}. Comparing the two NLO predictions, Sherpa and FxFx, we see that the FxFx prediction of leading b-jet p_{T} more closely matches the data past ≈ 100 GeV. Given that these two predictions differ in the number of partons included at LO vs. NLO, parton shower model, and merging scheme, it will be important to investigate each choice separately in future studies. All three predictions overestimate the leading c-jet p_{T} distribution in data, but Sherpa performs the best. For both b-jet and c-jet events, all three generator predictions exhibit a different shape at low p_{T}, with underestimations at $p_{\mathrm{T}} \lesssim 30$ GeV followed by overestimation at slightly higher p_{T}.

Figure 10.3 shows the differential cross sections as functions of the Z boson p_{T}. All three generators underestimate the b-jet data beyond $p_{\mathrm{T}} \gtrsim 50$. The agreement is worse than it is for leading b-jet p_{T}; one possible explanation is that, whereas modelling of the leading b-jet p_{T} relies more heavily on a proper NLO calculation to take into account the hard b-jet emission, modelling p_{T}^{Z} requires accounting for all of the QCD radiation against which the Z recoils, and so the inclusion of an NLO ME calculation may not be sufficient for an overall agreement with the data for p_{T}^{Z}. In addition, despite all three generators being in good agreement with the data for c-jet events up to 500 GeV, the predicted shapes differ from the data for both b-jet and c-jet events in a similar manner to the differences seen in Figure 10.2: a noticeable peak appears at ≈ 20 GeV for b-jet events and ≈ 50 GeV for c-jet events. This suggests that the three generators all struggle to model the soft QCD radiative effects well.

Figure 10.4 shows the differential cross sections as functions of the leading heavy-flavor jet x_F. While the b-jet predictions are quite similar in both lepton channels, there is more variance in the muon channel c-jet distribution, including the two bins noted in the previous Chapter as having anomalously high uncertainties. MadGraph exhibits the best agreement with the data for b-jets at $x_F < 0.15$, while the FxFx sample performs better for $0.15 < x_F < 0.3$. For c-jets, on the other hand (the impetus for including this observable in the first place), FxFx is in agreement with the data at low x_F while MadGraph is not. Overall, the lack of an NLO ME calculation does not appear to be as detrimental to the MadGraph prediction as it seems it is for leading heavy-flavor

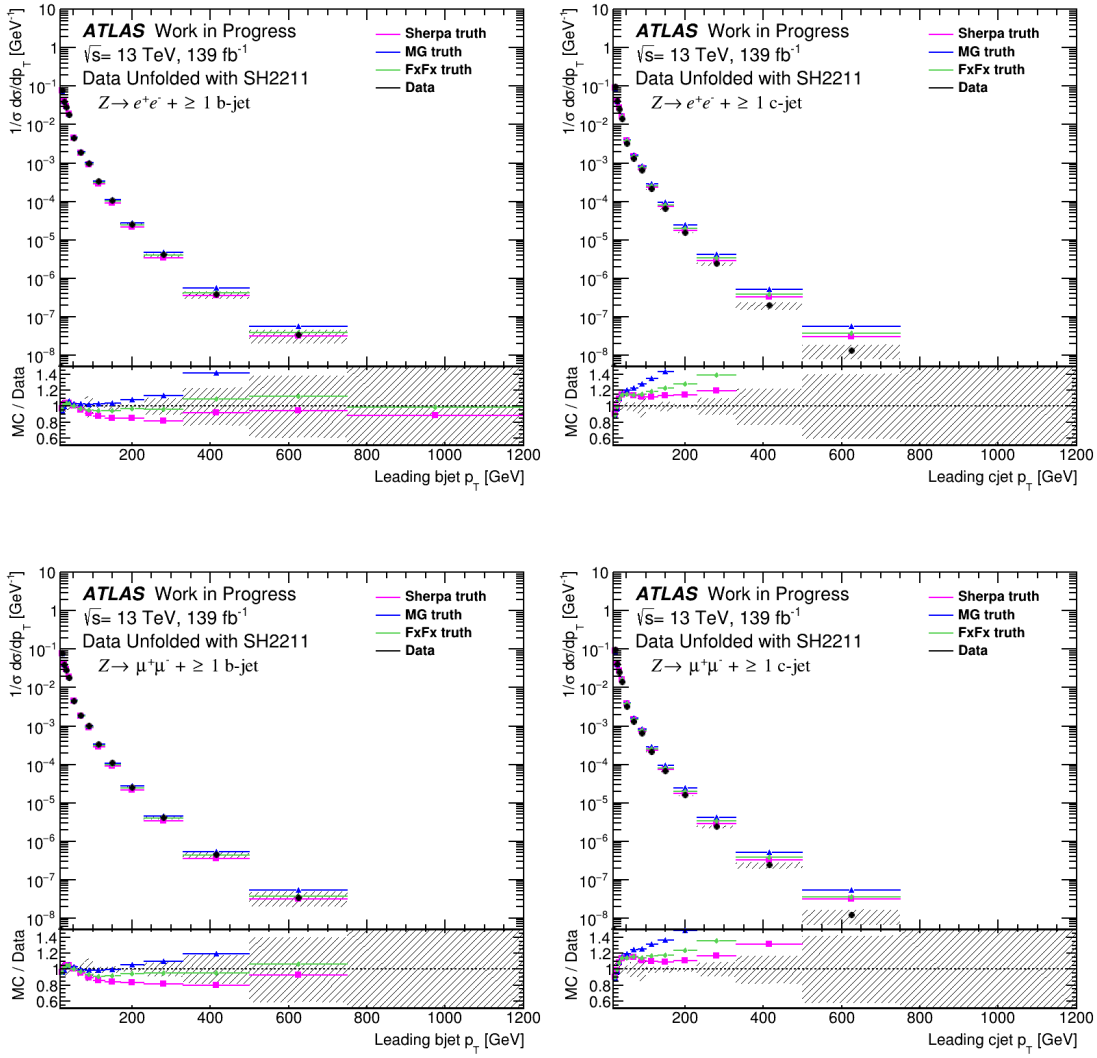

Figure 10.2: Measured $Z \to l^+l^- + \geq 1$ b-jet (left column) and $Z \to l^+l^- + \geq 1$ c-jet (right column) differential cross-sections as functions of the leading heavy-flavor jet p_{T}, in the electron (top row) and muon (bottom row) channels. The unfolded data is compared to the theoretical predictions of the Sherpa, MadGraph, and FxFx samples. The hatched bands in the lower plot correspond to the total uncertainty.

p_{T}. It is arguable that, because the heavy-flavor jet p_{T} is normalized by the scale process in the definition of x_F, x_F is less sensitive to the order of the ME calculation (which, as seen above, plays a large role in the modelling of hard QCD radiation). This is not surprising, as the x_F variable is designed to be sensitive to the PDF being used. Differences in PDFs are not tested well by the samples used in this analysis, however, given the lack of diversity in the PDF sets used (NNPDF 3.0 NNLO for Sherpa and MadGraph, NNPDF 3.1 NNLO for FxFx).[1]

Figure 10.5 shows the differential cross-sections as functions of the angular separation ΔR between the Z boson and the leading b-jets and c-jets. For both lepton channels, all three theoretical

[1]The main difference between the NNPDF 3.1 and 3.0 NNLO PDF sets, according to the collaboration [130], is the use of an inclusive HERA combined structure function measurement as opposed to separate HERA-I and ZEUS/H1HERA-II measurements.

