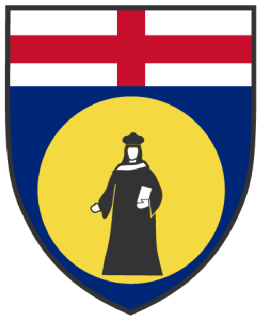
Università di Genova

Dipartimento di Fisica

Dottorato in Fisica e Nanoscienze – XXXVII Ciclo



An exploration of heavy flavour jets:

From phenomenology to precision measurements at the ATLAS Experiment

Candidato:

Alberto Lorenzo Rescia

Tutor:

Prof. Federico Sforza

Revisori Esterni:

Prof.ssa Evelin Meoni

Prof. Matt LeBlanc

Aprile 2026

This work received funding from the Deutsches Elektronen-Synchrotron.

Dai diamanti non nasce niente
Dal letame nascono i fior

Fabrizio de André

Abstract

This thesis provides a panoramic view of heavy flavour jets at hadron colliders, diving into their phenomenology, identification, and production. After a brief overview of the Standard Model and the phenomenology of Quantum Chromodynamics, we begin by considering the problem of defining heavy flavour jets in higher order QCD predictions, and analyse a set of possible solutions. We then turn to more experimental aspects, describing the ATLAS Detector and the algorithms employed by the ATLAS Collaboration to identify heavy flavour jets within the detector. Lastly, we present a measurement of Z boson production in association with a pair of bottom or light jets, focusing both on jet kinematics, as well as jet substructure. In particular, we consider the ratio of the Lund Jet Plane for these jets, uncovering the dead cone effect for the first time at the ATLAS Experiment.

Contents

Introduction **13**

1 V+jets and jet substructure **15**
1.1 The Standard Model . . . 15
1.1.1 A brief overview of the universe . . . 15
1.1.2 The Standard Model in a nutshell . . . 17
1.1.3 Spontaneous symmetry breaking . . . 21
1.1.4 Fermion Masses . . . 23
1.2 Physics at colliders . . . 23
1.2.1 Anatomy of a cross section . . . 23
1.2.2 Running of the coupling constant and perturbative QCD . . . 24
1.2.3 QCD radiation in the initial and final state . . . 25
1.2.4 Parton Distribution Functions . . . 27
1.2.5 Final state evolution . . . 28
1.2.6 Monte Carlo event generators . . . 29
1.2.7 Hadronisation . . . 30
1.3 Jets . . . 30
1.3.1 Jet algorithms . . . 31
1.3.2 Jets and IRC safety . . . 33
1.4 Jet substructure . . . 33
1.4.1 Jet grooming . . . 33
1.4.2 Heavy flavour and JSS . . . 34
1.4.3 The primary Lund Jet Plane . . . 35
1.4.4 Jet angularities . . . 37
1.5 State-of-the-art measurements . . . 38
1.5.1 Z+jets . . . 38
1.5.2 The Lund Jet Plane . . . 43
1.5.3 CMS measurement of groomed jet substructure and fragmentation observables . . . 49
1.5.4 Measurements of jet angularities . . . 52

2 Jet flavour labelling **55**
2.1 Current flavour labelling strategies . . . 56
2.2 The next generation of algorithms . . . 58
2.2.1 SoftDrop flavour (SDF) . . . 58
2.2.2 Flavoured anti-k_t (CMP) . . . 58
2.2.3 Flavour dressing (GHS) . . . 59

2.2.4 Interleaved flavour neutralisation (IFN) . . . 60
2.3 Flavoured jet algorithms: A comparative study . . . 60
2.3.1 Bridging the gap from theory to experiment . . . 63
2.3.2 Flavour recombination scheme . . . 63
2.3.3 Neutralisation of flavour from gluon splittings . . . 64
2.3.4 NNLO fixed-order predictions . . . 67
2.3.5 Impact of hadron reconstruction on flavoured jet definitions . . . 70
2.3.6 Consequences for flavour tagging . . . 71
2.3.7 Jet substructure . . . 72
2.4 Conclusions . . . 72

3 The ATLAS Detector **77**
3.1 The Large Hadron Collider . . . 77
3.1.1 Luminosity . . . 78
3.1.2 Pile-up . . . 79
3.2 The ATLAS Detector . . . 80
3.2.1 Coordinate system . . . 80
3.2.2 Magnets . . . 81
3.2.3 Inner Detector . . . 82
3.2.4 Calorimeters . . . 83
3.2.5 Muon Spectrometer . . . 83
3.2.6 Trigger . . . 84
3.2.7 Detector operations . . . 85
3.2.8 Putting it all together . . . 85
3.3 Object reconstruction . . . 86
3.3.1 Tracks . . . 86
3.3.2 Vertices . . . 87
3.3.3 Electrons . . . 88
3.3.4 Muons . . . 88
3.3.5 Missing transverse momentum . . . 89
3.3.6 Jets . . . 89
3.3.7 Jet calibration . . . 90
3.3.8 Pile-up jet suppression . . . 92

4 Heavy flavour tagging in ATLAS **93**
4.1 Flavour tagging . . . 93
4.1.1 Flavour tagging in a nutshell . . . 93
4.1.2 DL1r . . . 95
4.2 The Xbb tagger . . . 98
4.3 Calibration . . . 103
4.3.1 $g \to b\bar{b}$ calibration . . . 103
4.3.2 Efficiency scale factor fit strategy . . . 104
4.3.3 Systematic uncertainties . . . 107
4.3.4 SV1 mass fit . . . 111
4.3.5 Final comments . . . 116
4.4 Conclusions . . . 116

5 Heavy flavour and jet substructure measurements in Z+jets ATLAS data **121**
5.1 Motivation and analysis outline . . . 122
5.2 Observables of interest . . . 123
5.3 Data and Monte Carlo samples . . . 124
5.3.1 Signal and background estimation . . . 124
5.4 Event selection . . . 126
5.5 Fiducial selections . . . 127
5.6 Reconstruction of JSS observables . . . 131
5.6.1 Heavy hadron reconstruction . . . 135
5.6.2 One-dimensional representation of the Lund Plane . . . 140
5.7 Backgrounds . . . 142
5.7.1 tt rejection . . . 142
5.7.2 tt data-driven estimation . . . 146
5.7.3 Flavour fit . . . 150
5.7.4 Background estimation . . . 154
5.8 Systematics . . . 159
5.8.1 Detector-level systematic uncertainties . . . 159
5.8.2 Background systematic uncertainties . . . 160
5.9 Detector-level results . . . 160
5.10 Unfolding . . . 171
5.10.1 Unfolding the Lund Jet Plane . . . 171
5.10.2 Unfolding systematic uncertainties . . . 174
5.11 Results . . . 177
5.11.1 Resolved kinematic observables . . . 177
5.11.2 Boosted kinematic observables . . . 180
5.11.3 Jet substructure observables . . . 182
5.11.4 The dead cone effect . . . 187

Conclusions **191**

Appendices **193**

A Low-level objects for Lund Jet Plane reconstruction **193**
A.1 Track Calo Clusters . . . 193
A.2 Unified Flow Objects . . . 193
A.3 Lund Plane reconstruction . . . 195

B MultiFold unfolding **199**

Bibliography **213**

Acknowledgements **215**

Introduction

Since its inception, the Standard Model of Particle Physics has been remarkably successful at describing the elementary interactions governing the Universe in which we live. Nonetheless, it remains important to rigorously test every aspect of the theory with ever greater precision to improve our modelling and ensure that new physics is not hiding in plain sight.

The Large Hadron Collider (LHC) at CERN has been running since 2008 for just this purpose. From searching for new phenomena, such as the existence of magnetic monopoles, to deepening our understanding of flavour physics, to pushing the precision frontier of known Standard Model processes, and precisely characterising the properties of the Higgs boson, the vast physics scope of experiments such as ATLAS has been designed to leave no stone unturned.

Quantum Chromodynamics (QCD) is one facet of the Standard Model particularly deserving of our attention. At hadron colliders like the LHC, the interactions which take place are dominated by the strong force. This motivation alone is sufficient to justify the further study of QCD, and the case is only amplified when considering the many yet unexplored details of the theory and the inherent limit Nature places on our ability to make accurate predictions.

Our ability to make predictions for processes involving b- and c-quarks, or heavy flavour quarks, is then further complicated by one aspect: the non-negligible mass. Precisely thanks to their mass, however, heavy flavour quarks represent key gateways to access the Higgs boson and as portals for possible new physics. For this reason, the study of Standard Model processes involving these quarks is imperative in order to characterise, as much as possible, the main source of background.

Z+heavy flavour jets is one Standard Model process of particular interest. Besides representing a background for VH production, the process is abundant at the TeV-scale energies of the LHC, and the Z boson acts as a standard candle for precision physics, as the leptonic decay channels are easily identifiable within particle detectors. This process allows us to readily identify heavy flavour jets in the final state with relatively low background contamination, allowing for the analysis of their inner structure to better understand the final state evolution of a collision event.

In this thesis, we explore three aspects of heavy flavour jets: phenomenology, reconstruction, and a measurement of heavy flavour jet production in association with a Z boson. In all three aspects original, and often innovative, contributions resulted from the studies reported. The thesis is divided as follows.

In Chapter 1, we provide an introduction to QCD in the context of hadron colliders, introduce the concept of jets, and explain how a final state evolves from quarks to jets. Also introduced is the field of jet substructure, which addresses how the QCD evolution of a final state influences the inner structure of jets. Lastly, we recount the latest measurements of jet substructure and V+jets production.

Chapter 2 is dedicated to understanding the shortcomings of the current methods of assigning flavour to jets in the context of theory predictions, specifically problems related to inconsistencies

across theory and experiment which arise in predictions at higher orders in perturbation theory involving heavy flavour quarks and hadrons. We then recount a comparative study of novel flavour jet algorithms which solve the problem of assigning flavour to jets, and report the findings.

In Chapter 3 we introduce the ATLAS Detector, and report how the final state of a proton-proton collision event is reconstructed.

In Chapter 4, we turn our attention to the identification of heavy flavour jets in a given final state, a technique known as *flavour tagging*. We describe the characteristics of heavy flavour jets which allow for their identification within a particle detector, and proceed to describe the algorithms developed for this purpose for both small- and large-radius jets. We then present a calibration of the boosted tagging algorithm employed by the ATLAS Collaboration, known as the Xbb tagger.

Finally, in Chapter 5 we conclude by describing a novel measurement of $Z + bb$ and $Z + jj$ production, performed using 140.1 fb^{-1} of proton-proton collision data collected at $\sqrt{s} = 13$ TeV by the ATLAS Experiment. The measurement focuses in particular on jet substructure and the differences between b-jets and light jets.

The work described in this thesis was carried out as part of wider collaborations. As such, it is worthwhile to specify my personal contributions. These include all studies related to the application of the algorithms described in Chapter 2 to experiment; the entire calibration of the Xbb tagger described in Chapter 4, and all aspects of the measurement described in Chapter 5.

Chapter 1

V+jets and jet substructure

This chapter provides an overview to the physics behind W or Z (V) boson production in association with jets and to jet substructure, the main topics of this thesis. In Section 1.1, we begin by giving a brief overview to the Standard Model as it is currently known. We then move on to discussing aspects more directly related to Quantum Chromodynamics (QCD) and physics at colliders in Section 1.2. Section 1.3 is dedicated to jet physics, and an introduction to jet substructure. Specific jet substructure observables relevant to this thesis are introduced in Section 1.4. Lastly, in Section 1.5, we provide an overview of some recent Z+jets and jet substructure measurements.

1.1 The Standard Model

1.1.1 A brief overview of the universe

To make the chapter more readable, we will briefly describe the phenomenology of the Standard Model [1], [2], our current best description of Nature. It categorises all particles, into two categories: fermions and bosons.

The gauge bosons are the force-carrying particles, responsible for mediating the strong (in the case of the gluons) and electroweak interactions (in the case of the $W^{\pm}$ boson, Z boson, and photon γ). Notably, the W and Z bosons are *massive*, leading to the short range of the weak interaction, whereas the other two are massless.

There is also the Higgs boson, which arises from the spontaneous symmetry breaking of the electroweak force. Unlike the other bosons, it is not responsible for mediating a force, but for giving mass to the other particles of the Standard Model. More details will follow in Section 1.1.3.

Fermions compose the matter sector of the Standard Model. They are divided into three generations, each with the same quantum numbers, but differing in their masses. The fermions are further divided into quarks and leptons. The main differentiator between the two is that the former are the sole particles in the matter sector which can interact via the strong force. Free quarks cannot be found in Nature, and they are always bound to other quarks via gluons.

Figure 1.1 shows all fermions and bosons currently belonging to the Standard Model, along with properties such as spin, mass, charge, etc.

In the next section, we describe the Standard Model Lagrangian in more detail.

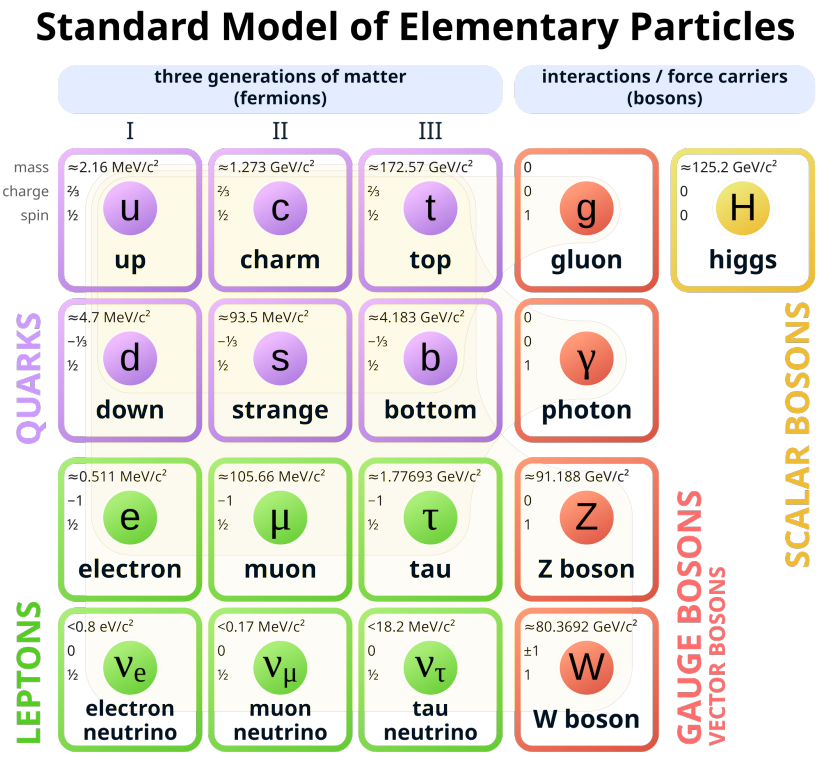


Figure 1.1: Our current knowledge of the particles described by the Standard Model of particle physics [3].

1.1.2 The Standard Model in a nutshell

The Standard Model Lagrangian

$$\begin{aligned}\mathcal{L}_{SM} = &-\frac{1}{4}F_{\mu\nu}F^{\mu\nu} \\ &+i\overline{\psi}D\psi + h.c. \\ &+\overline{\psi}_i y_i j \psi_j \phi + h.c. \\ &+|D_\mu \phi|^2 - V(\phi)\end{aligned} \tag{1.1}$$

is invariant under local gauge symmetries of the group $SU(3)_C \otimes SU(2)_L \otimes U(1)_Y$. The three gauge groups which compose the tensor product each represent one of the fundamental symmetries underlying the forces of Nature: colour, weak isospin, and hypercharge, each giving rise, respectively, to the strong, weak, and electromagnetic interactions described by the Standard Model Lagrangian.

Quantum Electrodynamics

A local $U(1)_{em}$ gauge symmetry leads to Quantum Electrodynamics (QED). This force describes interactions between charged particles and photons. Starting from a Dirac Lagrangian describing a free fermion

$$\mathcal{L}_D = \overline{\psi}(i\partial\!\!\!/ - m)\psi, \tag{1.2}$$

where ψ is the fermion field, $\bar{\psi}$ its adjoint, m the particle mass and $\partial\!\!\!/ = \partial_\mu \gamma^\mu$ the del operator, it is easy to see that $\mathcal{L}_D$ is invariant under global $U(1)_{em}$ transformations of the form $\psi' = \exp(i\alpha)\psi$ when α is a constant. In fact, the adjoint operation on the ψ' field leads to a factor which cancels with the term $\exp(ie\alpha)$.

In accordance with the gauge protocol, let us now promote the global symmetry of the Lagrangian to a local one

$$\psi' = \exp(ie\alpha(x))\psi, \tag{1.3}$$

where α is now a function which depends on the spacetime coordinates x and e is a constant representing the elementary charge. We can therefore define the covariant derivative

$$D_\mu \doteq \partial_\mu + ieA_\mu, \tag{1.4}$$

where A_μ is the gauge (photon) field, and require that the term $D_\mu\psi$ transform as ψ itself would

$$(D_\mu\psi) \to (D_\mu\psi)' = \{\partial_\mu + ieA'_\mu\}\psi' = [\exp(ie\alpha(x))]D_\mu\psi. \tag{1.5}$$

For the Lagrangian to remain invariant, A'_μ must equal the following expression:

$$A'_\mu = A_\mu - \partial_\mu \alpha(x). \tag{1.6}$$

We can then build a term to describe the free propagation of A_μ by computing the commutator of the covariant derivative

$$[D_\mu, D_\nu] = ie\{\partial_\mu A - \partial_\nu A) \equiv ieF_{\mu\nu} \tag{1.7}$$

where $F_{\mu\nu}$ is a tensor of the field A_μ defined from the commutator, i.e.

$$F_{\mu\nu} = -\frac{i}{e}[D_\mu, D_\nu]. \tag{1.8}$$

In this way, we have all the ingredients needed to construct the QED Lagrangian

$$\mathcal{L}_{QED} = -\frac{1}{4}F_{\mu\nu}F^{\mu\nu} + \overline{\psi}(i\not{D} - m)\psi. \tag{1.9}$$

The first term in Eq. 1.9 represents the kinetic term which describes the free propagation of the photon field and the second term is the Dirac Lagrangian from which we started, with the del operator replaced by the covariant derivative to ensure invariance under the local gauge symmetry from Eq. 1.3.

We would like to note that the method described is not the only way in which the QED Lagrangian (or that of any other fundamental force) can be derived. The principal advantage of the gauge protocol is that is does not assume any prior knowledge of the field A_μ, and thus allows for the derivation of the gauge field in a general way. We know from classical electrodynamics, however, that A_μ represents the electromagnetic potential. One could thus use this prior knowledge to derive 1.9 by means of the minimal coupling. This procedure accounts for the electromagnetic potential in the four-momentum of the fermion, leading to a modification of Eq. 1.2 as follows

$$\mathcal{L} = \overline{\psi}(i\not{\partial} - e\not{A} - m)\psi. \tag{1.10}$$

One can then develop the additional terms of the Lagrangian, impose the local $U(1)_{em}$ gauge symmetry, and derive the necessary transformation of the field A_μ needed to respect the symmetry.

Quantum Chromodynamics

In a similar fashion, it is possible to derive the Lagrangian for strong interactions. In this case, we consider the spinor q_f

$$q_f = \begin{bmatrix} q_f^1 \\ q_f^2 \\ q_f^3 \end{bmatrix}, \tag{1.11}$$

where the subscript f indicates quark flavour and the superscript is the colour index. This index can take three values, each corresponding to one of the three possible (anti-)colours which a quark may have: red, green, and blue. The free Lagrangian for these fields

$$\mathcal{L}_D = \sum_f \overline{q}_f \left(i\not{\partial} - m_f\right) q_f \tag{1.12}$$

is invariant under the transformation

$$q_f \to q_f' = \exp\left[i\theta_a \frac{\lambda^a}{2}\right] q_f, \tag{1.13}$$

where λ_a is the generator of the $SU(3)$ symmetry group and θ_a is the parameter relative to the generator λ_a.

If we define the covariant derivative as before, promote the symmetry to a local one, and follow through with the gauge protocol, we find that

$$\mathcal{L}_{QCD} = -\frac{1}{4}G^a_{\mu\nu}G^{\mu\nu}_a + \sum_f \overline{q}_f \left(i\not{D} - m_f\right) q_f, \tag{1.14}$$

in analogy with QED. In the fundamental representation, $SU(3)$ has eight generators, known as the Gell-Mann matrices, corresponding to the eight colour combinations ascribed to the gauge

(gluon) field G_μ. Gluons thus carry two colours, and a gluon exchange corresponds to an interaction which alters the colour of the quarks involved.

There is one important difference with respect to QED, however: the $SU(3)$ symmetry group, unlike the $U(1)$ group, is non-abelian. A direct consequence of this fact is the self-coupling of the gluon, namely the 3- and 4-point gluon self-interaction, as can be seen from

$$G^a_\mu \to (G^a_\mu)' = G^a_\mu - \frac{1}{g_s}\partial_\mu \theta^a(x) - f^{abc}\partial_\mu \theta_b(x) G_{\mu c}, \tag{1.15}$$

where g_s is the coupling constant for the strong interaction and f^{abc} the structure constant of $SU(3)$, defined as

$$\left[\frac{\lambda_a}{2}, \frac{\lambda_b}{2}\right] = i f_{abc}\frac{\lambda_c}{2}. \tag{1.16}$$

Electroweak Unification

The gauge protocol can also be applied to derive a Lagrangian describing weak interactions. The transformation acts on a weak-isospin doublet

$$\psi_L = \begin{pmatrix} \nu_\ell \\ \ell^- \end{pmatrix} \tag{1.17}$$

composed of the left-handed states of any two fermions belonging to the same generation (the right-handed state ℓ^+_R lives as a singlet). For the sake of simplicity, we have considered a charged lepton and neutrino.

We can apply the local gauge transformations

$$\psi_L \to \psi'_L = \exp\Big[i\frac{\tau_j}{2}\alpha^j(x)\Big]\psi \tag{1.18}$$

$$\ell^+_R \to \ell^{+\prime}_R = \ell^+_R, \tag{1.19}$$

where the matrices τ_j are the generators of the $SU(2)$ gauge group and $\alpha^j(x)$ represents the function which introduces local spacetime dependence to the transformation.

After following the gauge protocol, we can derive the gauge fields W^1_μ, W^2_μ, and W^3_μ. Linear combinations of W^1_μ and W^2_μ are able to correctly describe charged current interactions attributed to the $W^\pm$ bosons

$$W^\pm_\mu = \frac{1}{\sqrt{2}}(W^1_\mu \mp i W^2_\mu), \tag{1.20}$$

although W^3_μ is *not* able to describe the neutral current, as this field would only couple to left-handed fermions, and it is known from experiment that the neutral current couples to both left-handed and right-handed fermions.

To reconcile this fact, electromagnetic interactions, the only other interaction mediated by a neutral boson, can be included together with a description of the weak force. The Z boson and γ can then be derived from a linear combination of the neutral fields. The gauge group to consider is thus $SU(2)_L \otimes U(1)_Y$.

Let us introduce a quantity known as *hypercharge*

$$Y = 2(Q - I_3), \tag{1.21}$$

where Q represents electric charge and I_3 the weak isospin. The hypercharge is defined in such a way that it distinguishes components of the weak isospin doublet in Eq. 1.17 from right-handed states. For completeness, note that

$$Y(\nu_\ell) = -1 \tag{1.22}$$
$$Y(\ell^-) = -1 \tag{1.23}$$
$$Y(\ell^+) = -2. \tag{1.24}$$

The full transformation under the new gauge group becomes

$$\psi_L \to \psi'_L = \exp\left[iy_1\beta(x)\right] \exp\left[i\frac{\tau_j}{2}\alpha^j(x)\right] \psi_L \tag{1.25}$$
$$\ell_R \to \ell'_R = \exp\left[iy_2\beta(x)\right] \ell_R, \tag{1.26}$$

where y_1 represents the value of hypercharge for the left-handed states and y_2 for the right-handed states. The covariant derivatives become

$$D_\mu\psi_L(x) = \left[\partial_\mu + ig\frac{\tau_j}{2}W^j_\mu(x) + ig'\frac{y_1}{2}B_\mu(x)\right]\psi_L(x) \tag{1.27}$$
$$D_\mu\ell_R = \left[\partial_\mu ig'\frac{y_2}{2}B_\mu(x)\right]\ell_R(x), \tag{1.28}$$

where g and g' are the two coupling constants for the left- and right-handed states.

From Eq. 1.27, it is clear that we now have four gauge fields, the three W^j_μ fields from before and B_μ. The fields W^1_μ and W^2_μ combine as in Eq. 1.20 to form the $W^\pm$ bosons; the Z boson and photon must be obtained through an appropriate linear combination in the neutral sector.

We can consider a generic rotation by an angle θ_W, known as the weak mixing angle

$$\begin{bmatrix} \cos\theta_W & \sin\theta_W \\ -\sin\theta_W & \cos\theta_W \end{bmatrix} \begin{bmatrix} B_\mu \\ W^3_\mu \end{bmatrix} = \begin{bmatrix} A_\mu \\ Z_\mu \end{bmatrix}. \tag{1.29}$$

The interaction Lagrangian for the neutral sector in terms of Z^μ and A^μ can then be written as

$$\begin{aligned} \mathcal{L}_{NC} = &\overline{\psi}\gamma_\mu\left\{g\sin\theta_W\frac{\tau_3}{2} + g'\cos\theta_W\frac{Y(\psi)}{2}\right\}\psi A^\mu \\ &+ \overline{\psi}\gamma_\mu\left\{g\cos\theta_W\frac{\tau_3}{2} - g'\sin\theta_W\frac{Y(\psi)}{2}\right\}\psi Z^\mu. \end{aligned} \tag{1.30}$$

If we continue to assume as in Eq. 1.17 to be dealing with a charged lepton and neutrino, we can compare the first term in Eq. 1.30 to the interaction term in Eq. 1.9 to find that

$$-e = g\sin\theta_W\frac{\tau_3}{2} + g'\cos\theta_W\frac{Y(\psi_e)}{2}, \tag{1.31}$$

or, after specifying τ_3 and $Y(\psi_e)$

$$g\sin\theta_W = g'\cos\theta_W = e. \tag{1.32}$$

We have successfully found an expression relating the couplings g and g' to the electric charge e through the weak mixing angle. The weak mixing angle is a free parameter of the Standard Model, and can only be determined experimentally. At the mass of the Z boson, $\sin^2\theta_W$ has been measured to be [4]

$$\sin^2\theta_W(M_Z) = 0.23129(4) \pm 1.7 \times 10^{-4}. \tag{1.33}$$

From Eq. 1.30, we can also see that the coupling of the Z boson corresponds to

$$g_Z = \frac{e}{\sin\theta_w \cos\theta_W}\{\frac{1}{2}\tau_3 - Q\sin^2\theta_W\}. \tag{1.34}$$

We must specify the terms in the matrix to derive the numerical value of the coupling. For ν_L it is straightforward to show that the coupling $g_{Z\nu} = \frac{e}{2\sin\theta_W\cos\theta_W}$. For the lepton ℓ^-, we must decompose the spinor ψ into the left-handed and right-handed components by means of the projection operator $P_{L/R} = \frac{1}{2}(1 \mp \gamma^5)$, and, accordingly, consider the projections of the matrix in 1.34. When doing so, we find that

$$g_{Z\ell_L} = \frac{e}{2\sin\theta_W\cos\theta_W}[I_3(\ell_L)(1 + 4Q_\ell \sin^2\theta_W)] \tag{1.35}$$

$$g_{Z\ell_R} = \frac{e}{2\sin\theta_W\cos\theta_W}[I_3(\ell_R)]. \tag{1.36}$$

where I_3 as before indicates the weak isospin and Q_ℓ is the charge of the lepton in question.

We correctly find that the coupling constant differs for left-handed and right-handed spinors. We can also note that, if we were to rewrite the second term of Eq. 1.30, the use of the projection operators would lead to the V-A structure we expect of the neutral current. We can therefore fully reproduce the phenomenology of the weak interaction.

1.1.3 Spontaneous symmetry breaking

The careful reader would have surely noticed that none of the Lagrangians considered in the previous section contain massive gauge fields. However, in nature we know that the $W^\pm$ and Z bosons are massive. How do we reconcile this fact?

The solution is not as simple as adding in a mass term by hand à la Dirac[1]. This is due to the fact that a mass term of this form, i.e.

$$\mathcal{L} = -\frac{1}{4}F^{\mu\nu}F_{\mu\nu} - m^2 A^\mu A_\mu, \tag{1.37}$$

violates gauge symmetry. A more elegant solution is thus required.

Higgs et al. found this solution for us, which earned Higgs the Nobel Prize in 2012. The mechanism, aptly known as the Brout-Englert-Higgs-Hagen-Guralnik-Kibble (BEH) Mechanism, introduces a scalar field (the Higgs field) to the Standard Model and a potential (the Higgs potential) to induce spontaneous breaking of the $SU(2) \otimes U(1)$ gauge symmetry (SSB), leading to the mass of the gauge bosons.

The BEH Mechanism is based on a Lagrangian of the form

$$\mathcal{L} = (\partial_\mu\phi)^\dagger(\partial^\mu\phi) - V(\phi, \phi^\dagger) \tag{1.38}$$

with the potential V defined as

$$V(\phi, \phi^\dagger) = \frac{\mu^2}{2}\phi^\dagger\phi - \frac{\lambda}{4}(\phi^\dagger\phi)^2. \tag{1.39}$$

Here, ϕ is an $SU(2)$ doublet composed of two complex scalar fields

$$\phi(x) = \begin{bmatrix}\phi^+ \\ \phi^0\end{bmatrix} = \begin{bmatrix}\frac{1}{\sqrt{2}}(\phi_3 + i\phi_4) \\ \frac{1}{\sqrt{2}}(\phi_1 + i\phi_2)\end{bmatrix}, \tag{1.40}$$

[1] This is known in literature as the Proca action.

m is the mass of the field ϕ, and λ is a free parameter. Note that the Lagrangian in Eq. 1.38 is invariant under the $SU(2) \otimes U(1)$ gauge symmetry.

We can minimise the potential to find the ground state of ϕ

$$\frac{\partial V}{\partial |\phi|} = \mu^2 |\phi| + \lambda |\phi|^3. \tag{1.41}$$

If $\lambda, \mu^2 > 0$, the only minimum is for $\phi = 0$. Instead, if $\mu^2 < 0$, there is a maximum at $\phi = 0$ and an infinite set of degenerate minimum states found at $\langle 0| \phi^\dagger \phi |0\rangle = -\sqrt{m^2/\lambda} \equiv v$. This quantity is known as the *vacuum expectation value*.

We can choose *ad libitum* any one of these states to define our ground state, say

$$\phi_0 = \frac{1}{\sqrt{2}} \begin{bmatrix} 0 \\ v \end{bmatrix}. \tag{1.42}$$

It is easy to show that, although the Lagrangian in Eq. 1.38 remains invariant under the $SU(2) \otimes U(1)$, *the ground state does not respect the same symmetry!* Specifically, it is the action of the generators τ^i and y from Eq. 1.25 which lead to a variation of the field.

We would now like to find a way to give mass to the $W^\pm$ and Z bosons, while leaving the photon massless. The Goldstone theorem teaches us that for each broken generator, there exists a corresponding massless particle. Naïvely, we have just stated that there are in fact *four* broken generators. This problem is quickly resolved through a linear combination

$$Q = \frac{1}{2}\tau_3 - \frac{1}{2} y, \tag{1.43}$$

$$Q' = \frac{1}{2}\tau_3 + \frac{1}{2} y. \tag{1.44}$$

Now, Q is unbroken while τ_1, τ_2 and Q' remain broken. The corresponding massless particles are not physical, but are said to be "eaten" by the gauge bosons, giving rise to their mass.

Consider a small perturbation of the vacuum state ϕ_0

$$\phi(x) = \frac{1}{\sqrt{2}} \exp\left[\frac{i\tau_j \xi_j(x)}{2}\right] \begin{bmatrix} 0 \\ v + H(x) \end{bmatrix}, \tag{1.45}$$

where $\xi_j(x)$ is a parameter dependent on the spacetime coordinates and $H(x)$ is a small perturbation to the field ϕ_0 from Eq. 1.42 corresponding to the Higgs boson. The exponential term can be gauged away by means of the transformation

$$\phi(x) \to \phi'(x) = \exp\left[-\frac{iT_j \xi_j(x)}{2}\right] \phi. \tag{1.46}$$

Now, we can follow through with the gauge protocol and consider the term

$$\begin{aligned} (D_\mu \phi)^\dagger (D_\mu \phi) = \frac{g^2}{4} \left(v^2 + 2vH + H^2\right) W_\mu^+ W^{-\mu} \\ + \frac{1}{8}(v+H)^2 \left(g^2 W_\mu^3 W^{3\mu} - 2gg' W_\mu^3 B^\mu + g'^2 B_\mu B^\mu\right). \end{aligned} \tag{1.47}$$

Here, it is easy to read off all the information that we've sought from the start: the mass of the $W^\pm$ is found to be $m_W^2 = \frac{1}{4} g^2 v^2$, that of the Z boson (after electroweak unification) is $M_Z = \frac{1}{4} v^2 (g^2 + g'^2)$, and the photon field A_μ remains massless. Moreover, we find no interaction term between the H and A_μ, and we find that the coupling of the Higgs boson and gauge bosons is, as expected, proportional to the square of their masses. The problem of massive gauge bosons in the Standard Model is thus solved.

1.1.4 Fermion Masses

Fermion mass terms also fail to respect the gauge symmetry of the Standard Model Lagrangian. Specifically, since the $SU(2)_L$ symmetry acts differently on left-handed and right-handed spinors, terms of the form $m\bar{\psi}\psi$ are clearly not invariant after decomposition into their respective components.

Once again, the Higgs field can solve this problem. Let us consider transformations of the type found in Eq. 1.18. When applied to the Higgs field ϕ, we find that, to first order,

$$\phi \to \phi' = \exp\left[ig\frac{\tau_j}{2}\alpha^j(x)\right]\phi \approx [1 + ig\frac{\tau_j}{2}\alpha^j(x)]\phi. \tag{1.48}$$

Likewise, the same transformation on the field $\bar{\psi}_L$ leads to

$$\bar{\psi}_L \to \overline{(\psi_L)'} = \bar{\psi}_L \exp\left[ig\frac{\tau_j}{2}\alpha^j(x)\right] \approx \bar{\psi}_L[1 - ig\frac{\tau_j}{2}\alpha^j(x)]. \tag{1.49}$$

The quantity $\bar{\psi}_L\phi$ is therefore invariant under this transformation. The same holds true if we consider $\bar{\psi}_L\phi\psi_R$, as ψ_R does not transform under $SU(2)_L$.

We can thus write the Yukawa Lagrangian

$$\mathcal{L}_Y = -k(\bar{\psi}_L\phi\psi_R + \bar{\psi}_R\bar{\phi}\psi_L), \tag{1.50}$$

where k is a coupling constant. Equation 1.50 respects the full gauge symmetry of the Standard Model Lagrangian, and gives mass to particles and antiparticles. We can specify the terms to some species of fermion f to find that $m_f = k_f v/\sqrt{2}$, where k_f is the Yukawa coupling for f. Note that the coupling is arbitrary and no fermion mass is explicitly predicted by the Standard Model.

1.2 Physics at colliders

The LHC is a hadron collider. This means that the dominant interaction is QCD. It therefore is worth taking the time to look more deeply into this fundamental force. This section is dedicated to the ingredients needed to calculate a cross-section for a given process at the LHC, starting from the so-called "master equation" containing everything needed for this calculation, and moving on the individual components, such as parton distribution functions, perturbative QCD, parton showers, and hadronisation.

1.2.1 Anatomy of a cross section

Hadrons are not elementary particles. When two hadrons collide at high energies, it is not the hadron as a whole which collides, but the individual *partons* which make them up. The short-range physics which dominate the partonic interaction occur on a much shorter time-scale than the long-range interactions which proceed and precede it. As it occurs at short range, the partonic interaction is said to be *perturbative*, meaning that it occurs at a short-enough distance (or, equivalently, high enough energy scale) to allow for a series expansion in terms of α_s[2] and can therefore be calculated from first principles.

The differential cross section for a generic observable X from a collision at the LHC is said to be *factorised*, as it can be directly calculated from the underlying partonic process and from functions describing the non-perturbative physics involved in the interaction which must be derived from data.

[2]For additional details, cfr. Section 1.2.2.

The hadronic cross section is given by the following formula:

$$\frac{d\sigma}{dX} = \sum_{j,k} \int dx_1 dx_2 f_j(x_1, Q^2) f_k(x_2, Q^2) \frac{d\hat{\sigma}_{jk}(x_1 P_1, x_2 P_2, Q^2, \mu_R^2, \mu_F^2)}{d\hat{X}} F(\hat{X} \to X; Q^2, \mu_F^2). \tag{1.51}$$

Here, we are stating that the differential cross section for a hadronic observable $d\sigma/dX$ is given by the cross section of the partonic observable $d\hat{\sigma}/d\hat{X}$. This partonic cross section depends on Q^2, the centre-of-mass energy of the system. There is also dependence on the fraction x_i of momentum of the hadron P_i carried by the colliding partons. The integration over x_1 and x_2 guarantees that all possible momentum fractions are considered. Finally, the renormalisation scale μ_R and the factorisation scale μ_F respectively establish the thresholds at which ultraviolet divergences are subtracted [5], and at which the short-distance physics which enter the partonic cross section calculation are separated from the long-distance interactions.

There is also a dependence on the exact species of partons which collide, j and k. Protons, as hadrons, are dynamic systems composed of quarks and gluons. At valence level, they are composed of three quarks, u, u, and d, but this picture quickly fades as interactions between the quarks lead to a sea of gluons and gluon splittings lead to the production of other quark flavours. As there are multiple partonic interactions which may lead to the same final state, it is necessary to sum over all of these possible interactions.

The exact nature of the partons to be considered is determined by Parton Distribution Functions (PDFs), or $f(x, Q^2)$ in Eq. 1.51. These establish, at a given energy scale Q^2, which are the parton species found within a given hadron carrying momentum fraction x of the total hadron momentum.

Once the partons have been extracted and the hard scattering has occurred, the particles produced by the partonic interaction evolve towards lower-energy states. The first, perturbative part of the evolution is characterised by successive collinear branchings of quarks and gluons, as described by the Dokshitzer-Gribov-Lipatov-Altarelli-Parisi (DGLAP) equations [6], [7].

Lastly, the function $F(\hat{X} \to X; Q^2, \mu_F^2)$ describes the long-distance, non-perturbative evolution from a certain partonic final state $\hat{X}$ to a measurable hadronic state X. This function can either be a Fragmentation Function (FF), describing the production of specific hadron species, or a jet function which instead describes the production of groupings of particles known as *jets*. More on jets will follow in Section 1.3. Besides the specific partonic and hadronic states, F depends on the energy scale of the process and the factorisation scale.

The next sections describe the aforementioned components of Eq. 1.51 in more detail.

1.2.2 Running of the coupling constant and perturbative QCD

Our excursion begins with what exactly constitutes perturbative QCD. The underlying determinant is the running of the coupling constant, α_s [2].

The strength of an interaction governed by a given force is determined by the coupling constant. Despite the name, the coupling is anything but constant. Instead, it varies as a function of energy. This variation can be understood as the (anti-)"screening" of a charge by the self-energy of the photon, gluon, or any other gauge boson.

For QED, an intuitive explanation is the following: as one approaches a charge, there are fewer virtual electron-positron pairs which may screen this charge. In a scattering experiment, one can approach the charge by reaching higher and higher energies. Thus, as the photon energy increases, the value of the coupling increases.

More rigorously, the photon self-energy diverges and one needs *renormalisation* to regularise this divergence. Sparing the details, the regularisation procedure relies on the fact that, experimentally, the electric charge (and thus the coupling) is measured to be finite. If we know the electric charge at some scale $q^2 = \mu^2$, where q^2 is the energy of the virtual photon, the one-loop photon self-energy corrections $\Pi(q^2)$ and $\Pi(\mu^2)$ are both separately divergent. However, their difference

$$\Pi(q^2) - \Pi(\mu^2) \approx \frac{1}{12\pi^2} \ln\left(\frac{q^2}{\mu^2}\right) \tag{1.52}$$

is finite. This can be inserted into the expression for the coupling to give

$$\alpha(q^2) = \frac{\alpha(\mu^2)}{1 - \frac{\alpha(\mu^2)}{3\pi} \ln\left(\frac{q^2}{\mu^2}\right)}. \tag{1.53}$$

In QCD, things are complicated by the gluon self-interaction. This introduces additional contributions to the gluon self-energy which completely alters the behaviour.

The analogous expression for α_s turns out to be

$$\alpha_s = \frac{\alpha(\mu^2)}{1 + B\frac{\alpha_s(\mu^2)}{3\pi} \ln\left(\frac{q^2}{\mu^2}\right)} \tag{1.54}$$

where B is a quantity which depends on the number of fermionic and bosonic loops which are relevant to the calculation. The general expression depends on the number of quark flavours N_f and the number of colours N_c

$$B = \frac{11N_c - 2N_f}{12\pi}. \tag{1.55}$$

The behaviour of the coupling constant fundamentally depends on the number of flavours and colours. In our Universe, since $N_c = 3$ and $N_f \leq 6$ (depending on which quark flavours are accessible at a given energy scale), $B > 0$. This leads to an *anti-screening* of the colour charge at high energies and implies that QCD is only perturbative at high energies, a property known as *asymptotic freedom*. The scale at which this transition occurs is determined by the Landau pole $\Lambda \sim 300$ MeV.

Figure 1.2 shows the running of the coupling constants for the electromagnetic, weak, and strong interactions.

1.2.3 QCD radiation in the initial and final state

In this section we wish to introduce the QCD processes which influence the initial and final states of an event. Before moving on to Sections 1.2.4 and 1.2.5, we find it helpful to introduce the splitting functions and the DGLAP equations.

If we consider the radiation of a gluon subsequent to the production of a $q\bar{q}$ pair, we find that [4]

$$\frac{d\sigma_{q\bar{q}g}}{d\cos\theta dz} \approx \sigma_{q\bar{q}} C_F \frac{\alpha_s}{2\pi} \frac{2}{\sin^2\theta} \frac{1 + (1-z)^2}{z}, \tag{1.56}$$

which is proportional to the cross-section for the $q\bar{q}$ production process, the colour-factor $C_F = (N_C^2 - 1)/2N_c$, the strong coupling $\alpha_s = g_s^2/4\pi$, and the energy fraction of the gluon z. From here, it is already clear that the expression diverges in two circumstances: soft and collinear

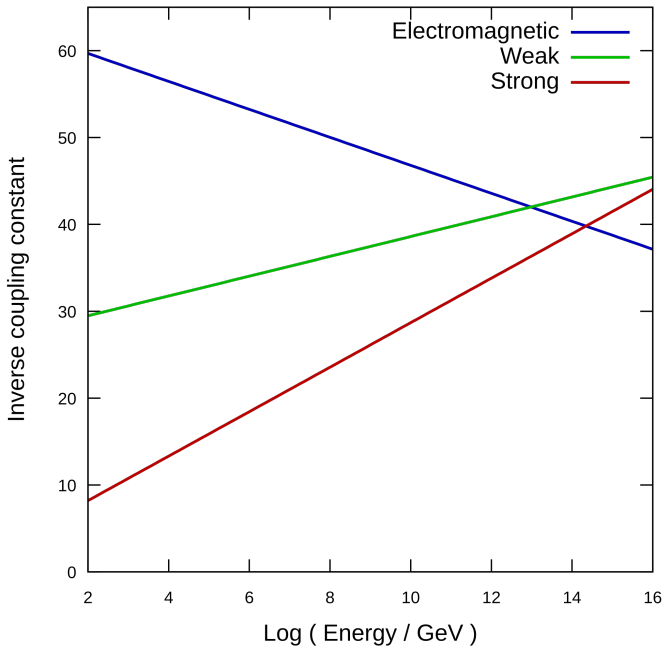


Figure 1.2: The inverse values of the coupling constants for the various forces described the Standard Model at different energy scales [8].

gluon emissions. A final state quark is thus very likely to emit a gluon which is either soft, collinear, or both.

The term $C_F \frac{1+(1-z)^2}{z}$ in Eq. 1.56 is known as the *splitting function* P_{qq}, which, in this case, can be interpreted as the probability of finding a quark emitting a collinear gluon with momentum fraction z. Analogous splitting functions exist for other types of emissions [9],

$$P_{qq}(z) = C_F \frac{1+z^2}{1-z} \tag{1.57}$$

$$P_{gg}(z) = C_A \frac{z^4+1+(1-z^4)^2}{z(1-z)} \tag{1.58}$$

$$P_{pg} = T_R(z^2+(1-z)^2) \tag{1.59}$$

where $C_A = N_C$ is the gluon colour factor and T_R is a colour factor fixed to $1/2$ by convention. These splitting functions are all reported at leading order.

The probability of a transition $a \to bc$ occurring can be written in terms of the individual splitting functions,

$$d\mathcal{P} = \frac{\alpha_s}{2\pi} \frac{dQ^2}{Q^2} P_{a\to bc}(z) dz, \tag{1.60}$$

known as the DGLAP equations.

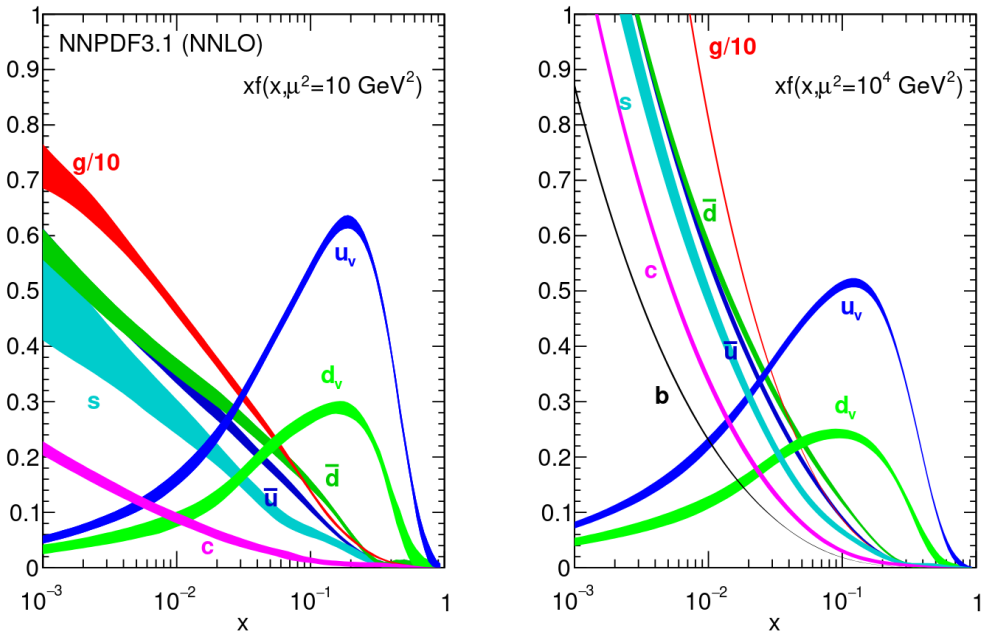


Figure 1.3: The NNLO proton PDFs in the NNPDF3.1 fit for quarks and gluons at 10 GeV (left) and 10^4 GeV (right) [10].

1.2.4 Parton Distribution Functions

When protons collide at high energies, it is not the protons themselves which interact, but their constituents, known as partons. The proton is a dynamic system; its composition is described by *parton distribution functions*, $f_j(x, Q^2)$, which give the probability of extracting a parton of species j carrying fraction x of the total hadron momentum at energy scale Q^2.

PDFs cannot be calculated. They must be extracted from data. This is done by reverse engineering the measured cross section, given by Eq. 1.51 to deconvolve the PDF contribution.

Several collaborations actively fit PDFs using old and new data, taken from both the LHC as well as Tevatron, HERA, and other colliders. Some of the best data for PDF fits actually comes from Deep Inelastic Scattering experiments, where electrons, muons or neutrinos are used to probe the proton. The advantage in this case is the simpler form of the master equation, which contains only one PDF.

Figure 1.3 shows the PDFs obtained from two such fits carried out by the NNPDF Collaboration, evaluated at $Q^2 = 10$ GeV^2 and $Q^2 = 10^4$ GeV^2 [10].

At distances of the order of a hadron radius, i.e. ~ 1 fm, QCD interactions of the form $a \to bc$, where a, b, and c are partons, are plentiful due to the running of the coupling constant. Considering this dynamicity, it is clear that the system is in constant evolution and that the PDFs for the partons a, b, and c are affected by these interactions. Fortunately, we can describe this evolution with the DGLAP equations [7]:

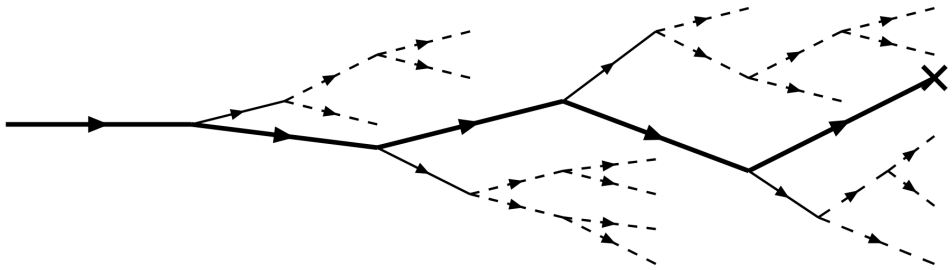


Figure 1.4: A cascade of successive branchings in the final state generated by a $\gamma^*/Z \to q\bar{q}$ process [9].

$$\frac{dq_f(x,Q^2)}{d\log Q^2} = \frac{\alpha_s}{2\pi}\int_x^1 \frac{dy}{y}\left[q_f(y,Q^2)P_{qq}\left(\frac{x}{y}\right) + g(y,Q^2)P_{qg}\left(\frac{x}{y}\right)\right];$$
$$\frac{dg(x,Q^2)}{d\log Q^2} = \frac{\alpha_s}{2\pi}\int_x^1 \frac{dy}{y}\left[\sum_f q_f(y,Q^2)P_{gq}\left(\frac{x}{y}\right) + g(y,Q^2)P_{gg}\left(\frac{x}{y}\right)\right]. \quad (1.61)$$

Here, both the PDF q_f for a quark of flavour f and the gluon PDF g are considered. In practice, the DGLAP equations in the form above describe the PDF evolution in Q^2. This allows us to measure PDFs at a given energy scale and evolve them to different values of Q^2, eliminating the need to scan in variables other than x.

1.2.5 Final state evolution

Partons in the final state are typically produced off-shell by the hard scattering process and need a way to shed their virtuality. This occurs through subsequent collinear emissions. Equation 1.60 describes a generic transition $a \to bc$ caused by QCD interactions. It can be applied in sequence to describe further splittings, as shown in Figure 1.4.

The expression for the emission of the $(n+1)^{\text{th}}$ gluon, considering all pairs i, j of external lines in the Feynman diagram, can be written as

$$d\sigma_{n+1} = d\sigma_n \frac{d\omega}{\omega}\frac{d\Omega}{2\pi}\frac{\alpha_s}{2\pi}\sum_{ij} C_{ij}W_{ij}, \quad (1.62)$$

where ω is the energy of the emitted gluon and $d\Omega$ is the element of solid angle for the emitted gluon [11]. The sum over the emitted pairs includes a radiation function W_{ij} weighted by the colour factor C_{ij}. The expression for the radiation function is as follows:

$$W_{ij} = \frac{\omega^2 p_i \cdot p_j}{(p_i \cdot q)(p_j \cdot q)}, \quad (1.63)$$

where q refers to the gluon momentum and p_i and p_j refer to the momenta of the quark before and after the emission, respectively.

In the massless case, Eq. 1.63 becomes

$$W_{ij} = \frac{1-\cos\theta_{ij}}{(1-\cos\theta_{iq})(1-\cos\theta_{jq})}. \quad (1.64)$$

We can then rewrite $W_{ij} = W_{ij}^{[i]} + W_{ij}^{[j]}$ with

$$W_{ij}^{[i]} = \frac{1}{2}\left(W_{ij} + \frac{1}{1-\cos\theta_{iq}} - \frac{1}{1-\cos\theta_{jq}}\right). \tag{1.65}$$

Integrating each expression $W_{ij}^{[i]}$ and $W_{ij}^{[j]}$ in the azimuthal angle ϕ, we find that

$$\int_0^{2\pi} \frac{d\phi_{iq}}{2\pi} W_{ij}^{[i]} = \begin{cases} \frac{1}{1-\cos\theta_{iq}}, & \text{if } \theta_{iq} < \theta_{ij} \\ 0, & \text{otherwise.} \end{cases} \tag{1.66}$$

This property is known as *angular ordering*, and implies that, on average, the gluon radiation is confined to a cone centred on i and extending up to j. Each subsequent emission is then confined to a smaller and smaller cone, leading to an overall conical structure of the entire sequence of successive branchings. Each subsequent emission is also softer than the last, as the emitting quark sheds virtuality.

1.2.6 Monte Carlo event generators

As discussed, there are many aspects to consider when using Eq. 1.51 to calculate a cross section at hadron colliders. Some, like the partonic cross section, are fully calculable. In other cases, such as for the collinear final state evolution, the underlying physics are known and single (or few) emissions can also be calculated explicitly. It is however difficult to carry out a fully calculation for the generic $2 \to n$ process by hand. Final steps such as hadronisation are not fully understood, and thus rely exclusively on models.

As even the known parts of a given cross section calculation can be complex, Monte Carlo (MC) event generators are used to calculate and simulate the entire event. The use of generators begins with Matrix Element (ME) calculations which use MC methods to solve the integrals required to calculate the hard scattering of partons. These calculations occur at fixed-order in α_s. Parton showers (PS) are then used to simulate the soft and collinear emissions in the final state. In this approach, the probability of a splitting, for example $q \to qg$, is subdivided into time intervals. In each interval, the probability of a splitting occurring is modelled as $\mathcal{P}(\text{splitting}) = 1 - \mathcal{P}(\text{no splitting})$. This leads to an exponential factor in the DLGAP equations

$$d\mathcal{P}_{a\to bc} = \frac{\alpha_s}{2\pi}\frac{dQ^2}{Q^2} P_{a\to bc} dz \ \exp\left(-\sum_{b,c}\int_{Q^2}^{Q^2_{\max}} \frac{dQ'^2}{dQ'^2}\int \frac{\alpha_s}{2\pi} P_{a\to bc}(z')\, dz'\right), \tag{1.67}$$

known as the *Sudakov factor* [9]. Equation 1.67 is none other than the DGLAP equation from Eq. 1.60 multiplied by the probability of no splitting $a \to bc$ occurring across the entire downward energy scale evolution being modelled, from $Q^2_{\max} \to Q^2$.

The variable of evolution used in the DGLAP equations is not unique. The equations can be equivalently formulated in terms of angle of emission, known as angular-ordered showers [12], or transverse momentum, known as dipole showers [13]. These are the most common choices used by parton shower algorithms.

Once the energy scale Q^2 is low enough, of the order ~ 1 GeV, the process terminates and a hadronisation model is implemented.

Matching and merging The parton shower approximates successive soft and collinear emissions of quarks and gluons. This is equivalent to calculating MEs at higher and higher orders. The ME element calculation is generally preferred, as it is more accurate thanks to the expansion

in terms of α_s which allows one to quantify the degree of precision. The PS, on the other hand, is a simulation and comes with no guaranteed degree of precision, despite the universality of the DGLAP equations and statistical reliability of the Sudakov factor which provide a level of physical accuracy.

When calculating MEs at higher (fixed) orders, a *matching* procedure with the PS is required to avoid double-counting of orders of emissions already accounted for. Likewise, *merging* schemes combine different MEs generated with multiple partons in the final states with the PS and each other. Examples of matching procedures include the MLM [14] and FxFx [15] approaches, while a popular merging procedure is the CKKW-L [16] approach. The exact details of these procedures are beyond the scope of this thesis.

1.2.7 Hadronisation

The parton shower describes the perturbative QCD evolution of the final state, down to about ~ 1 GeV. At lower energies, due to the running of the coupling constant α_s, the physics becomes non-perturbative and hadronisation occurs.

Hadronisation is the process of hadron formation from the quarks in the final state. We know from experimental observation that free quarks do not exist in nature and that only $SU(3)$ colour singlet states are observed. This fact is known as *colour confinement*.

As two quarks in a $q\bar{q}$ pair are pulled apart, they exchange virtual gluons. The colour field is confined to a tube between the two quarks. This tube is uniform and characterised by a constant energy density, leading to a linear potential $V(r) \sim \kappa r$, where κ is a constant roughly equal to the size of a hadron, i.e. $\kappa \sim 1$ GeV/fm.

As the potential is linear, the energy in the field increases as the distance between the quarks r increases. At some point, this leads to the production of a new $q'\bar{q}'$ pair from the vacuum. This splits the colour flux tube in two. This process repeats, leading to the production of n hadrons by the creation of $n-1$ new $q_i\bar{q}_i$ pairs.

This picture of hadronisation described above forms the basis of the Lund string model used by Monte Carlo generators [17]. Another model, the cluster mode[18], describes the same physics but through a different lens. In short, the cluster model is based on *QCD pre-confinement*, or namely the fact that the components of colour singlet pairs end up near to each other in phase space [19]. These pairs are the basis of clusters, which are just highly excited hadronic states. The spectrum of clusters formed is universal, in the sense that it is independent of the hard process and energy of the collision. Hadrons then emerge from the decay of these clusters.

1.3 Jets

It is feasible to calculate cross-sections for exclusive (or semi-inclusive) production of some hadrons in the final state. To do so, one needs to consider Fragmentation Functions $D_f(z)$, or the probability of a quark of flavour f will hadronise into a hadron of species D with fraction z of the total energy available. Like PDFs, Fragmentation Functions must be determined from data and cannot be calculated from first principles due to their non-perturbative nature.

It is impractical, however, to explicitly calculate final states with more than just a few hadrons. There are just too many hadrons present in the final state and, more importantly, it is experimentally impossible to explicitly determine the exact species of each hadron.

The solution to this problem is jets. Broadly speaking, jets are an agglomeration of particles roughly produced in the same direction, specifically those particles radiated by a hard parton produced in a scattering event. They can be defined at multiple levels: at parton level, jets are defined using the particles produced in the hard process and parton shower; at hadron level, jets

are defined from the hadrons obtained after modelling hadronisation. Experimentally, jets are defined based on tracks and calorimeter clusters in a given event. This section is dedicated to understanding how they are defined and their use.

1.3.1 Jet algorithms

The idea of a jet is to collect all of the radiation from a hard parton and treat it, together with the emitter, as a singular object. This allows us to circumvent the problem of identifying individual hadron species, and leads to more accurate measurements and predictions by defining a proxy to the final state parton, accounting for all radiation.

Intuitively, considering the angular ordering of the parton shower, jets form a conical structure. As such, jet definitions broadly aim to reproduce this shape.

Below, three of the main families of jet reconstruction algorithms are reported. These days, the generalised k_t algorithms are the go-to algorithm for most purposes.

Cone algorithms

The first jet algorithm, attributed to Sterman and Weinberg [20], was used to form jets produced in e^+e^- collisions. It relied on two parameters, ε and δ. If an event has a fraction of energy of at least $1-\varepsilon$ concentrated within two cones of half-angle δ, then that event is said to contain two jets. The Sterman-Weinberg algorithm is the first example of a general class of algorithms known as *cone algorithms*.

Cone algorithms generally work by either fixing a cone of radius R around a high energy seed particle and summing the contribution of all particles within that cone to assign a jet momentum, or iteratively repeating this procedure, substituting the seed particle with the jet formed at each step. These approaches are known as *fixed cone* and *iterative cone* algorithms, respectively.

JADE

The JADE algorithm [21], introduced by the JADE Collaboration [22], is the first example of a sequential recombination algorithm. This class of algorithms iteratively combines particles based on a distance parameter. In the case of JADE, an e^+e^- collision experiment, the distance

$$y_{ij} = \frac{2E_iE_j(1-\cos\theta_{ij})}{Q^2} \tag{1.68}$$

was computed between all pairs of particles (i,j). Here, E_i refers to the energy of the particle i[3], θ_{ij} to the angle between the two particles, and Q^2 to the hard scale of the process, corresponding to the centre-of-mass energy. The algorithm identifies the minimum distance $y_{\min} = \min(y_{ij})$ and, if this falls below some threshold $y_{\min} < y_{\text{cut}}$, the particles i and j are combined in a jet. The algorithm then iterates until all possible jets are identified.

Generalised k_t algorithm

In most cases, an algorithm from the family of k_t algorithms is used to define a jet. These are sequential recombination algorithms and are again based on a set of distances:

[3]For hadron colliders, the distance measures used by jet algorithms rely on the transverse momentum of the particles, rather than the energy.

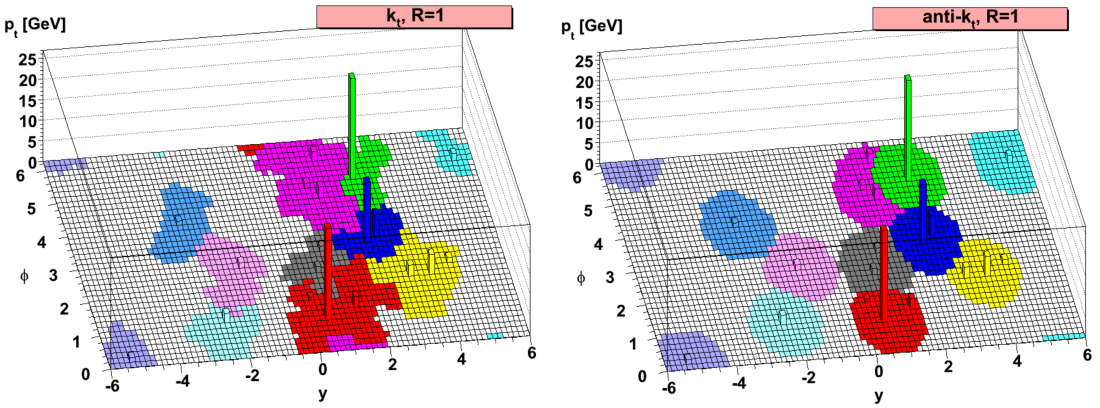


Figure 1.5: The jet shape for jets clustered with the k_t algorithm (left) vs. for those clustered with the anti-k_t algorithm (right) [23].

$$d_{ij} = \min(k_{ti}^{2p}, k_{tj}^{2p}) \frac{\Delta_{ij}^2}{R^2} \tag{1.69}$$

$$d_{iB} = k_{ti}^{2p}. \tag{1.70}$$

Here, d_{ij} is the distance between two particles i and j, with transverse momentum k_{ti} and k_{tj}, respectively. $\Delta^2 = (y_i - y_j)^2 + (\phi_i - \phi_j)^2$ is the distance between the two particles in the rapidity-azimuthal angle $(y - \phi)$ plane[4], while p and R are parameters, with R representing the radius of the jet we wish to cluster. d_{iB}, on the other hand, represents the *beam distance*, and is a reference condition used by the algorithm.

The condition for clustering two particles together is as follows: if $d_{ij} < d_{iB}$, the particles i and j are combined.

The parameter p can take any value, though three cases are notable enough to deserve mention:

- $p = 1$, also known as the k_t algorithm
- $p = 0$, the Cambridge-Aachen (C/A) algorithm
- $p = -1$, or the anti-k_t algorithm.

By varying p, the momentum-dependent weight on the distance d_{ij} varies, modifying the algorithm accordingly. The most common choice nowadays is the anti-k_t algorithm. This is because it is the algorithm which most matches the definition of a jet, i.e. it is the most conic as can be seen in Figure 1.5. The reason for this shape is the $1/k_t^2$ weight applied to the distances. In fact, this weight causes soft particles to be clustered with nearby hard particles rather than with other soft particles, as would occur in the case of the k_t algorithm. This property is known as *soft resilience*.

Recombination scheme

The *recombination scheme* specifies how to combine two particles when clustering jets. The most commonly used recombination scheme for jets produced by hadron colliders is the E-scheme,

[4] These quantities are defined in Section 3.2.1.

where the four-momenta are simply added $p_i + p_j = p_{ij}$. Another notable scheme is the Winner-Takes-All or WTA-scheme [24], where the four-momentum of the particle with the highest p_{T} of the two to be combined is assigned to the jet.

1.3.2 Jets and IRC safety

An important property for jets to maintain is infrared and collinear (IRC) safety. Namely, this means that a jet is ideally defined in such a way that it is insensitive to soft and collinear radiation.

Consider once again Eq. 1.56. Gluons radiated during the parton shower tend to be soft and collinear due to the presence of the divergences. Jets are, by design, meant to contain the radiation off of a hard parton. It is therefore desirable that a given algorithm behaves in this way.

Mathematically, this property translates to the condition

$$\mathcal{O}(X; p_1, \ldots, p_n, p_{n+1} \to 0) \to \mathcal{O}(X; p_1, \ldots, p_n) \tag{1.71}$$

$$\mathcal{O}(X; p_1, \ldots, p_n \parallel p_{n+1}) \to \mathcal{O}(X; p_1, \ldots, p_n + p_{n+1}). \tag{1.72}$$

Equation 1.71 states that, if a particle p_{n+1} is emitted with four-momentum tending to 0, then the jet function $\mathcal{O}$ behaves as if this particle were not present. Likewise, Eq. 1.72 states that, if the particle p_{n+1} is emitted collinear to the particle p_n, then the jet algorithm treats these as one particle with momentum $p_n + p_{n+1}$.

The good news is that the k_t algorithms in use today *are* in fact IRC safe. This is a major factor as to why the algorithm is as widespread as it is. Other aspects, such as IRC safe flavour association, is not yet so common. We will touch upon this in Chapter 2.

1.4 Jet substructure

Once we have defined our jets based on some algorithm, it is interesting to study the internal structure of these jets. This field of study is known as *jet substructure* (JSS).

Jet substructure is a burgeoning field which has its roots in the start of the LHC programme. The aim of the field is to understand the radiation pattern within jets. This is directly tied to the origin of said jets, as jets originating from a top quark, for example, will have a different structure from those originating from a Z boson. This is illustrated in Figure 1.6. This is particularly useful in boosted regimes, when jets are produced at very high transverse momentum. In these cases, large-radius jets are used to fully contain the decay products of the particle giving origin to the jet.

The simplest jet substructure observable is the jet mass, defined as

$$m^2 = \left(\sum_i k_i \right)^2, \tag{1.73}$$

where the sum runs over all particles i with momenta k_i in the jet. Other examples include the jet thrust [26], N-subjettiness [27], and the Lund Jet Plane [28]. We discuss this last observable in further detail in Section 1.4.3.

1.4.1 Jet grooming

To aid in the study of jet substructure, it is often useful to *groom* jets. Grooming is a technique by which jets are cleaned of soft radiation, which may have its origins from particles

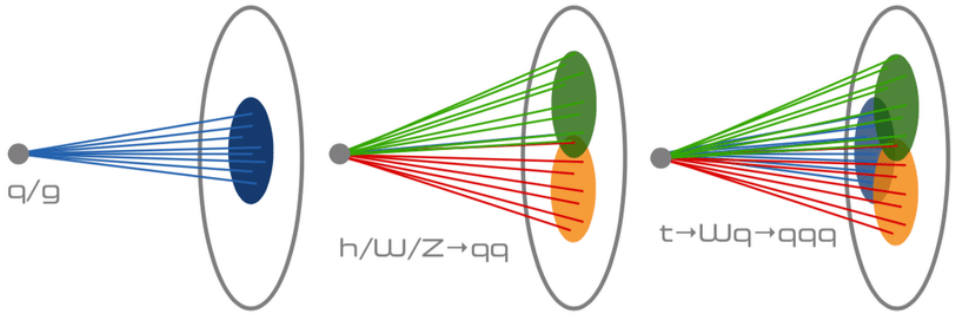


Figure 1.6: An illustration of the substructure of jets originating from quarks/gluons, $W/Z/H$ bosons, and top quarks [25]. Jets from quarks and gluon (left) will tend to have a single prong structure. Jets formed by other particles can have multiple prongs, corresponding to the number of quarks produced when these particles decay, such as $Z \to q\bar{q}$ (centre) and $t \to W^+q$, with further decays of the W boson (right).

other than the seed of the jet. It leaves intact the core constituents, allowing for more accurate predictions to be made.

One notable example of a grooming algorithm is the SoftDrop algorithm [29]. It works by reclustering a given jet with the C/A algorithm in order to obtain the angular-ordered clustering history (emissions) of the jet. Starting from the softest branch, the algorithm grooms away soft emissions until the following condition is met

$$\frac{\min(p_{\mathrm{T}i}, p_{\mathrm{T}j})}{p_{\mathrm{T}i} + p_{\mathrm{T}j}} = z_{\mathrm{cut}} \left(\frac{\Delta R_{ij}}{R}\right)^{\beta} \tag{1.74}$$

where $p_{\mathrm{T}i}$ and $p_{\mathrm{T}j}$ refer to the p_{T} of the subjets i and j, respectively, and ΔR_{ij} is the distance between the two in the $y - \phi$ plane. z_{cut} and β are parameters which regulate the emissions to be cut. The variable z refers to the relative p_{T} of the subjet i with respect to that of both i and j, $p_{\mathrm{T}i}/(p_{\mathrm{T}i} + p_{\mathrm{T}j})$. If the removal condition is never satisfied, the jet can either be removed or left as is.

Another notable grooming technique is known as *trimming* [30]. Trimming involves reclustering a jet's constituents into jets of smaller radius. If the p_{T} fraction carried by one of these smaller-radius jets is below a certain threshold, i.e. 5%, it and all its constituents are removed. The remaining constituents are clustered back into a jet of the original radius.

1.4.2 Heavy flavour and JSS

Jet substructure can give insight on the particle which gives origin to a jet. For instance, quark/gluon jet discrimination is possible due to the fact that gluon-initiated jets will have a higher particle multiplicity due to the large colour factor and more wide-angle emissions. Differences among quark-initiated jets are also possible, arising mainly from the presence of a heavy flavour quark. If the seed particle has a significant mass, such as in the case of c-quarks and b-quarks, the mass scale intrinsically affects the evolution of the final state.

If we consider the probability of emitting a gluon carrying energy fraction z at an angle θ,

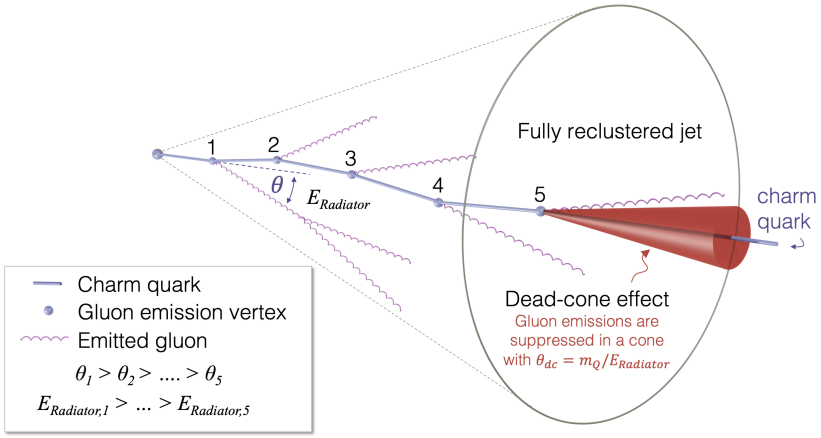


Figure 1.7: A schematic representation of the dead cone effect initiated by a charm quark [32].

the full expression with mass is found to be

$$P(z,\theta^2)dzd\theta^2 = \frac{\alpha_s C_F}{2\pi}\frac{dz}{z}\frac{d\theta^2}{\theta^2 + \frac{m^2}{E^2}}. \tag{1.75}$$

When $m \to 0$, we can read the soft and collinear divergences with which we are now familiar. Instead, if $m \not\approx 0$, we find that the collinear divergence is now *shielded* by the mass of the quark. This is known as the *dead cone effect* [31] and it is illustrated in Figure 1.7. The emission process is no longer scale invariant, and within a cone of angle $\theta_D^2 = m^2/E^2$, the emission of gluons is suppressed.

It is useful to note that, to first order, the dead cone angle can be equivalently seen as depending on m/p_{T} ratio. This can be seen by writing the four-momenta of the quarks and gluon involved in the emission in terms of $(p_{\mathrm{T}}, y, \phi, m)$. The dead cone condition predicting a suppression for $\theta \lesssim \theta_D$ thus becomes $\Delta R \lesssim m/p_{\mathrm{T}}$, where the p_{T} refers to that of the quark after emission.

1.4.3 The primary Lund Jet Plane

The primary Lund Jet Plane is the main observable which this thesis focuses on. As the name would suggest, the Lund Jet Plane is tightly linked to parton shower development. The precursors to the Lund Plane were Lund diagrams, used to describe the kinematic distribution of radiation from a QCD dipole within a y-k_t plane [33]. The kinematically allowed region of phase space forms a triangle bounded by $|y| < \ln(\sqrt{s}/k_t)$, leading to the plane's characteristic shape.

As the field of jet substructure began to take hold, the Lund diagram provided a way to visualise the effect of various techniques used to simplify analytic calculations, such as grooming and trimming. The description was later adapted for jets in [28]. In this work, the authors used the plane to create a two-dimensional representation of the clustering history, and therefore structure, of a jet.

The primary Lund Jet Plane is formed starting from an anti-k_t jet, which is then reclustered using the Cambridge-Aachen algorithm. After reclustering, the jet is de-clustered, working backwards through the C/A algorithm following the hardest branch at each splitting, to obtain a list of the branchings from the base of the jet. The variables $\ln(k_t/\mathrm{GeV})$ and $\ln(1/\Delta)$ where

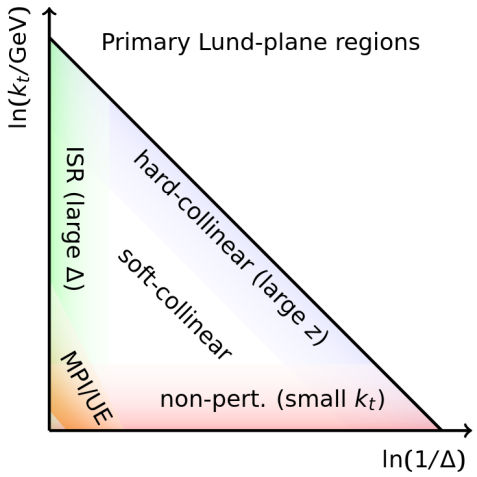


Figure 1.8: The kinematic regions within the primary Lund Jet Plane [28].

$$\begin{cases} \Delta^2 = (y_a - y_b)^2 + (\phi_a - \phi_b)^2 \\ k_t = p_{\mathrm{T}b}\Delta \end{cases} \tag{1.76}$$

are plotted for each splitting. Here, a and b refer to the two branches after the splitting, with $p_{\mathrm{T}a} > p_{\mathrm{T}b}$. In the limit $p_{\mathrm{T}a} \gg p_{\mathrm{T}b}$ and $\Delta \ll 1$, k_t refers to the transverse momentum of subjet b with respect to the emitter $a + b$.

The lack of momentum weighting which characterises the C/A algorithm allows for the algorithm to mimic the angular ordering of $1 \to 2$ splittings in QCD. Anti-k_t jets are not suitable as they are lacking this key angular-ordering which mimics the perturbative evolution of the final state and which allows us, in some sense, to work "backwards" through the shower. As each emission is plotted in the Lund Plane, the perturbative QCD evolution down to the hadronisation scale emerges, as each emission is found at smaller angles and lower k_t than the last.

Figure 1.8 shows the various kinematic regions identifiable within the primary Lund Jet Plane. Towards the left edge of the plane, an area characterised by large angle initial state radiation is found. A hard-collinear region where the two splittings have roughly equal p_{T} follows the diagonal edge of the triangle. This edge, in fact, is due to the kinematic limit $k_t < \frac{1}{2}p_{\mathrm{T},jet}\Delta$. In the bottom left corner, there is a region which originates from multiparton interactions and the underlying event, whose emissions within the jet are characterised by a large angle and small k_t. A non-perturbative region is found at small k_t, and lastly soft-collinear radiation falls towards the centre of the plane.

For the sake of completeness, we would like to mention that a splitting is characterised not just by the (Δ, k_t) coordinates, but also by

$$\begin{cases} z = \frac{p_{\mathrm{T}b}}{p_{\mathrm{T}a}+p_{\mathrm{T}b}} \\ m^2 = (p_a + p_b)^2 \\ \psi = \tan^{-1}\frac{y_b - y_a}{\phi_b - \phi_a}. \end{cases} \tag{1.77}$$

There exist also alternative representations of the Lund Plane, where instead of $\ln(k_t/\mathrm{GeV})$, $\ln(1/z)$ is shown. These other Lund coordinates have been used to train graph neural network based taggers, such as LundNet [34].

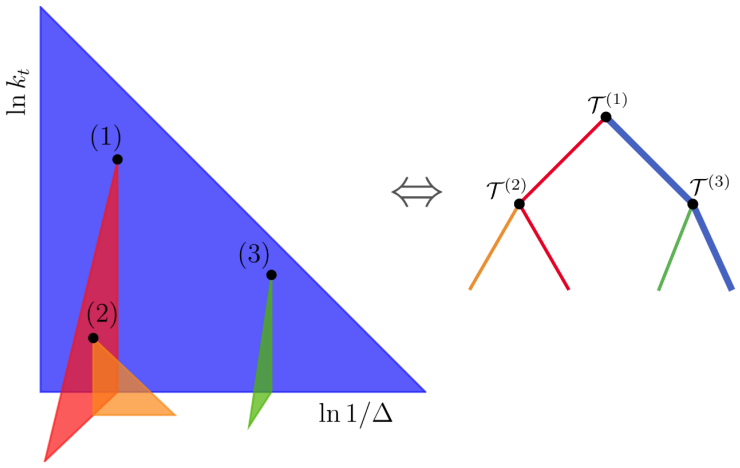


Figure 1.9: The primary Lund Plane for a hypothetical jet with emissions (1) and (3), and the secondary Lund Plane (2) from branch (1) and corresponding tree structure of the jet showing this evolution. Each splitting can be characterised by the set of Lund variables $\mathcal{T}$, given by Eq. 1.76 and Eq. 1.77 [34].

It is also possible to define Lund Planes beyond the primary, by considering further splittings off of the emitted subjets when present. This leads to the definition of secondary planes, tertiary planes, etc. This process is shown in Figure 1.9.

