THÈSE

en vue de l'obtention du grade de

Docteur en Physique de l'Univers
délivré par Sorbonne Université

École doctorale des Sciences de la Terre et de l'Environnement et Physique de l'Univers de Paris (ED 560)

Laboratoire de Physique Nucléaire et des Hautes Énergies - Équipe Masse et Interactions Fondamentales

---

Étalonnage, mesure des sections efficaces de production et étude de la sous-structure des jets avec les données d'ATLAS, perspectives au FCC-ee

---

par **Line Delagrange**

Dirigée par Bogdan MALAESCU

Présentée et soutenue publiquement le
Devant la commission d'examen formée de :

Mme.	Delphine HARDIN	*Sorbonne Université*	Présidente du jury
Mme.	Katerina LIPKA	*University of Wuppertal, DESY*	Rapportrice
M.	Peter LOCH	*University of Arizona*	Rapporteur
M.	Fabrice BALLI	*CEA Saclay*	Examinateur
M.	Bogdan MALAESCU	*CNRS*	Directeur de thèse

Contents

Introduction **9**

1 The Theoretical Framework **13**
1.1 The Standard Model of Particle Physics . . . 13
1.1.1 Quantum Field Theory . . . 13
1.1.1.1 Notations and conventions . . . 14
1.1.1.2 Lagrangian formalism and symmetries . . . 14
1.1.1.3 Gauge Theory . . . 15
1.1.2 Standard Model . . . 15
1.1.2.1 Symmetries . . . 15
1.1.2.2 Spontaneous symmetry breaking . . . 16
1.1.2.3 Field content of the Standard Model . . . 16
1.1.2.4 Mathematical definition of the Standard Model . . . 17
1.2 Quantum Chromodynamics . . . 18
1.2.1 The QCD Lagrangian . . . 18
1.2.1.1 Expression . . . 18
1.2.1.2 Matrix elements . . . 19
1.2.1.3 Renormalization . . . 20
1.2.1.4 The strong coupling constant . . . 20
1.2.2 From partons to hadrons . . . 20
1.2.2.1 Perturbative QCD . . . 21
1.2.2.2 Non-perturbative QCD . . . 24
1.2.2.3 The whole picture . . . 26
1.2.3 Monte Carlo Simulations . . . 27
1.2.3.1 Event Generation . . . 28
1.2.3.2 Parton-level prediction . . . 28
1.2.4 Jets . . . 29
1.2.4.1 Jet definition . . . 29
1.2.4.2 Jet algorithms . . . 30
1.2.4.3 Jets in hadron colliders . . . 32
1.2.4.4 Jet substructure . . . 33

2 The ATLAS experiment and object reconstruction in ATLAS **35**
2.1 LHC at CERN . . . 35
2.1.1 The acceleration chain . . . 36
2.1.2 Data-taking . . . 37
2.2 The ATLAS experiment . . . 38

2.2.1 The inner detector . . . 40
2.2.2 The calorimeter system . . . 42
2.2.2.1 Principle . . . 42
2.2.2.2 Layout . . . 42
2.2.2.3 Electromagnetic calorimeter . . . 43
2.2.2.4 Hadronic calorimeter (HCAL) . . . 44
2.2.3 The muon spectrometer . . . 45
2.2.4 The trigger system . . . 46
2.3 Reconstruction and data processing in ATLAS . . . 47
2.3.1 General object reconstruction . . . 47
2.3.1.1 Tracks . . . 47
2.3.1.2 Vertices . . . 48
2.3.1.3 Electrons and photons . . . 49
2.3.1.4 Others . . . 49
2.3.2 Jet reconstruction . . . 50
2.3.2.1 Topological clusters . . . 50
2.3.2.2 Local hadronic cell weighting . . . 51
2.3.2.3 PFlow objects . . . 51
2.3.2.4 Constituent-level pile-up rejection . . . 52
2.3.2.5 Jet reconstruction . . . 52
2.3.3 Jet Cleaning . . . 54
2.3.3.1 Jet Vertex Tagger . . . 54
2.3.3.2 Event level selection . . . 54
2.3.4 MC simulations . . . 56
2.3.5 The ATLAS data flow . . . 57

3 The Jet η-intercalibration 59
3.1 Current status of jet calibration in ATLAS . . . 60
3.1.1 Reconstruction and origin correction . . . 61
3.1.2 Pile-up corrections . . . 61
3.1.3 MC absolute calibration (MCJES) . . . 62
3.1.4 Global sequential calibration (GSC) . . . 63
3.1.5 Residual in situ calibration . . . 65
3.1.5.1 η-intercalibration . . . 67
3.1.5.2 Z/γ+ jet method . . . 67
3.1.5.3 Multijet balance . . . 68
3.1.5.4 Systematic uncertainties . . . 69
3.1.5.5 Combination of the in situ measurements . . . 70
3.1.5.6 Single-particle response calibration . . . 71
3.1.6 JES uncertainties application and reduction scheme . . . 73
3.1.7 The Jet Energy Resolution (JER) . . . 74
3.1.7.1 JER measurement using dijet events . . . 74
3.1.7.2 Noise measurement . . . 75
3.1.7.3 Combination and application . . . 77
3.2 Detailed description of the η-intercalibration . . . 78
3.2.1 Event selection . . . 78

3.2.1.1 Datasets . . . 78
3.2.1.2 Basic event selection and triggers . . . 79
3.2.1.3 Dijet selection . . . 79
3.2.1.4 Binning . . . 80
3.2.2 Derivation of the intercalibration factors . . . 82
3.2.2.1 Determining the mean asymmetry . . . 82
3.2.2.2 The central reference method . . . 84
3.2.2.3 The matrix method . . . 84
3.2.3 Derivation of the calibration factors . . . 87
3.2.4 Uncertainties . . . 88
3.2.4.1 Systematic uncertainties . . . 88
3.2.4.2 Statistical uncertainties . . . 90
3.3 Improvements to the η-intercalibration method . . . 91
3.3.1 Recovering constraints . . . 91
3.3.1.1 Discarded constraints . . . 91
3.3.1.2 Asymmetry binning . . . 93
3.3.1.3 Fit range . . . 94
3.3.1.4 Results . . . 94
3.3.2 The global χ^2/NDF test . . . 95
3.3.2.1 Principle . . . 95
3.3.2.2 Results . . . 95
3.3.2.3 Perspective . . . 96
3.3.3 The bootstrap method for the evaluation of statistical uncertainties 97
3.3.3.1 Principle . . . 97
3.3.3.2 For the η-intercalibration . . . 98
3.3.3.3 Results . . . 98
3.3.3.4 Application in physics analysis . . . 100
3.4 Deriving the official η-intercalibration corrections and uncertainties . . . 103
3.4.1 Pre-recommendations phase-I . . . 103
3.4.2 Pre-recommendations phase-II . . . 104
3.4.2.1 Trigger efficiency study . . . 105
3.4.2.2 Phase-II pre-recommendations results . . . 106
3.5 Conclusion . . . 107

4 The Inclusive Jet Cross-Section Analysis 111
4.1 The inclusive jet double-differential cross-section measurement . . . 111
4.2 Dataset . . . 112
4.2.1 Data . . . 112
4.2.2 MC simulations . . . 114
4.3 Unfolding . . . 114
4.3.1 Motivation . . . 114
4.3.2 General unfolding strategy . . . 114
4.3.3 The IDS method . . . 115
4.3.4 Data-driven closure test . . . 117
4.3.5 Application to the inclusive jet cross-section analysis . . . 118
4.3.5.1 Object definition . . . 118

4.3.5.2 Data-driven closure test . . . 118
4.4 Experimental uncertainties . . . 119
4.4.1 Statistical uncertainties . . . 120
4.4.1.1 Evaluation through the bootstrap method . . . 120
4.4.1.2 Correlation with the η-intercalibration factors . . . 121
4.4.2 Systematic uncertainties . . . 122
4.4.2.1 Missing Tile Modules . . . 122
4.4.2.2 JES uncertainties . . . 122
4.4.2.3 JER uncertainties . . . 126
4.4.2.4 Other uncertainties . . . 126
4.5 Results . . . 127
4.5.1 Unfolded data and uncertainties . . . 127
4.5.2 Comparison to theory predictions . . . 132
4.5.2.1 Theory predictions . . . 132
4.5.2.2 Uncertainties . . . 132
4.5.2.3 Smoothing of the predicted jet cross-section spectra . . . 133
4.5.2.4 Quantitative comparison with data . . . 134
4.5.2.5 QCD interpretation . . . 135
4.6 Conclusion . . . 135

5 The Lund Jet Plane Analysis 137
5.1 The Lund Jet Plane . . . 137
5.1.1 Principle . . . 137
5.1.2 Construction . . . 139
5.1.3 Analytical properties and applications . . . 140
5.2 Dataset . . . 141
5.2.1 Data . . . 141
5.2.2 Monte Carlo simulations . . . 142
5.3 Unfolding . . . 142
5.3.1 IDS vs IBU regularisation . . . 142
5.3.1.1 IBU regularisation principle . . . 142
5.3.1.2 Comparison with IDS . . . 143
5.3.2 Choice of the number of unfolding iterations . . . 143
5.3.2.1 Data-driven non-closure test . . . 143
5.3.2.2 Statistical uncertainties . . . 145
5.3.2.3 Final choice . . . 145
5.4 Experimental uncertainties . . . 146
5.4.1 Statistical uncertainties . . . 147
5.4.2 Systematic uncertainties . . . 147
5.4.2.1 JES and JER uncertainties . . . 147
5.4.2.2 Tracking related uncertainties . . . 147
5.4.2.3 Modelling uncertainty . . . 149
5.4.2.4 Pile-up reweighting . . . 150
5.4.2.5 Other uncertainties . . . 150
5.5 Results . . . 155
5.6 Conclusion and perspective . . . 155

6 Jet studies in an FCC-ee environment **157**
6.1 The Future Circular Collider . . . 157
6.1.1 Project . . . 157
6.1.2 Detector concepts . . . 158
6.1.3 Working in an e^+e^- environment . . . 159
6.1.3.1 Advantages . . . 159
6.1.3.2 Challenges . . . 160
6.1.4 Planned schedule . . . 161
6.2 Sensitivity to α_s at the FCC-ee . . . 162
6.2.1 Motivations . . . 162
6.2.2 Dataset and event selection . . . 162
6.2.3 $R_{3/2}$ . . . 163
6.2.3.1 Definition . . . 163
6.2.3.2 Dependence on jet algorithm variables . . . 163
6.2.3.3 Perspective at FCC-ee . . . 164
6.2.3.4 Discussions . . . 166
6.2.4 Lund Jet Plane . . . 166
6.3 Conclusion and perspective . . . 167

Conclusion **169**

Bibliography **171**

A Appendix for the jet η-intercalibration **185**
A.1 Sources of uncertainties in the Jet Energy Scale . . . 185
A.2 Trigger Strategy for data 2022 . . . 187

B Appendix for the Lund jet plane measurement **189**

Remerciements

Cette thèse présente le fruit de mon travail de recherche réalisé au cours de ces trois dernières années. Elle n'aurait pas pu voir le jour sans le soutien et l'aide précieuse de nombreuses personnes que je souhaite remercier ici.

Avant toute chose, je souhaite exprimer ma profonde gratitude à Bogdan Malaescu. Au-delà de l'encadrement scientifique, il m'a formée à devenir chercheuse, en m'inculquant la rigueur intellectuelle, le goût des choses bien faites et de la précision, qui resteront pour moi des repères précieux tout au long de ma carrière. Son exigence, son enthousiasme, ses conseils attentifs et sa disponibilité m'ont permis de progresser tant sur le plan scientifique que personnel.

Je remercie également Luc Poggioli, avec qui j'ai partagé de nombreuses discussions passionnantes sur l'avenir de la physique des particules. Nos échanges, qu'ils portent sur des aspects techniques, théoriques ou même géopolitiques, m'ont ouvert de nouvelles perspectives sur ce que pourrait être la recherche dans les prochaines décennies.

Je souhaite exprimer ma reconnaissance à Katerina Lipka, Peter Loch, Fabrice Balli et Delphine Hardin, qui ont pris le temps de lire se manuscrit avec attention et d'évaluer mon travail avec bienveillance.

Je remercie chaleureusement l'ensemble des membres du laboratoire pour le cadre de travail dans lequel j'ai pu mener cette thèse dans les meilleures conditions. Je remercie en particulier les membres du groupe ATLAS, Vera et Tristan pour les puzzles, et Pauline ma marraine du labo.

Je tiens aussi à exprimer ma gratitude à tous les professeurs qui m'ont accompagnée au long de ma scolarité. Une mention particulière revient à Mr Schwarz, qui m'a fait découvrir le Large Hadron Rap en cours de physique de première, qui est à l'origine de mon intérêt pour la physique des collisioneurs.

Je souhaite également remercier mes collègues au CERN avec qui j'ai eu la chance de collaborer: Tancredi, Daniel, Magda, Tae, Reina, Matt, Fabrice, Jeff, Pierre-Antoine. Leurs retours constructifs, leurs suggestions et leur confiance dans mon travail ont été essentiels pour progresser et m'ont donné l'assurance nécessaire pour m'affirmer en tant que jeune chercheuse. Je souhaite tout particulièrement remercier Louis Ginabat et Lata Panwar pour leur présence et leur aide au quotidien quand je me suis lancée dans ce

travail de recherche.

Je remercie également tous les doctorants et post-doctorants du laboratoire, qui ont contribué à rendre cette expérience agréable au quotidien, avec une mention spéciale pour Claire, Romain et Paul, Laura, Artur et Enya.

À mes amis en dehors du LPNHE, qui ont marqué mon parcours, au lycée, en classes préparatoires, à l'ENS et en master, je veux dire un immense merci. Je pense en particulier à Lucie, Marie D., Flavie, Fanny, Héloïse, Camille, Manon, Antoine, Jean-Baptiste, Madeleine, Cassandra, Solal, Lucien, Yanis, Alexis, Marie H. Merci d'avoir rendu ces années inoubliables.

Je n'oublie pas les copines de mon équipe de rugby qui m'ont offert une échappatoire en dehors du travail: merci de m'avoir donné la possibilité de collisionner autre chose que des particules.

À mes parents, je veux exprimer ma reconnaissance la plus profonde. Merci d'avoir cru en moi depuis toujours, de m'avoir soutenue sans relâche et d'avoir cultivé en moi la confiance et la curiosité qui m'ont guidée tout au long de ce parcours.

Enfin, à Anaëlle, je dédie ces dernières lignes. Merci pour ton soutien indéfectible, ta patience et ton écoute. Merci d'avoir été à mes côtés, dans les réussites comme dans les difficultés. Rien de tout cela n'aurait eu la même saveur sans toi.

Abstract

Abstract:

The Standard Model (SM) of particle physics describes the fundamental particles that make up matter, along with three of the four fundamental forces. So far, it has been very successful, with its predictions confirmed by a wide range of experimental results, but it does not provide explanations for several open questions in fundamental physics. This thesis explores a strategy for testing the SM: comparing its predictions with high-precision measurements. It focuses specifically on jets, an essential component of the Quantum Chromodynamics (QCD) sector of the SM, allowing to probe the latter at the smallest distances accessible in laboratory. The analyses are based on data collected by the ATLAS experiment at CERN, using proton–proton collisions at a centre-of-mass energy of 13 TeV, recorded during Run 2 of the LHC. Jets are among the most prominent observables at the LHC. They are collimated sprays of particles resulting from the hadronisation of quarks or gluons, leaving tracks in the tracking detectors and depositing energy in the calorimeters. To reconstruct the originating jet, these signals are clustered into what are known as "reconstructed jets". However, reconstructed jets differ from the original particle-level jets due to detector effects. To ensure meaningful comparisons between theoretical predictions and experimental results, these effects must be corrected through a precise calibration process and the unfolding of the resolution effects.