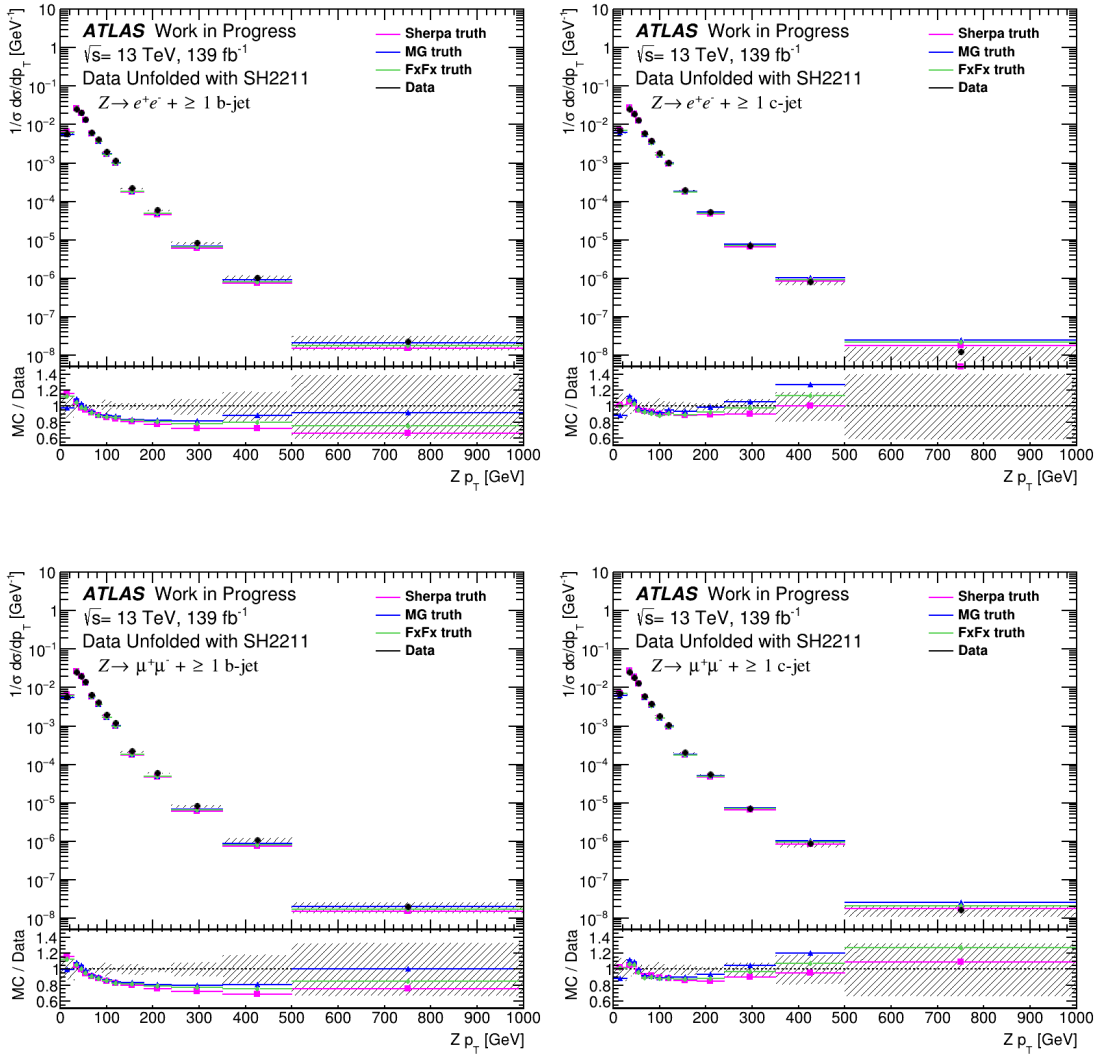

Figure 10.3: Measured $Z \to l^+l^- + \geq 1$ b-jet (left column) and $Z \to l^+l^- + \geq 1$ c-jet (right column) differential cross-sections as functions of p_{T}^{Z}, in the electron (top row) and muon (bottom row) channels. The unfolded data is compared to the theoretical predictions of the SHERPA, MADGRAPH, and FxFx samples. The hatched bands in the lower plot correspond to the total uncertainty.

predictions are in good agreement with the data for all values of ΔR, including at small separation values, where there is the greatest sensitivity to QCD radiation. The angular separation between the Z and the heavy jet is correlated to the heavy-flavor jet p_{T}, and so the excellent agreement for $\Delta R(Z,c\text{-jet})$, despite the poor agreement for the leading c-jet p_{T}, is noteworthy.