1.4.4 Jet angularities

The generalised angularities [35], [36], henceforth referred to as just *angularities*, are a class of observables designed for quark/gluon jet discrimination. Starting from an anti-k_t jet, the jet constituents are reclustered with the C/A algorithm and the WTA recombination scheme. The jet angularities are then defined as:

$$\lambda_\alpha^\kappa = \sum_i \left(\frac{p_{\mathrm{T}i}}{p_{\mathrm{T}jet}}\right)^\kappa \left(\frac{\Delta R_i}{R_0}\right)^\alpha \tag{1.78}$$

where the sum runs over all constituents of a jet of radius R_0, and ΔR_i indicates the distance between the i-th constituent and the jet axis. λ_α^κ is IR-safe for $\kappa \geq 1$, and we shall assume $\kappa = 1$ going forward, unless otherwise specified.

A few angularities are distinguished enough to have their own name. These include λ_1, known as the *jet width* and λ_2, known as the *jet thrust*, amongst others. The idea behind these and other angularities is to emphasise various aspects of a jet's substructure in order to highlight features typical of quark or gluon jets. By tuning the β parameter, one can accentuate either wide-angle or collinear radiation, while tuning the κ parameter has the same effect on harder or softer hadrons.

Recently, a proposal for modified angularities has been put forth [37]. The variants are defined with the same reclustering step as Eq. 1.78. From here, two vectors are defined from the WTA

axis:

$$n_0 = (\cosh y, \cos\phi, \sin\phi, \sinh y) \tag{1.79}$$

$$n = \left(\frac{m_{T,n}}{p_{\mathrm{T},n}}\cosh y, \cos\phi, \sin\phi, \frac{m_{T,n}}{p_{\mathrm{T},n}}\sinh y\right) \tag{1.80}$$

where ϕ and y are the azimuthal angle and rapidity of the WTA axis, and $m_{T,n} = p_{\mathrm{T},n} + m_n$ is the transverse mass of the particle aligned with the WTA axis. The (E, p_x, p_y, p_z) convention is used to define the vectors.

The variants are then defined as

$$\dot{\lambda}_0^\alpha = \sum_i \frac{p_{\mathrm{T},i}}{p_\mathrm{T}}\left(\frac{2p_i\cdot n_0}{p_{\mathrm{T},i}R^2}\right)^{\frac{\alpha}{2}} \qquad \dot{\lambda}^\alpha = \sum_i \frac{p_{\mathrm{T},i}}{p_\mathrm{T}}\left(\frac{2p_i\cdot n}{p_{\mathrm{T},i}R^2}\right)^{\frac{\alpha}{2}} \tag{1.81}$$

$$\hat{\lambda}_0^\alpha = \sum_{i\neq n} \frac{p_{\mathrm{T},i}}{p_\mathrm{T}}\left(\frac{2p_i\cdot n_0}{p_{\mathrm{T},i}R^2}\right)^{\frac{\alpha}{2}} \qquad \hat{\lambda}^\alpha = \sum_{i\neq n} \frac{p_{\mathrm{T},i}}{p_\mathrm{T}}\left(\frac{2p_i\cdot n}{p_{\mathrm{T},i}R^2}\right)^{\frac{\alpha}{2}}. \tag{1.82}$$

where the sums in Eq. 1.82 exclude the particle aligned with the WTA axis. These variants are specifically designed to bring to light the mass of the particle aligned with the WTA-axis. For jets at hadron level seeded by heavy flavour quarks, if the decay of the heavy hadron is suppressed, the particle aligned with the WTA axis will be a hadron containing the same flavour of quark which seeded the jet in accordance with the leading particle effect. As such, the observables defined in Eq. 1.81 will contain an *explicit mass dependence*. In the quasi-collinear limit, for example,

$$\dot{\lambda}_0^\alpha \approx \sum_{i\neq n} \frac{p_{\mathrm{T},i}}{p_\mathrm{T}}\left(\frac{m_i^2}{p_{\mathrm{T},i}^2R^2} + \frac{\Delta R_i^2}{R^2}\right)^{\frac{\alpha}{2}} + \left(\frac{m_n^2}{p_{\mathrm{T},n}^2R^2}\right)^{\frac{\alpha}{2}}. \tag{1.83}$$

The end result is a set of observables (Eq. 1.81) sensitive to quark mass. Figure 1.10 shows the expected distributions of jet angularities as defined in Eqs. 1.78, 1.81, and 1.82 for ungroomed jets of radius $R = 0.4$ with transverse momentum above 750 GeV from proton-proton collisions at centre-of-mass energy $\sqrt{s} = 13.6$ TeV. Mass effects are clearly visible for the $\dot{\lambda}$ variants, defined by including the particle aligned with the WTA axis in the summation.

The measurement described in this thesis focuses on the jet angularities as they are defined in Eq. 1.78. Only this observable will be unfolded. Regardless, the modified angularities are to be used to convince the reader of the successful reconstruction of the heavy flavour hadron from its decay products.

1.5 State-of-the-art measurements

In this section, we give an overview of the current state-of-the-art of the measurements relevant to this thesis. These include measurements of Z+heavy flavour jet production and of jet substructure observables.

1.5.1 Z+jets

Z+jets, or, more generally, V+jets, is one of the most abundant class of processes at the LHC, as shown in Figure 1.11. As an intrinsically NLO process, it is ideally suited for studies of

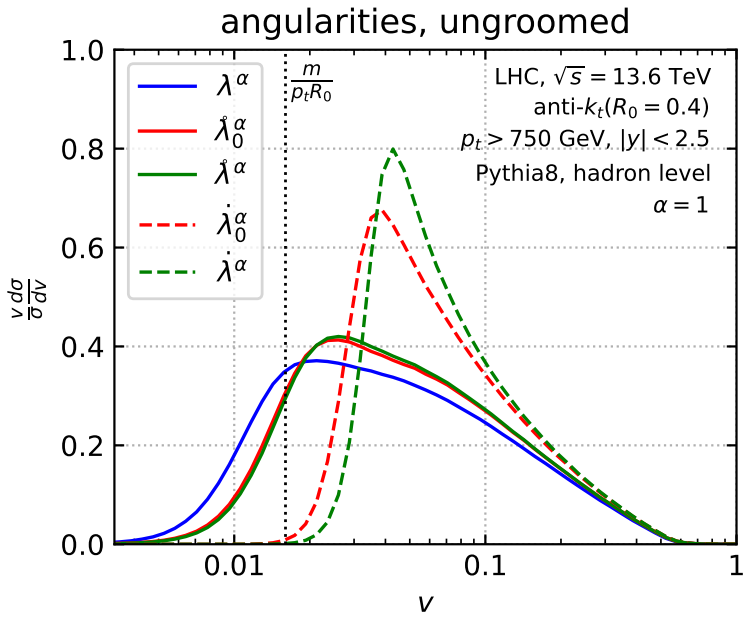


Figure 1.10: The jet angularities and mass-dependent variants calculated on anti-k_t jets of radius $R = 0.4$ at high transverse momentum with $\alpha = 1$. The line at $v = m/p_{\mathrm{T}}R_0$ represents the nominal value of dead cone angle [37].

perturbative QCD in a wide range of both kinematics and jet multiplicity. In addition to this, the leptonic decay channel of the Z provides a clear experimental signature of the process.

Z+jets comprises roughly 20-30% of the total Z production cross section. Each additional jet corresponds to an additional emission, and is thus suppressed by an order of α_s in the matrix element. This in and of itself constitutes a way to extract α_s [39].

The high precision that can be reached with Z+jets provides important data for PDF fits, allowing for their constraint and a decrease in modelling uncertainty in future measurements. Z+jets also constitutes an important background for Higgs measurements, especially those which focus on Higgsstrahlung production channel [4], and for searches for new physics beyond the Standard Model.

Figure 1.12 shows a Feynman diagram for the Z+jets production process.

Z+jets measurement

Z+jets production has been studied at the ATLAS Experiment at various centre-of-mass energies, including 7 TeV [41] and 8 TeV [42]. Here we focus on Z+jets production at $\sqrt{s} =$ 13 TeV collisions with 3.16 fb^{-1} of data collected during Run 2 in 2015 [43]. The measurement focuses on final states with up to 7 jets of radius $R = 0.4$. Jets are required to have transverse momentum $p_{\mathrm{T}} > 30$ GeV and to be found in the central region of the detector. The differential cross section is measured with respect to several different kinematic observables, such as the transverse momentum and rapidity.

Results are compared to predictions obtained with a various MC generators. These include LO predictions from MG5_aMC@NLO [44] with CKKWL merging with a Pythia8 [45]

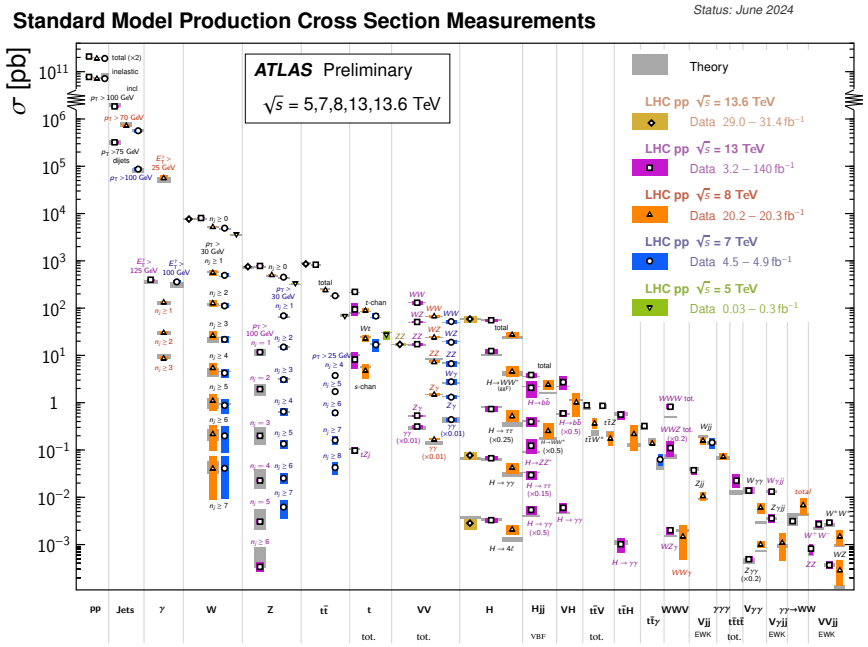


Figure 1.11: The production cross section for the main processes observed by ATLAS from proton-proton collisions at $\sqrt{s} = 7, 8, 13, 13.6$ TeV. The measured Z+jets cross sections are shown for number of jets n_j [38].

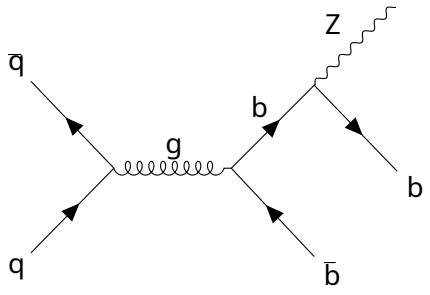


Figure 1.12: A Feynman diagram for Z production in association with b-quarks [40].

(MG5_aMC+PY8 CKKWL) parton shower; another set of MG5_aMC@NLO predictions, this time with FxFx merging and a Pythia8 shower (MG5_aMC+PY8 FxFx); and predictions with the ALPGEN [46] ME generator with a Pythia6 [47] parton shower (ALPGEN+PY6). Also included are NLO predictions from Sherpa 2.2.1 [48] (Sherpa 2.2) and a NLO fixed-order prediction obtained with BLACKHAT [49] generator and showered with Sherpa.

Up to 4 jets can be included in the ME calculation; afterwards, it is necessary to predict the production cross section perturbatively from the parton shower. Predictions are in agreement with data up to $N_{\text{jets}} = 4$. Afterwards, results diverge. Results and predictions for the cross section are shown in Figure 1.13.

We note that there exist several other Z+jets measurements published by other experimental collaborations, such as one carried out at $\sqrt{s} = 13$ TeV by the CMS Collaboration [50]. We also note that there is another ATLAS measurement of Z+jets at $\sqrt{s} = 13$ TeV which which focuses on high p_{T} jets above 100 GeV [51].

$Z + b$-jets

For Z+jets production with heavy flavours, considerations such as the *flavour scheme* (FS) start to become relevant. The flavour scheme indicates which quarks are found within the proton, i.e. whether a given flavour is to be considered as part of the PDF evolution and or if it can only be produced perturbatively. Specifically, a 5FS considers a massless b-quark PDF calculated from the evolution of the gluon PDF above the b mass threshold. The massless approximation is valid if $Q^2 \gg m_b^2$. As the b-quark PDF is calculated, Feynman diagrams for partonic cross sections with b-quarks in the initial state are accessible. A 4FS, on the other hand, only considers the perturbative production of b-quarks in the final state from hard gluon splittings. Bottom quarks are treated as massive, so they are not considered within the proton. The only contribution thus comes from the ME, though additional contributions from gluon splittings in the parton shower are possible in both cases.

ATLAS has published a measurement of $Z+ \geq 1\,b$-jet production at $\sqrt{s} = 7$ TeV [52] and more recently published another measurement focusing on final states with both $Z+ \geq 1\,b$-jet and $Z+ \geq 2\,b$-jets [53], which we describe below.

The measurement is carried out on Run 2 data, corresponding to 139 fb^{-1} and a centre-of-mass energy $\sqrt{s} = 13$ TeV. Differential cross sections as a function of variables such as p_{T} of the Z boson, p_{T} of the leading b-jet, and di-b-jet mass m_{jj}, and the angular separation between b-jets $\Delta\phi_{bb}$ are obtained. Results are compared to a number of NLO predictions, such as MG5_aMC+PY8 with FxFx merging with both the 4FS and 5FS, Sherpa 2.2, as well as to fixed-order predictions obtained at NLO and NNLO. These are corrected for gluon splittings, as will be described in Chapter 2.

Figure 1.14 shows the cross section as a function of the leading b-jet $p_{\text{T},b}$ and $\Delta\phi_{bb}$. For $p_{\text{T},b}$, predictions for the 5FS agree with data better than predictions obtained with the 4FS. Both the NLO and NNLO fixed-order predictions model the data well. For $\Delta\phi_{bb}$, all predictions agree with data within uncertainties, though the 4FS predictions tend to underestimate the data for collinear and back-to-back jets.

$Z + bb$ boosted

At very high transverse momenta, or boosted regimes, it is possible to execute tests of the Standard Model in extreme phase space. For the $Z+bb$ process in particular, the boosted regime allows us to probe hard and collinear physics.

An ATLAS Run 2 measurement considers precisely this final state in the boosted regime, mimicking as much as possible the selections for similar boosted Higgs analyses in order to precisely

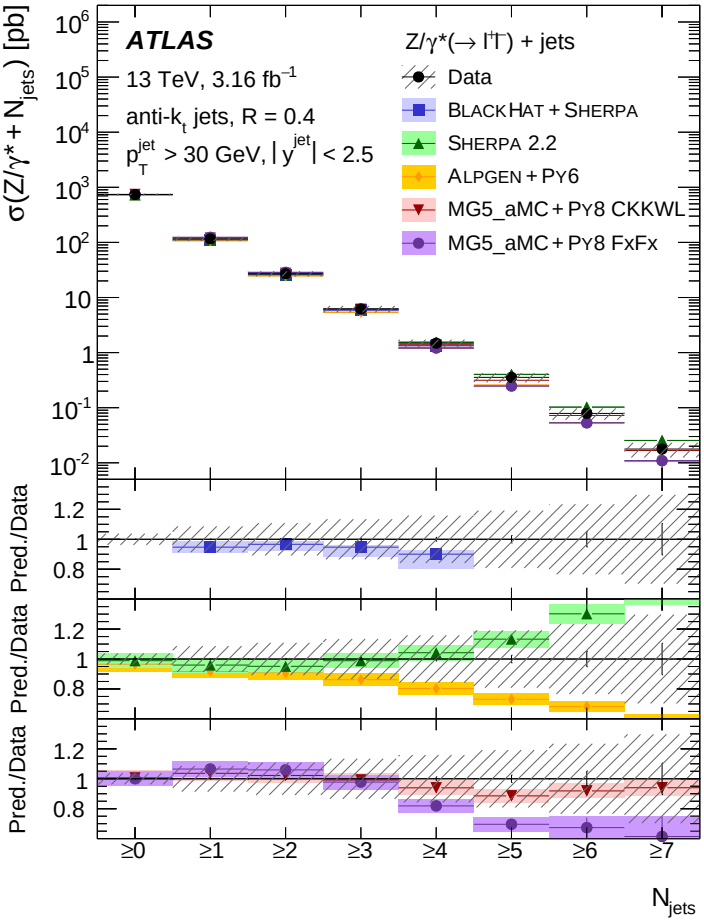


Figure 1.13: The Z+jets production cross section as a function of N_{jets} compared to predictions obtained with MG5_aMC+PY8 CKKWL, MG5_aMC+PY8 FxFx, ALPGEN+PY6, Sherpa 2.2, and BLACKHAT+Sherpa [43].

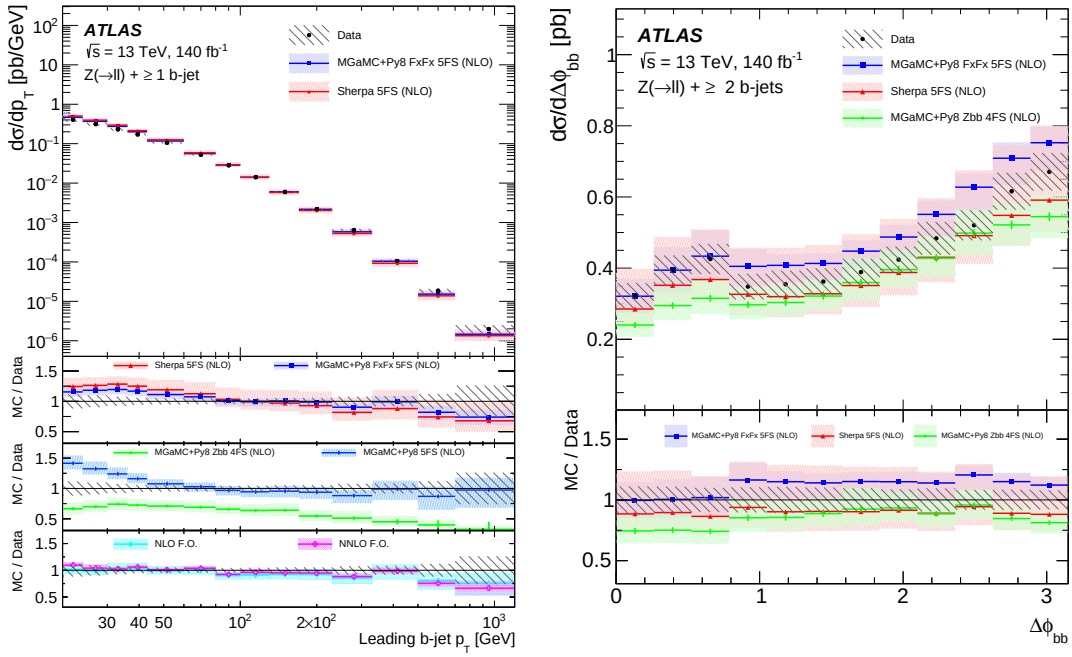


Figure 1.14: The differential cross section for the leading b-jet p_{T} (left) and for the angular separation between di-b-jets (right) compared to predictions [53].

.

characterise this principal background, composed mainly of gluon splittings. The measurement utilises 36.1 fb^{-1} of data collected at $\sqrt{s} = 13$ TeV [54].

The measurement studies large-R jets of radius $R = 1.0$ produced in the central region of the detector with $p_{\mathrm{T}} > 200$ GeV. Jets are tagged through the use of small-R jets of radius $R = 0.2$ with $p_{\mathrm{T}} > 10$ GeV, defined with only track constituents. The $R = 0.2$ jets are identified as b-jets by the flavour tagging algorithm MV2c10 [55] and, if two such jets are associated to a large-R jet using the ghost association criterion[5], then the large-R jet is considered a b-jet. Subclasses of flavour are determined by the number of associated small-R jets.

Results are compared to NLO predictions obtained with Sherpa 2.2.1 with the 5FS, as well as Sherpa 2.2.10 with the 5FS and 4FS variations, and with fusing, which mixes elements of both. Predictions with MadGraph5_aMC@NLO+Pythia8 were also obtained with 4FS and 5FS variations.

Results for the differential cross section as a function of the large-R jet p_{T}, mass, and angular separation $\Delta R(b, b)$ of the small-R b-jets is shown in Figure 1.15. The Sherpa 2.2.10 and MadGraph5_aMC@NLO predictions in the 5FS describe the data well. The 4FS Sherpa 2.2.10 and 5FS Sherpa 2.2.10 predictions underestimate the total cross section, while the Sherpa 2.2.10 fusing prediction is close to the Sherpa 2.2.10 5FS estimate.

1.5.2 The Lund Jet Plane

Since its inception in 2016, a number of measurements of the Lund Jet Plane and related observables have been carried out. The Lund Plane has been measured on inclusive jets by

[5] Cfr. Section 2.1 for a detailed explanation of ghost association.

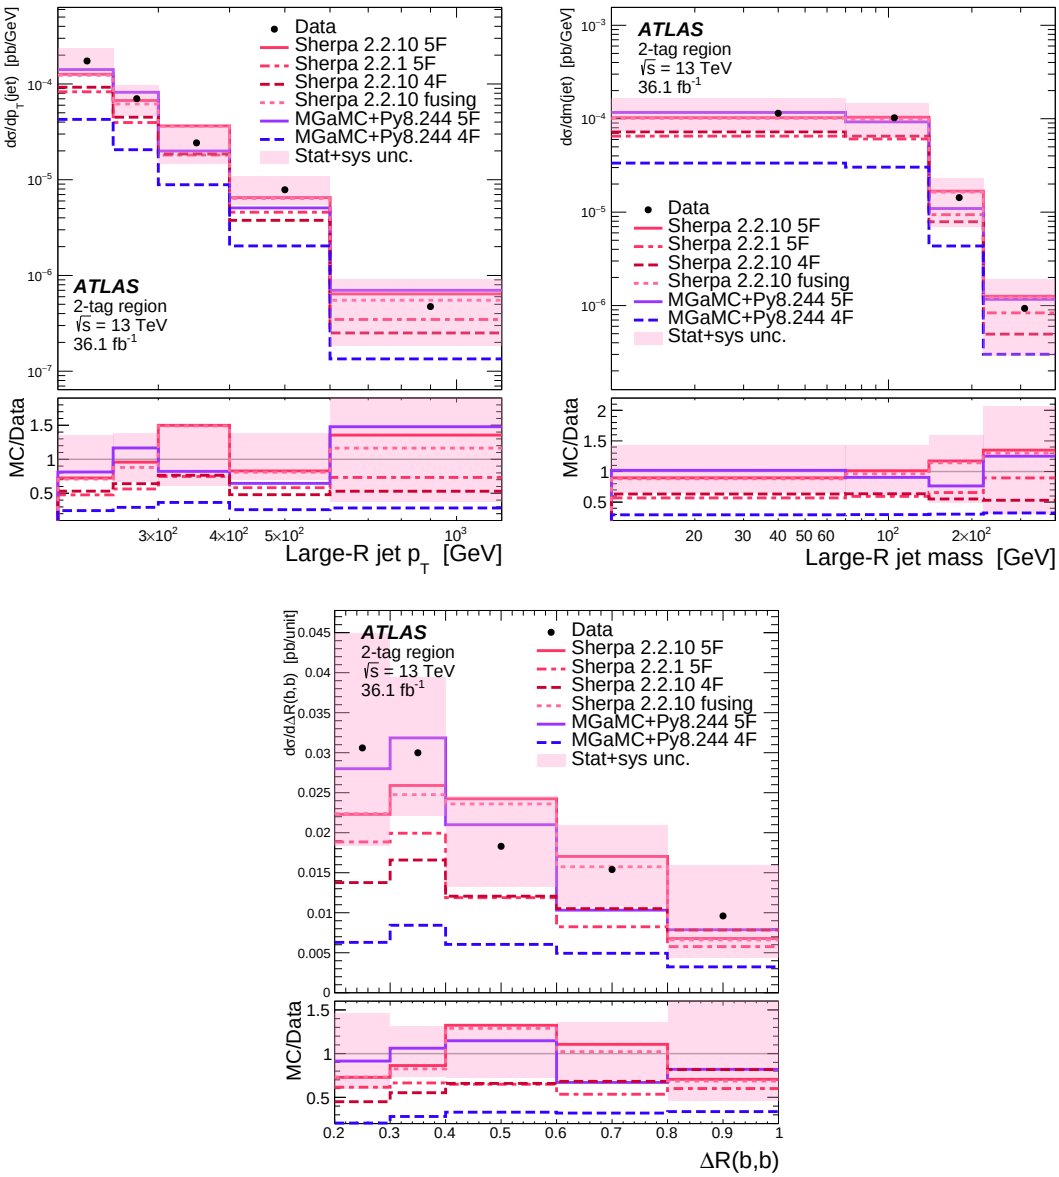


Figure 1.15: The differential cross sections as a function of large-R jet p_{T}, jet mass, and angular separation of the small-R b-jets for Z+bb events [54].

the ATLAS, CMS, and ALICE Collaborations [56]–[58] and on heavy flavour jets by the LHCb Collaboration [59]. The ATLAS and CMS measurements focus on high p_{T} jets, whereas the ALICE measurement looks at jets in a lower p_{T} range. In this section, we describe in detail the ATLAS and ALICE measurements. We then move on to discussing a measurement of the Lund Jet Plane in heavy flavour jets[59], and observations of the dead cone effect [32], [60].

For the sake of completeness, we also wish to mention a recent ATLAS measurement of the Lund Jet Plane on top and W jets[61].

ATLAS measurement on inclusive jets

The first measurement of the Lund Jet Plane is attributed to the ATLAS Collaboration [56]. It utilises all proton-proton data from Run 2 collected at $\sqrt{s} = 13$ TeV, corresponding to 139 fb^{-1}.

The measurement focuses on final state with balanced di-jet events, of radius $R = 0.4$ in the central region. The leading jet[6] is required to have a $p_{\mathrm{T}}^{\mathrm{leading}} > 675$ GeV, while the sub-leading jet is required to have a transverse momentum no less than 1.5× that of the leading jet $p_{\mathrm{T}}^{\mathrm{subleading}} > p_{\mathrm{T}}^{\mathrm{leading}}/1.5$. This facilitates the interpretation of the event as a $2 \to 2$ scattering process within QCD.

The Lund Plane is measured exclusively on the *charged* component of the jet. This is because of the higher resolution provided by the tracking system at ATLAS compare to the calorimeter. For more details, see Chapter 3. Tracks are required to have a $p_{\mathrm{T}} > 500$ MeV and to lie in the central region of the detector. They are associated to jets by a geometric matching criterion based on the distance in the pseudorapidity-azimuthal angle plane (η, ϕ)[7]. Specifically, the distance $\Delta R(\text{track, jet}) < 0.4$, where the cut-off corresponds to the jet radius. Tracks are clustered into jets using the C/A algorithm. The Lund Jet Plane is then measured on the jets clustered with track constituents. The unfolded particle-level Lund Jet Plane is shown in Figure 1.16.

Results are compared to predictions obtained with MC generators including: LO predictions obtained with Pythia v8.230, NLO predictions obtained with PowHeg [62]–[64] and showered with Pythia 8.230 (PowHeg+Pythia), NLO predictions obtained Sherpa 2.2.5 with either the AHADIC [18], [65], a cluster hadronisation model, or the Lund string hadronisation model, and NLO predictions obtained with Herwig 7.1.3 [66] with either the angle-ordered or dipole parton shower. Figure 1.17 shows ratios of the Lund Plane as predicted by these different generators, specifically by varying the LO/NLO ME in Pythia vs. PowHeg+Pythia, the Herwig angle-ordered and dipole showers, and the AHADIC and Lund string hadronisation models in Sherpa. Clear differences arise in the latter two cases, particularly in the regions dominated by hard radiation in the case of the parton shower variation, and in the collinear region in the case of the hadronisation model variation.

The comparison to data is carried out by slicing the Lund Jet Plane and focusing on specific intervals, corresponding to fixing an interval in ΔR and varying z or vice versa. Figure 1.18 shows two such slices, by varying z in the region between $0.67 < \ln(R/\Delta R) < 1.00$ or by varying ΔR in the region between $1.80 < \ln(1/z) < 2.08$. No single prediction is able to fully describe the data.

Similar selections are also used to measure the Lund subjet multiplicity, corresponding to the number of emissions in the Lund Plane. Specifically, the number of emissions in just the primary Lund Plane were counted $N_{\mathrm{Lund}}^{\mathrm{Primary}}$, as well as the total number of emissions within all Lund Planes N_{Lund}. The main difference compared to the measurement described above is a

[6] The terms *leading jet* and *subleading jet* refer to the jets with the highest and second highest transverse momentum in the event, respectively.

[7] For a definition of these quantities, see Section 3.2.1.

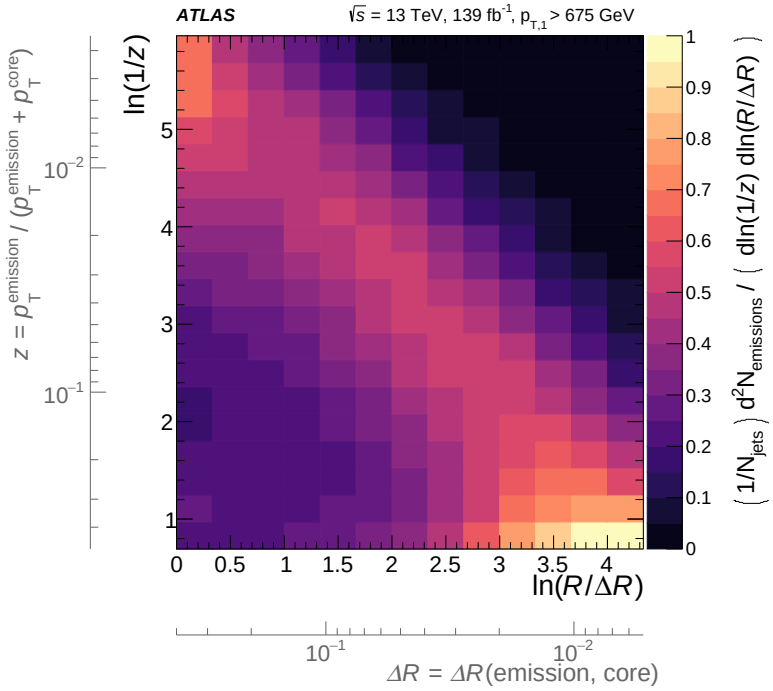


Figure 1.16: The particle-level Lund Jet Plane for the charged component of a jet, normalised to the number of jets which pass the event selection [56].

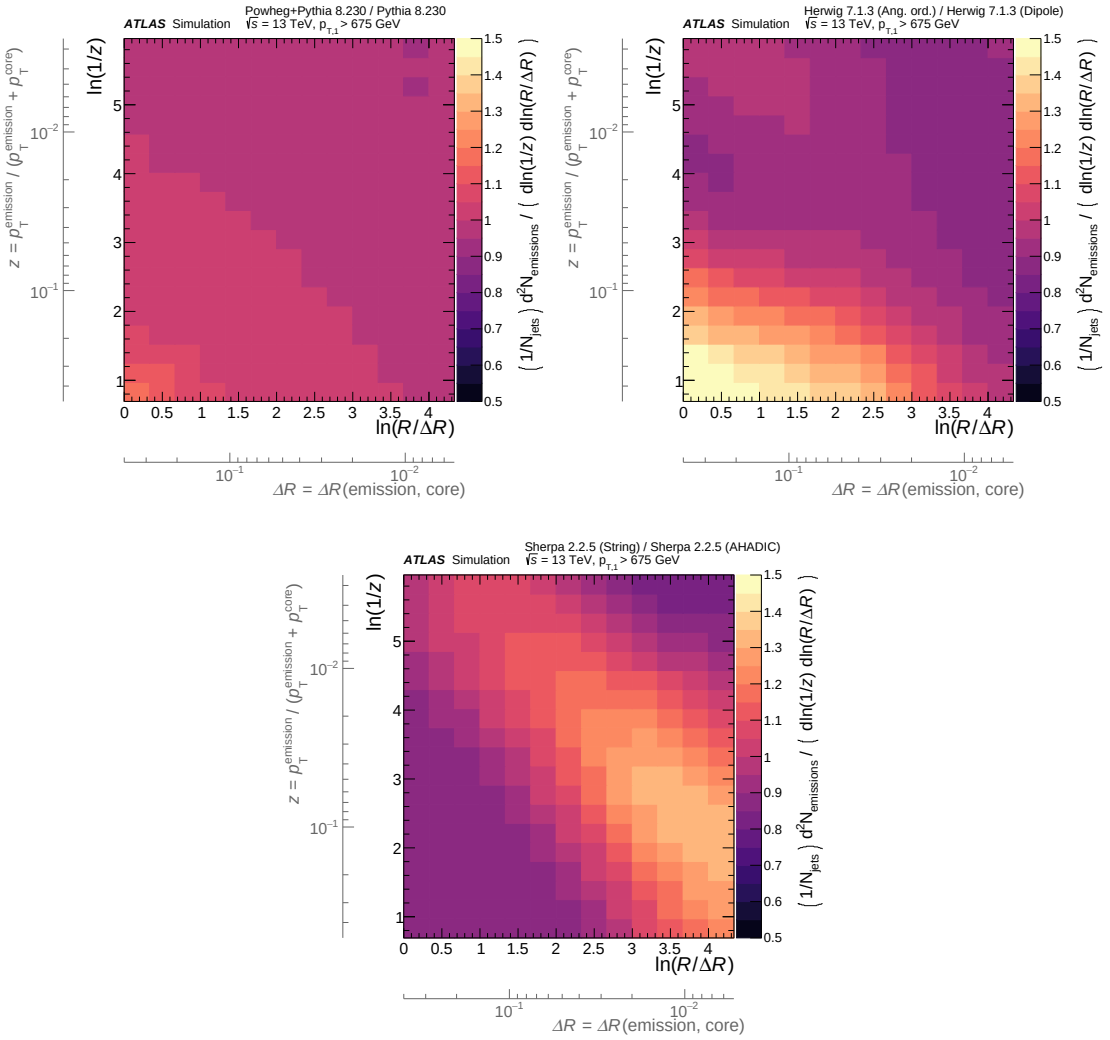


Figure 1.17: The ratio of the Lund Jet Planes in di-jet events obtained by either varying the order in perturbation theory at which events were generated (top left), the parton shower model (top right), or the hadronisation model (bottom) [56].

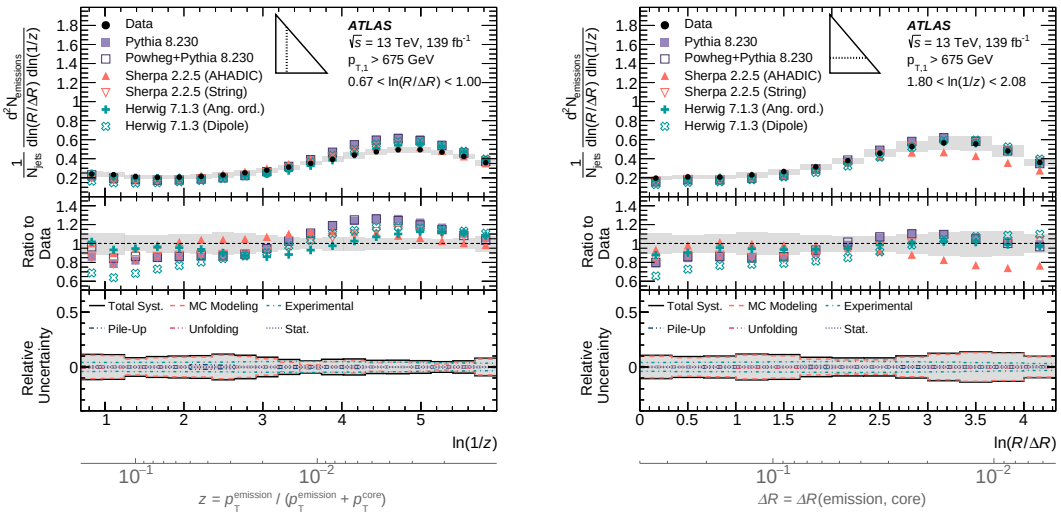


Figure 1.18: The results of the di-jet measurement of the Lund Jet Plane on inclusive jets compared to various predictions [56].

decrease in the p_{T} threshold, down to $p_{\mathrm{T}} > 120$ GeV for both jets. We limit ourselves to citing this result [67].

ALICE measurement on inclusive jets

ALICE has also looked at the Lund Jet Plane in inclusive jets, but in a much lower p_{T} range than the ATLAS measurement described above. The measurement focuses on charged particle jets of radius $R = 0.4$ in the p_{T} range 20-120 GeV produced in the central region in minimum bias events from proton-proton collisions at $\sqrt{s} = 13$ TeV. The data corresponds to an integrated luminosity of 22.55 nb^{-1} [58].

Results are compared to simulations carried out with Pythia8 with the Monash [68] tune, Herwig7, and Sherpa 2.2.8 with the Lund string and AHADIC hadronisation models. Figure 1.19 shows the results and the predictions for the projections of the Lund Plane onto the $\ln(R/\Delta R)$ and $\ln(k_t/\mathrm{GeV})$ axes in the perturbative region of the plane and for collinear splittings, respectively. For narrow splittings, it is found that the generators disagree at high k_t by up to 50%. Agreement is generally within 20%, though no single prediction fully agrees with the data.

LHCb measurement of the light jet and b-jet Lund Plane and their ratio

The LHCb Collaboration has recently published a measurement of the Lund Plane in light quark enriched and b-initiated jets [59] using data collected at $\sqrt{s} = 13$ TeV and corresponding to 5.4 fb^{-1} of integrated luminosity. The light quark enriched sample corresponds to a selection of Z+jets events[8], while the b-initiated sample contains jets containing a fully reconstructed $B^{\pm}$ meson. This is achieved by reconstructing the $B^{\pm} \to J/\psi(\to \mu^+\mu^-)\pi^{\pm}$ decay channel.

As LHCb is a forward detector, the measurement exclusively considers jets of radius $R = 0.5$ produced in the forward region, between a rapidity range of $2.5 < y < 4.0$. Jets are also required

[8]The inclusive multi-jet analyses described above also contain gluon-initiated jets. These are suppressed when considering Z+jets.

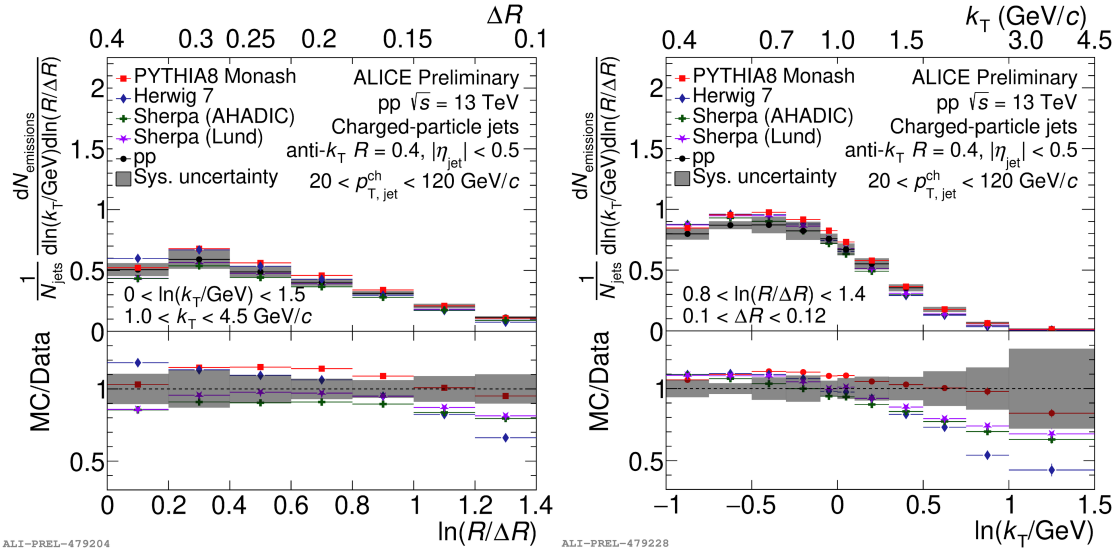


Figure 1.19: The Lund Jet Plane projected onto the $\ln(R/\Delta R)$ axis in a perturbative region ($k_t > 1.0$ GeV) (left) and onto the $\ln(k_t/\text{GeV})$ axis in a region charachterised by collinear splittings ($0.1 < \Delta R < 0.12$) (right) [58].

to have a $p_{\mathrm{T}} > 20$ GeV. A number of jet identification requirements are used to suppress noise and misidentified leptons. The Lund Jet Plane is measured on jets reconstructed with both charged and neutral objects. When constructing the Lund Jet Plane for $B^{\pm}$ jets, the C/A clustering step is undertaken with the WTA recombination scheme to ensure that the hadron is found within the hardest branch of the jet.

Results are compared to a sample simulated in Pythia8 with the LHCb tune [69]. In the results, shown in Figure 1.20, the ratio of the projection onto the $\ln(R/\Delta R)$ axis of the Lund Plane for b/light jets is also shown. This ratio, as a function of the angular distance ΔR, is sufficient to observe the dead cone effect. At smaller and smaller angles, the b-jet Lund Plane is expected to hollow out due to the effect. No such depletion of emissions is expected for light jets. The dead cone effect is thus visible as a decrease in the ratio, at high values of $R/\Delta R$. The comparison of simulation to data also reveals some mismodelling, especially for collinear splittings in $B^{\pm}$ jets.

We note that the LHCb measurement of $B^{\pm}$-jets is not the first measurement to see the dead cone effect. That honour goes to an ALICE measurement from 2022. This result focused instead on c-jets, and fully reconstructed the D^0 meson to identify c-jets [32].

1.5.3 CMS measurement of groomed jet substructure and fragmentation observables

CMS and ATLAS do not have the particle identification capabilities of ALICE and LHCb. As a result, they are generally not capable of fully reconstructing particle decay chains. We will now describe another b-jet substructure measurement, carried out by the CMS Collaboration, which uses flavour tagging techniques reliant on the identification of features typical of heavy flavour jets. These will be described in detail for ATLAS in Chapter 4.

The CMS result focuses on the measurement of substructure observables such as the SoftDrop

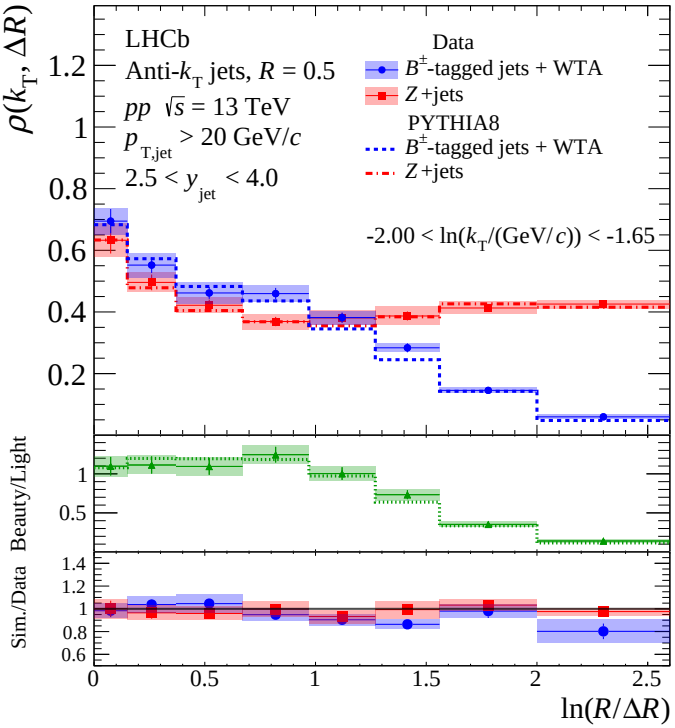


Figure 1.20: The projection onto the $\ln(R/\Delta R)$ Lund Plane for light quark enriched jets and $B^{\pm}$ jets for forward jets and their ratio, along with a comparison to predictions obtained with Pythia8 [59].

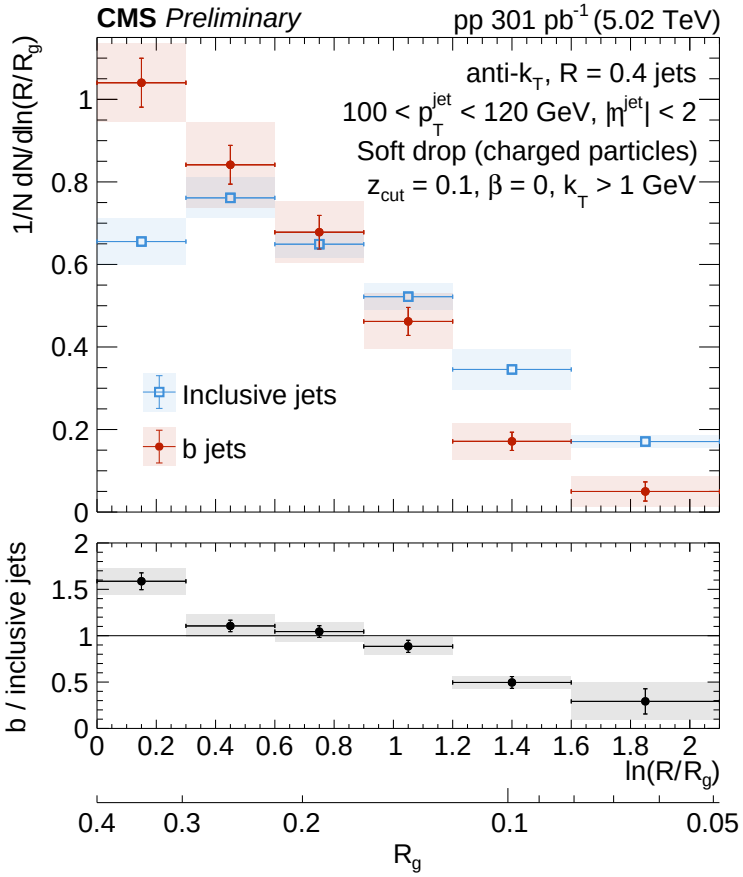


Figure 1.21: The ratio of b-jets with charged b-hadron momentum reconstruction and of inclusive jets in multi-jet events [60].

groomed jet radius R_g, groomed momentum fraction z_g and jet fragmentation function $z_{b,ch}$ [60]. Additionally, the dead cone effect is sought in the ratio of b-jets and an inclusive jet sample. The measurement is carried out on 301 pb^{-1} of data from the Run 2 low pile-up proton-proton run at $\sqrt{s} = 5.02$ TeV. Jets of radius $R = 0.4$ in the central region are identified in multi-jet final states in a p_{T} range between 100-120 GeV. The substructure observables are measured on the charged component of the jets only.

Jets originating from bottom quarks are identified using ParticleNet [70]. The b-hadron in the jets is partially reconstructed from their charged decay products. Tracks originating from the b-hadron decay are identified through a Boosted Decision Tree. The four-momenta of these tracks are summed to reconstruct the charged portion of the b-hadron momentum. A suppression of gluon emissions is found in the ratio, as shown in Figure 1.21. This suppression is consistent with the dead cone effect.

A similar measurement, also carried out by the CMS Collaboration, focuses instead on jets containing D^0 hadrons [71].

1.5.4 Measurements of jet angularities

The other substructure observable featuring in this thesis is the jet angularities. The phenomenology of such observables has been extensively studied at colliders for processes such as Z+jets [72] at the LHC, and they have been measured in proton-proton collisions by the ATLAS, CMS and ALICE Collaborations on a variety of energies and final states including di-jet, $t\bar{t}$, D^0-tagged jets and Z+jets [73]–[77]. The ALICE Collaboration has additionally measured the angularities in jets arising from Pb-Pb collisions [78].

This thesis describes the first ATLAS measurement of the b-jet angularities and of the angularities in quark-enriched jets from Z+jets. As a result, we briefly describe the CMS measurement of the angularities on Z+jets events [75].

Jet angularities from di-jet and Z+jet events in CMS

The measurement in question analyses 35.9 fb^{-1} of data from proton-proton collisions at $\sqrt{s} = 13$ TeV. The generalised angularities $\lambda^1_{0.5}$, λ^1_1, λ^1_2, λ^0_0, and λ^2_0 are measured on groomed and ungroomed jets of radii $R = 0.4$ and $R = 0.8$ with only charged particles and with all constituents as a function of p_{T}. Final states identified as Z+jets provide a quark-enriched sample of jets, while di-jet events are used to extract sample enriched with gluon-initiated jets. Jets are required to be found in the central region of the detector and to have $p_{\mathrm{T}} > 30$ GeV. The two leading jets j_1 and j_2 in di-jet events are required to be balanced, satisfying $(p_{\mathrm{T}}^{j_1} - p_{\mathrm{T}}^{j_2})/(p_{\mathrm{T}}^{j_1} - p_{\mathrm{T}}^{j_2}) < 0.3$. The same asymmetry requirement is imposed on j_1 and the Z candidate in the Z+jets region, along with an isolation criterion for j_1. For jets in di-jet events with $p_{\mathrm{T}} < 1000$ Gev, the most central (lowest $|y|$) di-jet is used for the gluon-enriched sample. This jet is referred to the central di-jet, while the other is referred to as the forward di-jet.

Results are compared to simulations obtained at LO with MadGraph5_aMC@NLO v2.2.2 showered with Pythia8 (MG5+Pythia8) or with Herwig++ [79], [80]. Figure 1.22 shows the results for λ^1_1 in the Z+jets and central di-jet regions for jets with p_{T} in the range 120-150 GeV. The agreement to data in noticeably worse in the central di-jet region for both sets of predictions. In the Z+jets region, the MG5+Pythia8 predictions provide the best description of the data.

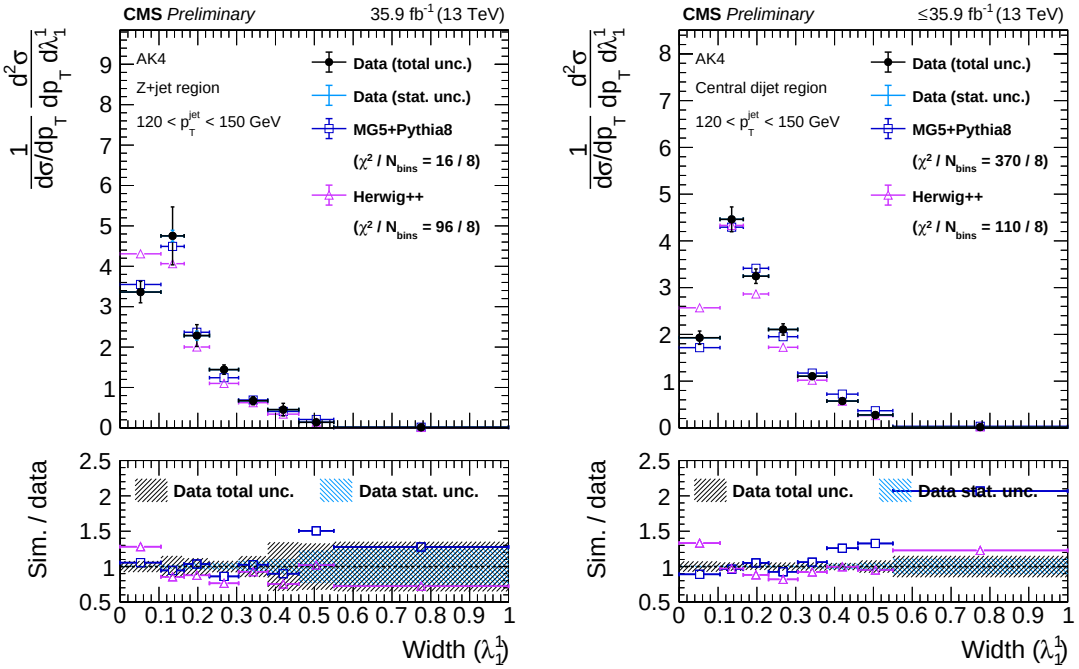


Figure 1.22: The λ_1^1 angularity measured for ungroomed jets with all constituents in the p_{T} range 120-150 GeV in Z+jet events (left) and in central di-jets (right) by CMS [75].

Chapter 2

Jet flavour labelling

In theoretical predictions, the classification of heavy flavour jets, or *flavour labelling*, is critical under multiple facets. First, it is a crucial component of fixed-order calculations involving heavy flavour production. At NNLO and beyond, defining a jet's flavour is highly non-trivial, as effects such as gluon splittings can "contaminate" a jet with heavy flavours which do not stem directly from the matrix element. This can be in the form of a soft gluon splitting into a heavy flavour pair within a jet, a wide-angle gluon splitting leading to one of the quarks entering a jet, or a splitting of the form $b \to gb$, which reduces the b-quark (and subsequent hadron) p_{T} below the threshold for consideration. Similar effects can also arise in simulation, such as for the parton shower evolution.

Were we to assign flavour naïvely, based solely on the presence of heavy flavour quarks (at parton level) or heavy flavour hadrons within a given jet, such effects would lead to flavour mislabelling. For precision calculations and comparison to data, the ability to identify spurious flavour arising from gluon splittings is paramount.

Accurate flavour labelling is also an important component of the MC simulation used in experiments for the development of algorithms capable of recognising the flavour of detected jets, known as *flavour tagging algorithms*. To develop a flavour tagging algorithm, one needs to identify at particle-level jets containing heavy flavours to then, at reconstruction-level, be able to spot within the detector simulation those features which allow for the identification of these jets in data. Here, there is some merit in an inclusive definition of flavour, allowing for these contaminations: if one is solely interested in fragmentation studies, it is irrelevant to consider whether the fragmenting heavy flavour quark arises from the matrix element or parton shower. In addition, it is impossible to distinguish these two cases experimentally. However, mislabelling of jets when developing flavour tagging algorithms can lead to bias since, for example, wide-angle $b\bar{b}$ pairs can lead to the contamination of c-jets, considering the fact that b-flavour typically takes priority over c. More details will follow in Section 2.1.

In recent years, a number of novel IRC-safe jet clustering algorithms which take flavour into account have been developed to solve this inconsistency. Collectively, these are known as *flavour jet algorithms*. In this chapter, we describe these algorithms and evaluate their impact on several theoretical and experimental aspects of flavour labelling. We begin in Section 2.1 with a description of the flavour labelling strategies currently in use at experiments such as ATLAS and CMS, before moving on to describe the novel flavour jet algorithms in Section 2.2. A detailed comparison of the performance of these algorithms, principally on Z+jets final states, is then carried out in Section 2.3.

This chapter is based off of work initiated at the 2023 Les Houches Workshop [81] which I

took part in and results from a follow-up study I published in [82].

2.1 Current flavour labelling strategies

Historically, three different strategies for flavour labelling have been used. These are as follows:

1. **Anti-k_t labelling:** A jet is labelled as a heavy flavour jet if a particle containing bottom or charm is found within it, as determined by the jet clustering algorithm (typically anti-k_t). This particle may be the quark itself, if the labelling is done at parton level, or else a hadron whose decay has been suppressed [83].

2. **Geometric matching:** This is the strategy used to label jets within ATLAS. It is applied exclusively at hadron level. A jet is labelled as b if a b-hadron with $p_{\mathrm{T}} > 5$ GeV is found within a radius of $R < 0.3$ of the jet's axis. The hadron must be one of the initial hadrons formed after hadronisation. The jet is otherwise labelled as a c-jet if no b-hadron is found, but a c-hadron is found satisfying the same requirements [84].

3. **Ghost matching or association:** This strategy is again applied exclusively at hadron level. The four-momentum of the heavy flavour hadrons found within an event are suppressed to near-zero. The jet clustering algorithm is then applied. If the ghost hadron is found as a constituent of the reclustered jet, the jet is labelled as the flavour corresponding to the hadron species. B flavour takes precedence over c flavour [85].

As mentioned, however, these flavour labelling strategies are not robust to spurious flavour contributions arising from gluon splittings at NNLO and beyond. For instance, a wide-angle gluon splitting, as is shown in Figure 2.1 leads to the misidentification of a light jet from the ME as flavoured. These methods of labelling also jeopardise the IRC safety of the jet definition, as the number of flavoured particles in a jet, which determine its flavour label, is not an IRC safe observable.

To resolve this issue, a number of flavour jet clustering algorithms have been proposed. The first such algorithm, known as *Flavour-k_t*, was proposed in 2006 by Banfi, Salam and Zanderighi [83], [87]. This algorithm was never widely adopted as it clusters jets with k_t-like kinematics, and thus more often leads to soft particles clustering together rather than with nearby hard particles[1] [23].

More recent proposals have succeeded in maintaining jets with anti-k_t like kinematics. In general, these work by identifying flavour pairs which originate from gluon splittings across the entire event. This is evaluated based on a criterion which, if satisfied, leads to the cancellation of flavour in the registry of the algorithm. The resulting flavour from these cancellations is what is assigned to the jet, which is clustered with the anti-k_t algorithm.

The method of combining sources of flavour content within a jet is known as the *flavour recombination scheme*. Here, flavour content typically refers to flavoured particles, which can be either heavy flavour partons or hadrons. A few different flavour recombination schemes are in use. These include: *any flavour*, *net flavour*, and *mod-2 flavour*. The any flavour recombination scheme, used in the strategies described above, defines a jet as flavoured as long as one flavoured particle is associated to it. The other two schemes combine flavour particles as the jet clustering takes place, taking either the net sum of the flavour content, defined as 1 for heavy flavour

[1] Equivalently, it can be said that k_t-like kinematics more often leads to irregularly shaped, rather than conical, jets.

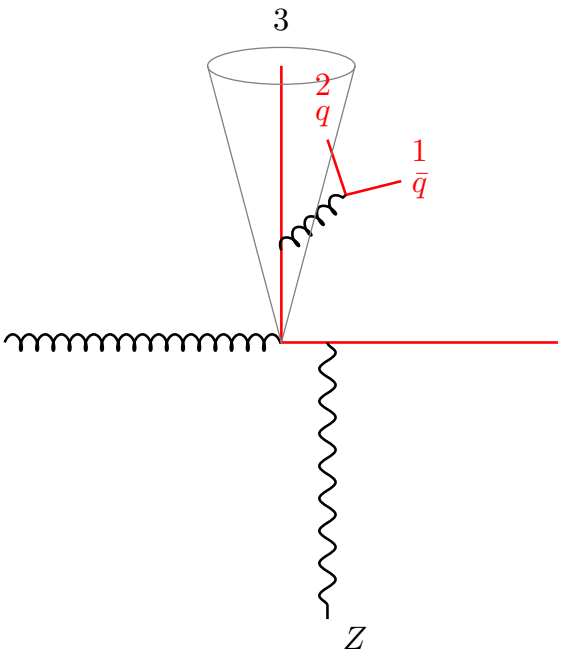


Figure 2.1: A wide-angle gluon splitting to heavy flavours (1,2) in a Z+jets event which contaminates the flavour of jet 3 [86].

particles and -1 for heavy flavour anti-particles, or the sum modulo 2. The recombination schemes are outlined in Table 2.1.

The any flavour recombination scheme suffers from the IRC-safety issues associated with counting flavoured particles described above. The flavour assigned based on the net flavour and mod-2 recombination schemes, on the other hand, does not change due to a $g \to b\bar{b}$ splitting, curing the IRC safety malady. On the other hand, the any flavour and mod-2 flavour recombination schemes have the benefit of being the most realistic, i.e. they are the ones most which can be most readily implemented in experiment. The mod-2 flavour recombination scheme has the added benefit of being robust against $B - \bar{B}$ oscillations when applied at hadron level.

Scheme	**Consider a set of particles flavoured if ...**
any flavour	$\sum_i \lvert f_i \rvert > 0$
net flavour	$\sum_i f_i \neq 0$
mod-2 flavour	$\sum_i \lvert f_i \rvert \equiv 1 \mod 2$

Table 2.1: The different flavour recombination schemes used to assign flavour to a jet. Flavoured particles are assigned a value of $f_i = +1$, while flavoured anti-particles are assigned a value of $f_i = -1$. For unflavoured particles, $f_i = 0$. The sum runs over all particles i under consideration whose flavour should be combined [82].

2.2 The next generation of algorithms

We are now ready to describe the new generation of flavour jet algorithms in more detail. There are four such algorithms. They can all make use of either the net flavour or mod-2 flavour recombination schemes to maintain IRC safety. Throughout this chapter, the net flavour recombination scheme has been used unless otherwise stated. When jets at hadron level are discussed, it is understood that the flavour jet algorithms require undecayed heavy hadrons in order to meaningfully ascribe flavour to jets.

2.2.1 SoftDrop flavour (SDF)

The SoftDrop flavour algorithm is based on the SoftDrop grooming algorithm. It uses the SoftDrop criterion (Eq. 1.74) to determine whether a flavour pair originates from a gluon splitting [88].

Starting from a jet clustered with an algorithm of choice, the SDF algorithm proceeds as follows:

1. The jet's constituents are reclustered using the JADE algorithm, following the distance metric given in Eq. 1.68.

2. At each step in the clustering procedure, the particles being clustered together are required to satisfy the SoftDrop criterion.

3. If the grooming step passes, sum the flavour content in the jet according to a recombination scheme and proceed with clustering algorithm..

4. If the grooming step fails, groom the soft branch and proceed with the clustering algorithm.

The choice of the JADE algorithm for the reclustering step is fundamental due to certain orderings of partons in generalised k_t jets at NNLO whose flavour cannot be made IRC safe with SoftDrop grooming. The JADE algorithm clusters soft particles together first, and they can then be groomed away with SoftDrop. The JADE reclustering step, however, necessarily entails a modification of anti-k_t kinematics, assuming the initial jet considered was clustered with that algorithm. The modification is modest and often negligible in practice.

To ensure IRC safety, the SoftDrop criterion must be applied with parameters $\beta > 0$ and $z_{\text{cut}} < \frac{1}{2}$. The values of $z_{\text{cut}} = 0.1$ and $\beta = 1$ or $\beta = 2$ are typically chosen.

2.2.2 Flavoured anti-k_t (CMP)

The flavoured anti-k_t algorithm, or CMP [89][2] from the authors' names (**C**zakon, **M**itov, **P**oncelet), modifies the distance measure used in anti-k_t (Eq. 1.70) in order to capture soft, flavoured particles. Specifically, the distance measure d_{ij} is modified through a damping function S_{ij}:

$$d_{ij}^{(\text{flavoured})} = d_{ij}^{(\text{standard})} \times \begin{cases} S_{ij}, & \text{if } i,j \text{ have nonzero flavour of opposite sign and magnitude,} \\ 1, & \text{otherwise.} \end{cases} \tag{2.1}$$

[2]Technically, the variant of the algorithm defined with the parameters listed in Eq. 2.2 is referred to as CMPΩ. We refer to it as the CMP algorithm as we make no distinction between the variants.

The damping function is applied exclusively to flavoured particles and is required to vanish for soft quark pairs, compared to the scale of the hard process. The function chosen which satisfies these properties is:

$$\begin{aligned} S_{ij} &= 1 - \theta\left(1 - \kappa_{ij}\right) \cos\left(\frac{\pi}{2}\kappa_{ij}\right) \\ \text{with} \quad \kappa_{ij} &\equiv \frac{1}{a}\sqrt{\frac{\Omega_{ij}^2}{\Delta R^2}}\frac{p_{t,i}^2 + p_{t,j}^2}{2p_{t,\max}^2} \\ \text{and} \quad \Omega_{ij}^2 &= 2\left[\frac{1}{\omega^2}\left(\cosh(\omega\Delta y_{ij}) - 1\right) - \left(\cos\Delta\phi_{ij} - 1\right)\right]. \end{aligned} \tag{2.2}$$

Here, θ is the Heaviside function, $p_{\mathrm{T,max}}$ is some hard scale, a is a parameter used for tuning, and ω is a parameter which regulates the distance measure in such a way as to ensure IRC safety. The values of $a = 0.1$, $\omega = 2$ are adopted throughout and $p_{\mathrm{T,max}}$ is dynamically set to the p_{T} of the hardest subjet at each step in the clustering procedure.

With such a damping function, the CMP algorithm reduces to the anti-k_t algorithm in the absence of soft flavour pairs, and modifies the kinematics slightly when they are present. We can thus speak of *approximate* anti-k_t kinematics, though the modifications are often negligible.

2.2.3 Flavour dressing (GHS)

As the name implies, the flavour dressing (or GHS, from **G**auld **H**uss **S**tagnitto) algorithm [90] aims to decorate a set of jets with flavour. The algorithm starts with a set of flavour-agnostic jets clustered with an algorithm of choice and all particles which make up a given event. If the particle p_i is a constituent of jet j_k, the distance $d(p_i, j_k)$ is saved. The distance between particles p_i and p_j $d(p_i, p_j)$ is saved as well if both particles are flavoured or if one is flavoured and they are both associated to the same jet. The distances between the objects are then evaluated, and the following considerations are made:

- If $d(p_i, p_j) < d(p_i, j_k)$, the particle $k_{ij} = p_i + p_j$ is formed by summing the four-momenta and flavour content of the two particles. In the record of saved distances, p_i and p_j are replaced by k_{ij}.
- If instead $d(p_i, p_j) > d(p_i, j_k)$, the particle p_i is assigned to jet j_k and all other records containing p_i are discarded.
- If the beam distance $d(p_i, B)$ is smaller than both $d(p_i, p_j)$ and $d(p_i, j_k)$, all saved distances containing p_i are discarded.

The distances considered are:

$$d(p_i, p_k) = \max(p_{\mathrm{T}i}, p_{\mathrm{T}k})^{\alpha} \min(p_{\mathrm{T}i}, p_{\mathrm{T}k})^{2-\alpha}\Omega_{ik}^2 \tag{2.3}$$

$$d(p_i, B_{\pm}) = \max\left(p_{\mathrm{T}_i}^{\alpha}, p_{\mathrm{T}_{B_{\pm}}}^{\alpha}(y_i)\right) \min\left(p_{\mathrm{T}_i}^{2-\alpha}, p_{\mathrm{T}_{B_{\pm}}}^{2-\alpha}(y_i)\right) \tag{2.4}$$

where $B_{\pm}$ takes into account the direction of positive/negative rapidity and

$$p_{\mathrm{T}B_{\pm}} = \max\left(p_{\mathrm{T}i}^{\alpha}, p_{\mathrm{T}B_{\pm}}^{\alpha}(y_i)\right) \min\left(p_{\mathrm{T}i}^{2-\alpha}, p_{\mathrm{T}B_{\pm}}^{2-\alpha}(y_i)\right), \tag{2.5}$$

with α a parameter set to 1, $\Theta(0) = 1/2$, and $\Delta y_{j_k} = y_{j_k} - y$. Ω_{ik}^2 takes the same form as in Eq. 2.2. The parameter ω is set to 2 as default. To ensure IRC safety, the parameters must be chosen such that $\omega > 2 - \alpha$.

This procedure is iterated until the set of saved distances is empty. The flavour of a jet is then the sum of that of the particles belonging to it. The flavoured jets obtained in this way maintain the initial kinematics of the flavour-agnostic algorithm, such as anti-k_t. This algorithm has the added benefit of potentially being applicable experimentally, as the flavoured particles can be proxies such as reconstructed secondary vertices.

2.2.4 Interleaved flavour neutralisation (IFN)

The interleaved flavour neutralisation algorithm [86] can be applied as part of any jet clustering algorithm. When clustering two particles i and j, with i flavoured and $p_{\mathrm{T}i} < p_{\mathrm{T}j}$, the algorithm searches for flavour to cancel throughout the entire event. Given a pseudojet i, the distance measure described in Eq. 2.3 and a threshold $u_{\max}$, and a set C off flavoured particles potentially stemming from soft gluon splittings together with i (neutralisation candidates), the algorithm proceeds as follows:

1. A set L of distances u_{ik} for all elements $k \in C$ such that $u_{ik} < u_{\max}$ is formed.
2. The element k corresponding to $\inf(L)$ is identified.
3. If k contains no flavour that can be neutralised with i, remove u_{ik} from L and return to step 2.
4. If k is found to have flavour which can be neutralised with i, check first to see if there are no other particles whose flavour can be neutralised with k. Apply this algorithm recursively considering k in the place of i.
5. If k and i are still found to be compatible candidates, neutralise as much flavour in i as possible with that of k.
6. If i is flavourless, exit. If not, remove the element u_{ik} from L and return to step 2.
7. If the list of candidates is exhausted, combine the particles i and j. Any remaining flavour in i is to be summed in accordance with some flavour recombination scheme.

The IFN algorithm is only called if i is the softer particle being combined and i is flavoured. If j is the flavoured particle, the pseudojet p_{ij} simply takes the flavour of j. If this is not the case, the list L of neutralisation candidates should include all pseudojets with flavour, including any flavoured jets already identified by the algorithm.

The result is a set of jets with the exact kinematics of the algorithm used to cluster them, usually anti-k_t. The algorithm depends on same two parameters as the GHS algorithm, (α, ω), and the same choice of $(1, 2)$ is possible, though the default choice is $(2, 1)$. The clustering procedure with flavour neutralisation is illustrated in Figure 2.2.

2.3 Flavoured jet algorithms: A comparative study

At the 2023 Les Houches Workshop, a detailed study of the differences between the flavour jet algorithms was initiated, from various points of views. Les Houches has a tradition of bringing together the theory and experimental communities to devise new ideas with which to tackle the challenges facing the Standard Model. Flavour labelling is one such topic which could benefit from the collaboration between these communities. This section is dedicated to selected findings from those studies, published in [82]. I was actively involved in discussions which led up to this work since their inception, and more directly to the results published in Sections 2.3.5, 2.3.6, and 2.3.7.

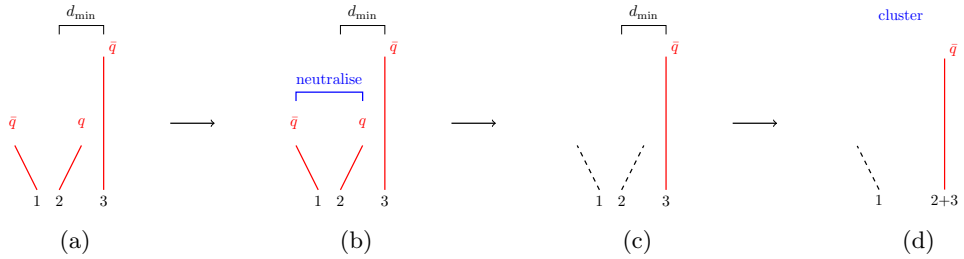


Figure 2.2: The steps in flavour neutralisation in the IFN algorithm. In Figure 2.2a, particles 1 and 2 are a heavy flavour pair arising from a soft gluon splitting, and particles 2 and 3 are meant to be clustered. In Figure 2.2b, it is found that the particles 1 and 2 satisfy the distance 2.3, and so their flavour is neutralised as shown in Figure 2.2c. The clustering can thus proceed as in Figure 2.2d, combining particles 2 and 3 with no flavour [86].

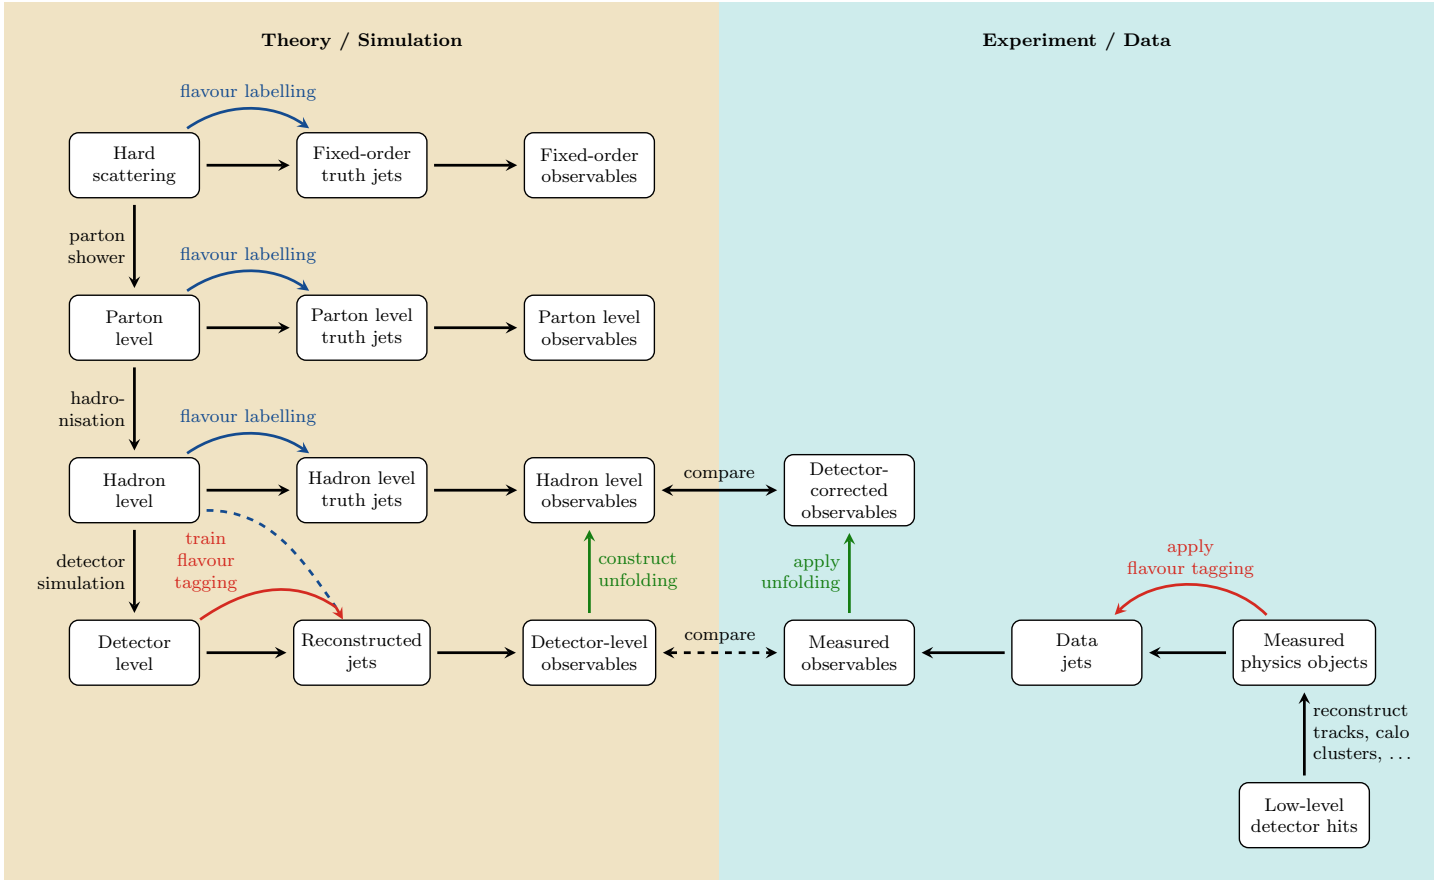


Figure 2.3: A simplified overview of a generic jet-based analysis, illustrating both the main steps for calculating theoretical predictions and for analysing experimental data. The coloured arrows show the propagation of information about flavour [82].

2.3.1 Bridging the gap from theory to experiment

To better understand the close link between the experimental and theoretical definitions of jets, it is useful to describe the typical way in which jet predictions and measurements are compared between theory and experiment. The simplified picture of a generic jet analysis is illustrated in the flowchart given in Figure 2.3. This chart displays how information is translated from first-principles theory and from low-level information recorded by the particle detectors through different stages of simulation and experimental analysis to quantities that can be compared directly. The left-hand side of Figure 2.3 sketches how theory predictions are obtained, starting from the description of events at the level of hard scattering. Events are successively translated to parton level by applying a parton shower, then to the hadron level through hadronisation and finally to the detector-level by feeding the hadronic final state through a simulation of an experimental detector's response. In principle, we can analyse the events at any of these stages, construct jets by applying a jet algorithm to the events, and calculate observables such as fiducial cross sections or distributions from these objects. Jets constructed from simulated events at parton or hadron levels are called *particle-level* or *truth* jets. With them, we can calculate fixed-order, parton level, or hadron level observables, respectively. Jets obtained from simulated detector-level events are called *detector-level* or *reconstructed jets*, and we use them to calculate detector-level observables.

On the right-hand side of Figure 2.3, low-level detector objects (tracker hits, energy deposits in calorimeter cells, etc.) are used to reconstruct higher-level physics objects (tracks, topological cell clusters, displaced vertices, ParticleFlow objects, etc.)[3]. We call jets constructed from these physics objects *data jets*, and observables calculated from them *measured observables*. There is therefore a level of parity between reconstructed jets and data jets, as well as detector-level observables and measured observables, which are at the same level but reconstructed using either simulation or experimental data, respectively.

The predictions for detector-level observables can be compared to measured observables (dashed black arrow in the flowchart), and this is done both in the context of measurements of Standard Model processes and in the context of searches for new physics. Since these quantities are detector specific, unfolding is often used to translate detector-specific measurements to hadron level (green arrows in the flowchart), making them more universal. Unfolding requires constructing a response matrix by comparing predictions for hadron level observables to detector-level observables derived from the same simulated events. This unfolding procedure is then applied to the measured observables to obtain detector-corrected observables. These unfolded results can finally be compared to predictions for hadron level observables, as well as to results from other experiments

2.3.2 Flavour recombination scheme

Before discussing the IRC safe flavoured jet algorithms, we wish to examine the impact of the choice of flavour recombination scheme. To ensure IRC safety from higher-order effects in QCD, either the mod-2 or net flavour recombination schemes must be employed, rather than the any flavour scheme typically utilised in experimental analyses.

To understand the effect of modifying the choice of flavour recombination scheme, a study on hadron level simulated data was carried out. Here, $pp \to Z(\mu^+\mu^-)$+jet matrix elements are calculated at NLO and showered using Herwig7 [91] and Sherpa [92] at $\sqrt{s} = 13$ TeV. Two different Herwig7 showers are carried out, one using the Lund string hadronisation model, and the other using the cluster hadronisation model. Jets are clustered suppressing the decays of the

[3]This reconstruction will be described in detail in Chapter 3.

heavy hadrons and with a radius $R = 0.5$. They are required to have $p_{\mathrm{T}} > 30$ GeV and rapidity $|y| < 2.4$. The transverse momentum of the individual muons (p_{T}^{μ}) and the di-muon pair ($p_{\mathrm{T}}^{\mu\mu}$) must both be greater than 20 GeV. The mass of the di-muon pair must be compatible with that of the Z boson, $77 < m^{\mu\mu}/\mathrm{GeV} < 111$. Individual muons are also required to be in the central region $|y^{\mu}| < 2.4$. b-jets and c-jets are identified using ghost and cone labelling with the any flavour recombination scheme, and with anti-k_t labelling using the mod-2 flavour recombination scheme.