I have worked on the η-intercalibration, a step in the jet calibration procedure that aims to ensure a uniform response across the full detector acceptance. I derived the calibration on Run 2 and Run 3 data, developing a new trigger strategy for the latter. To improve the robustness of the method, I extended the overconstrained system used to extract the calibration factors by incorporating additional terms and implementing a compatibility test among them. In addition, I developed a more reliable method for estimating the statistical uncertainties on the calibration factors. This approach allows for the study of their correlations.

I have contributed to the differential measurement of the inclusive jet production cross section. A precise quantitative comparison of the measured spectra with theoretical prediction allows the extraction of the strong coupling constant α_s. My work focused on the unfolding of the data, a procedure that statistically corrects for detector effects to recover the particle-level distribution. I evaluated the unfolding bias, introduced by differences in shape between data and simulation and propagated the systematic uncertainties through

the unfolding process. In addition, I propagated the corrected statistical uncertainties of the η-intercalibration factors, which I previously derived, to this measurement.

I have worked on the measurement of the Lund Jet Plane (LJP), a powerful jet substructure observable that maps the phase space of emissions inside a jet in terms of their momentum fraction and opening angle. The LJP is sensitive to α_s and its running with energy, making it a valuable tool for precision QCD studies. In this analysis, I was responsible for the unfolding of the data and the propagation of systematic uncertainties. I also evaluated the unfolding bias and used this information to optimise the parameters of the unfolding procedure.

Finally, I conducted jet studies in the context of the Future Circular Collider in its electron–positron collision phase (FCC-ee). The clean and well-controlled experimental environment of the FCC-ee provides an ideal setting for high-precision analyses of jet production and evolution. In this framework, I studied the projected sensitivity to α_s using jet-based observables: $R_{3/2}$ the ratio of cross sections for events with at least three jets to those with at least two jets, and the LJP.

Résumé :

Le modèle standard (MS) de la physique des particules décrit les particules fondamentales qui composent la matière, ainsi que trois des quatre forces fondamentales. Jusqu'à présent, il a connu un succès remarquable, ses prédictions étant confirmées par beaucoup de résultats expérimentaux, mais il ne fournit pas d'explications pour plusieurs questions ouvertes en physique fondamentale. Cette thèse explore une stratégie pour tester le MS: comparer ses prédictions théoriques avec des mesures de haute précision. Elle se concentre spécifiquement sur les jets, une composante essentielle de la chromodynamique quantique (QCD) du MS, permettant de sonder ces derniers aux plus petites distances accessibles en laboratoire. Les analyses sont basées sur les données collectées par l'expérience ATLAS au CERN, utilisant des collisions proton-proton à une énergie au centre de masse de 13 TeV, enregistrées pendant le Run 2. Les jets sont de gerbes collimatées de particules résultant de l'hadronisation de quarks ou de gluons, qui déposent des traces dans les trajectographes et de l'énergie dans les calorimètres. Pour reconstruire le jet de particule, ces signaux sont regroupés dans ce que l'on appelle les "jets reconstruits". Cependant, les jets reconstruits diffèrent des jets de particules en raison d'effets de détecteurs. Pour garantir des comparaisons fiables entre les prédictions théoriques et les résultats expérimentaux, ces effets doivent être corrigés par un processus d'étalonnage précis et l'unfolding des effets de résolution.

J'ai travaillé sur l'intercalibration en η, une étape de l'étalonnage des jets qui vise à assurer un étalonnage uniforme sur l'ensemble du détecteur. J'ai obtenu l'étalonnage sur les données du Run 2 et Run 3. Pour améliorer la robustesse de la méthode, j'ai étendu le système surcontraint utilisé pour extraire les facteurs d'étalonnage en incorporant des termes supplémentaires et en mettant en place un test de leur compatibilité. En outre, j'ai développé une méthode plus fiable pour estimer les incertitudes statistiques sur les facteurs d'étalonnage. Cette approche permet d'étudier leurs corrélations.

J'ai contribué à la mesure différentielle de la section efficace de production de jets. Une comparaison quantitative précise des spectres mesurés avec les prédictions théoriques per-

met d'extraire la constante de couplage fort, α_s. Mon travail s'est concentré sur l'unfolding des données, une procédure qui corrige statistiquement les effets du détecteur pour obtenir la distribution au niveau des particules. J'ai évalué le biais d'unfolding, introduit par les différences de forme entre les données et la simulation, et j'ai propagé les incertitudes systématiques à travers l'unfolding. En outre, j'ai propagé les incertitudes statistiques corrigées des facteurs d'intercalibration en η, que j'avais précédemment dérivées, à cette mesure.

J'ai travaillé sur la mesure du Lund Jet Plane (LJP), qui décrit la sous-strucure des jets en cartographiant l'espace de phase des émissions à l'intérieur d'un jet en termes de leur fraction de quantité de mouvement et de leur angle d'ouverture. Le LJP est sensible à α_s et à son évolution avec l'énergie. Dans cette analyse, j'étais responsable de l'unfolding des données et de la propagation des incertitudes systématiques. J'ai aussi évalué le biais de l'unfolding et l'ai utilisé pour optimiser les paramètres de la procédure d'unfolding.

Enfin, j'ai mené des études sur les jets dans le contexte du futur collisionneur circulaire dans sa phase de collision électron-positron (FCC-ee). L'environnement expérimental propre et bien contrôlé du FCC-ee offre un cadre idéal pour des analyses de haute précision de la production et de l'évolution des jets. Dans ce cadre, j'ai éstimé la sensibilité à α_s en utilisant des observables basées sur les jets: $R_{3/2}$ le rapport des sections efficaces pour les événements avec au moins trois jets à ceux avec au moins deux jets, et le LJP.

Introduction

The Standard Model (SM) of particle physics describes the fundamental particles that make up matter, along with three of the four fundamental forces (**Chapter 1**). So far, it has been remarkably successful, with its predictions confirmed by a wide range of experimental results. However, it does not provide explanations for several open questions in fundamental physics, such as the nature of dark matter, the matter–antimatter asymmetry in the universe, or the origin of neutrino masses. It is widely believed that the SM is a low-energy approximation of a more fundamental underlying theory.

There are essentially two complementary approaches to experimentally searching for physics beyond the Standard Model (BSM). The first relies on high-precision measurements and their comparison with extremely accurate experimental results to state-of-the-art theoretical predictions. Subtle deviations from SM predictions, caused by the virtual effects of BSM particles in quantum loops, can provide indirect evidence for their existence. The second approach involves directly probing higher energies, where new particles could be produced and observed. Collider physics plays a central role in both strategies. It consists of accelerating particles to extremely high energies and smashing them together. The energy of the collision can convert into mass, potentially producing new particles, while precision detectors surrounding the collision point record the particles resulting from the collision. Analysing the properties and distributions of these particles enable to test predictions of the SM and search for new particles. Colliders can reach the highest energies available in laboratory, and they provide enormous statistic enabling high-precision measurements. In general, hadron colliders are used for the exploratory phase of collider physics, as they allow for probing higher centre-of-mass energies, while lepton colliders are preferred for precision studies due to their cleaner experimental environment.

Historically, these two approaches have been used alternatively, with precision studies often providing hints about where to search during the next exploratory phase. For instance, measurements at LEP (Large Electron-Positron Collider) provided indirect evidence suggesting the mass range of the Higgs boson. This insight guided the LHC (Large Hadron Collider) experiments, which confirmed the Higgs boson's existence years later. These two experiments took place at CERN, the European Organization for Nuclear Research, located on the border between Switzerland and France. CERN is an intergovernmental organisation that provide the particle accelerators and other infrastructure needed for high-energy physics research. It hosts numerous experiments conducted by large interna-

tional collaborations. Currently, CERN operates the LHC, a 27-km circular accelerator that collides protons at centre-of-mass energies ranging from 7 to 13.6 TeV. The LHC hosts four major experiments, including the ATLAS (A Toroidal LHC ApparatuS) experiment, on which most of this thesis is based (**Chapter 2**). Looking to the future, CERN has declared the Future Circular Collider (FCC) as its next flagship project. In its first phase, the FCC will collide electrons and positrons at various center-of-mass energies, each targeting specific physics processes (FCC-ee). This setup will offer an exceptionally clean experimental environment and deliver a vast amount of high-precision data. Then, the same infrastructure will be used to collide hadrons at unprecedented energies of up to 100 TeV (FCC-hh). A portion of this thesis is dedicated to studies related to the FCC-ee program.

More broadly, this thesis focuses on the precision approach to experimental particle physics, using jets as key probes of the Quantum Chromodynamics (QCD) sector of the Standard Model. QCD is the theory describing the strong interaction, which acts on particles that carry colour charge (quarks and gluons) and is mediated by the exchange of gluons. It is notably responsible for the cohesion of atomic nuclei. Jets provide access to the dynamics of quarks and gluons at the smallest distance scales reachable in the laboratory and are among the most prominent observables at the LHC. They are collimated sprays of particles resulting from the hadronisation of a high-energy quark or gluon. They leave tracks in the tracking detectors and energy in the calorimeters. To reconstruct the originating jet, these signals are clustered into what are known as "reconstructed jets", to obtain the properties of the initial quark or gluon. However, reconstructed jets differ from their corresponding particle-level jets, which represent the true hadronic final state, as a result of detector effects such as limited resolution and inefficiencies. To ensure meaningful comparisons between theoretical predictions and experimental results, these effects must be corrected through a precise calibration process and the unfolding of the resolution effects.

On the calibration side, I have worked on the η-intercalibration, a step in the jet calibration procedure that aims to ensure a uniform response across the full detector acceptance (**Chapter 3**). I derived the η-intercalibration for ATLAS data. To improve the robustness of the method, I extended the overconstrained system used to extract the calibration factors by incorporating additional terms and implementing a compatibility test among them. In addition, I developed a more reliable method for estimating the statistical uncertainties on the calibration factors. This approach allows for the study of their correlations.

I have contributed to the inclusive jet production cross section measurement (**Chapter 4**). It is a double-differential measurements that count every single jet above a certain transverse momentum p_{T} threshold in bins of p_{T} and absolute rapidity $|y|$ (related to the polar angle in the detector). A precise quantitative comparison of the measured spectra with theoretical prediction allows the extraction of the coupling constant of the strong interaction, α_{s}. My work focused on the unfolding of the data, a procedure that statistically corrects for detector effects to recover the particle-level distribution. I evaluated the unfolding bias, introduced by differences in shape between data and simulation, and I propagated the systematic uncertainties through the unfolding process. In addition, I propagated the corrected statistical uncertainties and correlations of the η-intercalibration

factors, which I previously derived, to this measurement.

I have also worked on the measurement of the Lund Jet Plane (LJP), a powerful jet substructure observable that maps the phase space of emissions inside a jet in terms of their momentum fraction and opening angle (**Chapter 5**). The density and normalisation LJP are sensitive to α_s and its running with energy, making it a valuable tool for precision QCD studies. In this analysis, I was responsible for the unfolding of the data and the propagation of systematic uncertainties. I also evaluated the unfolding bias and used this information to optimise the parameters of the unfolding procedure.

Finally, I conducted jet studies in the context of the FCC-ee. The clean and well-controlled experimental environment of the FCC-ee provides an ideal setting for high-precision analyses of jet production and evolution (**Chapter 6**). In this framework, I studied the projected sensitivity to α_s using two jet-based observables: $R_{3/2}$ the ratio of cross sections for events with at least three jets to those with at least two jets, and the Lund Jet Plane.

Chapter 1

The Theoretical Framework

Introduction

The Standard Model (SM) of particle physics is the theoretical framework that describes all known elementary particles and the fundamental forces that govern their interactions. It is formulated within the framework of quantum field theory (QFT), which combines the principles of quantum mechanics and special relativity to describe particles as excitations of underlying fields. Developed in the second half of the twentieth century, it has been extraordinarily successful, providing predictions that have been confirmed with remarkable precision by many experiments, including those conducted at the Large Hadron Collider (LHC).

Despite its success, the SM is not complete, and several key phenomena remain unexplained. For instance, it does not account for the observed asymmetry between matter and antimatter in the universe, nor does it incorporate a quantum theory of gravity. It lacks a viable candidate for dark matter, a form of matter that does not interact with light and that has gravitational influence on galaxies and cosmic structures, and no explanation for dark energy, the mysterious component driving the accelerated expansion of the universe. Additionally, the Standard Model assumes neutrinos to be massless, but experimental evidence shows that neutrinos oscillate between different types (flavours), implying that they must have mass.

While the Standard Model remains a keystone of modern physics, it is clear that a more complete theory is needed. A precise study of the SM parameters and any deviations from its theoretical predictions can provide crucial hints of physics beyond the Standard Model (BSM), potentially revealing new fundamental particles or interactions.

1.1 The Standard Model of Particle Physics

1.1.1 Quantum Field Theory

Quantum Field Theory (QFT) emerged from the necessity of a consistent combination of Quantum Mechanics (QM) and Special Relativity (SR), two successful frameworks

developped in the early XX$^{\text{th}}$ century to adress the behavior of matter and energy at small scales and high speeds, respectively. On the one hand, QM introduces wave functions and operators to describe physical systems, explaining phenomena such as superposition, quantisation and the indetermination principle. On the other hand, SR reformulates the laws of physics to ensure they remain covariant under changes of inertial reference frames. It describes space and time as components of a four-dimensional spacetime and uses Lorentz-invariant tensors to represent physical quantities such that they respect this symmetry.

1.1.1.1 Notations and conventions

Throughout this chapter, the following notations are used:

- Covariant notation: a four-vector quantity depending on space-time coordinates is denoted by X_μ ($\mu = 0, 1, 2, 3$, similarly for other Greek indices). By convention, X_0 corresponds to the time component and $\mathbf{X} = (X_1, X_2, X_3)$ to the space components.
- Derivatives: the partial derivative with respect to the μ^{th} coordinate $\partial/\partial x_\mu$ is abbreviated as ∂_μ.
- Summation conventions: repeated Greek indices imply Einstein summation, i.e., an implicit sum over spacetime indices: $X_\mu Y^\mu = X^\mu Y_\mu = X_0 \cdot Y_0 - \mathbf{X} \cdot \mathbf{Y}$. For repeated Latin indices, the convention is a simple summation on space coordinates: $X_a Y^a = X^a Y_a = \mathbf{X} \cdot \mathbf{Y}$.

In addition, the natural units are used: $\hbar = c = 1$.

One of the most commonly used four-vectors is the four-momentum $P^\mu = (E, \vec{p})$, where E is the energy and p the momentum of the object described. Its mass is the norm of the four-momentum: $m^2 = P^2 = P_\mu P^\mu = E^2 - \vec{p}^2$.