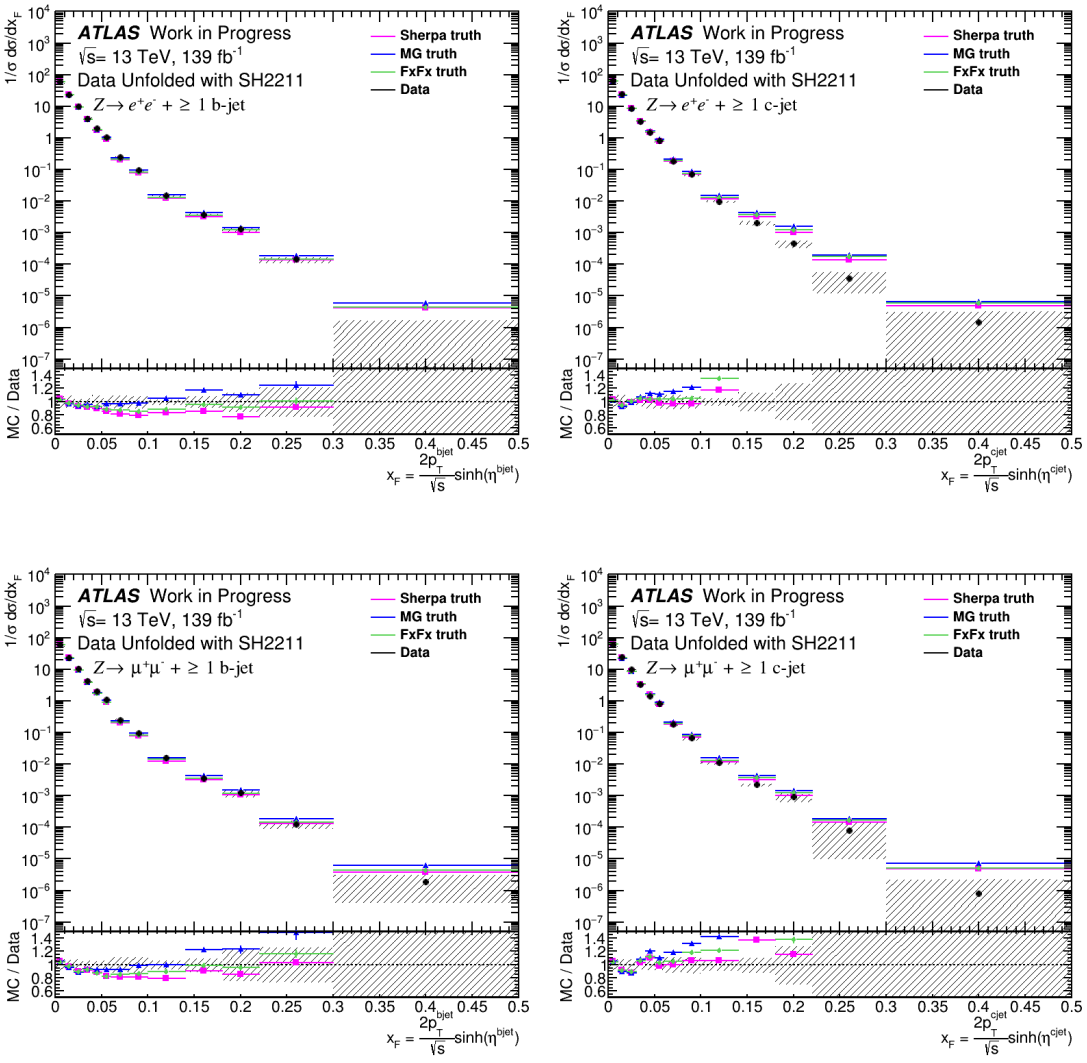

Figure 10.4: Measured $Z \to l^+l^- + \geq 1$ b-jet (left column) and $Z \to l^+l^- + \geq 1$ c-jet (right column) differential cross-sections as functions of the leading heavy-flavor jet x_F, in the electron (top row) and muon (bottom row) channels. The unfolded data is compared to the theoretical predictions of the SHERPA, MADGRAPH, and FxFx samples. The hatched bands in the lower plot correspond to the total uncertainty.

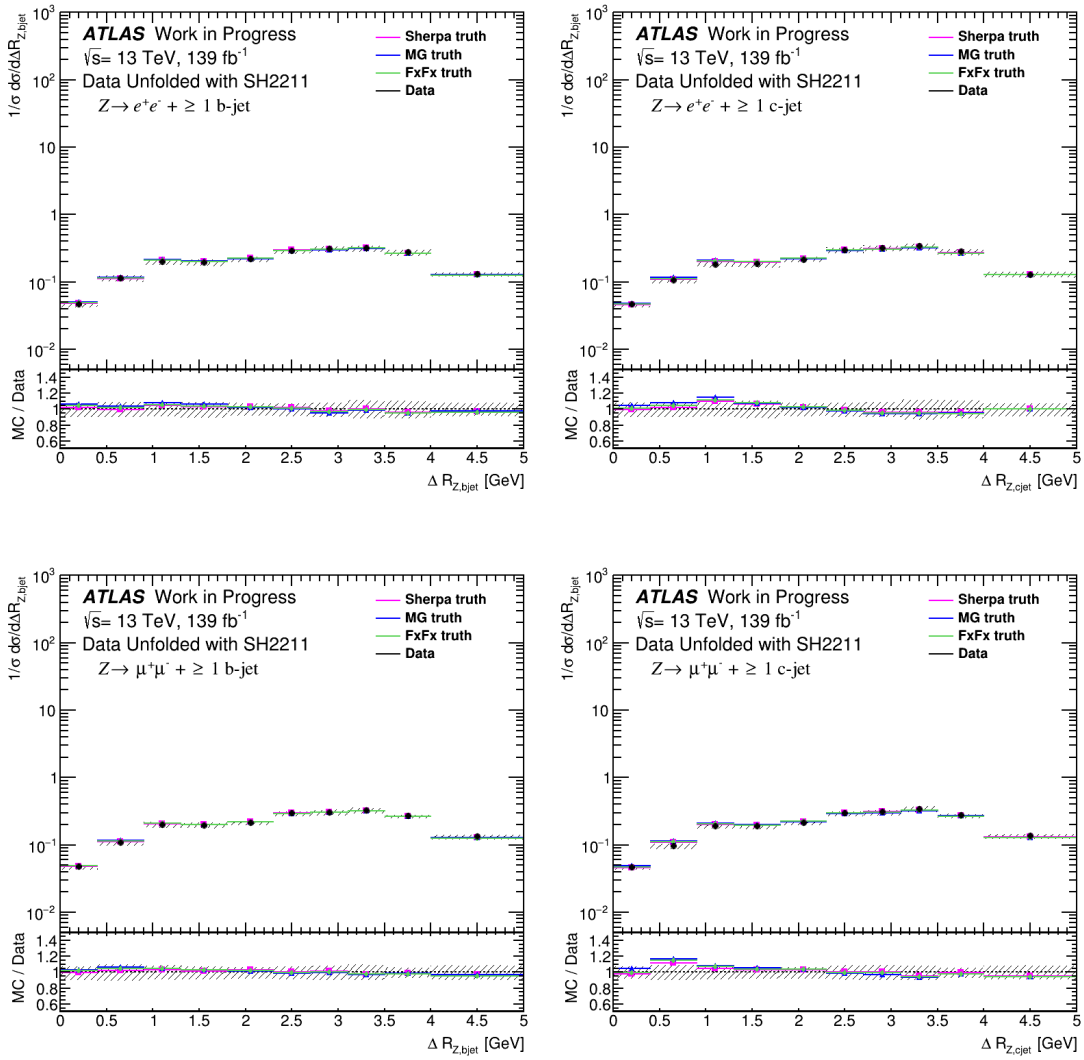

Figure 10.5: Measured $Z \to l^+l^- + \geq 1$ b-jet (left column) and $Z \to l^+l^- + \geq 1$ c-jet (right column) differential cross-sections as functions of ΔR(Z,heavy-flavor jet), in the electron (top row) and muon (bottom row) channels. The unfolded data is compared to the theoretical predictions of the SHERPA, MADGRAPH, and FxFx samples. The hatched bands in the lower plot correspond to the total uncertainty.