Figure 2.4 shows the leading b/c-jet distributions for these predictions with anti-k_t, ghost, and cone labelling. For anti-k_t labelling, the mod-2 flavour recombination scheme is used, while the any flavour recombination scheme is used in the other cases. The decay of heavy hadrons is inhibited in all cases.

We note how, in these circumstances, the ghost labelling algorithm is equivalent to the anti-k_t algorithm with the any flavour recombination scheme. This allows us to compare the schemes in use. For both the Sherpa and Herwig7 showers, we find a reduction in the number of jets labelled as b/c due to the elimination of heavy flavour pairs stemming from soft gluon splittings. This reduction increases with jet p_{T}, since the number of soft gluon splitting increases as the energy scale increases. The effect is also more pronounced for c-jets, as the number of $g \to c\bar{c}$ splittings is expected to be greater than that of $g \to b\bar{b}$ splittings due to the masses of the particles, and for jets labelled with ghost labelling. In this case, the difference compared to cone labelling is attributed to the labelling itself: in cone labelling, a radius of $\Delta R < 0.3$ between the hadron and jet is considered, whereas for ghost labelling the radius is effectively equal to that of the jet radius.

The labelling is effectively identical for both hadronisation models considered for the Herwig7 showers. When comparing parton/hadron level, some differences arise between jets labelled with ghost/cone and anti-k_t. These differences are more marked for c-jets.

2.3.3 Neutralisation of flavour from gluon splittings

We turn our attention now to the differences between the various flavour jet algorithms. Figure 2.5 shows LO+PS Z+jet events produced with selections similar to those in Section 2.3.2 using Pythia8 v8.306 [93] without hadronisation. It compares the number of b-jets found using the flavour jet algorithms with certain configurations of parameters to those found with anti-k_t labelling. It is clear that the number of b-jets labelled is lowest for the IFN algorithm with $\alpha = 2$. This effect is attributed to the neutralisation of the most heavy flavour pairs coming from soft gluon splittings. To prove this, Figure 2.6 shows a series of scatter plots of a subset of the generated Z+jet events with at least two b-quarks at parton level and leading b-jet $p_{\mathrm{T}} > 200$ GeV, as decided by anti-k_t labelling. On the x-axis, the angular distance between the b-quarks is shown, while on the y-axis the p_{T} fraction of the jet carried by the leading b-quark in the jet $p_{\mathrm{T},b}/p_{\mathrm{T},b\text{-jet}}$ is reported. Events where the flavour content originates from wide-angle, soft gluon splittings populate mostly the bottom-left corner. The exact truth origin of all events in the plot is shown in Figure 2.6a, showing true $Z + b$ events, $Z + q$ events containing mostly light quarks where b-jets tend to originate from gluon splittings, and $Z + g$ events, where b-jets arise from both soft and hard gluon splittings. The other figures show the events labelled as b-jets by anti-k_t and each individual flavour jet algorithm in grey, or just by anti-k_t in red. It is clear that IFN, as declared, is better at identifying genuine b-jets than the other algorithms.

The results shown in Figure 2.5 are valid also at NLO, as shown in the top panels of Figure 2.7 for b/c-jets. In this case, Sherpa and Herwig7 are used as parton showers along with the mod-2 flavour recombination scheme. The figure also shows how the results remain valid at hadron level independently of the hadronisation model chosen. When comparing the results for the

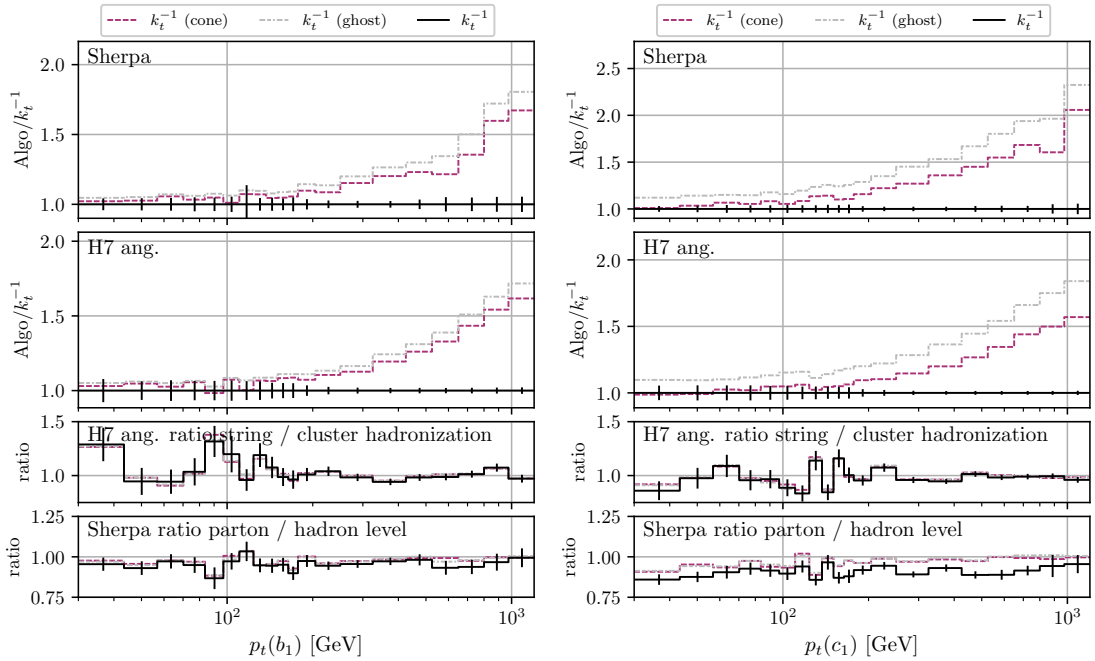


Figure 2.4: The leading jet p_{T} spectrum from NLO+PS predictions obtained with Sherpa and Herwig7 for $pp \to Z + b$ (left) and $pp \to Z + c$ (right) in central kinematics. The spectrum is compared for jets clustered with the anti-k_t algorithm and whose flavour is labelled with cone, ghost, and anti-k_t labelling. The any flavour recombination scheme is used for the ghost and cone algorithms, while the mod-2 recombination scheme is considered for the labels found using the anti-k_t algorithm. Additionally, differences between the cluster and string hadronisation model for Herwig7, as well as parton and hadron level for Sherpa, are shown [82].

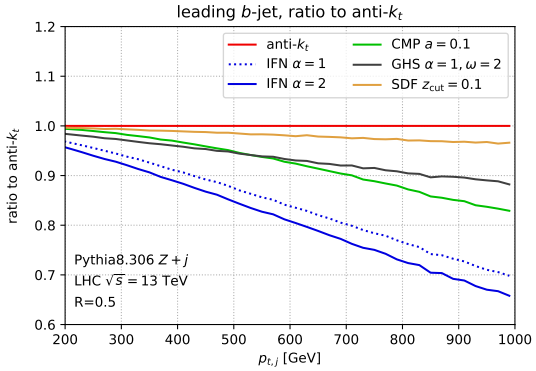
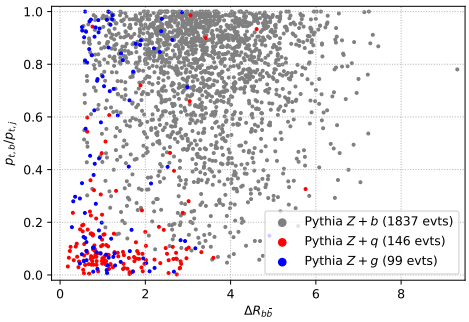
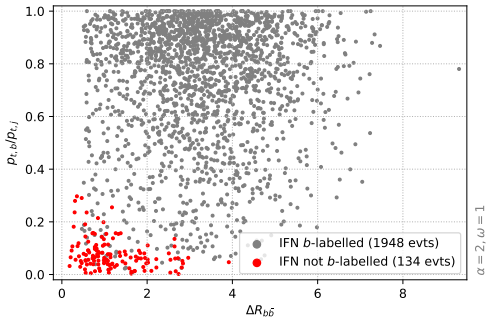


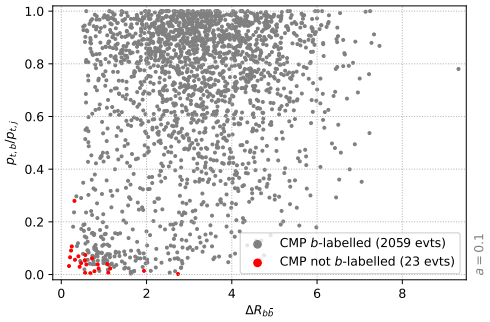
Figure 2.5: The ratio to net-flavour anti-k_t (in red) of IFN ($\alpha = 1$ and $\alpha = 2$, blue), CMP ($a = 0.1$, green), GHS ($\alpha = 1$, $\omega = 2$, black) and SDF ($\beta = 2$, $z_{\mathrm{cut}} = 0.1$, yellow) for the transverse momentum of the leading b-labelled jet $p_{\mathrm{T},b\text{-jet}}$ [82].



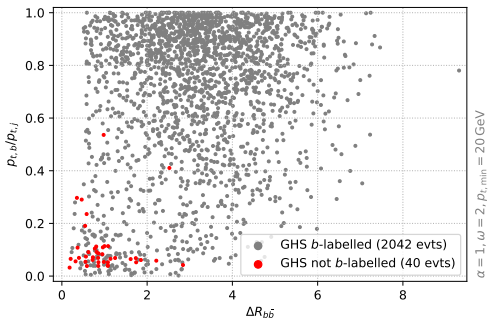
(a) Truth



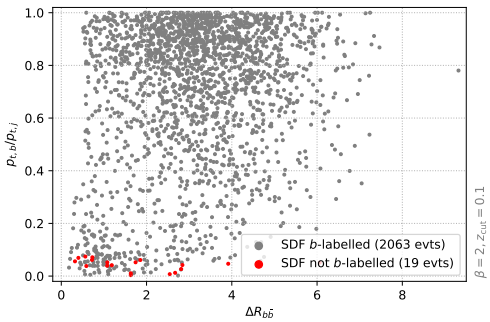
(b) IFN



(c) CMP



(d) GHS



(e) SDFlav

Figure 2.6: Z+jets events with at least 2 b-quarks and a b-jet with $p_{\mathrm{T}} > 200$ GeV as identified with anti-k_t as a function of $\Delta R_{b\bar{b}}$ and $p_{\mathrm{T},b}/p_{\mathrm{T},b\text{-jet}}$. In Figure 2.6a the truth origin of these events is shown. In Figures 2.6b–2.6e, the label assigned by the selected jet algorithm vs. anti-k_t is shown, with grey events labelled as b by both net flavour anti-k_t and the algorithm under consideration and red events labelled as b by net flavour anti-k_t, but not by the flavoured algorithm [82].

Sherpa sample at hadron and parton level, differences in the trends observed for the flavour jet algorithms arise for high p_{T} c-jets.

2.3.4 NNLO fixed-order predictions

The flavoured jet algorithms are studied in the context of fixed-order predictions at NNLO. The process studied is $pp \to W^+(\to e^+\nu_e)H(\to b\bar{b})$.

Events must contain two b-labelled jets of radius $R = 0.4$ with $p_{\mathrm{T}} > 25$ GeV and $|\eta| < 2.5$. The Higgs boson is reconstructed by selecting the pair of b-labelled jets with invariant mass closest to the mass of the Higgs boson, 125 GeV. The positron is required to have $p_{\mathrm{T}} > 15$ GeV and $|\eta| < 2.5$. Predictions are obtained using the 5FS.

Figure 2.8 shows the leading b-jet p_{T} spectrum (p_{T,b_1}) and the p_{T} of the Higgs boson ($p_{\mathrm{T},H(b\bar{b})}$) resulting from the calculation compared to the four flavoured jet algorithms described in Section 2.2 and flavour-k_t. The kinematic differences between this last algorithm and the others are apparent. Consistently with the results presented in Section 2.3.3, the IFN algorithm is still found to identify fewer flavoured jets than the other anti-k_t-like algorithms.

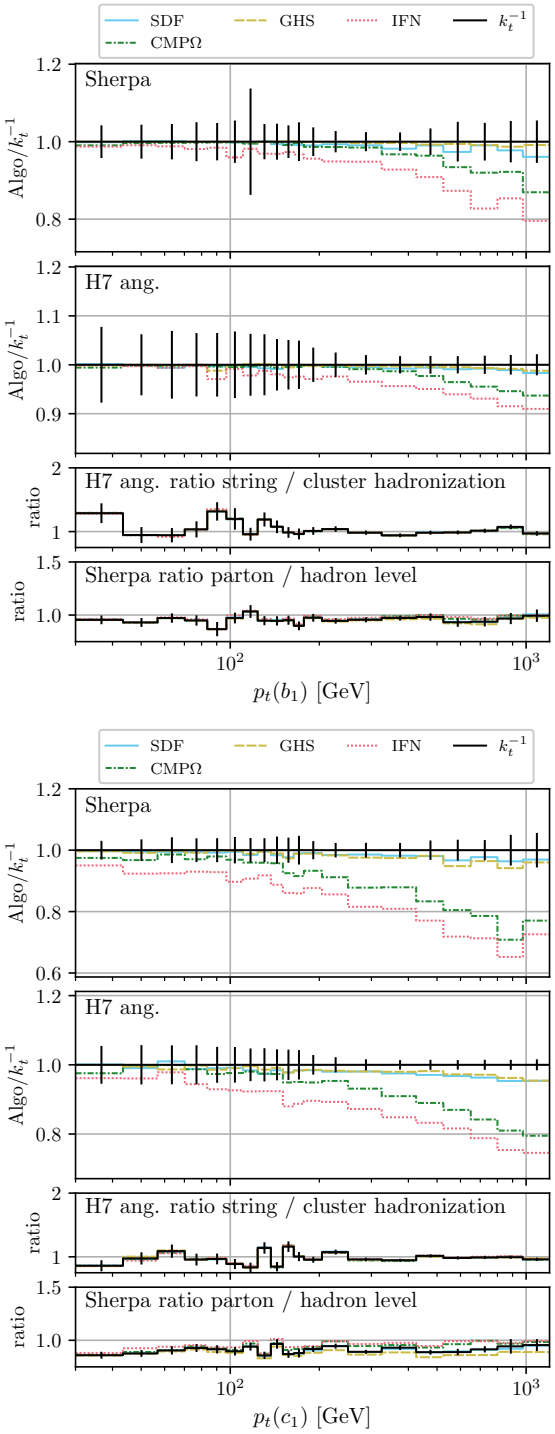


Figure 2.7: NLO+PS predictions from Sherpa and Herwig7 for $pp \to Z + b/c$ (top/bottom) in central kinematics. Different jet algorithms (CMP (green), SDF (light blue), GHS (yellow), IFN (red) are compared for the transverse momentum of the leading b/c-jet. Additionally, differences between the cluster and string hadronisation model for Herwig7, as well as parton and hadron level for Sherpa, are shown. The vertical bars indicate statistical uncertainties [82].

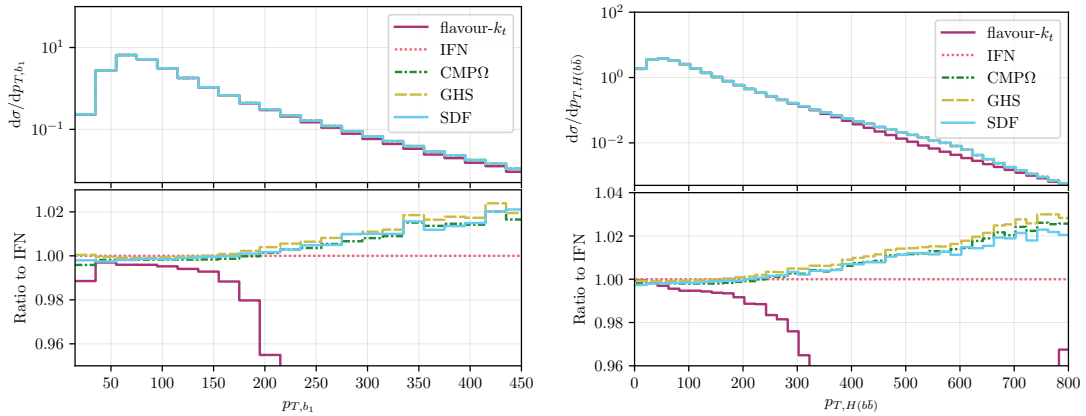


Figure 2.8: The p_{T,b_1} (left) and $p_{\mathrm{T},H(b\bar{b})}$ (right) spectra for the flavour-k_t and the flavour jet algorithms resulting from an NNLO fixed-order calculation of $pp \to W^+(\to e^+\nu_e)H(\to b\bar{b})$ [82].

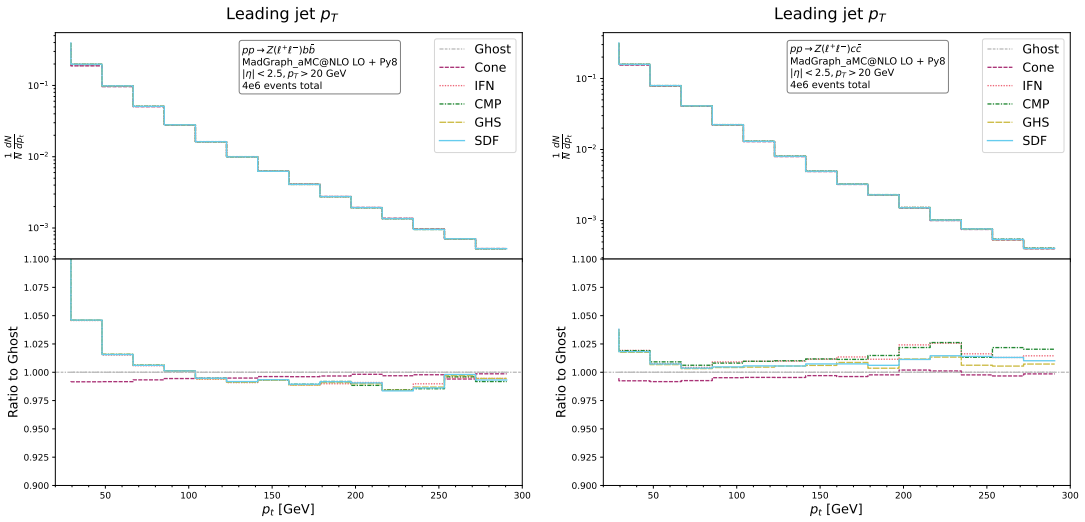


Figure 2.9: The p_{T} distribution of leading b-jets (left) and leading c-jets (right) identified with the experimental ghost and cone labelling schemes with stable final state particles and with the flavour jet algorithms with undecayed heavy hadrons. Ratios are compared to the distribution for jets obtained with ghost labelling [82].

2.3.5 Impact of hadron reconstruction on flavoured jet definitions

We consider now the effect of the constituents used in the jet definition on flavour labelling. In experimental collaborations, such as ATLAS and CMS, flavour labelling is done on jets whose fiducial definitions match the experimental reality as closely as possible using cone or ghost labelling. This means that the final states considered when clustering jets consist of all stable, visible particles in the event, and flavour is assigned to these jets. This is in contrast to the jet definitions used by theorists where, in many instances, the jets are not clustered in the most realistic manner, for instance when suppressing hadron decay for anti-k_t labelling.

This difference has critical implications for jet kinematics. Consider hadron level $Z + b\bar{b}$ and $Z + c\bar{c}$ events generated at LO with MadGraph_aMCNLO v3.5.8 [44] and showered with Pythia8 v8.313. Jets are clustered following the prescription used by the ATLAS AntiKt04TruthWZJet collections, namely with all stable, visible particles within $|\eta| < 2.5$ and with radius $R = 0.4$, after having removed the Z boson decay products. An exception is made for jets clustered with the flavour jet algorithms. In this case, heavy hadron decays are artificially suppressed as required by the algorithms. The undecayed hadrons are treated as stable. Jets are required to fall within the central region, i.e $|y| < 2.5$, and have $p_{\mathrm{T}} > 20$ GeV. $Z \to l^+l^-$ events are considered, where the individual leptons are required to have $p_{\mathrm{T}}^{\ell} > 27$ GeV, $|\eta_{\ell}| < 2.5$ and di-lepton mass $76 < m_{ll}/\mathrm{GeV} < 106$. Events with at least one b/c-jet are selected. The two leading b/c-jets are selected when multiple labelled jets are found.

As can be seen in Figure 2.9, there are significant kinematic differences between the jets labelled as b/c with ghost and cone labelling and with the flavour jet algorithms. These differences are accentuated at low p_{T}.

This is attributed to the difference in constituents. Figure 2.10 shows how, if heavy hadron decays are inhibited in all labelling schemes, the kinematic differences disappear and the expected behaviour, i.e. the cancellation of heavy flavour pairs by the flavour jet algorithms, is observed throughout the entire p_{T} spectrum.

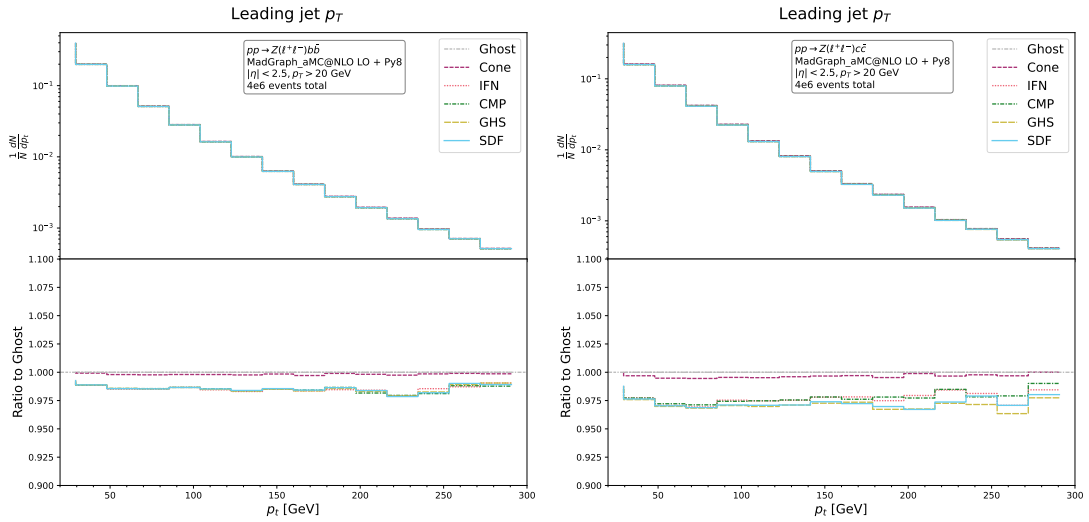


Figure 2.10: The p_{T} distribution of leading b-jets (left) and leading c-jet (right) identified with the ghost cone labelling schemes and the flavoured jet algorithms, all with undecayed heavy hadrons. Ratios are compared to the distribution for jets obtained with ghost labelling [82].

These results highlight the importance of harmonisation. Ideally, the new flavour algorithms would be adopted by experimental collaborations, at least in the context of precision physics in order to more accurately predict relevant cross sections for the production of heavy flavour particles, correcting for effects such as gluon splittings. If this is to be done, the collaborations must either revisit their truth-level jet definitions in order to take into account kinematic effects arising from the difference in constituents, or implement an additional unfolding-like step to correct for the differences that arise when using these novel labelling schemes.

2.3.6 Consequences for flavour tagging

Other practical implications of the inadequacy of the current flavour labelling strategies in use in experiments can be found in the development of flavour tagging algorithms, in particular in cases when simulated events are used for the training of machine learning based classifiers.

Let us consider one of the samples used for this scope by the ATLAS Collaboration, $Z' \to c\bar{c}$[4]. The events are generated at LO and showered with Pythia8, and jets are required to have $|\eta| < 2.5$ and $p_{\mathrm{T}} > 250$ GeV. In all cases, heavy hadron decays are suppressed.

In Figure 2.11, we show how for intermediate values of p_{T} in this range, between 1.5 and 4 TeV, there is a surplus of jets labelled as c-jets by the flavour jet algorithms and anti-k_t labelling scheme compared to the cone algorithm. If instead we consider b-jet labelling, the expected behaviour is found, namely a surplus of jets labelled by the cone algorithm throughout the entire p_{T} spectrum.

By considering a dual "bc" label for jets labelled with the cone algorithm, we uncover the origin of this c-jet anomaly. As shown in Figure 2.11c, $g \to b\bar{b}$ splittings contaminate the c-jets with b flavour. When labelling b-jets, there is no issue as all b flavour in these events originates from gluon splittings. For c-jets, on the other hand, algorithms which more aggressively remove

[4] Z' refers to a particle with the same properties as the Z boson but with an artificially high mass. It is used to derive distributions of jets at high p_{T}.

gluon splittings tend to assign the correct flavour, while those which less aggressively remove these splittings will mislabel some c-jets as b-jets. When considering jets either labelled as c or with the dual flavour, the expected behaviour for the cone algorithm returns, as the preivously missing c-jets are now counted.

This again highlights a shortcoming in flavour labelling used by the ATLAS Collaboration. It is necessary to take into account multiple labels when developing a training sample to avoid misidentifying hard c-jets as b-jets. In these cases, the label assigned is wrong even if one is purely interested in fragmentation studies, as these are jets where the p_{T} of the c-quark can be far greater than than that of the b-quark arising from the splitting.

2.3.7 Jet substructure

Finally, we investigate the effects of flavour labelling on jet substructure, specifically concentrating on the Lund Jet Plane and the jet angularities.

The investigation of the Lund Jet Plane is carried out as part of the study described in Section 2.3.5. The aim is to understand if, correcting for the effect of jet constituents, there are significant differences in the substructure found by the algorithms and the current experimental labelling strategies (ghost). To account for the difference in the constituents, only jets with undecayed heavy hadrons are considered.

Figure 2.12 shows the ratio of the Lund Jet Planes for jets labelled according to the IFN algorithm and the ghost algorithm. In the principal part of the plane, there are no significant differences for b-jets, and slight differences of the order of a few percent for c-jets at high k_t. In less populated regions of the plane, jets labelled with ghost are characterised by more emissions, up to over 10% in the region of the plane populated by wide-angle, hard emissions.

In Figure 2.13, we show the jet angularities spectrum predicted as part of the study carried out in Section 2.3.2. The MC generation is carried out at NLO+PS without hadronisation, and sensitivity to the heavy flavour quark mass is investigated by setting it to zero. Angularities with $\alpha = 0.5$ are shown for b-jets and c-jets. In the top panel, the differential cross section for the angularities for Sherpa, the Herwig7 angular ordered shower with quark mass, and the Herwig7 dipole shower without quark masses is shown for anti-k_t labelled jets with the mod-2 recombination scheme. The ratio of the flavoured jets, again with the mod-2 recombination scheme, is then shown to the anti-k_t labelled jets for the three parton showers.

The predictions obtained with the Sherpa and Herwig7 angular ordered showers agree with each other, and the distribution of the angularities is found to be shifted towards larger values compared to the predictions obtained with the massless splitting functions. Small differences are seen between the various flavoured jets algorithms, but these are enhanced when considering c-jets. The excess of jets labelled with the ghost algorithm is consistent with wide-angle radiation entering the jet. No differences are observed between the spectrum predicted by the algorithms for b-jets and c-jets when the massless dipole shower is considered.

2.4 Conclusions

In this chapter, we've shown how gluon splittings at NNLO and beyond can introduce flavour to events which leads to jet mislabelling. We have introduced new flavour jet algorithms which resolve this problem by recognising and neutralising flavour stemming from these gluon splittings. We have shown how these algorithms are effective both at parton level and hadron level, and highlighted areas in experiment, namely in flavour tagging training, where the effect of soft gluon splittings to $b\bar{b}$ pairs can lead to the misidentification of jets in data due to the prioritisation of a much softer b-hadron with respect to the hard c-hadron arising from the hard process.

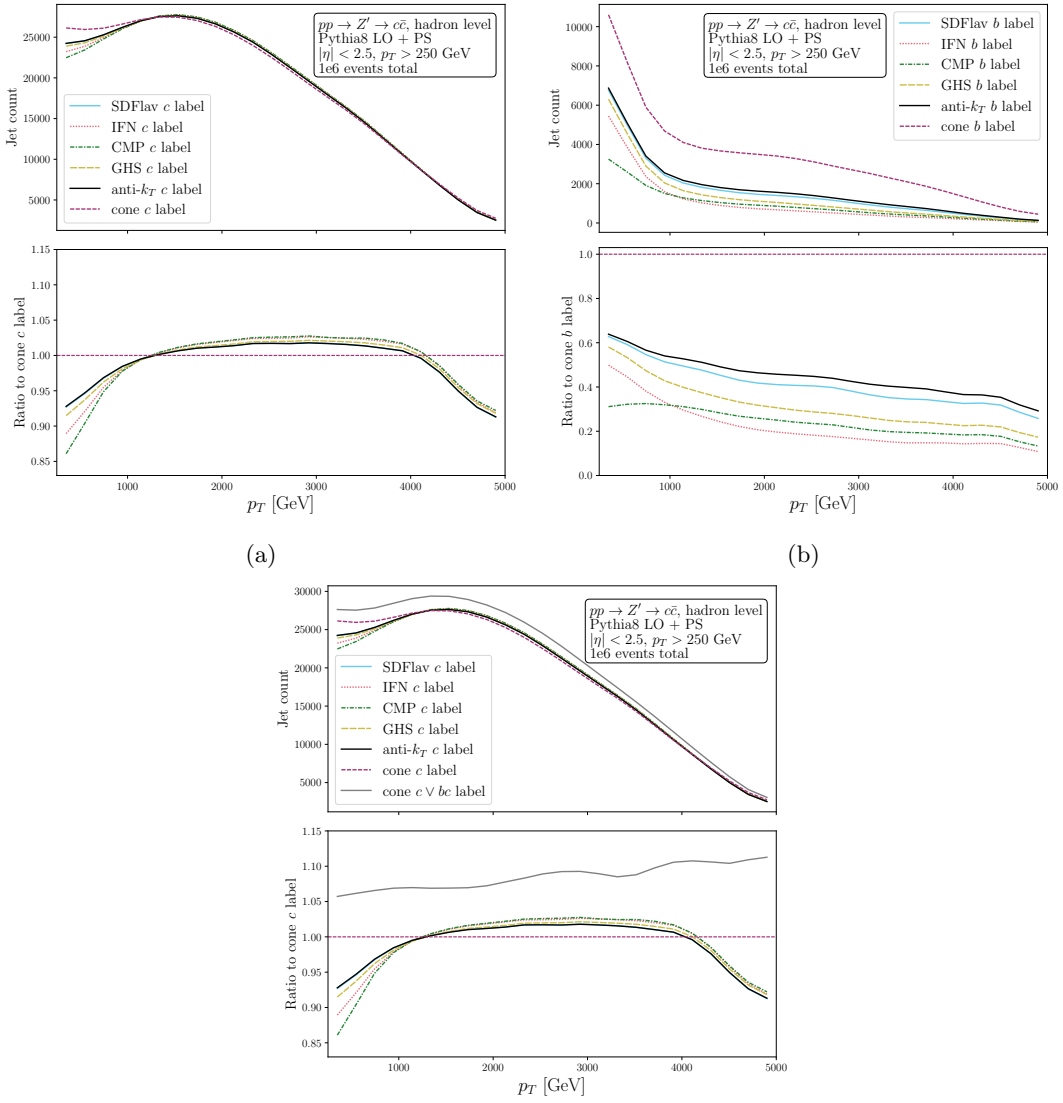


Figure 2.11: The p_{T} spectra of c-jets (top left) and b-jets (top right) found in $pp \to Z' \to c\bar{c}$ events, and the same p_{T} spectrum of c-jets compared against jets labelled as either c or "bc" with the cone algorithm (bottom) [82].

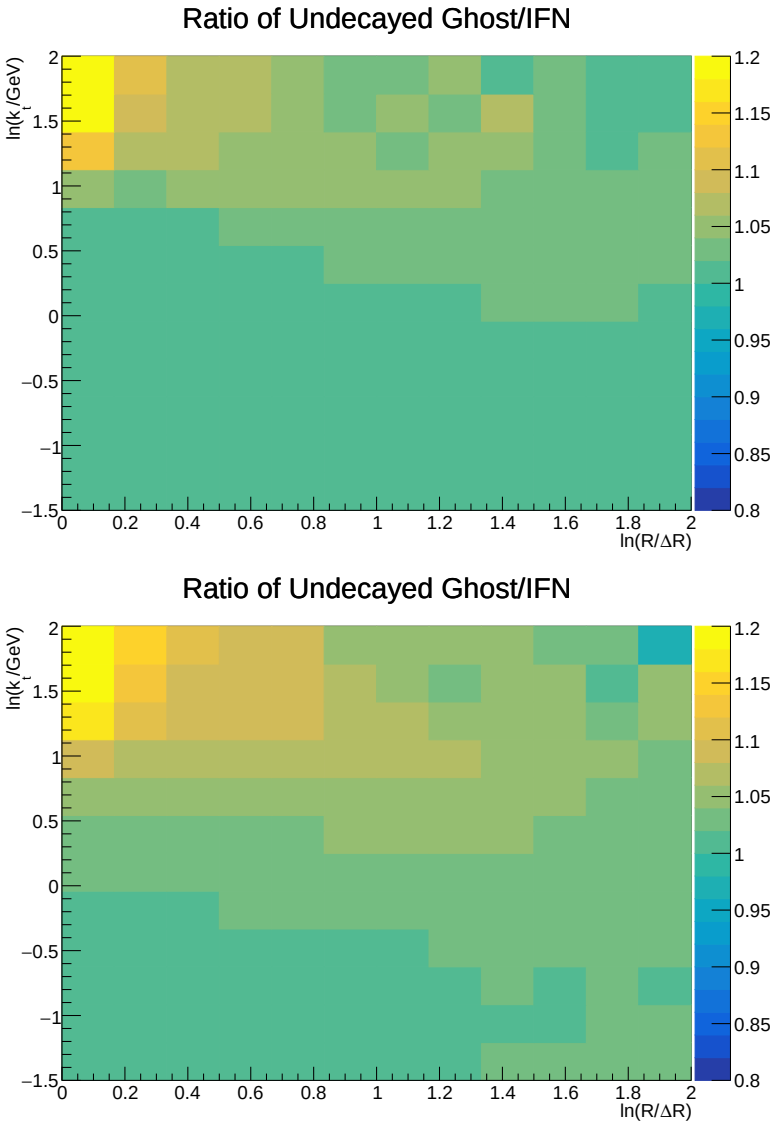


Figure 2.12: The ratio of the Lund Jet Plane for leading b/c-jets (top/bottom) identified with ghost labelling containing undecayed b/c hadrons and those identified with IFN.

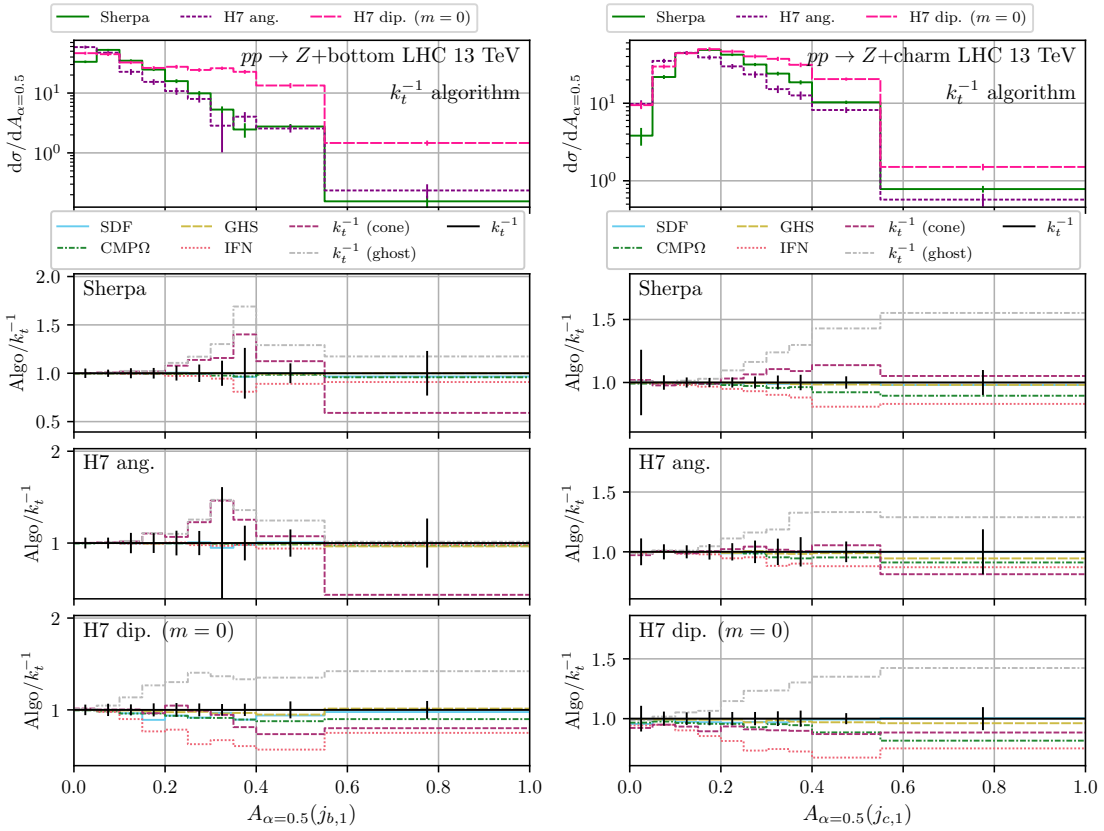


Figure 2.13: The Les Houches angularity $A_{\alpha=0.5}$ of the leading b-jet (left) and c-jet (right) for NLO+PS predictions obtained at parton level with massive Sherpa and Herwig7 angular ordered showers and for a massless Herwig7 dipole shower [82].

Additionally, we have identified the need for harmonisation in the definition of jet constituents between theory and experiment required to avoid intrinsic kinematic differences in jets which in turn affect which jets are identified as flavoured or not. Lastly, we have studied the effects of the novel flavour jet algorithms on jet substructure, specifically the jet angularities and the Lund Jet Plane, and shown how they are able to reduce the number of wide-angle emissions in jets arising from gluon splittings.

Chapter 3

The ATLAS Detector

The ATLAS Detector was designed for the purpose of reconstructing the full kinematics of a given final state after a collision event at the Large Hadron Collider (LHC). This is no easy task: with forty million bunch crossings per second, state-of-the-art technology is needed to meet the challenge.

In this chapter, we describe the many layers which compose the ATLAS Detector and illustrate how they come together to become the marvel of engineering that we know and love. We first discuss the acceleration chain at the LHC in Section 3.1 and then the ATLAS Detector in detail in Section 3.2. We conclude the chapter in Section 3.3 with a description of how the various objects relevant to this thesis, such as tracks, jets, and leptons, are reconstructed.

3.1 The Large Hadron Collider

Let us begin our journey from the underlying facility: the LHC. The LHC represents the largest particle accelerator ever built and currently in operation. It provides collisions for the four major experiments, ATLAS, CMS, LHCb and ALICE, along with a number of secondary experiments such as LHCf, TOTEM, MoEDAL and FASER.

The LHC can be best described as an accelerator complex rather than a single collider. In fact, before protons (or heavy ions) are injected into the LHC, they are accelerated in stages which rely in part on the previous colliders built at CERN. We briefly describe these stages below.

First of all, protons are sourced from a flask of hydrogen gas. The gas is ionised with an electric field to isolate the protons. Next, the protons are accelerated with a linear accelerator. Up until 2020, this was the Linac 2, though later it was the Linac 4. This step brings the energy up to 50 MeV. The protons are then sent into the Proton Synchrotron Booster (PSB) which, besides further accelerating the particles to 1.4 GeV, also separates the proton beam into bunches. The protons are then injected into the Proton Synchrotron which brings the energy up to 26 GeV, followed by the Super Proton Synchrotron (SPS) which accelerates the protons to 450 GeV.

Lastly, the protons are injected into the LHC. This machine is a 27 km underground ring which is responsible for accelerating the particles to the final collision energy and maintaining that energy until collisions deplete the beam, at which point more protons are injected. The proton bunches are injected into two beam pipes and travel in opposing directions. At the four interaction points where the main experiments are located, the bunches cross every 25 ns allowing for collisions to occur.

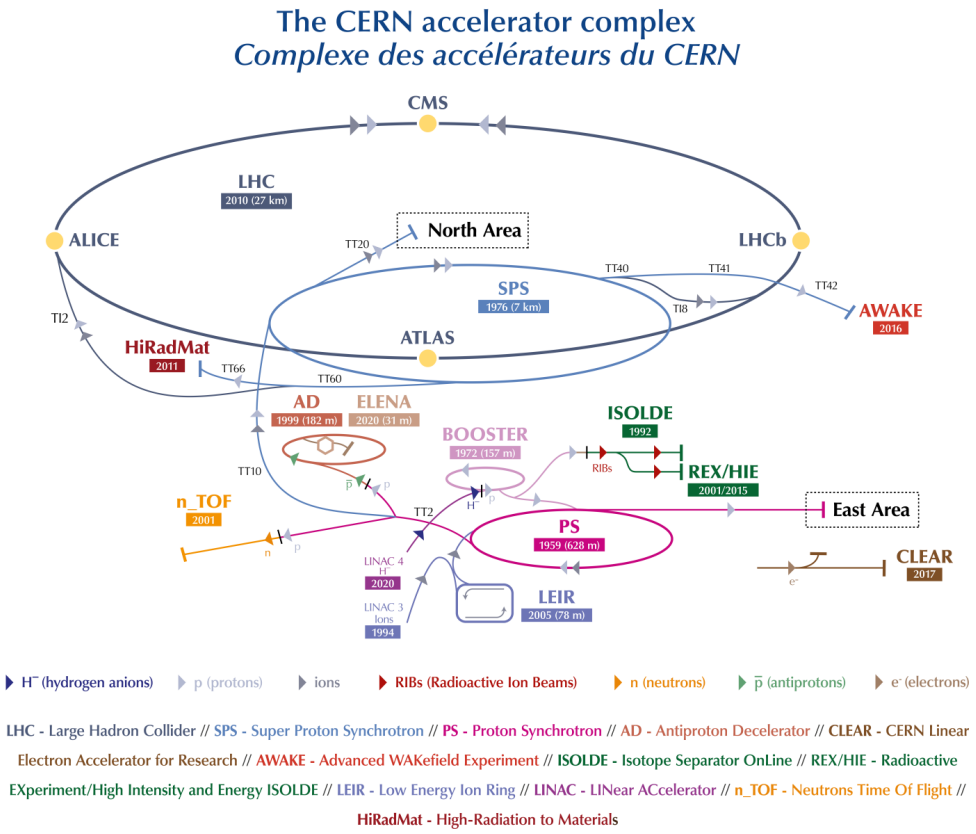


Figure 3.1: The full accelerator complex at CERN [94].

The main advantage of a circular collider is to maximise the collision energy. During the first run of data taking (Run 1) from 2010-2013, the LHC collided protons at $\sqrt{s} = 7$ TeV or at $\sqrt{s} = 8$ TeV. In Run 2, from 2015-2018, protons were collided at $\sqrt{s} = 13$ TeV. Since 2022, Run 3 has been underway and protons are being collided at $\sqrt{s} = 13.6$ TeV. In the future, during Runs 4 and 5, the centre-of-mass energy may reach up to 14 TeV, depending on the magnets.

Figure 3.1 shows the full accelerator complex at CERN, including all major and minor experiments currently underway.

3.1.1 Luminosity

The luminosity is an important quantity which keeps count of how many collisions have occurred. The instantaneous luminosity

$$L_{\text{instant}} = \frac{N_1 N_2 f N_b}{4\pi\sigma_x\sigma_y}, \tag{3.1}$$

determined by beam parameters such as the number of protons in the two colliding bunches N_1 and N_2, the beam frequency f, the number of bunches N_b, and the beam width in the x and y directions σ_x and σ_y, specifies how many opportunities there are for collisions. It is measured in

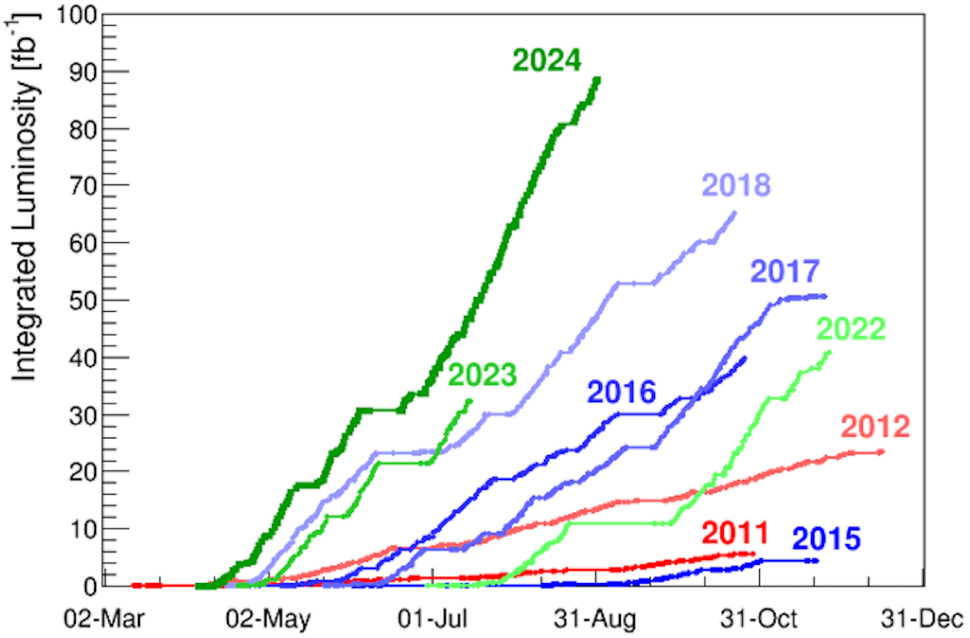


Figure 3.2: The total integrated luminosity provided by the LHC during each year of operation [95]

.

cm^{-2}s^{-1}. When combined with the cross section for a certain process σ, we can determine the rate at which that process occurs

$$\frac{dN}{dt} = \sigma L_{\text{instant}}. \tag{3.2}$$

If we integrate the instantaneous luminosity over time $L = \int L_{\text{instant}} dt$ we get the *integrated luminosity*, which allows us to get the total number of collisions which have occurred N

$$N = \sigma L. \tag{3.3}$$

Figure 3.2 shows the total integrated luminosity provided by the LHC during each year of operation, up to 2024. The results for individual experiments may vary slightly based on any dead time experienced by the detector. For Run 2, the data taking period comprising the measurement discussed in this thesis, a total of 140.1 fb^{-1} data was recorded at the ATLAS Experiment.

Starting from 2030, the High-Lumi LHC programme is expected to increase the instantaneous luminosity delivered by the LHC to a value of 5×10^{-34}cm^{-2}s^{-1}, or a factor of 5 greater than the nominal value in the original LHC design [96]. This is expected to deliver a total integrated luminosity of 3000 fb^{-1} during Runs 4 and 5 [97]. This means that about 90% of collisions are yet to be delivered, as of 2025!

3.1.2 Pile-up

Over time, the instantaneous luminosity at the LHC has increased. As it has increased, so too has the average pile-up.

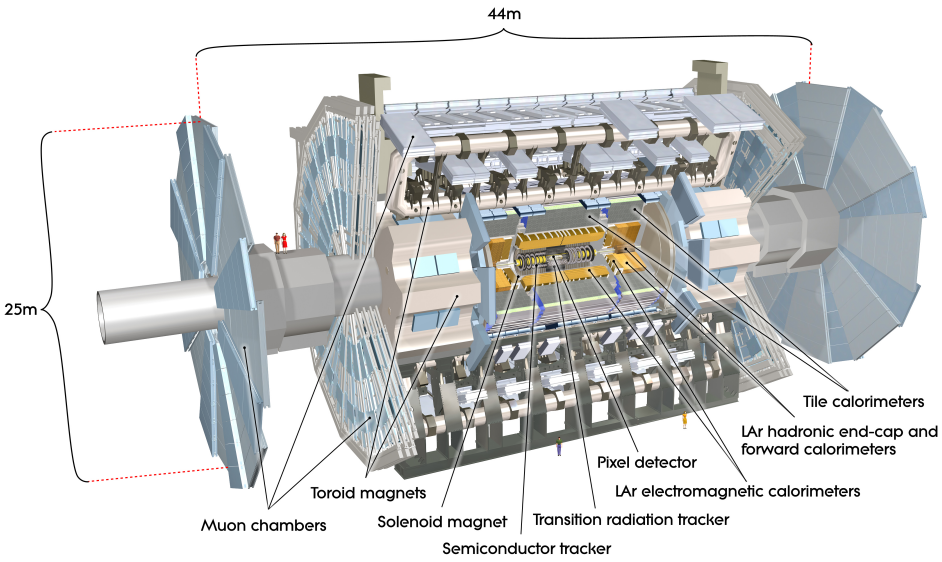


Figure 3.3: The ATLAS Detector with all subsystems. [100].

Usually when proton bunches collide, at most one or two hard, head-on scattering events between protons will occur. These are the collisions which are the most interesting. However, there are often many underlying soft, lateral collisions between protons in the bunches. This is known as *pile-up*.

In Run 1, the average pile-up $\langle\mu\rangle$ per event recorded in ATLAS was roughly 20 collisions/interaction [98]. This increased to $\langle\mu\rangle = 33$ in Run 2 [99], and was as high as $\langle\mu\rangle = 62$ in Run 3 during 2024 data taking. During Runs 4 and 5, pile-up is expected to increase to $\mathcal{O}(200)$ collisions/interaction.

3.2 The ATLAS Detector

ATLAS is a general-purpose detector built for multiple scopes. These include discovering the Higgs boson and searching for new physics phenomena, but also probing the Standard Model to the maximal precision and in extreme phase space. The ATLAS precision physics programme provides valuable data which has been able to constrain models which extend the Standard Model beyond the energy ranges we can currently reach.

In this section, we describe the design of the ATLAS Detector (shown in its full glory in Figure 3.3) during Run 2 to ascertain how it is capable of delivering such impressive results.

3.2.1 Coordinate system

Let us start by describing the coordinate system used at the experiment. ATLAS uses a right-handed cylindrical coordinate system. The z-axis is pointed along the beam, while the x-axis points towards the centre of the LHC ring and the y-axis points upwards towards the sky. In the transverse plane, polar coordinates (r, ϕ) are used. ϕ is the azimuthal angle around the z-axis, and r gives the distance from the beamline in the $x - z$ plane.

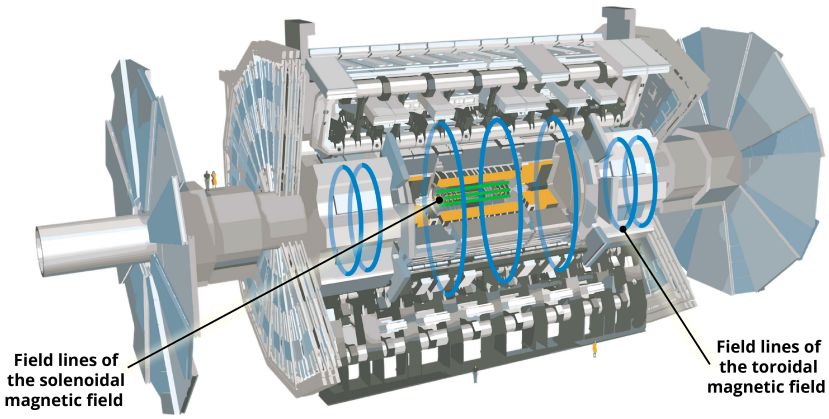


Figure 3.4: The solenoidal and toroidal magnetic fields generated within the ATLAS Detector [101].

For the sake of convenience, rather than r, the *rapidity* y is used, defined as

$$y = \frac{1}{2} \ln \left(\frac{E + p_z}{E - p_z} \right), \tag{3.4}$$

where E and p_z refer to the energy and the momentum along the z-axis for a given particle, respectively. At high energies, when $p \gg m$, the *pseudorapidity* η can be used equivalently

$$\eta = -\ln \left(\tan(\theta/2) \right), \tag{3.5}$$

where θ is the angle with respect to the z-axis. The advantage of using η rather than y is that η involves purely geometrical information, instead of four-momentum coordinates of a specific particle.

Kinematics are closed in the transverse plane. This is due to the fact that, at high energy hadron colliders, it is the partons which collide rather than the hadron itself. The partons are each carrying an unknown fraction x of the total momentum, but conservation of momentum ensures that this quantity is conserved in the transverse plane.

3.2.2 Magnets

The magnetic field is a key feature of the detector since it allows for the measurement of particles' charge and momentum through their deflection via the Lorentz force. There are two magnet systems present at ATLAS. The first is a 2 T solenoidal magnetic field the along the z-direction, fully surrounding the Inner Detector. The second is a toroidal magnet system integrated in the Muon Spectrometer on the outermost part of the detector. There is one toroidal magnet in the central barrel and two in the end-caps. The field strength produced by the toroidal magnet varies across the detector, between 0.5 T in the central region, and 1 T in the forward end-cap regions. The acronym "ATLAS" originally stood for "A Toroidal LHC ApparatuS". The toroidal magnets therefore can be seen as the defining feature of the detector. The magnetic field lines within the ATLAS Detector are show in Figure 3.4.

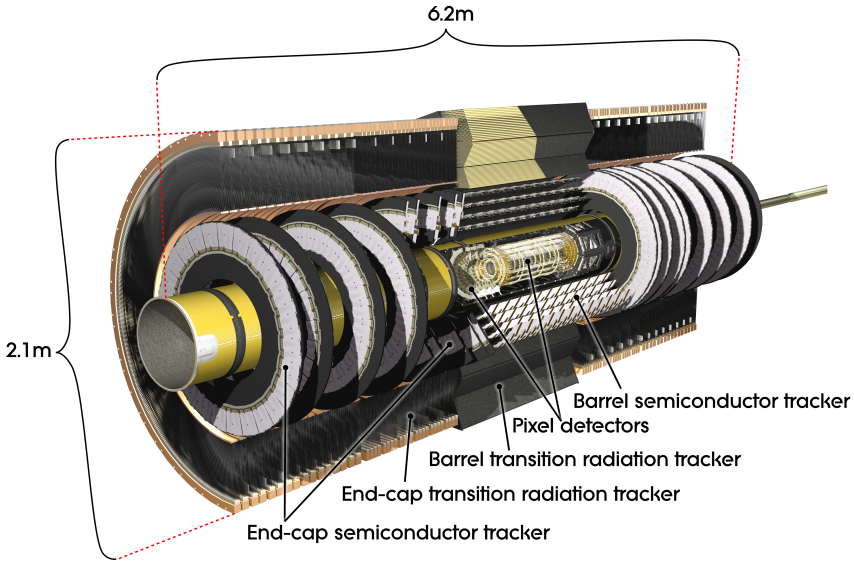


Figure 3.5: The Inner Detector and all of its subsystems [102].

3.2.3 Inner Detector

The innermost layer of the ATLAS Detector is known as the *Inner Detector* (ID). Covering a range up to $|\eta| < 2.5$, it is designed to track charged particles as they moved through the detector, deviated by a magnetic field. The presence of the magnetic field allows quantities such as the charge and momentum of the particle to be determined.

The ID is composed of three subsystems: the Pixel Detector, the SemiConductor Tracker (SCT) and the Transition Radiation Tracker (TRT). The ID and its subsystems, which we describe below, are shown in Figure 3.5.

Pixel Detector

The silicon Pixel Detector is the closest to the beamline – just 3.3 cm away – and thus has been designed to provide the highest granularity. It consists of four cylindrical layers, as well as two end-caps with 3 layers each. The innermost layer of the central barrel of the subdetector is known as the Insertable B-Layer (IBL). This part of the Pixel Detector has the highest granularity at 50 μm in the transverse plane and 250 μm along the z-axis. The other three layers also have a resolution of 50 μm in the transverse plane, but a wider 400 μm along the z-axis. Thanks to the geometry of the detector, this leads to a total resolution of 8-10 μm in the transverse plane and 40-115 μm in the z direction. On average, a charged particle will leave 4 hits in the Pixel Detector.

SCT

The SCT is also a silicon detector, though it is based on strip technology rather than pixels. Each strip has a length of 80 μm, and they are combined into modules composed of two sensors placed at an angle of 40 mrad with respect to each other. In this way, a module allows for two-dimensional reconstruction of the direction of the charged particle. Strip modules are arranged

into four cylindrical layers and two end-caps with nine layers of modules each. The SCT accuracy is required to be $\sigma(r\phi) = 17$ μm and $\sigma(rz) = 580$ μm [103].

TRT

The TRT is the outermost part of the ID. It is a gaseous detector composed of cylindrical (straw) tubes with a diameter of 4 mm and is characterised by an intrinsic accuracy of 130 μm per straw. The TRT provides particle identification capabilities. The empty space between the straws is filled with polymer fibres. As a particle traverses the boundary between the media, it emits radiation. The frequency of the radiation depends on the Lorentz boost of the emitting particle $\gamma \sim E/m$. As such, the TRT can distinguish particles of different masses, provided they have the same momenta. Generally, this equates to distinguishing electrons from pions within the momentum range $1 - 100$ GeV^2.

3.2.4 Calorimeters

ATLAS has two calorimeters: the Electromagnetic Calorimeter (ECAL), designed principally to detect energy deposits by electrons, positrons and photons, and the Hadronic Calorimeter (HCAL), designed to detect energy deposits by hadrons.

The calorimeters are nearly hermetic, covering the full span of ϕ and $|\eta| < 4.9$, in order to contain as much of the energy of the event as possible. Both calorimeters are *sampling calorimeters*. They feature an active layer where the energy is deposited, and a high atomic number passive layer which induces showers. The calorimeters measure only a fraction of the energy deposited, which is however proportional to the total energy of the particle responsible for the deposit.

The ECAL uses lead as a passive material and liquid argon (LAr) as the active material. The HCAL, on the other hand, uses varying materials. In the HCAL end-cap and forward calorimeters, liquid argon is used as the active material. In the forward calorimeters, copper and tungsten are used as passive materials, while in the end-cap it is just copper. In the barrel, the tile calorimeter uses steel as a passive material and plastic scintillator as the active material.

The region between the barrel and end-cap calorimeters is known as the *crack region*. It is found between $1.37 < |\eta| < 1.52$. As this region has a particularly high amount of passive material, electrons and photons at these pseudorapidities are often excluded from analyses [104].

Both electromagnetic and hadronic showers have a characteristic length. In electromagnetic showers, this is X_0, or the length at which the shower energy is depleted by a factor $1/e$. In hadronic showers, the characteristic length is the average distance travelled by a hadron between nuclear interactions λ. In these units, the ECAL has a depth $X_0 = 22$ and the HCAL has a depth $\lambda = 10$. Therefore, both are able to contain a majority of the energy deposited.

The ECAL is required to have a resolution of $\sigma_E/E = 10\%/\sqrt{E} \oplus 0.7\%$, where E is the energy deposited in the calorimeter. Likewise, the HCAL is required to have a resolution of $\sigma_E/E = 50\%/\sqrt{E} \oplus 3\%$ in the barrel and end-caps [105].

The ATLAS calorimeters are represented in Figure 3.6 .

3.2.5 Muon Spectrometer

Muons produced by colliders are typically minimum ionising particles; they deposit little energy in the detector and are able to fly through and escape. Thankfully, they do deposit *some* energy, appearing as tracks which are then reconnected to hits in the outermost layer of the detector: the Muon Spectrometer (MS).

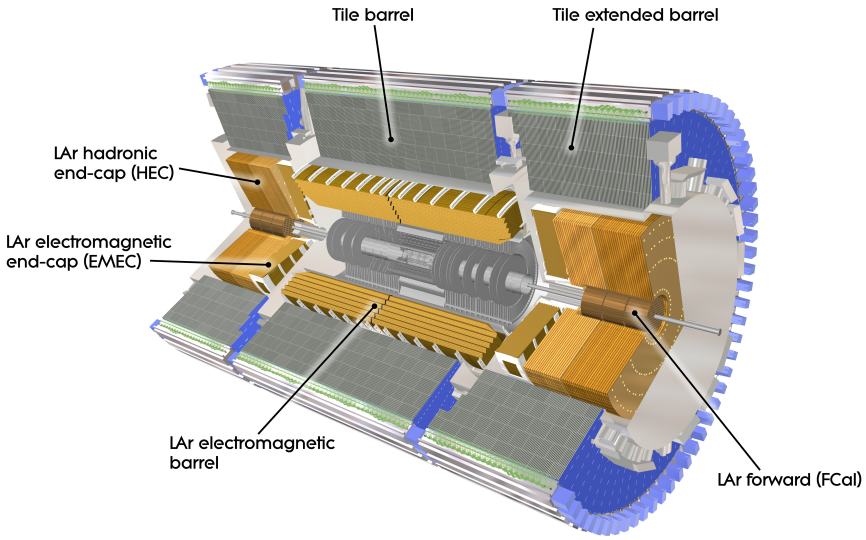


Figure 3.6: The ATLAS ECAL and HCAL systems [106].

The Monitored Drift Tube (MDT) allows for precision tracking. It spans a region up to $|\eta| < 2.7$ and is composed of multiple layers. In the central region, up to $|\eta| < 2.0$, the MDT provides a resolution of 35 μm. In the high occupancy region, between $2.0 < |\eta| < 2.7$, the Cathode Strip Chambers (CSCs) provide coverage. These are multi-wire proportional chambers which provide a resolution of 60 μm in the r direction and 5 mm in the ϕ direction. The overall momentum resolution of the MS depends on the p_{T} of the muons detected: for muons of $p_{\mathrm{T}} = 100$ GeV, the resolution $\delta p/p$ is of the order of 3%; for muons of $p_{\mathrm{T}} = 1$ TeV it is of the order of 10% [107], [108]. At high momentum, the resolution normalised to transverse momentum ($\sigma_{p_{\mathrm{T}}}/p_{\mathrm{T}}$) is proportional to $\sigma_{p_{\mathrm{T}}}/p_{\mathrm{T}} = \sigma_{\mathrm{res}} \times p_{\mathrm{T}}$, where σ_{res} is the single-hit spatial resolution, which has a design value $\sigma_{\mathrm{res}} = 80\ \mu m$ [109].

The Muon Spectrometer also features a built-in trigger system . Three layers of Resistive Plate Chambers (RPCs) are used in the central barrel region and three or four layers of Thin Gap Chambers are in the end-cap regions. These two systems provide a hardware trigger by detecting the coincidence of the hits in the different layers when a muon passes through the detector. They also allow a measurement of the trajectory of the particle.

The Muon Spectrometer in its entirety is shown in Figure 3.7.

3.2.6 Trigger

The rate of collisions at the LHC is 40 MHz. As such, an extraordinary amount of data is produced, more than we are capable of storing. It is therefore necessary to put in place a trigger system to select and save only the most interesting collisions. The ATLAS trigger system is composed of the Level 1 Trigger (L1) and the High-Level Trigger (HLT). The L1 Trigger is hardware-based system which, using signatures from the MS and calorimeters, decides if an event contains high p_{T} particles. It reduces the rate of events to roughly 100 kHz. These interesting events which pass are saved in a buffer, and are then fed to the HLT, which operates on a dedicated farm. The HLT has access to information from the entire detector as well as regions

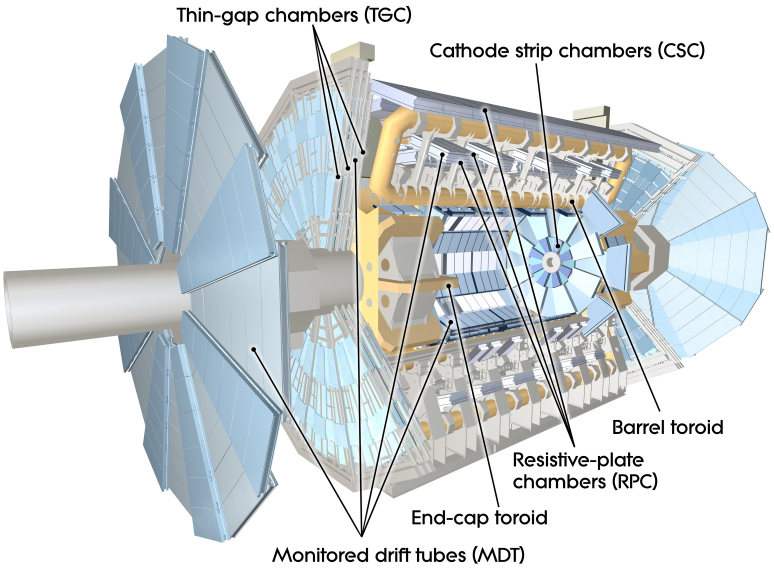


Figure 3.7: The ATLAS Muon Spectrometer [110].

of interest identified by the L1 Trigger. It uses software to further skim the number of events, down to a rate of 1.5 kHz. These then get written to disk as data.

3.2.7 Detector operations

It takes a dedicated team to run the ATLAS Detector. During data taking periods, typically from March-October, a myriad of people are responsible for monitoring each subsystem in real time, both day and night. Some may be in the control room, ensuring that the detector is functioning as expected. Others are behind the scenes, updating trigger keys based on information coming from the LHC, ensuring each subsystem of the detector is ready for the plan of the day, and others still are monitoring data quality to identify that the events saved are indeed useful.

Personally, I was involved in Run Control and Trigger Operations. I was responsible for monitoring detector dead time, warnings and errors in real-time and reacting to them appropriately. I also monitored trigger rates to anticipate any problems in the various detector subsystems, and adjusted the triggers to account for the decrease in instantaneous luminosity as the protons in the beam were exhausted, and prepared the triggers and detector for each new fill from the LHC.

3.2.8 Putting it all together

The combination of various subsystems allow for the identification of the majority of particles that pass through the ATLAS Detector.

Charged particles are tracked in the ID. Muons will also produce hits in the MS, allowing for the identification of tracks corresponding to muons. Electrons and positrons will deposit their energy in the ECAL. Pions, kaons and other charged hadrons will instead deposit their energy in the HCAL. Photons will only deposit energy in the ECAL. Neutrinos will pass through the detector without leaving a trace. The signatures that various particles leave in the detector are schematically represented in Figure 3.8.

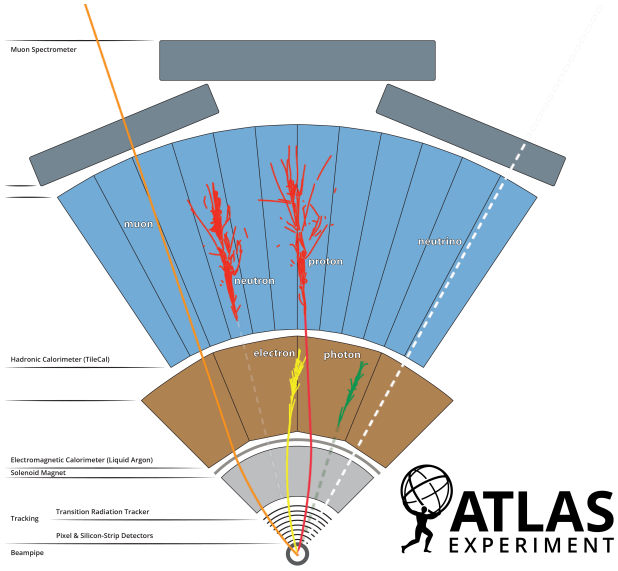


Figure 3.8: The signatures left by different types of particles in the ATLAS Detector [111].

The hits (or lack thereof) produced in the different layers allow for the construction of what is known as *physics objects*. These are objects such as electrons, photons, and muons, which directly correspond to particles, though they can also be more abstract concepts such as tracks, $E_{\mathrm{T}}^{\mathrm{miss}}$ and jets. Data analysis is carried on physics objects, and calibration ensures that any detector effects are corrected for before carrying out the analysis.

3.3 Object reconstruction

In this section, we will discuss how the relevant physics objects for this thesis are reconstructed. These are the basis for the selections for the measurement described in Chapter 5.

The measurement focuses in particular on final states involving Z+jets, with the Z decaying leptonically. As a result, the reconstruction of electrons and muons is key to the measurement.

Jets in the final state are reconstructed with different radii – $R = 0.4$, $R = 1.0$, but also with variable radius – in order to identify events with resolved and boosted topologies. Jets of interest lie across a wide p_{T} range, from 20 GeV, up to and exceeding 1 TeV.

Jets from heavy flavour are identified using flavour fagging techniques. These methods have a dedicated introduction in Chapter 4.

Lastly, as one of the main focuses of the measurement is on jet substructure, tracks are of particular importance due to their higher resolution when compared to the calorimeter.

3.3.1 Tracks

Tracks are formed by matching hits in the different layers of the ID. Two algorithms are used for reconstruction. The first is an inside-out algorithm which reconstructs tracks starting from any series of three or more hits along the same axis in the various layers of the Pixel and SCT detectors. The track is then extrapolated to the end of the SCT detector, accounting for

the magnetic field. Multiple candidates are considered in cases where ambiguity is present; the final selection is made based on criteria such as the fit quality or the number of missing hits, which would penalise the candidate in question. The tracks which pass this step are extended to the TRT and combined with the hits registered in that subdetector. To cross-check the result of the track fitter, a complementary outside-in algorithm is run, known as back-tracking. This time, the algorithm starts from hits in the TRT and gradually adds hits in the silicon layers to form tracks. This algorithm is also able to trace tracks coming from photon conversions, decay chains of long-lived particles such as the K_s, and other charged particles produced in a way which would not produce hits in the silicon layers. Tracks are required to have a resolution of $\sigma_{p_T}/p_T = 0.05\%/p_T \oplus 1\%$. During Run 2, specifically during 2015 data taking, the resolution reached $\sigma_{p_T}/p_T = 0.038\%/p_T \oplus 1.5\%$ [105].

Track coordinates

In the ID, tracks are immersed in a solenoidal magnetic field, leading to helicoidal trajectories. It is thus useful to use a helicoidal coordinate system to define the trajectory in space. The parameters used to define the trajectory are $(d_0, z_0, \phi, \theta, q/p)$. These correspond to, in order, the transverse (d_0) and longitudinal (z_0) impact parameters, or the distance between the closest point of the trajectory of the track and a reference point in the transverse plane or along the z-axis; the azimuthal (ϕ) and polar (θ) angles of the track at the point of closest approach to the reference point; and the ratio q/p of the charge of the reconstructed track divided by the magnitude of its momentum. The reference point is usually taken to be the beamspot position, or the average point of pp interactions.

Track quality

Tracks used for analyses must be within the acceptance of the ID ($|\eta| < 2.5$) and have a $p_T > 500$ MeV. Additional quality requirements can sometimes be placed to filter out fake tracks. Examples of these are

- **Loose Track Selection:** Tracks must have $p_T > 500$ MeV cut and $|\eta| < 2.5$. There must be at least eight hits in the silicon layers. There can be no more than one shared module between the pixel detector and SCT, no more than one hole in the pixel detector and no more than two holes total in the pixel layers.

- **Tight Track Selection:** Tracks must have at least nine hits in the silicon layers if their pseudorapidity $|\eta| \leq 1.65$ and at least eleven hits otherwise. There can be no holes in the pixel layer, and the sum of hits in the IBL and the B-Layer must be greater than zero, if both hits are expected. All other requirements are the same as in the Loose selection.

3.3.2 Vertices

Vertices are identified as points from which tracks originate. There are two types of vertices. Primary vertices are located roughly where the pp interaction occurs and are the point from which the event originates. Secondary vertices can be found anywhere in the ID and are the point from which particle decays, such as decays of heavy hadrons, occur.

Primary vertex candidates must have at least two tracks associated to them with $p_T > 500$ MeV. The primary vertex is identified as the point with the highest p_T^2 sum of contributing tracks. Each event must contain at least one primary vertex.

Track-to-Vertex Association

To reduce the number of tracks fro pile-up interactions, selections can be applied to increase the probability that a track comes from the primary vertex. These are:

- **Loose working point:** This selection corresponds to requiring that the track be used in any track fit, and that the $\Delta z \sin\theta < 3$ mm. Here, $\Delta z = |z_0 + v_z - z_{\mathrm{BS}}|$, where v_z and z_{BS} are the z-coordinates of the primary vertex and beamspot, respectively..
- **Nominal working point:** The track must have $\Delta z \sin\theta < 3$ mm and $|d_0| < 2$ mm.
- **Tight working point:** The track must have $\Delta z \sin\theta < 0.5$ mm and $|d_0| < 0.5$ mm.

3.3.3 Electrons

Electron candidates are identified from energy clusters in the ECAL that are linked to reconstructed tracks in the ID. Candidates must be found within the calorimeter, i.e. within $|\eta| < 2.47$, excluding the transition region between $1.37 < |\eta| < 1.52$.

Candidates must be associated to the primary vertex. This is done through a cut on the longitudinal and transverse impact parameters $|z_0 \sin(\theta)| < 0.5$ mm and $|d_0/\sigma(d_0)| < 5.0$.

In cases in which electron candidates cannot be unequivocally identified as electrons, they must meet criteria regarding the shape of the electromagnetic shower, track quality, and track alignment with respect to the calorimeter cluster to distinguish them from photons. This is done via a likelihood based fit, with a "Tight" working point [112], corresponding to a cut on the likelihood discriminant, the requirement that the energy/momentum ratio $E/p < 10$, and requiring that the primary track have $p_{\mathrm{T}} > 2$ GeV. Furthermore, electron candidates are rejected if a two-track conversion vertex in the silicon layers is reconstructed with momentum closer to that of the calorimeter cluster than of the primary track.

Finally, in order to qualify as electrons, candidates are required to be meet isolation criteria satisfying the Tight_VarRad isolation working point. This equates to having both the calorimeter energy sum and p_{T} sum of the tracks within a variable radius cone (up to a maximum radius of $\Delta R = 0.2$) centred around the electron less than $0.06 E_{\mathrm{T}}$, where E_{T} indicates the transverse energy.

3.3.4 Muons

Muon candidates are identified by correlating tracks found within the ID to tracks or track segments in the MS. They must meet impact parameter criteria of $|z_0 \sin(\theta)| < 0.5$ mm and $|d_0/\sigma(d_0)| < 3.0$, and have $|\eta| < 2.5$.

Muon candidates must pass the "Medium" identification working point [113], based on quality requirements on the tracks in the ID and the MS. Candidates must be so-called combined muons, i.e. they are required to have hits in both the ID and MS. Specifically, they must have at least three hits in at least two MDT layers, except for tracks in the $|\eta| < 0.1$ region where candidates are allowed to have hits in at least one MDT layer but may have no more than one hole in the MDT layers. To reduce contamination from hadrons misidentified as muons, the q/p significance must be $(q/p)/\sigma(q/p) < 7$.

Lastly, muons must pass the "FCTight_FixedRad" working point to pass isolation requirements, i.e. the p_{T} sum of tracks in a variable radius cone (maximum radius $\Delta R = 0.3$) centred around the muon must be no more than 0.04 times the muon p_{T}.

Muons, when reconstructed only in the MS, must have a resolution of $\sigma_{p_T}/p_T = 10$ at $p_T = 1$ TeV. This was matched in 2015 data taking, and, when combined with information from the tracker, a resolution of $\sigma_{p_T}/p_T = 7$ at $p_T = 1$ TeV was obtained [105].

3.3.5 Missing transverse momentum

Not all particles leave a visible trace in the detector. Neutrinos, due to their low interaction cross section, traverse the entirety of the detector without any energy deposit.

Their presence can still be inferred through the *missing transverse momentum*, or $\vec{p}_T^{\,miss}$[1] [114]. Since kinematics in the transverse plane are closed, the vector sum of the transverse momentum in this plane is expected to be null. The amount missing to render this true is the missing transverse momentum. It corresponds to a two-dimensional vector, with magnitude p_T^{miss} and direction ϕ^{miss}.

Neutrinos are not the only source of missing transverse energy. Other sources include cosmic muons from outside the event, beam halo muons, and invisible particles from beyond the Standard Model, if they exist and are produced. The presence of missing transverse momentum is thus critical in many searches for new physics.

3.3.6 Jets

Jets in ATLAS are clustered with the anti-k_t algorithm. However, rather than using particles, physics objects are used instead. The types of jets and the objects used to cluster them are described below.

VR trackjets

Variable radius (VR) trackjets [115] are composed of tracks only. They are characterised by a radius parameter which varies with the p_T of the jet. The radius parameter varies as follows:

$$R \to R_{eff} p_{T,i} = \frac{\rho}{p_{T,i}} \tag{3.6}$$

where R_{eff} is an effective radius parameter dependent on the transverse momentum of the proto-jet $p_{T,i}$. ρ is a parameter which determines the speed with which the effective radius decreases. It is set to $\rho = 30$ GeV. The jet radius varies according to 3.6, up to (down to) a maximum (minimum) radius of $R = 0.4$ ($R = 0.02$). Overall, the radius parameter is given by

$$R = \max(0.02, \min(0.4, \frac{\rho}{p_T}). \tag{3.7}$$

LCTopo Jets

LCTopo jets [116] are are formed from locally-calibrated topological (LCTopo) clusters of energy from the hadronic calorimeter. To construct these, first calorimeter cells noise over four times greater than the expected value are identified. These are known as seed cells. Adjacent cells to the seed cell are iteratively grouped together to form a cluster, which is calibrated at the electromagnetic scale.

This calibration, however, does not account for the different response of the ECAL to hadrons. Jets with different fractions of hadronic constituents would thus produce different responses. In addition to this, one must account for detector inefficiencies in both the ECAL and HCAL.

[1] This quantity is often referred to as *missing transverse energy* or E_T^{miss}.

To correct for these effects, a cell-by-cell calibration of the topoclusters, known as Local Cell Weighting (LCW), is applied. The calibrated clusters are then used as input objects for the jets.

LCTopo jets are clustered with a radius of either $R = 0.4$ or $R = 1.0$, depending on if they are small-R jets or large-R jets. The latter jets are also trimmed with parameters $R_{\text{sub}} = 0.2$ and $f_{\text{cut}} = 0.05$, forming the AntiKt10LCTopoTrimmedPtFrac5SmallR20Jets collection. The trimming procedure is described in Section 1.4.

PFlow Jets

Particle flow objects (PFOs) are constructed by matching tracks to calorimeter clusters, and subtracting the contribution of the energy deposit in the calorimeter due to the particles from which the tracks originate.

The algorithm works as follows: a set of "well-measured" tracks with a track fit of $\chi^2/dof < 5$, where "dof" stands for "degrees of freedom" and $p_{\text{T}} < 100$ GeV is identified. Tracks are matched to a single topocluster in the calorimeter. For each track/topocluster pair, the probability that the particle giving origin to the track deposits its energy in more than one topocluster is determined and, if this is found to be the case, additional topoclusters are added to the system considered so as to reconstruct the full shower. A cell-by-cell subtraction of the energy of the particle then occurs, and a list of tracks, topoclusters, and charged particle energy-subtracted topoclusters is then returned. These are the PFOs.