1.1.1.2 Lagrangian formalism and symmetries

Classical mechanics can be formulated as a least-action principle, which states that a physical system evolves between two configurations in such a way that the action is stationary (typically minimised). The action S is the integral of the Lagrangian L over time :

$$S = \int dt\, L = \int d^4x\, \mathcal{L}(x),$$

where $\mathcal{L}(x)$ is the Lagrangian density, representing the difference between the kinetic and potential energy densities of a system.

This Lagrangian approach can be adapted to QFT. Contrary to classical mechanics, quantum physics cannot determine simultaneously and with precision the position and momentum of a particle at the same time. As a result, the content of a system cannot be described by a point-like particle following definite trajectories. Instead, their physical state is represented by complex fields that depend on spacetime coordinates, commonly denoted $\phi(x_\mu)$. A field is a superposition of all possible momentum modes, it can describe all possible configurations of one or more particles in a given momentum state. In this approach, particles are simply field excitations. The squared amplitude of this quantity,

$|\phi(x_\mu)|^2$, is interpreted as a presence probability density at a given point in space and time. The Lagrangian density in field theory is then expressed in terms of $\phi(x_\mu)$, its space-time derivatives $\partial_\mu\phi(x_\mu)$ and the coordinates x_μ. From the least-action principle can be derived the Euler-Lagrange equation, which gives the equations of motion:

$$\frac{\partial\mathcal{L}}{\partial\phi} - \partial_\mu\left(\frac{\partial\mathcal{L}}{\partial(\partial_\mu\phi)}\right) = 0.$$

1.1.1.3 Gauge Theory

A Lagrangian may be invariant under some transformation $\phi \to \phi + \delta\phi$. Such an invariance is called a symmetry of the Lagrangian. According to Noether's theorem [1], every continuous symmetry of the Lagrangian corresponds to an associated conserved quantity whose value remains constant over time when the equations of motion are satisfied. For example, in classical mechanics, invariance under spatial translations leads to the conservation of momentum, while invariance under time translations is responsible for the conservation of energy. Symmetries that depend on the space-time coordinates are called gauge or local symmetries. Otherwise, the symmetries are referred to as global symmetries.

To ensure the Lagrangian remains invariant under local symmetry transformations, new fields must be introduced. They are called gauge fields, and they transform in such a way as to cancel the extra terms introduced by the spacetime dependence. In the Lagrangian, gauge fields couple to the matter fields, and these coupling terms can be interpreted as interaction terms. Physically, gauge fields correspond to the mediators of fundamental interactions. To preserve gauge invariance, those gauge fields cannot have a mass.

1.1.2 Standard Model

The Standard Model was built by looking at the observed symmetries of the systems, and writing the most general Lagrangian which follows such symmetries.

1.1.2.1 Symmetries

The Standard Model Lagrangian respects several fundamental principles that are described by mathematical groups, each representing different types of symmetries and interactions. The Standard Model is invariant under the Poincaré group, a global space-time symmetry that includes translations, rotations and Lorentz boosts (change of reference frame moving at constant velocity). This symmetry implies the conservation of energy, momentum, angular momentum and spin.

Local symmetries in the Standard Model have been motivated by experimental evidence for the conservation of certain quantum numbers. The most important gauge groups in the Standard Model are the following:

- The $SU(3)_\mathrm{C}$, which is the symmetry group of the strong interaction (Quantum Chromodynamics, QCD). It describes how quarks interact via the gluons, the carriers of the strong force. $SU(3)$ stands for Special Unitary group of degree 3, meaning it consists of 3x3 unitary matrices with determinant 1.

- The $SU(2)_{\mathrm{L}}$, which is the symmetry group of the weak interaction. It governs how particles interact via the W and Z bosons, the mediators of the weak force. $SU(2)$ refers to a group of 2x2 matrices and describes how doublets (pairs of particles) like leptons and quarks transform.
- The $U(1)_{\mathrm{Y}}$, which is the symmetry group of the hypercharge. It represents the electromagnetic and weak charge properties of particles. $U(1)$ consists of complex numbers of unit magnitude, representing phase rotations. It combines with $SU(2)_{\mathrm{L}}$ to form the electroweak interaction.

The Standard Model is based on those three independent symmetry groups, each acting on particles associated with different types of charges or quantum numbers. It is written as $SU(3)_{\mathrm{C}} \times SU(2)_{\mathrm{L}}$ or $U(1)_{\mathrm{Y}}$ where "$\times$" denotes the direct product. This means that each group acts independently on its corresponding quantum numbers, and the full symmetry transformation is simply the combination of the individual transformations.

1.1.2.2 Spontaneous symmetry breaking

A symmetry is said to be spontaneously broken when it is a symmetry of the Lagrangian but not of the ground state of the system. A classic example is a pencil balanced vertically on a table: the system has rotational symmetry around the vertical axis, but this upright position is unstable. When the pencil falls due to a small disturbance, it picks a specific direction, breaking the symmetry. According to Goldstone's theorem [2], a massless state (Goldstone Boson) appears for every broken symmetry.

In the Standard Model, the $SU(2)_{\mathrm{L}} \times U(1)_{\mathrm{Y}}$ symmetry is spontaneously broken by the Higgs field, which is an $SU(2)$ doublet. The Goldstone boson, hence created, is "eaten" by the gauge bosons (it becomes the longitudinal mode of the gauge boson associated with the broken generator). The resulting boson has an extra degree of freedom and thus is allowed to have mass. This mechanism has been postulated by Guralnik, Hagen and Kibble [3], Brout and Englert [4], and Higgs [5], and it is known as the Brout-Englert-Higgs mechanism.

The unbroken symmetry left after electroweak symmetry breaking is $U(1)_{\mathrm{Q}}$, with the electric charge being the associated conserved quantity, and the remaining gauge symmetry of the Standard Model is $SU(3)_{\mathrm{C}} \times U(1)_{\mathrm{Q}}$.

1.1.2.3 Field content of the Standard Model

To express the Standard Model Lagrangian, the field content of the SM has to be defined. It starts with the observed fermions (particles of spin 1/2), referring to the usual six quarks (u, d, s, c, t, b), three charged leptons (e, μ, τ) and three neutral leptons $(\nu_e, \nu_\mu, \nu_\tau)$. Each fermion field transforms under one or more of the Standard Model gauge groups: $SU(3)_{\mathrm{C}} \times SU(2)_{\mathrm{L}}$ or $U(1)_{\mathrm{Y}}$. Every fermion also has an associated antiparticle, with the same properties but opposite quantum numbers and (potentially) a different mass.

The gauge symmetries give rise to bosons (particles of integer spin), which mediate the fundamental interactions. From the $SU(3)_{\mathrm{C}}$ symmetry arise eight gauge bosons (the gluons) which carry colour charge. The electroweak sector, described by the $SU(2)_{\mathrm{L}} \times$

$U(1)_Y$ symmetry, generates the B, W^1, W^2 and W^3 bosons. The Higgs boson is a complex scalar $SU(2)_L$ doublet that plays a central role in the Higgs mechanism. This mechanism produces three Goldstone bosons, which are absorbed by the weak bosons to form the massive $Z^0, W^+$and W^-. The photon γ emerges from the mixing of the B and W^3 fields. The remaining physical scalar degree of freedom is the Higgs boson h, a spin-0 particle with no electric or colour charge.

Each gauge boson couples to the fermions carrying the corresponding charge:

- Gluons couple to all quarks, and to themselves,
- The photon couples to all electrically charged fermions,
- The $Z^0, W^+$and W^- bosons couple to all fermions.

The properties of all particles of the Standard Model are summarised in Figure 1.1.

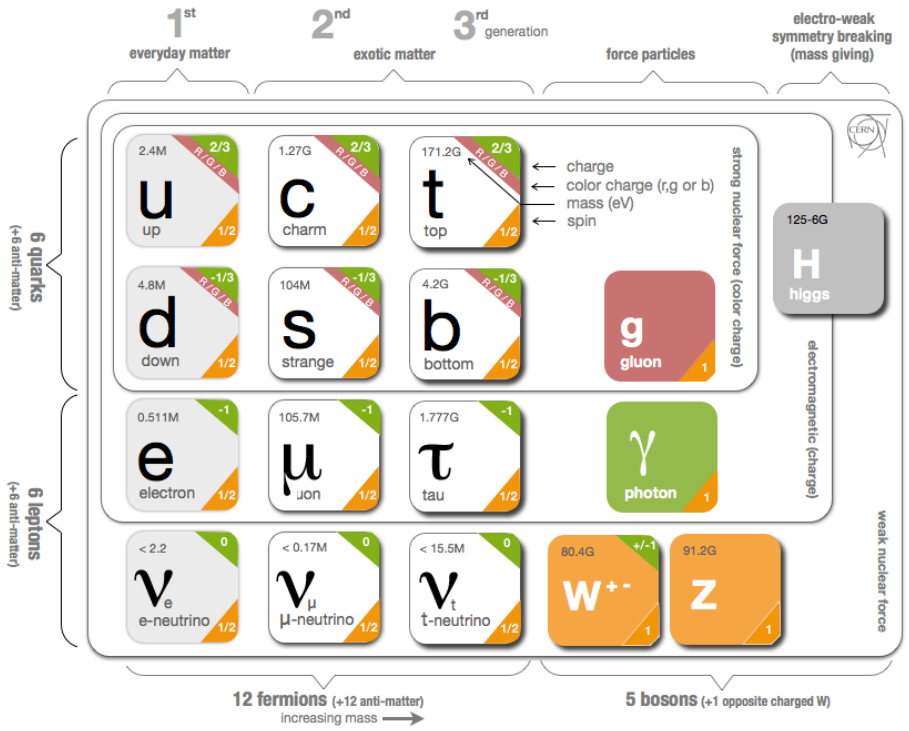


Figure 1.1: Summary of the particles composing the Standard Model and their properties. Reproduced from [6].

1.1.2.4 Mathematical definition of the Standard Model

The Standard Model Lagrangian is split into multiple parts:

$$\mathcal{L}_{\text{SM}} = \mathcal{L}_{\text{gauge}} + \mathcal{L}_{\text{QCD}} + \mathcal{L}_{\text{EW}} + \mathcal{L}_{\text{Higgs}} + \mathcal{L}_{\text{Yukawa}}.$$

- $\mathcal{L}_{\text{gauge}}$ contains the free propagating term and self-interaction of the gauge fields,
- $\mathcal{L}_{\text{QCD}}$ and $\mathcal{L}_{\text{EW}}$ describe the interactions between the fermions and the gauge bosons of the strong and electroweak sectors,

- $\mathcal{L}_{\text{Higgs}}$ includes the kinetic term, mass term, and self-interactions of the Higgs field, as well as its couplings to the electroweak gauge bosons (from which the $Z^0, W^+$and W^- gain their mass)
- $\mathcal{L}_{\text{Yukawa}}$ encodes the Yukawa interactions between the Higgs field and the fermions, which generate the fermion masses after electroweak symmetry breaking.

The Standard Model, as it stands, includes 19 free parameters, meaning they are not predicted by the above theoretical formulation and can only be determined by experimental measurements. The rest of this chapter focuses on the QCD sector of the Standard Model, which underpins the processes studied throughout this thesis.

1.2 Quantum Chromodynamics

Quantum Chromodynamics is a non-Abelian local gauge theory, introduced by Yang and Mills in 1954 [7]. It describes the strong interaction that governs particles that carry colour charge, that is, fields that transform under $SU(3)_{\text{C}}$. It includes six matter fermions (the quarks) and eight gauge fields (the gluons), which mediate the strong force and themselves carry colour charge. Most of this section's content is inspired by the QCD lectures of Gavin Salam (2024 European School of High Energy Physics), and Matteo Cacciari (lectures for STEP'UP doctoral school, 2023), as well as from Reference [8].

1.2.1 The QCD Lagrangian

1.2.1.1 Expression

In the following, the gluons are described by the gauge field $\mathcal{A}^C_\mu$. The $SU(3)_{\text{C}}$ gauge symmetry has eight generators t^A corresponding to the eight gluons $\mathcal{A}^C_\mu$. These generators obey the relation: $[t^A, t^B] = if_{ABC}t^C$ where the f_{ABC} are the structure constant of $SU(3)$, fully antisymmetric in their indices. The free fermion fields $\psi_a = (\psi_1, \psi_2, \psi_3)$, represent the three colour components of a quark field, and g_s is the QCD coupling constant. The field tensor is given by:

$$F^A_{\mu\nu} = \partial_\mu \mathcal{A}^A_\nu - \partial_\nu \mathcal{A}^A_\mu - g_s f_{ABC} \mathcal{A}^B_\mu \mathcal{A}^C_\nu.$$

The first term is Abelian, it resembles the field tensor of the electroweak interaction. The second term is the non-Abelian part of the interaction. It causes the Lagrangian to contain terms involving three or four gluon fields, meaning that gluons can interact directly with one another. The self-interaction of gluons is a specificity of QCD and has direct consequences on the properties of the strong force.

The QCD part of the Standard Model Lagrangian is expressed in the following way:

$$\mathcal{L}_{\text{QCD}} = -\frac{1}{4} F^A_{\mu\nu} F^{A\,\mu\nu} + \bar{\psi}_a \left(i\gamma^\mu \partial_\mu \delta_{ab} - m\right) \psi_b - g_s \bar{\psi}_a \gamma^\mu t^C_{ab} \mathcal{A}^C_\mu \psi_b.$$

The first term describes the kinetic energy of the gluon fields and their self-interactions. The second term represents the free propagation and mass of the quark fields. The third term is the interaction term, which describes how quarks interact by exchanging gluons.

1.2.1.2 Matrix elements

In QFT, physical observables such as scattering cross-sections and decay rates are computed from the S-matrix (scattering matrix), which encodes the probabilities to go from an initial to a final state. In particular, the cross-section quantifies the probability of a scattering process occurring between incoming particles. It has dimensions of an area. These probabilities are obtained by applying Feynman rules, derived from the interaction terms in the QCD Lagrangian, to the corresponding Feynman diagrams. The Feynman rules translate each vertex and propagator into mathematical expressions involving momenta and quantum numbers. The resulting expressions are called matrix elements. The probability for a process is determined by squaring the sum of all contributing matrix elements and integrating the result over the final-state phase space.

Three interaction vertices can be deduced from the QCD Lagrangian. The Feynman rules are represented in Figure 1.2. The matrix elements associated with the vertices represented in Figure 1.2b and Figure 1.2c are proportional to g_s, while it is proportional to g_s^2 in the case of the vertex of Figure 1.2d.