Chapter 11

Conclusion and Future Work

This thesis presents a measurement of the cross section of Z bosons in association with one or more heavy-flavor (b or c) jets. The purpose of the measurement is to better understand QCD phenomena and their importance in the scattering processes observed in hadron collider experiments, such as the splitting of gluons into massive quarks, the accuracy of differing pQCD predictions, and the underlying structure of the proton. Specifically, b-jet measurements help test the strength of 4FNS and 5FNS predictions in different regions of phase space, while c-jet measurements explore the possibility of intrinsic charm.

The results presented here will be included in the final publication of an ongoing ATLAS analysis focusing on Z+HF jets. The fitting and unfolding techniques, in particular, will be adapted and improved upon for the analysis, which hopes to measure not only Z+b and Z+c events but also Z+bb, Z+cc, and Z+b/Z+c ratios. One significant improvement planned for the ATLAS analysis is the estimation of the background modelling PDF uncertainty through variated weights on the existing samples, accounting for differences in PDF through 1σ variations in the eigenvalues of the Hessian matrix. By applying differing event weights to the sample samples, the PDF systematic uncertainty, taken as the relative difference between the nominal and variated distributions, will be decoupled from systematic uncertainties related to the chosen parton shower model, matching/merging scheme, etc. which are currently included in the uncertainty quoted in this thesis. Other improvements which are planned or in development include a data-driven $t\bar{t}$ estimate, to better understand the top background; additional observables to test different regions of phase space, such as those mentioned at the end of Chapter 6; and additional theoretical predictions to compare the unfolded data against, including those derived from the 4FNS; no centrally-produced 4FNS samples were available in time to be included in this analysis.

With the techniques developed and presented in this thesis, the full ATLAS measurement should provide an even more robust addition to our understanding of QCD phenomena in heavy-quark

production. Nevertheless, the analysis techniques and results discussed in this thesis, which will form the bulk of the ATLAS paper, already illuminate many avenues of future study.

Bibliography

[1] ATLAS Collaboration, "Observation of $H \to b\bar{b}$ decays and VH production with the ATLAS detector," *Phys. Lett. B*, 786: 59–86, Sep. 2018. arXiv: 1808.08238 [hep-ex].

[2] ATLAS Collaboration, "Measurement of the associated production of a Higgs boson decaying into b-quarks with a vector boson at high transverse momentum in pp collisions at $\sqrt{s}$ =13 TeV," *Phys. Lett. B*, 816: 136204, May 2021. arXiv: 2008.02508 [hep-ex].

[3] ATLAS Collaboration, "A search for $B - L$ R-parity-violating top squarks in $\sqrt{s}$=13 TeV pp collisions with the ATLAS experiment," *Phys. Rev. D*, 97(3): 032003, Feb. 2018. arXiv: 1710.05544 [hep-ex].

[4] ATLAS Collaboration, "Search for Higgs boson decays into a pair of pseudoscalar particles in the $bb\mu\mu$ final state with the ATLAS detector in pp collisions at $\sqrt{s}$=13 TeV," *Phys. Rev. D*, 105(1): 012006, Jan. 2022. arXiv: 2110.00313 [hep-ex].

[5] ATLAS Collaboration, "Measurements of the production cross-section for a Z boson in association with b-jets in proton-proton collisions at $\sqrt{s}$ = 13 TeV with the ATLAS detector," *JHEP*, 2020(7), Jul. 2020. arXiv: 2003.11960 [hep-ex].

[6] CMS Collaboration, "Measurement of the production cross section for Z + b jets in proton-proton collisions at $\sqrt{s}$ = 13 TeV," 2021. arXiv: 2112.09659 [hep-ex].

[7] MissMJ and Cush, *Standard Model of Elementary Particle Physics*, (Own work by uploader, PBS NOVA, Fermilab, Office of Science, United States Department of Energy, Particle Data Group). URL: https://en.wikipedia.org/wiki/File:Standard_Model_of_Elementary_Particles.svg.

[8] F. Halzen and A. D. Martin, *Quarks and Leptons: An Introductory Course in Modern Particle Physics*. Wiley, 1984.

[9] S. Weinberg, *The Quantum Theory of Fields*. Cambridge University Press, 1995.

[10] J. C. Ward, "An Identity in Quantum Electrodynamics," *Phys. Rev.*, 78(2): 182, 1950.

[11] Y. Takahashi, "On the Generalized Ward Identity," *Nuovo Cim*, 6(2): 371–375, 1957.

[12] M. Thomson, *Modern Particle Physics*. Cambridge University Press, 2013.

[13] R. K. Ellis, W. J. Sterling, and B. R. Webber, *QCD and Collider Physics*. Cambridge University Press, 1996.

[14] Perkins, D. H., *Introduction to High Energy Physics*. Cambridge University Press, 2000.

[15] D. J. Gross and F. Wilczek, "Ultraviolet Behavior of Non-Abelian Gauge Theories," *Phys. Rev. Lett.*, 30(26): 1343–1346, 1973.

[16] S. Bethke, "The 2009 world average of α_s," *Eur. Phys. J. C*, 64(4): 689–703, Oct. 2009. arXiv: 0908.1135 [hep-ph].

[17] C. Quigg, "Elementary Particle Physics: Discoveries, Insights, and Tools," in *Quarks, Quasars and Quandries, Ed. Aubrecht, G., American Association of Physics Teachers*, 1987.

[18] J. D. Bjorken, "Current Algebra at Small Distances," in *Proceedings of the International School of Physics Enrico Fermi Course XLI, J. Steinberger, ed., Academic Press, New York*, 1967, pp. 55–81.

[19] J. I. Friedman and H. W. Kendall, "Deep inelastic electron scattering," *Annual Review of Nuclear Science*, 22(1): 203–254, 1972. URL: https://www.annualreviews.org/doi/abs/10.1146/annurev.ns.22.120172.001223.

[20] E. D. Bloom et al., "Observed behavior of highly inelastic electron-proton scattering," *Phys. Rev. Lett.*, 23(16): 935–939, Oct. 1969. URL: https://link.aps.org/doi/10.1103/PhysRevLett.23.935.

[21] Y. L. Dokshitzer, "Calculation of the Structure Functions for Deep Inelastic Scattering and e^+e^- Annihilation by Perturbation Theory in Quantum Chromodynamics," *Sov. Phys. JETP*, 46: 641–653, 1977. URL: https://inspirehep.net/literature/126153.

[22] V. N. Gribov and L. N. Lipatov, "Deep inelastic *ep* scattering in perturbation theory," *Sov. J. Nucl. Phys.*, 15: 438–450, 1972. URL: https://inspirehep.net/literature/73449.

[23] G. Altarelli and G. Parisi, "Asymptotic freedom in parton language," *Nuclear Physics B*, 126(2): 298–318, 1977. URL: https://inspirehep.net/literature/119585.