The resulting PFOs can be either *charged*, consisting of tracks, whether matched or not, and with or without subtraction, or *neutral*, consisting of unmatched topoclusters and those which remain after the energy subtraction.

Charged PFOs must be matched to the primary vertex. This step is taken to reduce pile-up contributions. The algorithm applied is known as *Charged Hadron Subtraction* and consists of a cut of $|z_0 \sin\theta| < 2$ mm on the PFOs.

By combining information from both tracks and topoclusters, particle flow jets (PFlow jets) are able to improve the angular and energy resolution of LCTopo jets, especially at low p_{T}. The radius parameters chosen for PFlow jets are either $R = 0.4$ for small-R jets, or $R = 1.0$ for large-R jets.

3.3.7 Jet calibration

Jets are corrected for detector effects in order to restore their energy and resolution to the particle-level value. This is known as the *jet energy scale* (JES) calibration and *jet energy resolution* (JER) calibrations [117], [118].

The individual jet energy response is defined as $r = E_{\text{reco}}/E_{\text{true}}$. Due to the stochastic nature of the shower and detector effects, it is often useful to consider the mode of the resulting distribution $\mathcal{R} = \text{mode}(r_{\text{e}})$. This quantity is known as the *average response* or, more often, just the *response*.

The aim of the calibration is to bring the response to unity. This is achieved in various steps. First the jet is corrected for any pile-up effects following a procedure known as in the SoftKiller algorithm[2]. Next, the jet four-momentum is corrected to that at particle-level via a MC based calibration. This is done by geometrically matching reco level jets to their truth counterparts by requiring that they be found within $\Delta R = 0.3$ of each other. At truth level, no other jet with $p_{\text{T}} > 7$ GeV may be present within $\Delta R = 1.0$, though at reco level this reduces to $\Delta R = 0.6$. Di-jet events are considered.

[2]The SoftKiller algorithm is described in Appendix A.

The average response is then calculated in bins of E^{true} and η_{det}, where η_{det} refers to the jet eta with respect to the geometrical centre of the detector. The response is then parametrised as a function of E^{reco}.

To reduce the dependence of the response on shower fluctuations, such as those due to the intrinsic differences in quark and gluon initiated jets, the Global Sequential Recombination [119] is applied. This leads to a reduction in the response dependence on jet flavour.

Lastly, an in situ calibration is applied in order to correct for the fact that simulated events are used for the calibration, rather than data. For example, this reference object can be a Z boson which decays leptonically. In this case, Z+jet events are used for the in situ calibration, and the hadronic recoil is balanced against the Z.

From this, the double-ratio of the response in simulation and data

$$c = \frac{\mathcal{R}_{\text{data}}}{\mathcal{R}_{\text{sim}}} \tag{3.8}$$

is obtained as a function of reference object p_{T}. Its p_{T} dependence with respect to the jets is then found, and the JES calibration is complete.

The JER calibration, on the other hand, is done similarly to the in situ calibration step in the JES calibration, as it also relies on the precise knowledge of the p_{T} of the jet. Specifically, the jet resolution is parametrised as follows:

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

where N is a noise term representing the contribution of detector noise to the resolution, S is a stochastic term representing the statistical fluctuations in the energy deposit, and C is a constant term accounting for fluctuations which are constant as a function of p_{T}, such as the energy deposited in the passive layer of the calorimeter. The terms in Eq. 3.9 are summed in quadrature.

To determine the JER, balanced di-jet events are typically used, with data taken in a well-calibrated region of the detector. One jet is taken as a probe and the other as reference. An asymmetry between the p_{T} of the two jets is then defined

$$\mathcal{A} = \frac{p_{\text{T}}^{\text{probe}} - p_{\text{T}}^{\text{ref}}}{p_{\text{T}}^{\text{avg}}} \tag{3.10}$$

where $p_{\text{T}}^{\text{avg}} = (p_{\text{T}}^{\text{probe}} + p_{\text{T}}^{\text{ref}})/2$. The jet resolution can be found from the standard deviation of the asymmetry

$$\sigma_{\mathcal{A}} = \frac{\sigma_{p_{\text{T}}}^{\text{probe}} \oplus \sigma_{p_{\text{T}}}^{\text{ref}}}{\langle p_{\text{T}}^{\text{avg}} \rangle} = \left\langle \frac{\sigma_{p_{\text{T}}}^{\text{probe}}}{p_{\text{T}}^{\text{avg}}} \right\rangle \oplus \left\langle \frac{\sigma_{p_{\text{T}}}^{\text{ref}}}{p_{\text{T}}^{\text{avg}}} \right\rangle \tag{3.11}$$

The value of N is also estimated by the random cones (RC) method. Circular regions corresponding to the area of anti-k_t jets with radius $R = 0.4$ are selected at random values of ϕ and opposite values of η. The difference in the energy scale of the constituents within these two areas $\Delta p_{\text{T}}^{\text{RC}}$ allows for the determination of the random fluctuation of energy deposits. This process is repeated for robustness, giving a value of the noise term due to pile-up

$$N^{PU} = \frac{\sigma_{\text{RC}}}{2\sqrt{2}}. \tag{3.12}$$

The electronic noise is then estimated from dedicated simulations without any pile-up contributions. This term is labelled $N^{\mu=0}$. The final noise contribution is given by the sum in quadrature of these terms $N = N^{PU} \oplus N^{\mu=0}$.

3.3.8 Pile-up jet suppression

Suppression of pile-up jets is fundamental to identifying jets in the final state stemming from the hard scattering rather than secondary interactions. ATLAS employs two algorithms which carry out this task: the jet vertex fraction (JVF) and jet vertex tagger (JVT) [120].

The jet vertex fraction is defined as the scalar sum of the p_{T} of tracks in a jet associated to the primary vertex PV_0 divided by the scalar sum of p_{T} of all tracks associated to the jet from vertices PV_n

$$\mathrm{JVF} = \frac{\sum_k p_{\mathrm{T}}^{\mathrm{trk},k}(\mathrm{PV}_0)}{\sum_l p_{\mathrm{T}}^{\mathrm{trk},l}(\mathrm{PV}_0) + \sum_{n\geq 1}\sum_l p_{\mathrm{T}}^{\mathrm{trk},l}(\mathrm{PV}_n)}. \tag{3.13}$$

The jet vertex tagger (JVT) is a more sophisticated algorithm which tags jets as originating from the primary vertex or from pile-up vertices. It is based on a two-dimensional likelihood discriminant obtained by considering two variables: $R_{p_{\mathrm{T}}}$ and corrJVF. $R_{p_{\mathrm{T}}}$ is defined as the scalar sum of the p_{T} of tracks associated to a jet originating from the primary vertex, divided by the jet's p_{T}

$$R_{p_{\mathrm{T}}} = \sum_{i\in jet} \frac{p_{\mathrm{T}i}(\mathrm{PV})}{p_{\mathrm{T}}^{\mathrm{jet}}}. \tag{3.14}$$

The corrJVF variable is related to the jet vertex fraction (JVF) which estimates the p_{T} fraction of a jet that is associated to the primary vertex. It is corrected to account for the number of reconstructed vertices in the an event, specifically the dependence on this quantity of the average scalar sum of p_{T} of tracks associated to a jet $\langle p_{\mathrm{T}}^{\mathrm{PU}}\rangle$. It is defined as

$$\mathrm{corrJVF} = \frac{\sum_k p_{\mathrm{T}}^{\mathrm{trk}^k}(\mathrm{PV}_0)}{\sum_l p_{\mathrm{T}}^{\mathrm{trk}^l}(\mathrm{PV}_0) + \sum_{n\geq 1}\frac{\sum_l p_{\mathrm{T}}^{\mathrm{trk}^l}(\mathrm{PV}_n)}{(k\cdot n_{\mathrm{trk}}^{\mathrm{PU}})}}, \tag{3.15}$$

where $k \cdot n_{\mathrm{trk}}^{\mathrm{PU}}$ is the factor which corrects for the linear increase of $\langle p_{\mathrm{T}}^{\mathrm{PU}}\rangle$ with the total number of tracks from pile-up vertices in an event $n_{\mathrm{trk}}^{\mathrm{PU}}$. The quantity k is just a scalar set to 0.01.

Chapter 4

Heavy flavour tagging in ATLAS

In order to identify final states containing heavy flavours, experiments such as ATLAS have developed flavour tagging techniques based on properties of heavy flavour jets which distinguish them from light jets. In Section 4.1, we describe these properties and the flavour tagging algorithm implemented by the ATLAS Collaboration in Run 2, DL1r.

The higher centre-of-mass energy available in Run 2 also necessitated the development of a dedicated boosted tagger to identify b-jets in topologies for which particles such as, but not limited to, the Higgs boson decay to b-quarks at a small angular distance from each other. The resulting b-jets are contained within a large-R jet of radius $R = 1.0$, and the tagging algorithm is applied directly to this jet. This tagger is known as the *Xbb tagger*, and is introduced in Section 4.2. The efficiency scale factor calibration of this tagger on multi-jet events, to which I personally contributed, is then described in Section 4.3.

4.1 Flavour tagging

This thesis focuses mainly on the substructure of heavy flavour jets. In order to study these jets at colliders, we must first be able to identify them in a given final state. Experimentally, this process is known as *flavour tagging*. This section focuses on the description of the ATLAS DL1r tagger and its performance.

4.1.1 Flavour tagging in a nutshell

Jets, as they are defined, are merely a collection of particles grouped together by an algorithm. In order to discern more information of the underlying physics which led to their formation, it is often useful to know which particle gave origin to the jet.

Flavour tagging algorithms are one of the first solutions devised to exploit substructure properties of jets to understand more about their formation. Specifically, the aim of flavour tagging is to understand whether jets originate from heavy quarks or light quarks/gluons. Jets falling in the latter case are collectively known as *light jets*.

Flavour tagging relies on the long lifetime of heavy flavour hadrons. Bottom hadrons in particular have the lifetime of the order of 1 ps, which, at the relativistic speeds at which they travel when produced at colliders, translates to a macroscopic distance traversed before the decay occurs, allowing for the identification of the decay vertex in the detector. Tracks originating from this vertex are characterised by a large transverse impact parameter as they are not aligned with

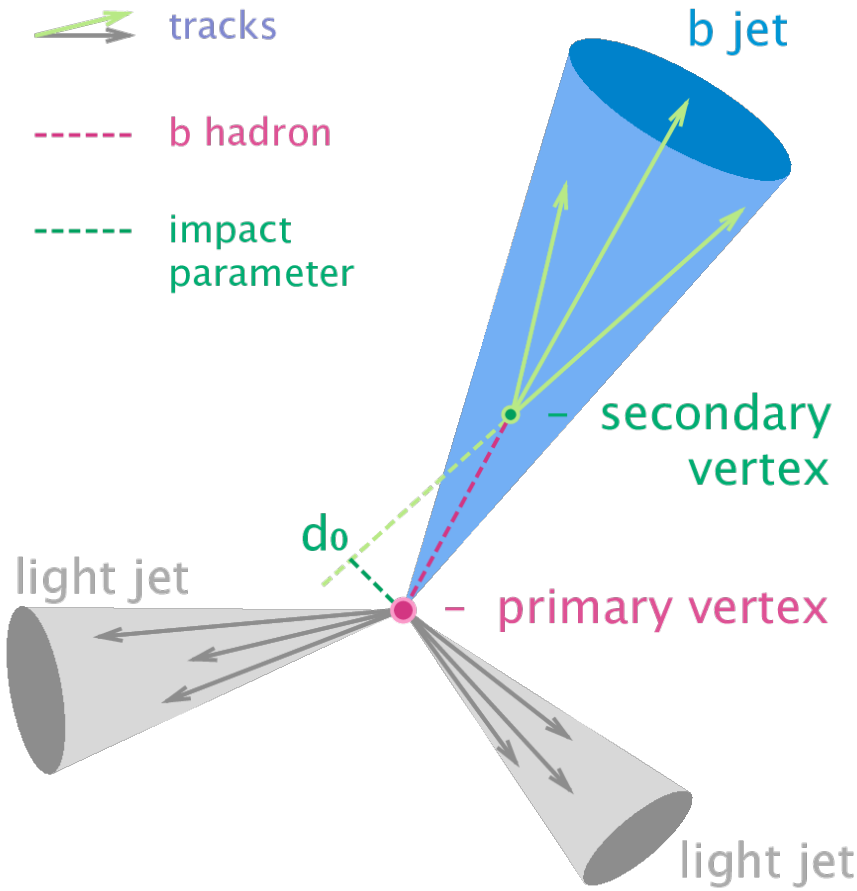


Figure 4.1: The properties of a secondary vertex from the decay of a b-hadron which allow for flavour tagging [121].

the primary vertex. A representation of the formation of the secondary vertex is shown in Figure 4.1.

Various other effects also contribute to the identification of heavy flavour jets. Focusing again on b-jets, the high charged particle multiplicity of the principal decay channels facilitates the identification of the secondary vertex. This vertex is characterised by a mass of the order of 5 GeV, or that of the b-hadron before its decay. Due to the leading particle effect, the heavy hadron (and consequently its decay products) also carries a significant fraction of the momentum of the jet. All these features are also present in c-jets, though they are not as pronounced as in b-jets, rendering the task of flavour tagging more difficult. Ideal light jets, on the other hand, lack these features. They consist mainly of tracks which point directly back to the primary vertex. In practice, however, many factors can lead to the misidentification of light jets as heavy flavour jets. Early flavour tagging algorithms, such as those used during Run 1 of ATLAS, exploited these geometric and topological properties and identified b-jets via a cut-based approach. The approach was not sophisticated enough for the identification of c-jets. The advent of machine learning brought about a revolution in the world of flavour tagging. Next generation taggers, such as the DL1 series, fed the output of several low-level algorithms and kinematic properties of jets

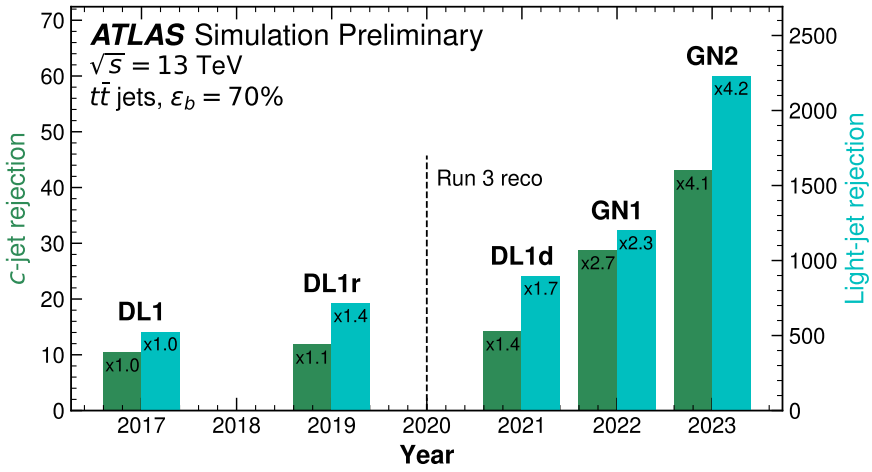


Figure 4.2: The improvement over Run 2 and Run 3 of the ATLAS flavour tagging algorithm, specifically different iterations of the DL1 series and the GN series. The tagger performance is evaluated on $t\bar{t}$ events at a 70% b-tagging efficiency [122].

to a neural network to tag jets. Tagging c-jets became possible at this point. The next revolution was brought about by the use of graph neural networks. The GN tagger series represents a jet as a graph, where each track is a node. Kinematic information of the track, together with low-level detector information such as the number of hits or holes in the layers of the inner detector, are used to train the algorithm. This strategy proved successful, and led to a marked increase in tagging performance for Run 3. The rejection of flavour tagging algorithms used by ATLAS for $t\bar{t}$ events with a 70% b-tagging working point in shown in Figure 4.2. Specifically, three iterations of the DL1 series is shown, together with the first two iterations of the GN series, with DL1 used as a reference point. The increase in tagging performance over time is striking, with a notable jump in the GN series.

4.1.2 DL1r

The flavour tagging algorithm used in the measurement which constitutes the bulk of this thesis is DL1r [123]. DL1r is part of the DL1 series, i.e. the second generation of flavour tagging algorithms used by ATLAS. It is based on the output of a high-level neural network which takes in several low-level tagging algorithms as well as kinematic information of the jet in question. With this information it decides whether a jet is b-, c-, or light. The low-level algorithms in question are described below.

IP2D & IP3D

IP2D and IP3D [124] are two algorithms based on the track impact parameter (IP). The former considers only the transverse IP (d_0), while the latter considers the longitudinal IP (z_0) as well. Specifically, the *signed IP significances* are used, where the sign refers to the position of a given track with respect to the axis of the jet that belongs to, with a positive (negative) sign referring to a track which is ahead (behind) the jet axis. The value of the signed IP is divided by the error of the value itself to obtain its significance. s_{d_0} and s_{z_0} are used to refer to the transverse and longitudinal IP significances, respectively.

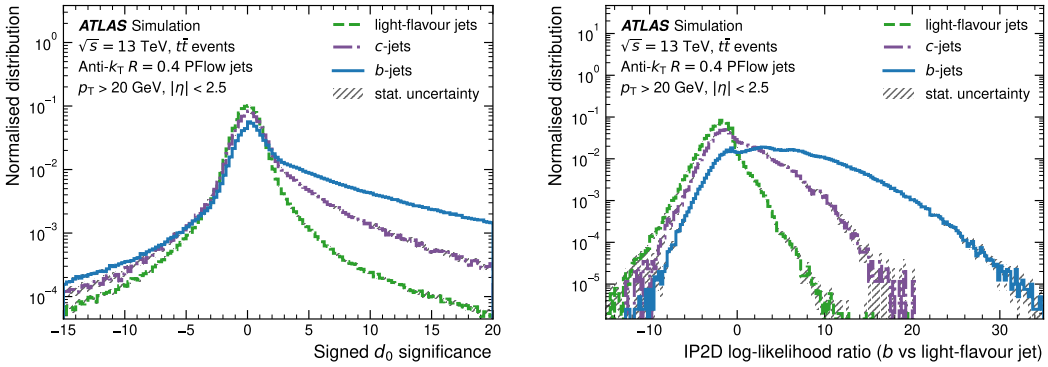


Figure 4.3: The signed transverse IP distribution for tracks from b-, c- and light jets coming from $t\bar{t}$ events (left), together with the log-likelihood ratio IP2D computed from it (right) [123].

Tracks coming from secondary vertices are characterised by large values of the signed IP significance as they are not aligned with the primary vertex of the event.

MC simulations of b- and c-jets from $t\bar{t}$ and Z' events are used to populate reference histograms of the signed IP significances, S_i. These are used to derive probability density functions, from which log-likelihood ratios are calculated to distinguish each category – b-, c-, and light jets – from each other. As an example, for the b- vs. light jet hypothesis, this ratio is

$$\sum_i p_b(S_i)/p_{\text{light}}(S_i), \tag{4.1}$$

where the sum runs on all tracks i in the jet. $p_b(S_i)$ and $p_{\text{light}}(S_i)$ represent the template probability density functions for the two hypotheses. Flavour probabilities are assumed to be independent.

The s_{d_0} value for tracks from b-, c- and light jets with $p_{\text{T}} > 20$ GeV from $t\bar{t}$ events is shown in Figure 4.3, together with the IP2D discriminant for the b-jet vs. light jet flavour hypothesis. The discrimination power of the variable is clearly seen.

RNNIP

The IP-based algorithms described above assume that each track's IP significance is independent from that of every other track. In practice, this is not the case. If a track with a large IP is identified from a secondary vertex stemming from a heavy flavour hadron decay, it is more likely to find at least one other track with a large IP significance within the jet. To account for this and to improve the modelling of the properties of heavy flavour jets, recurrent neural networks (RNNs) are used to learn the correlation between tracks in a given jet. The algorithm used is known as RNNIP [125].

A vector of track features, such as s_{d_0} and s_{z_0}, the p_{T} fraction carried by the track $p_{\text{T}}^{\text{frac}}$, the distance to the jet axis $\Delta R(\text{track}, \text{jet})$, and the hit multiplicity in the various ID layers, are used to train the RNN. The vector is ordered by decreasing values of $|s_{d_0}|$ to highlight the importance of this feature. The p_{T} spectra of the b-jet and c-jet distributions are also reweighted to that of the light jet distribution to avoid training the network on features specifically related to the sample considered or flavour features.

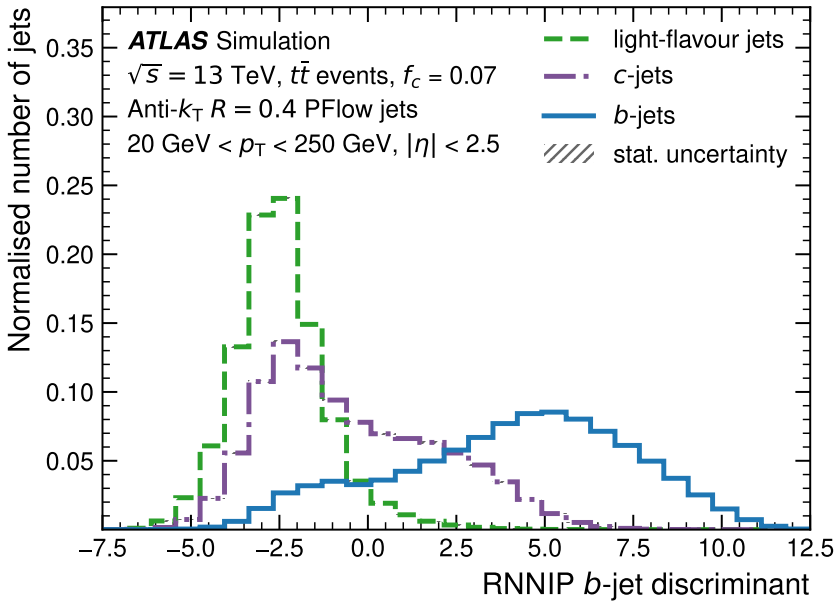


Figure 4.4: The RNNIP flavour tagging discriminant for b-, c- and light jets [123].

The output of RNNIP

$$D_{\text{RNNIP}} = \ln\left(\frac{p_b}{f_c \cdot p_c + (1 - f_c) \cdot p_{\text{light}}}\right) \tag{4.2}$$

is used as a discriminant, where f_c represents the fraction of c-jets in the sample, which is optimally determined to be $f_c = 0.07$. The distribution of the flavour tagging discriminant derived from the RNN for jets from a $t\bar{t}$ sample is shown in Figure 4.4.

SV1

Another way to identify heavy flavour jets is to explicitly reconstruct the secondary vertex. The secondary-vertex-tagging algorithm (SV1) [126] does just that. Using the twenty-five highest p_{T} tracks within the heavy flavour jet candidate, the algorithm identifies all possible two-track vertices. Vertices which are found to be compatible with particles such as the K_s and Λ are rejected, along with photon conversions and hadron conversions. The algorithm iteratively tries to fit a single secondary vertex from the remaining vertex candidates. For each iteration, a χ^2 test is used to evaluate the track-to-vertex matching for all tracks. The track with the largest χ^2 is then removed from the fit. This process is repeated until a vertex with a small enough χ^2 and mass below 6 GeV is found.

JetFitter

JetFitter [127] is a multi-vertex secondary vertex finder which makes use of the topological structure of heavy hadron decays to reconstruct a b-hadron's full decay chain. Using a Kalman filter, the algorithm identifies a line on which the b-hadron and subsequent c-hadron decays occur,

allowing for the reconstruction of the b-hadron's trajectory and the vertex positions with even just a single track from each decay. The discrimination of c-jets from b-jets and light-flavour jets is improved by specifically targeting jets that have a single reconstructed secondary vertex at a distance similar to the b-hadron decay vertex in b-jets and with intermediate charged decay multiplicity.

DL1r

The outputs of the of the low-level algorithms described above are combined with the p_{T} and $|\eta|$ of the jet to train a high-level neural network, known as DL1(r). The difference between the two lies in the inclusion of the RNNIP output, which is present in DL1r but not in DL1. Table 4.1 provides an overview of all variables included in the neural network.

DL1(r) is trained on a sample consisting in 70% of jets from a $t\bar{t}$ production process, while the remaining portion consists of jets from a $Z' \to q\bar{q}$ sample, where the Z' corresponds to a heavy resonance. The Z' is set to decay to b-quarks, c-quarks, and light quarks with a branching ratio set to 1/3 for each decay channel. The two samples are used to generate jets across a wide p_{T} spectrum, with the Z' sample being included specifically to increase the number of jets with $p_{\mathrm{T}} > 250$ GeV. The p_{T} spectrum of the jets used for training is reweighted in such a way that the final p_{T} distribution is uniform in order to avoid biases. Jets can be of different types, such as PFlow jets of radius $R = 0.4$ or VR trackjets.

The resulting output discriminant from the tagger, D_{DL1r}, is defined analogously to Eq. 4.2. In this case, however, the c-jet fraction is set to $f_c = 0.018$. f_b is instead set to 0.2.

Figure 4.5 shows the output of DL1r for b-jets, c-jets, and light PFlow jets in the p_{T} range $20 < p_{\mathrm{T}} < 250$ GeV. Clear discrimination of b-jets is visible. Figure 4.6 shows the rejection of c- and light jets as a function of b-tagging efficiency for DL1r compared to that for DL1 and MV2c10 [128], a tagging algorithm based on a Boosted Decision Tree used by ATLAS during the early years of Run 2. A marked improvement in both the c-jet rejection and light jet rejection can be seen.

To use the DL1r tagger, several working points are calibrated. These are the 60%, 70%, 77%, and 85% working points, which correspond to a b-jet acceptance of the same percentage. Lower working points are characterised by a higher c-jet and light jet rejection, resulting in a higher purity of b-jets. The c-jet and light jet rejections for each working point are shown in Table 4.2.

Performance on VR Trackjets The performances of the DL1r tagger described above are valid for $R = 0.4$ PFlow jets. The tagger is also trained on other jets in use by ATLAS, including VR trackjets. We briefly describe the performance on these jets, as they are relevant for Section 4.2 and the measurement in Chapter 5.

The DL1r tagger for VR trackjets is calibrated to the same working points as for PFlow jets. The value of the c-jet fraction in the training sample remains $f_c = 0.30$. Table 4.3 shows the c-jet and light jet rejections for each working point.

4.2 The Xbb tagger

The Xbb tagger is a tool optimised for the identification of $X \to b\bar{b}$ final states in boosted topologies. Here, X is typically a Higgs boson, though the tool is designed in such a way as to tag boosted $b\bar{b}$ pairs arising from a generic particle [129].

The tool is developed as a neural network trained on the large-R jet kinematics and on information stemming from VR trackjets used to resolve the individual b-hadron decays. Jets from the AntiKt10LCTopoTrimmedPtFrac5SmallR20Jets collection are used for training. These

Input	Variable	Description
Kinematics	p_T	Jet p_T
	η	Jet $\|\eta\|$
IP2D, IP3D	$\log(P_b/P_{light})$	Likelihood ratio of the b-jet to light-flavour jet hypotheses
	$\log(P_b/P_c)$	Likelihood ratio of the b-jet to c-jet hypotheses
	$\log(P_c/P_{light})$	Likelihood ratio of the c-jet to light-flavour jet hypotheses
RNNIP	P_b	b-jet probability
	P_c	c-jet probability
	P_{light}	light-flavour jet probability
SV1	m(SV)	Invariant mass of tracks at the secondary vertex assuming pion mass
	f_E(SV)	Jet energy fraction of the tracks associated with the secondary vertex
	N_{TrkAtVtx}(SV)	Number of tracks used in the secondary vertex
	N_{2TrkVtx}(SV)	Number of two-track vertex candidates
	L_{xy}(SV)	Transverse distance between the primary and secondary vertices
	L_{xyz}(SV)	Distance between the primary and secondary vertices
	S_{xyz}(SV)	Distance between the primary and secondary vertices divided by its uncertainty
	$\Delta R(\vec{p}_{\text{jet}}, \vec{p}_{\text{vtx}})$(SV)	ΔR between the jet axis and the direction of the secondary vertex relative to the primary vertex.
JetFitter	m(JF)	Invariant mass of tracks from displaced vertices
	f_E(JF)	Jet energy fraction of the tracks associated with the displaced vertices
	$\Delta R(\vec{p}_{\text{jet}}, \vec{p}_{\text{vtx}})$(JF)	ΔR between the jet axis and the vectorial sum of momenta of all tracks attached to displaced vertices
	S_{xyz}(JF)	Significance of the average distance between PV and displaced vertices
	N_{TrkAtVtx}(JF)	Number of tracks from multi-prong displaced vertices
	N_{2TrkVtx}(JF)	Number of two-track vertex candidates (prior to decay chain fit)
	$N_{\text{1-trk vertices}}$(JF)	Number of single-prong displaced vertices
	$N_{\geq\text{2-trk vertices}}$(JF)	Number of multi-prong displaced vertices
	$L_{xyz}(2^{\text{nd}})$(JF)	Distance of 2nd vertex from PV
	$L_{xy}(2^{\text{nd}})$(JF)	Transverse displacement of the 2nd vertex
	$m_{\text{Trk}}(2^{\text{nd}})$(JF)	Invariant mass of tracks associated with the 2nd vertex
	$E(2^{\text{nd}})$(JF)	Energy of the tracks associated with the 2nd vertex
	$f_E(2^{\text{nd}})$(JF)	Jet energy fraction of the tracks associated with the 2nd vertex
	$N_{\text{TrkAtVtx}}(2^{\text{nd}})$(JF)	Number of tracks associated with the 2nd vertex
	$\eta_{\text{trk}}^{\text{min,max,avg}}(2^{\text{nd}})$(JF)	Min., max. and avg. pseudorapidity of tracks at the 2nd vertex

Table 4.1: The input variables used to train the DL1r tagging algorithm [123]

Working point	c-jet rejection	Light jet rejection
60%	29	1155
70%	10	417
77%	5	170
85%	2	38

Table 4.2: The c-jet and light jet rejections for the DL1r tagger trained on PFlow jets for each calibrated working point.

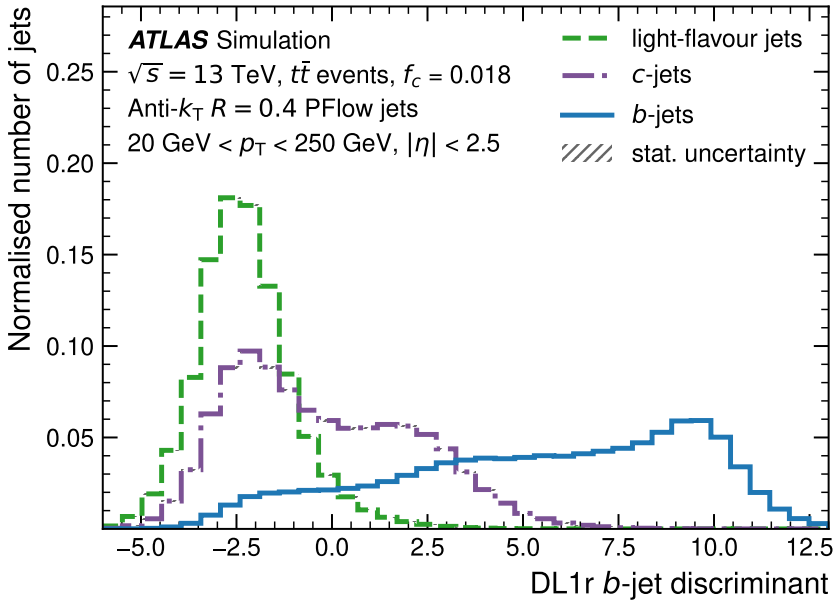


Figure 4.5: The DL1r discriminant for b-jets, c-jets, and light jets [123].

jets must have a p_{T} in the range $250 < p_{\mathrm{T}} < 3000$ GeV and $|\eta| < 2.0$. VR trackjets are ghost-associated[1] to the large-R jet. They must have $p_{\mathrm{T}} > 7$ GeV, $|\eta| < 2.5$ and at least two track constituents.

Several input variables are used to train the Xbb tagger. These include the probability of a given trackjet associated to the large-R jet of being a b-, c-, and light jet (p_b, p_c, p_{light}) for up to three (if available) VR trackjets. Each individual probability is decided by applying the DL1r tagger on the trackjet. The list of training variables is concluded with kinematic information of the large-R jet, such as its p_{T} and η.

Jets from $H \to b\bar{b}$, $Z' \to t\bar{t}$ and multi-jet samples are used to train the Xbb tagger. The $H \to b\bar{b}$ sample is treated as signal, while the others are treated as background. In the signal sample, a Higgs boson is required to be ghost-matched to the large-R jet. The three MC samples used for training are built in such a way as to have identical p_{T} distributions and a flat mass distribution, in order to avoid bias from the constraint of the Higgs mass resonance.

The Xbb tagger output is a discriminant which gives the likelihood of the large-R jet coming

[1] This procedure is described in Section 2.1.

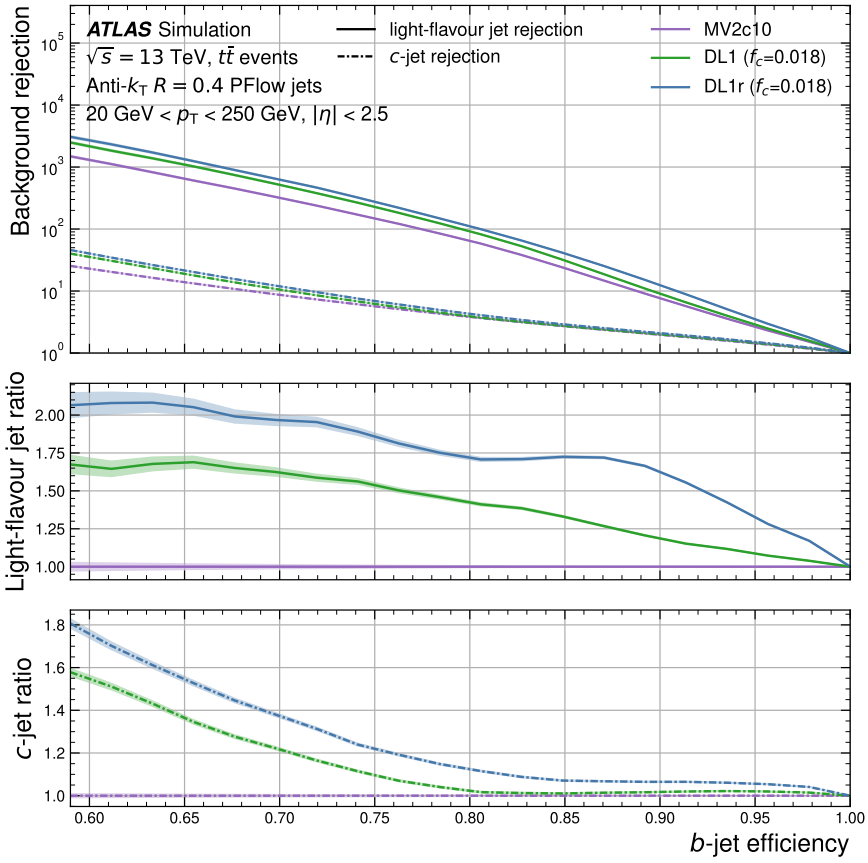


Figure 4.6: The c- and light jet rejections vs. b-tagging efficiencies of the DL1r, DL1, and MV2c10 taggers [123].

Working point	c-jet rejection	Light jet rejection
60%	28	1140
70%	10	400
77%	5	148
85%	2	39

Table 4.3: The c-jet and light jet rejections for the DL1r tagger trained on VR trackjets for each calibrated working point.

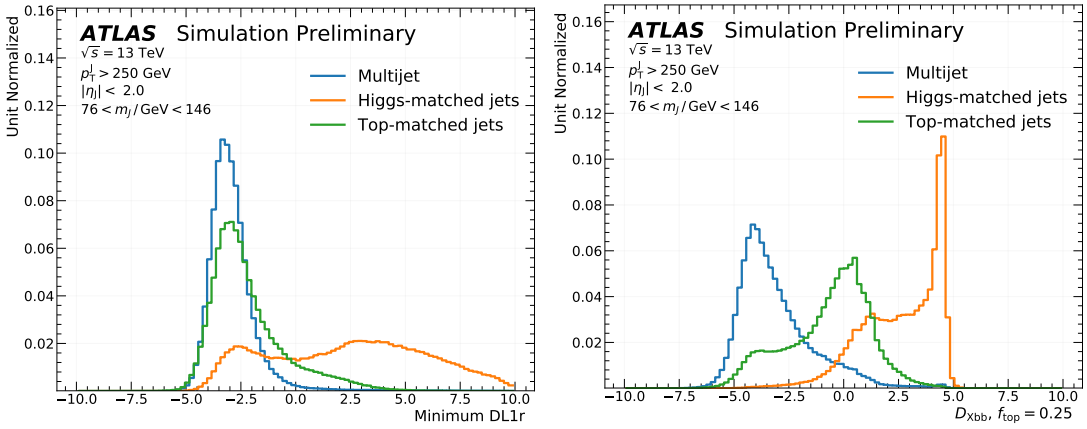


Figure 4.7: The Xbb discriminant distribution for Higgs, top and multi-jet samples when both VR trackjets are required to pass a minimum DL1r threshold (left) vs. for the dedicated Xbb tagger (right) [130].

from a Higgs boson decaying to a pair of b-quarks vs. $t\bar{t}$ or multi-jet. In particular, the probability of each outcome is quantified in the values p_{Higgs}, p_{top}, and $p_{\text{multi-jet}}$, respectively. The tagger output is then defined analogously to Eq. 4.2:

$$D_{Xbb} = \ln\left(\frac{p_{\text{Higgs}}}{f_{\text{top}} \cdot p_{\text{top}} + (1 - f_{\text{top}}) \cdot p_{\text{multi-jet}}}\right), \tag{4.3}$$

where $f_{\text{top}} = 0.25$. Figure 4.7 shows the D_{Xbb} output for jets with $p_{\text{T}} > 250$ GeV and mass $76 < m_J/\text{GeV} < 146$ from Higgs, top, and multi-jet samples compared to the output obtained by directly tagging the two leading VR trackjets with DL1r.

Large-R jets are considered tagged when either the D_{Xbb} discriminant or the minimum DL1r/MV2c discriminant for the two leading VR trackjets surpasses a given threshold. Three working points are defined: 50%, 60%, and 70%, corresponding to appropriate cuts on the discriminants leading to an efficiency in the signal selection equal to the same percentage.

At the 60% working point, for jets with $p_{\text{T}} > 500$ GeV, the multi-jet and top rejections are found to be 92 and 31, respectively, for the D_{Xbb} discriminant. The multi-jet rejection when double-b-tagging with DL1r is comparable, though the D_{Xbb} discriminant shows an improvement of a factor of approximately 1.6 when comparing the top rejections. In general, it is found that the performance of the Xbb tagger when compared to other taggers improves as jet p_{T} increases. This is due to the more boosted topology, allowing the large-R jet to fully contain more of the decay and the fact the topology is more distinct at higher p_{T}, compared to the resolved topologies used to train the other taggers.

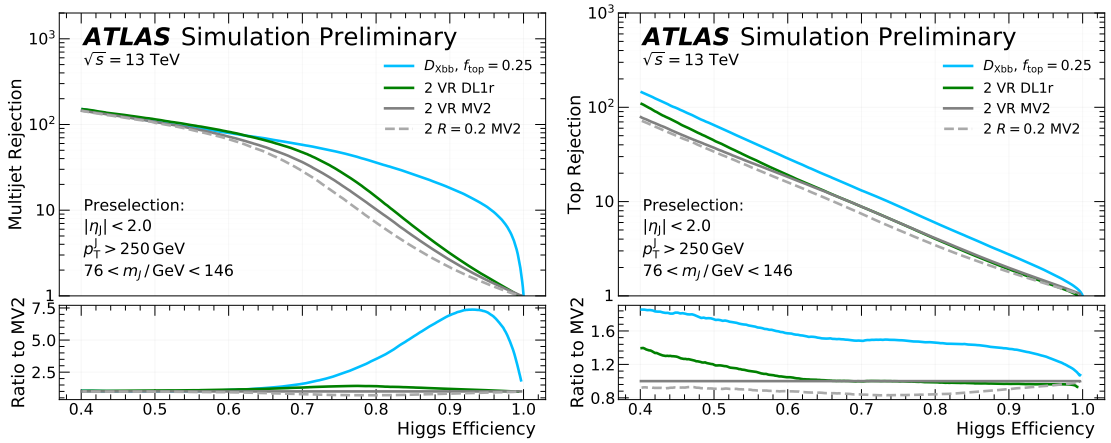


Figure 4.8: The multi-jet (left) and $t\bar{t}$ (right) rejection as a function of signal efficiency of the Xbb tagger compared to other taggers, such as DL1r and MV2 applied to VR and fixed-radius trackjets [130].

Figure 4.8 shows the Xbb tagger rejection of multi-jet and top jets as a function of signal efficiency for large-R jets with $p_{\mathrm{T}} > 250$ GeV. The rejection is compared to double-b-tagging the VR trackjets with the DL1r and MV2 algorithms, as well as double-b-tagging trackjets with a fixed radius of $R = 0.2$ with MV2. The rejections as a function of jet p_{T} for a fixed 60% WP is then shown in Figure 4.9.

4.3 Calibration

As the Xbb tagger is trained exclusively on simulated MC events, it must be calibrated to be used in physics analyses. The calibration ensures that the efficiency of the tagger on signal and background processes matches between data and simulation. This work constitutes my ATLAS Qualification Project.

The calibration is carried out on four final states. The signal efficiency $X \to b\bar{b}$ is calibrated on $Z(\to b\bar{b}) + \text{jets}$ and $Z(\to b\bar{b})\gamma$ events, while the top and multi-jet backgrounds are calibrated using $t\bar{t}$ and $g \to b\bar{b}$ events. My work focused on the $g \to b\bar{b}$ calibration, and is described below. For details regarding the other calibrations, see Reference [131].

4.3.1 $g \to b\bar{b}$ calibration

The calibration focuses on di-jet events where the presence of b-hadrons in large-R jets is enhanced in data by requiring a soft muon within the jet area, signalling a heavy flavour semi-leptonic decay. The selected sample is then compared to simulation using a multicomponent template fit based on an observable sensitive to the $g \to b\bar{b}$ vs. $g \to q\bar{q}$ separation, where q indicates a light-flavour quark.

The full Run 2 dataset is used for the calibration. At the MC level, $2 \to 2$ QCD interactions are simulated at LO using Pythia8 [132]. The parton shower is then simulated with Pythia8 and EvtGen [133] using the A14 tune [134] and the NNPDF2.3 LO PDF set [135]. An inclusive sample is simulated along with a sample in which the events are required to contain a muon at truth level. This muon is required to have $p_{\mathrm{T}} > 3$ GeV and $|\eta| < 2.8$.

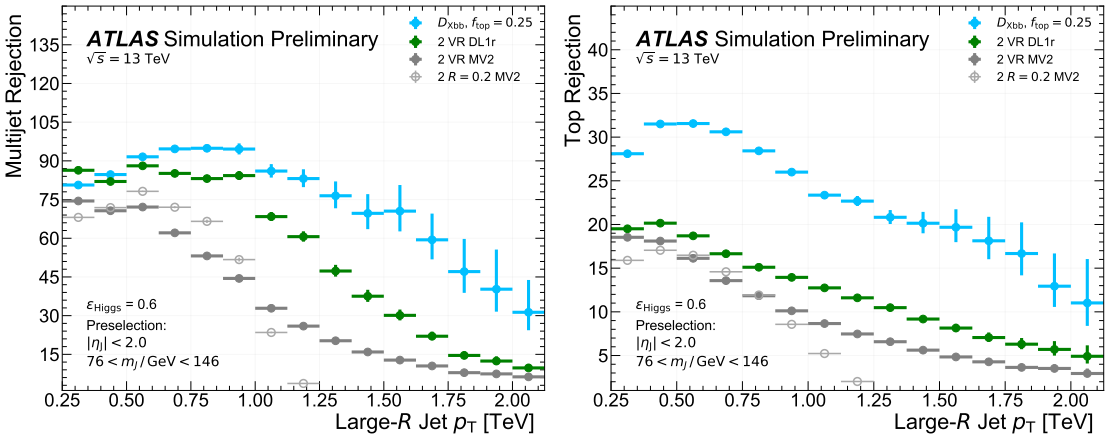


Figure 4.9: The multi-jet (left) and $t\bar{t}$ rejection as a function of large-R jet p_{T} of the Xbb tagger compared to other taggers, such as DL1r and MV2 applied to VR and fixed-radius trackjets [130].

Objects are selected as described above: a large-R jet satisfying the selection criteria of the Xbb tagger (described in Section 4.2) is identified, and up to three VR trackjets are ghost-matched to the jet, henceforth referred to as *subjets*. One subjet is required to contain a muon matching the appropriate selection criteria to preferentially favour jets containing heavy flavours. This subjet is known as the *muon jet*. If more than one subjet contains a muon, the one with the highest p_{T} is called the muon jet.

Each large-R jet is assigned a label based on the subjets it contains. As up to three subjets are considered, the large-R jet label contains three letters corresponding to the truth-level flavour content of the subjets found within it. The order of the large-R jet label corresponds to the p_{T} ordering of the subjets.

Subjets are labelled as follows: the label B is assigned if a b-hadron is ghost-matched to the subjet, C if a c-hadron is ghost-matched, L if no heavy hadrons are ghost-matched, and x if a subjet of that index is not present in the large-R jet.

The flavour labels are used to determine the flavour category of the large-R jet. The MC-to-data fit is carried out in each flavour category of the large-R jet. To reduce the number of large-R jet flavour categories considered, a simplified labelling scheme which groups together different three-flavour categories is applied to the large-R jet, as shown in Table 4.4. The large-R jet can be labelled as BB, BL, CC, CL or LL.

4.3.2 Efficiency scale factor fit strategy

Given the importance and complexity of the fit procedure for the extraction of the $g \to b\bar{b}$ scale factors, the methodology is described below.

After labelling the large-R jet, a template of a flavour-sensitive observable is made for fitting MC to data. For this calibration, the mean signed d_0 significance ($\langle s_{d_0} \rangle$) is chosen, where the mean is calculated on the three highest p_{T} tracks associated to the subjet in question. The average is specifically chosen to reduce the influence of outliers in light-flavour jets, such as mismodelled tracks or K_s decays.

The fit is carried out in several different p_{T} regions of the large-R jet: below 280 GeV, 280-310 GeV, 310-340 GeV, 340-380 GeV, 380-420 GeV, 420-460 GeV, 460-500 GeV, 500-550 GeV, 550-600 GeV, 600-750 GeV, 750-1000 GeV, and above 1000 GeV.

Large-R jet flavour category	**Subjet flavours**	**Description of flavour category**
BB	BBB, BBC, BBL, BBx, BCB, BLB, CBB, LBB	At least two trackjets contain b-hadrons
BL	BCC, BCL, BCx, CBC, CBL, CBx, CCB, BLC, BLL, BLx, LBC, LBL, LBx, CLB, LCB, LLB, Bxx	Exactly one trackjet contains b-hadrons
CC	CCC, CCL, CCx, CLC, LCC	At least two trackjets contain c-hadrons, and none contain b-hadrons
CL	CLL, CLx, LCL, LCx, LLC, Cxx	Exactly one trackjet contains c-hadrons, and none contain b-hadrons
LL	LLL, LLx, Lxx	No trackjet contains b- or c-hadrons

Table 4.4: The large-R jet flavour categories and corresponding subjet flavours. For the subjet labelling scheme, B, C, and L indicate a subjet containing a b-hadron, c-hadron, or light hadron, respectively, while x indicates a missing subjet. The subjet labelling scheme follows the p_{T} ordering of the subjets.

In each p_{T} bin, four sets of templates are considered in the fit. These are derived using two observables - namely the mean s_{d_0} of the muon jet and the leading non-muon jet - within two regions, i.e. a double-b-tagged (with DL1r) and a non-double-b-tagged large-R jet. A 60% b-tagging working point is chosen when defining these regions, and the large-R jet templates are labelled as defined in Table 4.4. The four sets of templates are fit simultaneously in each p_{T} bin, though the fits in the various p_{T} bins are independent.

The fit method chosen is a maximum likelihood fit. The aim is to extract the best-fit value of a correction factor $\vec{f}$, i.e. the vector of correction factors for each flavour category. This is known as the efficiency scale factor (SF), as is defined as

$$\mathrm{SF} = \frac{\epsilon_{\mathrm{data}}}{\epsilon_{\mathrm{MC}}}, \tag{4.4}$$

where ϵ indicates the tagging efficiency in data or MC. Since the efficiency is defined as the ratio of events which pass the selection over the total number of events $\epsilon = N^{\mathrm{pass}}/N^{\mathrm{total}}$, the SF can be equivalently written as the ratio of signal strength post- and pre-tagging

$$\mathrm{SF} = \mu_{\text{post-tag}}/\mu_{\text{pre-tag}}. \tag{4.5}$$

The flavour corrections are allowed to float freely, and are extracted from the fit. The efficiency scale factor in the tag region $\mu_{\mathrm{tag}} = \epsilon_{\mathrm{data}}/\epsilon_{\mathrm{MC}}$ corresponds to the flavour correction in the BB template in the tagged region. The anti-tag scale factor $\mu_{\text{anti-tag}} = \frac{1-\epsilon_{\mathrm{data}}}{1-\epsilon_{\mathrm{MC}}} = \frac{1-\mu_{\mathrm{tag}}\epsilon_{\mathrm{MC}}}{1-\epsilon_{\mathrm{MC}}}$ instead represents the difference in tagging rejection rates. The tag and anti-tag regions of the fit are thus required to derive the respective efficiency scale factors. The muon jet and non-muon jet are both considered as the shape distribution of the mean s_{d_0} is uncorrelated in these two regions, thus allowing for the fit to be carried out simultaneously by multiplying the relative likelihoods.

In each histogram bin i, the expectation value for the number of entries $E[n_i]$ is given by

$$E[n_i] = \prod_{xx} f_{xx} y_{xx,i} = \vec{f} \cdot \vec{y_i}, \tag{4.6}$$

where $y_{xx,i}$ is the nominal number of entries in the bin and f_{xx} is a correction factor applied to the flavour template xx. The template distributions depend on nuisance parameters (NPs) $\vec{\theta}$ with prior probability distributions determined from auxiliary measurements. The priors are assumed to follow a split-normal distribution $\mathcal{SN}$ with a width on each side given by a variation histogram. The likelihood function for the overall number of entries in each bin and in each template is thus given by

$$\mathcal{L}(\vec{f}, \vec{\theta}) = \prod_{i=1}^{N} e^{-(\vec{f}\cdot\vec{y_i})} \frac{(\vec{f}\cdot\vec{y_i})^{n_i}}{n_i!} \prod_{k} \mathcal{SN}(\theta_k), \tag{4.7}$$

where the vector of NPs is taken to be of length k. NPs are taken to be systematic uncertainties, such as modelling uncertainties and those related to the detectorc reconstruction.

A smoothing procedure is also applied in order to reduce statistical fluctuations which may impede the convergence of the fit. This is achieved by merging bins in which these fluctuations are present. Bins are merged when the statistical uncertainty is greater than the systematic uncertainty.

For a given bin with nominal entries N and systematic variation S, with relative errors δN and δS, the total error is given by $\delta M = \sqrt{\delta S^2 + \delta N^2}$, for uncorrelated variations, or $\delta M = \max(\delta N, \delta S)$ for correlated variations. If for a bin $|S - N| < \delta M$, this bin is merged with

the neighbouring bin with the highest value of $\delta M/M$. This is the rebinning step in the template definition.

To further reduce the statistical fluctuations of the systematic variations, for each pair of neighbouring bins b_j and b_{j-1} the value

$$X_{j-1,j} = \left| \frac{S_j - N_j}{N_j} - \frac{S_{j-1} - N_{j-1}}{N_{j-1}} \right| \tag{4.8}$$

and relative error

$$\delta X_{j-1,j} = \sqrt{\frac{\delta M_j^2}{N_j^2} + \frac{\delta M_{j-1}^2}{N_{j-1}^2}} \tag{4.9}$$

is calculated. If for some pair $\delta X < X$, the bins in the pair are merged. The systematic variations calculated in the merged bin, and propagated back to the original binning. The end result is an ease of tension in the fit.

Figure 4.10 shows the pre-fit templates considered for the fit, for the large-R jet p_{T} slice between 600-750 GeV. The post-fit agreement is shown in Figure 4.11.

The scale factors extracted from the fit in each p_{T} slice are shown in Figure 4.12. They are entirely below unity, indicating a lower di-jet $g \to b\bar{b}$ tagging efficiency in data. As indicated by the small uncertainties on the scale factor, some p_{T} bins are also highly constrained.

4.3.3 Systematic uncertainties

A number of systematic uncertainties come into play when fitting the MC simulation to data. These uncertainties also represent NPs for the fit.

Uncertainties fall within three main categories: experimental uncertainties related to the reconstruction of objects used for the selection, such as jets, tracks, and muons; theoretical uncertainties related to MC modelling; and uncertainties specific to the $g \to b\bar{b}$ calibration. These include uncertainties on the fragmentation of different species of b-hadrons with different lifetimes or on processes that produce tracks with large values of s_{d_0}, such as the production of K_s or Λ.

Besides the systematic uncertainties, extrapolation uncertainties are also considered as NPs. These consider the selection bias imposed by requiring a soft muon in the jet and by deriving the scale factors from $g \to b\bar{b}$ events. In general, users of the Xbb tagger would not require a soft muon within the jet, and would not apply the tagger to multi-jet events, but to a general $b\bar{b}$ resonance. The extrapolation uncertainties therefore account for any possible differences in the scale factors derived from this specific process and selection used for the calibration.

The categorisation of systematics is illustrated in full in Table 4.5. A full list of systematics considered is available in Table 4.6. This table also states whether a given systematic is applied to the template shape, normalisation, or both.

Figure 4.13 shows the pulls derived from the NPs in the template fit, in the p_{T} bin 600-750 GeV. These are defined as the difference between the post-fit ($\hat{\theta}$) and pre-fit (θ_0) values of a quantity in units of standard deviation $\Delta\theta$

$$\text{pull} = \frac{\hat{\theta} - \theta_0}{\Delta\theta}. \tag{4.10}$$

From the plot, the systematics with the most significant pulls are those related to the MC uncertainties on the definitions of the individual flavour templates. Some individual systematics related to the fake rate of tracks, initial/final state radiation, and jet modelling compose of the rest of the main contribution.

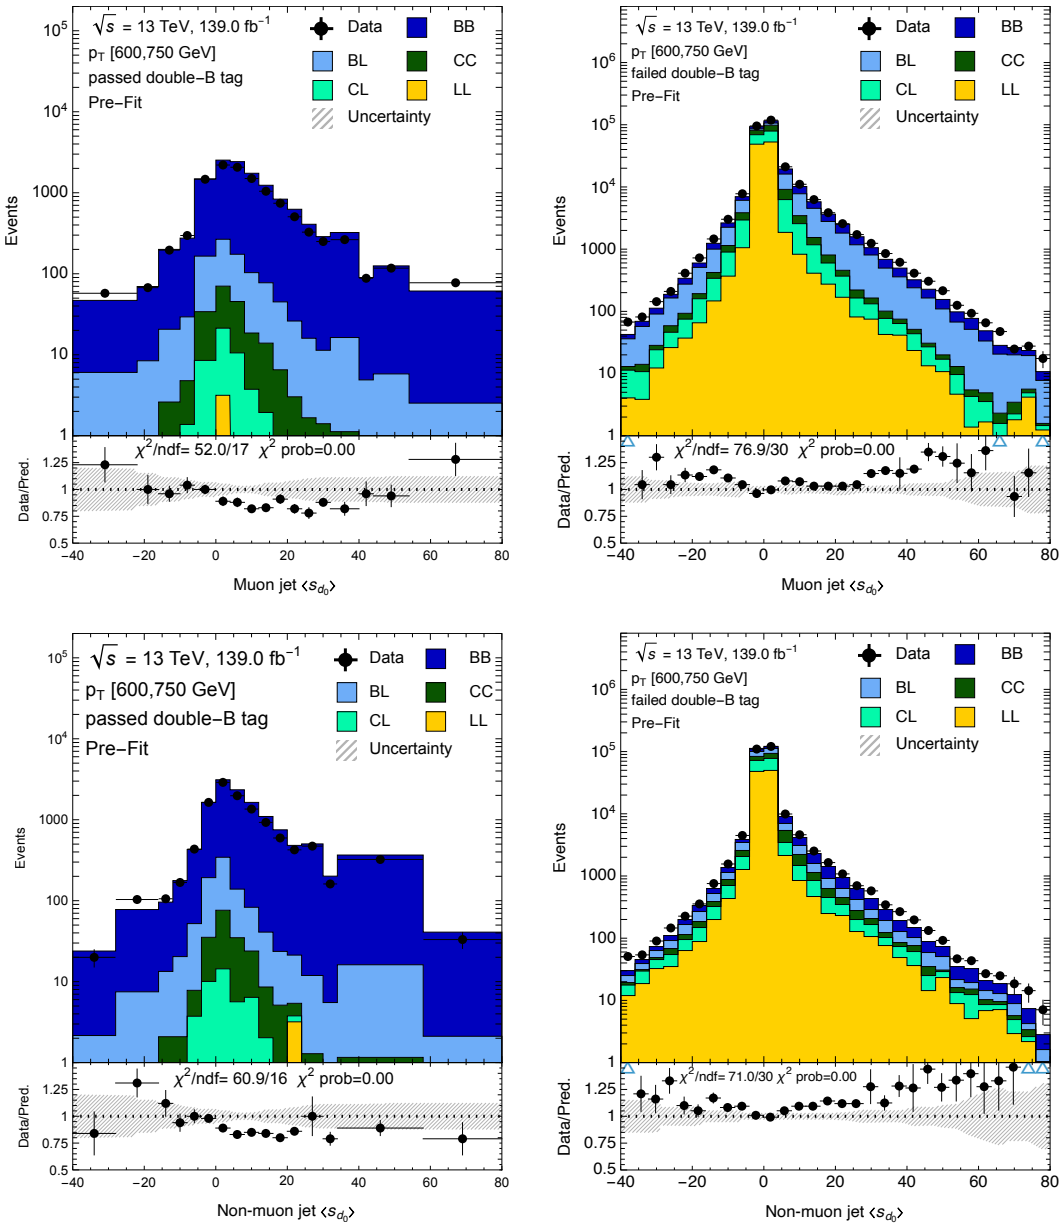


Figure 4.10: The pre-fit agreement in the four sets of templates used for the $g \to b\bar{b}$ calibration. The top row corresponds to the muon jet and the bottom row to the leading non-muon jet. The left column shows the double-b-tagged large-R jet and the right column the anti-tagged jet.

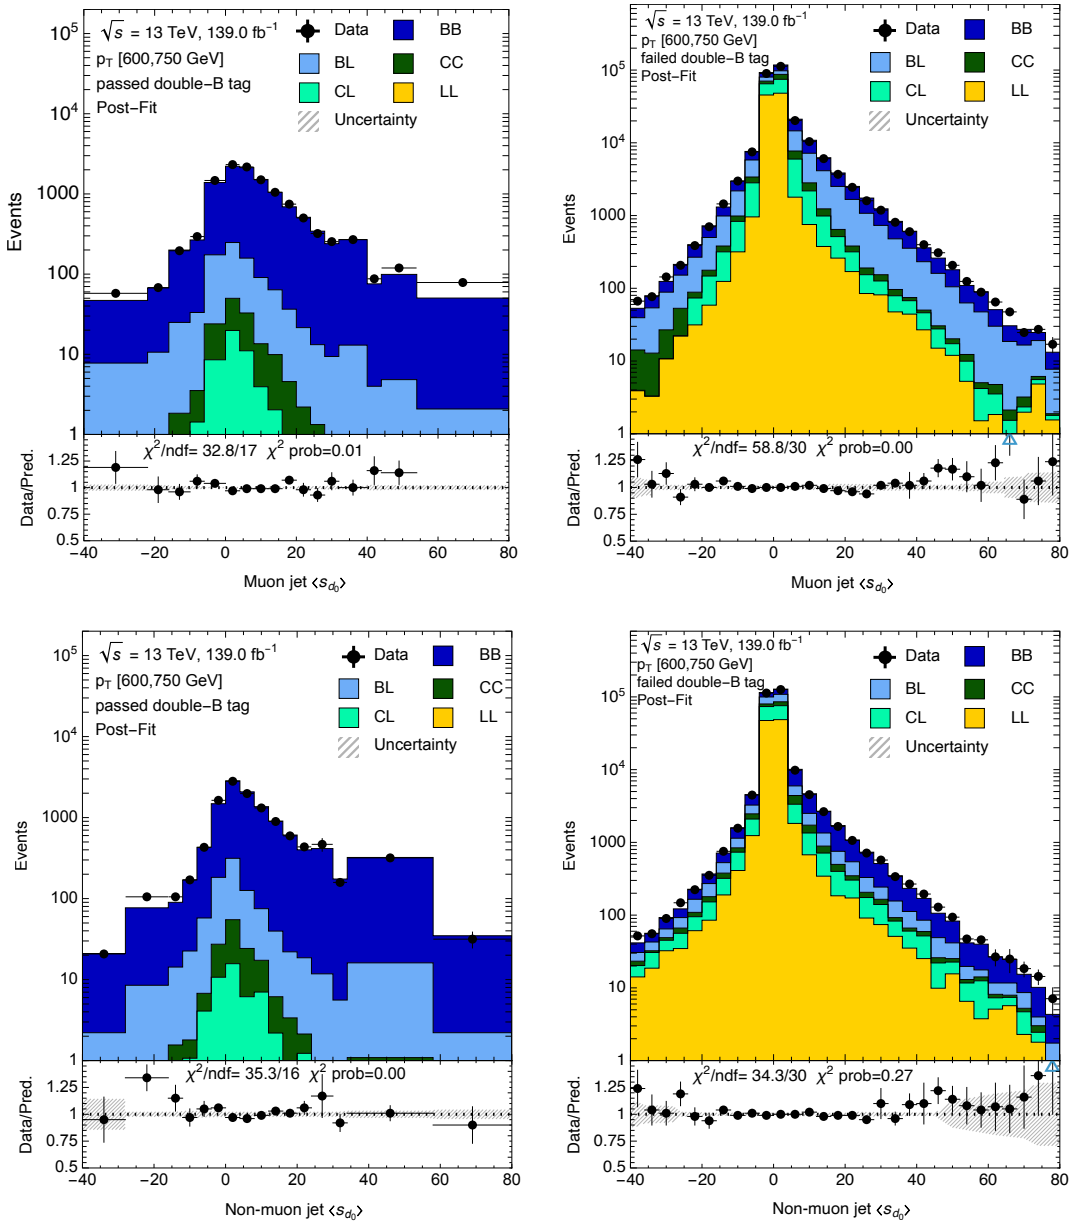


Figure 4.11: The post-fit agreement in the four sets of templates used for the $g \to b\bar{b}$ calibration. The top row corresponds to the muon jet and the bottom row to the leading non-muon jet. The left column shows the double-b-tagged large-R jet and the right column the anti-tagged jet.

Systematic Uncertainty	Brief Description
Template Uncertainties	
Fake secondary vertices	Uncertainty on the rate of processes which create large s_{d_0} tracks can in light jets.
Fake muons	Uncertainty on the rate of false-positive muon identification in b-jets.
b-hadron fractions	Uncertainty on the relative production rate of b-hadron species.
Muon Requirement	Uncertainty in the inclusive phase-space from the muon requirement used to derive scale factors.
$g \to bb$ to $X \to bb$	Uncertainty on extrapolating from $g \to bb$ decays to the general $X \to bb$ case.
Experimental Uncertainties	
Luminosity	Uncertainty on the full Run 2 integrated luminosity.
Pile-up Reweighting	Uncertainties on pile-up conditions applied when reweighting simulations to match data.
Jet Energy Scale (JES)	Uncertainty on the reconstruction of large-R jet energies from detector inputs. Applied as 30 independent NPs.
Jet Energy Resolution (JER)	Uncertainty on the precision of jet energy reconstruction.
Jet Mass Scale (JMS)	Uncertainty on jet mass reconstruction. Calculated separately from JES and applied as 6 independent NPs.
Jet Mass Resolution (JMR)	Uncertainty on the precision of jet mass reconstruction. Separate NPs used for Higgs jets and top jets.
Muon Reconstruction Efficiency	Uncertainties on the muon reconstruction efficiency and track-to-vertex association [113].
Muon Momentum Scale	Uncertainty on muon momentum reconstruction [113]. Includes separate uncertainties on the resolution of ID and MS tracks.
Sagitta Bias Correction	Uncertainties due to charge-dependent effects of detector misalignment [113].
Track reconstruction efficiency	Uncertainties on passive material in the ID and on the GEANT4 model used in simulation.
Track fake rate	Uncertainty on the rate of combinatorial fake tracks from large numbers of hits in the ID.
Track impact parameter resolution	Uncertainties based on the difference in d_0 and z_0 resolution between data and MC.
Theoretical Uncertainties	
Parton Shower	Uncertainty in the parton shower model is measured by comparing PYTHIA 8 and HERWIG 7.
Renormalization Scale	Uncertainties in renormalization and factorization scales, and in final state radiation (FSR) are assessed by sample weight variations in PYTHIA 8.

Table 4.5: Uncertainties on the derivation of b-tagging scale factors for the $g \to b\bar{b}$ calibration of the Xbb tagger.

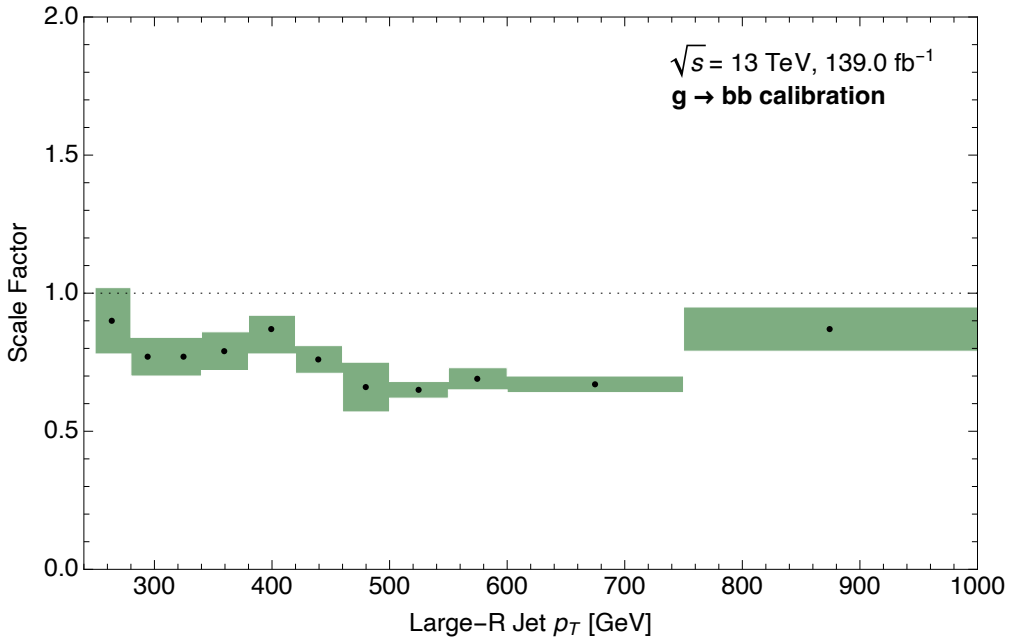


Figure 4.12: The scale factors in each p_{T} sliced extracted from the $g \to b\bar{b}$ fit.

The correlation matrix of the NPs is also shown in Figure 4.14. As can be seen, the most highly correlated NPs are those related the major systematics described above and the normalisation factors of the various flavour templates used for the fit.

4.3.4 SV1 mass fit

To validate the results obtained via the $\langle s_{d_0} \rangle$ template fit, the efficiency scale factors can be derived with a different set of templates constructed with an alternative flavour-sensitive variable to demonstrate the compatibility of the scale factors derived with the two methods. For this purpose, the secondary vertex (SV1) mass is chosen as the variable used to define the templates.

The SV1 mass corresponds to the mass of the tracks coming from the secondary vertex, as found by the vertex finder described in Section 4.1.2. To reiterate, if it were possible to associate all particles coming from a heavy hadron decay to the secondary vertex, this would be found to have a mass corresponding to that of the heavy hadron, of the order of 5 GeV for b-hadrons. In practice, only charged particles can be associated to the secondary vertex, and there is always the possibility of fake tracks and inefficiencies, so the reconstructed vertex tends to have values below 5 GeV.

Aside from the choice of variable, the details of the fit are otherwise identical. We will thus limit ourselves to the results, with relevant comments when necessary.

In Figures 4.15 and 4.16, we see the pre-fit and post-fit plot for the templates defined with the SV1 mass variable. As compared to the mean s_{d_0}, we can see how the SV1 mass distribution is overall smoother, missing the significant peak centred at zero which was present in the s_{d_0} fit. Some flavour templates, particularly LL in the non-muon jet in the fail region, do portray peaks, however.

The fit converges and the scale factors obtained in all large-R jet p_{T} regions is shown in Figure 4.17. These are shown to be compatible with those shown in Figure 4.12. They are

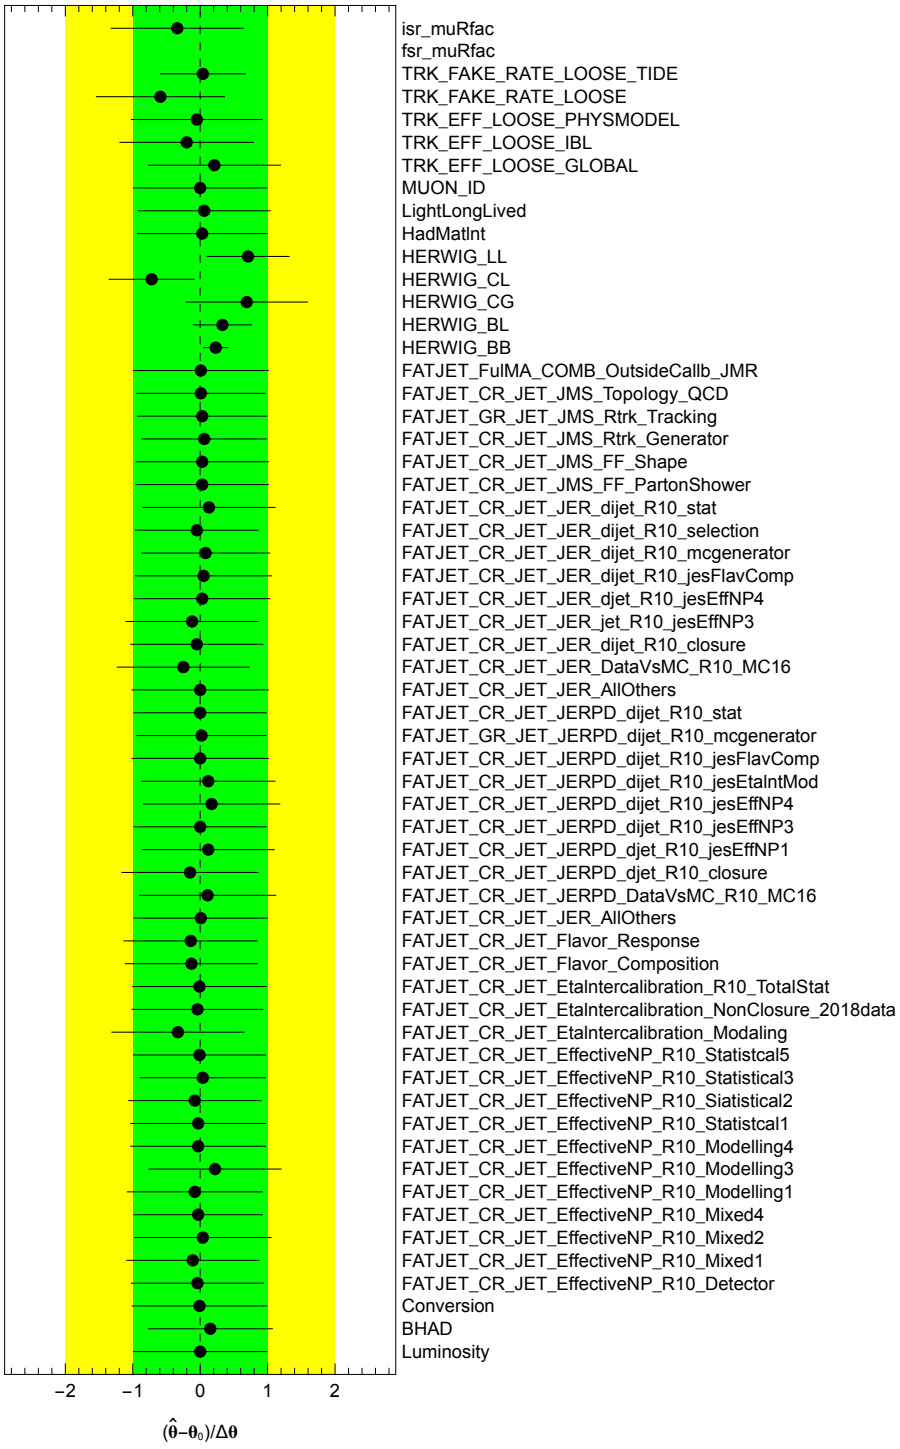


Figure 4.13: The pulls of the systematic variations affecting the template fit, in the large-R jet p_{T} region between $600 < p_{\mathrm{T}}/\mathrm{GeV} < 750$.

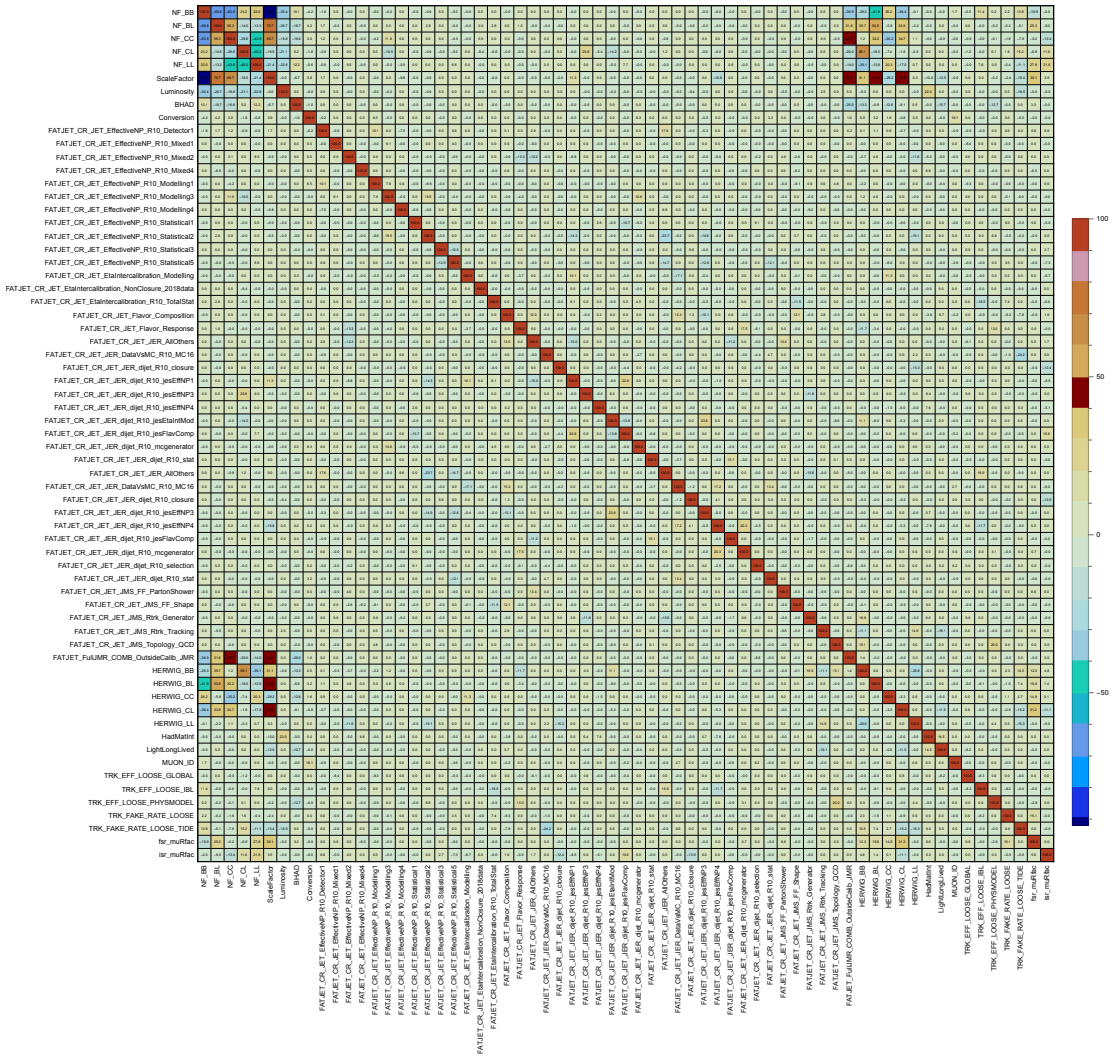


Figure 4.14: The correlation matrix of the NPs from the template fit, in the large-R jet p_{T} region between $600 < p_{\mathrm{T}}/\mathrm{GeV} < 750$.