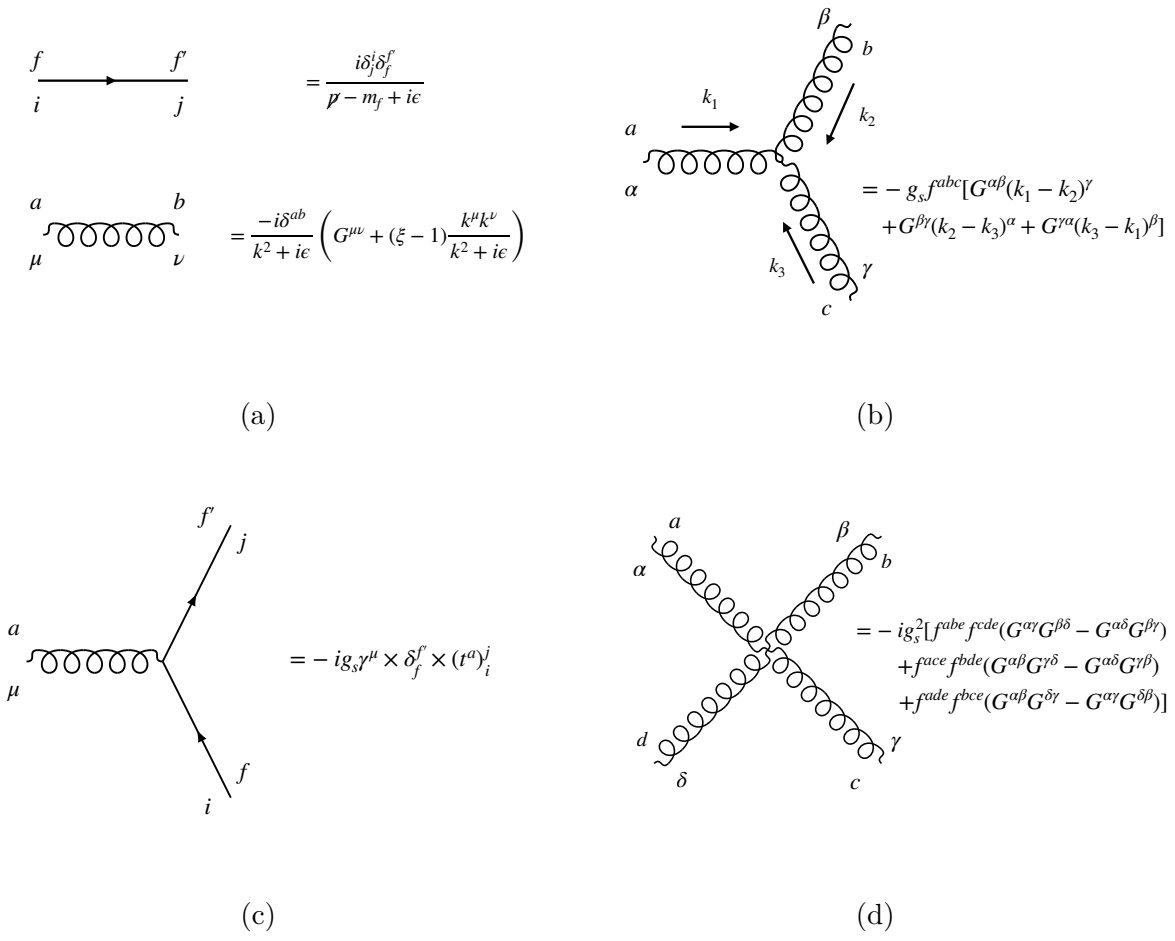


Figure 1.2: Feynman rules for QCD. (a) represent the propagator for the quark and the gluon, while (b), (c) and (d) are the vertices of QCD.

1.2.1.3 Renormalization

When calculating matrix elements for diagrams that include loops, one must integrate over all possible values of the virtual momenta circulating in the loop. This leads to ultraviolet (UV) divergences, caused by the arbitrarily large momenta running in the loop. To handle these infinities, a procedure called renormalisation is introduced. It consists of redefining the parameters of the theory by absorbing the divergences into them at a chosen energy scale, called the renormalisation scale μ_{R}. This allows finite predictions to be extracted from the theory. This scale is unphysical and can be arbitrarily chosen, but physical observables must remain independent of it.

1.2.1.4 The strong coupling constant

The strength of the strong interaction is characterised by $\alpha_{\mathrm{s}} = g_s^2/4\pi$, the strong coupling constant. Despite its name, α_{s} is not truly constant: it varies with the renormalisation scale μ_{R}. In physical predictions, the choice of μ_{R} is arbitrary, but it is usually chosen to be of the order of the characteristic energy scale Q of the process. The energy dependence of α_{s} is governed by the renormalisation group equation (RGE):

$$\mu_{\mathrm{R}}^2 \frac{\partial \alpha_s}{\partial \mu_{\mathrm{R}}^2} = \beta(\alpha_s), \qquad \beta(\alpha_s) = -\alpha_s^2 \left(b_0 + b_1 \alpha_s + b_2 \alpha_s^2 + \dots \right).$$

At leading order, the first coefficient in the perturbative expansion is $b_0 = \frac{11C_A - 2n_f}{12\pi}$, with n_f the number of quark flavours and C_A the number of colours. Because $b_0 > 0$, the first term of the perturbative expansion of $\beta(\alpha_s)$ is negative. This feature is a consequence of gluon self-interactions. It leads to a key property of QCD known as asymptotic freedom [9, 10]:

- At high energy scales, the coupling becomes small: quarks and gluons behave as if they were free.
- At low energy scales, the coupling becomes strong: quarks and gluons are strongly coupled and confined into bound states, the hadrons.

The energy dependence of α_{s}, known as the running of the coupling, is displayed in Figure 1.3, along with various measurements showing relatively good agreement.

Solving the RGE at leading order in the expansion of β yields:

$$\alpha_{\mathrm{s}}(\mu_{\mathrm{R}}^2) = \frac{1}{b_0 \ln(\mu_{\mathrm{R}}^2/\Lambda^2)},$$

where $\Lambda \simeq 0.2$ GeV is a constant of integration, known as Λ_{QCD}. It corresponds to the fundamental scale of QCD at which the perturbative coupling diverges. It is a free parameter which provides a boundary condition for the evolution equation of the coupling. Physically, Λ_{QCD} sets the scale for the hadron masses.

1.2.2 From partons to hadrons

Free quarks have never been observed, which is understood as a result of a long-distance confining behaviour of the strong force. Except for the top quark, which decays too

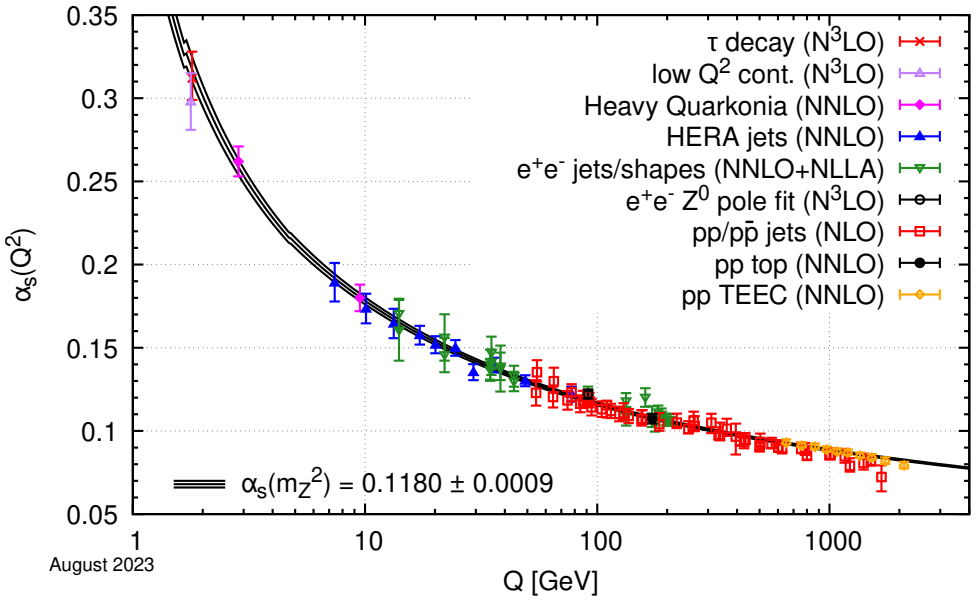


Figure 1.3: Summary of determinations of α_s as a function of the energy scale Q, compared to the running of the coupling predicted by the theory. Taken from [11].

rapidly, the quarks hadronise, i.e. they become part of a colour-neutral structure, the hadrons. A familiar example is the proton, which consists of two up quarks and one down quark. These quarks, whose quantum numbers combine to reproduce those of the proton, are referred to as valence quarks. However, the full structure of a hadron is more complex: in addition to the valence quarks, it contains gluons and quark–antiquark pairs, which are continuously created and annihilated via quantum fluctuations. Feynman introduced the parton model to describe this internal structure, in which hadrons are made of point-like constituents called partons. The consequence of asymptotic freedom is that in perturbative QCD calculations at high energies (assuring α_s sufficiently small), partons can be treated as free particles.

1.2.2.1 Perturbative QCD

Perturbation theory At high-energy scales ($Q >> \Lambda_{\mathrm{QCD}}$), typically in the high-energy proton-proton collisions of the LHC, α_s is small. This allows QCD processes that involve high-momentum transfer (usually referred to as "hard processes") to be described using perturbation theory, where observables are expressed as a power series in α_s:

$$\sigma = \sigma_0 + \alpha_s\sigma_1 + \alpha_s^2\sigma_2 + \alpha_s^3\sigma_3 + ...$$

In the language of Feynman diagrams, each power of α_s corresponds to the emission of an additional parton in the final state, or the apparition of a loop involving a parton. The leading order (LO) contribution (σ_0, or born-level observable) comes from the simplest Feynman diagrams with the fewest QCD vertices. The next-to-leading order (NLO) includes one additional power of α_s, typically from diagrams with one loop or additional radiation, and so on for the next-to-next-to-leading order (NNLO) and beyond.

The contributions from parton emissions are labelled "real" while the contributions from

loops are labelled "virtual". An example of the building blocks for a simple process is shown in Figure 1.4.

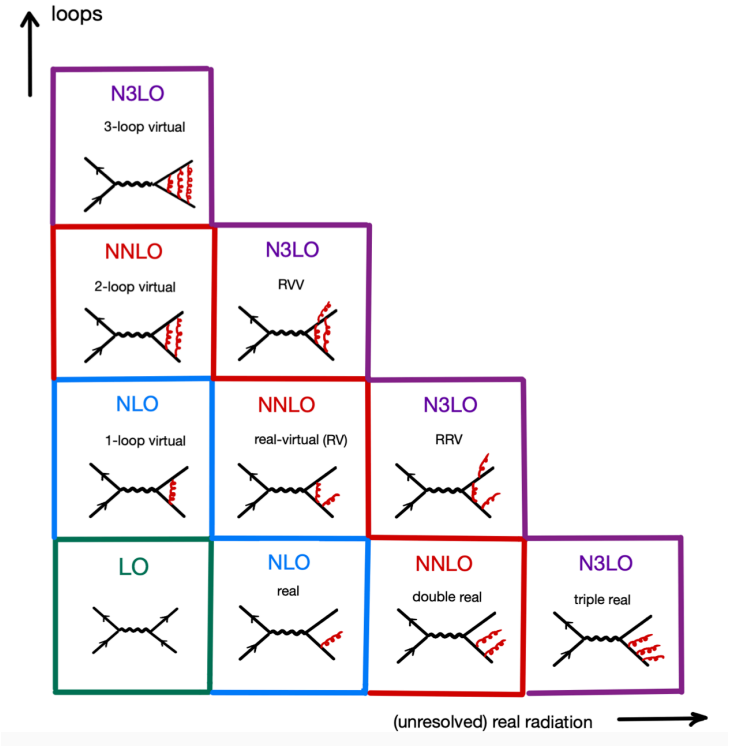


Figure 1.4: Building blocks of an N^{x}LO calculation for processes where the leading order is at tree level. Taken from [8].

Soft & collinear gluon emission When a soft gluon is emitted from a scattered parton (i.e., a gluon with a small momentum compared to that of the emitting parton), the matrix element for the process factorises into a hard part describing the underlying hard scattering and a soft factor dS representing the soft-gluon emission. With θ the angle between the gluon and the emitting parton, ϕ the azimuthal angle and E the energy of the emitted gluon, the soft factor can be expressed as:

$$dS = \frac{2\alpha_s C_F}{\pi} \frac{dE}{E} \frac{d\theta}{\sin\theta} \frac{d\phi}{2\pi}.$$

This expression diverges when integrating for $E \to 0$ (infrared divergence), and for $\theta \to 0$ and $\theta \to \pi$ (collinear divergence). Virtual diagrams also exhibit analogous infrared and collinear (IRC) singularities. However, these divergences cancel when real and virtual corrections are combined in properly defined observables [12–14]. This can be understood in the following way: soft gluons are emitted on a long timescale ($1/E\theta$) relative to the collision ($1/Q$), and therefore cannot influence the cross-section of the process studied. Such observables are said to be infrared and collinear safe (IRC safe); they are insensitive to soft emission or collinear splitting.

Large logarithms and resummation Fixed-order perturbation theory is reliable while $\alpha_s(Q)$ remains small (with Q the energy of the quark emitting the gluon), but

it loses its predictive power near the non-perturbative scale Λ_{QCD}. A cut-off on the transverse momenta of the gluon emitted, below the non-perturbative threshold Λ_{QCD}, has to be added. Defining $Q_0 \geq \Lambda_{\mathrm{QCD}}$, Q is required to be above Q_0, and the angle of emission θ is required to be above Q_0/E. Integrating dS between Q_0 and Q gives the probability for a quark of energy Q to emit a gluon:

$$P_g = \frac{2\alpha_{\mathrm{s}} C_F}{\pi} \int_{Q_0}^{Q} \frac{dE}{E} \int_{Q_0/E}^{1} \frac{d\theta}{\theta} \int_{0}^{2\pi} \frac{d\phi}{2\pi} = \frac{\alpha_{\mathrm{s}} C_F}{\pi} \ln^2(Q/Q_0) + \mathcal{O}(\alpha_{\mathrm{s}} \ln^2(Q/Q_0)).$$

Because each additional soft or collinear emission can add one logarithm $L = \ln(Q/Q_0)$, every power of α_{s} brings the possibility of up to two logarithms, yielding the perturbative series:

$$\sigma(Q) = \sigma_0 + \alpha_{\mathrm{s}}(\sigma_{12} L^2 + \sigma_{11} L + ...) + \alpha_{\mathrm{s}}^2(\sigma_{24} L^4 + \sigma_{23} L^3 + ...) + ...,$$

with $L^n = \ln^n(Q/Q_0)$. Unfortunately, when the cut-off is small, the logarithm is large and compensates for the smallness of α_{s}. All-order resummation is a reorganisation of the above perturbative series to sum those double logarithms to infinity. For many observables, the resummed expression exponentiates, yielding the Sudakov form factor:

$$\Delta(Q, Q_0) = \exp\left[-\frac{\alpha_s\, C_F}{\pi}\, \ln^2\left(\frac{Q}{Q_0}\right) \right].$$

In other words, exponentiation reflects the Poisson probability of multiple independent emissions: the Sudakov form factor represents the probability that no emission above the cut Q_0 occurs. Because probabilities are bounded, the large-log growth is automatically controlled. The fixed-order equation and the resumed equation are complementary, and ultimately, the most accurate theoretical prediction for the observable is achieved by combining the two approaches.

The resummation procedure is a powerful tool for predicting one observable with high accuracy; however, when calculations must extend to a wide range of kinematic variables, a more flexible approach is needed.

Parton showers A highly energetic parton initiates a cascade of successive branchings: it can emit a gluon, and that gluon may in turn radiate another gluon or split into a quark-antiquark pair, and every new parton produced can continue to branch. This is known as a parton shower, and it can be viewed as higher-order corrections to the hard process. At each step, the probability of a further splitting is governed by the Sudakov form factor.

The parton shower algorithms rely on the simulation of the emission of an arbitrary number of particles. Namely, the p_{T} distribution of the first emission can be deduced from the probability of not emitting gluons above a certain p_{T}, using the Sudakov form factor. A random number generator can be used to sample that p_{T} distribution, and the process can be repeated recursively until the p_{T} of the emission falls below the non-perturbative cut-off Q_0. The parton shower algorithms can be thought of as a numerical implementation of a renormalisation-group equation that connects two scales, Q and Q_0, at which hadronisation takes over.