[24] J. Gayler, "Proton structure function measurements from HERA," in *32nd International Symposium on Multiparticle Dynamics*, 2002, pp. 274–279. arXiv: 0211051 [hep-ex].

[25] A. Martin, W. Stirling, R. Thorne, and G. Watt, "Inclusive Jet Production, Parton Distributions, and the Search for New Physics," *Eur. Phys. J. C*, 63(2): 189–285, 2009. arXiv: 0901.0002 [hep-ph].

[26] D. S. et al., "Inclusive Jet Production, Parton Distributions, and the Search for New Physics," *JHEP*, 2003(10): 046, 2003. arXiv: hep-ph/0303013 [hep-ph].

[27] C. Amsler et al., (Particle Data Group) *Physics Letters B667*, 1, 2008. URL: https://pdg.lbl.gov/index.html.

[28] S. Höche, *Introduction to parton-shower event generators*, 2015. arXiv: 1411.4085 [hep-ph].

[29] V. V. Sudakov, "Vertex parts at very high-energies in quantum electrodynamics," *Sov. Phys. JETP*, 3: 65–71, 1956.

[30] pharisaical35. URL: https://pharisaical35.rssing.com/chan-18244518/all_p11.html.

[31] E. Mobs, "The CERN accelerator complex - 2019," 2019. URL: https://cds.cern.ch/record/2684277.

[32] ATLAS Collaboration, "The ATLAS Experiment at the CERN Large Hadron Collider," *JINST*, 3(08): S08003, 2008, also published by CERN Geneva in 2010. URL: https://inspirehep.net/literature/796888.

[33] CMS Collaboration, "The CMS Experiment at the CERN LHC," *JINST*, 3(08): S08004, 2008, also published by CERN Geneva. URL: https://inspirehep.net/literature/796887.

[34] LHCb Collaboration, "The LHCb Detector at the LHC," *JINST*, 3(08): S08005, 2008, also published by CERN Geneva. URL: https://inspirehep.net/literature/796248.

[35] ALICE Collaboration, "The ALICE Experiment at the CERN LHC," *JINST*, 2008, also published by CERN Geneva. URL: https://inspirehep.net/literature/796251.

[36] R. C. Manzano and X. C. Vidal, *Taking a closer look at LHC*. URL: https://www.lhc-closer.es/taking_a_closer_look_at_lhc/.

[37] CERN, *LHC the guide*. 2017. URL: https://cds.cern.ch/record/2255762/files/CERN-Brochure-2017-002-Eng.pdf.

[38] J. Stirling, "Parton luminosity and cross section plots," URL: http://www.hep.ph.ic.ac.uk/~wstirlin/plots/plots.html.

[39] OpenStax, *College Physics*. OpenStax CNX. URL: https://opentextbc.ca/openstaxcollegephysics/.

[40] ATLAS Collaboration, "Luminosity Public Results - Run 2," URL: https://twiki.cern.ch/twiki/bin/view/AtlasPublic/LuminosityPublicResultsRun2.

[41] R. Steerenberg, *LHC Report: Protons: mission accomplished*. Oct. 2018. URL: https://home.cern/news/news/physics/lhc-report-protons-mission-accomplished.

[42] ATLAS Collaboration, "Standard Model Summary Plots July 2021," 2021, ATL-PHYS-PUB-2021-032. URL: http://cdsweb.cern.ch/record/2777014.

[43] "Summary Plots for Heavy Particle Searches and Long-lived Particle Searches - March 2022," CERN, Geneva, Tech. Rep., Apr. 2022. URL: http://cds.cern.ch/record/2805714.

[44] K. Fabian, "Signal Formation Processes in Micromegas Detectors and Quality Control for large size Detector Construction for the ATLAS New Small Wheel," 2017. arXiv: 1708.01624 [physics.ins-det].

[45] ATLAS Collaboration, "Study of the material of the ATLAS inner detector for Run 2 of the LHC," *JINST*, 12(12): P12009, Dec. 2017. arXiv: 1707.02826 [hep-ex].

[46] ATLAS Collaboration, "The ATLAS Inner Detector commissioning and calibration," *Eur. Phys. J. C*, 70(3): 787–821, Aug. 2010. arXiv: 1004.5293 [hep-ex].

[47] ATLAS Collaboration, "Production and integration of the ATLAS Insertable B-Layer," *JINST*, 13(05): T05008, May 2018. arXiv: 1803.00844 [physics.ins-det].

[48] ATLAS Collaboration, "The Silicon microstrip sensors of the ATLAS semiconductor tracker," *Nucl. Instrum. Meth. A*, 578(1): 98–118, 2007. URL: https://inspirehep.net/literature/755578.

[49] ATLAS Collaboration, "Operation and performance of the ATLAS semiconductor tracker," *JINST*, 9(08): P08009, Aug. 2014. arXiv: 1404.7473 [hep-ex].

[50] ATLAS TRT Collaboration, "The ATLAS Transition Radiation Tracker (TRT) proportional drift tube: Design and performance," *JINST*, 3: P02013, 2008. URL: https://inspirehep.net/literature/780935.

[51] ATLAS Collaboration, "Performance of the ATLAS Transition Radiation Tracker in Run 1 of the LHC: tracker properties," *JINST*, 12(05): P05002, May 2017. arXiv: 1702.06473 [hep-ex].

[52] J. Pequenão, "Computer generated image of the ATLAS inner detector," 2008. URL: https://cds.cern.ch/record/1095926.

[53] A. Yamamoto et al., "The ATLAS central solenoid," *Nucl. Instrum. Meth. A*, 584(1): 53–74, 2008. URL: https://www.sciencedirect.com/science/article/pii/S0168900207020414.

[54] N. Ch. Benekos et al., "ATLAS Inner Detector Performance," CERN, Geneva, Tech. Rep., 12 2003. URL: https://cds.cern.ch/record/688762.

[55] ATLAS Collaboration, *ATLAS Liquid-Argon Calorimeter: Technical Design Report*, Technical design report. ATLAS. Geneva: CERN, 1996. URL: http://cds.cern.ch/record/331061.

[56] ATLAS Collaboration, *ATLAS Tile Calorimeter: Technical Design Report*, Technical design report. ATLAS. Geneva: CERN, 1996. URL: http://cds.cern.ch/record/331062.

[57] ATLAS Collaboration, "Study of the material of the ATLAS inner detector for run 2 of the LHC," *JINST*, 12(12): P12009, Dec. 2017. arXiv: 1707.02826 [hep-ex].

[58] J. Pequenão, "Computer Generated image of the ATLAS calorimeter," 2008. URL: https://cds.cern.ch/record/1095927.

[59] F. Cavallari, "Performance of calorimeters at the LHC," *J. Phys. Conf. Ser.*, 293: 012001, Apr. 2011. URL: https://doi.org/10.1088/1742-6596/293/1/012001.