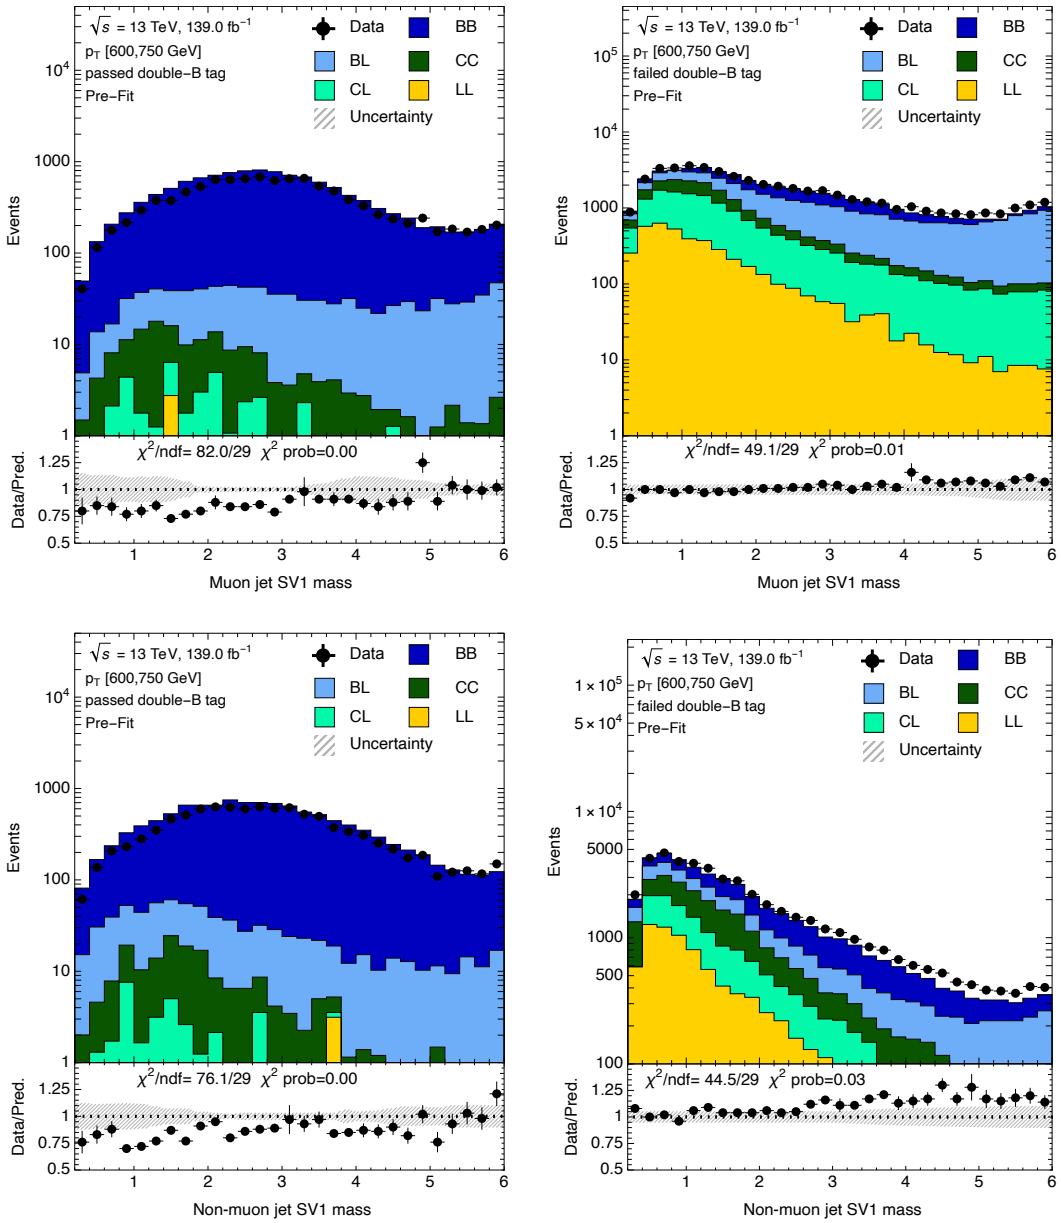


Figure 4.15: The pre-fit agreement in the four sets of templates used for the $g \to b\bar{b}$ calibration with the SV1 mass variable. The top row corresponds to the muon jet and the bottom row to the leading non-muon jet. The left column shows the double-b-tagged large-R jet and the right column the anti-tagged jet.

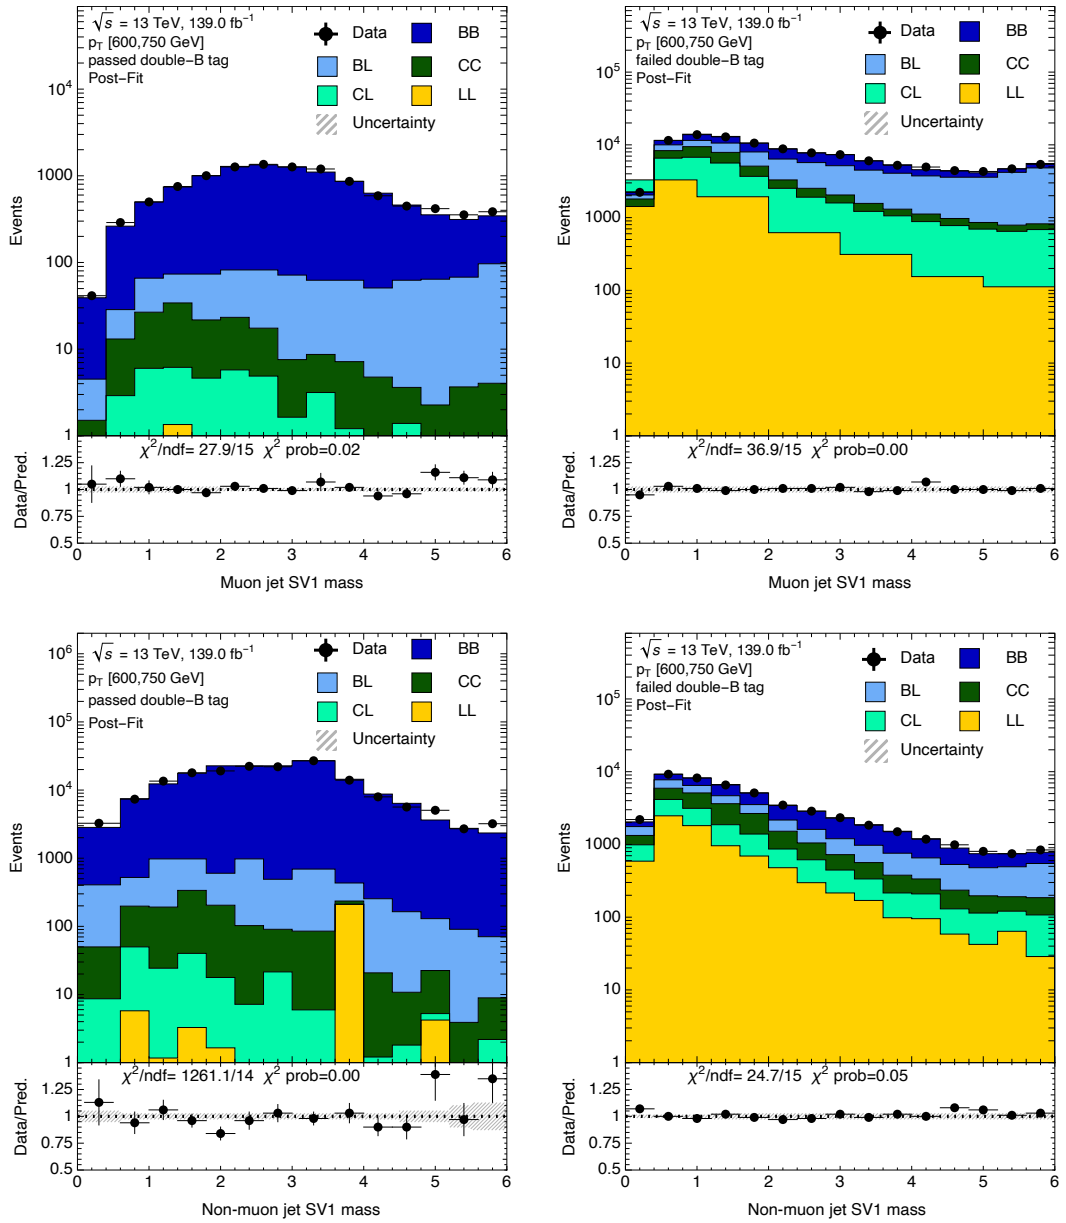


Figure 4.16: The post-fit agreement in the four sets of templates used for the $g \to b\bar{b}$ calibration with the SV1 mass variable. The top row corresponds to the muon jet and the bottom row to the leading non-muon jet. The left column shows the double-b-tagged large-R jet and the right column the anti-tagged jet.

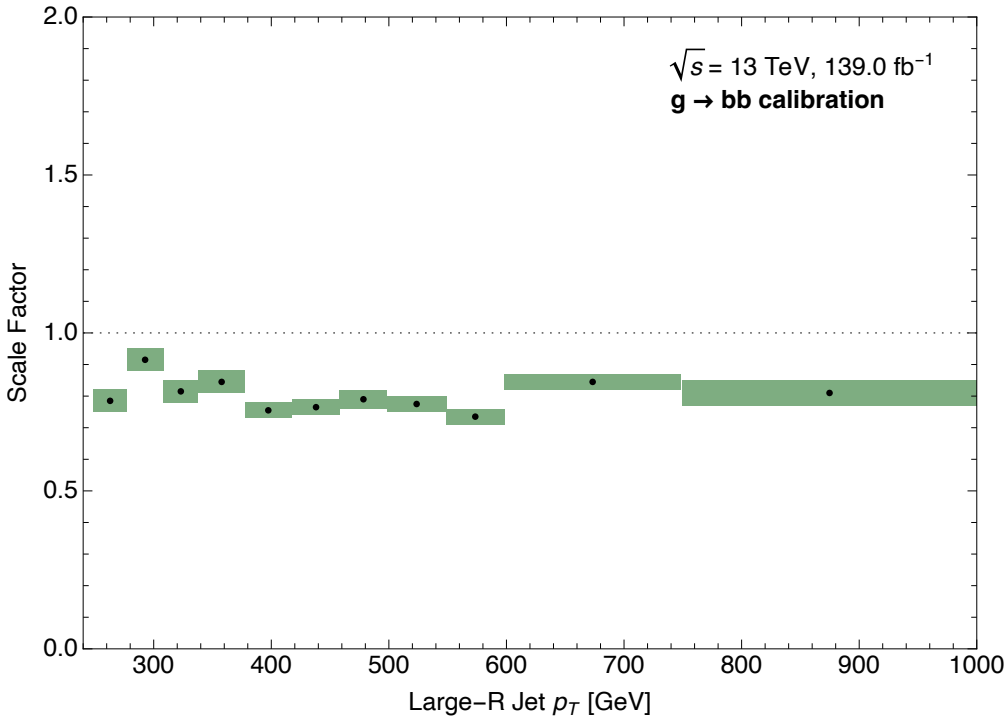


Figure 4.17: The scale factors in each p_{T} sliced extracted from the $g \to b\bar{b}$ fit with the SV1 mass.

all below unity, and generally compatible within uncertainties. Figure 4.18 shows the scale factors derived by both the mean s_{d_0} fit and the SV1 mass fit. These are shown to be generally compatible within uncertainties, with a few exceptions such as for the 280-310 GeV p_{T} bin, the 500-550 GeV p_{T} bin, and the 600-750 GeV p_{T} bin.

Figure 4.19 shows the pulls for the NPs of the SV1 mass fit. These are the same as those found for the mean s_{d_0} fit.

4.3.5 Final comments

All in all, the Xbb tagger calibration of $g \to b\bar{b}$ events was successfully completed. The convergence of the fit was proven and results were corroborated with alternate variables.

Some final studies remained on the effect of the smoothing procedure on the fit, and, especially in the SV1 mass fit, the optimisation of the binning. Additional studies, such as an optimisation of the flavour template definition, were also carried out, but did not find their way into this thesis.

4.4 Conclusions

In this chapter we have outlined the flavour tagging methods implemented by the ATLAS Experiment. Starting from a general description of flavour tagging, we moved on to describe the low-level algorithms which feed into the DL1r tagger, to then show the performances of this tagger. We then described the Xbb tagger, a dedicated $b\bar{b}$ resonance tagger, optimised for the Higgs boson. Finally, we described the calibration of this tagger on multi-jet events, showing how

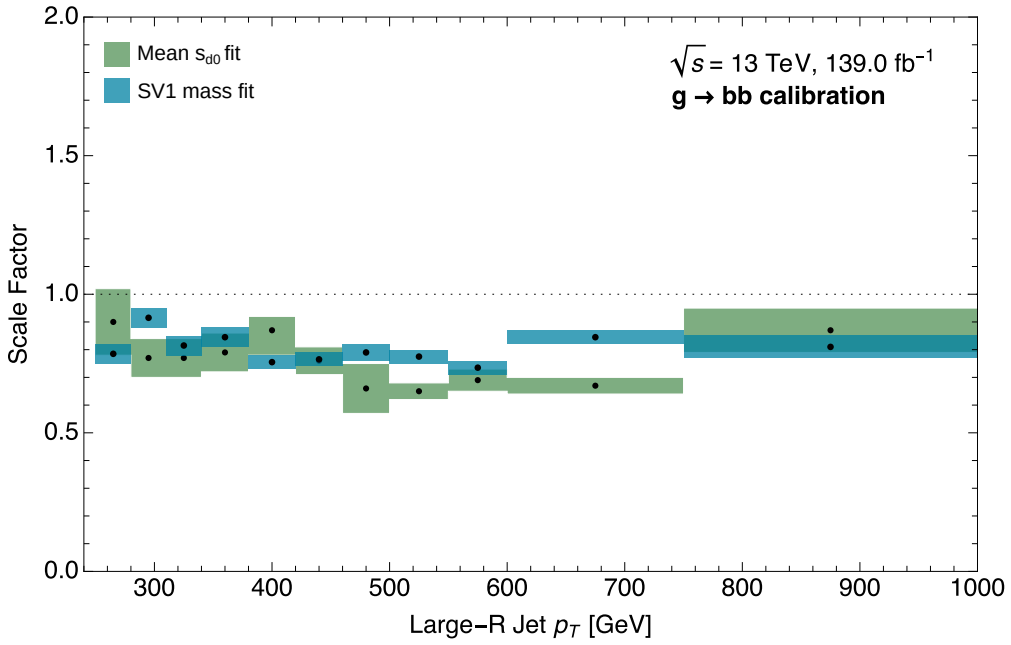


Figure 4.18: The scale factors in each p_{T} sliced extracted from the $g \to b\bar{b}$ fit with templates defined with mean s_{d_0} compared to those obtained with templates defined with the SV1 mass.

the maximum likelihood template fits based on two different flavour-sensitive observables, the mean s_{d_0} and SV1 mass, can provide the efficiency scale factors needed to match the performance of the tagger on data to that on simulation.

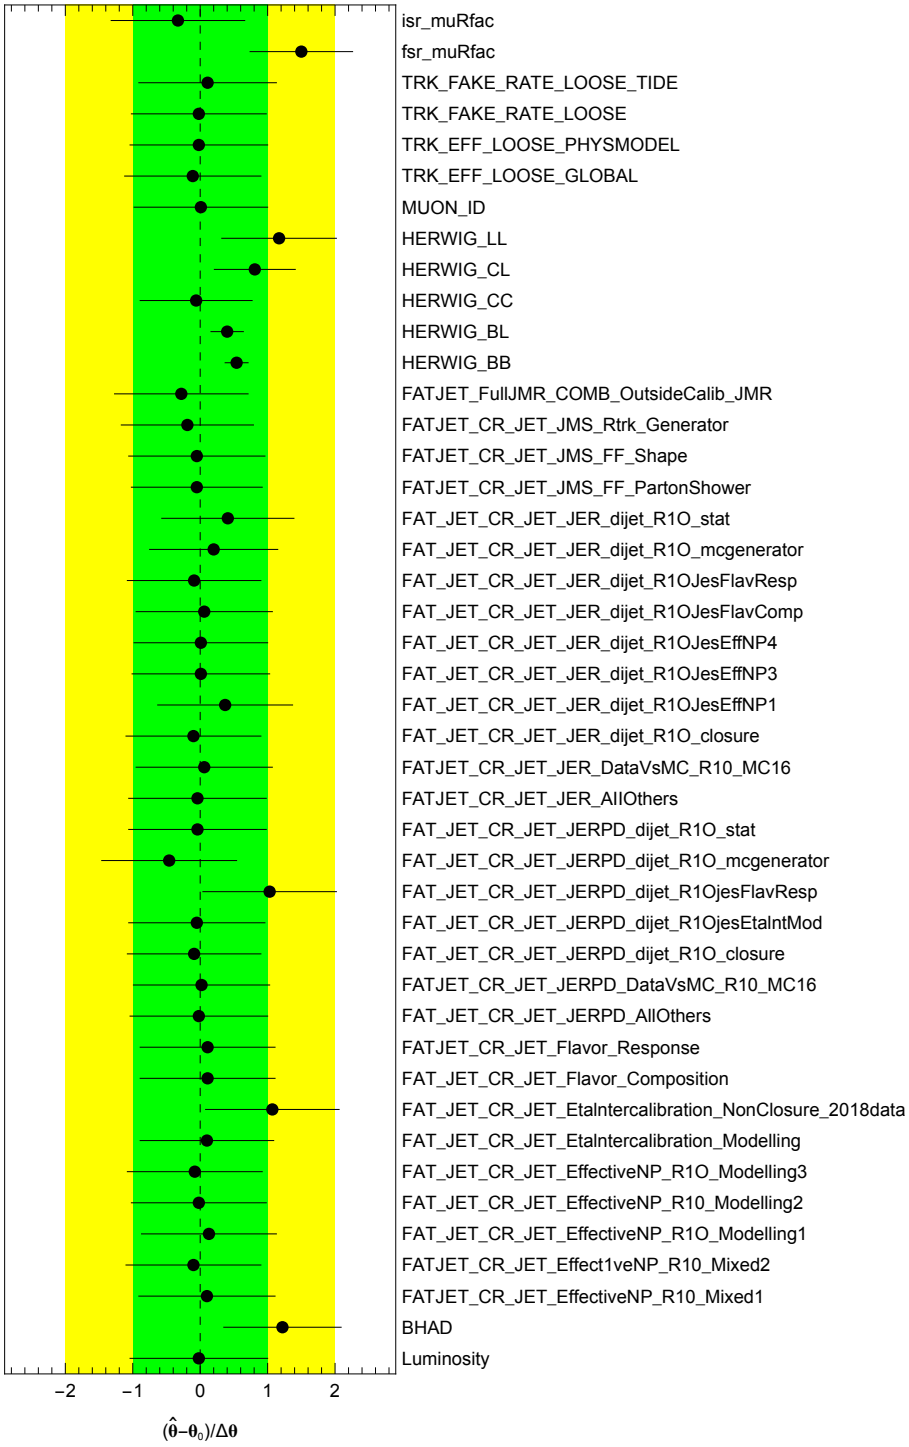


Figure 4.19: The pulls of the systematic variations affecting the SV1 mass template fit, in the large-R jet p_{T} region between $600 < p_{\mathrm{T}}/\mathrm{GeV} < 750$.

Name	Regions applied	Norm/Shape
JET_EffectiveNP_R10_GrestTerm JET_EffectiveNP_R10_Pi4 JET_EffectiveNP_R10_2 JET_EffectiveNP_R10_3 JET_EffectiveNP_R10_4 JET_EffectiveNP_R10_5	all templates, fully correlated	shape only
JET_EtaIntercalib_Modelling JET_EtaIntercalib_NonClosure_2018data JET_EtaIntercalib_R10_TotalStat	all templates, fully correlated	shape only
JET_Flavour_Composition	all templates, fully correlated	shape only
MUON_ID MUON_SAGITTA_RESBIAS MUON_SCALE MUON_EFF_TTVA_SYS	all templates, fully correlated	shape only
TRK_RES_P2_MEAS TRK_RES_P2_ERR TRK_RES_D0_DEAD TRK_RES_Z0_DEAD	all templates, fully correlated	shape only
TRK_EFF_LOOSE_GLOBAL TRK_EFF_LOOSE_T1 TRK_EFF_LOOSE_PHYSMODEL TRK_EFF_LOOSE_PP0 TRK_EFF_LOOSE_TIDE TRK_FAKE_RATE_LOOSE_PP0 TRK_FAKE_RATE_LOOSE_TIDE	all templates, fully correlated	shape only
isr_muRfac_65__fsr_muRfac_1 isr_muRfac_1__fsr_muRfac_0925 isr_muRfac_2__fsr_muRfac_1 isr_muRfac_1__fsr_muRfac_2	all templates, fully correlated	shape only
Conversion	all templates, fully correlated	normalisation
HadMat	all templates, fully correlated	normalisation
LightLongLived	electron/tau flavour categories, only for BB and BL	normalisation
BHAD	all templates, fully correlated	shape only
Herwig	Merging categories and flavours, only for BB and BL	shape only

Table 4.6: Names of the uncertainties included in the final template fit with the corresponding values and usage.

Chapter 5

Heavy flavour and jet substructure measurements in Z+jets ATLAS data

So far in this work, we have explored various aspects of heavy flavour jets. The journey began in Chapter 2, where we have seen how flavour is assigned to particle-level jets in predictions, highlighted issues these methods face at NNLO, and studied new methods of obtaining flavoured jets in an IRC-safe manner for predictions and how these may be applied to an experimental context.

We then turned our attention to the identification of flavour at experiment, reviewing the features of heavy flavour jets and the techniques enacted by the ATLAS Experiment to discern their appearance in the detector in both boosted and resolved topologies. This was followed by a detailed exposé of the calibration of the boosted Xbb tagger, revealing how the experimental techniques are tuned in order to be faithfully applied to data.

We now pivot to an application of the experimental technique, consisting in a measurement of Z+jet final states with heavy flavour jets, and of the resulting jet substructure. The analysis focuses on boosted and resolved topologies of both $Z + bb$ and $Z + jj$ final states to study jets in a wide kinematic range with distinct internal structures and radii.

This chapter introduces the measurement in detail. First, a motivation for the measurement and introduction to the analysis strategy are provided in Section 5.1. To follow, in Section 5.2 we define the observables whose cross sections we aim to measure.

The data and Monte Carlo samples used in the measurement and predictions are described in Section 5.3, while the event selections in the signal regions (SRs) considered are defined in Section 5.4 and in Section 5.5. Subsequently, in Section 5.6, we describe the reconstruction of the jet substructure observables, namely the Lund Jet Plane and jet angularities, focusing in particular on the b-hadron reconstruction strategy.

We proceed to detail the background suppression strategies in Section 5.7, and outline the systematics which affect the measurement in Section 5.8. Detector-level comparisons of the data and MC predictions are then provided in Section 5.9.

We conclude by delineating the unfolding strategy in Section 5.10, and show the latest results in Section 5.11.

5.1 Motivation and analysis outline

Z+jets production represents one of the premier tests of QCD, as experimentally it is a readily identifiable final state thanks to the Z boson tag, and each jet produced is a direct result of an emission. This provides the ability to test perturbative QCD in a wide kinematic range.

Considering Z+heavy flavour jet final states further allows one to probe the structure of the proton and test predictions involving heavy flavour production, whose resummation is especially complicated by the presence of the mass term. $Z + b\bar{b}$ in particular allows one to investigate gluon splittings to heavy quark pairs. This also represents one of the main backgrounds for Higgs measurements which focus on ZH production, as the $H \to b\bar{b}$ decay channel is the most probable, representing roughly 58% of the branching ratio for a Higgs mass of 125 GeV [4].

This measurement aims to study the properties of the internal structure of b-jets, highlighting in particular the differences when compared to light jets. Currently, little is known about the radiation pattern within these jets, and, in fact, this thesis describes the first ever measurement of b-jet substructure at the ATLAS Experiment, and the first ever measurement of light jet substructure across a wide p_{T} range.

The analysis strategy has been developed to facilitate the achievement of these lofty aims. This begins with the choice of the Z+jet final state: besides serving as a useful tag thanks to the easily identifiable leptonic decay channels of the Z boson leading to a simplified trigger selection, the final state is characterised by a low non-QCD background. The presence of two jets – either bb or jj – in the final states further reduces the background contamination. Jets are either b-tagged or anti-b-tagged, i.e. are required to fail the b-tag selection, with a tight 70% working point, further enriching the purity of the signal selection.

The four signal regions of the measurement are listed below:

- **Resolved 2B:** A resolved final state with $Z+ \geq 2$ b-jets. Jets considered are of radius $R = 0.4$ and reconstructed with PFlow objects.

- **Resolved 0B:** A resolved final state with $Z+ \geq 2$ j, where j is an anti-tagged light jet, which fails the b-tag selection. The jets considered are of radius $R = 0.4$ and reconstructed with PFlow objects.

- **Boosted 2B:** A boosted final state characterised by a large-R jet of radius $R = 1.0$ containing at least two b-tagged VR trackjets.

- **Boosted 0B:** A boosted final state characterised by a large-R jet of radius $R = 1.0$ containing at least two anti-b-tagged VR trackjets.

The main backgrounds remaining in the four signal regions after these initial selections consist in $t\bar{t}$ events and other Z+jets events. In the Resolved 2B SR, the $t\bar{t}$ background is significantly reduced by the application of a dedicated neural network developed for signal/background discrimination. In all four SRs, the $t\bar{t}$ background contribution is then estimated via a data-driven method reliant on the definition of a semi-leptonic $e^{\pm}\mu^{\mp}$ control region.

Other Z+jets backgrounds are estimated through a profile-likelihood fit based on flavour templates. These are defined based on the DL1r tagger output score for both jets, binned into discrete quantiles based on the tagging efficiency.

Finally, measured events are corrected for detector effects such as reconstruction efficiency and resolution, in order to obtain particle-level cross section measurements.

5.2 Observables of interest

As shown in Section 1.5, kinematic observables for Z+jets have been measured several times by multiple experiments. In particular, an ATLAS measurement from 2024 already focuses on the differential cross section of kinematic observables such as the p_{T} and di-jet mass using an event selection which exactly matches our Resolved 2B final state [53]. Similarly, a another ATLAS measurement using partial Run 2 data observes a final state similar to our Boosted 2B SR [54].

For this reason, this measurement focuses particularly on jet substructure observables, such as the Lund Jet Plane and the jet angularities. The emphasis on jet substructure for heavy flavour jets and for low p_{T} light jets is novel in ATLAS. This also represents the first time that the Lund Jet Plane for large-R jets in a final state enriched by gluon splittings[1].

Some kinematic variables, however, are still measured. This is done to update previous measurements, such that cited in Reference [54], to the full Run 2 dataset, but also to include novel observables which were not previously considered.

We specify all observables measured by the signal regions in which they are measured below:

Resolved 2B:

- Transverse momentum of the leading b-jet pair, $p_{\mathrm{T},bb}$;
- Invariant mass of the leading b-jet pair, M_{bb};
- Ratio of the transverse momentum of the leading b-jet pair over their invariant mass, $p_{\mathrm{T},bb}/M_{bb}$;
- Angular distance between the two leading b-jets ΔR_{bb};
- The jet width λ^1_1 (hereby denoted as λ) for the leading b-jet;
- The two-dimensional primary Lund Jet Plane of the leading b-jet.

Resolved 0B:

- Transverse momentum of the leading jet pair, $p_{\mathrm{T},jj}$;
- Invariant mass of the leading jet pair, M_{jj}
- Ratio of the transverse momentum of the leading jet pair over their invariant mass, $p_{\mathrm{T},jj}/M_{jj}$;
- Angular distance between the two leading jets ΔR_{jj};
- The jet width λ^1_1 (hereby denoted as λ) for the leading jet;
- The two-dimensional primary Lund Jet Plane of the leading jet.

Boosted 2B:

- Transverse momentum of the large-R jet $p_{\mathrm{T,large\text{-}R}}$;
- Mass of the large-R jet $M_{\text{large-R}}$;
- Angular distance between the two b-hadrons reconstructed as subjets within the large-R jet $\Delta R_{bb,\text{large-R}}$;
- The two-dimensional primary Lund Jet plane of the large-R jet.

[1] The ATLAS Collaboration has already measured the Lund Jet Plane for large-R jets from top and W bosons.

Boosted 0B:

- Transverse momentum of the large-R jet $p_{\mathrm{T,large\text{-}R}}$;
- Mass of the large-R jet $M_{\text{large-R}}$;
- The two-dimensional primary Lund Jet plane of the large-R jet.

We also note that a number of kinematic observables, such as the leading jet p_{T} and di-lepton mass $M_{\ell\ell}$, though not measured, are still reconstructed and used as control variables.

For all distributions considered, the data in the di-electron and di-muon channels are summed after background subtraction. Unfolding is then carried out using the Iterative Bayesian Unfolding (IBU) procedure using a fiducial phase space definition close to the detector-level one.

The Lund Plane distributions in the Resolved 2B and 0B signal regions are further divided into jet p_{T} slices. Three such slices are considered, between 20-60 GeV, 60-100 GeV, and above 100 GeV.

5.3 Data and Monte Carlo samples

This analysis makes use of data recorded with the ATLAS Experiment during Run 2, corresponding to the years 2015-2018. The events consist of proton-proton collisions at $\sqrt{s}$ =13 TeV during stable beam conditions with all subsystems of the detector operational and passing all data quality requirements. Overall, the data amounts to 140.1 fb^{-1}, with an uncertainty on the integrated luminosity of 0.8% [136].

The events analysed are selected using a single high p_{T} electron or muon required to pass a given trigger chain (cfr. Section 3.2.6). Each trigger chain imposes a different p_{T} threshold and isolation requirement for the leptons at both track and calorimeter level, with the isolation requirement loosening as the p_{T} threshold increases. The full list of high-level triggers considered is found in Table 5.1. Events must pass exactly one trigger.

5.3.1 Signal and background estimation

Signal and background processes are simulated using MC event generators. Z+jets represents both a signal for this measurement as well as a background, stemming from mistagged jets. This section is dedicated to the details of the simulation of this key sample.

Matrix elements for Z+jets with b-, c-, and light jet filters are simulated at NLO with MadGraph5_aMC@NLO v2.6.5. The ME is calculated in the 5FS, and explicitly includes the production of up to three partons. Generation is carried out with the NNPDF3.1$_{\mathrm{LUX}}$QED [137] PDF set with $\alpha_s = 0.118$. The Z boson is required to decay to charged leptons, and the invariant mass of the di-lepton system is forced to be greater than 40 GeV. Jets are generated with a $p_{\mathrm{T}} > 10$ GeV.

Events are then showered with Pythia8 v8.245 [132] parton shower with the A14 tune [134] and the NNPDF 2.3 LO PDF set with $\alpha_s = 0.130$ [135]. The different jet multiplicities are combined with the FxFx merging scheme [15], [44]. The decays of heavy flavour particles are managed by EvtGen 1.7.0 [133], while any QED radiation is simulated within Pythia8.

An alternative set of predictions for the signal samples is obtained with the ATLAS configuration of Sherpa 2.2.11 [138]. Using the Comix [139] and OpenLoops [140]–[142] libraries, matrix elements for the emission of up to two partons are computed at NLO accuracy, and they are computed at LO accuracy for the emission of up to five partons. Predictions are obtained in the 5FS and use the Hessian NNPDF 3.0 NNLO PDF sets with $\alpha_s = 0.118$ [143]. The Sherpa parton

Muon channel	
2015	HLT_mu20_iloose_L1MU15 HLT_mu50
2016–2018	HLT_mu26_ivarmedium HLT_mu50
Electron channel	
2015	HLT_e24_lhmedium_L1EM20VH HLT_e60_lhmedium HLT_e120_lhloose
2016–2018	HLT_e26_lhtight_nod0_ivarloose HLT_e60_lhmedium_nod0 HLT_e140_lhloose_nod0 HLT_e300_etcut

Table 5.1: The high-level triggers (HLT) used to select events in the muon (mu) and electron (e) channels of the measurement in each year of data taking. The names of each trigger are representative. After the "HLT" string, the particle type is listed, followed by the p_{T} threshold imposed. The third string in the muon triggers and the last string in the electron triggers specifies the isolation requirements. The strings beginning with "L1" specify the Level 1 trigger which the event must pass, while the remaining information signals any multivariate selections imposed.

shower [144] is used to simulate the final state evolution with the cluster hadronisation model [65]. Different jet multiplicities are merged with the MEPS@NLO prescription [137], [145]–[147] with a merging scale set to 20 GeV.

Background Monte Carlo samples

Several background sources are considered in this analysis and are estimated using MC simulation, proceeding similarly as for what was described for the signal samples, but simulating a different set of physics processes: $t\bar{t}$ and single-top production, $Z(\tau^+\tau^-)$+jets, VH production, and VV production. Below we describe the simulation parameters of each.

- ***$t\bar{t}$* :** The $t\bar{t}$ background is simulated using PowhegBox v2 at NLO and the NNPDF 3.0 NNLO PDF set [62], [63], [148]. The h_{damp} parameter – a damping factor regulating the matching of the ME to the parton shower and, consequently, the recoil of the $t\bar{t}$ system against high p_{T} radiation – is set to 1.5 times the top mass [149]. The sample is showered in Pythia8 v8.230 with the A14 tune and the NNPDF 2.3 LO PDF set. The sample is normalised to the NNLO cross-section calculation with next-to-next-to-leading logarithmic (NNLL) resummation of soft gluon terms, as calculated with Top++2.0 [149], [150].

- **Single-top:** Single-top production, either in association with a W boson, or in the s- and t-channels, is also calculated at NLO with PowhegBox v2 and showered with Pythia8 with the NNPDF 2.3 LO PDF set. Diagrams overlapping with the $t\bar{t}$ production are removed [151].

- ***VH* :** VH final states are also simulated with PowhegBox and Pythia8 with the NNPDF 3.0 NNLO PDF set.

- ***VV* :** Semi-leptonic di-boson VV final states, where one of the vector bosons decays hadronically and the other leptonically, are simulated with Sherpa v2.2.1 [152] with NLO ME for one parton emissions, and LO ME for up to three parton emissions. The NNPDF 3.0 NNLO PDF set is used in this calculation. The MEs are then showered with Sherpa 2.2.11 with the MEPS@NLO prescription [137], [145]–[147].

- ***Z*($\tau^+\tau^-$) :** Final states containing a Z boson decaying to τ leptons are generated in the same way as other $Z(\ell^+\ell^-)$ signal and background samples.

In all samples, the effect of pile-up is accounted for by overlaying simulated minimum-bias events. In particular, single-, double-, and non-diffractive proton-proton processes are simulated in Pythia8 v8.186 with the A3 tune [153] and and the NNPDF 2.3 LO PDF set.

All MC samples are processed in a simulation of the ATLAS Detector based on GEANT4 [154], [155]. The resulting physics objects are reconstructed using the same algorithms used for data.

5.4 Event selection

This measurement considers $Z + b\bar{b}$ and $Z + jj$ final states in both resolved and boosted topologies, making for a total of four signal regions: Resolved 2B, Resolved 0B, Boosted 2B and Boosted 0B. This section is dedicated to the selections applied in the four regions.

Lepton selections Lepton selections are in common to all signal regions. Events must contain exactly two same-flavour, opposite-sign (SFOS) leptons. Each individual lepton must have a $p_{\mathrm{T}} > 27$ GeV due to trigger thresholds, while the mass of the di-lepton pair must be found to be compatible with that of the Z boson, namely in the range 76-106 GeV. Muons are required to be located within the pseudorapidity range $|\eta| < 2.5$, while electrons must be found within $|\eta| < 1.37$ or $1.52 < |\eta| < 2.47$[2].

Leptons are also required to satisfy selections based on the impact parameters d_0 and z_0. Specifically, they must have $|z_0 \sin\theta| < 0.5$ mm and $|d_0/\sigma(d_0)| < 5.0$, while muons must have $|z_0 \sin\theta| < 0.5$ mm and $|d_0/\sigma(d_0)| < 3.0$. Lastly, both electrons and muons must meet isolation criteria, given by the "Tight" working point for electrons and "FCTight_FixedRad" working point for muons. These working points are defined in Sections 3.3.3 and 3.3.4.

The remaining selections focus on jets and their natures depends on the signal region considered.

Resolved SRs In the Resolved SRs, anti-k_t jets of radius $R = 0.4$ clustered with PFlow objects are considered. Jets must have a $p_{\mathrm{T}} > 20$ GeV and be found in the central region, within rapidity $|y| < 2.5$. Events in the Resolved 2B SR must contain at least 2 b-jets, tagged with the DL1r algorithm with the 70% working point. If a greater number of jets satisfy this criterion, the two leading jets are considered. In the Resolved 0B SR, events must contain at least 2 anti-b-tagged jets and no b-jets, still identified with the DL1r algorithm with the 70% working point. The two SRs are thus completely orthogonal to each other.

An overlap removal procedure is applied to ensure that leptons and jets are uniquely identified. Preselected jets found to be within $\Delta R < 0.2$ of a lepton are removed as they are likely to be misidentified leptons. In a second step, leptons found within $\Delta R < 0.4$ of jets which pass the full selections are removed.

[2]The electron pseudorapidity selection accounts for the presence of the transition region in the ECAL.

Boosted SRs Events passing the Boosted signal region selections must contain at least one AntiKt10LCTopoTrimmedPtFrac5SmallR20Jet. This is a jet collection of radius $R = 1.0$ produced using LCTopo objects. The jets are trimmed with parameters $R_{\text{sub}} = 0.2$ and $f_{\text{cut}} = 0.05$, as described in Section 3.3.6. Jets are required to have $p_{\text{T}} > 200$ GeV and to be found in the central region, with $|y| < 2.0$. These are be referred to as large-R jets through the rest of this chapter.

Large-R jets are tagged through the use of VR trackjets. Events must have at least two VR trackjets with $p_{\text{T}} > 10$ GeV and $|y| < 2.5$. VR trackjets are associated to large-R jets through a geometric matching criterion. In order to be considered associated, a VR trackjet must be within a distance of $\Delta R < 1.0$ from the large-R jet in the rapidity-azimuth plane.

A large-R jet is considered double-b-tagged (and thus belonging to the Boosted 2B region) if at least two b-tagged VR trackjets are associated to it. VR trackjets are tagged with the DL1r algorithm with the 70% working point. Likewise, large-R jets in the Boosted 0B SR must have at least two anti-b-tagged VR trackjets associated to it, and there must be no b-jets in the event. The highest p_{T} double-tagged large-R jet is considered for further analysis.

An analogous overlap removal procedure as in the Resolved SRs is applied. In this case, however, the radius in the second OR step is increased to 1.0 to match the large-R jet radius.

Unlike the 2B and 0B SRs, the Resolved and Boosted SRs are not orthogonal; it is possible for a given event to pass both selections.

Table 5.2 outlines the selections in all SRs.

5.5 Fiducial selections

The fiducial regions are defined to be as close as possible to the reconstruction-level selection applied on data. Electrons and muons stemming from the decay of the Z boson are selected identically to their detector-level selections, with two same-flavour, opposite-sign leptons, each with $p_{\text{T}} > 27$ GeV and $|\eta| < 2.5$. The di-lepton mass $M_{\ell\ell}$ must be compatible with that of the Z boson $76 < M_{\ell\ell}/\text{GeV} < 106$.

Jet selections depend on the signal region considered. The AntiKt4TruthDressedWZJets collection is used in the Resolved SRs. These contain jets clustered using stable, visible truth particles within $|\eta| < 5.0$. All particles are considered with the exception of:

- Prompt leptons from W, Z, or top decays
- Prompt neutrinos produced in the ME[3]
- Prompt photons from $H \to \gamma\gamma$ decays
- Photons not from hadron decays found to be within a cone of $\Delta R < 0.1$ centred around any charged leptons. These are dressing photons originating from final state radiation off of the leptons.

Jets are required to have $p_{\text{T}} > 20$ GeV, $|y| < 2.5$, and be isolated from the leptons. Specifically, jets are removed if a lepton is found within their radius.

Jets are then labelled as b-jets if a b-hadron with $p_{\text{T}} > 5$ GeV is found within a distance of $\Delta R < 0.3$ of the jet axis. Jets are labelled as light jets if no b- or c-hadrons with transverse momentum above the threshold are associated to it.

In the Boosted SRs, the AntiKt10TruthTrimmedPtFrac5SmallR20Jets are considered. These are formed with all visible, stable particles with the exception of muons and are trimmed with

[3] Neutrinos from other sources, such as muon or hadron decays, are considered.

Leptons	
2 SFOS leptons $\ell^+\ell^-$ $p_{\mathrm{T}}^{\ell} > 27$ GeV, $76 < M_{\ell\ell}/\mathrm{GeV} < 106$	
Resolved SRs	
≥ 2 anti-k_t $R = 0.4$ PFlow jets $p_{\mathrm{T}}^{\mathrm{jet}} > 20$ GeV $\lvert y_{\mathrm{jet}}\rvert < 2.5$ & $\lvert\eta_{\mathrm{jet}}\rvert < 2.4$	
Resolved 2B	**Resolved 0B**
≥ 2 b-tagged jets with $\mathrm{DL1r}_{\mathrm{PFlow}}$ 70% WP	Exactly 0 b-tagged jets with $\mathrm{DL1r}_{\mathrm{PFlow}}$ 70% WP
Boosted SRs	
≥ 1 $R = 1.0$ LCTopo jet $p_{\mathrm{T}}^{\mathrm{jet}} > 200$ GeV, $\lvert y_{\mathrm{jet}}\rvert < 2.0$ ≥ 2 VR Trackjets matched to large-R jet $p_{\mathrm{T}}^{\mathrm{VR}} > 10$ GeV $\lvert y_{\mathrm{VR}}\rvert < 2.5$ & $\lvert\eta_{\mathrm{VR}}\rvert < 2.4$	
Boosted 2B	**Boosted 0B**
≥ 2 b-tagged jets with $\mathrm{DL1r}_{\mathrm{VR}}$ 70% WP	Exactly 0 b-tagged jets with $\mathrm{DL1r}_{\mathrm{VR}}$ 70% WP

Table 5.2: An overview of the event selections applied in each of the four signal regions.

the same parameters as in the detector-level case. A large-R jet passes the pre-selection if it has $p_{\mathrm{T}} > 200$ GeV, $|y| < 2.0$, and does not contain a lepton within its radius. The fiducial Boosted 2B selection is passed if at least two b-hadrons with $p_{\mathrm{T}} > 5$ GeV are ghost matched to the jet. A large-R jet instead passes the fiducial Boosted 0B selection if no b-hadrons or c-hadrons are ghost matched to the jet. Details on the jet-hadron matching schemes used can be found in Section 2.1.

Table 5.3 summarises the fiducial selections in the Resolved and Boosted SRs.

Leptons	
Exactly 2 SFOS leptons $\ell^+\ell^-$ from a Z decay $p_{\mathrm{T}}^{\ell} > 27$ GeV, $76 < M_{\ell\ell}/\mathrm{GeV} < 106$	
Resolved SRs	
≥ 2 anti-k_t $R = 0.4$ TruthWZ jets $p_{\mathrm{T}}^{\mathrm{jet}} > 20$ GeV, $\lvert y_{\mathrm{jet}}\rvert < 2.5$ $\Delta R(\mathrm{jet}, \ell) > 0.4$	
Resolved 2B	**Resolved 0B**
≥ 2 b-jets labelled with cone matching b-hadron $p_{\mathrm{T}} > 5$ GeV	Exactly 0 b-jets and c-jets labelled with ghost matching b/c-hadron $p_{\mathrm{T}} > 5$ GeV
Boosted SRs	
≥ 1 $R = 1.0$ trimmed jets $p_{\mathrm{T}}^{\mathrm{jet}} > 200$ GeV, $\lvert y_{\mathrm{jet}}\rvert < 2.0$, $\Delta R(\mathrm{jet}, \ell) > 1.0$ Trimming parameters: $R_{\mathrm{sub}} = 0.2$, $f_{\mathrm{cut}} = 0.05$	
Boosted 2B	**Boosted 0B**
≥ 2 b-hadrons with $p_{\mathrm{T}} > 5$ GeV ghost matched to the large-R jet	Exactly 0 b/c-hadrons with $p_{\mathrm{T}} > 5$ GeV ghost matched to the large-R jet

Table 5.3: An overview of the event selections applied in each of the four fiducial regions.

5.6 Reconstruction of JSS observables

Due to the higher resolution of the ATLAS Inner Detector, the Lund Jet Plane and jet angularities are reconstructed from tracks. An initial selection of tracks is made from those satisfying the requirements described in Section 3.3.1.

The JSS observables are formed by matching tracks[4] to PFlow jets in the Resolved SRs and trimmed LCTopo jets in the Boosted SRs. The matching criterion is geometric, requiring that a track be found within a distance equal to the jet radius in the $y-\phi$ plane. After matching, tracks found to be associated to a jet are clustered into an anti-k_t jet of the same radius as the input jet in consideration and, in the case of the Boosted SRs, the trimming algorithm is applied to the newly-clustered large-R trackjet. The Lund Jet Plane or jet angularities are then calculated according to their definitions.

Tracks with $p_{\mathrm{T}} > 900$ MeV satisfying the "Loose" Track Quality Selection (TS) and "Loose" Track-to-Vertex Association (VA) working points are considered. These selections are optimised to ensure the highest number of matches between the number of emissions in a given jet at particle- and detector-level[5]. It is important to note that this does not necessarily ensure that the number of fakes and inefficiencies is minimised, just that, per happenstance, the number of fakes and inefficiencies is such that the total number of emissions in the primary Lund Plane is the same between particle- and detector-level. The scan of different working points for a fixed track p_{T} threshold of 900 MeV is shown in Figure 5.1.

In Figure 5.2, we justify the choice of the 900 MeV threshold. It is shown how, as the track p_{T} threshold increases, so too does the frequency of jets having the same number of emissions at both particle- and detector-level. However, as the p_{T} threshold increases, it is important to note how more and more tracks associated to the jet fail to reach this threshold, as can be seen by the normalisation of each distribution. This effect is more significant for jets with low p_{T}. As a compromise, the threshold is set to 900 MeV.

These patterns in matching were verified across a wide range of $R = 0.4$ jet p_{T} intervals, for large-R jets, and comparisons were made for b-jets and light jets. We find that the same patterns hold across types of jets considered. For instance, in Figure 5.3, one can see that the pattern in track working points for a fixed track p_{T} cut shown in Figure 5.1 holds for boosted large-R jets. Similarly, in Figure 5.5 we see that the pattern in matching when varying the track p_{T} cut holds independently of the jet p_{T} considered. We also see that the frequency of jets having the same number of emissions across particle- and detector-level increases as the jet p_{T} decreases. This is due to the discrete nature of the variable and the fact that low p_{T} jets tend to have a lower number of emissions in the primary Lund Plane. Thus, there is a higher probability of finding the same number of emissions at both particle- and detector-level. Lastly, in Figure 5.4 we show that the pattern for track working points holds for light jets. If these light jets are further matched across particle- and detector-level to suppress pile-up jets, the frequency of the primary Lund Plane having the same number of emissions is remarkably similar to that of b-jets.

For a small minority of jets (< 1 %), no tracks are matched to the jet. This can happen indistinctively to the leading and subleading jet in both the Resolved 2B and Resolved 0B SRs, though it occurs more often to subleading jets and more often in the Resolved 0B SR. In these cases, we always require that the leading jet have at least one matched track to allow for the reconstruction of the substructure observables.

[4] A study detailing the motive behind the choice of tracks is available in Appendix A.

[5] This is done to reduce bias introduced in the definition of the response matrix used in the unfolding procedure. See Section 5.10 for details.

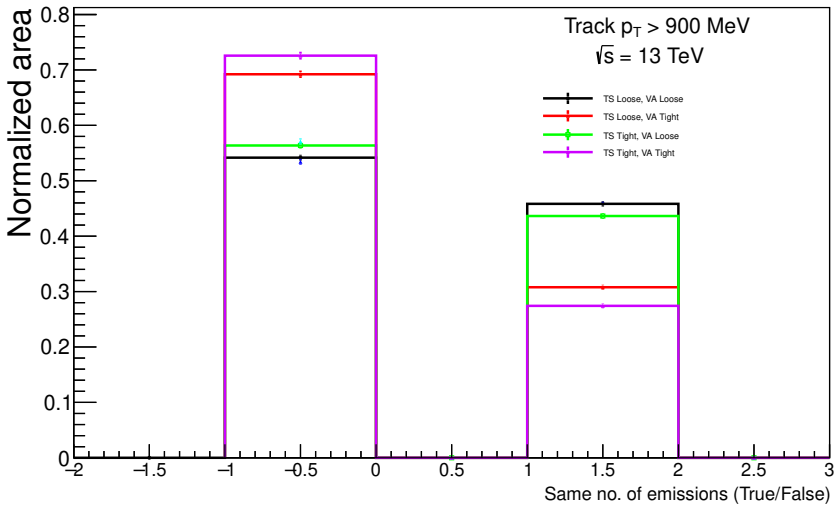


Figure 5.1: The frequency with which the primary Lund Plane for jets at particle- and detector-level have the same number of emissions as a function of Track Selection and Track-to-Vertex Association working points. The value of zero indicates that the two Lund Planes at have a different number of emissions, while the value of one indicates that they have the same value. The evaluation is made for track with $p_{\mathrm{T}} > 900$ MeV.

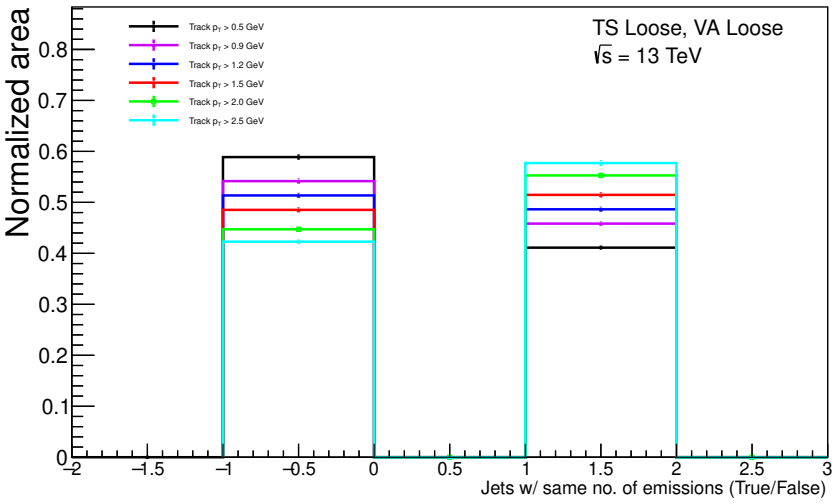


Figure 5.2: The frequency with which the primary Lund Plane for jets at particle- and detector-level have the same number of emissions as a function of the track p_t threshold. The value of zero indicates that the two Lund Planes have a different number of emissions, while the value of one indicates that they have the same value. The Track Selection and Track-to-Vertex Association working points are both fixed to "Loose".

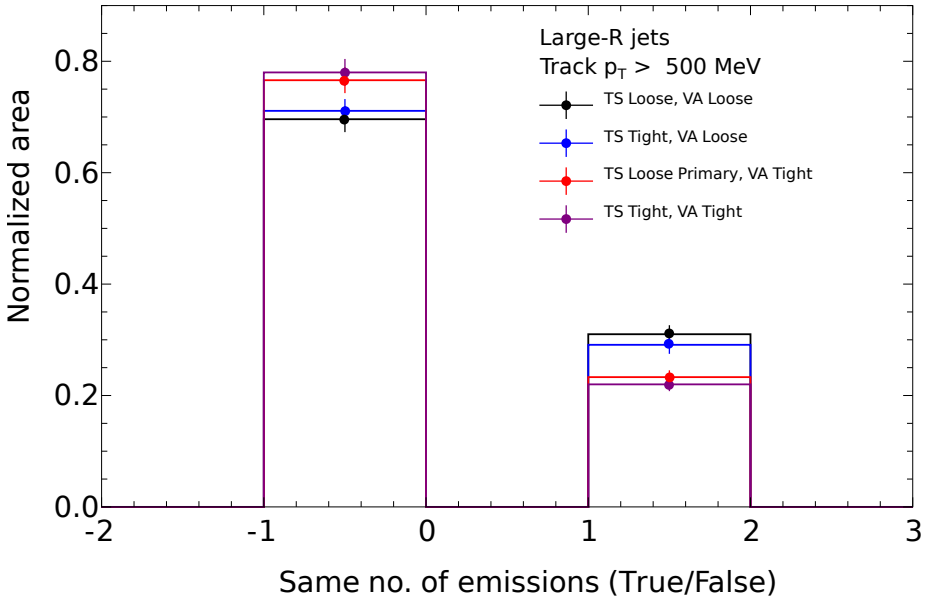


Figure 5.3: The frequency with which the primary Lund Plane for boosted large-R jets at particle- and detector-level have the same number of emissions as a function of Track Selection and Track-to-Vertex Association working points. The value of zero indicates that the two Lund Planes at have a different number of emissions, while the value of one indicates that they have the same value. The evaluation is made for track with $p_{\mathrm{T}} > 500$ MeV.

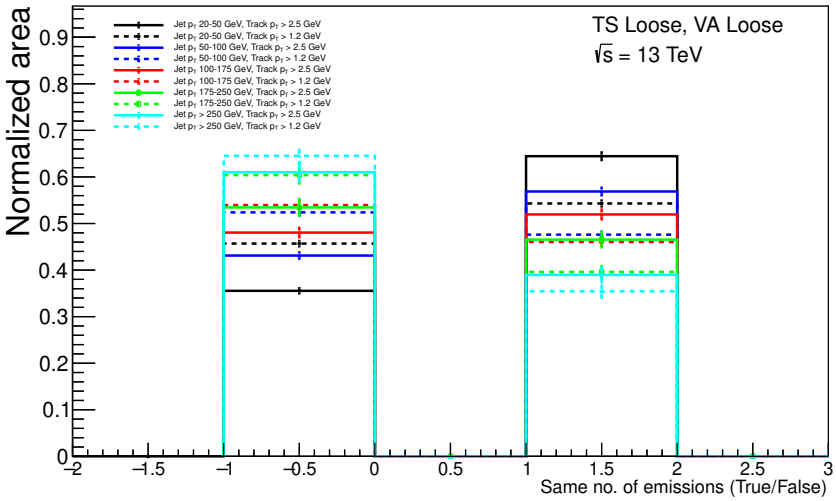


Figure 5.4: The frequency with which the primary Lund Plane for jets at particle- and detector-level have the same number of emissions for a variety of jet p_{T} intervals. Two different track p_{T} cuts are considered, 1.2 GeV and 2.5 GeV. The value of zero indicates that the two Lund Planes at have a different number of emissions, while the value of one indicates that they have the same value. The evaluation is made for a "Loose" Track Selection and "Loose" Track-to-Vertex Associations working points.

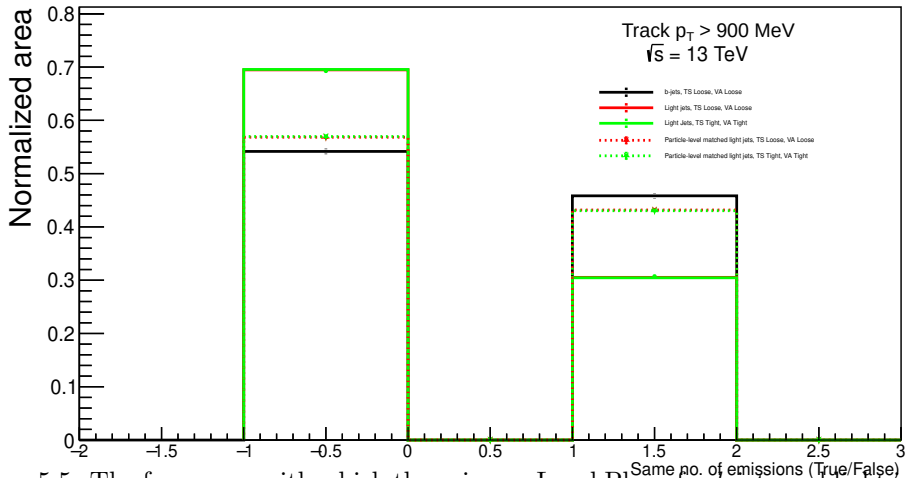


Figure 5.5: The frequency with which the primary Lund Plane for b-jets and light jets at particle- and detector-level have the same number of emissions as a function of Track Selection and Track-to-Vertex Association working points. The value of zero indicates that the two Lund Planes at have a different number of emissions, while the value of one indicates that they have the same value. The evaluation is made for track with $p_{\mathrm{T}} > 500$ MeV.

5.6.1 Heavy hadron reconstruction

A quark in the final state will radiate gluons, with each subsequent gluon being emitted at a smaller angle. The radiation may in turn emit more radiation, leading to the tree-like structure we call a jet. This structure is mostly preserved after hadronisation.

Heavy flavour hadrons, however, decay into light hadrons via the weak interaction. To a jet clustering algorithm, the decay products rightly appear as just another emission as there is no way to distinguish one from the other in data. To get around this issue, theorists often publish predictions for processes at either parton level or artificially inhibit the decay of the heavy flavour hadron if at hadron level. Indeed, in this analysis we also choose to inhibit the decay of heavy hadrons in the fiducial definition of the Lund Jet Plane.

The effect of inhibiting the b-hadron decay in the particle-level definition of the Lund Jet Plane is shown in Figure 5.6 for b-jets between 60-100 GeV. It is clear how the presence of the decay products leads to an excess of emissions in the small-angle region.

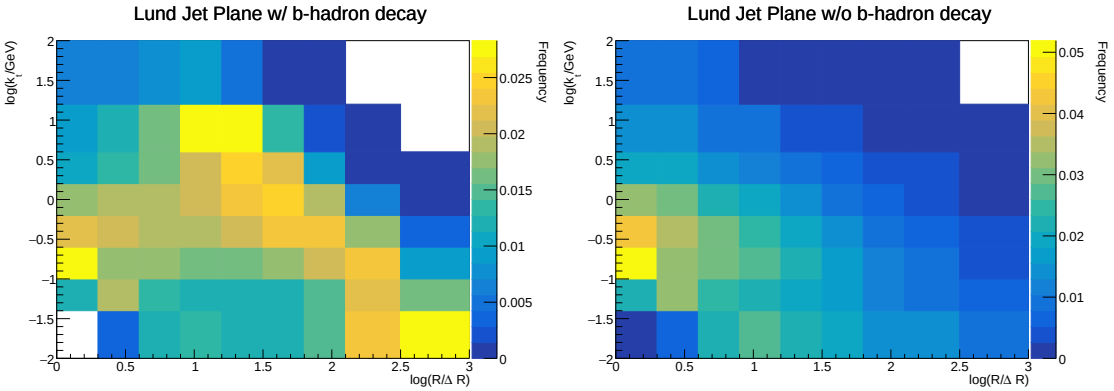


Figure 5.6: The particle-level b-jet Lund Jet Plane defined with and without the decay products of the b-hadron. The area is normalised to unity.

In this measurement, we opt to reconstruct the heavy flavour hadron at detector-level from its charged decay products with a cut-based strategy. The b-hadron is reconstructed through a selection on the track signed impact parameter significance s_{d_0}, introduced in Section 4.1.2. In Figure 5.7 the s_{d_0} distribution for tracks associated to PFlow jets in our selection tagged as b-jets or light jets is shown. From here, it is clear that the observable chosen has discrimination potential.

To reconstruct the b-hadron from its charged decay products, a detailed study was undertaken to optimise the track selection. The reconstructed b-hadron four-momentum was optimised against the particle-level b-hadron partially reconstructed from only its charged decay products. This quantity is hereby referred to as the *particle-level charged four-momentum*.

We find that tracks satisfying the requirement $s_{d_0} > 0.5$ or $s_{d_0} < -4$ best reproduces the particle-level charged four-momentum of the b-hadron. In cases in which no track in a b-jet is found to satisfy these requirements, the leading track is assumed to originate from the decay, in accordance with the leading particle effect. In this way, no additional requirement beyond b-tagging is placed on the jets whose substructure we analyse.

Figure 5.8 shows, for each reconstructed b-jet in the Resolved 2B SR, the particle-level charged four-momentum versus the reconstructed b-hadron four-momentum from tracks. The charged b-hadron momentum is defined from the charged decay products of the b-hadron with

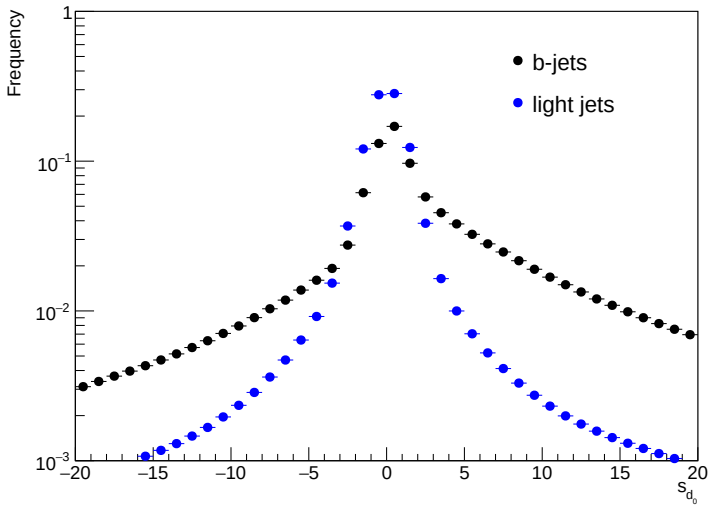


Figure 5.7: The distribution of the signed d_0 significance variable for tracks in PFlow jets tagged as b-jets (black) and the same distribution for tracks associated in PFlow jets tagged as light jets (blue).

$p_{\mathrm{T}} > 900$ MeV to match the track selections. The migration is studied in simulation for cases in which the particle-level and detector-level jet fall within a distance of $\Delta R < 0.4$, and for which exactly one undecayed b-hadron is matched to the particle-level jet within a distance of $\Delta R < 0.4$. In a minority of cases (< 1 %) no charged decay products match the selection criteria. These cases are not considered.

From Figure 5.8 it is clear how the track selection criterion can effectively identify those tracks stemming from decay products, as the transverse momentum, rapidity and azimuthal angle of the charged component of the b-hadron is aligned with that of the b-hadron reconstructed from tracks.

Figure 5.9, on the other hand, shows the migration of the reconstructed b-hadron four-momentum when compared against the full, undecayed particle-level b-hadron. From here, it is clear that, although the p_{T} of the reconstructed b-hadron tends to underestimate that of the undecayed truth b-hadron, the y and ϕ distributions remain aligned between particle- and detector-level.

Up to now, the discussion of b-hadron reconstruction has been focused on tracks associated to PFlow jets. Although there is no reason to expect why this should not hold true for tracks associated to boosted jets, the effectiveness of the reconstruction was also checked in this case.

In the Boosted 2B SR, the s_{d_0} distribution for tracks associated to VR trackjets tagged as b-jets is studied. Figure 5.10 shows the distribution of this variable for tracks associated to b-jets in the Boosted 2B SR vs. the same distribution in the Resolved 2B SR. No significant difference is observed. For this reason, the same track selections are implemented in the Boosted 2B SR.

Tables 5.5 and 5.4 summarise the detector-level and particle-level selections and heavy hadron reconstruction strategies used to reconstruct the Lund Jet Plane and jet angularities in the four SRs of the analysis, respectively.

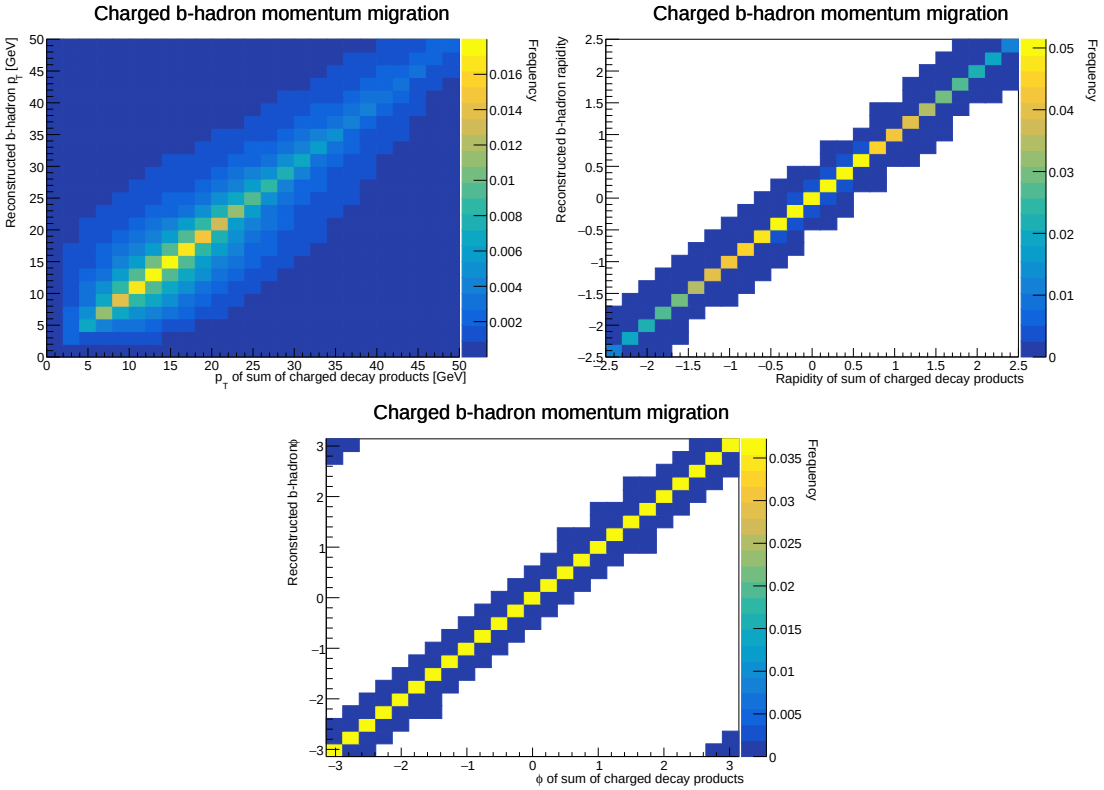


Figure 5.8: The p_{T}, η, and ϕ components of the particle-level charged b-hadron momentum obtained via the sum of the heavy hadron decay products vs. the same components for the detector-level reconstructed b-hadron.

Particle-level selections
Jet constituents with $p_{\mathrm{T}} > 900$ MeV
Resolved & Boosted 2B regions
Inhibit decay of b-hadron(s) used to label the jet

Table 5.4: The particle-level selections used to reconstruct the jet substructure observables and heavy hadrons in the four SRs.

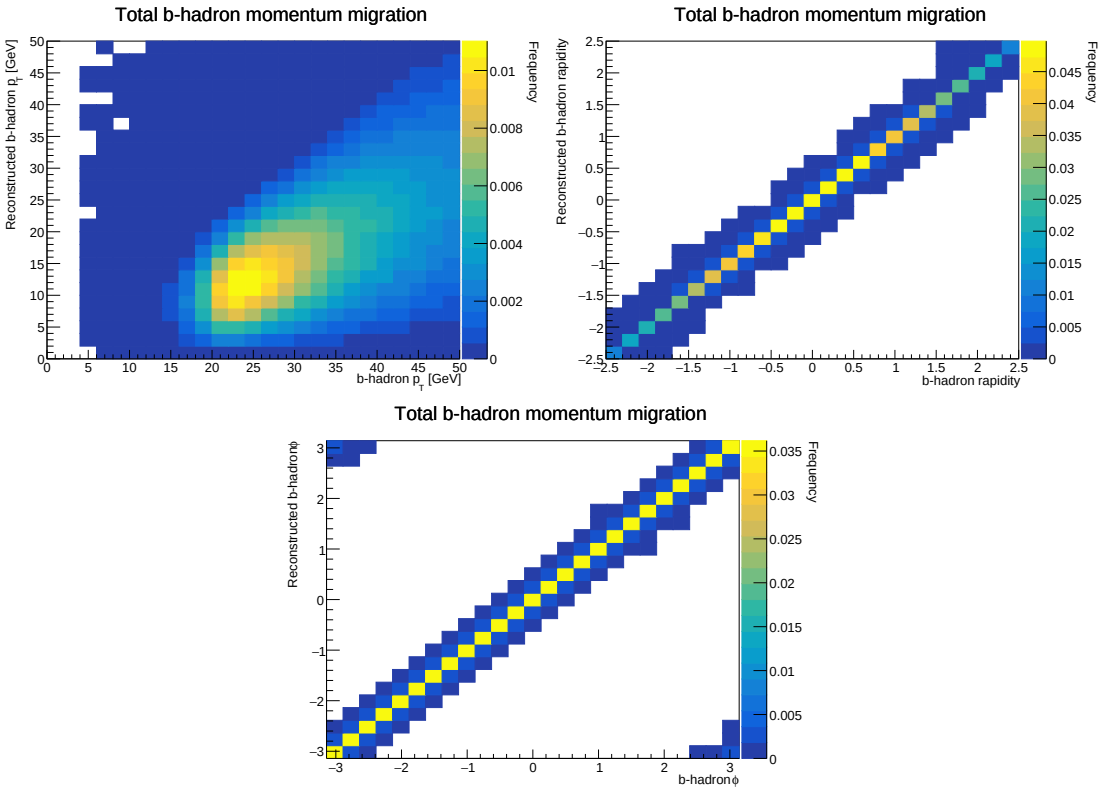


Figure 5.9: The reconstructed b-hadron p_{T}, y, and ϕ vs. that of the undecayed particle-level b-hadron.

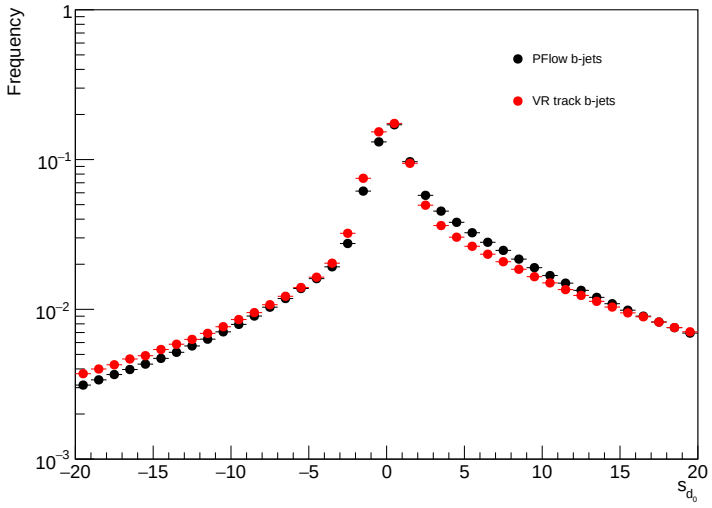


Figure 5.10: The s_{d_0} variable for tracks associated to PFlow jets tagged as b-jets (black) and to VR trackjets tagged as b-jets (red).

Detector-level selections	
Tracks with $p_{\mathrm{T}} > 900$ GeV Pass "Loose" Track-to-Vertex Association Pass "Loose" Track Quality Selection	
Resolved 2B	**Resolved 0B**
ΔR(track, b-jet) < 0.4 Combine tracks within b-jet with $s_{d_0} < -4$ or $s_{d_0} > 0.5$	ΔR(track, jet) < 0.4
Boosted 2B	**Boosted 0B**
ΔR(track, large-R jet) < 1.0 Combine tracks within same VR trackjet with $s_{d_0} < -4$ or $s_{d_0} > 0.5$	ΔR(track, large-R jet) < 1.0

Table 5.5: The detector-level selections used to reconstruct the jet substructure observables and heavy hadrons in the four SRs.

5.6.2 One-dimensional representation of the Lund Plane

The Lund Jet plane is a two-dimensional distribution. To ease the visual analysis of the agreement between reconstructed data and estimated signal and background contributions, and to simply the unfolding procedure, a so-called "unrolling" procedure of the distribution is implemented. This procedure calls for the assignment of a unique number to each bin in the two-dimensional histogram of the Lund Plane, including overflow and underflow bins. The data are then plotted in a one-dimensional histogram, where on the abscissa the bin number is indicated.

Bins are numbered sequentially row-by-row, from the bottom to the top of the two-dimensional histogram. Two additional bins are added to each row to account for underflow and overflow values of $\ln(R/\Delta R)$. These are added to the left and right of each row, respectively. Likewise, two additional rows are added to the one-dimensional mapping to account for the underflow and overflow in $\ln(k_t/\text{GeV})$, the former at the bottom of the Lund Plane and the latter at the top. Overall, if the two-dimensional primary Lund Plane is defined with N_x and N_y bins on the x- and y-axis, the one-dimensional variant contains $(N_x + 2) \times (N_y + 2)$ total bins. An example is shown in Figure 5.11.

The unrolling operation is bijective, allowing one to return to the two-dimensional representation of the Lund Plane at any point.

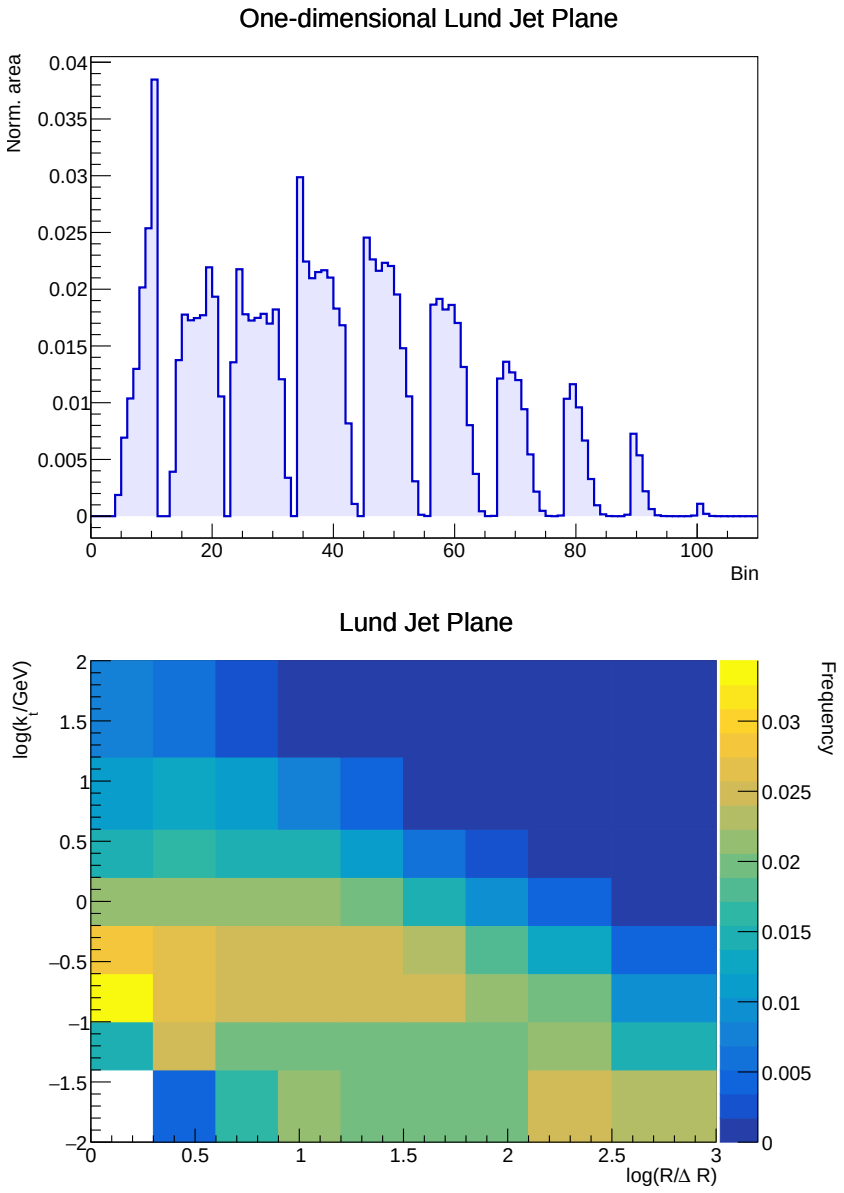


Figure 5.11: An example of the output of the Lund Plane unrolling procedure of simulated leading jets between 60-100 GeV passing the Resolved 0B selections and the equivalent two-dimensional distribution. The characteristic rise and fall seen is a result of the scan across rows of the two dimensional Lund Plane. The area under the curve is normalised to unity.

5.7 Backgrounds

In Section 5.3, we provide a list of the backgrounds which contribute to the various signal regions. Most of these, particularly VH, VV, single-top, multi-jet, and $Z(\tau^+\tau^-)$ are small and can estimated directly from MC, or, such as in the case of the multi-jet background, are considered to be negligible based on previous results [53]. The most significant backgrounds are instead $t\bar{t}$ and Z+jets with mistagged jets. These must be addressed with more sophisticated techniques. In the case of the $t\bar{t}$ background, we designed an ad-hoc neural network to reject this background. The remaining $t\bar{t}$ background is then estimated with data-driven methods. These are described in Sections 5.7.1 and 5.7.2, respectively.

To combat the Z+jets background, a dedicated flavour fit on a flavour-sensitive observable, namely the DL1r score, is carried out. The aim of this fit is to constrain the shape and normalisation of each observable measured. In Section 5.7.3, we describe this fit.

Lastly, in Section 5.7.4, we provide the final estimate of all background processes to each SR.

5.7.1 $t\bar{t}$ rejection

To suppress the large $t\bar{t}$ background, a dedicated dense neural network (NN) is used to differentiate Z+jets events from $t\bar{t}$ events. The NN is trained on pre-tagged jets, and is applied exclusively in the Resolved 2B SR.

The NN takes in input a variety of kinematic variables related to the jets, leptons, missing transverse momentum, etc. A full list of the input variables is provided in Table 5.6. Details on the NN architecture, written in the Keras library [156], is provided in Table 5.7. The hyperparameters are listed in Table 5.8.

Figure 5.12 shows the discrimination potential of a selection of input variables – namely the di-jet transverse momentum p_{T}^{jj}, missing transverse momentum $p_{\mathrm{T}}^{\mathrm{miss}}$, di-lepton mass $M_{\ell\ell}$ around the Z boson peak, and angular distance between the leptons $\Delta R(\ell^0, \ell^1)$ – for the $Z + b\bar{b}$ signal and $t\bar{t}$ background. Also shown is the combined signal+background distribution, its mean value, and the root mean squared (RMS). The signal/background discrimination potential of these observables, especially the di-lepton mass, is evident.

Figure 5.13 shows the normalised output of the NN evaluated on $t\bar{t}$ and $Z + b\bar{b}$ events and the relative cumulative integrals (rejections). Additionally, the Receiver Operator Characteristic (ROC) curve showing the background rejection vs. signal efficiency is shown.

Table 5.9 shows the scores and rejections for a selection of working points considered. The final score used to cut on the NN output is 0.490, corresponding to an 85% working point and an 85% rejection of the $t\bar{t}$ background. Ideally, at the working point defined, the signal efficiency is high enough to avoid any significant loss of signal events. This would circumvent the issue of needing to define a cut equivalent to the effect of the NN in the fiducial selection.