There are two main types of parton shower algorithms:

- Angular Ordered: Emissions are ordered by decreasing emission angle relative to the initial parton's direction. This approach is motivated by the pattern of colour coherence in QCD: soft gluon emissions from a system of partons tend to interfere destructively outside a certain angular region.
- Dipole: Each quark is connected to one colour partner, forming a dipole which splits into two dipoles by emitting a gluon. The masses and momenta of the produced dipoles are determined by the kinematics of the gluon emission. The splitting procedure is performed until the non-perturbative scale is reached.

1.2.2.2 Non-perturbative QCD

While perturbative QCD provides a framework to describe the short-distance physics in the hard interaction, it is not sufficient to fully describe the complex structure of events observed in detectors. Phenomenological models are required to account for the long-distance, non-perturbative aspects of QCD, allowing for a more realistic description of what is recorded by the detectors.

Factorisation & Parton Distribution Function The QCD factorisation theorem states that, for a given process, one can separate physical effects happening at different scales. This key ingredient of QCD has previously been encountered in Section 1.2.2.1 in the context of soft-collinear gluon emission. A particularly important version of this theorem, known as collinear factorisation, allows the separation of the non-perturbative dynamics confining partons within hadrons from the perturbative short-distance interactions responsible for the hard scattering. This separation is valid up to a correction of order $(\Lambda_{\mathrm{QCD}}/Q)^p$, where p is process-dependent, usually larger than one. The factorisation requires the choice of a factorisation scale, μ_{F}, which defines the boundary between perturbative and non-perturbative physics.

In the context of hadron colliders, the factorised form of an observable can be written as a convolution of the hard matrix element and Parton Distribution Functions (PDF). A PDF $f_i(x, \mu_{\mathrm{F}}^2)$ can be interpreted as the probability density for finding a parton of type i that carries a momentum fraction x of the parent momentum of a hadron. The evolution of the PDFs with μ_{F} is governed by the Dokshitzer-Gribov-Lipatov-Altarelli-Parisi (DGLAP) equations [15–17].

The PDFs cannot be calculated using perturbative QCD; instead, they are extracted from global fits to experimental data and evolved to different energy scales using the DGLAP equations. In addition, they are universal; they can be calculated using a specific process and used to predict another. Experimentally determined PDFs are available from various groups worldwide, using data from many experiments. PDF sets are available at different orders in perturbation theory, based on the order of the predictions of the cross sections employed in the fits. An example of PDF for two values of the factorisation scale is displayed in Figure 1.5.

Hadronisation As the shower evolves down to energy scales near the QCD infrared cut-off Λ_{QCD}, the strong coupling α_{s} grows large. Once $\alpha_{\mathrm{s}} \sim 1$, perturbation theory breaks down and the system enters the non-perturbative regime. In this phase, colour

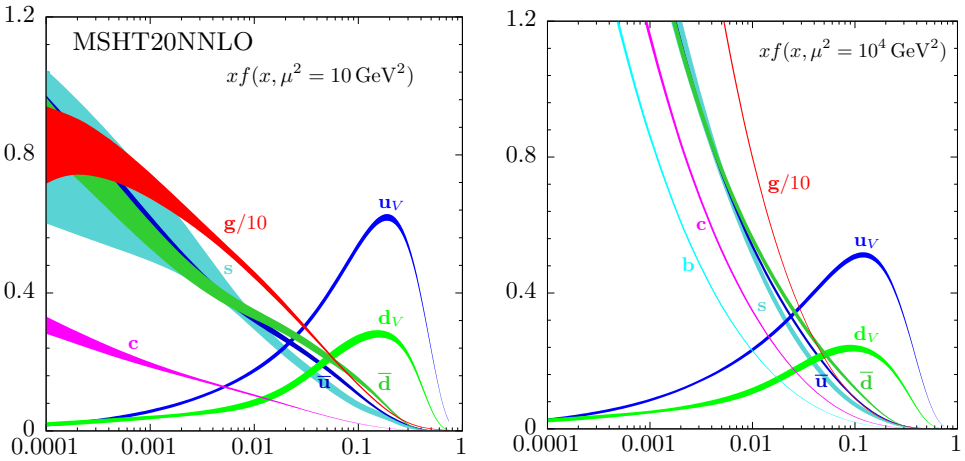


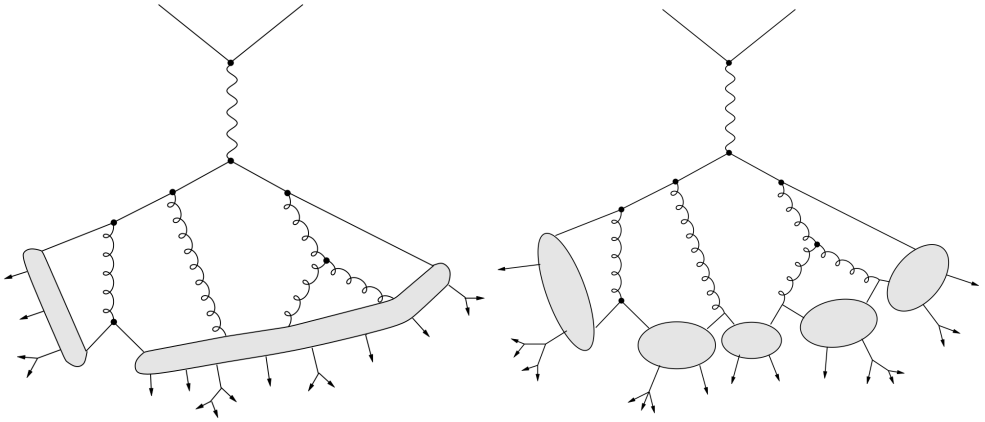
Figure 1.5: $xf_i(x, \mu_F^2)$ as a function of x for $\mu_F^2 = 10$ GeV^2 on the left, and $\mu_F^2 = 10^4$ GeV^2 on the right. Taken from [11].

is confined and the partons combine into colour-neutral hadrons. Hadronisation is the process which leads to the formation of the final-state hadrons from quarks and gluons. To this day, only QCD-inspired phenomenological models can describe this transition. These models are tuned to experimental data and attached to the perturbative parton shower, providing a realistic description of the final-state hadron spectrum. There are two main fragmentation models: Lund strings and clusters.

The Lund string fragmentation model [18] involves stretching a colour "string" across quarks and gluons, and breaking it up into hadrons. It relies on the observation through lattice QCD (non-perturbative QCD calculated on a grid of points) that the potential energy of two colour sources grows linearly with their separation due to gluon self-interaction. In this picture, the colour field is analogous to the potential energy of a string of tension $\kappa \sim 1$ GeV/fm. After a high-energy collision, the quark and antiquark (and any radiated gluons, which appear as a transverse "kinks" in the otherwise one-dimensional string) fly apart, stretching the string and therefore increasing the potential energy. Since gluons carry two colour charges, they are part of two strings and lose energy twice as fast as quarks. A $q\bar{q}$ pair is produced when this energy reaches approximately a hadron mass. This string-breaking process repeats until the available energy is exhausted, converting the kinetic energy of the partons into a chain of quark–antiquark pairs that subsequently combine into colour-neutral hadrons. A schematic representation of the string model is shown in Figure 1.6a.

The cluster model [19] breaks each gluon into a $q\bar{q}$ pair and then groups quarks and antiquarks into colourless "clusters", which then give the hadrons. It exploits the property of preconfinement: regardless of the hard process or shower history, a perturbative QCD cascade evolved down to some scale Q_0 naturally organises its partons into colourless clusters whose invariant-mass spectrum is universal and falls rapidly at large masses. In this scheme, only the number of clusters depends on the scale of the hard subprocess. In most implementations, the remaining gluons are split "by hand" into a $q\bar{q}$ pair, producing a new set of clusters. The clusters are identified as "proto-hadrons" that decay into the final-state hadrons. The many low-mass clusters are turned into hadrons through simple

two-body decays governed by mass and available phase space, while the few high-mass clusters undergo a string-like sequence of splittings until their daughters are light enough to hadronise directly, yielding the final-state hadrons seen in the detector. A schematic representation of the cluster model is shown in Figure 1.6b.



(a) (b)

Figure 1.6: Schematic representation of the string (a) and cluster (b) hadronisation models. Taken from [20].

Underlying event & Multi-Parton Interactions In hadron colliders, multiple low-energy collisions can occur alongside the primary hard scattering. These may arise either from secondary interactions within the same proton–proton collision or from interactions involving other protons in the colliding bunches. These secondary interactions, which do not directly involve the partons participating in the hard subprocess, are collectively referred to as the underlying event (UE). The UE is a significant source of background for most QCD analyses and must be properly accounted for in precision measurements.

The additional scatterings of partons in the incoming protons can be modelled as additional perturbative gluon exchanges, called Multi-parton interactions (MPI). These soft interactions lead to the production of more low-energy hadrons in the final state. As such, MPIs are essential for accurately describing QCD event structure in hadron collisions.

Hadron decay Most hadrons are unstable and decay before reaching the detector. Their decay products (other hadrons and leptons) are the particles that are ultimately detected. The lifetime of hadrons and the probability of each decay mode (branching fraction) are governed by quantum field theory. However, not all particles, excited states, or decay modes have been experimentally measured, so models and assumptions are used to describe the full decay chain.

1.2.2.3 The whole picture

QCD processes at hadron colliders involve a wide range of energy scales, which makes their theoretical description particularly challenging. The hard scattering typically occurs at the TeV scale, where the strong coupling constant is small, and the colliding protons themselves are bound states with masses of order 1 GeV, where the coupling

becomes large. As a result, QCD events at colliders combine both perturbative and non-perturbative physics, requiring factorisation approaches and phenomenological models to describe processes from parton-level interactions to the observable hadronic final states. Figure 1.7 illustrates the multitude of processes involved in a QCD event at a hadron collider.

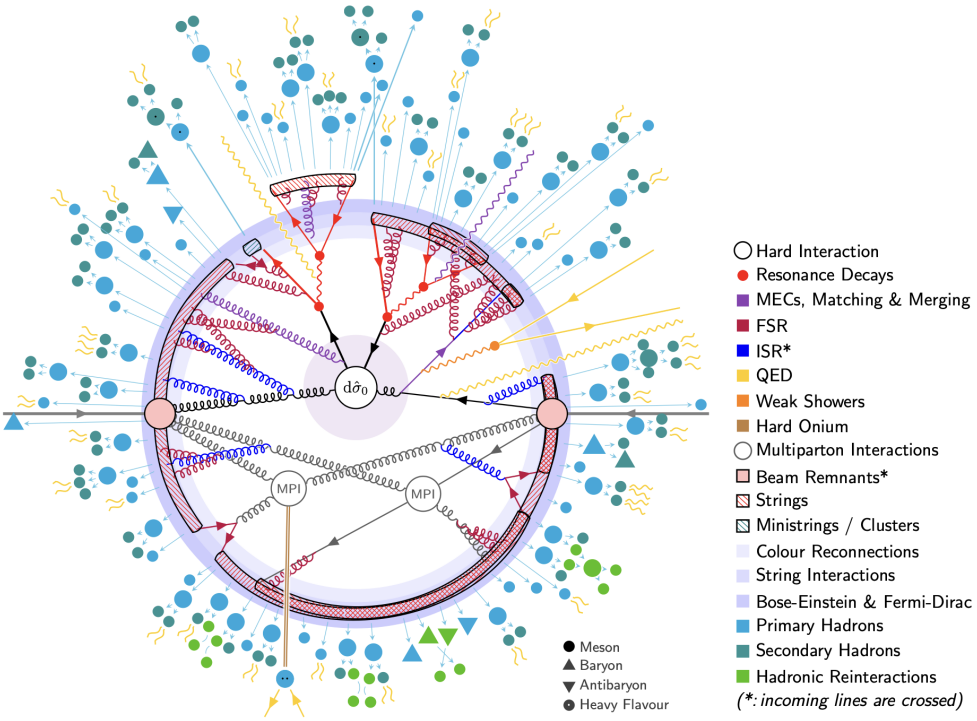


Figure 1.7: Schematic of the structure of a $pp \to t\bar{t}$ event, as modelled by Pythia. The further from the centre, the softer and more non-perturbative the processes is. Taken from [21].

In contrast, QCD studies at electron–positron colliders benefit from a much cleaner environment. The initial state is point-like and precisely known, and there is no UE, MPI or PDF to account for. This allows for a more direct connection between theory and measurements, particularly in jet studies.

1.2.3 Monte Carlo Simulations

Simulated data is a key ingredient of experimental high-energy physics: it makes the link between theory and experiment, and is used at every stage of an analysis. During detector design, simulations help ensure that key signals can be detected. Once the experiment is running, they are essential for optimising event selection. In data analysis, they are used to model detector effects and backgrounds. In precision measurements, simulations provide the theoretical expectations of final state observables, which are compared to data in order to extract parameters of the theory.

Simulations rely on Monte Carlo (MC) methods, which use random numbers to sample possible outcomes according to the quantum mechanical probabilities predicted by

the theory. These techniques are employed both for simulating collision events and for performing numerical integrations in theoretical calculations of final state observables.

1.2.3.1 Event Generation

Event generation using the Monte Carlo method consists of three sequential steps, performed by different programs. Within each step, a set of rules is defined to evolve the event from the previous step to the next.

The first step consists of computing matrix elements from the input Lagrangian, whether it be the SM Lagrangian or an extended one that includes physics beyond the Standard Model. The matrix elements are then combined with PDFs and sampled in phase space. Events are generated so that regions of phase space with higher cross-sections are sampled more frequently, or are assigned weights that reflect their probability of occurrence.

Then, in the second step, another program takes over to describe the showering associated with the process, the presence of additional interactions in the same hadron–hadron collision (MPI), the hadronisation and hadron decay. This step involves non-perturbative QCD processes, which are implemented in Monte Carlo event generators through phenomenological models tuned to experimental data.

Multiple General-purpose Monte Carlo generators (GPMC) are used by the high-energy physics community.

- Pythia 8 [21–23] calculates matrix elements at LO, uses a dipole-style showering model, and uses the Lund string model for hadronisation.
- Sherpa [24] generates LO matrix elements, has a dipole showering, and has two possible hadronisation models: a cluster one and a Lund string one.
- Herwig 7 [25–27] generates events either at LO or NLO, has two possible shower models, either angular-ordered or dipole, and uses a cluster hadronisation model.

Some generators, such as Powheg [28–30] or Madgraph [31], only deal with matrix element. They can be interfaced with GPMC, like Pythia 8 or Herwig 7. The diversity of GPMC allows for testing various models for the physics predictions and identifying the one that best describes the data. In analyses, several GPCMs are used to evaluate modelling uncertainties by comparing their predictions.

1.2.3.2 Parton-level prediction

Monte Carlo techniques are also used to predict parton-level observables, such as cross-sections. They typically generate matrix elements at fixed order, compute the variable of interest and discard them, without simulating the non-perturbative effects. NLOJET++ [32–34] and NNLOJET++ [35, 36], generate NLO and NNLO matrix elements, respectively.