[60] ATLAS Collaboration, *ATLAS Muon Spectrometer: Technical Design Report*. Geneva: CERN, 1997. URL: https://cds.cern.ch/record/331068.

[61] J. Pequenão, "Computer Generated image of the ATLAS Muon subsystem," 2008. URL: https://cds.cern.ch/record/1095929.

[62] ATLAS Collaboration, "Resolution of the ATLAS muon spectrometer monitored drift tubes in LHC Run 2," *JINST*, 14(09): P09011, Sep. 2019. arXiv: 1906.12226 [hep-ex].

[63] A. Salvucci, "Measurement of muon momentum resolution of the ATLAS detector," *EPJ Web of Conferences*, 28: 12039, 2012. arXiv: 1201.4704 [physics.ins-det].

[64] ATLAS Collaboration, "Standalone Vertex Finding in the ATLAS Muon Spectrometer," *JINST*, 9(02): P02001, 2014. arXiv: 1311.7070 [physics.ins-det].

[65] A. Ruiz Martínez and the ATLAS Collaboration, *The Run-2 ATLAS Trigger System*. Geneva: CERN, Feb. 2016. URL: https://cds.cern.ch/record/2133909.

[66] ATLAS Collaboration, "Operation of the ATLAS trigger system in Run 2," *JINST*, 15(10): P10004, Oct. 2020. arXiv: 2007.12539 [physics.ins-det].

[67] ATLAS Collaboration, "Approved Plots DAQ," URL: https://twiki.cern.ch/twiki/bin/view/AtlasPublic/ApprovedPlotsDAQ.

[68] ATLAS Collaboration, *ATLAS Level-1 Trigger: Technical Design Report*, Technical design report. ATLAS. Geneva: CERN, 1998. URL: http://cds.cern.ch/record/381429.

[69] P. Jenni, M. Nessi, M. Nordberg, and K. Smith, *ATLAS High-Level Trigger, Data-Acquisition and Cntrols: Technical Design Report*, Technical design report. ATLAS. Geneva: CERN, 2003. URL: https://cds.cern.ch/record/616089.

[70] J. Pequenão, "Event Cross Section in a computer generated image of the ATLAS detector," 2008. URL: https://cds.cern.ch/record/1096081.

[71] G. Pownall, "QCD analysis and measurement of W boson production in association with jets at the ATLAS detector," Presented 15 Jan 2021. URL: https://cds.cern.ch/record/2774953.

[72] ATLAS Collaboration, "Reconstruction of primary vertices at the ATLAS experiment in Run 1 proton–proton collisions at the LHC," *Eur. Phys. J. C*, 77(332), 2017. arXiv: 1611.10235 [physics.ins-det].

[73] ATLAS Collaboration, "Performance of the ATLAS track reconstruction algorithms in dense environments in LHC Run 2," *Eur. Phys. J. C*, 77(673), 2017. arXiv: 1704.07983 [hep-ex].

[74] R. Frühwirth, "Application of Kalman filtering to track and vertex fitting," *Nucl. Instrum. Meth. A*, 262(2): 444–450, 1987. URL: https://www.sciencedirect.com/science/article/pii/0168900287908874.

[75] ATLAS Collaboration, "Alignment of the ATLAS Inner Detector in Run-2," *Eur. Phys. J. C*, 80(1994), 2020. arXiv: 2007.07624 [hep-ex].

[76] ATLAS Collaboration, "Topological cell clustering in the ATLAS calorimeters and its performance in LHC Run 1," *Eur. Phys. J. C*, 77(7), Oct. 2017. arXiv: 1603.02934 [hep-ex].

[77] ATLAS Collaboration, "Jet energy measurement with the ATLAS detector in proton-proton collisions at $\sqrt{s} = 7$ TeV," *Eur. Phys. J. C*, 73(3): (2304), Mar. 2013. arXiv: 1112.6426 [hep-ex].

[78] P. Giovannini, "Local Hadronic Calibration with ATLAS," 2010. URL: https://cds.cern.ch/record/1284955.

[79] S. Agostinelli et al., "GEANT4–a simulation toolkit," *Nucl. Instrum. Meth. A*, 506(3): 250–303, 2003. URL: https://inspirehep.net/literature/593382.

[80] ATLAS Collaboration, "Muon reconstruction performance of the ATLAS detector in proton–proton collision data at $\sqrt{s}$ =13 TeV," *Eur. Phys. J. C*, 76(292), 2016. arXiv: 1603.05598 [hep-ex].

[81] ATLAS Collaboration, "Muon reconstruction and identification efficiency in ATLAS using the full Run 2 *pp* collision data set at $\sqrt{s}$ =13 TeV," *Eur. Phys. J. C*, 81(578), 2021. arXiv: 2012.00578 [hep-ex].

[82] J. Illingworth and J. Kittler, "A survey of the Hough Transform," *Computer Vision, Graphics, and Image Processing*, 44(1): 87–116, 1988. URL: https://www.sciencedirect.com/science/article/pii/S0734189X88800331.

[83] Y. He and Y. Okumura, "An extension of Associative Memory approach to tracking with a drift-tube detector using timing information and its demonstration for HL-LHC ATLAS muon trigger," CERN, Geneva, Tech. Rep., Oct. 2020. URL: https://cds.cern.ch/record/2742382.

[84] ATLAS Collaboration, "Performance of the ATLAS muon triggers in Run 2," *JINST*, 15(9): P09015, Sep. 2020. arXiv: 2004.13447 [physics.ins-det].

[85] 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," *Eur. Phys. J. C*, 79(639), 2019. arXiv: 1902.04655 [physics.ins-det].

[86] ATLAS Collaboration, "Performance of electron and photon triggers in ATLAS during LHC Run 2," *Eur. Phys. J. C*, 80(47), 2020. arXiv: 1909.00761 [hep-ex].

[87] W. L. et al., "Calorimeter Clustering Algorithms: Description and Performance," CERN, Geneva, Tech. Rep., Apr. 2008. URL: https://cds.cern.ch/record/1099735.

[88] ATLAS Collaboration, "Electron and photon performance measurements with the ATLAS detector using the 2015—2017 LHC proton-proton collision data," 14(12): P12006, Dec. 2019. arXiv: 1908.00005 [hep-ex].

[89] ATLAS Collaboration, "Improved electron reconstruction in ATLAS using the Gaussian Sum Filter-based model for bremsstrahlung," 2021, ATLAS-CONF-2012-047. URL: https://cds.cern.ch/record/1449796.

[90] A. Hoecker et al., *TMVA - Toolkit for Multivariate Data Analysis*, 2007. arXiv: physics/0703039 [physics.data-an].

[91] ATLAS Collaboration, "Electron and photon energy calibration with the ATLAS detector using 2015–2016 LHC proton-proton collision data," *JINST*, 14(03): P03017, 2019. arXiv: 1812.03848 [hep-ex].