Variable	Description
$p_{\mathrm{T}}^{\mathrm{jet}_0}$	Leading jet p_{T}
$p_{\mathrm{T}}^{\mathrm{jet}_1}$	Subleading jet p_{T}
p_{T}^{jj}	di-jet p_{T}
M_{jj}	di-jet mass
$\Delta R(j^0, j^1)$	Distance between jets
$p_{\mathrm{T}}^{\ell^0}$	Leading lepton p_{T}
$p_{\mathrm{T}}^{\ell^1}$	Subleading lepton p_{T}
$p_{\mathrm{T}}^{\ell\ell}$	di-lepton p_{T}
$p_z^{\ell\ell}$	di-lepton momentum along the z-axis
$M_{\ell\ell}$	di-lepton mass
$\Delta R(\ell^0, \ell^1)$	Distance between leptons
$p_{\mathrm{T}}^{\mathrm{miss}}$	Missing transverse momentum
$p_{\mathrm{T}}^{\mathrm{vis}}$	p_{T} of the $(j^0, j^1, \ell^0, \ell^1)$ system
p_z^{vis}	Momentum along the z-axis of the $(j^0, j^1, \ell^0, \ell^1)$ system
M_{vis}	Mass of the $(j^0, j^1, \ell^0, \ell^1)$ system
$M_{t,\mathrm{vis}}$	Transverse mass of the $(j^0, j^1, \ell^0, \ell^1)$ system
$p_{\mathrm{T}}^{\mathrm{soft}}$	$p_{\mathrm{T}}^{\mathrm{miss}} - p_{\mathrm{T}}^{\mathrm{vis}}$
$M_{j^0\ell^0}$	Mass of the (j^0, ℓ^0) system
$M_{j^1\ell^1}$	Mass of the (j^1, ℓ^1) system
h_t	Scalar sum of p_{T} of all particles
$\Delta\phi(p_{\mathrm{T}}^{\mathrm{miss}}, j^0)$	Angular distance of missing transverse momentum and j^0
$\Delta\phi(p_{\mathrm{T}}^{\mathrm{miss}}, j^1)$	Angular distance of missing transverse momentum and j^1
$\Delta\phi(p_{\mathrm{T}}^{\mathrm{miss}}, \ell^0)$	Angular distance of missing transverse momentum and ℓ^0
$\Delta\phi(p_{\mathrm{T}}^{\mathrm{miss}}, \ell^1)$	Angular distance of missing transverse momentum and ℓ^1
$\Delta\phi(p_{\mathrm{T}}^{\mathrm{miss}}, j^0 + j^1)$	Angular distance of missing transverse momentum and the di-jet system
$\Delta\phi(p_{\mathrm{T}}^{\mathrm{miss}}, \ell^0 + \ell^1)$	Angular distance of missing transverse momentum and the di-lepton system

Table 5.6: The input variables of the dense neural network trained for the $t\bar{t}$ background rejection.

Layer	No. of nodes	Activation function
1	40	ReLU
2	50	ReLU
3	20	ReLU
4	1	sigmoid

Table 5.7: The architecture of the dense neural network trained for the $t\bar{t}$ background rejection.

Hyperparameter	Setting
Training sample size	460000
Learning rate	10^{-3}
Batch size	10^{3}
Max epochs	75

Table 5.8: The settings of the hyperparameters used for training the neural network for the $t\bar{t}$ background rejection.

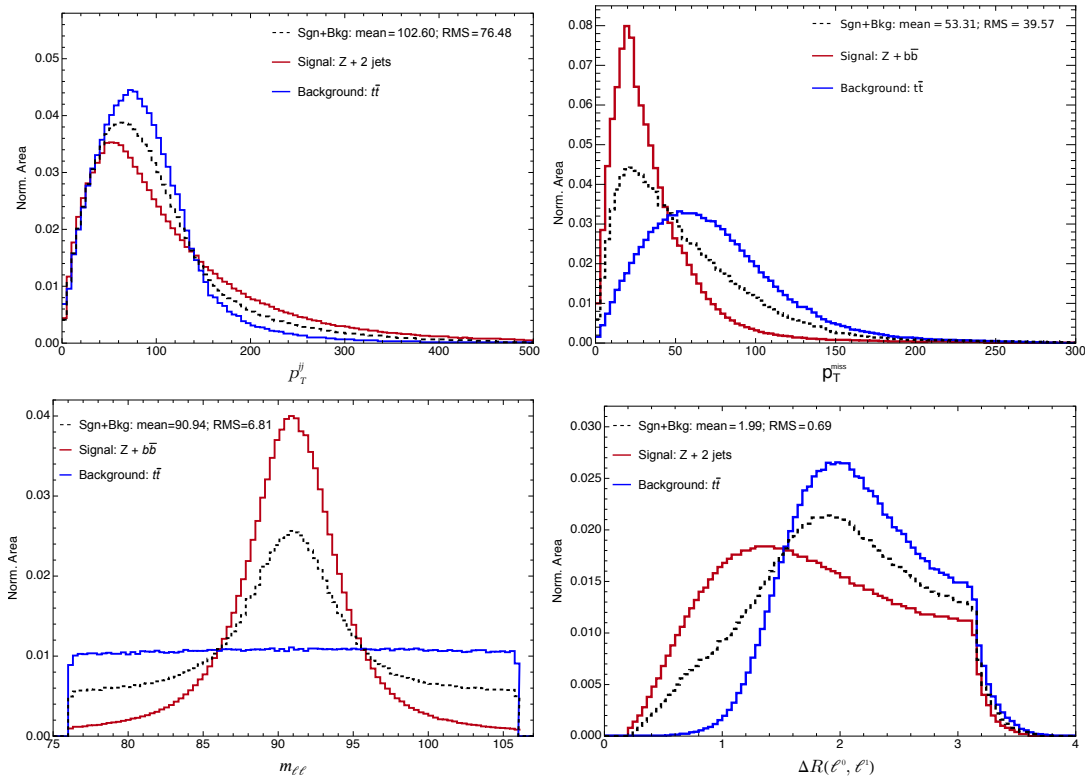


Figure 5.12: The signal and background distributions for the p_{T}^{jj}, $p_{\mathrm{T}}^{\mathrm{miss}}$, $M_{\ell\ell}$, and $\Delta R(\ell^0, \ell^1)$ variables used to train the neural network.

Signal efficiency	Background rejection	Output score
98.3%	46.6%	0.01
96.6%	58.8%	0.05
94.8%	66.7%	0.11
90.0%	78.6%	0.31
85.3%	84.7%	0.49

Table 5.9: A selection of signal efficiency working points for the NN designed for $t\bar{t}$ rejection and their relative background rejections and output scores.

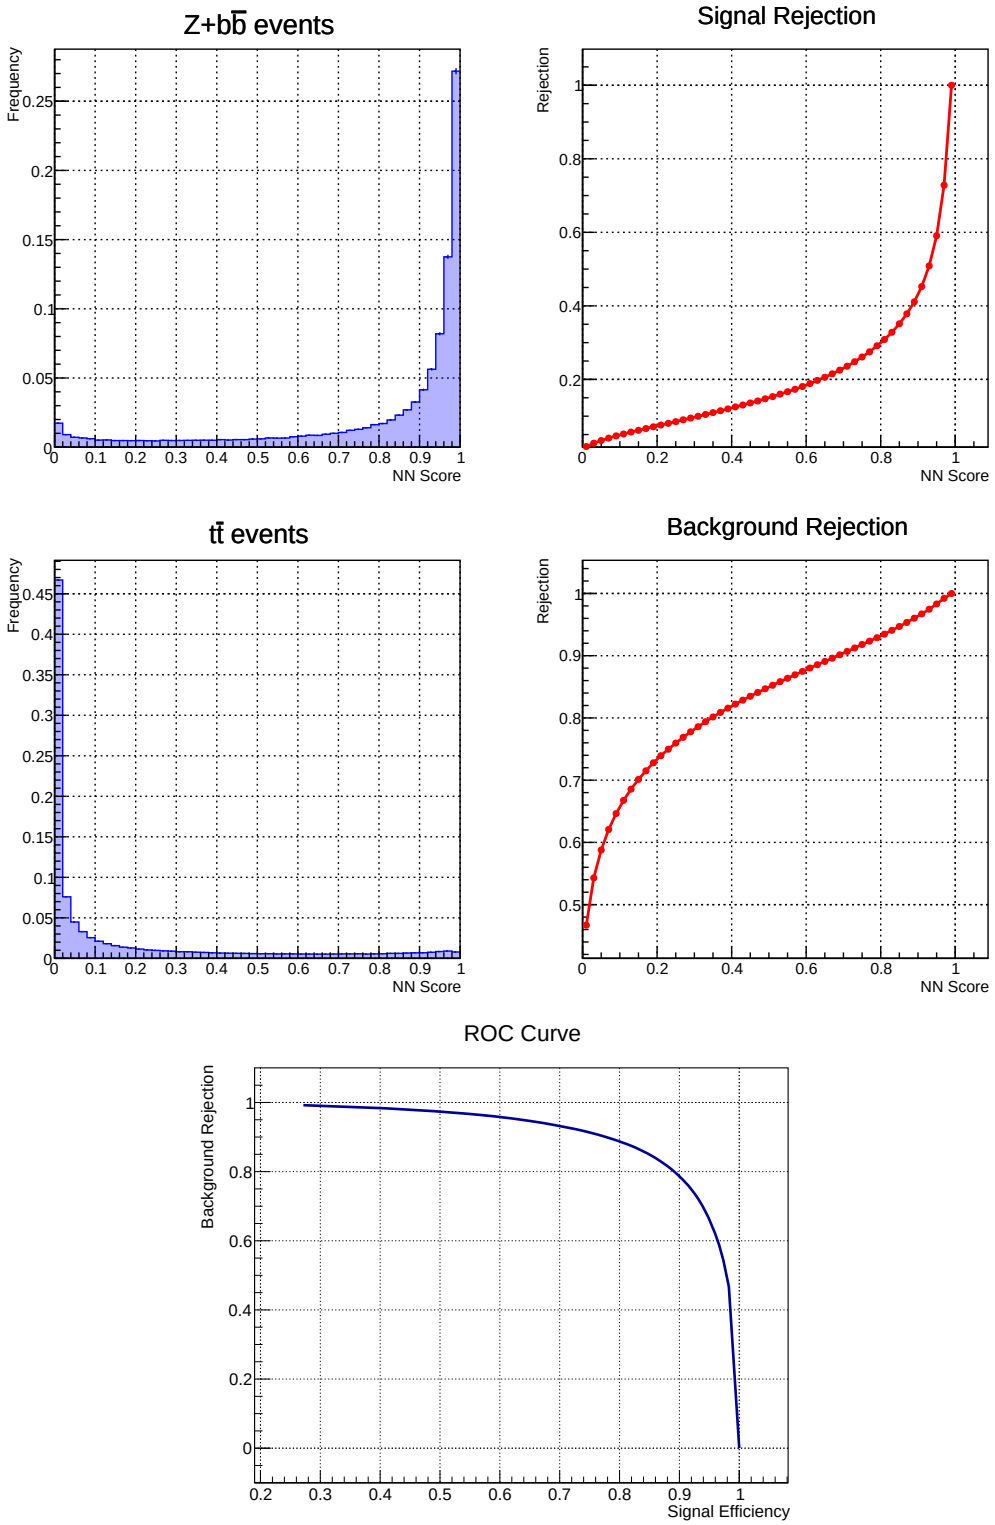


Figure 5.13: The NN output scores for the $Z + b\bar{b}$ signal (top left) and $t\bar{t}$ background (centre left), with their relative integrals (top right and centre right). The ROC curve is shown in the bottom row. The area under the curve (AUC) for the ROC curve is equal to 0.797.

5.7.2 $t\bar{t}$ data-driven estimation

As the $t\bar{t}$ background is highly sensitive to modelling uncertainties in the MC, a data-driven estimation is carried out. To this aim, a $t\bar{t}$ enriched control region (CR) is defined by requiring an opposite-sign electron and muon pair in the selections. This CR is chosen because the event kinematics for di-leptonic $t\bar{t}$ decays is identical in the e^+e^-, $\mu^+\mu^-$, and $e^\pm\mu^\mp$ channels. Therefore, the CR is characterised by very a high purity $t\bar{t}$ selection ($> 99\%$), yet remains orthogonal to the SRs considered in the measurement, allowing us to use the MC to derive transfer factors from the $e^\pm\mu^\mp$ CR to the $\ell^+\ell^-$ regions which can then be applied to data.

First, the backgrounds of the $t\bar{t}$ selection are subtracted from the CR region in the data sample. These are estimated from simulation. The result is an expectation of the number of $t\bar{t}$ events in data which pass the CR selections.

Transfer factors are then derived for each observable. These are obtained from simulation, and correspond to the ratio of each observable in the SR to that in the CR. The transfer factors are then used to multiply the background-subtracted CR distributions in data, mapping the $t\bar{t}$ estimation to the SR.

The transfer factors from the di-electron and di-muon channels for a selection of observables from each signal region considered are shown in Figures 5.14–5.17. Transfer factors from the di-muon channel are generally greater than those from the di-electron channel.

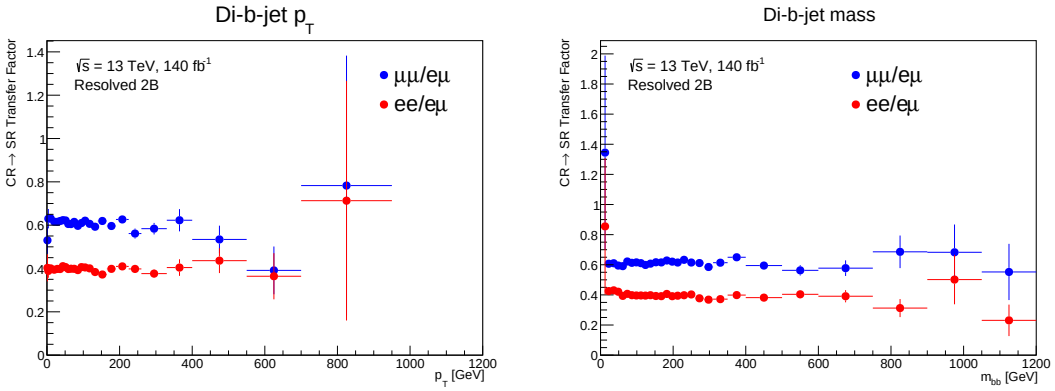


Figure 5.14: The transfer factors for the di-electron and di-muon channels for the di-b-jet p_{T} (left) and di-b-jet mass (right) in the Resolved 2B signal region.

The difference between the data-driven estimation and MC estimation of the $t\bar{t}$ background is shown in Figures 5.18–5.21 for the same selection of observables. The data-driven estimation of the background is consistently about 10% higher in the bulk of the Resolved SRs, with peaks above 20%.

Lastly, in Table 5.10 we report the exact values of the estimated number of $t\bar{t}$ events for both the data-driven and MC estimations for all years of data taking in all signal regions in the combined di-lepton channel. The table with all backgrounds is available in Section 5.7.4.

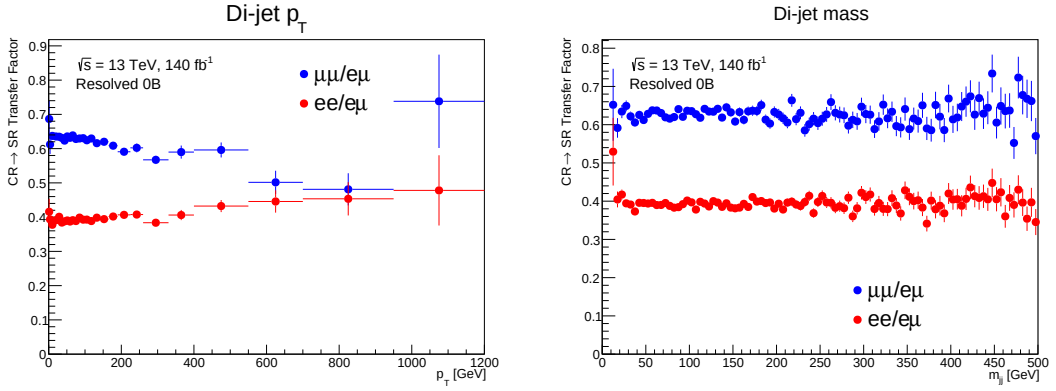


Figure 5.15: The transfer factors for the di-electron and di-muon channels for the di-jet p_{T} (left) and di-jet mass (right) in the Resolved 0B signal region.

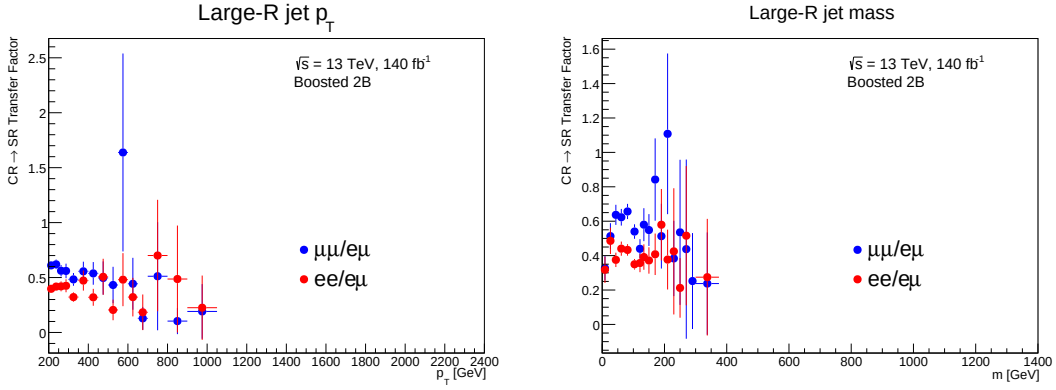


Figure 5.16: The transfer factors for the di-electron and di-muon channels for the large-R jet p_{T} (left) and large-R jet mass (right) in the Boosted 2B signal region.

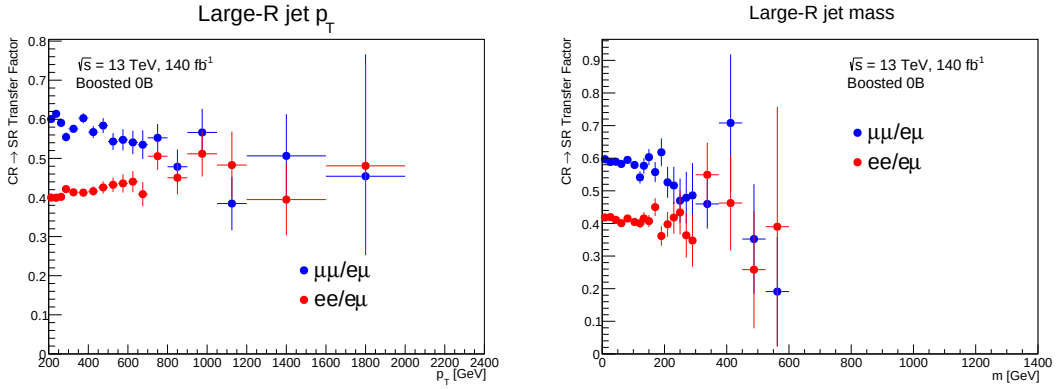


Figure 5.17: The transfer factors for the di-electron and di-muon channels for the large-R jet p_{T} (left) and large-R jet mass (right) Boosted 0B signal region.

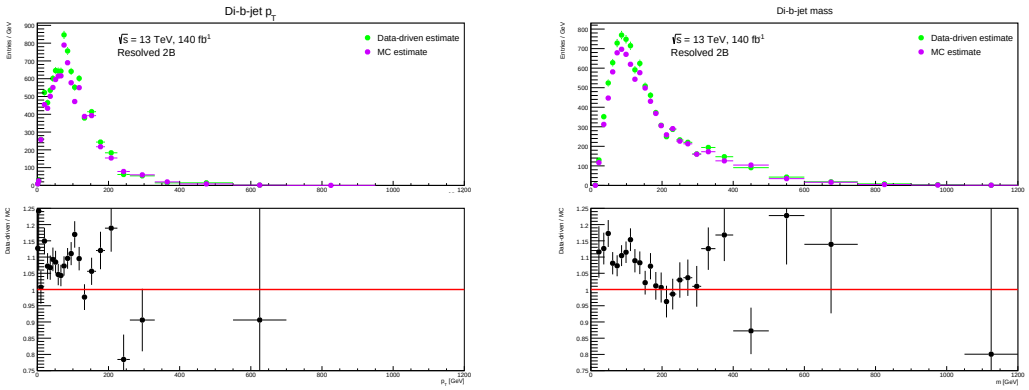


Figure 5.18: A comparison of the $t\bar{t}$ background estimation from data and simulation for the di-b-jet p_{T} (left) and di-b-jet mass (right) in the Resolved 2B signal region.

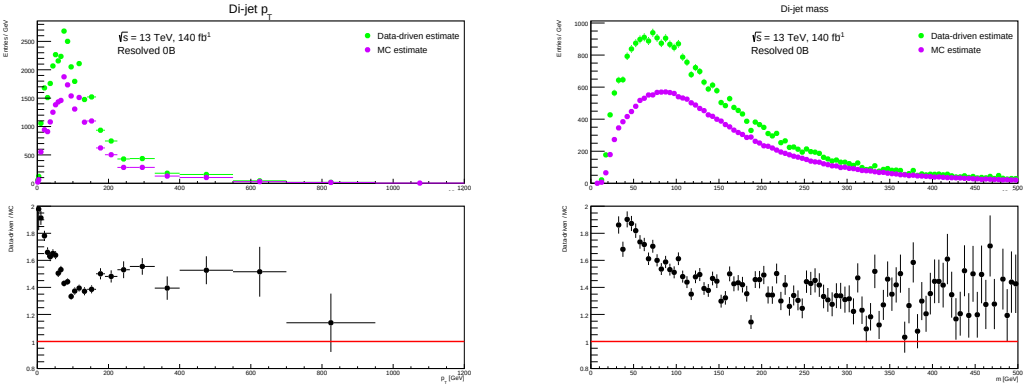


Figure 5.19: A comparison of the $t\bar{t}$ background estimation from data and simulation for the leading jet p_{T}, di-jet p_{T} (left) and di-jet mass (right) in the Resolved 0B signal region.

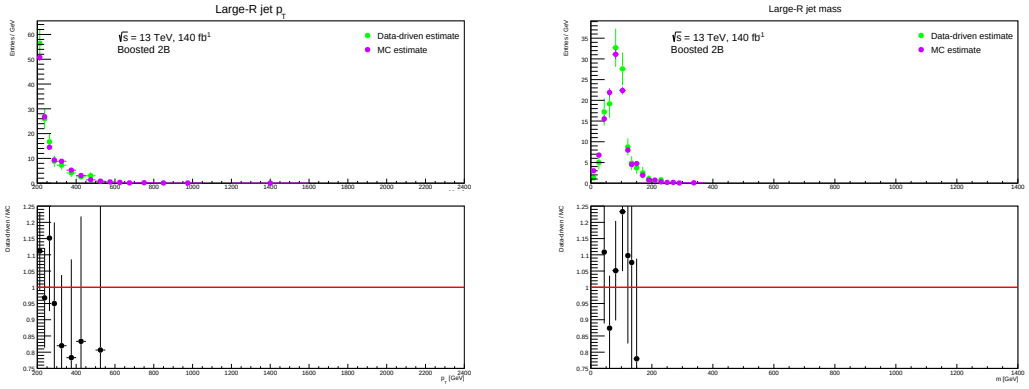


Figure 5.20: A comparison of the $t\bar{t}$ background estimation from data and simulation for the large-R jet p_{T} (left) and large-R jet mass (right) in the Boosted 2B signal region.

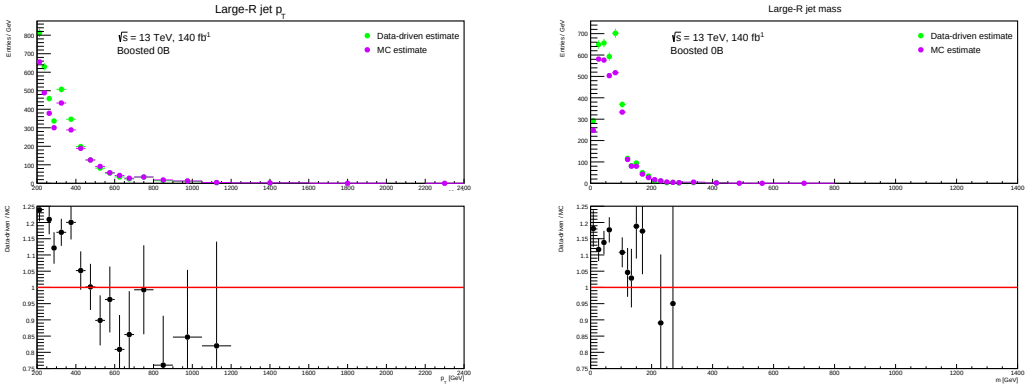


Figure 5.21: A comparison of the $t\bar{t}$ background estimation from data and simulation for the large-R jet p_{T} (left) and large-R jet mass (right) in the Boosted 0B signal region.

Signal region	Data-driven $t\bar{t}$ estimate (events)	MC $t\bar{t}$ estimate (events)
Resolved 2B	9116 ± 47	8454 ± 19
Resolved 0B	31989 ± 144	21204 ± 31
Boosted 2B	126 ± 9	122 ± 2
Boosted 0B	3672 ± 47	3151 ± 12

Table 5.10: The expected yields from $t\bar{t}$ events in the combined di-lepton channel with a data-driven estimation and a Monte Carlo estimation in the four signal regions of the measurement. The errors considered are statistical only.

Bin	Jet 1 Quantile	Jet 2 Quantile
1	100-85%	100-85%
2	100-85%	85-77%
3	100-85%	77-70%
4	100-85%	70-60%
5	100-85%	< 60 %
6	85-77%	85-77%
7	85-77%	77-70%
8	85-77%	70-60%
9	85-77%	< 60 %
10	77-70%	77-70%
11	77-70%	70-60%
12	77-70%	< 60 %
13	70-60%	70-60%
14	70-60%	< 60 %
15	< 60 %	< 60 %

Table 5.11: The bin numbering scheme for the flavour template based on the b-tagging efficiency of the two jets in the resolved signal regions.

5.7.3 Flavour fit

To improve the modelling of the Z+jets backgrounds in the signal regions, another profile likelihood fit[6] to data is carried out to extract their distributions. For each SR, templates based on a flavour-sensitive observable are derived from MC simulations. The chosen flavour-sensitive observable is the DL1r score. The fit is then carried out simultaneously in di-electron and di-muon channels to increase the statistical power.

The flavour template is based on quantiles of the DL1r score, defined based on the b-tagging efficiency. These are 100-85%, 85-77%, 77-70%, 70-60%, and below 60%. The score of PFlow jets is used in the Resolved SRs, while the score for VR trackjets is used in the Boosted SRs. The template thus consists of 15 bins, accounting for all possible unique combinations of jets falling into the five DL1r score quantiles. An outline of the template scheme is provided in Table 5.11. In the 2B regions, only bins 13-15 are populated, while in the 0B regions, only bins 1-3, 6, 7, and 10 are populated.

In each SR, the fit includes a number of floating parameters corresponding to the normalisation scale factors of the Z+jets signal and backgrounds. The exact number of scale factors considered depends on the combination of the background categories. In general, the categories *NB*, *1B*, *NC*, *1C*, and *L* are used. These correspond to the processes $Z + bb$, $Z + b$, $Z + cc$, $Z + c$ and $Z + j$, respectively. In the 1B ($Z + b$) category, one jet is of b flavour and the other jet can be of flavour c or light, while in the 1C ($Z + c$) category one jet is of c flavour and the other jet is of light flavour. In the L ($Z + j$) category, all jets are light. Although the Z+jets signal normalisation is considered as a floating parameter in the fit, the scale factor derived is not used as the aim of the measurement is to measure exactly this contribution.

In the Resolved and Boosted 2B SRs, the flavour categories considered in the fit correspond to NB and BxCL, obtained by merging all other categories; in the Resolved and Boosted 0B SRs, the categories considered are L, NC, and 1C, while the categories containing b-flavour are

[6] For a description of the profile likelihood fit procedure, see 4.3.2.

fixed to the MC estimation due to their small contribution.

Systematic uncertainties are propagated with the offset method. The fit is re-run for each systematic variation and different scale factors are extracted each time. These scale factors are then propagated to the unfolding procedure.

We present below the nominal-only results of the flavour fit. The pre-fit and post-fit distributions of the di-jet quantiles are shown in Figures 5.22 and 5.23, respectively. For the post-fit distribution, only the background contributions are scaled. The extracted scale factors in all flavour categories in all SRs are listed in Table ??.

Flavour Region	Scale Factor	Error
Resolved 2B		
NB	0.8777	±0.0042
BxCL	1.3529	±0.00354
Resolved 0B		
L	0.8887	±0.0006
NC	1.0047	±0.0190
1C	1.2453	±0.0083
Boosted 2B		
NB	1.0843	±0.0481
BxCL	2.6560	±0.4571
Boosted 0B		
L	0.9655	±0.0045
NC	1.3089	±0.1528
1C	1.4097	±0.0629

Table 5.12: The scale factors and relative errors for each combination of flavour categories in each signal region resulting from the flavour fit to suppress the Z+jets background.

$Z + c$ background contribution

There is an important remark to be made about the c-jet fraction in the Z+jets background in the Resolved 0B SR. As can be seen in Figure 5.22 and 5.23, the c-jet fraction grows significantly as the b-tagging working point tightens. This justifies the choice of analysing only the substructure of the leading jet in this region: it is found that the increase in the c-jet fraction is significantly higher for the subleading jet. By analysing only the leading jet, this background is effectively suppressed.

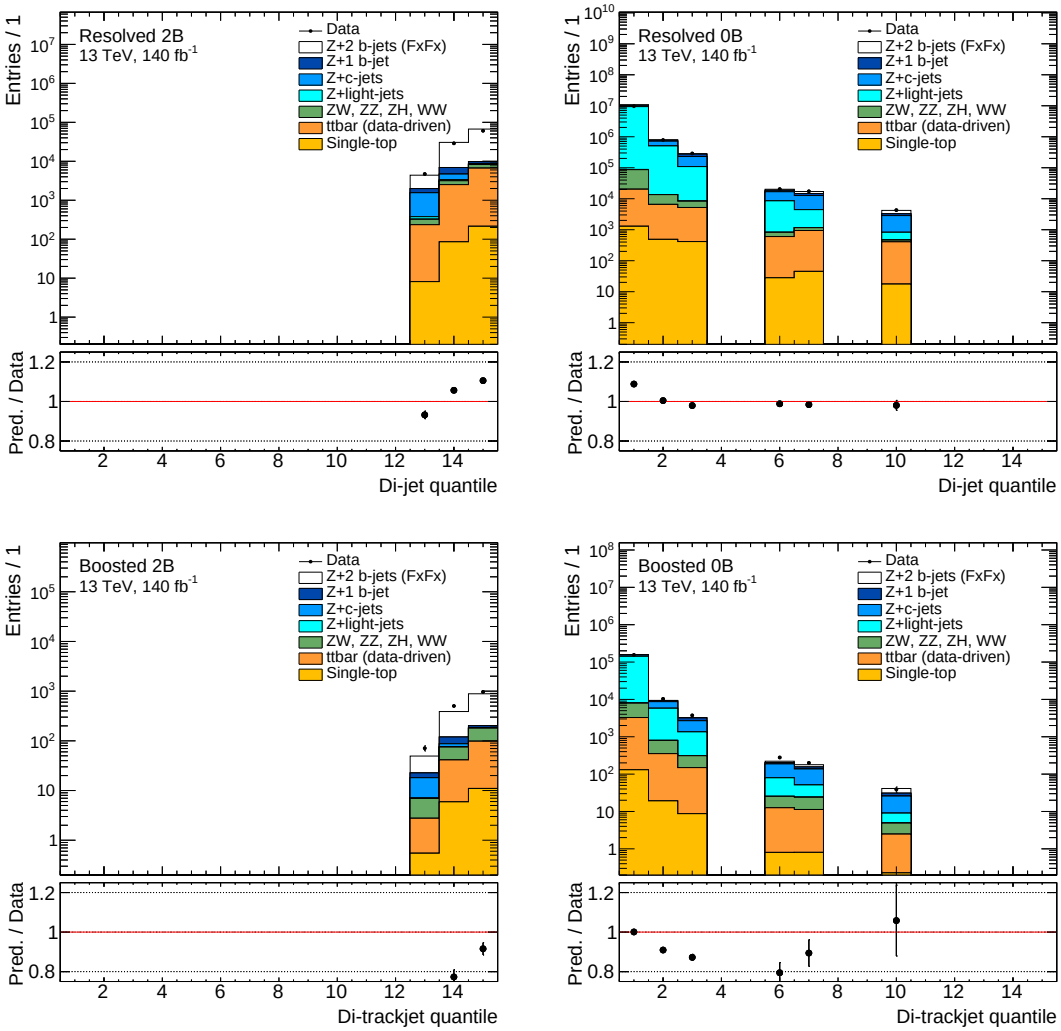


Figure 5.22: The pre-fit distributions of the di-jet quantiles determined based on the DL1r score in the four signal regions.

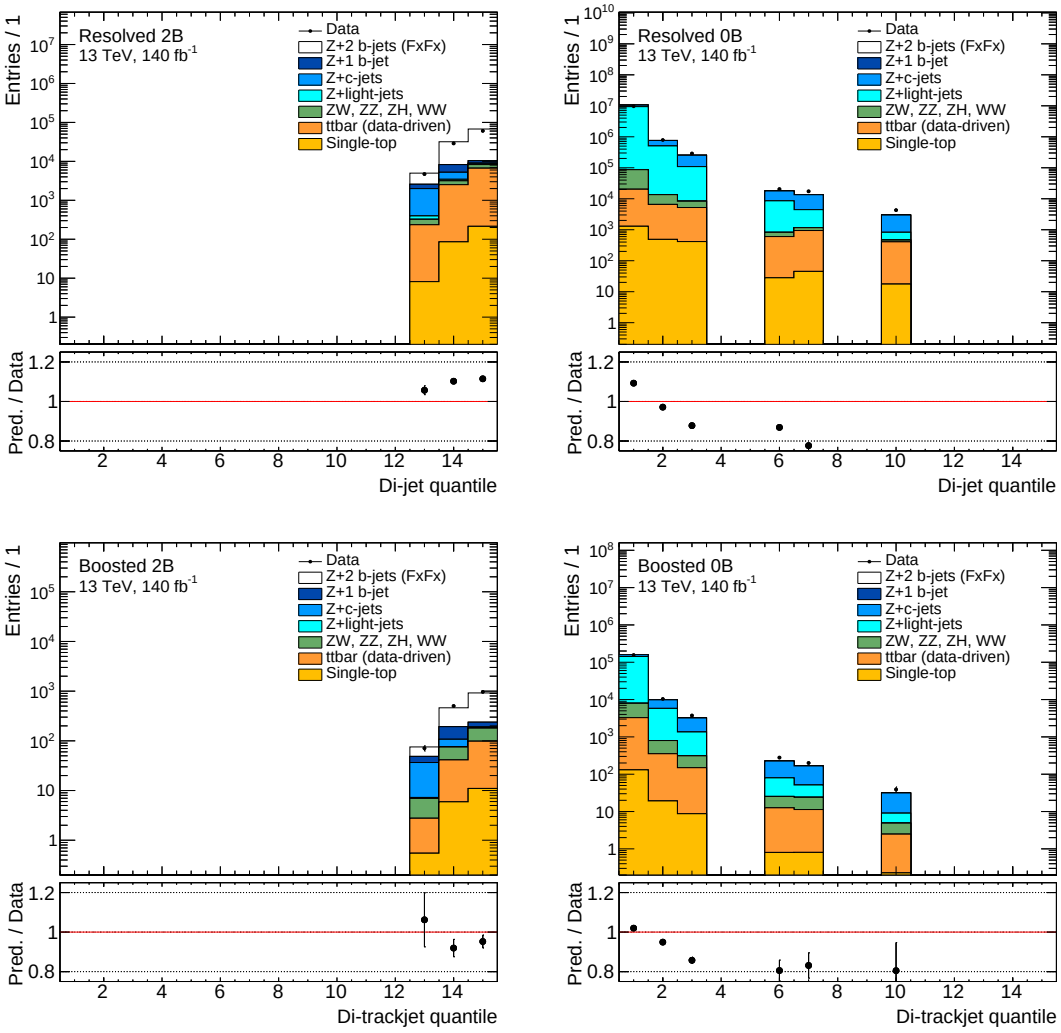


Figure 5.23: The post-fit distributions of the di-jet quantiles determined based on the DL1r score in the four signal regions.

5.7.4 Background estimation

In this section we limit ourselves to reporting the full yields for all signal and background processes in all four SRs for the di-lepton combination. These are reported in Table 5.13.

As can be seen, the two jet selection and tight b-tagging working point lead to high signal purity, particularly in the Resolved SRs. The data in Resolved 2B, Resolved 0B, Boosted 2B, and Boosted 0B is estimated to be composed of 80%, 84%, 67%, and 80% of signal, respectively. The low purity in the Boosted 2B region is caused in part by the low statistics in that region.

In Tables 5.14-5.16, the number of emissions in the primary Lund Jet Plane found in data and the predicted number of emissions in all SRs are listed. In the Resolved SRs, the number of emissions for leading jets is subdivided into slices of jet p_{T}, while in the Boosted SRs, the number of emissions in the large-R jet is shown. Also shown are the estimates for the number of emissions per jet. For this estimate, only jets with at least one track associated to them, i.e. those for which it is possible to measure the substructure, are considered.

Within each signal region and p_{T} slice, the number of emissions/jet is generally comparable for all signal and background processes. Any outliers are due to extremely low predicted yields, as can be seen by cross-referencing with Table 5.13. It is found that as the jet p_{T} increases, the number of emissions/jet also increases. Light jets have more emissions/jet than b-jets. However, this is not surprising considering the b-hadron reconstruction applied in the Resolved 2B region. Indeed, for the $Z + b$-jet and $Z + bb$ backgrounds in the Resolved 0B region, where no b-hadron reconstruction occurs, the predicted number of emissions/jet for b-jets is actually higher than for light jets, reflecting the greater particle multiplicity due to the decay of the heavy hadron.

	Resolved 2B Estimate (events)	**Resolved 0B** Estimate (events)	**Boosted 2B** Estimate (events)	**Boosted 0B** Estimate (events)
Data	94365 ± 307	10867200 ± 3300	1541 ± 39	171032 ± 414
$Z+ \geq 2$ b-jets	83328.5 ± 233.0	32063.7 ± 185.9	980.00 ± 6.72	438.11 ± 4.56
$Z + 1$ b-jet	5005.74 ± 76.86	250326 ± 578	144.35 ± 5.02	2572.23 ± 12.36
$Z + c$-jets	3862 ± 88	1548830 ± 1650	70.48 ± 4.87	22298.1 ± 52.1
$Z+$ light jets	602.63 ± 53.39	10058300 ± 5400	1.61 ± 1.04	141378 ± 150
$Z \to \tau^+\tau$	4.68 ± 2.84	1449.41 ± 58.91	0.25 ± 0.15	26.59 ± 2.06
$ZW/ZZ/ZH$	2326.75 ± 9.68	77444 ± 62	119.56 ± 2.17	5323.5 ± 14.9
$t\bar{t}$	9118.01 ± 74.40	31979.8 ± 144.9	125.34 ± 8.65	3672.23 ± 47.08
single-top	308.09 ± 6.56	2307.01 ± 18.75	17.47 ± 1.64	162.49 ± 4.88

Table 5.13: The observed and predicted number of events passing the selections for each signal region and the estimated contribution of the backgrounds. The values correspond to the di-lepton combination, and include the data-driven estimation for $t\bar{t}$.

	Resolved 2B					
	$20 \leq p_{\mathrm{T}}/\mathrm{GeV} < 60$		$60 \leq p_{\mathrm{T}}/\mathrm{GeV} < 100$		$p_{\mathrm{T}} \geq 100\mathrm{GeV}$	
	Total emissions	Emissions/jet	Total emissions	Emissions/jet	Total emissions	Emissions/jet
Data	67283±259	1.32±0.008	48981±221	1.88±0.014	44485±211	2.568±0.023
$Z+ \geq 2$ b-jets	56606.564±218.023	1.181±0.007	38156.002±128.027	1.751±0.010	33321.527±75.388	2.449±0.011
$Z + 1$ b-jet	4012.055±76.762	1.291±0.038	2192.471±39.161	1.971±0.061	2194.972±23.574	2.793±0.06
$Z + c$-jets	3772.692±93.361	1.395±0.054	1639.654±41.919	2.231±0.106	1267.179±22.06	2.997±0.109
Z+light jets	513.453±54.647	1.005±0.146	102.536±17.752	1.634±0.474	39.108±6.716	1.344±0.490
$Z \to \tau^+\tau^-$	7.024±3.029	2.393±2.310	2.335±1.368	1.349±1.162	1.579±1.008	81.541±2803.756
$ZW/ZZ/ZH$	1477.581±7.9	1.251±0.010	1230.494±7.067	1.775±0.017	1245.918±6.718	2.549±0.026
$t\bar{t}$	4333.312±51.470	1.438±0.027	6451.017±63.262	1.87±0.031	6695.861±63.788	2.522±0.045
single-top	145.053±4.407	1.218±0.055	171.482±4.988	1.766±0.086	225.854±5.665	2.457±0.114

Table 5.14: The observed and predicted number of emissions in the leading jet in the p_{T} slices 20-60 GeV, 60-100 GeV, and > 100 GeV in the Resolved 2B signal region and the estimated composition of the background contributions. Also shown are the observed and predicted number of emissions per jet. The values correspond to the di-lepton combination, and include the data-driven estimation for $t\bar{t}$.

	Resolved 0B					
	$20 \leq p_{\mathrm{T}}/\mathrm{GeV} < 60$		$60 \leq p_{\mathrm{T}}/\mathrm{GeV} < 100$		$p_{\mathrm{T}} \geq 100\mathrm{GeV}$	
	Total emissions	Emissions/jet	Total emissions	Emissions/jet	Total emissions	Emissions/jet
Data	18581239±4311	2.645±0.001	8619260±2936	4.005±0.003	8202805±2864	5.281±0.005
$Z+ \geq 2$ b-jets	56604.811±273.834	2.781±0.027	28207.233±129.444	4.13±0.043	26585.459±73.03	5.564±0.04
$Z + 1$ b-jet	461912.768±840.494	2.714±0.010	191539.359±350.697	4.157±0.018	176577.331±202.759	5.554±0.017
$Z + c$-jets	2673623.979±2451.425	2.829±0.005	1389932.869±1177.353	4.133±0.008	1408709.124±730.282	5.487±0.007
Z+light jets	17315919.58±7665.492	2.584±0.002	7449037.065±3367.023	3.958±0.004	6930429.392±2067.716	5.311±0.004
$Z \to \tau^+\tau^-$	2378.316±86.802	2.588±0.178	1147.021±37.971	4.031±0.307	1147.892±25.721	5.403±0.314
$ZW/ZZ/ZH$	94175.056±68.153	2.978±0.004	89886.703±67.693	3.904±0.007	122076.653±74.276	5.408±0.008
$t\bar{t}$	36281.737±155.678	3.169±0.028	46192.038±172.221	4.298±0.037	53711.749±184.143	5.562±0.049
single-top	3138.527±21.686	3.222±0.046	3304.92±22.416	4.393±0.069	3286.937±22.601	5.772±0.103

Table 5.15: The observed and predicted number of emissions in the leading jet in the p_{T} slices 20-60 GeV, 60-100 GeV, and > 100 GeV in the Resolved 0B signal region and the estimated composition of the background contributions. Also shown are the observed and predicted number of emissions per jet. The values correspond to the di-lepton combination, and include the data-driven estimation for $t\bar{t}$.

	Boosted 2B		**Boosted 0B**	
	Total emissions	Emissions/jet	Total emissions	Emissions/jet
Data	5883±77	3.818±0.109	1117161±1057	6.566±0.017
$Z+ \geq 2$ b-jets	3473.521±58.937	3.544±0.027	2641.342±51.394	6.966±0.085
$Z + 1$ b-jet	559.601±23.656	3.877±0.151	17211.669±131.193	6.756±0.035
$Z + c$-jets	296.583±17.222	4.208±0.321	148753.669±385.686	6.692±0.017
Z+light jets	5.083±2.255	3.148±2.386	937745.814±968.373	6.664±0.008
$Z \to \tau^+\tau^-$	1.181±1.087	4.637±3.004	164.193±12.814	6.194±0.518
$ZW/ZZ/ZH$	419.784±20.489	3.511±0.072	34065.964±184.57	6.433±0.019
$t\bar{t}$	477.672±21.856	3.82±0.297	24741.815±157.295	6.798±0.093
single-top	58.785±7.667	3.365±0.358	1155.458±33.992	7.202±0.233

Table 5.16: The observed and predicted number of emissions in the large-R jet in the Boosted 2B and Boosted 0B signal region and the estimated composition of the background contributions. Also shown are the observed and predicted number of emissions per jet. The values correspond to the di-lepton combination, and include the data-driven estimation for $t\bar{t}$.

5.8 Systematics

This section is dedicated to outlining the systematic uncertainties which afflict the measurement. These can stem from imperfect knowledge and simulation of the detector or imperfect calibration of reconstructed physics objects (detector-level systematic uncertainties), but also from incomplete knowledge of the contribution of background processes and inexact theoretical understanding.

5.8.1 Detector-level systematic uncertainties

The detector-level uncertainties described in this section are propagated through the flavour fit described in Section 5.7.3 and the unfolding procedure described in Section 5.10.

Each uncertainty is evaluated twice for all signal and background components in the Monte Carlo: once with an upwards variation of one standard deviation $+1\sigma$, and once with a downwards variation of -1σ. This leads to two alternative distributions for the event yields with respect to the nominal result.

When both the upwards and downwards variations result in a fluctuation in the same direction, the resulting variation is assumed to be of equal magnitude in both directions. This is known as symmetrisation. Asymmetric variations, on the other hand, are maintained separate though the propagation of the uncertainties.

Below we outline the categories of the systematic variations considered.

- **Lumi:** The luminosity uncertainty (Lumi) on the full Run 2 dataset, equivalent to 140.1 fb^{-1}, is equal to $\pm 0.8\%$. The luminosity is determined through a measurement by dedicated detectors such as LUCID [157]. The measurement is then calibrated through van der Meer scans, performed once per year, and the calibration is propagated to the physics dataset. A detailed description of this procedure is available in Reference [136].

- **Lepton:** These uncertainties affect the electron and muon triggers, reconstruction, identification, isolation efficiency, momentum scale, and momentum resolution. For muons, uncertainties on aspects such as the track-to-vertex association and sagitta-bias corrections are also included. The reconstruction, identification and isolation efficiency corrections are obtained via scale factors applied to the MC to match data. The momentum scale and resolution uncertainties are derived through a calibration.

- **MET:** Uncertainties associated to the missing transverse energy (MET) affect the momentum scale and resolution. They are derived from soft-term contributions from tracks, where the soft-terms refer to energy deposits in the calorimeter not associated to jets. The hard terms are accounted for by electron, muon, and jet uncertainties.

- **PRW:** These uncertainties are related to the reweighting of pile-up vertices in simulation to match the data.

- **TRK:** Tracking uncertainties (TRK) are due to inclusive tracking efficiencies, track efficiency inside jets, and the reconstruction fake rate. The first of these is related to the uncertainties in the modelling of the particle interaction with matter. Their effect is determined by varying the material configurations and changing the model of the interaction. The efficiency inside jets uncertainty comes from the dense tracking environment in high p_{T} jets, which can decrease the tracking resolution. The reconstruction fake rate is calculated through a calibration which utilises a control region enriched by fake tracks to estimate the real and fake track fractions.

- **Jet:** Uncertainties related to jets include those stemming from the jet energy scale (JES) and jet energy resolution (JER). These quantities are described in Section 3.3.7. JES uncertainties are found via correction factors derived from in situ measurements of certain reference objects. These include Z bosons, photons, and low p_{T} jet systems. The p_{T} of the jet is balanced against that of another object, leading to a correction factor which is then applied. Other calibrations are performed to derive the uncertainties, including an in situ calibration to balance forward jets against central jets, very high p_{T} jets, pile-up, mismodelling of the flavour response and composition, and jet vertex tagging efficiency.

- **BTagging:** These uncertainties are related to the flavour tagging efficiency. They are derived from the relative calibrations in enriched samples of the different categories of b, c, and light flavours, and propagated to the final calibration of the DL1r tagger. The analyses carried out to derive the flavour tagging efficiencies are typically at low p_{T} due to the higher available statistics, necessitating an extrapolation to higher p_{T} ranges. These extrapolations lead to an additional set of uncertainties.

5.8.2 Background systematic uncertainties

The background uncertainties account for variations of the yields due to Monte Carlo mismodelling and uncertainties in the background predictions. A number of MC generator modelling uncertainties are propagated through the aspects of the measurement which rely on simulation, such as MC modelled backgrounds, MC components of the flavour fit, as well as in the data-driven $t\bar{t}$ extraction.

Uncertainties on the Z+jets backgrounds are affected by QCD scale uncertainties. They are derived by varying the renormalisation and factorisation scales μ_R and μ_R by a factor of 0.5 and 2 while constraining $0.5 < \mu_R/\mu_F < 2$. The flavour fit is repeated for each of these variations and the error estimate is derived from an envelope of the post-fit results.

Other sources of error include modelling uncertainties, accounted for by varying the MC sample used. The Sherpa 2.2.11 [152] sample is to be used to evaluate these errors.

The $t\bar{t}$ background is afflicted by an extrapolation uncertainty from the control region to the signal region. The uncertainty in other background processes, namely di-boson, single-top, ZH, and $Z \to \tau^+\tau^-$, is given by the Pythagorean sum of the individual components. This is possible due to the small contribution of these backgrounds. For the di-boson processes, the QCD scale uncertainties are evaluated as for the Z+jets backgrounds.

Lastly, PDF uncertainties correspond to the root mean squared of the NNPDF 3.0 NNLO PDF replica sample [143], while α_s uncertainties are derived by varying the value of α_s in the PDF sample.

5.9 Detector-level results

In this section, we show the detector-level results with systematics for all observables in all SRs to analyse the data/MC agreement. Figures 5.24-5.26 show the kinematic observables, the jet angularities, and the Lund Jet Plane for jets in the Resolved 2B Region. Some kinematic observables display poor modelling, particularly $p_{\mathrm{T},bb}/M_{bb}$ and ΔR_{bb} in the central region.

The jet angularities λ do not agree well with data, particularly for small values. The comparison with the variation $\dot{\lambda}$, defined in Section 1.4.4, confirm the success of the b-hadron reconstruction, especially when compared for the equivalent plot in the Resolved 0B region (Figure 5.28).

The Lund Jet Plane is presented in a two-dimensional version for just data, and in the "unrolled" version for the purpose of MC comparison. The Lund Plane shows generally poor

agreement between data and MC, though this varies depending on the region of the Lund Plane considered and on the jet kinematics. It should also be noted how as the jet p_{T} increases, the phase space accessed by the emissions within the jet also increases.

Figures 5.27-5.29 show the kinematic observables, angularities, and Lund Jet Plane in the Resolved 0B Region. The agreement of the kinematic observables is significantly improved when compared to the Resolved 2B region, fully agreeing within uncertainties. The modelling of the angularities and the Lund Jet Plane remains poor.

The kinematic observables measured in the Boosted 2B region are shown in Figure 5.30, while the Lund Jet Plane for the large-R jet is shown in Figure 5.31. The predicted distributions of kinematic observables again do not agree with data, particularly for the jet mass. The modelling of the distance between subjets ΔR_{bb} agrees better with data when compared to the equivalent plot in the Resolved 2B region. For what concerns the Lund Plane, there is still poor agreement between data and MC throughout the entire span of the plane.

Lastly, the Boosted 0B kinematic plots are shown in Figures 5.32 and the Lund Jet Plane is shown in Figure 5.33. Again, the kinematic observables present significantly better agreement with data when compared to the Boosted 2B region. The Lund Plane, however, still shows generally poor agreement.

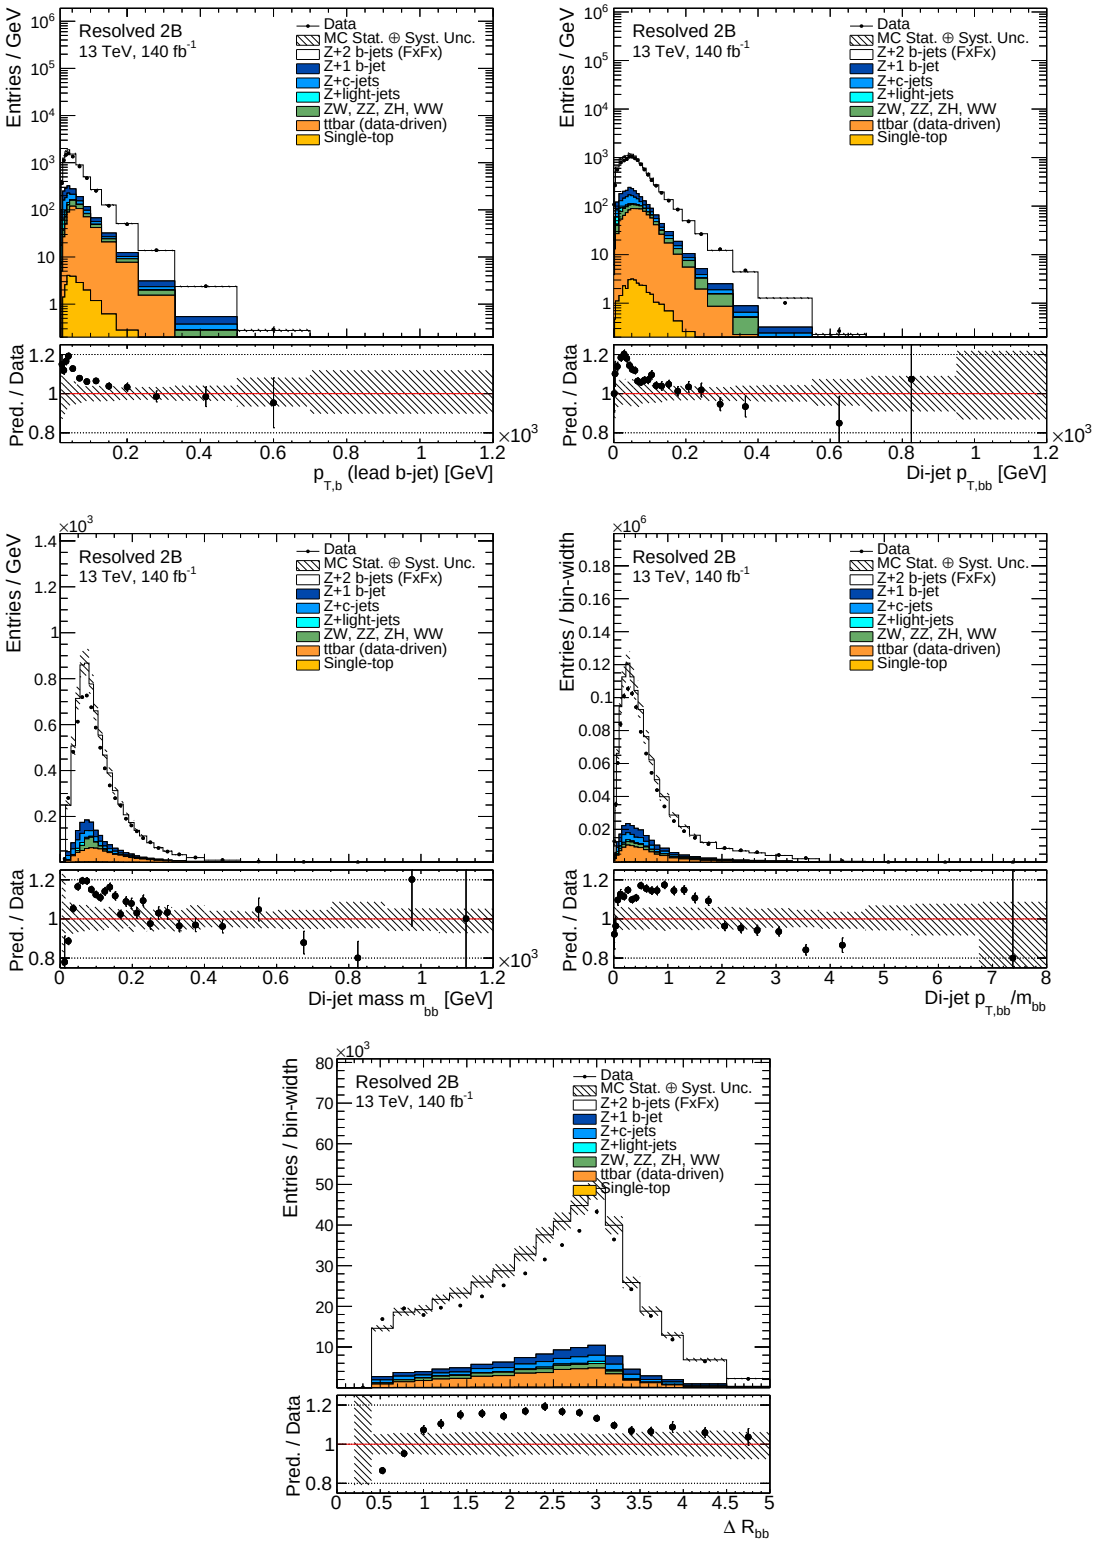


Figure 5.24: The data/MC agreement of the kinematic observables $p_{\mathrm{T},b}$, $p_{\mathrm{T},bb}$, M_{bb}, $p_{\mathrm{T},bb}/M_{bb}$ and ΔR_{bb} in the Resolved 2B SR.

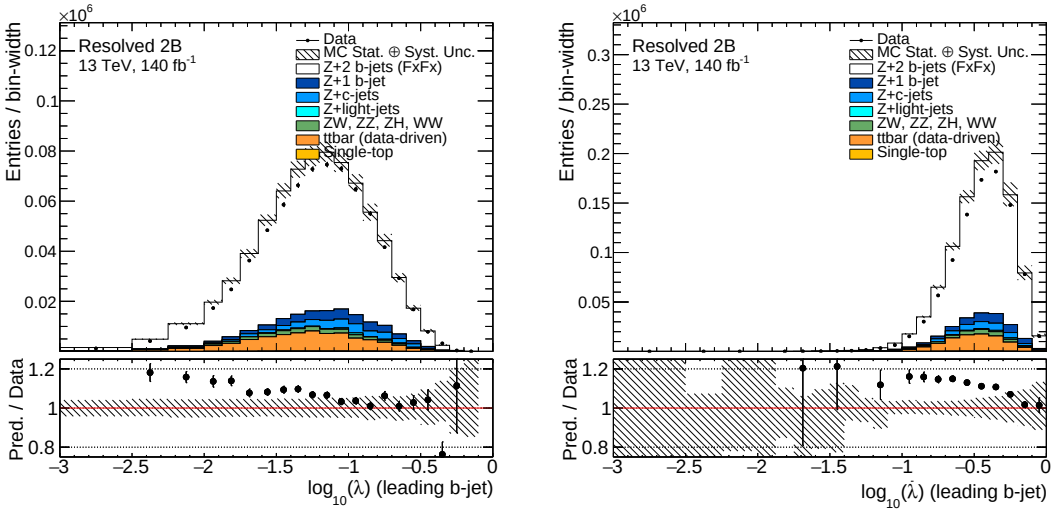


Figure 5.25: The logarithm of the jet angularities λ (left) and the $\dot{\lambda}$ variation (right) for the leading b-jet in the Resolved 2B SR.

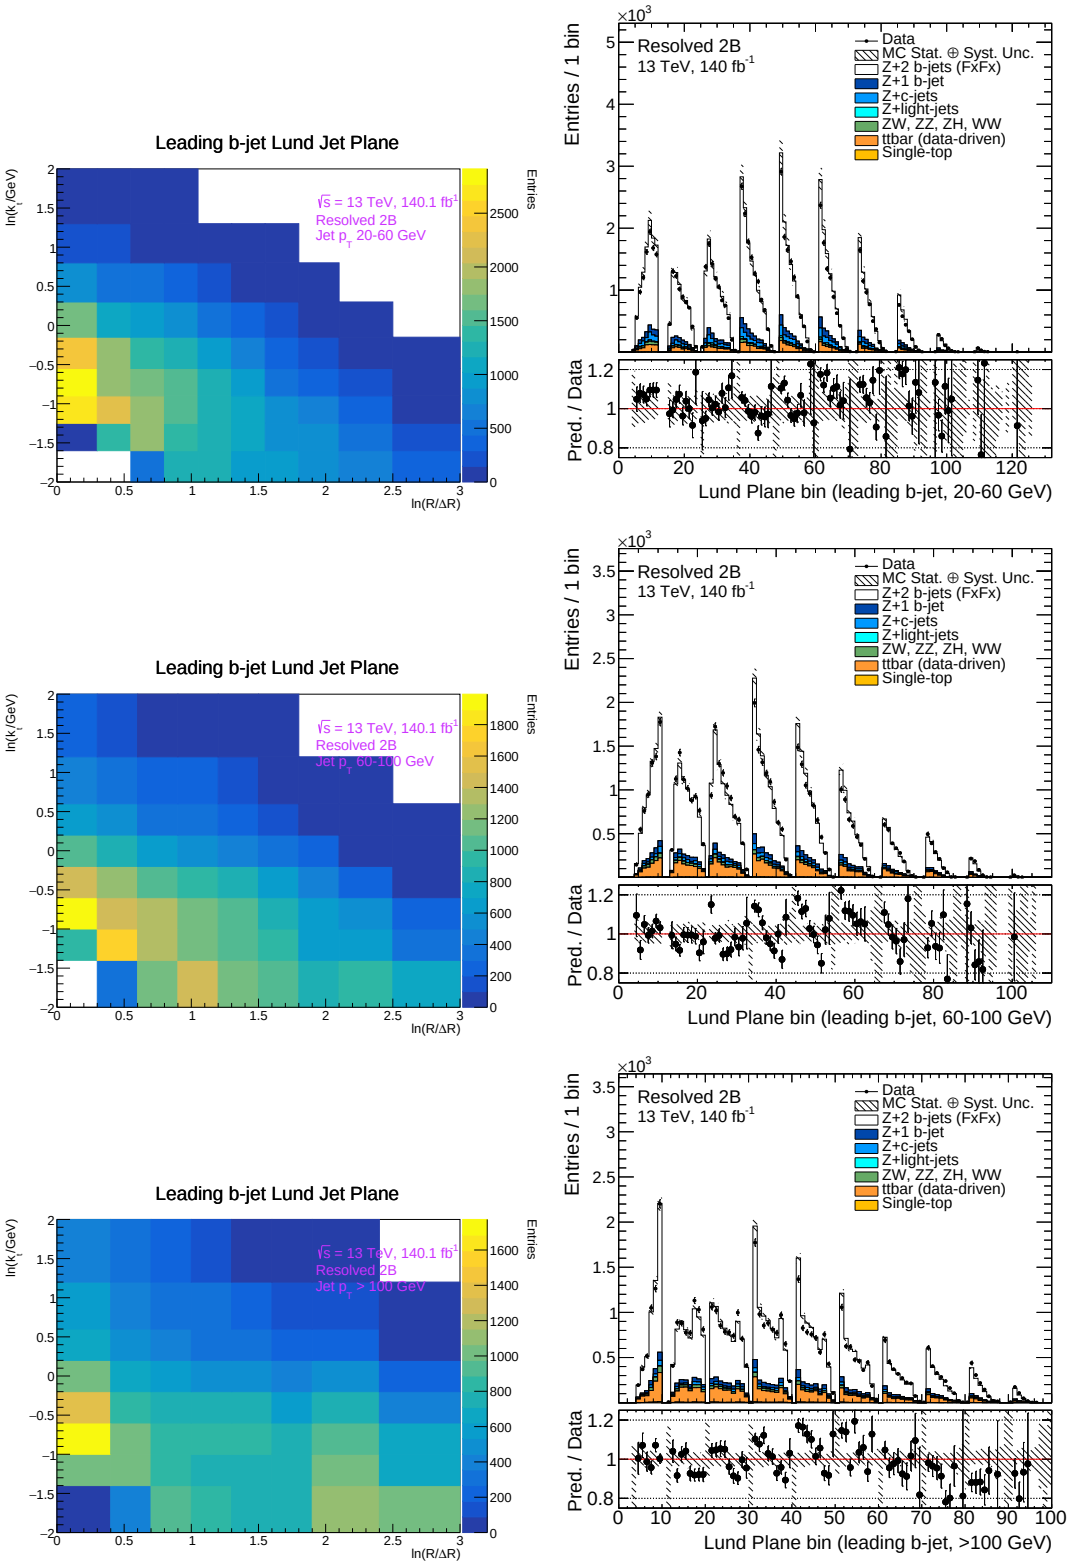


Figure 5.26: The Lund Jet Plane for the leading b-jet falling within the p_{T} intervals 20-60 GeV (top), 60-100 GeV (centre), and above 100 GeV (bottom) in the Resolved 2B SR. The left column shows the nominal-only two-dimensional Lund Jet Plane, while the right column shows the one-dimensional analogue.

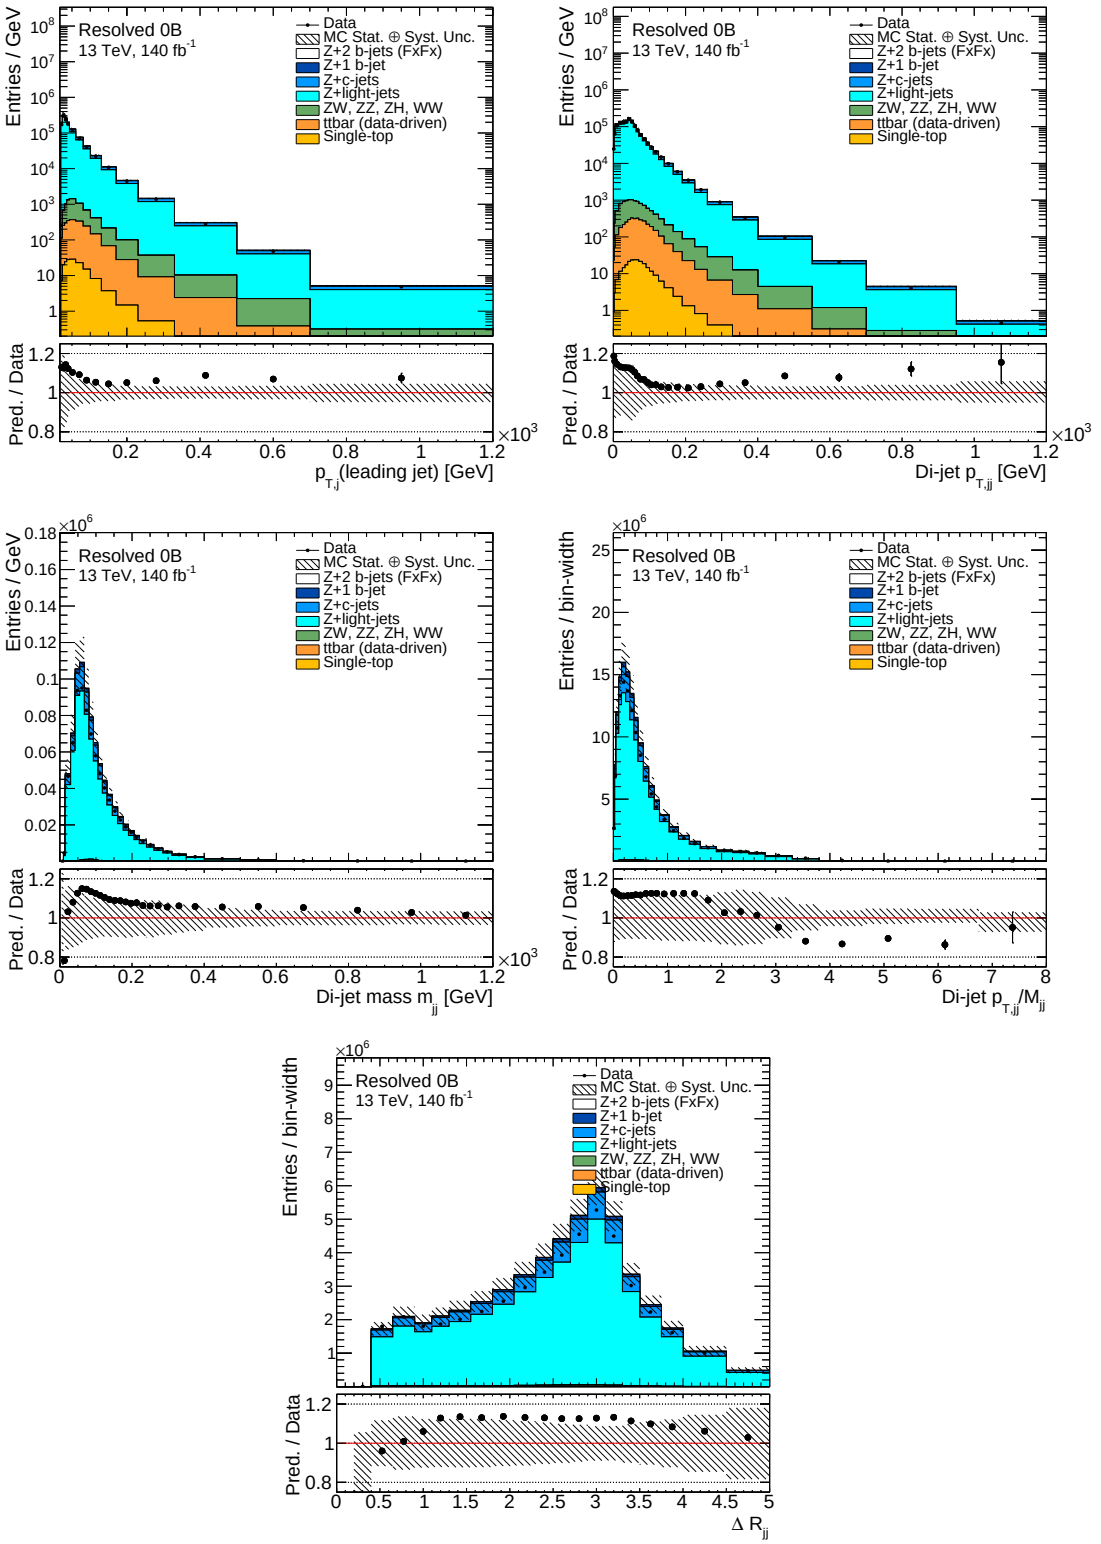


Figure 5.27: The data/MC agreement of the kinematic observables $p_{\mathrm{T},j}$, $p_{\mathrm{T},jj}$, M_{jj}, $p_{\mathrm{T},jj}/M_{jj}$ and ΔR_{jj} in the Resolved 0B SR.

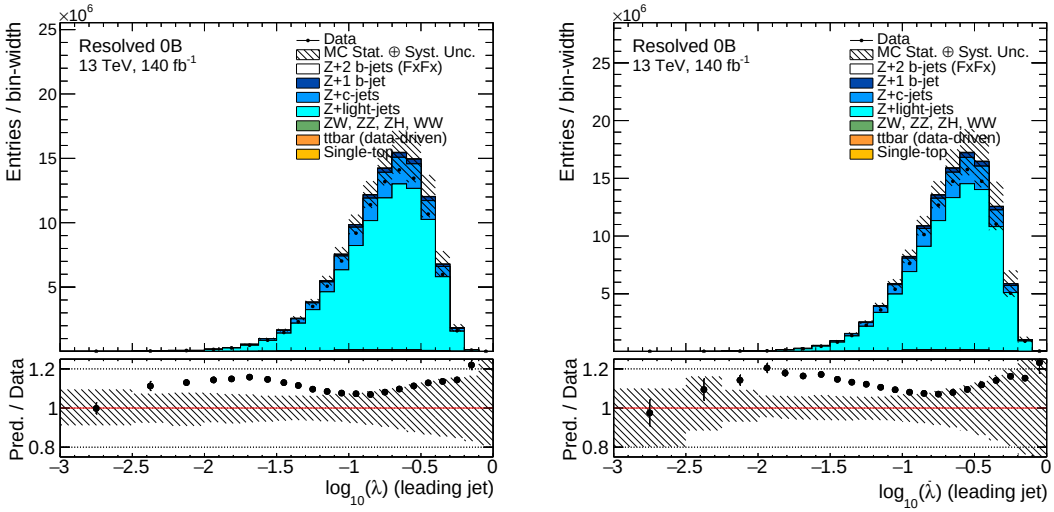


Figure 5.28: The logarithm of the jet angularities λ (left) and the $\dot{\lambda}$ variation (right) for the leading jet in the Resolved 0B SR.

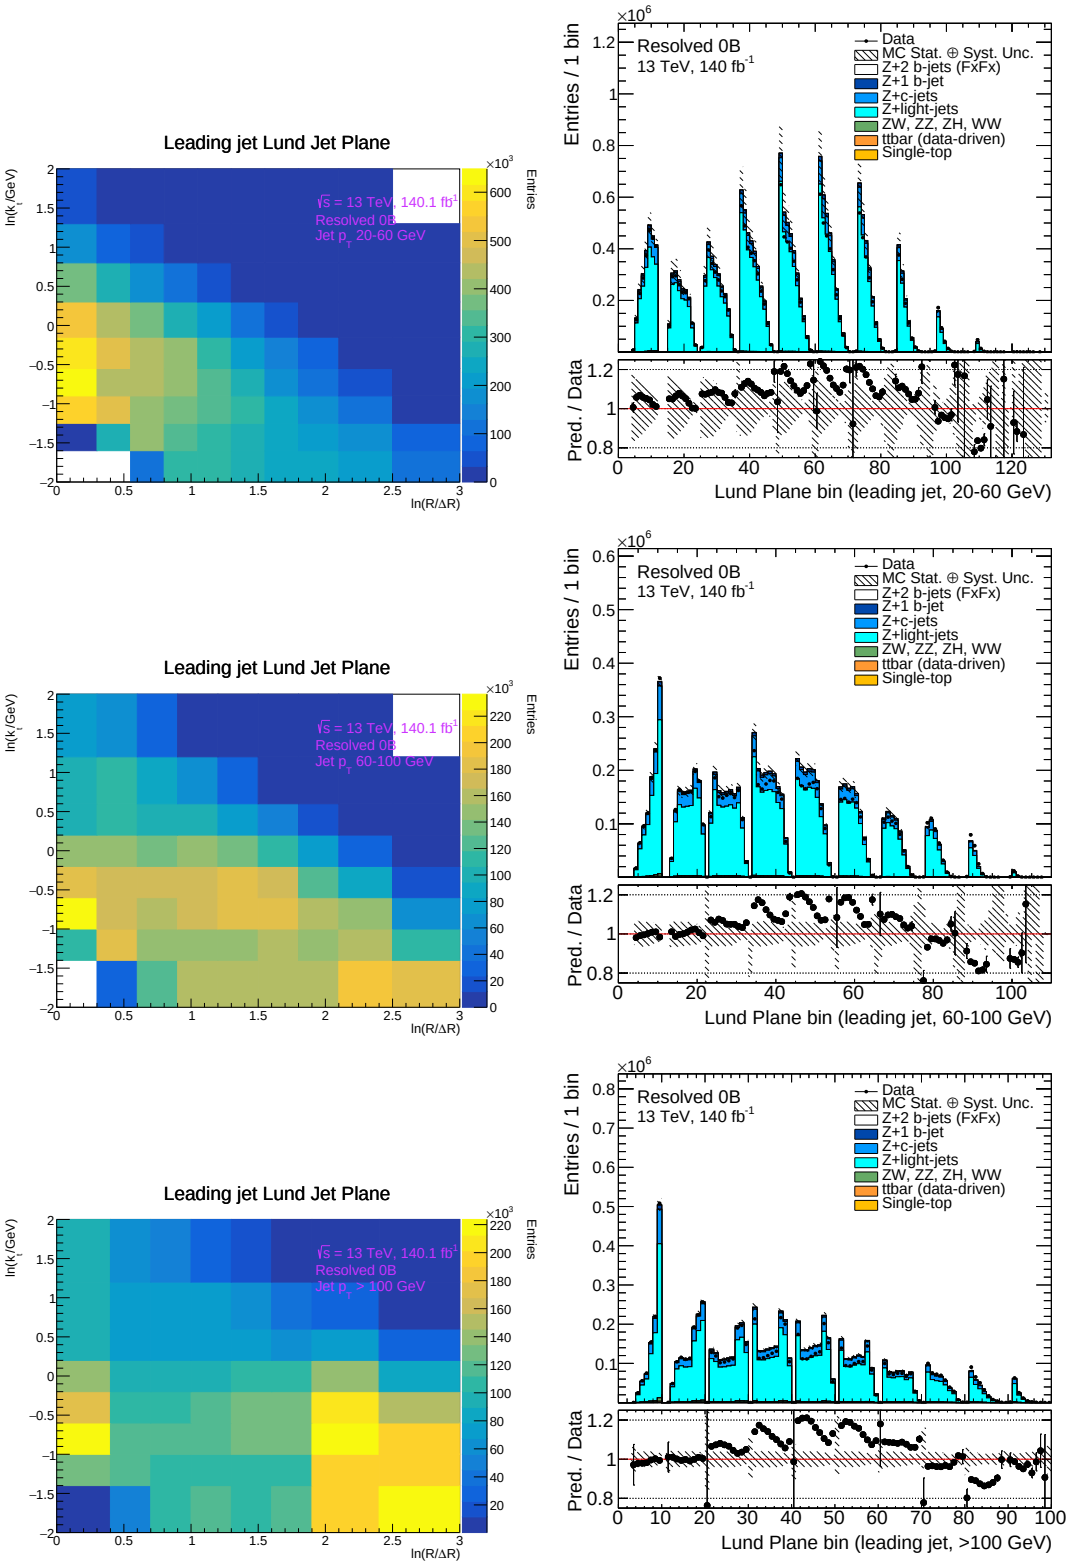


Figure 5.29: The Lund Jet Plane for the leading jet falling within the p_{T} intervals 20-60 GeV (top), 60-100 GeV (centre), and above 100 GeV (bottom) in the Resolved 0B SR. The left column shows the nominal-only two-dimensional Lund Jet Plane, while the right column shows the one-dimensional analogue.