Since non-perturbative effects are not simulated event-by-event, they have to be accounted for in addition to the partonic prediction. This is typically done with multiplicative factors, derived using GPMC with various configurations, as follows. The cross-section is first calculated with the hadronisation and underlying event switched off, using the final

partons in the event record. It is then recalculated using the full event simulation, which includes hadronisation and the underlying event, using stable particles for jet reconstruction. The correction factor is defined as the bin-by-bin ratio of the two.

Given the high precision of NNLO calculations, electroweak corrections become non-negligible and must be accounted for. They are due to higher-order contributions involving the exchange of electro-weakly interacting particles like γ, W, Z, and H. The correction is defined as the ratio of a hard-scattering prediction including tree-level effects of order α_s^2 and α^2, and $\alpha_s\alpha$ (from the interference of QCD and EW diagrams), and weak loop corrections of order $\alpha_s^2\alpha$ to the LO prediction. This multiplicative factors method also allows for taking multiple effects into account, the convolution to a different PDF set, changing renormalisation or factorisation scale, or using a different tune of the MC generators.

Since producing NNLO predictions is computationally heavy, the APPLfast project [37] was developed, building upon the APPLGRID framework [38], which provides similar capabilities at NLO. It enables fast a posteriori convolutions of the perturbative coefficients with different choices of PDFs and scale settings. This is achieved by computing and storing the matrix elements of the fixed-order QCD calculation in the form of three-dimensional interpolation grids. These grids depend on the renormalisation/factorisation scale and the two Bjorken-x variables of the initial-state partons, and are constructed for each bin of the observable and each partonic subprocess. This approach allows for fast and flexible theoretical evaluations without rerunning the full computation.

1.2.4 Jets

In scattering processes, once the collision has occurred, an energetic quark or gluon radiates further gluons, which in turn may branch again, each step carrying progressively less energy. This perturbative cascade continues until the branching scale falls into the non-perturbative regime, where colour is confined and partons combine into colour-neutral hadrons. The resulting spray of hadrons "remembers" the preceding parton branching history. Since the initial parton is boosted, its decay products are collimated in the laboratory frame. A jet corresponds to this collimated spray of particles and provides a way to map the hadronic final state onto an observable that reflects the underlying parton kinematics.

1.2.4.1 Jet definition

Jets are not directly observable objects, but proxies for sprays of hadrons originating from partons. Their definition depends on how they are reconstructed and the choice of inputs. These choices are arbitrary and influenced by the goals of the analysis, so there is no unique jet definition.

There is a set of rules for how to group entities into jets (jet algorithm), and how their momenta are combined to define the jet momentum (recombination scheme). These definitions must be simple to implement and resilient to QCD effects. That is, they should be IRC safe and yield well-defined and finite results at all orders in perturbation theory. The jets should further behave equivalently at different orders of the QCD evolution: at the

parton, hadron and detector levels. In addition, the input four-vectors should be clustered in a way that has physical meaning, such that the history of the clustering steps reflects a meaningful hierarchy of branchings within the jet, offering insight into the parton shower structure.

Jets can be reconstructed from various types of inputs. In experimental data, jets are often built from calorimeter energy deposits or other objects, combining information from multiple detector subsystems. In simulation, jets can be reconstructed from stable hadrons, charged tracks, or even partons before hadronisation.

The most widespread recombination scheme is the 4-vector recombination scheme. It consists of adding the 4-vectors of merged particles and produces massive jets.

1.2.4.2 Jet algorithms

Jet algorithms define rules for grouping particles into a jet. This requires an external parameter, the jet radius R, specifying up to which angle separate partons are recombined into a single jet. It also takes a $p_{\mathrm{T}}^{\mathrm{min}}$ parameter, which defines the minimum transverse momentum of a jet. The choice of R and $p_{\mathrm{T}}^{\mathrm{min}}$ values results of a compromise between the ability to resolve multijet physics, the loss of radiation from jets, the additional spurious jets and the contamination from pile-up.

The basic idea of jet algorithms is to invert the QCD branching process, clustering pairs which are closest in a metric defined by the divergence structure of the theory. Many jet algorithms exist that come with diverse properties. The most popular at hadron colliders are generalised k_{T} algorithms. These sequential recombination algorithms recombine iteratively pairs of particles with a bottom-up procedure. It starts by defining the distance between every pair of entities, d_{ij}, and the distance to the beam axis d_{iB}:

$$d_{ij} = \min p_{T,i}^{2p},\ p_{T,j}^{2p}\frac{\Delta R_{ij}^2}{R^2},\ d_{iB} = p_{T,i}^{2p},$$

with $\Delta R^2 = \Delta y^2 + \Delta\phi^2$ the geometric distance between i and j, and the exponent p depends on the algorithm considered. Then, the following iterative procedure is applied until no entities remain:

1. Find the minimum among all d_{ij} and d_{iB} values,
2. If the minimum is a d_{ij}, merge i and j, and return to step 1.
3. If the minimum is a d_{iB}, declare i to be jet, and remove it from the list, and return to step 1.

In the end, only jets with transverse momentum $p_{\mathrm{T}} > p_{\mathrm{T}}^{\mathrm{min}}$ are kept. At this stage, the jet p_{T} does not reflect that of the initiating parton yet: several corrections are still required to achieve accurate agreement. This is the focus of Chapter 3.

Sequential clustering algorithms are characterised by a parameter p, which determines the nature of the clustering:

- If $p = 1$, the algorithm is known as the k_{T} algorithm [39,40]. It favours the clustering of soft and collinear emissions and is robust against background contributions.

- If $p = 0$, it is the Cambridge/Aachen (C/A) algorithm [41, 42]. It clusters entities based on their angular separation, independently of their energies. This makes it particularly relevant for studying jet substructure.
- If $p = -1$, it is the anti-k_{T} algorithm [43]. It clusters hard particles first, causing jets to grow around energetic cores. This leads to regular, conical jet shapes, which are advantageous in experimental particle physics for detector-level corrections. Anti-k_{T} is the algorithm of choice at LHC for jet finding.

The effects of jet clustering using different jet algorithms on the same input are shown in Figure 1.8.

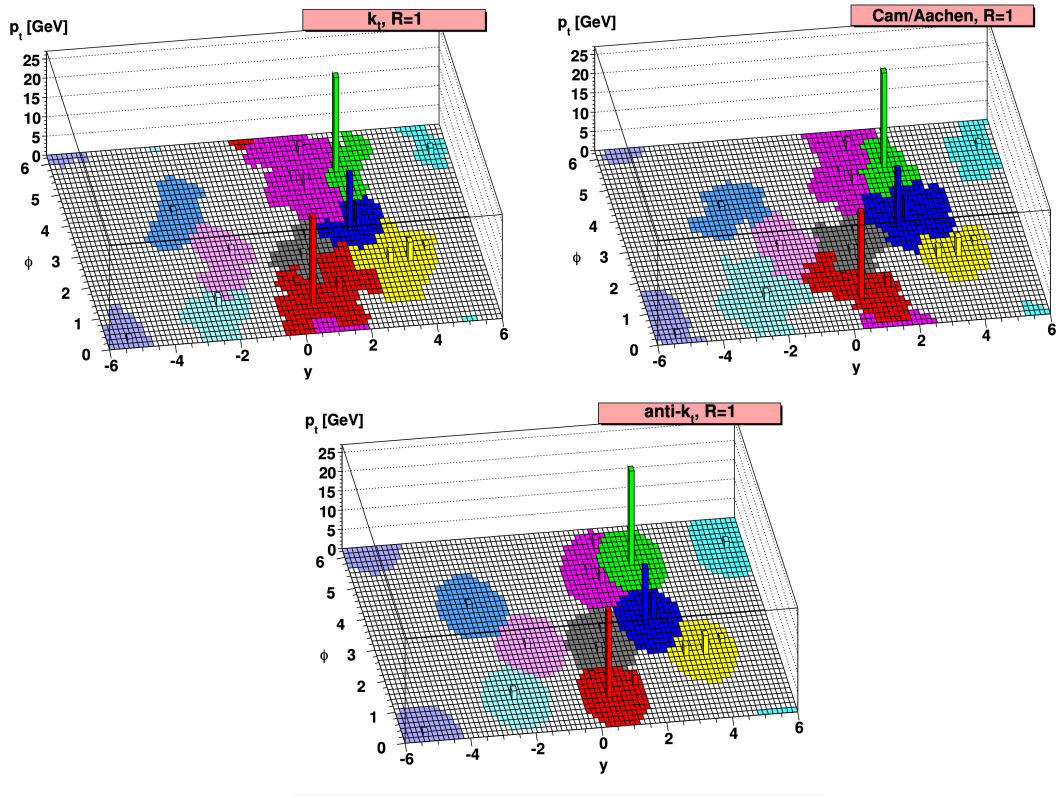


Figure 1.8: Jets clustered from the same inputs (parton-level event together with many soft emissions) with different jet algorithms: k_{T} (top left), Cambridge-Aachen (top right), anti-k_{T} (bottom). Reproduced from [44].

Sequential clustering algorithms were originally developed for lepton colliders, where no preferred beam axis exists and there is no pile-up. A prominent example is the Durham algorithm [45], widely used at LEP, the e^+e^- collider that preceded the LHC. It clusters nearby particles sequentially, but unlike hadron collider algorithms, it does not include a beam–particle distance. The distance between entities is modified to:

$$d_{ij} = 2\min(E_i^2, E_j^2)\,(1 - \cos\theta_{ij}),$$

where $\theta_{i,j}$ is the angle between the entities.

An alternative is the generalised k_{T} algorithm [46], which modifies the distance measure and introduces a beam–entity distance:

$$d_{ij} = \min(E_i^{2p}, E_j^{2p}) \frac{1 - \cos\theta_{ij}}{1 - \cos R}, \; d_{iB} = E_i^{2p}.$$

For values of $R \leq \pi$, the generalised k_{T} algorithm behaves in analogy with the pp algorithms. In particular, choosing $p = 1$ provides a simple, infrared and collinear safe way of obtaining a cone-like algorithm. Another alternative is the Valencia algorithm [47], which uses the following distance definitions:

$$d_{ij} = \min(E_i^{2p}, E_j^{2p}) \frac{1 - \cos\theta_{ij}}{R^2}, \; d_{iB} = p_{T,i}^{2p}.$$

For $p = 1$, the Valencia algorithm behaves similarly to the k_{T} algorithm of hadron colliders. Both those algorithms penalise relatively isolated particles in the forward and backwards directions, which are likely due to background processes.

1.2.4.3 Jets in hadron colliders

Jets are essential tools both for testing our understanding and predictions of high-energy QCD processes, and for identifying the partonic structure of decays of heavy particles such as top quarks and the W, Z, and Higgs bosons. In hadron colliders, jets are ubiquitous and often dominate the final state, making them not only central objects of interest but also a major source of background in many analyses. Figure 1.9 shows the production cross-sections of various processes at the LHC, which are largely dominated by the production of jets. The evaluation of the inclusive jet cross-section production measurement is the subject of Chapter 4.

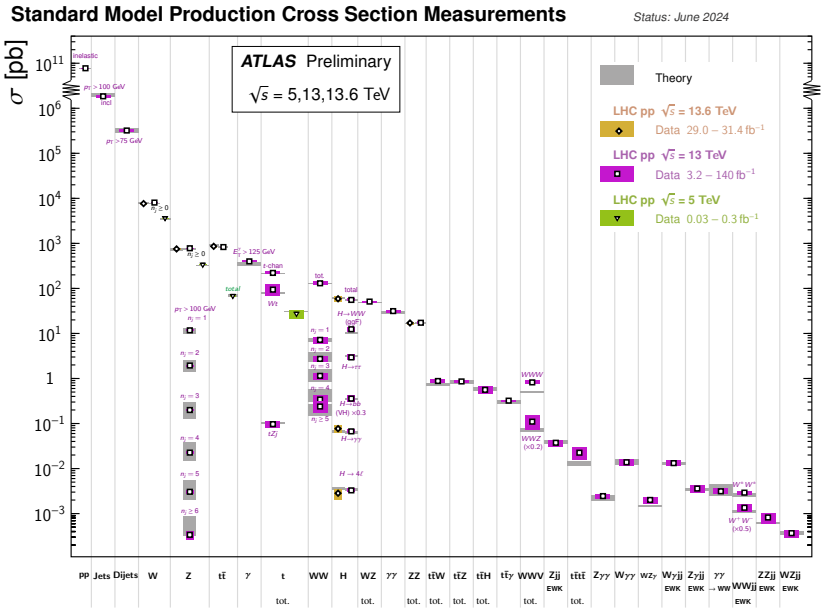
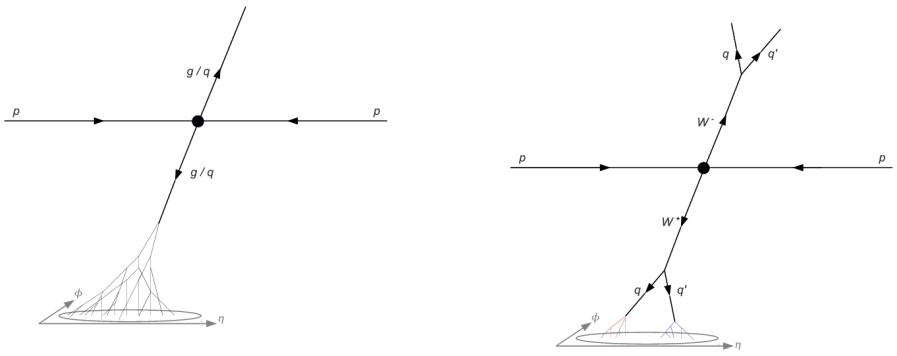


Figure 1.9: Summary of several Standard Model production cross-section measurements. Taken from [48].

The diversity of jet algorithms and parameters enables capturing the wide variety of physics processes probed at colliders. For example, smaller values of the radius parameter R are useful for resolving dense final states, such as those with many nearby jets. They are also less susceptible to contamination from pile-up. On the contrary, larger R values are employed to reconstruct boosted heavy objects whose decay products are collimated into a single, wide jet. These large-R jets can be further probed with jet substructure techniques to identify their internal decay patterns and distinguish them from QCD background jets.

1.2.4.4 Jet substructure

Jets are objects that gather multiple particles originating from different QCD processes occurring at various energy scales. Studying their substructure allows for probing the internal dynamics of the jet. It can give information on the hard process that initiated it. For example, gluon-initiated jets tend to be broader and contain more particles than quark-initiated jets due to the larger QCD colour factor associated with gluons, which leads to more radiation. Jets initiated by heavy boosted particles decaying into a quark pair often exhibit a two- or three-prong substructure. Two jets from different processes are illustrated in Figure 1.10, showing the difference in internal structure and energy distribution of a jet initiated with two different processes.



(a) (b)

Figure 1.10: Schematic representation of a jet from a QCD event (a) and from the decay of a boosted W boson (b). Taken from [49].

Several jet substructure variables are used to characterise the internal structure of jets. N-subjetiness [49] quantifies to what degree a jet can be regarded as containing N subjets. It does so by checking how closely the particles inside the jet are grouped around N candidate axes. The more tightly clustered the particles are around these axes, the smaller the N-subjetiness value. A N-subjetiness close to zero indicates that the jet considered is likely to have N subjets. N-subjetiness and the ratio of them are often used to distinguish QCD jets from jets resulting from the decay of a boosted particle.