[92] M. Cacciari, G. P. Salam, and G. Soyez, "The anti-k_t jet clustering algorithm," *JHEP*, 2008(04): 063, 2008. arXiv: 0802.1189 [hep-ph].

[93] M. Cacciari, G. P. Salam, and G. Soyez, "FastJet user manual," *Eur. Phys. J. C*, 72(3), Mar. 2012. arXiv: 1111.6097 [hep-ph].

[94] P. Schieferdecker, "Jet Algorithms," 2009. URL: https://twiki.cern.ch/twiki/pub/Sandbox/Lecture/Philipp_Schieferdeckers_Lecture.pdf.

[95] ATLAS Collaboration, "Jet reconstruction and performance using particle flow with the ATLAS Detector," *Eur. Phys. J. C*, 77(7), Oct. 2017. arXiv: 1703.10485 [hep-ex].

[96] ATLAS Collaboration, "Jet energy scale and resolution measured in proton-proton collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector," *Eur. Phys. J. C*, 81(689), 2021. arXiv: 2007.02645 [hep-ex].

[97] R. A. Ong, "Measurement of the B hadron lifetime," Jan. 1987. URL: https://www.osti.gov/biblio/5014940.

[98] N. Bartosik, *B-tagging diagram*. 2016, (Own work by uploader). URL: https://commons.wikimedia.org/wiki/File:B-tagging_diagram.png.

[99] ATLAS Collaboration, "ATLAS b-jet identification performance and efficiency measurement with $t\bar{t}$ events in pp collisions at $\sqrt{s} = 13$ TeV," *Eur. Phys. J. C*, 79(11), Sep. 2019. arXiv: 1907.05120 [hep-ex].

[100] ATLAS Collaboration, "Monte Carlo to Monte Carlo scale factors for flavour tagging efficiency calibration," 2020. URL: https://cds.cern.ch/record/2718610/files/ATL-PHYS-PUB-2020-009.pdf.

[101] S. Dawson et al., *Handbook of LHC Higgs Cross Sections: 1. Inclusive Observables*, 2011. arXiv: 1101.0593 [hep-ph].

[102] ATLAS Collaboration, "Measurement of the associated production of a Higgs boson decaying into b-quarks with a vector boson at high transverse momentum in pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector," *Physics Letters B*, 816: 136204, 2021. arXiv: 2008.02508 [hep-ex].

[103] F. Maltoni, G. Ridolfi, and M. Ubiali, "b-initiated processes at the LHC: a reappraisal," *JHEP*, 2012(7), 2012. arXiv: 1203.6393 [hep-ph].

[104] ATLAS Collaboration, "Measurement of differential cross sections of isolated-photon plus heavy-flavour jet production in pp collisions at $\sqrt{s}$=8 TeV using the ATLAS detector," *Phys. Lett. B*, 776: 295–317, Jan. 2018. arXiv: 1710.09560 [hep-ex].

[105] S. J. Brodsky et al., "The Intrinsic Charm of the Proton," *Phys. Lett. B*, 93(4): 451–455, 1980. URL: https://inspirehep.net/literature/152807.

[106] S. J. Brodsky, C. Peterson, and N. Sakai, "Intrinsic heavy-quark states," *Phys. Rev. D*, 23(11): 2745–2757, Jun. 1981. URL: https://link.aps.org/doi/10.1103/PhysRevD.23.2745.

[107] P-H. Beauchemin et al., "Search for intrinsic charm in vector boson production accompanied by heavy-flavor jets," *Phys. Rev. D*, 92(3): 034014, Aug. 2015. arXiv: 1410.2616 [hep-ph].

[108] ATLAS Collaboration, *Good Run Lists for Analysis Run 2*. URL: https://twiki.cern.ch/twiki/bin/viewauth/AtlasProtected/GoodRunListsForAnalysisRun2.

[109] JVT Collaboration, "JVTCalibration," URL: https://twiki.cern.ch/twiki/bin/view/AtlasProtected/PileupJetRecommendations.

[110] ATLAS Collaboration, "Performance of pile-up mitigation techniques for jets in pp collisions at $\sqrt{s} = 8$ TeV using the ATLAS detector," *Eur. Phys. J. C*, 76(11), Oct. 2016. arXiv: 1510.03823 [hep-ex].

[111] "Sherpa Manual Version 2.2.11," Accessed 17 May 2022. URL: https://sherpa.hepforge.org/doc/SHERPA-MC-2.2.11.html.

[112] ATLAS Collaboration, "Modelling and computational improvements to the simulation of single vector-boson plus jet processes for the atlas experiment," 2021. arXiv: 2112.09588 [hep-ex].

[113] T. Gleisberg and S. Höche, "Comix, a new matrix element generator," *JHEP*, 2008(12): 039, Dec. 2008. arXiv: 0808.3674 [hep-ph].

[114] F Buccioni et al., "OpenLoops 2," *Eur. Phys. J. C*, 79(10), Oct. 2019. arXiv: 1907.13071 [hep-ph].

[115] R. D. Ball et al., "Parton distributions for the LHC run II," *JHEP*, 2015(4), Apr. 2015. arXiv: 1410.8849 [hep-ph].

[116] S. Schumann and F. Krauss, "A parton shower algorithm based on Catani-Seymour dipole factorisation," *JHEP*, 2008(03): 038, Mar. 2008. arXiv: 0709.1027 [hep-ph].

[117] S. Catani and M. Seymour, "A general algorithm for calculating jet cross sections in NLO QCD," *Nuclear Physics B*, 485(1-2): 291–419, Feb. 1997. arXiv: hep-ph/9605323.

[118] S. Catani et al., "The dipole formalism for next-to-leading order QCD calculations with massive partons," *Nuclear Physics B*, 627(1-2): 189–265, Apr. 2002. arXiv: hep-ph/0201036.

[119] S. Höche et al., "A critical appraisal of NLO+PS matching methods," *JHEP*, 2012(9), Sep. 2012. arXiv: 1111.1220 [hep-ph].

[120] S. Höche et al., "QCD matrix elements + parton shower: The NLO case," *JHEP*, 2013(04), Apr. 2013. arXiv: 1207.5030 [hep-ph].

[121] 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," *JHEP*, 2014(7), Jul. 2014. arXiv: 1405.0301 [hep-ph].

[122] J. Alwall et al., "MadGraph 5: Going beyond," *JHEP*, 2011(6), Jun. 2011. arXiv: 1106.0522 [hep-ph].

[123] S. Frixione and B. R. Webber, "Matching NLO QCD computations and parton shower simulations," *JHEP*, 2002(06): 029, Jun. 2002. arXiv: hep-ph/0204244 [hep-ph].

[124] T. Sjöstrand, S. Mrenna, and P. Skands, "A brief introduction to PYTHIA 8.1," *Computer Physics Communications*, 178(11): 852–867, Jun. 2008. arXiv: 0710.3820 [hep-ph].