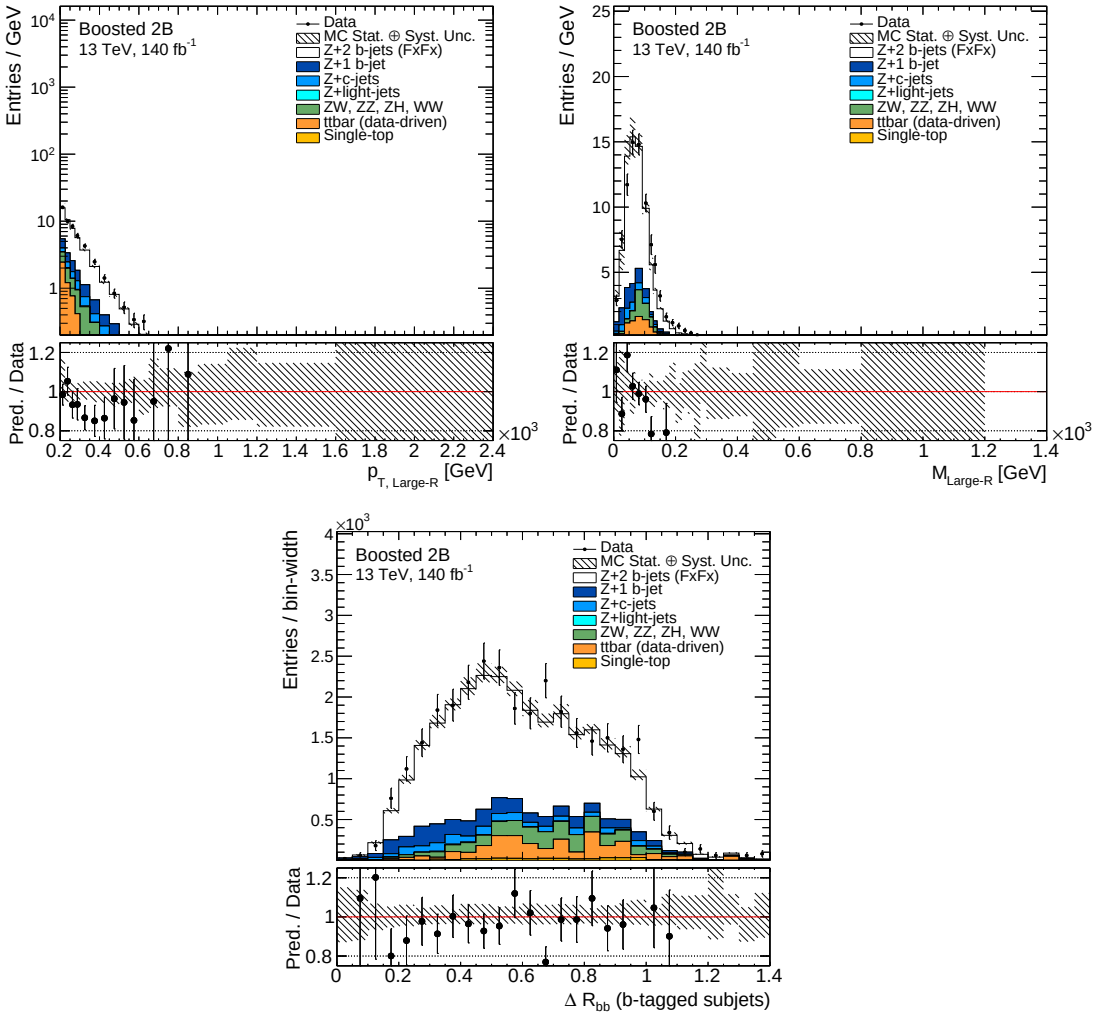


Figure 5.30: The data/MC agreement of the kinematic observables $p_{\mathrm{T,large-R}}$, $M_{\mathrm{large-R}}$, and $\Delta R_{bb,\mathrm{large-R}}$ in the Boosted 2B SR.

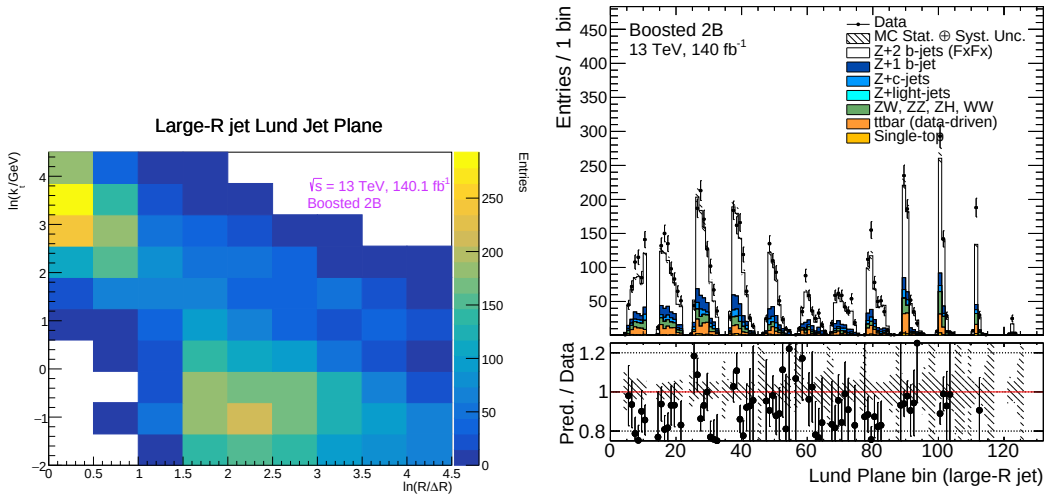


Figure 5.31: The Lund Jet Plane for the leading large-R jet in the Boosted 2B SR. The left column shows the nominal-only two-dimensional Lund Jet Plane, while the right column shows the one-dimensional analogue.

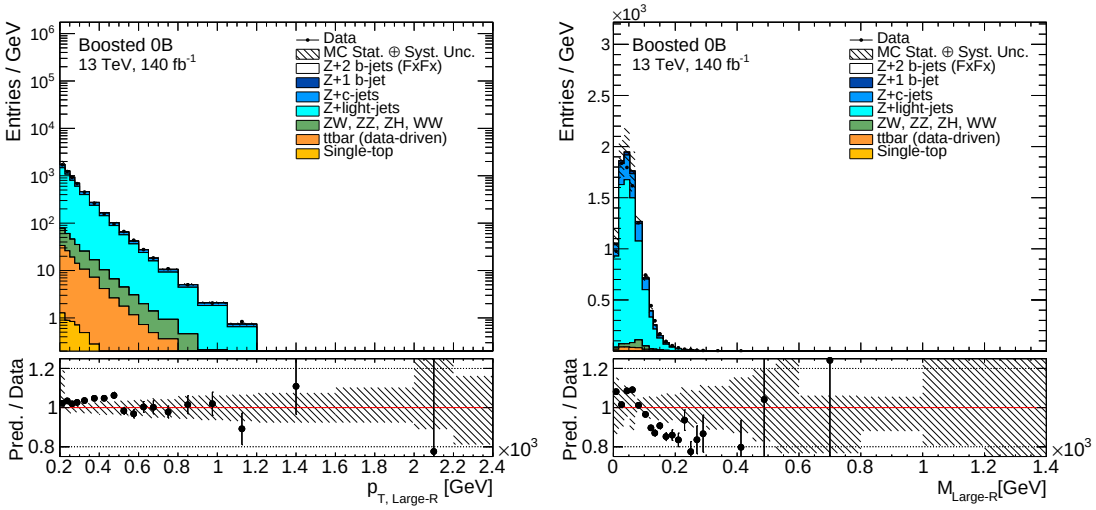


Figure 5.32: The data/MC agreement of the kinematic observables $p_{\mathrm{T,large-R}}$ and $M_{\mathrm{large-R}}$ in the Boosted 0B SR.

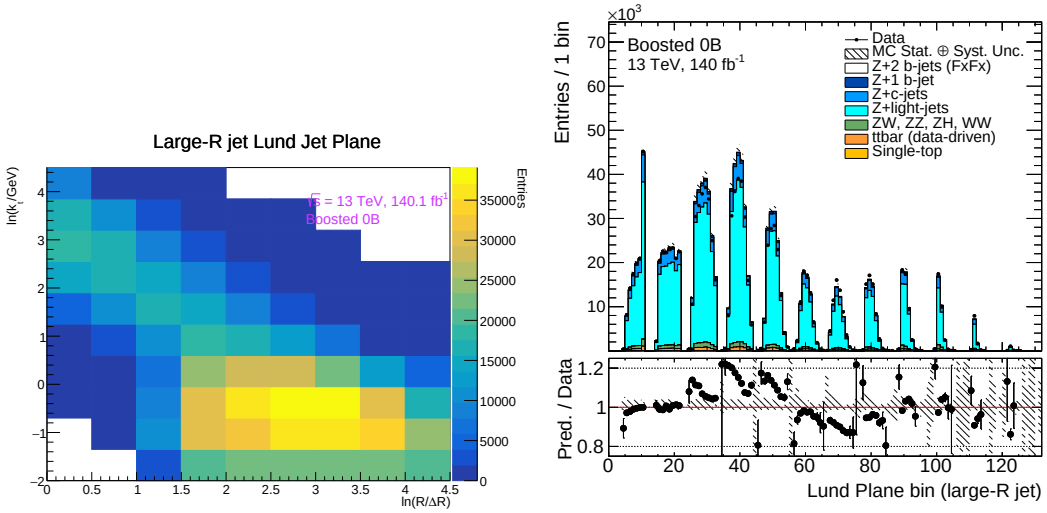


Figure 5.33: The Lund Jet Plane for the leading large-R jet in the Boosted 0B SR. The left column shows the nominal-only two-dimensional Lund Jet Plane, while the right column shows the one-dimensional analogue.

5.10 Unfolding

The unfolding procedure corrects the data for detector effects. Inevitably, due to the imperfect reconstruction that the machine offers, distortions arise in the "ideal" particle-level distributions of the observables measured. This passage from particle-level distributions to detector-level distributions is known as *folding*.

Unfolding is the inverse process: given a detector-level distribution of some observable, we wish to recover the particle-level distribution. Several unfolding techniques exist, each reliant on statistical methods to recover the true distribution. In the measurement described in this thesis, the *Iterative Bayesian Unfolding* (IBU) [158] procedure is implemented [7].

The problem of unfolding can be generally formulated in terms of the construction of a *response matrix* $\mathbf{M}$, which maps the particle-level values of a list of observables $\vec{x}$ to the detector-level values $\vec{y}$:

$$\mathbf{M}\vec{x} = \vec{y}. \tag{5.1}$$

In general, $\mathbf{M}$ is not invertible. The IBU procedure relies on Bayes' theorem to construct the matrix $\mathbf{M}$.

The inverse matrix $\mathbf{M}^{-1}$ can be written as

$$M^{-1}(x_i|y_j) = \frac{M(y_j|x_i)P_0(x_i)}{\sum_{k=1}^{N} M(y_j|x_k)P_0(x_k)}, \tag{5.2}$$

where $M(y_j|x_i)$ represents the probability that a given particle-level value which falls within the bin i of the distribution is reconstructed in bin j at detector-level. The quantity $P_0(x_i)$ is the prior distribution, and represents our knowledge of the true distribution of the observable. For the first iteration, the prior is set to the particle-level distribution obtained from the MC. The IBU procedure is then iteratively run, and each time the prior distribution is updated with the result of the previous step.

The iteration of the IBU procedure is done to reduce the bias which can be introduced by the prior distribution. On the other hand, each iteration introduces statistical fluctuations in the matrix $\mathbf{M}^{-1}$. For this reason, the number of iterations must be optimised to balance the two effects. In this measurement, the number of iterations is set to two.

Beyond bin migrations, the unfolding procedure must also account for fakes and inefficiencies. Fakes are events which pass the detector-level selections but not the particle-level selections, whereas inefficiencies pass the particle-level selections but have no correspondent detector-level value. By definition, these are not accounted for in the response matrix and its inverse.

A fake correction term is applied to correct for these effects. It is a multiplicative factor applied to the background-subtracted data and MC distributions before the unfolding procedure. Typically, the largest corrections occur at the edge of the phase space definitions.

The response matrices for a selection of observables in the four signal regions are shown in Figures 5.34-5.36. The fraction of entries which fall on the diagonal of the response matrices for each observable (purity) is found in Table 5.17 for the Resolved SRs and in Table 5.18 for the Boosted SRs.

5.10.1 Unfolding the Lund Jet Plane

At detector-level, the measurement of the Lund Jet Plane is limited to charged particles due to the higher resolution of the tracker. However, the pattern of emissions in the primary

[7] We also plan to develop alternative unfolding strategies, such as Multifold [159], so as to be able to improve upon the final result. More details on this strategy can be found in Appendix B.

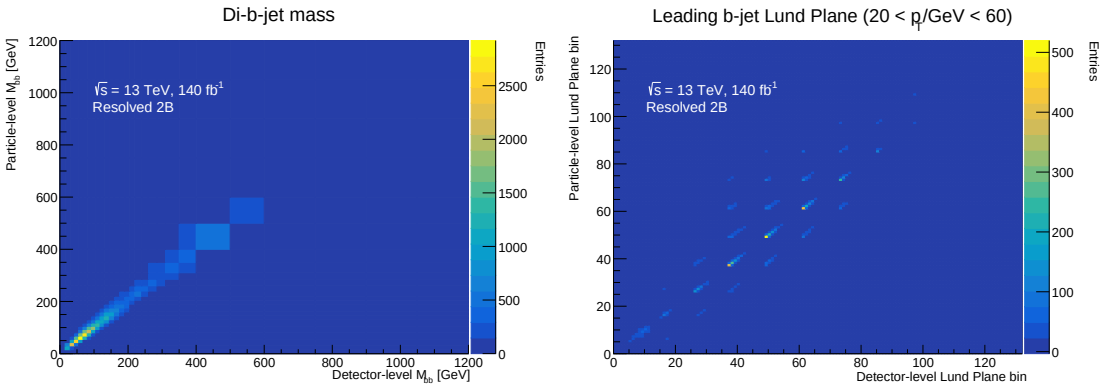


Figure 5.34: The response matrices for the b-jet mass (left) and one-dimensional leading jet Lund Plane in the 20-60 GeV p_{T} range (right) in the Resolved 2B SR.

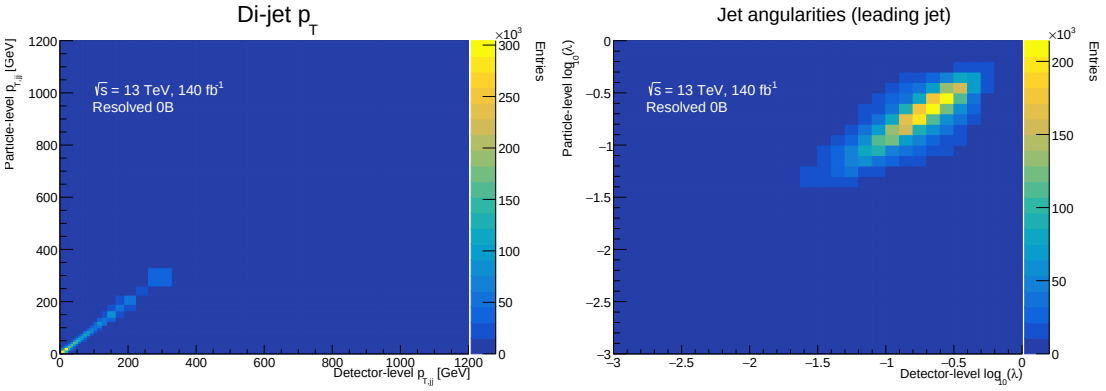


Figure 5.35: The response matrices for the di-jet p_{T} (left) and the logarithm of the jet angularity (right) in the Resolved 0B SR.

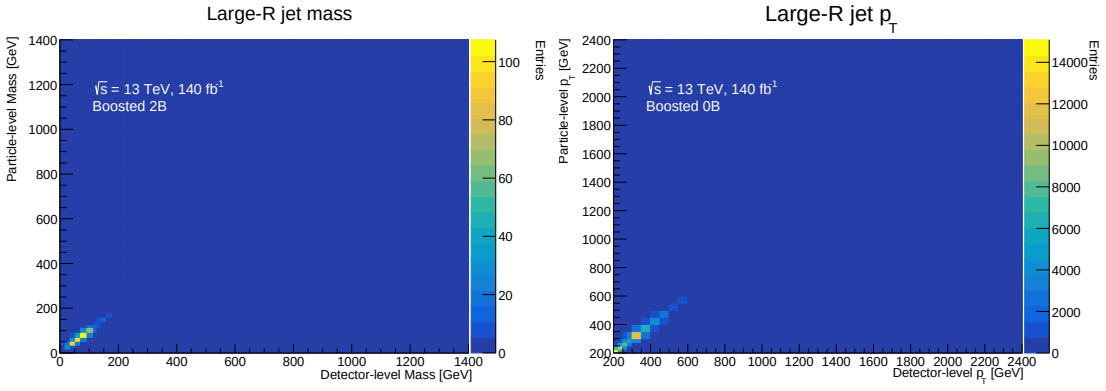


Figure 5.36: The response matrices for the large-R jet mass in the Boosted 2B SR (left) and for the large-R jet p_{T} in the Boosted 0B SR (right).

Observable	Purity Resolved 2B	Purity Resolved 0B
$p_{\mathrm{T},bb/jj}$	0.38	0.40
$M_{bb/jj}$	0.37	0.39
$p_{\mathrm{T},bb/jj}/M_{bb/jj}$	0.56	0.52
$\Delta R_{bb/jj}$	0.83	0.76
$\log_{10}(\lambda)$	0.11	0.25
Lund Jet Plane (leading jet, p_{T} 20-60 GeV)	0.23	0.33
Lund Jet Plane (leading jet, p_{T} 60-100 GeV)	0.28	0.35
Lund Jet Plane (leading jet, $p_{\mathrm{T}} \geq 100$ GeV)	0.31	0.37

Table 5.17: The fraction of entries on the diagonal of the response matrices for the observables measured in the Resolved 2B and Resolved 0B SRs.

Observable	Purity Boosted 2B	Purity Boosted 0B
$p_{\mathrm{large-R}}$	0.56	0.54
$M_{\mathrm{large-R}}$	0.57	0.52
ΔR_{bb}	0.50	-
Lund Jet Plane (large-R jet)	0.25	0.37

Table 5.18: The fraction of entries on the diagonal of the response matrices for the observables measured in the Boosted 2B and Boosted 0B SRs.

Lund Plane is roughly independent of charge, as it is a consequence of the approximate isospin symmetry of QCD[8]. The fiducial definition of the primary Lund Plane for the purpose of unfolding this data thus includes neutral particles. All particles in the fiducial definition must satisfy the 900 MeV cut in transverse momentum, as a variation in this cut would lead to differences in particle multiplicity which *can* affect the structure of the primary plane.

Figure 5.37 shows how the inclusion of the neutral particles in the fiducial definition of the Lund Plane does not significantly alter it. The greatest variation can be seen in the small-angle region and is due to differences in particle multiplicity, though the overall shape and structure of the plane is similar.

The Lund Jet Plane is represented as a vector of coordinates of the emissions for each jet present in an event. Constructing the response matrix for the observable is therefore not as straightforward as for scalar values.

First and foremost, the jet whose clustering history is examined must be present at both particle- and detector-level, just as for any other observable. As the migration is studied for emissions within a jet, we are also forced to require that jet kinematics roughly match between particle- and detector-level. Without imposing this requirement, we run the risk of studying the clustering histories of two completely different jets, leading to distortions in the response matrix. The matching is applied by looking at the position of the jet axis within the detector: the jets at particle- and detector-level must fall within one jet radius of each other within the $y - \phi$ plane.

The number of emissions within the primary Lund Jet Plane can then vary across particle- and detector-level. We must find a way to account for these variations, as there must be a one-

[8]This is not strictly true in the entire Lund Tree, as the difference in particle multiplicity between the definition of the Lund Plane with and without the neutral component of the jet leads to additional branchings in further planes beyond the primary.

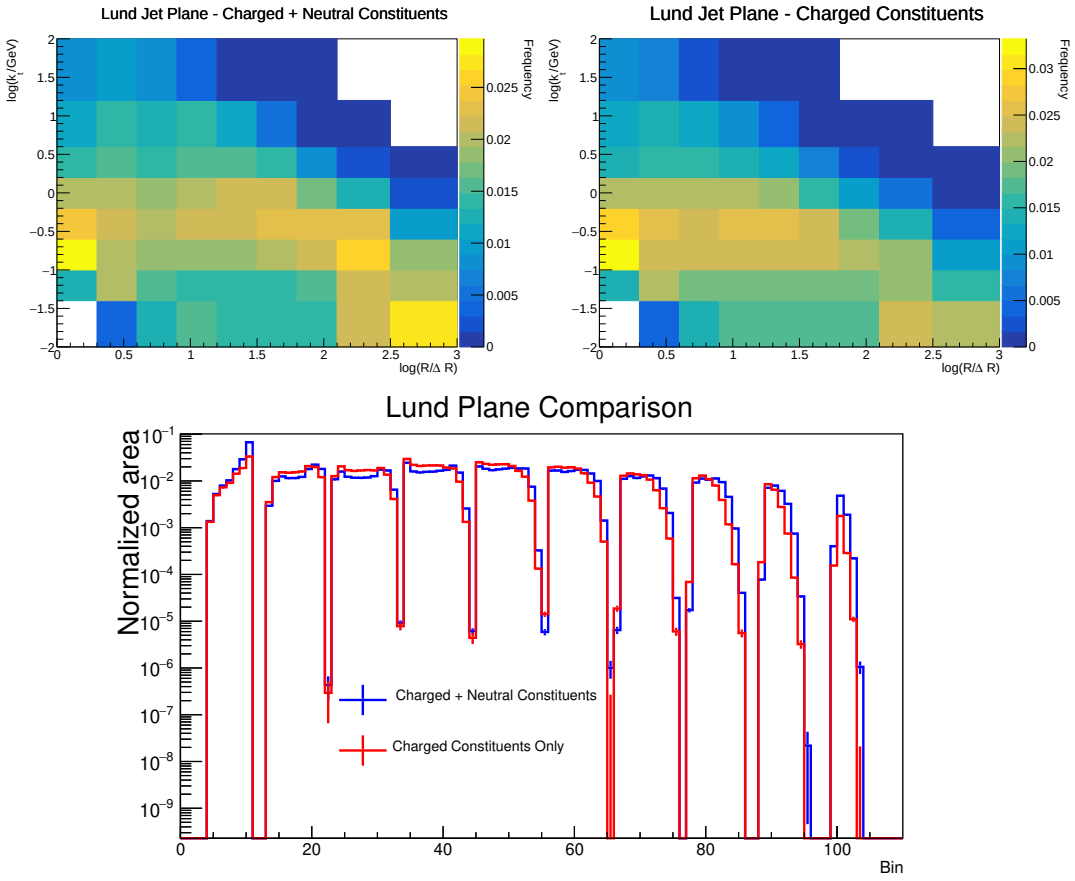


Figure 5.37: The particle-level leading jet Lund Plane satisfying the fiducial selection in the Resolved 0B signal region with p_{T} between 60-100 GeV. The Lund planes shown are defined with particles with $p_{\mathrm{T}} > 900$ MeV, though the inclusion of neutral particles varies.

to-one correspondence to construct the matrix. This is done by imposing a geometric matching of the emissions, in accordance with the strategy implemented in References [56] and [61]. Each emission across particle- and detector-level is required to fall within a distance of $\Delta R < 0.1$ in the $y-\phi$ plane. The order of the matching of the emissions, i.e. whether the jet clustering history is parsed starting from the initial emission or the final, has been shown to not matter.

The matching procedure, however, can lead to the introduction of bias as prior knowledge is input to form a response matrix with a more diagonal structure. To combat this, the matching procedure is only implemented when the number of emissions varies across particle- and detector-level. Additionally, the track selections to construct the Lund Jet Plane are optimised in order to maximise the amount of times the same number of emissions is present at particle- and detector-level, as described in Section 5.6.

5.10.2 Unfolding systematic uncertainties

Several sources of uncertainty also enter the unfolding procedure. These include uncertainties of statistical nature, intrinsic bias from the unfolding method itself, and modelling and theoretical

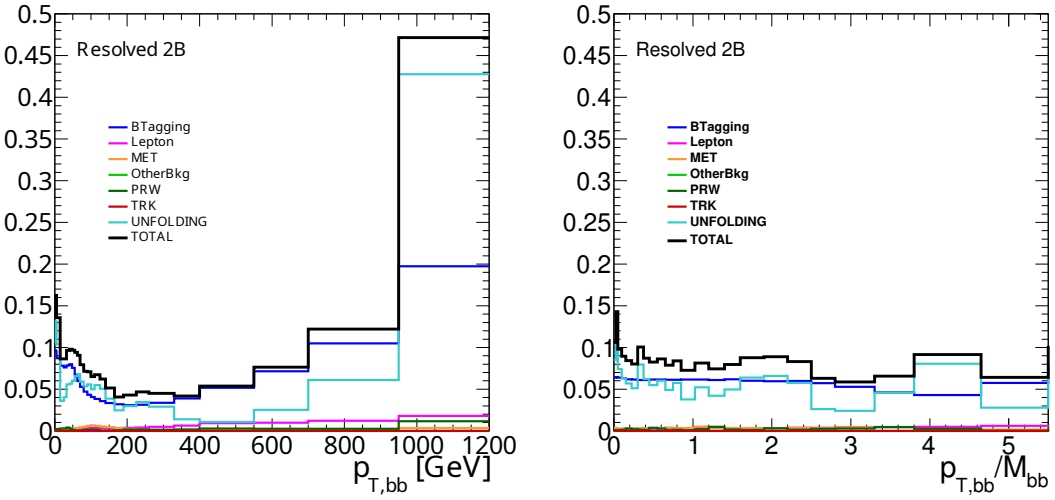


Figure 5.38: The systematics breakdown for the uncertainties in the $p_{\mathrm{T},bb}$ (left) and $p_{\mathrm{T},bb}/M_{bb}$ (right) in the Resolved 2B signal region.

uncertainties.

The limited statistics of the MG5_@aMC+Py8 FxFx samples (MG5+Py) can lead to distortions in the response matrix and a subsequent impact on the statistical error in the fake correction distributions. To correct for these effects, a toy MC is run one hundred times and the unfolding inputs are fluctuated randomly following a normal distribution. The error is then estimated from the root mean squared.

The choice of initial prior can also influence the final uncertainty. The dependence on this prior is determined by reweighting the detector-level distributions of each unfolded observable in the MG5_@aMC+Py8 FxFx samples to match data. The reweighted samples are then used as data and the unfolding procedure is run with the nominal response matrix. The uncertainty estimated in this way is known as the modelling-shape uncertainty. The treatment of this uncertainty is not yet included in the final results.

Lastly, signal modelling uncertainties carry over into the response matrixes, the reconstruction efficiency, and the unmatched events corrections. This uncertainty is estimated by comparing results obtained by varying the QCD scales and by running the unfolding procedure with different MC generators. Specifically, in this thesis results obtained with a conservative estimate of the signal modelling uncertainty are shown. This uncertainty is derived by running the unfolding procedure with response matrices obtained with the Sherpa sample. The difference with the nominal result obtained by unfolding with MG5_@aMC is taken as a systematic uncertainty.

The systematics breakdown related to unfolding is shown in Figures 5.38–5.40 for the four signal regions of the measurement. As can be seen, in the Resolved SRs, the dominant systematics are associated to b-tagging and unfolding. This holds true also in the Boosted SRs, with the additional contribution of PRW and Lepton uncertainties in some bins of the distributions shown.

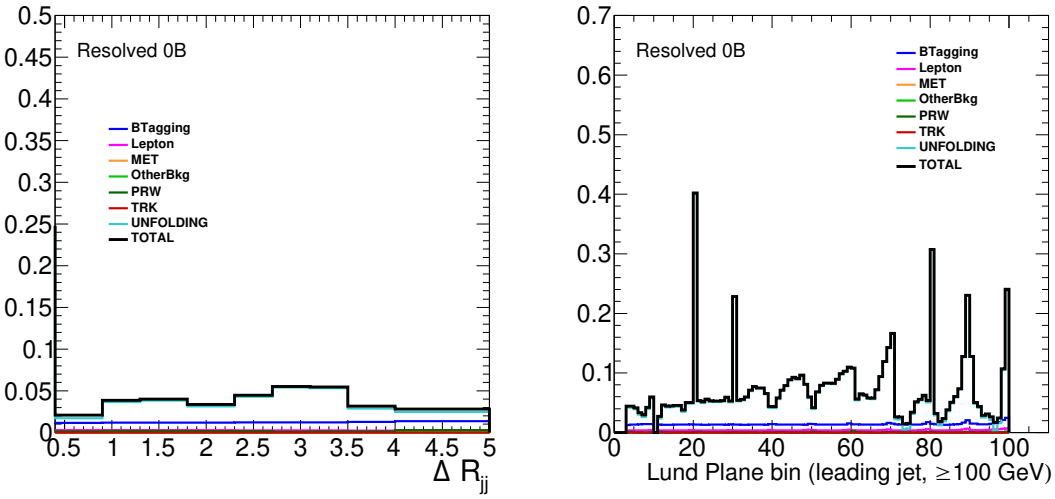


Figure 5.39: The systematics breakdown for the uncertainties in the ΔR_{jj} (left) and one-dimensional Lund Plane distribution in $p_{\mathrm{T}} \geq 100$ GeV interval (right) in the Resolved 0B signal region.

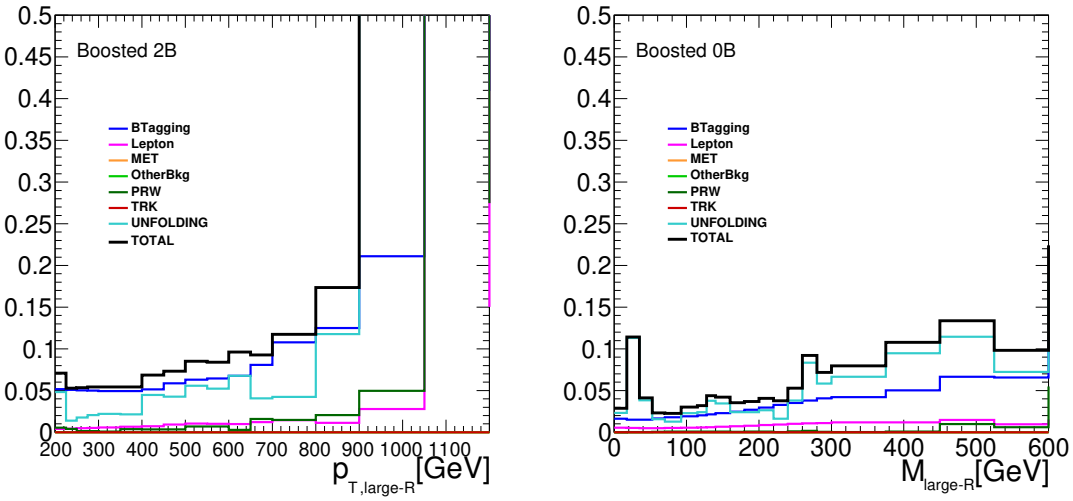


Figure 5.40: The systematics breakdown for the uncertainties of the large-R jet p_{T} distribution in the Boosted 2B signal region (left) and the large-R jet mass in the Boosted 0B signal region (right).

5.11 Results

This section is dedicated to presenting the results of the most up-to-date version of the measurement. The results presented are partial, as some observables have yet to be finalised for the unfolding[9] and the treatment of some systematics, such as the theory uncertainties and those related to the flavour fit, are yet to be included. Along with the results, predictions obtained with MG5+Py8 and Sherpa are also shown with the associated uncertainty on the MC statistics.

5.11.1 Resolved kinematic observables

We begin by considering the results in the Resolved signal regions. In Figures 5.41 and 5.42, we see the measured differential cross sections for the di-jet p_{T}, di-jet mass, the ratio between these two quantities, and the di-jet separation ΔR in the Resolved 2B and Resolved 0B regions, respectively.

In the Resolved 2B SR, we generally find good shape agreement between predictions and data in the majority of the spectrum of the observables, though issues with normalisation are present. These tend to be more pronounced for predictions obtained with Sherpa, and the MG5+Py8 predictions also tend to estimate higher cross sections compared to the Sherpa predictions.

A notable shoulder is present at values of $p_{\mathrm{T},bb} \sim 90$ GeV for both sets of predictions, after which the MG5+Py8 predictions tend to agree with data within uncertainties, while the Sherpa predictions systematically underestimate the cross section. There is also a notable peak in the predicted cross section for both MG5+Py8 and Sherpa between 400-600 GeV.

In the M_{bb} spectrum, again a shoulder is present in the MG5+Py8 prediction of the spectrum at the value of $M_{bb} \sim 100$ GeV. The shoulder marks a turning point in the M_{bb} spectrum: for values below 50 GeV, the cross section is underestimated, and subsequently overestimated up to values of 300 GeV, at which point it is once again underestimated. The Sherpa predictions are more flat, agreeing with data for values below 150 GeV, from which point the spectrum begins to fall and underestimate data more and more.

The $p_{\mathrm{T},bb}/M_{bb}$ MG5+Py8 predictions tend to agree with data within uncertainties in the lower end of the spectrum, but then tend to overestimate the cross section from values of 0.2 onwards, until a turning point at 1.8, from which they begin to underestimate the data. The Sherpa predictions for this observable are instead more flat, underestimating data but generally matching the shape.

The MG5+Py8 predictions underestimate the ΔR_{bb} spectrum for low values of the observable and overestimate it for high values, with a turning point around $\Delta R_{bb} \sim 1.0$. From that point, however, the shape of the predicted spectrum tends to match that of data, apart from normalisation factors. Despite underestimating the ΔR_{bb} spectrum for all values, a similar behaviour is found for the Sherpa predictions. The distributions for both set of predictions roughly flatten from values of $\Delta R_{bb} = 1.0$ onwards.

Unlike for b-jets, the results for light jets in Figure 5.42 show better agreement between MC and data for all kinematic observables. Sherpa predictions for the differential cross sections still tend to fall below MG5+Py8 predictions, though both show good shape agreement, even when the predictions do not agree with data within uncertainties. One notable exception is in the $p_{\mathrm{T},jj}/M_{jj}$ distribution for predictions obtained with MG5+Py8. Although the Sherpa predictions for this observable generally agree with data, the MG5+Py8 predictions underestimate the measured result from values of roughly $p_{\mathrm{T},jj}/M_{jj} \sim 3$ onward. Lastly, we wish to note that this measurement represents the first exclusive Z+jets results with anti-b-tagging in ATLAS.

[9] At the time of writing, optimisation of the unfolding strategy for the jet substructure observables in the Resolved SRs is underway.

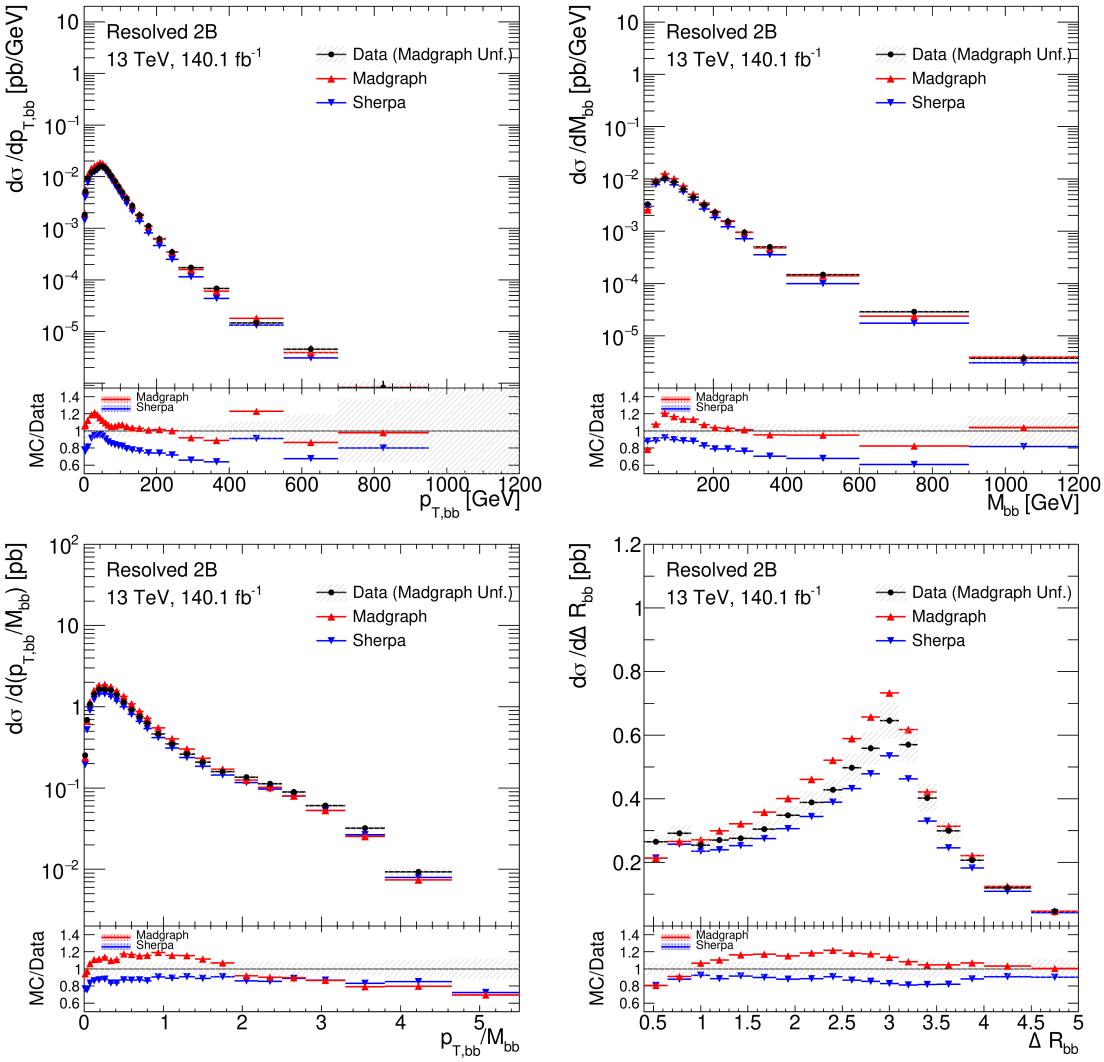


Figure 5.41: The unfolded results of the kinematic observables measured in the Resolved 2B signal region with comparison to predictions obtained with MG5+Py8 and Sherpa.

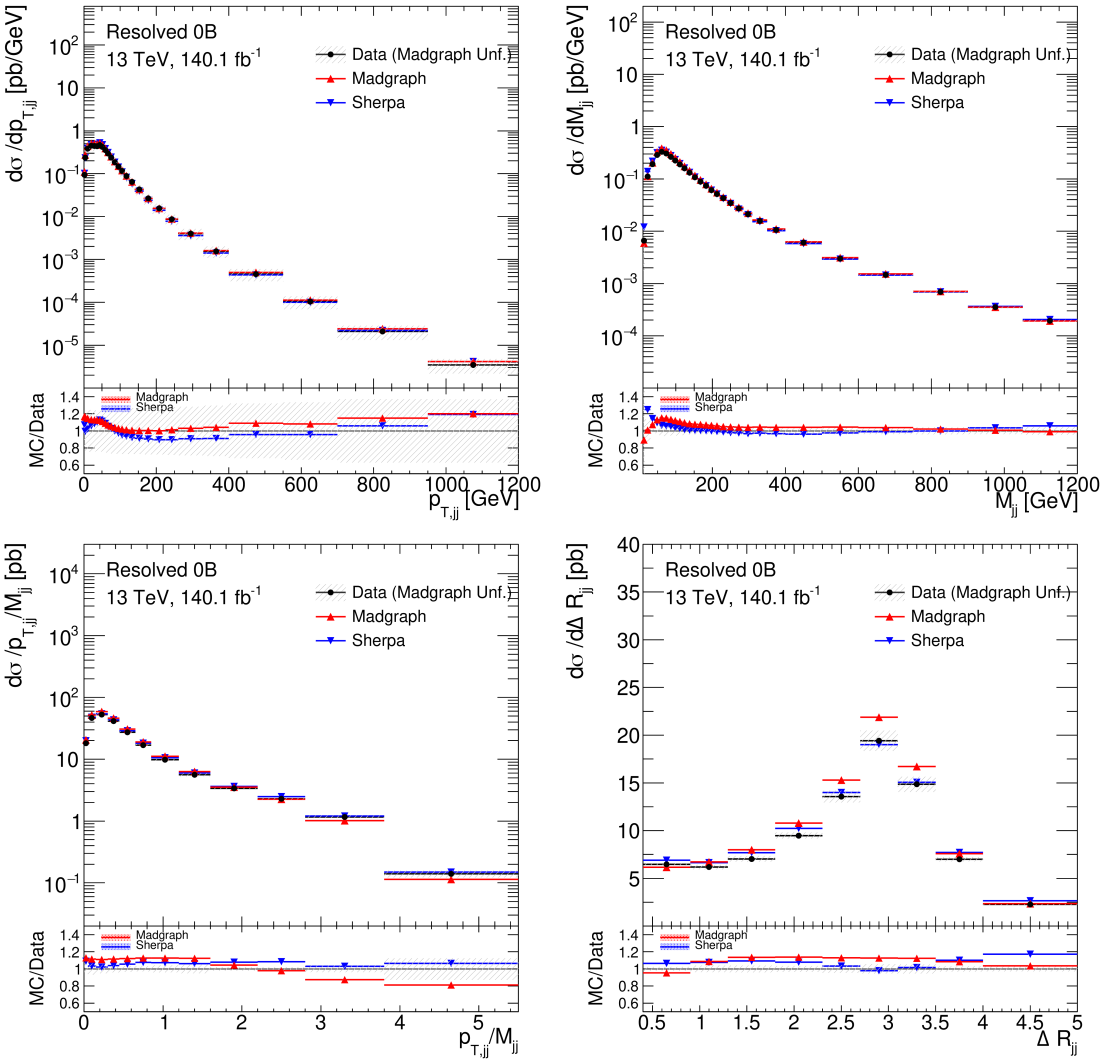


Figure 5.42: The unfolded results of the kinematic observables measured in the Resolved 0B signal region with comparison to predictions obtained with MG5+Py8 and Sherpa.

5.11.2 Boosted kinematic observables

In Figure 5.43, the results and comparisons to predictions are shown for the large-R jet p_{T} and mass in the Boosted 2B and Boosted 0B signal regions. For the double-b-tagged large-R jet, the cross section is systematically underestimated by both MC generators. Both sets of predictions exhibit poor shape agreement between with data as well. Agreement is also very poor for the large-R jet mass. The situation slightly improves for the angular distance of the b-jets, as there is agreement between data and predictions in some parts of the spectrum, and the shape is only distorted in some regions. Both sets of predictions underestimate the cross section of this observable.

In the Boosted 0B signal region, there is instead excellent agreement between predictions and data for the large-R jet p_{T} spectrum. There is also good agreement in the lower end of the large-R jet mass spectrum, particularly for the predictions obtained with Sherpa. The MG5+Py8 predictions only agree for values of jet mass below 100 GeV.

The discrepancy for the jet mass is universal, present both for b-jets and light jets. This points to inherent deficit in the simulation used to obtain predictions. The problem is only accentuated for b-jets, and mismodelling is also present in the large-R double-b-jet p_{T} spectrum, possibly caused by scale effects.

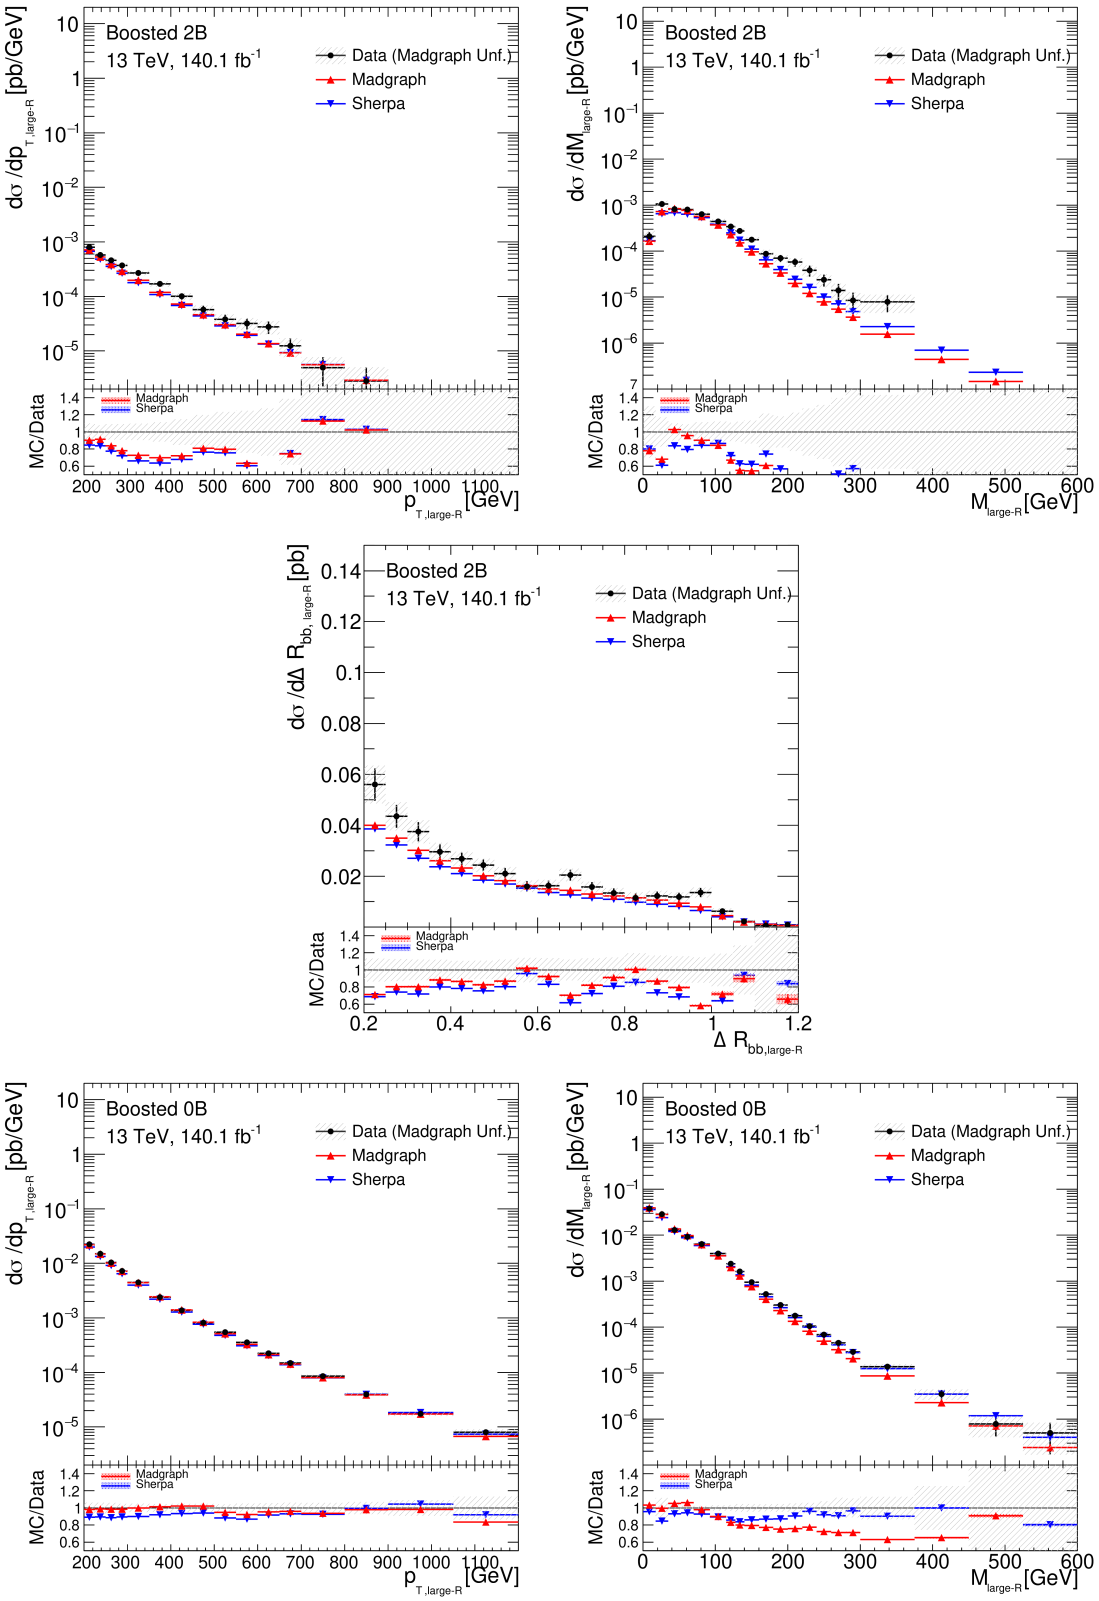


Figure 5.43: The unfolded results of the kinematic observables measured in the Boosted 2B and Boosted 0B signal region with comparison to predictions obtained with MG5+Py8 and Sherpa.

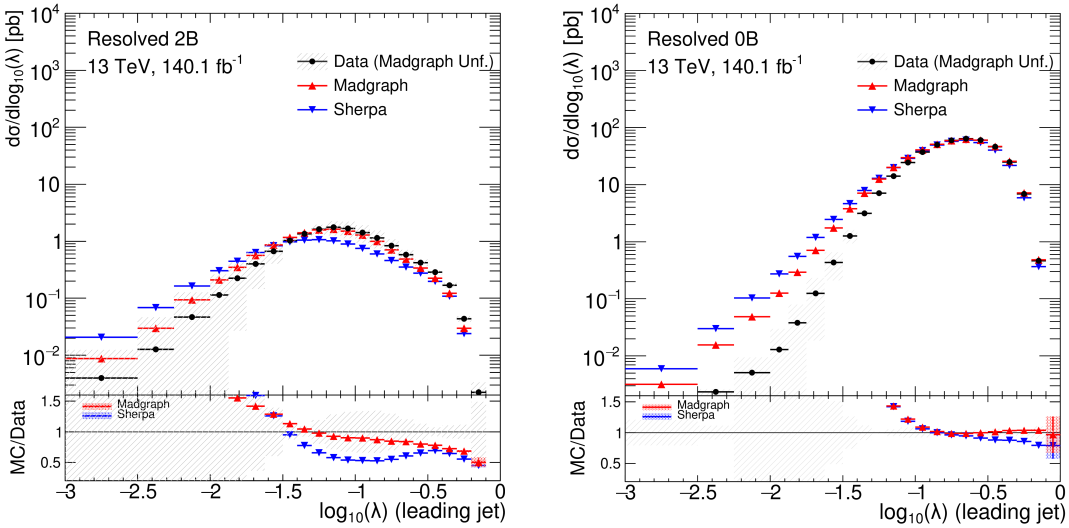


Figure 5.44: The unfolded results of the leading jet angularity differential cross section in the Resolved 2B (left) and Resolved 0B (right) signal regions and comparison to predictions.

5.11.3 Jet substructure observables

Turning to jet substructure, in Figure 5.44, we show the distributions for the logarithm of the leading jet angularity λ in both the Resolved 2B and Resolved 0B signal regions. Agreement between predictions and data is fair in the higher end of the spectra – from values of -1 onward – in the Resolved 0B SR, though notably poorer in the Resolved 2B SR. For this portion of the spectrum in both signal regions, however, the MG5+Py8 predictions agree better with data, generally within uncertainties. For values for $\log(\lambda) < -1$, and particularly below -1.5, the agreement dramatically worsens. Despite the large uncertainty, which reaches values above 50%, both sets of predictions overestimate the measured distribution. For b-jets, the MG5+Py8 predictions manage to fall within the uncertainty band, though for light jets both sets of predictions fail to agree with the data.

The small-λ region is populated by jets whose constituents are soft and collinear relative to the jet p_{T} and axis ($\lambda \to 0$). The mismodelling at low values of λ can possibly be attributed to soft gluon splittings.

The Lund Jet Planes in the 20-60 GeV, 60-100 GeV, and ≥ 100 GeV jet p_{T} slices for the leading b-jet in the Resolved 2B signal region are shown in Figure 5.45. Both the two-dimensional and one-dimensional representations are shown[10], though the former is shown without uncertainty.

In the Lund Plane, for jets in all p_{T} intervals, it is clear how the majority of emissions fall in the mid-to-low k_t region, between roughly $-1.25 < \ln(k_t/\mathrm{GeV}) < 0.25$. Most emissions also fall in the wide-angle region of the Lund Plane, i.e. at low values of $\ln(R/\Delta R)$. The relative number of emissions in the soft and collinear regions increases with jet p_{T}. The high k_t and collinear regions are more sparsely populated, while the upper-right corner representing hard, wide-angle splittings is mostly devoid of emissions.

When looking at the one-dimensional representations, we find that the Sherpa predictions do not agree with the measured cross sections for the Lund Planes throughout the entire spectrum

[10]The two representations of the Lund Plane are equivalent. For an explanation of how to pass from one to the other, see Section 5.6.2

in all p_{T} intervals. In middle k_t ranges, the Sherpa predictions tend to underestimate the amount of radiation and within the b-jets, and there is an angular dependence in the mismodelling as well. In low k_t regions, Sherpa predictions are more compatible with data, at times falling within uncertainty. This points to mismodelling in both the perturbative and non-perturbative regions of the Lund Plane.

MG5+Py8 predictions generally agree better with data, mostly falling within uncertaintiies. An angular dependence to the mismodelling is still present, particularly in the mid k_t ranges, but for low values of k_t both the shape and normalisation of the MG5+Py8 predictions agree with data. This agreement is best in the lowest p_{T} interval.

Predictions for the leading light jet Lund Planes in the Resolved 0B signal region (Figure 5.46) still show mismodelling, though it is not as extreme as for b-jets. For light jets in the 20-60 GeV p_{T} interval, Sherpa predictions tend to overestimate the measured cross sections at low k_t, though agreement improves at higher k_t. The predictions exhibit some angular dependence, though they generally fall within uncertainty. For higher p_{T} jets, the Sherpa predictions agree with data in the low k_t region, while the cross section in higher k_t regions is underestimated.

MG5+Py8 predictions for the light jet Lund Plane, on the other hand, tend to agree well with data in the low k_t region in all jet p_{T} intervals. A strong angular dependence, however, is present in the predictions for the cross sections in the mid-to-high k_t regions of the Lund Jet Plane.

Also to note is the significantly higher relative number of emissions in the soft and collinear region in the 60-100 GeV and $\geq$ 100 GeV light jet Lund Planes, compared both to softer light jets and to b-jets in general. This region is also highly non-perturbative, and sensitive to hadronisation effects.

Lastly, Figure 5.47 shows the large-R jet Lund Jet Planes in the Boosted 2B and Boosted 0B SRs. Here, the pattern of emissions is significantly different: although the bulk of the emissions still fall within the mid k_t and low k_t parts of the planes, the wide-angle region is cut out by the jet trimming and there is a significant number of emissions in the high k_t wide-angle region. Modelling of both MG5+Py8 is very poor throughout the entire spectrum for large-R jets in both the boosted signal regions.

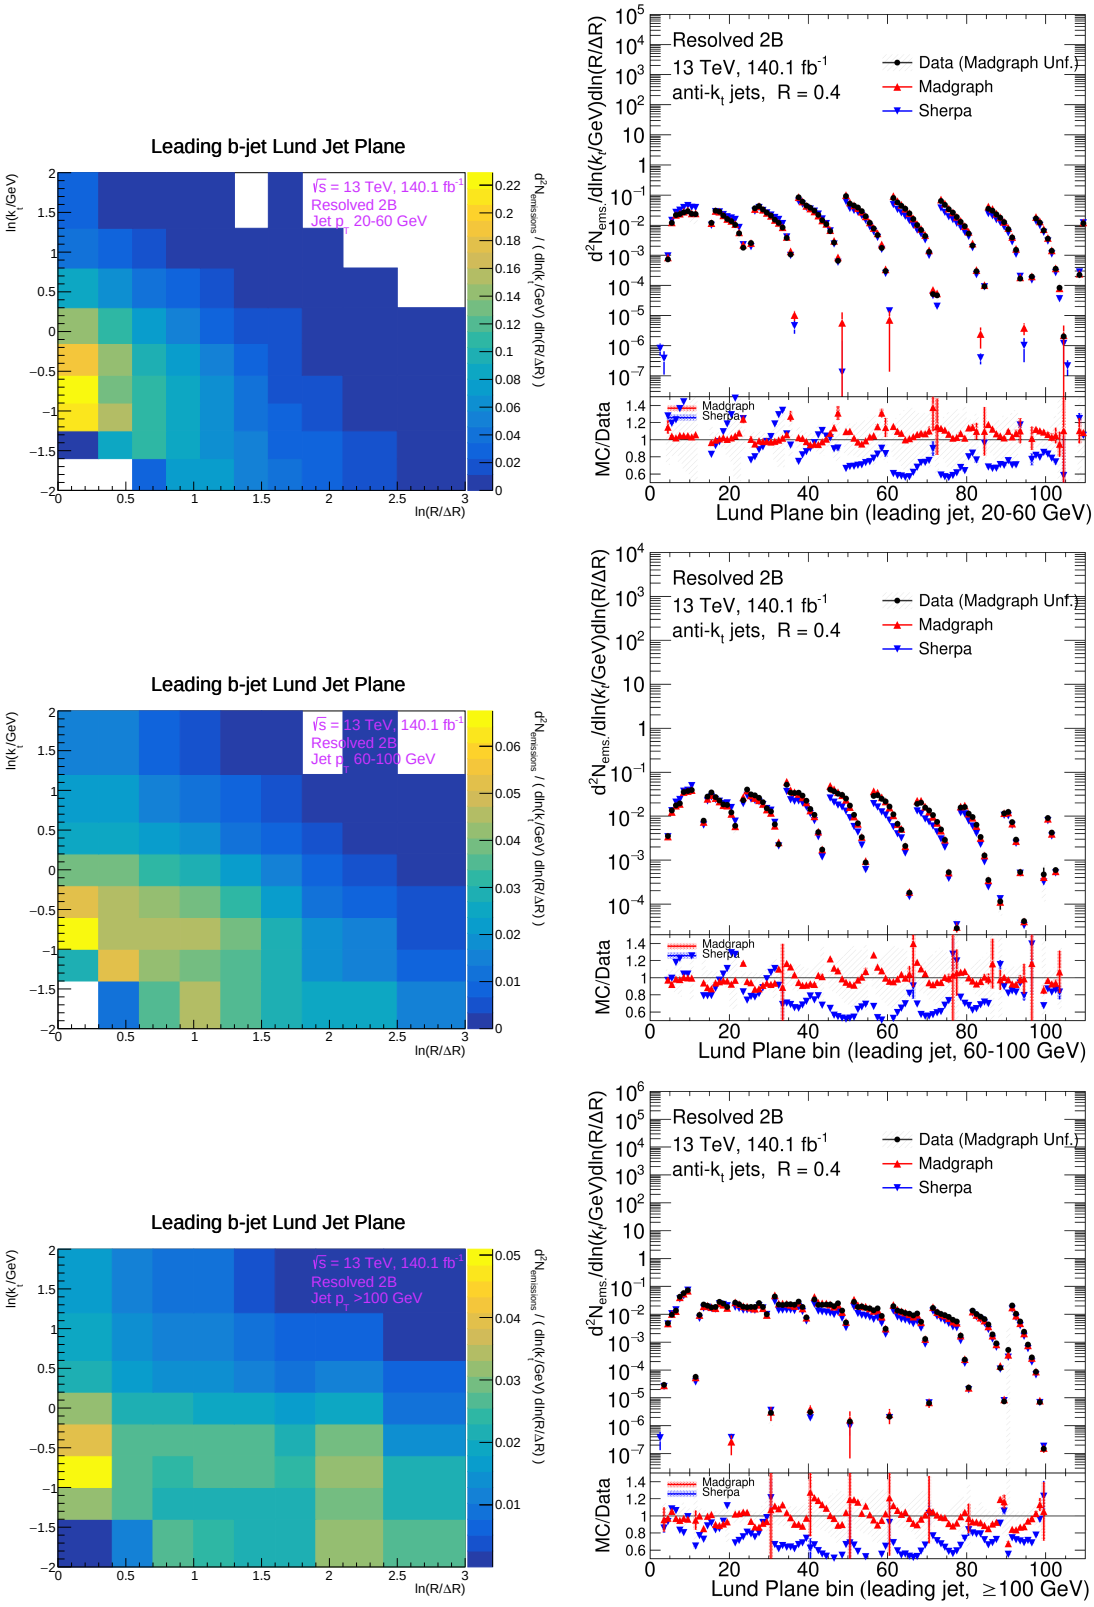


Figure 5.45: The unfolded results for the b-jet Lund Jet Plane in the 20-60 GeV (top), 60-100 GeV (centre), and ≥ 100 GeV (bottom) p_{T} intervals. The Lund Jet Plane is shown in both its two-dimensional and one-dimensional representations.

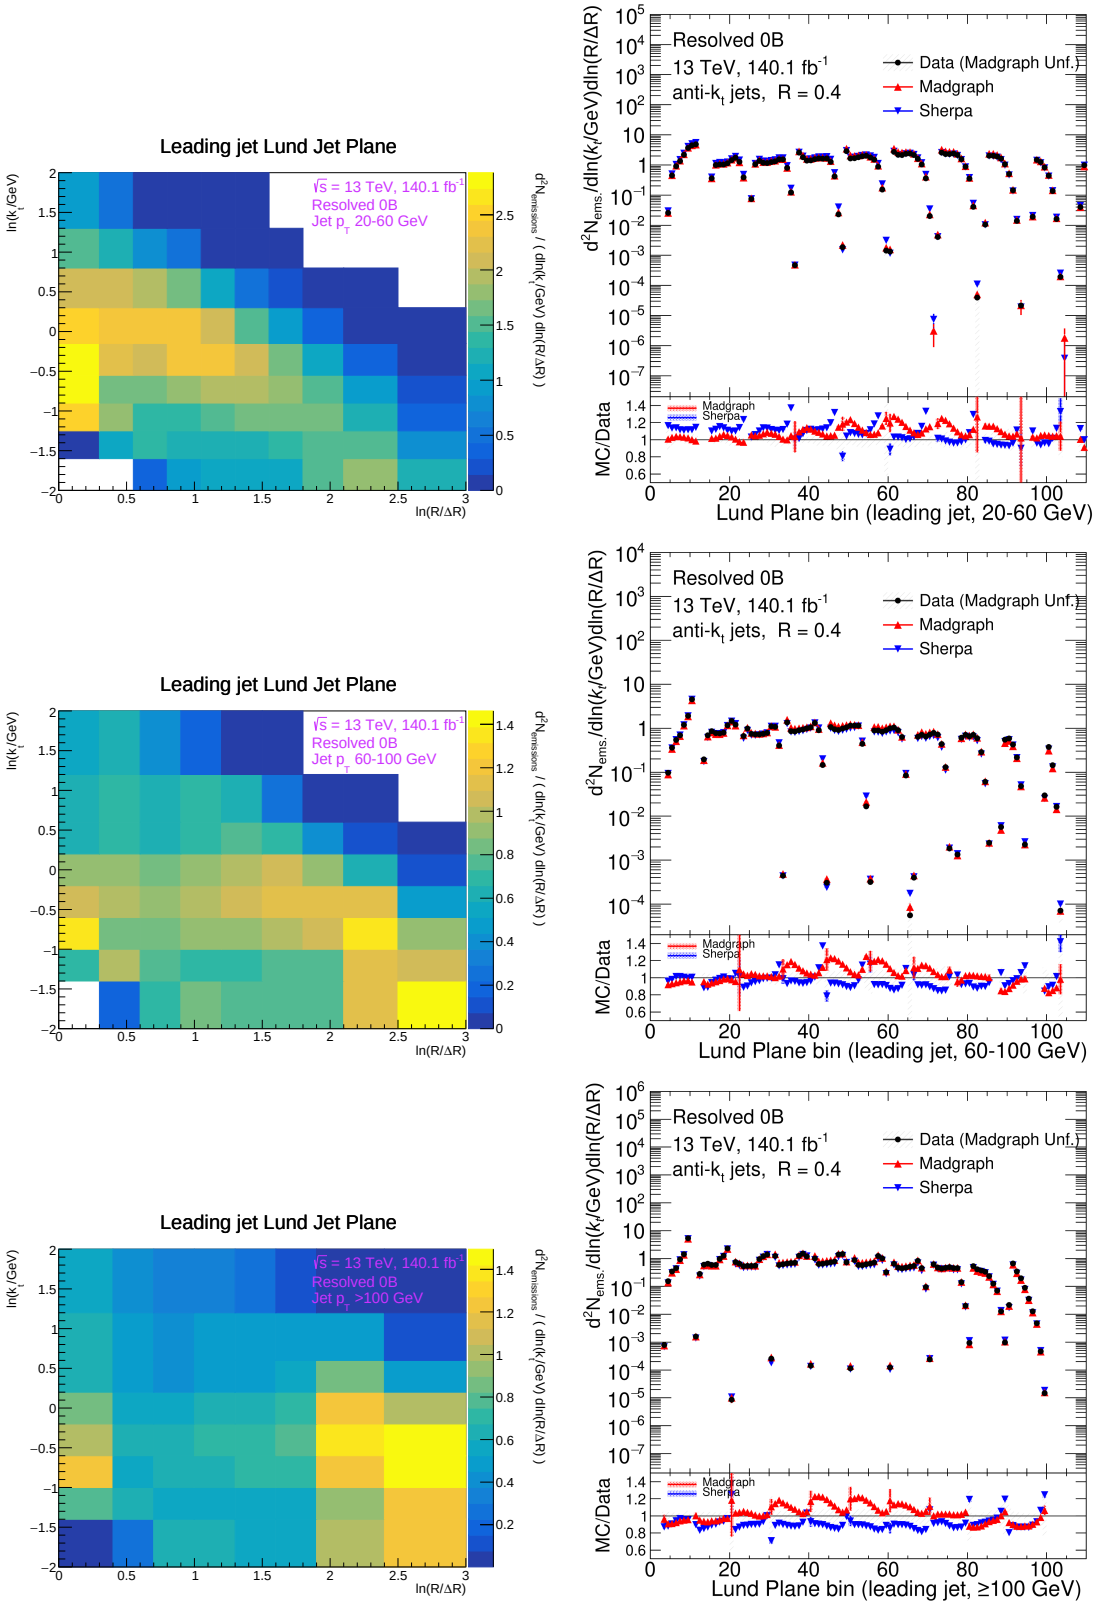


Figure 5.46: The unfolded results for the light jet Lund Jet Plane in the 20-60 GeV (top), 60-100 GeV (centre), and $\geq$ 100 GeV (bottom) p_{T} intervals. The Lund Jet Plane is shown in both its two-dimensional and one-dimensional representations.

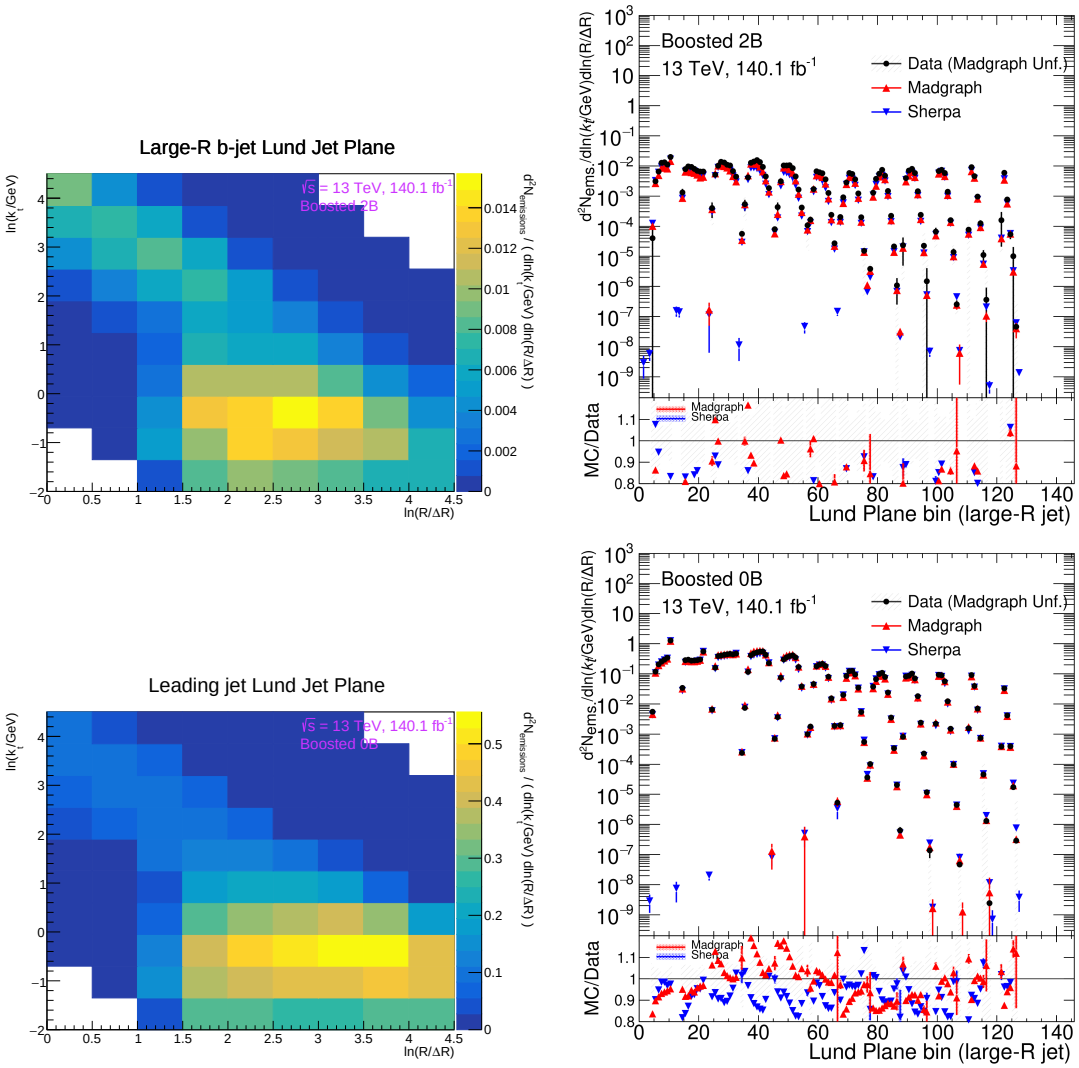


Figure 5.47: The unfolded results for the light jet Lund Jet Plane in the Boosted 2B (top) and Boosted 0B (bottom) SRs. The Lund Jet Plane is shown in both its two-dimensional and one-dimensional representations.

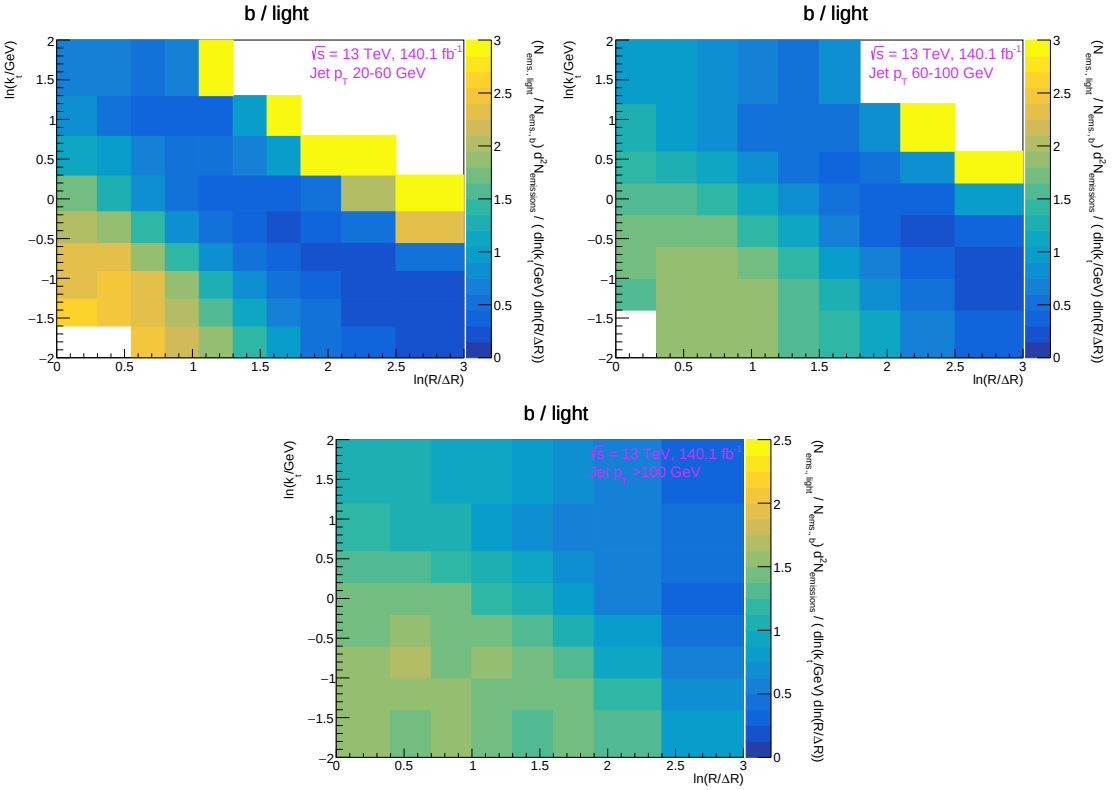


Figure 5.48: The ratio of the b-jet and light jet Lund Jet Planes in the 20-60 GeV (top left), 60-100 GeV (top right), and $\geq$ 100 GeV (bottom) p_{T} intervals.

5.11.4 The dead cone effect

In Figure 5.48, we directly compare the ratios of the b- and light leading jet Lund Planes normalised to the total number of emissions. When looking at the ratios, it is clear how emissions in b-jets mostly populate soft and wide-angle regions, while emissions in light jets more preferentially populate soft and collinear regions of the plane. This is especially true for jets in the lowest p_{T} range considered, and the effect is less pronounced for jets with higher p_{T}.

From the ratio of the projections of the b- and light jet Lund Planes onto the $\ln(R/\Delta R)$ axis, the dead cone effect can be observed. These are shown in Figure 5.49 for the 20-60 GeV jet p_{T} region, the 60-100 GeV region, and the $\geq$ 100 GeV p_{T} region. The projections do not include any cut on the value of k_t and are normalised to the total number of emissions in the Lund Jet Plane. This is done to reduce the number of observables measured, as the number of emissions can be taken directly from the measured Lund Plane, whereas another choice of normalisation, e.g. the number of jets, would require an additional measurement.

To calculate the systematic uncertainty on the ratio, each systematic is individually considered in order to account for correlations. In each bin, the upward and downward variations for b-jets are each divided by the matching variation for light jets. The deviation from the nominal value of the ratio is then found for each of the upward and downward ratio variations, and the greatest of the two (in absolute value) is taken as the uncertainty associated to that systematic. The uncertainties for each systematic are then summed in quadrature to find the total systematic

uncertainty. As the ratio of the projections is statistically limited, it is obtained by considering both the leading and subleading b-jets to reduce the uncertainty.

In the jet p_{T} slices, it can be clearly seen how, as splitting angle decreases, the discrepancy between light jets and b-jets also decreases. For wide-angle splittings, b-jets have a higher relative number of emissions than light jets. Moving towards smaller angles, there is a depletion in the number of emissions in the jets. This depletion is more rapid for b-jets than for light jets, as is predicted by QCD.

As the jet functions as a proxy for the quark in the final state, the dead cone angle can be approximated as $\theta_c \sim M_b/p_{\mathrm{T,\,jet}}$. For jets with $p_{\mathrm{T}} = 20$ GeV, the effect is expected to begin at values of $\ln(R/\theta_c) = 0.6$. This increases to $\ln(R/\theta_c) = 1.7$ for jets with $p_{\mathrm{T}} = 60$ GeV and to $\ln(R/\theta_c) = 2.3$ for jets with $p_{\mathrm{T}} = 100$ GeV. This result constitutes the first observation of the dead cone effect by the ATLAS Experiment[11].

[11] A precise statistical analysis of the observation with the null hypothesis in which there is no dead cone effect is in preparation for the final publication

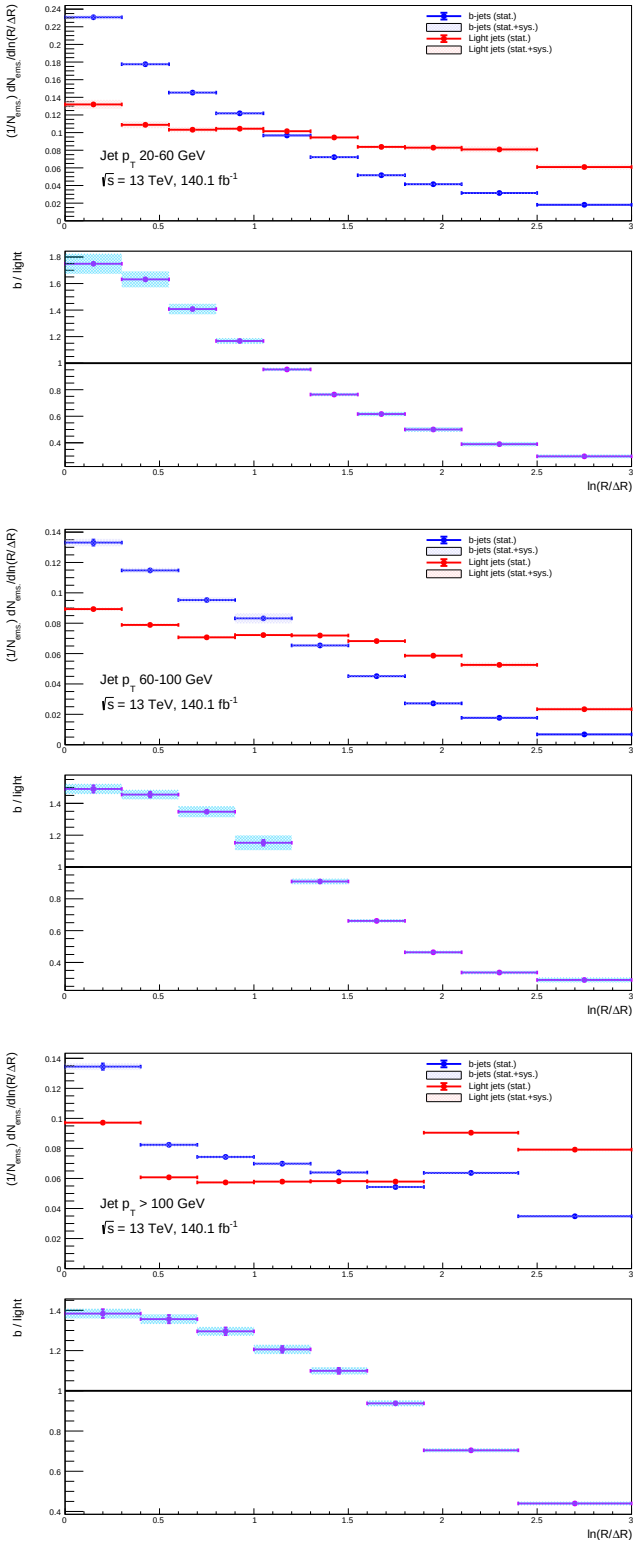


Figure 5.49: The projections onto the $\ln(R/\Delta R)$ axis of the perturbative component of the b-jet and light jet Lund Planes in the 20-60 GeV p_{T} interval (top), the 60-100 GeV p_{T} interval (centre), and the $\geq$ 100 GeV p_{T} interval (bottom), and the ratio of each projection.