The energy correlation function [50] is a subject-independent method used to identify N-pronged jets. This approach looks directly at the distribution of energy and angles between particles in the jet. The $(N + 1)$-point correlators identify N-pronged jets by evaluating how the energies and angles between groups of particles within the jet are

correlated. The energy correlation function as a jet substructure variable is also used to distinguish jets of various origins.

The Lund jet plane (LJP) [51] is another type of jet substructure observable, which maps the emissions inside a jet onto a two-dimensional space defined by their transverse momentum and angular separation. Unlike the other observables discussed here, the LJP focuses on the branching history of the jet. It is the subject of Chapter 5.

Chapter 2

The ATLAS experiment and object reconstruction in ATLAS

Introduction

Most of the particles predicted by the Standard Model are unstable. The only known stable free particles are the photon, electron, neutrinos, and proton. Due to the confinement property of QCD, free quarks cannot exist in isolation: they form bound states known as hadrons. Among these, only the proton is stable. Given enough time, any unstable particle will undergo physically possible and kinematically allowed interactions (those permitted by the Standard Model and constrained by conservation laws such as energy, electric charge, and baryon number). Because of energy conservation, these processes can only result in final states of lower energy (and thus lower mass). Consequently, if a decay is allowed, it will eventually occur, transforming the particle into lighter products.

To detect such unstable particles, they must first be produced. According to Einstein's famous relation $E = mc^2$, energy and mass are equivalent: high-energy collisions can be used to produce massive particles. This is the role of particle colliders: by accelerating and colliding stable particles at high energies, enough energy is concentrated at the interaction point to create heavier, unstable particles. Thus, the history of discoveries in experimental particle physics is closely tied to the progress in particle accelerator technology.

2.1 LHC at CERN

CERN is the European organisation for nuclear research, founded in 1954, located at the border of Switzerland and France, near Geneva. It is an international organisation that operates the largest particle physics laboratory in the world. Nowadays, more than 17,500 people of 110 nationalities contribute to CERN, from institutes in more than 70 countries. The most recent addition to CERN is the Large Hadron Collider (LHC) [52], a 27-kilometre-long circular accelerator that collides protons at the highest kinematic energies ever achieved in a laboratory setting.

2.1.1 The acceleration chain

CERN's accelerator complex consists of a series of machines, each designed to accelerate particle beams to progressively higher energies. Radio-frequency (RF) cavities accelerate the particles, while superconducting magnets bend and focus the beams along their trajectories. The beam is transferred from one accelerator to the next, with each stage boosting its energy further. Many of the accelerators in the chain also have their experimental areas, where lower-energy beams are used for various experiments.

In the LHC, the final element of this chain, each proton is accelerated to velocities approaching the speed of light. The two counter-rotating beams are brought into collision at four interaction points, where detectors are located to record the results. These proton–proton (pp) collisions occur at extremely high energies in the centre-of-mass frame (reaching the multi-TeV scale), where the centre-of-mass energy $\sqrt{s}$ represents the total energy available for particle production.

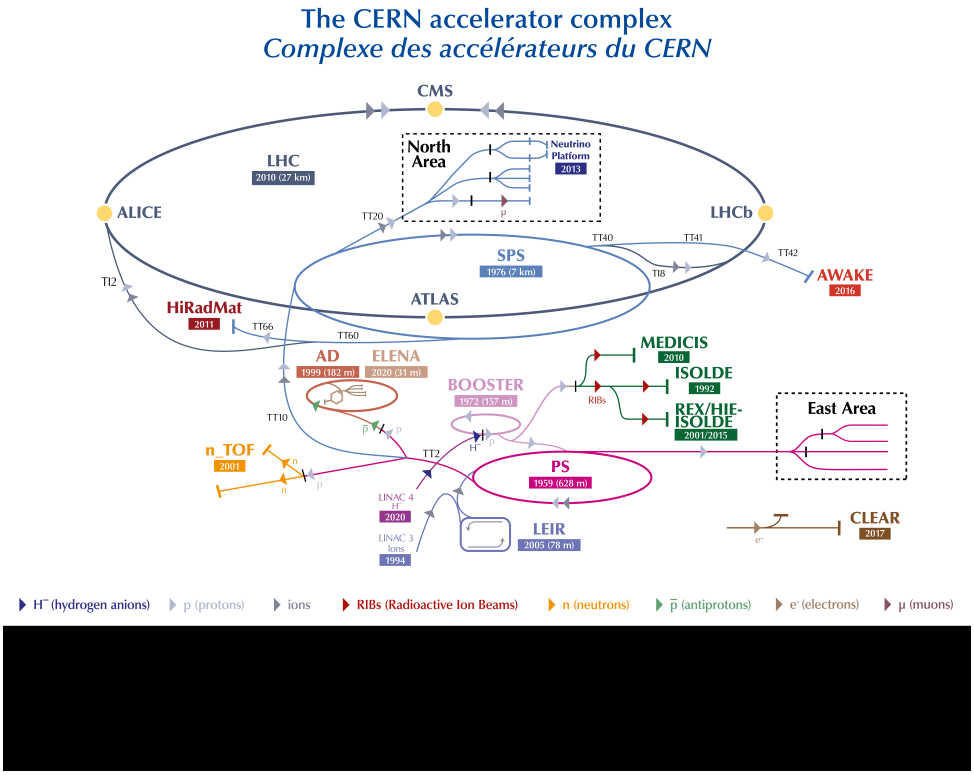


Figure 2.1: CERN accelerator complex. The commissioning dates of the accelerators are indicated. Protons circulating in the LHC are first accelerated in LINAC 2, then in the Booster, followed by the PS, then the SPS, and finally in the LHC. Taken from [53].

To record and study the products of high-energy collisions, detectors are placed at the four interaction points of the LHC. Each of these hosts a dedicated experiment, optimised for specific physics goals:

- ATLAS (A Toroidal LHC AppartuS) [54], and CMS (Compact Muon Solenoid) [55], are designed to address a broad range of physics topics, from precision measurements of Standard Model processes such as the Higgs boson and the top quark, to searches

for new phenomena beyond the Standard Model.

- LHCb (LHC beauty) [56], focuses on “b-physics”, i.e. the study of hadrons containing at least one b-quark, to test the fundamental symmetries of the Standard Model or their violation.
- ALICE (A Large Ion Collider Experiment) [57], dedicated to the study of the quark-gluon plasma, a state of matter that is likely to have occurred in the earliest instants of the Universe. It primarily analyses heavy-ion collisions.

2.1.2 Data-taking

The LHC has been colliding particles since 2008. Data-taking periods are organised in "runs", which extend over several years. Two runs have been completed so far, and one is ongoing:

- Run 1, during which pp collisions occured at $\sqrt{s} = 7$ TeV during 2010 and 2011, and 8 TeV in 2012.
- Run 2, during which pp collisions occured at $\sqrt{s} = 13$ TeV from 2015 to 2018.
- Run 3, which is ongoing since 2022, with pp collisions at $\sqrt{s} = 13.6$ TeV

During these periods, heavy-ion collisions have also been recorded, as well as lower-energy pp collisions. Between the runs, long shutdowns are scheduled to prepare, upgrade, or maintain the detectors. This thesis focuses on pp collisions recorded during Run 2 and Run 3.

To record rare events and accumulate sufficient statistics for precision measurements, the runs are designed to collect a large integrated luminosity. The instantaneous luminosity $\mathcal{L}$ represents the ability for the LHC to produce a certain number of events with a given cross-section σ per unit of time:

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

Since the instantaneous luminosity varies over time during a run, the integrated luminosity is defined as:

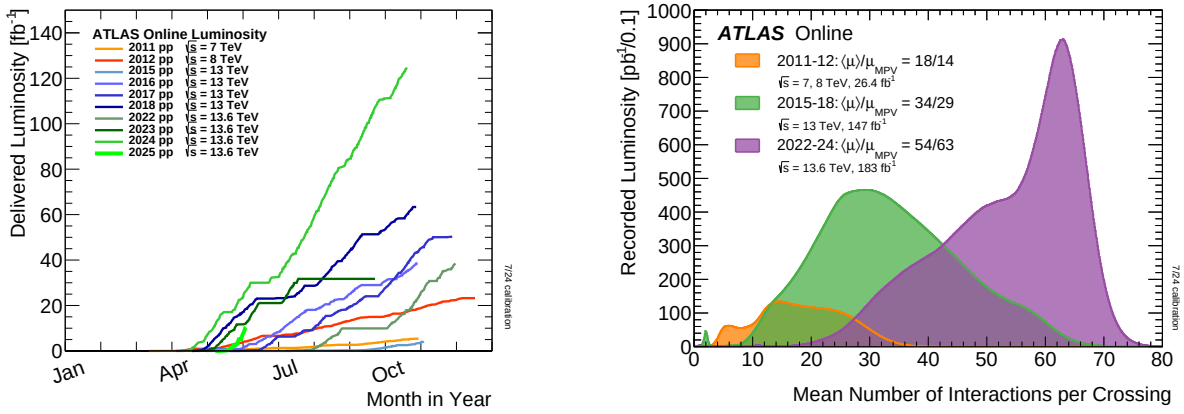
$$\mathcal{L}_{\text{int}} = \int \mathcal{L}\mathrm{d}t.$$

It represents the total dataset collected by an experiment and is typically expressed in units of inverse femtobarns, where barn is the unit of cross-sections (1 b $= 10^{-28}$ m^2). Thus, an inverse femtobarn corresponds to the number of events expected for a process with a cross-section of one femtobarn. The cumulative luminosity delivered to the ATLAS detector for pp collisions is presented in Figure 2.2a. During Run 2, ATLAS recorded 139 fb^{-1}, with a relative uncertainty of 0.83% [58]. A precise knowledge of the luminosity is important, particularly to measure precisely the cross-section of a specific process.

Maximising luminosity is a goal of the LHC’s operation strategy. It involves optimising various factors such as the number of bunches, the number of protons per bunch, and beam focusing at the interaction points. New upgrades are planned after the end of Run 3, to increase the instantaneous luminosity by a factor of five. This phase is called the High Luminosity LHC (HL-LHC) [59], and is expected to record 3 to 4 ab^{-1} of 13.6–14 TeV pp collisions, until $\sim$ 2041.

In the CERN accelerator chain, protons are grouped into bunches, each containing approximately 1.15×10^{11} protons. These bunches are spaced 25 nanoseconds apart. Since each bunch contains such a large number of protons, the probability of multiple interactions occurring in a single bunch crossing is high. This phenomenon, known as in-time pile-up, refers to additional pp collisions happening at the same time as the collision of interest. It is parametrised by μ, the mean number of interactions per bunch crossing, and is proportional to the instantaneous luminosity. The distribution of μ as measured by ATLAS in pp collisions, for Run 1, Run 2 and Run 3, is shown in Figure 2.2b. During the HL-LHC phase, the average number of interactions per bunch crossing is expected to reach 200.

In addition, due to the response time of the detectors (often longer than the interval between bunch crossings), signals from multiple consecutive bunches can overlap. This effect is referred to as out-of-time pile-up, where remnants from earlier or later collisions are recorded together with the event of interest. Pile-up affects event reconstruction in the detectors, making it more challenging to isolate and analyse the desired interactions.



(a) (b)

Figure 2.2: (a) Cumulative luminosity versus day delivered to ATLAS during stable beams and for high-energy pp collisions. (b) Luminosity-weighted distributions of the mean number of interactions per crossing for pp collisions recorded by ATLAS. $\langle\mu\rangle$ is the average μ value, and μ_{MPV} is the most probable value of μ for each year of data-taking. Both figures are taken from [60].

2.2 The ATLAS experiment

The ATLAS experiment [54] is a multipurpose particle detector with a forward–backward symmetric cylindrical geometry. It is the largest detector at CERN, measuring about 44 meters in length, 25 meters in diameter and weighing about 7000 tons. A schematic of the ATLAS detector is shown in Figure 2.3.

ATLAS uses a right-handed coordinate system with its origin at the nominal interaction point (IP) in the centre of the detector and the z-axis along the beam pipe. The x-axis points from the IP to the centre of the LHC ring, and the y-axis points upwards. Polar

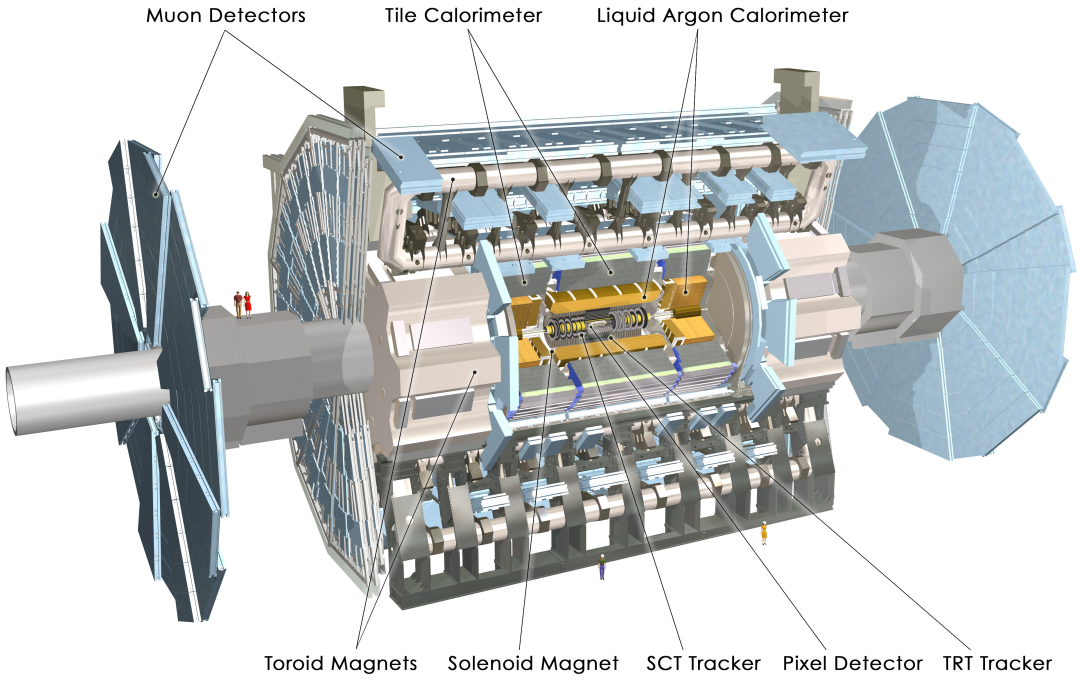


Figure 2.3: Cut-away view of the ATLAS detector showing its various components. Taken from [61].

coordinates (r, ϕ) are used in the transverse plane, ϕ being the azimuthal angle around the x-axis. In particle physics, the rapidity y is defined relative to the beam axis:

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

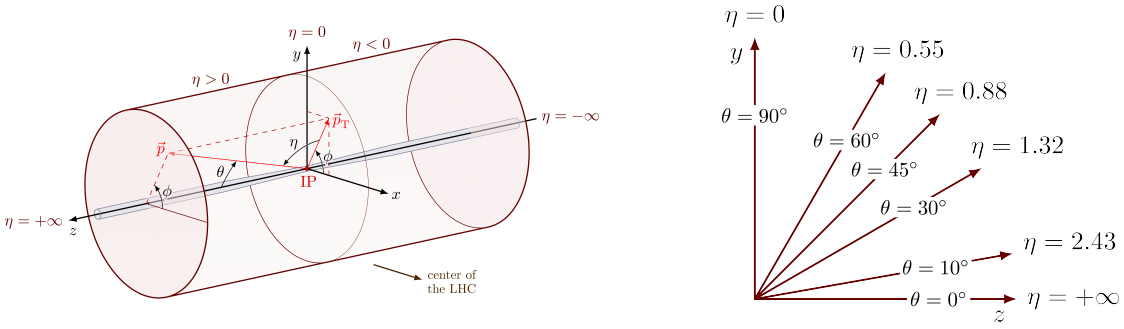
Most widely used is the pseudorapidity, defined in terms of the polar angle θ as

$$\eta = -\ln \tan(\theta/2).$$

It is equal to the rapidity y in the relativistic limit. Angular distance is measured in units of $\Delta R \equiv \sqrt{(\Delta y)^2 + (\Delta \phi)^2}$. The coordinate system is represented in Figure 2.4.