[125] "ATLAS Pythia 8 tunes to 7 TeV data," CERN, Geneva, Tech. Rep., Sep. 2014. URL: https://cds.cern.ch/record/1966419.

[126] R. D. Ball et al., "Parton distributions with LHC data," *Nuclear Physics B*, 867(2): 244–289, Feb. 2013. arXiv: 1207.1303 [hep-ph].

[127] L. Lönnblad, "Correcting the colour-dipole cascade model with fixed order matrix elements," *JHEP*, 2002(05): 046, May 2002. arXiv: hep-ph/0112284 [hep-ph].

[128] L. Lönnblad and S. Prestel, "Matching tree-level matrix elements with interleaved showers," *JHEP*, 2012(3), Mar. 2012. arXiv: 1109.4829 [hep-ph].

[129] V. Bertone et al., "Illuminating the photon content of the proton within a global PDF analysis," *SciPost Physics*, 5(1), Jul. 2018. arXiv: 1712.07053 [hep-ph].

[130] R. D. Ball et al., "Parton distributions from high-precision collider data," *Eur. Phys. J. C*, 77(10), Oct. 2017. arXiv: 1706.00428v2 [hep-ph].

[131] R. Frederix and S. Frixione, "Merging meets matching in MC@NLO," *JHEP*, 2012(12), Dec. 2012. arXiv: 1209.6215 [hep-ph].

[132] S. Frixione, G. Ridolfi, and P. Nason, "A positive-weight next-to-leading-order Monte Carlo for heavy flavour hadroproduction," *Journal of High Energy Physics*, 2007(09): 126, Sep. 2007. arXiv: 0707.3088 [hep-ph].

[133] S. Alioli et al., "A general framework for implementing NLO calculations in shower monte carlo programs: The POWHEG BOX," *JHEP*, 2010(06), Jun. 2010. arXiv: 1002.2581 [hep-ph].

[134] M. Czakon and A. Mitov, "Top++: A program for the calculation of the top-pair cross-section at hadron colliders," *Computer Physics Communications*, 185(11): 2930–2938, Nov. 2014. arXiv: 1112.5675 [hep-ph].

[135] S. Frixione et al., "Single-top hadroproduction in association with a W boson," *JHEP*, 2008(07): 029, Jul. 2008. arXiv: 0805.3067 [hep-ph].

[136] E. Bothmann et al., "Event generation with Sherpa 2.2," *SciPost Physics*, 7(3), 2019. arXiv: 1905.09127 [hep-ph].

[137] ATLAS Collaboration, "The Pythia 8 A3 tune description of ATLAS minimum bias and inelastic measurements incorporating the Donnachie-Landshoff diffractive model," CERN, Geneva, Tech. Rep., Aug. 2016. URL: https://cds.cern.ch/record/2206965.

[138] ATLAS Collaboration, "Characterisation and mitigation of beam-induced backgrounds observed in the ATLAS detector during the 2011 proton-proton run," *JINST*, 8(07): P07004, Jul. 2013. arXiv: 1303.0223 [hep-ex].

[139] ATLAS Collaboration, "Tagging and suppression of pileup jets with the ATLAS detector," 2014, ATLAS-CONF-2014-018. URL: https://cds.cern.ch/record/1700870.

[140] ATLAS Collaboration, "Recommendations of the Physics Objects and 2 Analysis Harmonisation Study Groups 2014," 2014, ATLAS Internal note. URL: https://cds.cern.ch/record/1700870.

[141] *RooFit*. URL: https://root.cern.ch/doc/master/group__Roofitmain.html.

[142] V. Blobel, *An Unfolding Method for High Energy Physics Experiments*, 2002. arXiv: hep-ex/0208022 [hep-ex].

[143] F. Spano, "Unfolding in particle physics: a window on solving inverse problems," *EPJ Web Conf.*, T. Delemontex and A. Lucotte, Eds., 55: 03002, 2013. URL: https://inspirehep.net/literature/1245933.

[144] G. Cowan, *Statistical Data Analysis*. Clarendon Press, 1998.

[145] L. Brenner et al., "Comparison of unfolding methods using RooFitUnfold," 2019. arXiv: 1910.14654 [physics.data-an].

[146] W. H. Richardson, "Bayesian-Based Iterative Method of Image Restoration," *J. Opt. Soc. Am.*, 62(1): 55–59, Jan. 1972. URL: http://opg.optica.org/abstract.cfm?URI=josa-62-1-55.

[147] L. B. Lucy, "An iterative technique for the rectification of observed distributions," *Astronomical Journal*, 79(6): 745–754, Jun. 1974. URL: https://ui.adsabs.harvard.edu/abs/1974AJ.....79..745L.

[148] G. D'Agostini, "A multidimensional unfolding method based on bayes' theorem," *Nucl. Instrum. Meth. A*, 362(2): 487–498, 1995. URL: https://inspirehep.net/literature/374574.

[149] G. Cowan, "A survey of unfolding methods for particle physics," *Conf. Proc. C*, 0203181: 248–257, 2002. URL: https://inspirehep.net/literature/599644.

[150] A. Höcker and V. Kartvelishvili, "SVD Approach to Data Unfolding," *Nucl. Instrum. Meth. A*, 372(3): 469–481, Apr. 1996. arXiv: hep-ph/9509307 [hep-ph].

[151] Georg-Johann, *Singular-Value-Decomposition*. 2010, (Own work by uploader). URL: https://commons.wikimedia.org/wiki/File:Singular-Value-Decomposition.svg.

[152] T. Adye, "Unfolding algorithms and tests using RooUnfold," 2011. arXiv: 1105.1160 [physics.data-an].

[153] C. Bittrich, "Evidence for Scattering of Electroweak Gauge Bosons in the $W^{\pm}Z$ Channel with the ATLAS Detector at the Large Hadron Collider," Presented 14 Jul 2020. URL: https://cds.cern.ch/record/2719126.

[154] ATLAS Collaboration, "Luminosity determination in pp collisions at $\sqrt{s} = 13$ TeV using the ATLAS detector at the LHC," Jun. 2019, ATLAS-CONF-2019-021. URL: https://inspirehep.net/literature/1737864.

[155] G. Avoni et al., "The new LUCID-2 detector for luminosity measurement and monitoring in ATLAS," *JINST*, 13(07): P07017, 2018. URL: https://inspirehep.net/literature/1684422.

[156] J. Davidson, *Econometric Theory*. Wiley-Blackwell, 2000, ch. Statistical Analysis of the Regression Model, pp. 17–36.

[157] ATLAS Collaboration, "Cross-section measurements for the production of a Z boson in association with high-transverse-momentum jets in pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector," 2022. arXiv: 2205.02597 [hep-ex].