Conclusions

In this thesis, we have sought to learn more about heavy flavour jets. By probing them from within, we have uncovered more about the phenomenology and the experimental reality of their production.

Beginning in Chapter 2, we saw how, from NNLO and beyond, soft and collinear gluon splittings can lead to the mislabelling of light jets as heavy flavour jets. We introduced a number of novel flavour jet algorithms which, by clustering together the jet constituents in a flavour-aware fashion, correct the issue by identifying heavy flavour pairs originating from gluon splittings and nullifying their flavour contribution.

We then studied the performance of these algorithms in detail, testing how they behave at parton level and hadron level on Z+jets final states. We found similar performance for all algorithms, but saw that the IFN algorithm is able to suppress the greatest amount of spurious flavour contributions from gluon splittings.

Although, by definition, these algorithms can only be applied to particle-level final states, we hope to employ them in experiment. Specifically, one area prime for their application is in the particle-level definitions of jets in fiducial final states defined at NNLO. By unfolding heavy flavour processes to jets defined with these algorithms, we can correct for heavy flavour introduced not by the hard scattering process itself, but in the final state evolution of the event, leading to more precise measurements.

To this aim, we studied the labelling of heavy flavour jets in $Z + bb$ final states with both more theory-like definitions of jet constituents, i.e. with undecayed heavy hadrons, and more experiment-like definitions, where the heavy hadron is allowed to decay to stable particles. We found that, especially for low p_{T} jets, the intrinsic differences introduced varying the constituent definitions far surpass any differences in the performance of the flavour jet algorithms. This highlights the need to harmonise or correct for the constituent definition across theory and experiment in order to make more accurate measurements and predictions.

Lastly, we saw how the introduction of heavy flavour pairs from gluon splittings can affect the training of flavour tagging algorithms. In particular, we saw how the injection of flavour from soft $g \to b\bar{b}$ splittings in TeV-scale $Z' \to c\bar{c}$ events leads to the mislabelling of the hard c-jets as b-jets. This sample is one of those used to train the flavour tagging algorithms employed by the ATLAS Experiment. Although flavour tagging algorithms are trained with the intention of identifying *all* jets containing heavy flavour irrespective of its origin, the vast scale differences at play are sufficient to justify the correction of this effect, lest a modelling bias be introduced in the flavour tagging training sample. Therefore, the applications of the novel flavour jet algorithms to experiment go beyond the mere improvement of fiducial jet definitions in NNLO predictions.

In Chapter 4, we studied the question of flavour tagging at the ATLAS Experiment in more detail. We recognised features within heavy flavour jets which facilitate their identification within particle detectors, and discerned how these features are exploited to develop algorithms specialised in distinguishing b- and c-jets in data. We considered tagging algorithms developed for

both resolved and boosted topologies, namely the DL1r tagger and the Xbb tagger, respectively.

We then described in detail the calibration of the Xbb tagger on $g \to b\bar{b}$ events. Specifically, we showed how the profile likelihood fit on templates defined based on flavour-sensitive observables such as the signed d_0 significance, can be used to derive data-to-MC efficiency scale factors to correct for the mismodelling of subjet flavours in simulation. We also showed how the use of an alternative template variable, the SV1 mass, leads to similar scale factors, attesting to the robustness of the fit method.

Although the specific version of the Xbb tagger discussed in this work has not been widely used in the ATLAS Experiment, the studies served as an important baseline for successive tagger calibrations. In particular, with the introduction of the GN2 tagger – the new tagger used by the ATLAS Collaboration – the light jet and c-jet rejections have increased by a factor 2.5 and 3, respectively. Work is now underway to calibrate GN2X, the successor to the Xbb tagger.

Last but not least, in Chapter 5 we presented a new ATLAS Run 2 measurement focusing on the previously unmeasured kinematic observables and the jet substructure of heavy flavour and light jets in a $Z + bb$ or $Z + jj$ final state.

Tracks are used to reconstruct jet substructure observables, namely the Lund Jet Plane and jet angularities, at detector-level. These are measured across a wide kinematic range of jets, from a p_{T} as low as 20 GeV and as high as several hundred GeV, and with both resolved and boosted topologies.

We have shown how, with the aid of a simple selection on tracks based on the signed d_0 significance, it is possible to reconstruct the b-hadron kinematics from the charged decay products. The simplicity of the discriminating variable chosen leads to the introduction of minimal bias, beyond that from the b-tagging of the jet itself.

The two jet final state contributes to a high purity selection in this measurement, further improved upon by the flavour fit to account for the Z+jets background from mistagged jets, the neural network based $t\bar{t}$ rejection strategy, and $t\bar{t}$ data-driven background estimation. This purity leads to good performance of the unfolding, which is even able to correct for the lack of neutral particles within the detector-level Lund Plane.

In general, unfolded results show good MC/data agreement for kinematic observables measured in the light jet signal regions of the analysis, but poor agreement for the heavy flavour signal regions. The jet angularities show good agreement both for b-jets and light jets, though the Lund Planes for both flavours of jet show strong mismodelling in several kinematic regions. From the projections of the measured Lund Jet Planes, it was also possible to observe the dead cone effect for the first time with the ATLAS Detector.

Overall, the unfolded results presented in this thesis provide valuable input for the tuning and further development of next-generation Monte Carlo event generators. In the near future, we wish to compare the final results to predictions at NNLO and to fixed-order calculations for the b-jet Lund Plane [160]. The comparison to predictions at this order in perturbation theory will allow for the comparison to jet definitions obtained with the novel flavour jet algorithms presented in Chapter 2. Data from this and other Lund Plane measurements carried out by the ATLAS Collaboration can also be used to test the universality of fragmentation functions. By exploring heavy flavour dynamics using jet substructure techniques, this work opens new directions for precision studies of QCD at the LHC and provides a foundation for future experimental and theoretical development.

Appendix A

Low-level objects for Lund Jet Plane reconstruction

A significant portion of this thesis is focused on the Lund Jet Plane. In Section 5.6, we explain exactly how the Lund Plane is reconstructed from tracks. We omit, however, a discussion of the various studies which culminated in this decision. This appendix is dedicated to detailing these studies.

Many different low-level objects were tested to see if the resulting Lund Plane built from these objects could adequately describe the equivalent particle-level observable. Besides tracks, these objects include LCTopo clusters (calo clusters) and PFlow objects (PFOs), described in Section 3.3.6, as well as Unified Flow Objects (UFOs), described below.

A.1 Track Calo Clusters

Before we begin our discussion on the reconstruction of UFOs, we must first introduce Track Calo Clusters (TCCs) [161]. At high p_{T}, the angular resolution of LCTopo jets is not sufficient for many studies, such as those involving jet substructure. Additionally, PFlow jets are optimised to improve resolution at low p_{T} as the p_{T} resolution of tracks degrades at high p_{T}, though the spatial resolution improves.

Track Calo Clusters offer a high p_{T} solution similar to that adopted by the PFlow algorithm. To build a TCC object, tracks are matched to calorimeter clusters. In the simple case of one track being matched to one cluster, TCC objects adopt the direction of the track and the energy of the calorimeter cluster. If instead no tracks are matched to a given cluster, the TCC simply adopts the four-momentum of the topocluster. More complex configurations, such as multiple track-cluster matches, require a proportional energy sharing between the tracks and calorimeter clusters. The details of the energy-sharing procedure are beyond the scope of this thesis and can be found in Reference [161].

A.2 Unified Flow Objects

Jets formed by Unified Flow Objects (UFOs) are the most recent collection introduced by ATLAS. They get the best of both worlds: improved resolution at low p_{T} as for PFlow jets and at high p_{T} as for TCC jets. UFOs are built starting from PFOs. The TCC algorithm is then applied on tracks with $p_{\mathrm{T}} > 100$ GeV not used for PFlow reconstruction, neutral PFOs, and

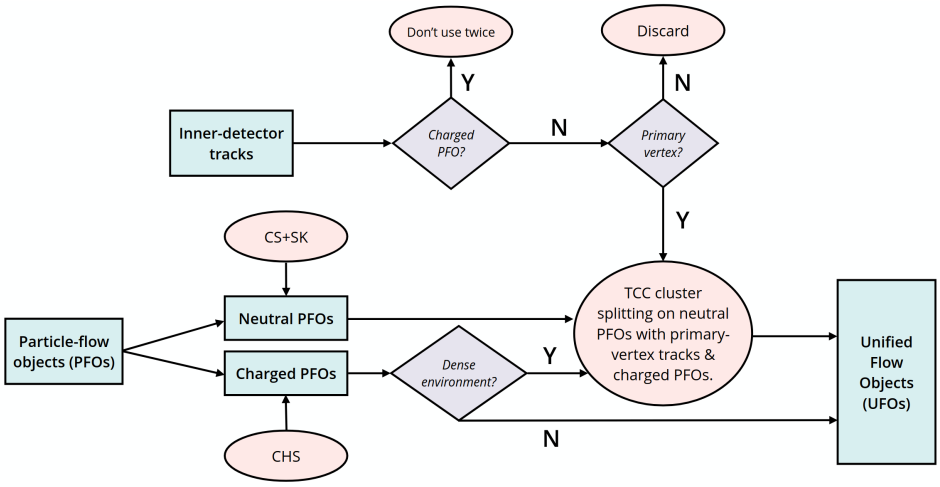


Figure A.1: A flowchart outlining the algorithm used to obtain UFOs [162].

unsubtracted charged PFOs. The resulting TCCs and subtracted charged PFOs are taken as the new UFO objects. A flowchart showing the full decision tree used to obtain the UFOs is shown in Figure A.1 A few grooming and pile-up mitigation techniques are applied: on neutral PFOs, Constituent Subtraction and SoftKiller are applied, while on charged PFOs Charged Hadron Subtraction is applied, as described in Section 3.3.6. The algorithms not previously introduced are briefly described below:

1. **Constituent subtraction:** Constituent subtraction is a pile-up mitigation technique which subtracts pile-up contributions from individual objects. The pile-up energy density contribution ρ is estimated on a per-event basis as

$$\rho = p_{\mathrm{T}}^{\mathrm{jet}} / A^{\mathrm{jet}} \tag{A.1}$$

   where A^{jet} is the area of a jet. Jets used for the estimation are clustered with the k_t algorithm with radius $R = 0.4$. Jets are reconstructed without a p_{T} requirement, but are required to be within $|\eta| < 2.0$ The median value of ρ for all jets is taken.

   Massless soft particles known as *ghosts* are overlaid uniformly in the $\eta - \phi$ plane in such a way that $p_{\mathrm{T}}^{\mathrm{ghost}} = 0.01 \times \rho$ in order to derive the expected contribution of pile-up in an area $\Delta\eta \times \Delta\phi = 0.1 \times 0.1$ of the plane. These ghosts form "negative" contributions which are subtracted iteratively from jet constituents in order to correct for pile-up.

2. **SoftKiller:** SoftKiller is another pile-up mitigation technique which applies a p_{T} cut on constituents, as most pile-up contributions are soft. This cut is applied in such a way that $\rho \sim 0$.

   Specifically, the $\eta - \phi$ plane is divided into a grid with length scale $\ell = 0.6$, optimised for the performance of the algorithm on $R = 0.4$ jets. The specific p_{T} cut applied on the constituents in the event is determined in such a way that half of the grid cells are empty after it is applied.

A.3 Lund Plane reconstruction

With the preliminary notions provided, we are now ready to document the relevant studies on the low-level objects utilised for the Lund Plane reconstruction.

The aim of this study is to optimise the reconstruction of the Lund Jet Plane. Ideally, minimal information loss occurs when passing from particle-level to detector-level, i.e. the radiation pattern is as similar as possible before and after reconstruction.

The final object selected ideally is also capable of reconstructing the Lund Plane in a wide kinematic range. Indeed, the measurement described in Chapter 5 requires reconstruction of jets with p_{T} as low as 20 GeV and as high as several hundred GeV.

The starting point for this study is the charged particle Lund Plane measurement published by the ATLAS Collaboration in 2020 [56]. As in the published measurement, detailed in Section 1.5.2, we consider balanced di-jet events and require two jets of radius $R = 0.4$ per event in the central region of the detector $|\eta| < 2.4$. We also require the leading jet to have a $p_{\mathrm{T}} > 650$ GeV and the subleading jet to have a p_{T} such that $p_{\mathrm{T}}^{\mathrm{leading}} < 1.5\, p_{\mathrm{T}}^{\mathrm{subleading}}$. Jet constituents, whether particles or reconstructed objects, are required to have a $p_{\mathrm{T}} > 500$ MeV.

Detector-level jets are defined based on the object they are built from. In the case of PFlow jets and UFO jets, the jet defined with this object is required to satisfy the kinematic requirements. The jet constituents with $p_{\mathrm{T}} > 500$ MeV are then taken to construct the Lund Jet Plane. In the case of tracks and calo clusters, the objects present in the event with $p_{\mathrm{T}} > 500$ MeV are clustered into jets of radius $R = 0.4$ and subsequently geometrically matched to the particle-level jet with a distance parameter of $\Delta R < 0.1$. The jet constituents from these matched jets are used to construct the Lund Jet Plane.

Figure A.2 shows the particle-level Lund Jet Plane compared to that obtained from the low-level objects considered: calo clusters, PFlow objects, tracks, and UFOs. As can be seen, only tracks and UFOs are able to reproduce the radiation pattern found within the particle-level Lund Jet Plane, eliminating the other objects from consideration.

To decide which object to use between tracks and UFOs, we turn to the one-dimensional Lund Planes defined with the same selections. As can be seen in Figure A.3, although UFO objects can generally recreate the shape of the Lund Plane, tracks have the better agreement when compared to solely the charged component of the particle-level Lund Plane.

It is important to note that the excellent agreement found between the Lund Planes defined with tracks and truth particles is due to two additional considerations: the aforementioned restriction to charged particles, and the application of the "Tight" Track Selection and Track-to-Vertex Association working points. The further optimisation of these working points are described in Section 5.6.

Lastly, there is an argument to be made in favour of the choice of UFOs: the addition of the neutral component. As they are defined with calorimeter information, neutral information is included. As mentioned in Section 5.6, isospin symmetry should ensure that there is little difference between the Lund Plane defined with just charged particles vs. that defined with charged and neutral particles. This had already been shown to be true for the relative jet mass $\rho = \log(m^2/p_{\mathrm{T}}^2)$ and the groomed jet radius r_g [163]. For this reason, and considering also the fact that no systematic uncertainties currently exist for UFO object, it was ultimately decided to opt for the track-based reconstruction.

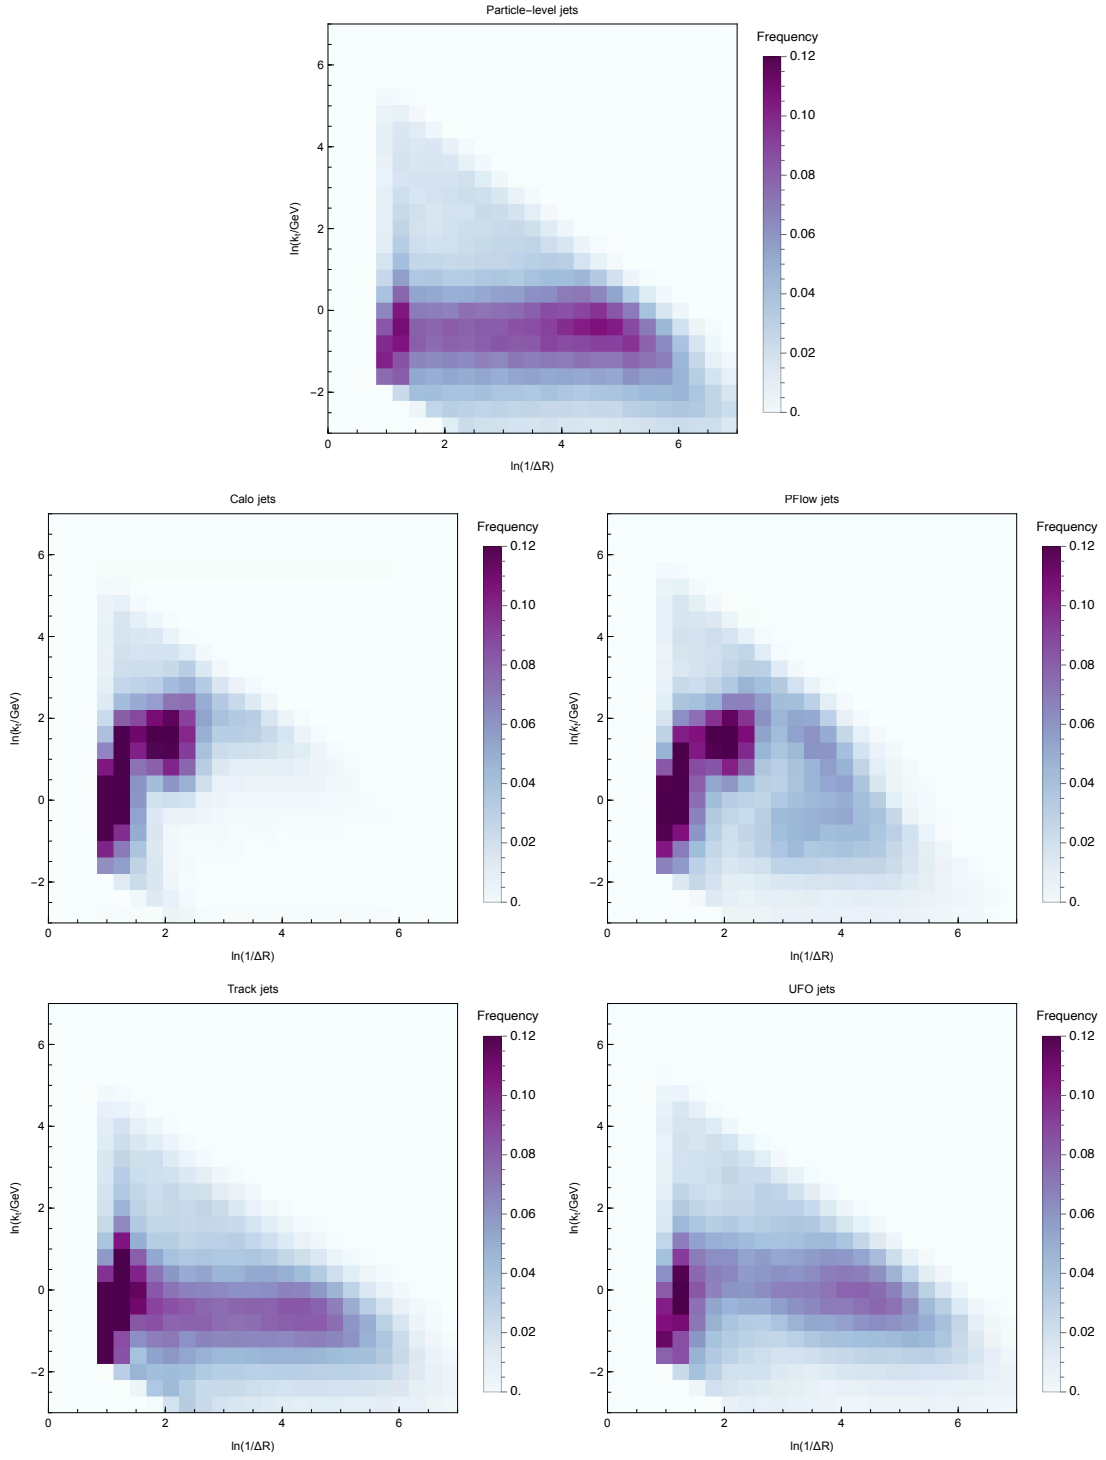


Figure A.2: The particle-level Lund Jet Plane for jets from balanced di-jet events with $p_{\mathrm{T}}^{\mathrm{leading}} > 650$ GeV (top) compared to the Lund Plane reconstructed from calo clusters (centre left), PFlow objects (centre right), tracks (bottom left) and UFOs (bottom right). All objects considered have a $p_{\mathrm{T}} > 500$ MeV.

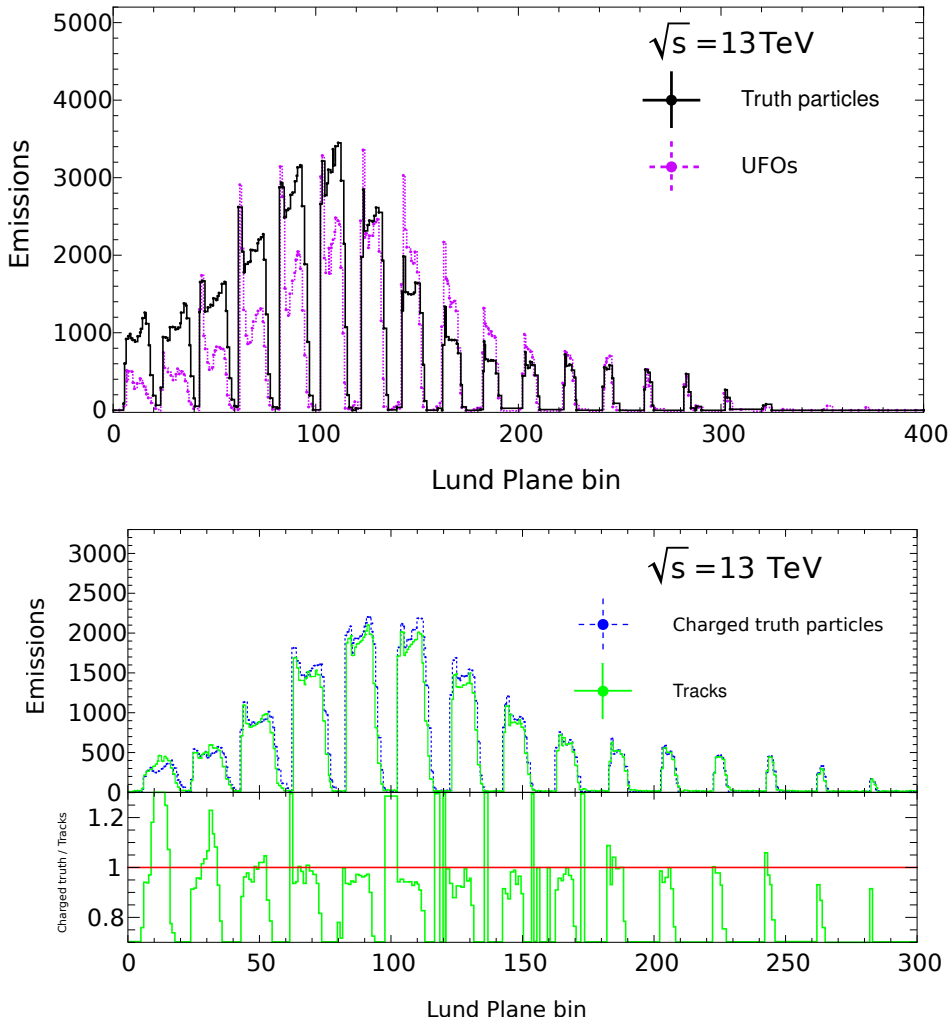


Figure A.3: The agreement between the one-dimensional particle-level Lund Plane obtained with charged and neutral particles vs. the detector-level Lund Plane obtained with UFOs (top) and the particle-level Lund Plane obtained with just charged particles and the detector-level Lund Plane obtained with tracks satisfying the Tight Track Selection and Tight Track-to-Vertex Association working points (bottom).

Appendix B

MultiFold unfolding

Iterative Bayesian Unfolding, described in Section 5.10, is a common unfolding procedure utilised in Standard Model measurements published by the ATLAS Collaboration. However, the procedure has a few limitations. Firstly, the unfolding is performed on previously specified observables. By definition, the response matrix must be constructed for each observable, so it is impossible to devise new quantities *a posteriori*. For the same reason, the binning for the distributions of each observables is also specified before the construction of the matrix, and cannot be subsequently altered. Since the observables must be necessarily binned, most measurements are performed as a function of a single observable, or at most in bins of two or three observables (e.g. double- or triple- differential cross sections). Lastly, if the response matrix depends on features that are not directly used in unfolding, the results may end up biased.

For these reasons, the OmniFold unfolding strategy was developed [159]. OmniFold is a machine learning based unfolding procedure which is capable of simultaneously correcting for detector effects in an arbitrary number of dimensions without any observable binning. The unfolded observables may also be combined after the fact to devise new observables, which come unfolded *gratis*.

The OmniFold unfolding procedure is described below. We start with a vector of particle-level observables $\vec{x}_p$ and detector-level observables $\vec{x}_r$ from a MC sample, and the vector of observables from data $\vec{x}$. First, the distributions of detector-level observables are reweighted to match data. This reweighting step is represented as a function $\omega(\vec{x}_r)$. The weights are then propagated to the particle-level distributions which are reweighted by a function $\nu(\vec{x}_p)$ to match the distribution $\omega(\vec{x}_r)$. The weights $\nu(\vec{x}_p)$ are then propagated back to detector-level in order to improve the agreement to data. The entire process is then iterated a predetermined number of times. The unfolded result can then be saved as a set of generated events with associated weights, or as the learned weighting function after n iterations $\nu_n(\vec{x}_p)$ and instructions for sampling from the particle-level distribution. The unfolding steps are represented in Figure B.1 with four iterations.

The reweighting functions are based on the output of classifiers based on machine learning methods. Specifically, the likelihood ratio

$$L[(\omega, X), (\omega', X')](x) = \frac{p_{(\omega,X)}(x)}{p_{(\omega',X')}(x)} \tag{B.1}$$

is approximated from the probability density $p_{(\omega,X)}(x)$ of x estimated from empirical weights ω and samples X using a classifier trained to distinguish (ω, X) from (ω', X'). Any method of estimating this likelihood ratio is valid.

OmniFold can be used in different variations: if a single observable is unfolded, the method

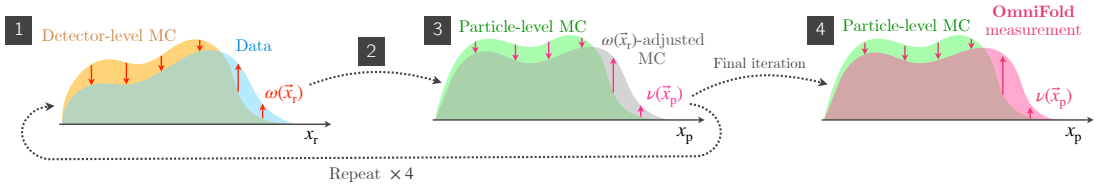


Figure B.1: A graphical representation of the reweighting steps which implemented by the OmniFold unfolding method [164].

is referred to as UniFold. If instead multiple observables are simultaneously unfolded, the term MultiFold is used. If all phase space information is used, allowing one to essentially unfold the entire event, the method is referred to as OmniFold.

The OmniFold method and related strategies have already implemented by a number of measurements carried out with data from ATLAS, CMS, LHCb, H1 and STAR experiments [164–169]. The ATLAS $Z + b\bar{b}/jj$ analysis team also plans on publishing results on jet substructure observables unfolded with OmniFold in the Resolved 2B and Resolved 0B signal regions.

Our measurement makes use of MultiFold to unfold the full Lund Tree, i.e. the full clustering history of the jet rather than just the primary Lund Jet Plane in two of the four signal regions. This would represent the first time that the full Lund Tree is measured. The MultiFold unfolding method is not applied to all four signal regions due to the limited statistics in the Boosted regions. It is performed as a complementary procedure in the Resolved regions as it is statistically more powerful than the IBU method. As in the procedure described in Section 5.10, we aim to unfold to the full Lund Tree defined with both charged and neutral particles.

In our measurement, the likelihood is estimated through a Graph Neural Network (GNN) coined as LundNet-9. LundNet-9 is based on LundNet [34], a GNN based on jet substructure designed for boosted top- and W-jet tagging, and for quark/gluon jet discrimination. In LundNet, each declustering is represented as a node, connected both to its emitter and its emissions. Each node is represented by a number of coordinates representing the declustering, either $(\ln(k_t), \ln(\Delta), \ln(z))$ in LundNet-3 or $(\ln(k_t), \ln(\Delta), \ln(z), \ln(m), \psi)$ in LundNet-5, as defined in Equations 1.76 and 1.77.

The nodes in LundNet-9 are instead represented by $(\ln(k_t), \ln(R/\Delta R), \ln(1/z))$, the typical information from LundNet-3, with the addition of information relating the emitter to the entire jet and kinematic information of the jet. These include $(\ln(k_{t,\,\mathrm{jet}}), \ln(R/\Delta R_{\mathrm{jet}}), \ln(1/z_{\mathrm{jet}}))$, namely the k_t, angular distance, and fraction of momentum carried by the emitter with respect to the jet, and $(\ln(p_{\mathrm{T},\,\mathrm{jet}}), \ln(E_{\mathrm{jet}}), \ln(m_{\mathrm{jet}}))$, or the transverse momentum, energy, and mass of the jet.

The performance of LundNet-9 is evaluated with a closure test, where the nominal MG5+Py8 Z+jets samples is unfolded to pseudodata from a Sherpa 2.2.11 sample. The performance of the unfolding procedure itself is evaluated through the triangular discriminator [36], [170]:

$$\Delta(\ln(1/z), \ln(R/\Delta R)) = \left(\frac{1}{2} \int_{LP} d\ln(1/z) d\ln(R/\Delta R) \frac{(p(\ln(1/z)) - q(\ln(R/\Delta R)))^2}{p(\ln(1/z)) + q(\ln(R/\Delta R))} \right) \times 10^3, \tag{B.2}$$

where $p(\ln(1/z))$ and $q(\ln(R/\Delta R))$ represent the probability distributions for the observables $\ln(1/z)$ and $\ln(R/\Delta R)$ in the MG5+Py8 and Sherpa 2.2.11 samples and the integral is carried out over the area of the Lund Plane.

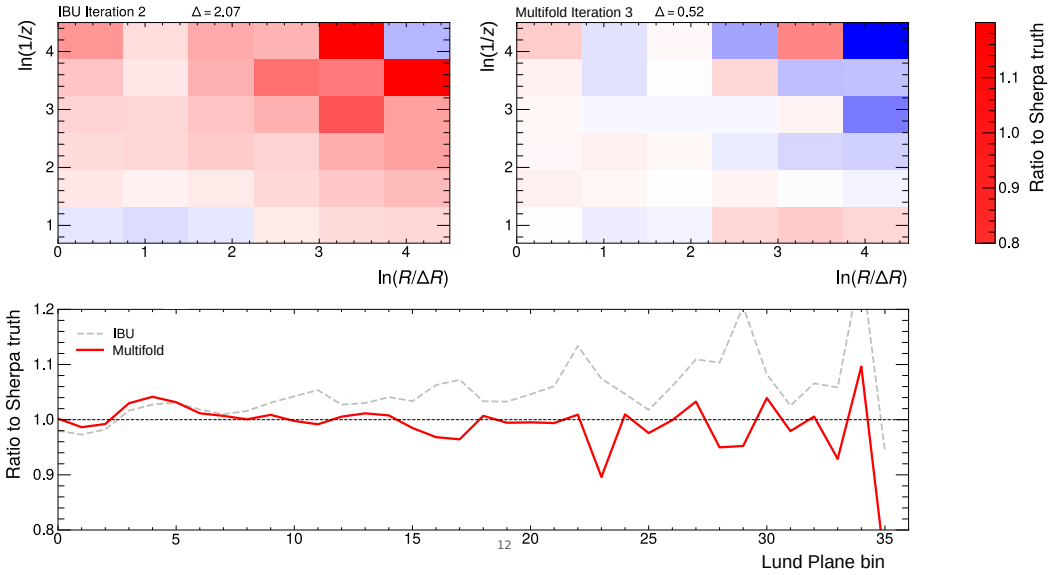


Figure B.2: The comparison of the two-dimensional and one-dimensional Lund Jet Planes obtained by applying IBU and MultiFold unfolding methods in the Resolved 0B SR. In both cases, a MG5+Py8 sample is unfolded to a Sherpa 2.2.11 "truth" sample.

Figure B.2 shows the ratio of the MG5+Py8 Lund Planes for jets in the Resolved 0B region obtained via IBU unfolding with two iterations and via the LundNet-9 based MultiFold method with three iterations. The respective values of the triangular discriminants are $\Delta_{\text{IBU}} = 2.07$ and $\Delta_{\text{MultiFold}} = 0.52$, respectively. Also shown are the one-dimensional comparisons of the unfolded results to the Sherpa 2.2.11 "truth" sample obtained by unrolling the Lund Jet Planes. As can be seen, the ratio for the MultiFold method is closer to unity throughout the majority of the populated regions of the Lund Plane, and the discriminator is approximately four times smaller, showcasing the power of the method.

One limitation of the MultiFold unfolding strategy in this measurement in its current form is the utilisation of a GNN architecture. As the GNN requires at least three track constituents in a jet in order to build an edge, this limits the number of jets whose substructure can be analysed. For the IBU unfolding method, the only requirement is that at least one track constituent be associated to the jet. For this reason, we find that the statistics in the Resolved 2B SR are insufficient to apply this strategy. We are, however, investigating the use of a simpler deep neural network (DNN) to see if the unfolding performance still improves when compared to IBU. The DNN architecture would remove all edge convolutions in the network, allowing us to return to the same selections as applied in the IBU case.

Bibliography

[1] L. H. Ryder, *Quantum Field Theory*. Cambridge University Press, Jun. 1996, ISBN: 978-0-521-47814-4, 978-1-139-63239-3, 978-0-521-23764-2. DOI: 10.1017/CBO9780511813900.

[2] M. Thomson, *Modern particle physics*. New York: Cambridge University Press, Oct. 2013, ISBN: 978-1-107-03426-6, 978-1-139-52536-7. DOI: 10.1017/CBO9781139525367.

[3] Wikimedia Commons, *File: Standard Model of Elementary Particles.svg — Wikimedia Commons, the free media repository*, 2025. https://commons.wikimedia.org/w/index.php?title=File:Standard_Model_of_Elementary_Particles.svg&oldid=1013094830.

[4] Particle Data Group, "Review of particle physics," *Phys. Rev. D*, vol. 110, no. 3, p. 030 001, 2024. DOI: 10.1103/PhysRevD.110.030001.

[5] G. 't Hooft and M. J. G. Veltman, "Regularization and Renormalization of Gauge Fields," *Nucl. Phys. B*, vol. 44, pp. 189–213, 1972. DOI: 10.1016/0550-3213(72)90279-9.

[6] 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*, vol. 46, pp. 641–653, 1977.

[7] G. Altarelli and G. Parisi, "Asymptotic Freedom in Parton Language," *Nucl. Phys. B*, vol. 126, pp. 298–318, 1977. DOI: 10.1016/0550-3213(77)90384-4.

[8] Wikimedia Commons, *File: Running coupling constants.svg — Wikimedia Commons, the free media repository*, 2020. https://commons.wikimedia.org/w/index.php?title=File:Running_coupling_constants.svg&oldid=507424081.

[9] T. Sjöstrand, "Monte Carlo Generators," in *2006 European School of High-Energy Physics*, Nov. 2006, pp. 51–74. arXiv: hep-ph/0611247.

[10] NNPDF Collaboration, "Parton distributions from high-precision collider data," *Eur. Phys. J. C*, vol. 77, no. 10, p. 663, 2017. DOI: 10.1140/epjc/s10052-017-5199-5. arXiv: 1706.00428 [hep-ph].

[11] R. K. Ellis, W. J. Stirling, and B. R. Webber, *QCD and collider physics*. Cambridge University Press, Feb. 2011, vol. 8, ISBN: 978-0-511-82328-2, 978-0-521-54589-1. DOI: 10.1017/CBO9780511628788.

[12] J.-B. Lee, M. R. Masouminia, M. H. Seymour, and U.-k. Yang, "Generalized angular-order parton showers in Herwig 7," *JHEP*, vol. 08, p. 064, 2024. DOI: 10.1007/JHEP08(2024)064. arXiv: 2312.13125 [hep-ph].

[13] M. Dinsdale, M. Ternick, and S. Weinzierl, "Parton showers from the dipole formalism," *Phys. Rev. D*, vol. 76, p. 094 003, 2007. DOI: 10.1103/PhysRevD.76.094003. arXiv: 0709.1026 [hep-ph].

[14] M. L. Mangano, M. Moretti, F. Piccinini, and M. Treccani, "Matching matrix elements and shower evolution for top-quark production in hadronic collisions," *JHEP*, vol. 01, p. 013, 2007. DOI: 10.1088/1126-6708/2007/01/013. arXiv: hep-ph/0611129.

[15] R. Frederix and S. Frixione, "Merging meets matching in MC@NLO," *JHEP*, vol. 12, p. 061, 2012. DOI: 10.1007/JHEP12(2012)061. arXiv: 1209.6215 [hep-ph].

[16] L. Lonnblad, "Correcting the color dipole cascade model with fixed order matrix elements," *JHEP*, vol. 05, p. 046, 2002. DOI: 10.1088/1126-6708/2002/05/046. arXiv: hep-ph/0112284.

[17] B. Andersson, G. Gustafson, G. Ingelman, and T. Sjöstrand, "Parton Fragmentation and String Dynamics," *Phys. Rept.*, vol. 97, pp. 31–145, 1983. DOI: 10.1016/0370-1573(83)90080-7.

[18] B. R. Webber, "A QCD Model for Jet Fragmentation Including Soft Gluon Interference," *Nucl. Phys. B*, vol. 238, pp. 492–528, 1984. DOI: 10.1016/0550-3213(84)90333-X.

[19] D. Amati and G. Veneziano, "Preconfinement as a Property of Perturbative QCD," *Phys. Lett. B*, vol. 83, pp. 87–92, 1979. DOI: 10.1016/0370-2693(79)90896-7.

[20] G. Sterman and S. Weinberg, "Jets from quantum chromodynamics," *Phys. Rev. Lett.*, vol. 39, pp. 1436–1439, 23 Dec. 1977. DOI: 10.1103/PhysRevLett.39.1436.

[21] F. Dreyer, "Precision physics at the large hadron collider," Theses, Université Pierre et Marie Curie - Paris VI, May 2016. https://theses.hal.science/tel-01407230.

[22] S. Bethke and A. Wagner, "The JADE Experiment at the PETRA e^+e^- collider – history, achievements and revival," *Eur. Phys. J. H*, vol. 47, p. 16, 2022. DOI: 10.1140/epjh/s13129-022-00047-8. arXiv: 2208.11076 [hep-ex].

[23] M. Cacciari, G. P. Salam, and G. Soyez, "The anti-k_t jet clustering algorithm," *JHEP*, vol. 04, p. 063, 2008. DOI: 10.1088/1126-6708/2008/04/063. arXiv: 0802.1189 [hep-ph].

[24] A. J. Larkoski, D. Neill, and J. Thaler, "Jet Shapes with the Broadening Axis," *JHEP*, vol. 04, p. 017, 2014. DOI: 10.1007/JHEP04(2014)017. arXiv: 1401.2158 [hep-ph].

[25] E. Moreno, O. Cerri, J. Duarte, *et al.*, "Jedi-net: A jet identification algorithm based on interaction networks," *The European Physical Journal C*, vol. 80, Jan. 2020. DOI: 10.1140/epjc/s10052-020-7608-4.

[26] E. Farhi, "A QCD Test for Jets," *Phys. Rev. Lett.*, vol. 39, pp. 1587–1588, 1977. DOI: 10.1103/PhysRevLett.39.1587.

[27] J. Thaler and K. Van Tilburg, "Identifying Boosted Objects with N-subjettiness," *JHEP*, vol. 03, p. 015, 2011. DOI: 10.1007/JHEP03(2011)015. arXiv: 1011.2268 [hep-ph].

[28] F. A. Dreyer, G. P. Salam, and G. Soyez, "The Lund Jet Plane," *JHEP*, vol. 12, p. 064, 2018. DOI: 10.1007/JHEP12(2018)064. arXiv: 1807.04758 [hep-ph].

[29] A. J. Larkoski, S. Marzani, G. Soyez, and J. Thaler, "Soft Drop," *JHEP*, vol. 05, p. 146, 2014. DOI: 10.1007/JHEP05(2014)146. arXiv: 1402.2657 [hep-ph].

[30] D. Krohn, J. Thaler, and L.-T. Wang, "Jet Trimming," *JHEP*, vol. 02, p. 084, 2010. DOI: 10.1007/JHEP02(2010)084. arXiv: 0912.1342 [hep-ph].

[31] Y. L. Dokshitzer, V. A. Khoze, and S. I. Troian, "On specific QCD properties of heavy quark fragmentation ('dead cone')," *J. Phys. G*, vol. 17, pp. 1602–1604, 1991. DOI: 10.1088/0954-3899/17/10/023.

[32] ALICE Collaboration, "Direct observation of the dead-cone effect in quantum chromodynamics," *Nature*, vol. 605, no. 7910, pp. 440–446, 2022. DOI: 10.1038/s41586-022-04572-w. arXiv: 2106.05713 [nucl-ex].

[33] B. Andersson, G. Gustafson, L. Lönnblad, and U. Pettersson, "Coherence effects in deep inelastic scattering," *Zeitschrift für Physik C Particles and Fields*, vol. 43, pp. 625–632, 1989. DOI: 10.1007/BF01550942.

[34] F. A. Dreyer and H. Qu, "Jet tagging in the Lund plane with graph networks," *JHEP*, vol. 03, p. 052, 2021. DOI: 10.1007/JHEP03(2021)052. arXiv: 2012.08526 [hep-ph].

[35] *Les Houches 2017: Physics at TeV Colliders Standard Model Working Group Report*, Mar. 2018. arXiv: 1803.07977 [hep-ph].

[36] P. Gras, S. Höche, D. Kar, *et al.*, "Systematics of quark/gluon tagging," *JHEP*, vol. 07, p. 091, 2017. DOI: 10.1007/JHEP07(2017)091. arXiv: 1704.03878 [hep-ph].

[37] P. K. Dhani, O. Fedkevych, A. Ghira, S. Marzani, and G. Soyez, "Heavy Flavour Jet Substructure," Oct. 2024. arXiv: 2410.05415 [hep-ph].

[38] "Standard Model Summary Plots June 2024," CERN, Geneva, Tech. Rep., 2024. https://cds.cern.ch/record/2903866.

[39] M. Johnson and D. Maître, "Strong coupling constant extraction from high-multiplicity Z+jets observables," *Phys. Rev. D*, vol. 97, no. 5, p. 054 013, 2018. DOI: 10.1103/PhysRevD.97.054013. arXiv: 1711.01408 [hep-ph].

[40] L. Cavallini, A. Coccaro, C. K. Khosa, *et al.*, "Tagging the Higgs boson decay to bottom quarks with colour-sensitive observables and the Lund jet plane," *Eur. Phys. J. C*, vol. 82, no. 5, p. 493, 2022. DOI: 10.1140/epjc/s10052-022-10447-1. arXiv: 2112.09650 [hep-ph].

[41] ATLAS Collaboration, "Measurement of the production cross section of jets in association with a Z boson in pp collisions at $\sqrt{s} = 7$ TeV with the ATLAS detector," *JHEP*, vol. 07, p. 032, 2013. DOI: 10.1007/JHEP07(2013)032. arXiv: 1304.7098 [hep-ex].

[42] ATLAS Collaboration, "Measurement of the inclusive cross-section for the production of jets in association with a Z boson in proton-proton collisions at 8 TeV using the ATLAS detector," *Eur. Phys. J. C*, vol. 79, no. 10, p. 847, 2019. DOI: 10.1140/epjc/s10052-019-7321-3. arXiv: 1907.06728 [hep-ex].

[43] ATLAS Collaboration, "Measurements of the production cross section of a Z boson in association with jets in pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector," *Eur. Phys. J. C*, vol. 77, no. 6, p. 361, 2017. DOI: 10.1140/epjc/s10052-017-4900-z. arXiv: 1702.05725 [hep-ex].

[44] J. Alwall, R. Frederix, S. Frixione, *et al.*, "The automated computation of tree-level and next-to-leading order differential cross sections, and their matching to parton shower simulations," *JHEP*, vol. 07, p. 079, 2014. DOI: 10.1007/JHEP07(2014)079. arXiv: 1405.0301 [hep-ph].

[45] T. Sjöstrand, S. Mrenna, and P. Z. Skands, "A Brief Introduction to PYTHIA 8.1," *Comput. Phys. Commun.*, vol. 178, pp. 852–867, 2008. DOI: 10.1016/j.cpc.2008.01.036. arXiv: 0710.3820 [hep-ph].

[46] M. L. Mangano, M. Moretti, F. Piccinini, R. Pittau, and A. D. Polosa, "ALPGEN, a generator for hard multiparton processes in hadronic collisions," *JHEP*, vol. 07, p. 001, 2003. DOI: 10.1088/1126-6708/2003/07/001. arXiv: hep-ph/0206293.

[47] T. Sjöstrand, S. Mrenna, and P. Z. Skands, "PYTHIA 6.4 Physics and Manual," *JHEP*, vol. 05, p. 026, 2006. DOI: 10.1088/1126-6708/2006/05/026. arXiv: hep-ph/0603175.

[48] T. Gleisberg, S. Hoeche, F. Krauss, *et al.*, "Event generation with SHERPA 1.1," *JHEP*, vol. 02, p. 007, 2009. DOI: 10.1088/1126-6708/2009/02/007. arXiv: 0811.4622 [hep-ph].

[49] C. F. Berger, Z. Bern, L. J. Dixon, *et al.*, "Next-to-Leading Order QCD Predictions for $Z, \gamma^* + 3$-Jet Distributions at the Tevatron," *Phys. Rev. D*, vol. 82, p. 074 002, 2010. DOI: 10.1103/PhysRevD.82.074002. arXiv: 1004.1659 [hep-ph].

[50] CMS Collaboration, "Measurement of differential cross sections for the production of a Z boson in association with jets in proton-proton collisions at $\sqrt{s} = 13$ TeV," *Phys. Rev. D*, vol. 108, p. 052 004, 2023. DOI: 10.1103/PhysRevD.108.052004. arXiv: 2205.02872 [hep-ex].

[51] 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," *JHEP*, vol. 06, p. 080, 2023. DOI: 10.1007/JHEP06(2023)080. arXiv: 2205.02597 [hep-ex].

[52] ATLAS Collaboration, "Measurement of differential production cross-sections for a Z boson in association with b-jets in 7 TeV proton-proton collisions with the ATLAS detector," *JHEP*, vol. 10, p. 141, 2014. DOI: 10.1007/JHEP10(2014)141. arXiv: 1407.3643 [hep-ex].

[53] ATLAS Collaboration, "Measurements of the production cross-section for a Z boson in association with b- or c-jets in proton–proton collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector," *Eur. Phys. J. C*, vol. 84, no. 9, p. 984, 2024. DOI: 10.1140/epjc/s10052-024-13159-w. arXiv: 2403.15093 [hep-ex].

[54] ATLAS Collaboration, "Measurement of cross-sections for production of a Z boson in association with a flavor-inclusive or doubly b-tagged large-radius jet in proton-proton collisions at $\sqrt{s} = 13$ TeV with the ATLAS experiment," *Phys. Rev. D*, vol. 108, no. 1, p. 012 022, 2023. DOI: 10.1103/PhysRevD.108.012022. arXiv: 2204.12355 [hep-ex].

[55] ATLAS Collaboration, "Performance of top-quark and W-boson tagging with ATLAS in Run 2 of the LHC," *Eur. Phys. J. C*, vol. 79, no. 5, p. 375, 2019. DOI: 10.1140/epjc/s10052-019-6847-8. arXiv: 1808.07858 [hep-ex].

[56] ATLAS Collaboration, "Measurement of the Lund Jet Plane Using Charged Particles in 13 TeV Proton-Proton Collisions with the ATLAS Detector," *Phys. Rev. Lett.*, vol. 124, no. 22, p. 222 002, 2020. DOI: 10.1103/PhysRevLett.124.222002. arXiv: 2004.03540 [hep-ex].

[57] CMS Collaboration, "Measurement of the primary Lund jet plane density in proton-proton collisions at $\sqrt{s} = 13$ TeV," *JHEP*, vol. 05, p. 116, 2024. DOI: 10.1007/JHEP05(2024)116. arXiv: 2312.16343 [hep-ex].

[58] "Physics Preliminary Summary: Measurement of the primary Lund plane density in pp collisions at $\sqrt{s} = 13$ TeV with ALICE," 2021. https://cds.cern.ch/record/2759456.

[59] LHCb Collaboration, "Measurement of the Lund plane for light- and beauty-quark jets," *Phys. Rev. D*, vol. 112, no. 7, p. 072 015, 2025. DOI: 10.1103/r16d-b4my. arXiv: 2505.23530 [hep-ex].

[60] CMS Collaboration, "Jet fragmentation function and groomed substructure of bottom quark jets in proton-proton collisions at 5.02 TeV," Nov. 2025. arXiv: 2511.10666 [hep-ex].

[61] ATLAS Collaboration, "Measurement of the Lund jet plane in hadronic decays of top quarks and W bosons with the ATLAS detector," Jul. 2024. arXiv: 2407.10879 [hep-ex].

[62] P. Nason, "A New method for combining NLO QCD with shower Monte Carlo algorithms," *JHEP*, vol. 11, p. 040, 2004. DOI: 10.1088/1126-6708/2004/11/040. arXiv: hep-ph/0409146.

[63] S. Frixione, P. Nason, and C. Oleari, "Matching NLO QCD computations with Parton Shower simulations: the POWHEG method," *JHEP*, vol. 11, p. 070, 2007. DOI: 10.1088/1126-6708/2007/11/070. arXiv: 0709.2092 [hep-ph].

[64] S. Alioli, K. Hamilton, P. Nason, C. Oleari, and E. Re, "Jet pair production in POWHEG," *JHEP*, vol. 04, p. 081, 2011. DOI: 10.1007/JHEP04(2011)081. arXiv: 1012.3380 [hep-ph].

[65] J.-C. Winter, F. Krauss, and G. Soff, "A Modified cluster hadronization model," *Eur. Phys. J. C*, vol. 36, pp. 381–395, 2004. DOI: 10.1140/epjc/s2004-01960-8. arXiv: hep-ph/0311085.

[66] J. Bellm *et al.*, "Herwig 7.1 Release Note," May 2017. arXiv: 1705.06919 [hep-ph].

[67] ATLAS Collaboration, "Measurements of Lund subjet multiplicities in 13 TeV proton-proton collisions with the ATLAS detector," *Phys. Lett. B*, vol. 859, p. 139 090, 2024. DOI: 10.1016/j.physletb.2024.139090. arXiv: 2402.13052 [hep-ex].

[68] P. Skands, S. Carrazza, and J. Rojo, "Tuning PYTHIA 8.1: the Monash 2013 Tune," *Eur. Phys. J. C*, vol. 74, no. 8, p. 3024, 2014. DOI: 10.1140/epjc/s10052-014-3024-y. arXiv: 1404.5630 [hep-ph].

[69] LHCb Collaboration, "Handling of the generation of primary events in Gauss, the LHCb simulation framework," in *2010 IEEE Nuclear Science Symposium, Medical Imaging Conference, and 17th Room Temperature Semiconductor Detectors Workshop*, 2010, pp. 1155–1161. DOI: 10.1109/NSSMIC.2010.5873949.

[70] H. Qu and L. Gouskos, "ParticleNet: Jet Tagging via Particle Clouds," *Phys. Rev. D*, vol. 101, no. 5, p. 056 019, 2020. DOI: 10.1103/PhysRevD.101.056019. arXiv: 1902.08570 [hep-ph].

[71] CMS Collaboration, "Exploring small-angle emissions in charm quark jets in proton-proton collisions at $\sqrt{s} = 5.02$ TeV," Jul. 2025. arXiv: 2507.13469 [nucl-ex].

[72] D. Reichelt, S. Caletti, O. Fedkevych, S. Marzani, S. Schumann, and G. Soyez, "Phenomenology of jet angularities at the LHC," *JHEP*, vol. 03, p. 131, 2022. DOI: 10.1007/JHEP03(2022)131. arXiv: 2112.09545 [hep-ph].

[73] ATLAS Collaboration, "Measurement of jet-substructure observables in top quark, W boson and light jet production in proton-proton collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector," *JHEP*, vol. 08, p. 033, 2019. DOI: 10.1007/JHEP08(2019)033. arXiv: 1903.02942 [hep-ex].

[74] ATLAS Collaboration, "Measurement of jet substructure in boosted $\mathrm{t\bar{t}}$ events with the ATLAS detector using 140 fb-1 of 13 TeV pp collisions," *Phys. Rev. D*, vol. 109, no. 11, p. 112 016, 2024. DOI: 10.1103/PhysRevD.109.112016. arXiv: 2312.03797 [hep-ex].

[75] A. Tumasyan *et al.*, "Study of quark and gluon jet substructure in Z+jet and dijet events from pp collisions," *JHEP*, vol. 01, p. 188, 2022. DOI: 10.1007/JHEP01(2022)188. arXiv: 2109.03340 [hep-ex].

[76] ALICE Collaboration, "Measurements of the groomed and ungroomed jet angularities in pp collisions at $\sqrt{s} = 5.02$ TeV," *JHEP*, vol. 05, p. 061, 2022. DOI: 10.1007/JHEP05(2022)061. arXiv: 2107.11303 [nucl-ex].

[77] P. Dhankher, "First measurement of jet angularities with D^0-meson tagged jets with ALICE," *PoS*, vol. HardProbes2023, p. 140, 2024. DOI: 10.22323/1.438.0140.

[78] ALICE Collaboration, "Medium-induced modification of groomed and ungroomed jet mass and angularities in Pb–Pb collisions at sNN=5.02 TeV," *Phys. Lett. B*, vol. 864, p. 139 409, 2025. DOI: 10.1016/j.physletb.2025.139409. arXiv: 2411.03106 [nucl-ex].

[79] J. Bellm *et al.*, "Herwig++ 2.7 Release Note," Oct. 2013. arXiv: 1310.6877 [hep-ph].

[80] M. Bahr *et al.*, "Herwig++ Physics and Manual," *Eur. Phys. J. C*, vol. 58, pp. 639–707, 2008. DOI: 10.1140/epjc/s10052-008-0798-9. arXiv: 0803.0883 [hep-ph].

[81] A. Huss, J. Huston, S. Jones, M. Pellen, and R. Röntsch, "Les Houches 2023 - Physics at TeV Colliders: Report on the Standard Model Precision Wishlist," Tech. Rep., 2025. arXiv: 2504.06689. https://cds.cern.ch/record/2929863.

[82] A. Behring *et al.*, "Flavoured jet algorithms: a comparative study," *JHEP*, vol. 09, p. 149, 2025. DOI: 10.1007/JHEP09(2025)149. arXiv: 2506.13449 [hep-ph].

[83] A. Banfi, G. P. Salam, and G. Zanderighi, "Infrared safe definition of jet flavor," *Eur. Phys. J. C*, vol. 47, pp. 113–124, 2006. DOI: 10.1140/epjc/s2006-02552-4. arXiv: hep-ph/0601139.

[84] D0 Collaboration, "*b*-Jet Identification in the D0 Experiment," *Nucl. Instrum. Meth. A*, vol. 620, pp. 490–517, 2010. DOI: 10.1016/j.nima.2010.03.118. arXiv: 1002.4224 [hep-ex].

[85] M. Cacciari and G. P. Salam, "Dispelling the N^3 myth for the k_t jet-finder," *Phys. Lett. B*, vol. 641, pp. 57–61, 2006. DOI: 10.1016/j.physletb.2006.08.037. arXiv: hep-ph/0512210.

[86] F. Caola, R. Grabarczyk, M. L. Hutt, G. P. Salam, L. Scyboz, and J. Thaler, "Flavored jets with exact anti-kt kinematics and tests of infrared and collinear safety," *Phys. Rev. D*, vol. 108, no. 9, p. 094 010, 2023. DOI: 10.1103/PhysRevD.108.094010. arXiv: 2306.07314 [hep-ph].

[87] A. Banfi, G. P. Salam, and G. Zanderighi, "Accurate QCD predictions for heavy-quark jets at the Tevatron and LHC," *JHEP*, vol. 07, p. 026, 2007. DOI: 10.1088/1126-6708/2007/07/026. arXiv: 0704.2999 [hep-ph].

[88] S. Caletti, A. J. Larkoski, S. Marzani, and D. Reichelt, "Practical jet flavour through NNLO," *Eur. Phys. J. C*, vol. 82, no. 7, p. 632, 2022. DOI: 10.1140/epjc/s10052-022-10568-7. arXiv: 2205.01109 [hep-ph].

[89] M. Czakon, A. Mitov, and R. Poncelet, "Infrared-safe flavoured anti-k_T jets," *JHEP*, vol. 04, p. 138, 2023. DOI: 10.1007/JHEP04(2023)138. arXiv: 2205.11879 [hep-ph].

[90] R. Gauld, A. Huss, and G. Stagnitto, "Flavor Identification of Reconstructed Hadronic Jets," *Phys. Rev. Lett.*, vol. 130, no. 16, p. 161 901, 2023. DOI: 10.1103/PhysRevLett.130.161901. arXiv: 2208.11138 [hep-ph].

[91] J. Bellm *et al.*, "Herwig 7.0/Herwig++ 3.0 release note," *Eur. Phys. J. C*, vol. 76, no. 4, p. 196, 2016. DOI: 10.1140/epjc/s10052-016-4018-8. arXiv: 1512.01178 [hep-ph].

[92] E. Bothmann *et al.*, "Event generation with Sherpa 3," *JHEP*, vol. 12, p. 156, 2024. DOI: 10.1007/JHEP12(2024)156. arXiv: 2410.22148 [hep-ph].

[93] C. Bierlich *et al.*, "A comprehensive guide to the physics and usage of PYTHIA 8.3," *SciPost Phys. Codeb.*, vol. 2022, p. 8, 2022. DOI: 10.21468/SciPostPhysCodeb.8. arXiv: 2203.11601 [hep-ph].

[94] H. Bartosik and G. Rumolo, "Performance of the LHC injector chain after the upgrade and potential development," in *Snowmass 2021*, Mar. 2022. arXiv: 2203.09202 [physics.acc-ph].

[95] https://home.cern/news/news/accelerators/accelerator-report-lhc-run-3-achieves-record-breaking-integrated-luminosity.

[96] "The High-Luminosity LHC Project," Tech. Rep., 2016. https://cds.cern.ch/record/2199189.

[97] O. Brüning and M. Zerlauth, "LHC operation and the High-Luminosity LHC upgrade project," Tech. Rep., 2025. arXiv: 2505.03535. https://cds.cern.ch/record/2935496.

[98] ATLAS Collaboration, "Constituent-level pile-up mitigation techniques in ATLAS," CERN, Geneva, Tech. Rep., 2017. https://cds.cern.ch/record/2281055.

[99] ATLAS Collaboration, "Primary Vertex identification using deep learning in ATLAS," CERN, Geneva, Tech. Rep., 2023. https://cds.cern.ch/record/2858348.

[100] J. Pequenao, *Computer generated image of the whole ATLAS detector*, 2008. https://cds.cern.ch/record/1095924.

[101] A. M. Rodriguez Vera and J. Antunes Pequenao, *ATLAS Detector Magnet System*, 2021. https://cds.cern.ch/record/2770604.

[102] J. Pequenao, *Computer generated image of the ATLAS inner detector*, 2008. https://cds.cern.ch/record/1095926.

[103] H. Sandaker, "SCT and TRT Performance from Cosmic Ray Runs," 2007. DOI: 10.5170/CERN-2007-007.168. https://cds.cern.ch/record/1089261.

[104] O. Ducu, "Electron efficiency in LHC Run-2 with the ATLAS experiment," *Int. J. Mod. Phys. A*, vol. 40, no. 08, p. 2 443 022, 2025. DOI: 10.1142/S0217751X2443022X. arXiv: 2412.19323 [hep-ex].

[105] C. Vittori, "Inclusive z boson production and in association with b-jets in proton-proton collisions at 13 tev with the atlas experiment," Ph.D. dissertation, University of Bologna, Apr. 2018. https://amsdottorato.unibo.it/id/eprint/8604/.

[106] J. Pequenao, *Computer Generated image of the ATLAS calorimeter*, 2008. https://cds.cern.ch/record/1095927.

[107] ATLAS Collaboration, *ATLAS muon spectrometer: Technical design report*. Geneva: CERN, 1997. https://cds.cern.ch/record/331068.

[108] ATLAS Collaboration, "Resolution of the ATLAS muon spectrometer monitored drift tubes in LHC Run 2," *JINST*, vol. 14, no. 09, P09011, 2019. DOI: 10.1088/1748-0221/14/09/P09011. arXiv: 1906.12226 [hep-ex].

[109] ATLAS Collaboration, "Muon reconstruction efficiency and momentum resolution of the ATLAS experiment in proton-proton collisions at $\sqrt{s} = 7$ TeV in 2010," *Eur. Phys. J. C*, vol. 74, no. 9, p. 3034, 2014. DOI: 10.1140/epjc/s10052-014-3034-9. arXiv: 1404.4562 [hep-ex].

[110] J. Pequenao, *Computer generated image of the ATLAS Muons subsystem*, 2008. https://cds.cern.ch/record/1095929.

[111] S. Mehlhase, "ATLAS detector slice (and particle visualisations)," 2021. https://cds.cern.ch/record/2770815.

[112] ATLAS Collaboration, "Electron and photon performance measurements with the ATLAS detector using the 2015–2017 LHC proton-proton collision data," *JINST*, vol. 14, no. 12, P12006, 2019. DOI: 10.1088/1748-0221/14/12/P12006. arXiv: 1908.00005 [hep-ex].

[113] ATLAS Collaboration, "Muon reconstruction performance of the ATLAS detector in proton–proton collision data at $\sqrt{s}$ =13 TeV," *Eur. Phys. J. C*, vol. 76, no. 5, p. 292, 2016. DOI: 10.1140/epjc/s10052-016-4120-y. arXiv: 1603.05598 [hep-ex].

[114] ATLAS Collaboration, "The performance of missing transverse momentum reconstruction and its significance with the ATLAS detector using 140 fb^{-1} of $\sqrt{s}$ = 13 TeV pp collisions," *Eur. Phys. J. C*, vol. 85, no. 6, p. 606, 2025. DOI: 10.1140/epjc/s10052-025-14062-8. arXiv: 2402.05858 [hep-ex].

[115] ATLAS Collaboration, "Variable Radius, Exclusive-k$_T$, and Center-of-Mass Subjet Reconstruction for Higgs($\to b\bar{b}$) Tagging in ATLAS," CERN, Geneva, Tech. Rep., 2017. https://cds.cern.ch/record/2268678.

[116] ATLAS Collaboration, "Topological cell clustering in the ATLAS calorimeters and its performance in LHC Run 1," *Eur. Phys. J. C*, vol. 77, p. 490, 2017. DOI: 10.1140/epjc/s10052-017-5004-5. arXiv: 1603.02934 [hep-ex].

[117] ATLAS Collaboration, "Jet energy scale measurements and their systematic uncertainties in proton-proton collisions at $\sqrt{s}$ = 13 TeV with the ATLAS detector," *Phys. Rev. D*, vol. 96, no. 7, p. 072 002, 2017. DOI: 10.1103/PhysRevD.96.072002. arXiv: 1703.09665 [hep-ex].

[118] 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*, vol. 81, no. 8, p. 689, 2021. DOI: 10.1140/epjc/s10052-021-09402-3. arXiv: 2007.02645 [hep-ex].

[119] "Jet global sequential corrections with the ATLAS detector in proton-proton collisions at sqrt(s) = 8 TeV," CERN, Geneva, Tech. Rep., 2015. https://cds.cern.ch/record/2001682.

[120] "Tagging and suppression of pileup jets with the ATLAS detector," CERN, Geneva, Tech. Rep., 2014. https://cds.cern.ch/record/1700870.

[121] Wikimedia Commons, *File: B-tagging diagram.png — Wikimedia Commons, the free media repository*, 2023. https://commons.wikimedia.org/w/index.php?title=File:B-tagging_diagram.png&oldid=809849110.

[122] A. Duperrin, "Flavour Tagging with Graph Neural Network with the ATLAS Detector," 2023. https://cds.cern.ch/record/2855275.

[123] ATLAS Collaboration, "ATLAS flavour-tagging algorithms for the LHC Run 2 pp collision dataset," *Eur. Phys. J. C*, vol. 83, no. 7, p. 681, 2023. DOI: 10.1140/epjc/s10052-023-11699-1. arXiv: 2211.16345 [physics.data-an].

[124] ATLAS Collaboration, "Optimisation and performance studies of the ATLAS b-tagging algorithms for the 2017-18 LHC run," CERN, Geneva, Tech. Rep., 2017. https://cds.cern.ch/record/2273281.

[125] ATLAS Collaboration, "Identification of Jets Containing b-Hadrons with Recurrent Neural Networks at the ATLAS Experiment," CERN, Geneva, Tech. Rep., 2017. https://cds.cern.ch/record/2255226.

[126] ATLAS Collaboration, "Secondary vertex finding for jet flavour identification with the ATLAS detector," CERN, Geneva, Tech. Rep., 2017. https://cds.cern.ch/record/2270366.

[127] ATLAS Collaboration, "Topological b-hadron decay reconstruction and identification of b-jets with the JetFitter package in the ATLAS experiment at the LHC," CERN, Geneva, Tech. Rep., 2018. https://cds.cern.ch/record/2645405.

[128] 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*, vol. 79, no. 11, p. 970, 2019. DOI: 10.1140/epjc/s10052-019-7450-8. arXiv: 1907.05120 [hep-ex].

[129] ATLAS Collaboration, "Anomaly detection search for new resonances decaying into a Higgs boson and a generic new particle X in hadronic final states using $\sqrt{s} = 13$ TeV pp collisions with the ATLAS detector," *Phys. Rev. D*, vol. 108, p. 052 009, 2023. DOI: 10.1103/PhysRevD.108.052009. arXiv: 2306.03637 [hep-ex].

[130] ATLAS Collaboration, "Identification of Boosted Higgs Bosons Decaying Into $b\bar{b}$ With Neural Networks and Variable Radius Subjets in ATLAS," CERN, Geneva, Tech. Rep., 2020. https://cds.cern.ch/record/2724739.

[131] ATLAS Collaboration, "Efficiency corrections for a tagger for boosted $H \to b\bar{b}$ decays in pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector," CERN, Geneva, Tech. Rep., 2021. https://cds.cern.ch/record/2777811.

[132] T. Sjöstrand, S. Ask, J. R. Christiansen, *et al.*, "An introduction to PYTHIA 8.2," *Comput. Phys. Commun.*, vol. 191, pp. 159–177, 2015. DOI: 10.1016/j.cpc.2015.01.024. arXiv: 1410.3012 [hep-ph].

[133] D. J. Lange, "The EvtGen particle decay simulation package," *Nucl. Instrum. Meth. A*, vol. 462, S. Erhan, P. Schlein, and Y. Rozen, Eds., pp. 152–155, 2001. DOI: 10.1016/S0168-9002(01)00089-4.

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

[135] R. D. Ball *et al.*, "Parton distributions with LHC data," *Nucl. Phys. B*, vol. 867, pp. 244–289, 2013. DOI: 10.1016/j.nuclphysb.2012.10.003. arXiv: 1207.1303 [hep-ph].

[136] ATLAS Collaboration, "Luminosity determination in pp collisions at $\sqrt{s} = 13$ TeV using the ATLAS detector at the LHC," *Eur. Phys. J. C*, vol. 83, no. 10, p. 982, 2023. DOI: 10.1140/epjc/s10052-023-11747-w. arXiv: 2212.09379 [hep-ex].

[137] S. Catani, F. Krauss, R. Kuhn, and B. R. Webber, "QCD matrix elements + parton showers," *JHEP*, vol. 11, p. 063, 2001. DOI: 10.1088/1126-6708/2001/11/063. arXiv: hep-ph/0109231.

[138] ATLAS Collaboration, "Modelling and computational improvements to the simulation of single vector-boson plus jet processes for the ATLAS experiment," *JHEP*, vol. 08, p. 089, 2022. DOI: 10.1007/JHEP08(2022)089. arXiv: 2112.09588 [hep-ex].

[139] T. Gleisberg and S. Hoeche, "Comix, a new matrix element generator," *JHEP*, vol. 12, p. 039, 2008. DOI: 10.1088/1126-6708/2008/12/039. arXiv: 0808.3674 [hep-ph].

[140] F. Buccioni, J.-N. Lang, J. M. Lindert, *et al.*, "OpenLoops 2," *Eur. Phys. J. C*, vol. 79, no. 10, p. 866, 2019. DOI: 10.1140/epjc/s10052-019-7306-2. arXiv: 1907.13071 [hep-ph].

[141] F. Cascioli, P. Maierhofer, and S. Pozzorini, "Scattering Amplitudes with Open Loops," *Phys. Rev. Lett.*, vol. 108, p. 111 601, 2012. DOI: 10.1103/PhysRevLett.108.111601. arXiv: 1111.5206 [hep-ph].

[142] A. Denner, S. Dittmaier, and L. Hofer, "Collier: a fortran-based Complex One-Loop LIbrary in Extended Regularizations," *Comput. Phys. Commun.*, vol. 212, pp. 220–238, 2017. DOI: 10.1016/j.cpc.2016.10.013. arXiv: 1604.06792 [hep-ph].

[143] NNPDF Collaboration, "Parton distributions for the LHC Run II," *JHEP*, vol. 04, p. 040, 2015. DOI: 10.1007/JHEP04(2015)040. arXiv: 1410.8849 [hep-ph].

[144] S. Schumann and F. Krauss, "A Parton shower algorithm based on Catani-Seymour dipole factorisation," *JHEP*, vol. 03, p. 038, 2008. DOI: 10.1088/1126-6708/2008/03/038. arXiv: 0709.1027 [hep-ph].

[145] S. Hoeche, F. Krauss, M. Schonherr, and F. Siegert, "A critical appraisal of NLO+PS matching methods," *JHEP*, vol. 09, p. 049, 2012. DOI: 10.1007/JHEP09(2012)049. arXiv: 1111.1220 [hep-ph].

[146] S. Hoeche, F. Krauss, M. Schonherr, and F. Siegert, "QCD matrix elements + parton showers: The NLO case," *JHEP*, vol. 04, p. 027, 2013. DOI: 10.1007/JHEP04(2013)027. arXiv: 1207.5030 [hep-ph].

[147] S. Hoeche, F. Krauss, S. Schumann, and F. Siegert, "QCD matrix elements and truncated showers," *JHEP*, vol. 05, p. 053, 2009. DOI: 10.1088/1126-6708/2009/05/053. arXiv: 0903.1219 [hep-ph].

[148] S. Frixione, P. Nason, and G. Ridolfi, "A Positive-weight next-to-leading-order Monte Carlo for heavy flavour hadroproduction," *JHEP*, vol. 09, p. 126, 2007. DOI: 10.1088/1126-6708/2007/09/126. arXiv: 0707.3088 [hep-ph].

[149] ATLAS Collaboration, "Studies on top-quark Monte Carlo modelling for Top2016," CERN, Geneva, Tech. Rep., 2016. https://cds.cern.ch/record/2216168.

[150] M. Czakon and A. Mitov, "Top++: A Program for the Calculation of the Top-Pair Cross-Section at Hadron Colliders," *Comput. Phys. Commun.*, vol. 185, p. 2930, 2014. DOI: 10.1016/j.cpc.2014.06.021. arXiv: 1112.5675 [hep-ph].

[151] S. Frixione, E. Laenen, P. Motylinski, B. R. Webber, and C. D. White, "Single-top hadroproduction in association with a W boson," *JHEP*, vol. 07, p. 029, 2008. DOI: 10.1088/1126-6708/2008/07/029. arXiv: 0805.3067 [hep-ph].

[152] E. Bothmann *et al.*, "Event Generation with Sherpa 2.2," *SciPost Phys.*, vol. 7, no. 3, p. 034, 2019. DOI: 10.21468/SciPostPhys.7.3.034. arXiv: 1905.09127 [hep-ph].

[153] 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., 2016. https://cds.cern.ch/record/2206965.

[154] Geant4 Collaboration, "GEANT4 - A Simulation Toolkit," *Nucl. Instrum. Meth. A*, vol. 506, pp. 250–303, 2003. DOI: 10.1016/S0168-9002(03)01368-8.

[155] ATLAS Collaboration, "The ATLAS Simulation Infrastructure," *Eur. Phys. J. C*, vol. 70, pp. 823–874, 2010. DOI: 10.1140/epjc/s10052-010-1429-9. arXiv: 1005.4568 [physics.ins-det].

[156] F. Chollet *et al.*, *Keras*, https://keras.io, 2015.

[157] G. Avoni *et al.*, "The new LUCID-2 detector for luminosity measurement and monitoring in ATLAS," *JINST*, vol. 13, no. 07, P07017, 2018. DOI: 10.1088/1748-0221/13/07/P07017.

[158] G. D'Agostini, "A Multidimensional unfolding method based on Bayes' theorem," *Nucl. Instrum. Meth. A*, vol. 362, pp. 487–498, 1995. DOI: 10.1016/0168-9002(95)00274-X.

[159] A. Andreassen, P. T. Komiske, E. M. Metodiev, B. Nachman, and J. Thaler, "OmniFold: A Method to Simultaneously Unfold All Observables," *Phys. Rev. Lett.*, vol. 124, no. 18, p. 182 001, 2020. DOI: 10.1103/PhysRevLett.124.182001. arXiv: 1911.09107 [hep-ph].

[160] A. Ghira, S. Marzani, and G. Soyez, "The Lund b-jet plane," Dec. 2025. arXiv: 2512.17408 [hep-ph].

[161] ATLAS Collaboration, "Improving jet substructure performance in ATLAS using Track-CaloClusters," CERN, Geneva, Tech. Rep., 2017. https://cds.cern.ch/record/2275636.

[162] ATLAS Collaboration, "Optimisation of large-radius jet reconstruction for the ATLAS detector in 13 TeV proton–proton collisions," *Eur. Phys. J. C*, vol. 81, no. 4, p. 334, 2021. DOI: 10.1140/epjc/s10052-021-09054-3. arXiv: 2009.04986 [hep-ex].

[163] ATLAS Collaboration, "Measurement of soft-drop jet observables in pp collisions with the ATLAS detector at $\sqrt{s}$ =13 TeV," *Phys. Rev. D*, vol. 101, no. 5, p. 052 007, 2020. DOI: 10.1103/PhysRevD.101.052007. arXiv: 1912.09837 [hep-ex].

[164] ATLAS Collaboration, "Simultaneous Unbinned Differential Cross-Section Measurement of Twenty-Four Z+jets Kinematic Observables with the ATLAS Detector," *Phys. Rev. Lett.*, vol. 133, no. 26, p. 261 803, 2024. DOI: 10.1103/PhysRevLett.133.261803. arXiv: 2405.20041 [hep-ex].

[165] P. T. Komiske, S. Kryhin, and J. Thaler, "Disentangling quarks and gluons in CMS open data," *Phys. Rev. D*, vol. 106, no. 9, p. 094 021, 2022. DOI: 10.1103/PhysRevD.106.094021. arXiv: 2205.04459 [hep-ph].

[166] LHCb Collaboration, "Multidifferential study of identified charged hadron distributions in Z-tagged jets in proton-proton collisions at $\sqrt{s}$ =13 TeV," *Phys. Rev. D*, vol. 108, p. L031103, 2023. DOI: 10.1103/PhysRevD.108.L031103. arXiv: 2208.11691 [hep-ex].

[167] H1 Collaboration, "Measurement of Lepton-Jet Correlation in Deep-Inelastic Scattering with the H1 Detector Using Machine Learning for Unfolding," *Phys. Rev. Lett.*, vol. 128, no. 13, p. 132 002, 2022. DOI: 10.1103/PhysRevLett.128.132002. arXiv: 2108.12376 [hep-ex].

[168] H1 Collaboration, "Unbinned deep learning jet substructure measurement in high Q2ep collisions at HERA," *Phys. Lett. B*, vol. 844, p. 138 101, 2023. DOI: 10.1016/j.physletb.2023.138101. arXiv: 2303.13620 [hep-ex].

[169] Y. Song, "Measurement of CollinearDrop jet mass and its correlation with SoftDrop groomed jet substructure observables in $\sqrt{s}$ = 200 GeV pp collisions by STAR," Jul. 2023. arXiv: 2307.07718 [nucl-ex].

[170] F. Topsoe, "Some inequalities for information divergence and related measures of discrimination," *IEEE Transactions on Information Theory*, vol. 46, no. 4, pp. 1602–1609, 2000. DOI: 10.1109/18.850703.

Ringraziamenti

Non potrei definire questa tesi completa senza una sezione dedicata a tutti coloro che mi hanno accompagnato in questo percorso. Come prima cosa, voglio ringraziarti, Fede, per tutto ciò che hai fatto per me in questi anni. Oltre ad essere una fonte inesauribile di conoscenza di fisica e di energia, sei sempre stato disponibilissimo e alla mano. Grazie a te so condurre un'analisi dalla A alla Z, ma anche quali sono i locali migliori di Genova/Londra e come usare un cucchiaino come cavatappi. Non sarei mai riuscito ad arrivare fino in fondo senza la tua guida, e te ne sarò per sempre grato.

Allo stesso modo, vorrei ringraziare Evelin, Matt e Simone per aver accettato di revisionare questa tesi e/o far parte della commissione. Grazie anche ad Evelin per avermi preso in considerazione per lavorare con lei.

Grazie anche a tutto il team di $Z + bb$: Radek, Toni, Tommaso, Hanfei, Mario e Giusy. Senza di voi non avrei mai scoperto la "bellezza" dei jet.

Grazie poi a Francesco Giuli, che considero un mentore al quale mi sono sempre potuto rivolgere per consigli, e a Ludovica, che ha dovuto subire tutti i miei sfoghi. Grazie anche alla mentore storica, Daniela.

Allargando gli orizzonti, questo dottorato mi ha permesso di conoscere persone da ogni parte del mondo, fra cui posti esotici come il Brasile e la Sardegna[1]. Obrigado aos meus amigos brasileiros Úrsula, Osvaldo, Isabela, Victoria, Julia, Leonardo, Bia e Pedro que me consideram um carioca honorário. A proposito di posti esotici, grazie ad Oleksii ed Ilia per avermi accompagnato nel Atacama.

Da buon piccione viaggiatore, ritorno sempre a casa. Grazie a Fra, Ric, Andrea, Giulia, Valeria e ancora Tommaso per avermi sempre accolto quando tornavo di sorpresa. Sarebbe stato bellissimo fare tutto il dottorato insieme a voi.

Non potrei non menzionare Max, Craig, Rebecca, Manuel, Savannah, Eleanor, Fiona, Gianpiero, Alessandro, Pepe, Nico ed Adri per le merende a base di chisme e le feste a tema.

Infine un ultimo grazie agli amici di sempre: Amo, Alle, Nino, Rubens, Marti, Virgi, Vale, Frollo, Mago, Matte e Uolt, e ai miei per esserci stati sempre.

[1] Cito la Sardegna solo per nominare Alessandra, che non saprei come includere altrimenti.