The ATLAS detector consists of several subdetectors:

- An inner tracking detector, used to reconstruct the trajectory and momentum of the charged particles. It is also used to identify the primary and secondary vertices and separate pile-up vertices from the hard-scatter. It is surrounded by a thin superconducting solenoid that provides a magnetic field that bends the trajectory of the charged particles.
- A calorimeter system, to measure the energy and position of electromagnetic particles and hadrons.
- A muon spectrometer, placed at the outermost part of the ATLAS detector, used to identify muons and measure their momenta. It features the largest superconducting toroid magnet, which provides a magnetic field to bend the trajectories of the muons.



(a) (b)

Figure 2.4: (a) ATLAS Coordinate system (b) Correspondence between polar angle and pseudorapidity. Adapted from [62].

The luminosity is measured mainly by the LUCID–2 [63] detector, which is located close to the beampipe.

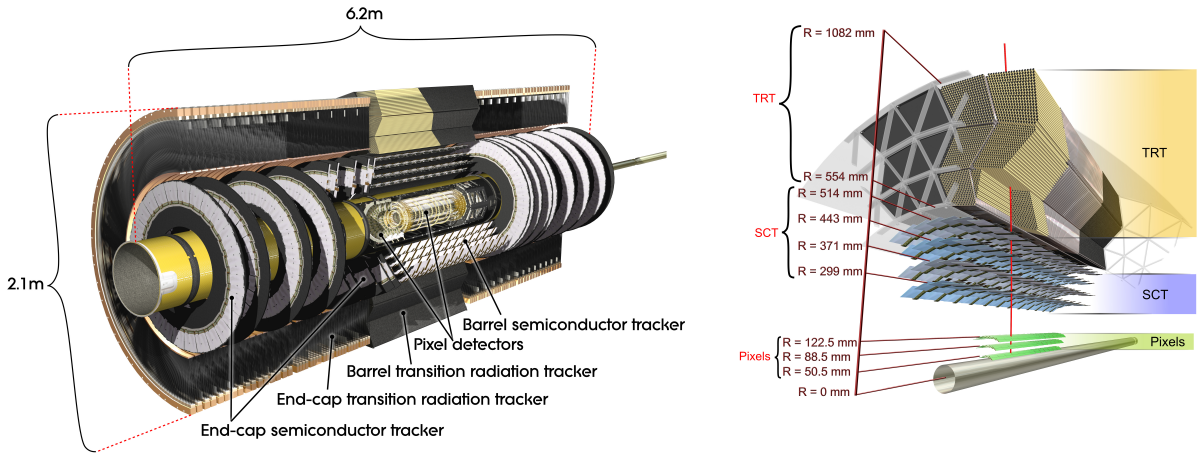
Some particles, such as neutrinos, traverse the entire detector without leaving a direct signal, as they only interact via the weak force. Their presence is inferred from the imbalance in the total transverse momentum measured in the event, referred to as the missing transverse energy $E_{\mathrm{T}}^{\mathrm{miss}}$. To ensure a reliable measure of $E_{\mathrm{T}}^{\mathrm{miss}}$, the detector has a near 4π coverage in solid angle, and has sufficient depth to contain and measure most particles.

2.2.1 The inner detector

The Inner Detector (ID) [64, 65] consists of a succession of layers of silicon sensors and gaseous detectors, immersed in a 2 T magnetic field from a solenoid magnet. It has a coverage up to $|\eta| < 2.5$, and is divided into three regions: the barrel, concentric around the beam axis, covering the $|\eta| \lesssim 1$ region, and the two end-caps, parallel to the beam axis, covering the $1 \lesssim |\eta| < 2.5$. The layout of the ID is shown in Figure 2.5.

When a charged particle traverses the ID, it ionises the detector medium and produces hits in the active elements. As the particle moves through the successive layers of the ID, it generates multiple hits along its path. These hits are then combined using pattern recognition and track-fitting algorithms to reconstruct the trajectory of the particle, referred to as a track. The ID has been designed to provide a p_{T} resolution of $\sigma_{p\mathrm{T}}/p_{\mathrm{T}} = 0.05\% \times p_{\mathrm{T}} \oplus 1\%$, and a transverse impact parameter resolution of 10 μm for high momentum particle in the central region. It is made of several subdetectors, described in the following from the closest to the farthest from the beam.

The Pixel Detector (PD) [67] consists of silicon pixel modules arranged in four concentric barrel layers and two endcaps of three disks each. In the pixels, incident particles generate electron-hole pairs along their trajectories. These charges drift through the silicon sensors and are read out by radiation-hard front-end chips, bump-bonded to the sensor. Hits in



(a) (b)

Figure 2.5: (a) Three-dimensional view and (b) transverse view of the inner detector. Taken from [66].

the pixel are read out if the signal exceeds a tunable threshold. The first barrel layer, the Insertable B-Layer (IBL) [68], achieves intrinsic accuracy of 10 μm in the (R, ϕ) plane and 66.5 μm along the longitudinal z axis. The remaining barrel layers are composed of silicon pixel modules that are segmented in R and z. The intrinsic accuracies of the three main barrel layers are 10 μm in the (R, ϕ) plane and 115 μm along the longitudinal z axis. For the endcaps, the intrinsic accuracies are 10 μm in the (R, ϕ) plane and 115 μm in the radial R direction. Thanks to its fine granularity, the PD achieves excellent spatial resolution of 12 μm in R and 50 μm in z.

The Semi-Conductor Tracker (SCT) consists of silicon micro-strip modules, using the same technology as the PD. The modules are arranged in four concentric barrel layers, which are approximately parallel to both the solenoidal magnetic field and the beam axis, as well as in two endcaps, each comprising nine disks with radially oriented strips. The micro-strips are read out by radiation-hard front-end chips. Similar to the PD, hits are recorded only if the signal height exceeds a preset threshold. The intrinsic accuracy of the barrel modules is 130 μm in the (R, ϕ) plane. The SCT does not provide position measurements along the z direction in the barrel.

The Transition Radiation Tracker (TRT) enables track-following up to $|\eta| = 2.0$. It is made of 4 mm diameter proportional drift tubes (straws) filled with a gas mixture of xenon, carbon dioxide and dioxygen, with a thin tungsten wire at its centre. As a charged particle traverses the gas, it ionises the medium, producing electrons that drift toward and are collected by the central wire. In the barrel region, the straws are oriented parallel to the beam axis, while in the end-cap region, they are arranged radially in wheels. Although the TRT has the lowest spatial resolution among the ID subsystems, it can provide up to 36 measurements per track, significantly enhancing track reconstruction. Additionally, the TRT offers electron identification through transition radiation, which arises when a relativistic particle crosses the boundary between materials with differing dielectric constants. This effect, which is more likely for particles with high E/m ratios,

enables the discrimination between electrons and pions. Transition radiation is produced in polypropylene fibres in the barrel region and foils in the endcaps, located between the straws.

2.2.2 The calorimeter system

2.2.2.1 Principle

Calorimeters are designed to measure the energy carried by incoming particles. As particles go through the material, they undergo interactions that produce showers. At high energy, photons mainly interact via electron-positron pair production, while electrons primarily lose energy through bremsstrahlung, in which they emit photons. These two processes lead to the development of an electromagnetic shower, a cascade of secondary electrons, positrons, and photons. Hadrons interact both strongly with atomic nuclei and electromagnetically. These interactions can result in the emission of secondary particles (pions, nucleons), which may then further interact with the material, producing a hadronic shower. A significant portion of hadronic showers includes an electromagnetic component, primarily from the decay of neutral pions ($\pi^0 \to \gamma\gamma$), which contributes to the overall energy deposition. The development and structure of those two types of showers differ significantly. These differences allow for the discrimination between electromagnetic and hadronic showers, for example, through their longitudinal and lateral profiles, as well as the fraction of energy deposited in each calorimeter layer. Moreover, the calorimeter response, defined as the average signal produced per unit of deposited energy, is generally lower for hadronic showers than for electromagnetic showers. This is because a fraction of the energy in hadronic showers is not detectable due to invisible energy losses such as nuclear breakup, nuclear excitation and neutrino production. The ATLAS calorimeter system is non-compensating, meaning it is not designed to provide an equal response to hadronic and electromagnetic showers.

The calorimeters of ATLAS are sampling calorimeters, composed of alternating layers of absorber and active materials. The absorber layers are used to initiate and develop the particle showers, while the active layers are used to measure the signal. The main advantage of sampling calorimeters is their intrinsic segmentation, which helps in particle reconstruction and identification. However, a key disadvantage is that part of the particle energy is lost in the absorber material and is not directly measured.

2.2.2.2 Layout

The calorimeter system [69] has been designed to be deep enough to contain the entire particle shower, and its coverage goes up to $|\eta| < 4.9$ to ensure that all the energy of the event can be measured. It is divided into two subsystems:

- The Electromagnetic calorimeter (ECAL) provides coverage in the region $|\eta| < 3.2$.
- The Hadronic calorimeters (HCAL), composed of the tile barrel and extended barrel modules (TileCal), covering the region $|\eta| < 1.7$, and the Hadronic End-Cap (HEC) calorimeters that extend the coverage to $1.5 < |\eta| < 3.2$. The Forward Calorimeters (FCal) cover the forward region of $3.1 < |\eta| < 4.9$.

The transition regions between different calorimeter subsystems pose challenges for jet calibration as a function of pseudorapidity η (discussed in Chapter 3). These regions include the interface between the barrel and endcap calorimeters ($|\eta| \sim 1.4$), between the HEC and the FCal ($|\eta| \sim 3.1$), and the boundary of the tracker coverage ($|\eta| \sim 2.5$).

Due to the presence of inactive material in front of the calorimeter system, the energy resolution is degraded. To mitigate this effect, a liquid argon presampler is placed inside the solenoid, covering $|\eta| < 1.8$. It provides a measurement of the energy deposited before the dead material. The structure of the calorimeter system is shown in Figure 2.6.

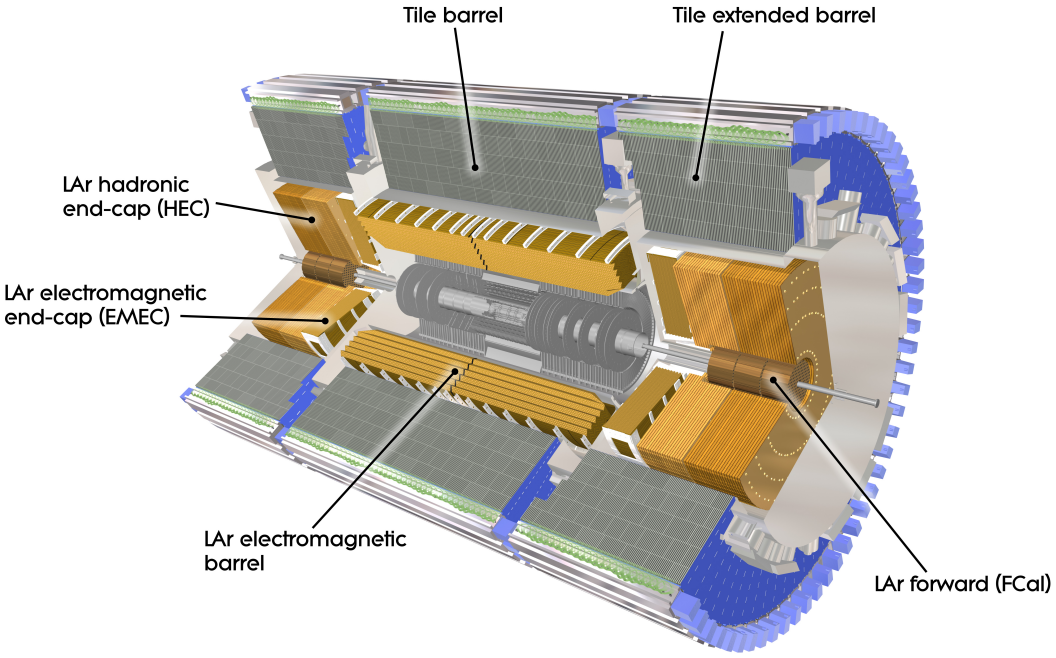


Figure 2.6: Layout of the ATLAS calorimeters. Taken from [70].

The energy resolution of calorimeters can be parametrised as:

$$\frac{\sigma(E)}{E} = \frac{N}{E} \oplus \frac{S}{\sqrt{E}} \oplus C,$$

where N is the noise term, accounting for electronic noise and the effects of pile-up; S is the stochastic term, which includes contributions from statistical fluctuations in the shower development, quantum fluctuations in the signal, and the impact of sampling; and C is the constant term, which reflects systematic effects such as imperfect calibration, detector geometry, energy loss before the calorimeter, and particles escaping the calorimeter.

2.2.2.3 Electromagnetic calorimeter

The Electromagnetic Calorimeter (ECAL) is a lead-liquid argon (LAr) sampling calorimeter with accordion geometry [71, 72]. When electrons or photons traverse the lead absorbers, they initiate an electromagnetic shower. The resulting secondary particles ionise

the liquid argon, and the charges drift toward the electrodes, producing a signal. Liquid argon is chosen for its very linear behaviour, its stability over time and its radiation-hardness. The accordion geometry provides full coverage in ϕ without azimuthal cracks, and allows a fast signal extraction.

The ECAL is divided into a barrel section ($|\eta| < 1.475$) and two end-cap components ($1.375 < |\eta| < 3.2$, EMEC). In the transition region between those two parts, a large amount of material is present before the first calorimeter layer. To mitigate the resulting degradation in energy measurement, a scintillator slab is inserted between the barrel and end-cap to complement the presampler, covering the region $1.0 < |\eta| < 1.6$. The barrel is segmented longitudinally into three layers in depth, while each end-cap is segmented into two. The first layer of the barrel is finely segmented in the η direction, with a granularity of $\Delta\eta \times \Delta\phi = 0.0031 \times 0.025$, optimised for precision measurements of electrons and photons. The second layer, which collects most of the electromagnetic shower energy, has an angular resolution of $\Delta\eta \times \Delta\phi = 0.025 \times 0.025$. The third layer is designed to measure the tails of the showers, and has an angular resolution of $\Delta\eta \times \Delta\phi = 0.05 \times 0.025$. The coarser granularity of the two outer layers is sufficient to satisfy the physics requirements for jet reconstruction. In the end cap, the lateral granularity is coarser, and the spatial resolution varies, ranging from $\Delta\eta \times \Delta\phi = 0.0031 \times 0.1$ to $\Delta\eta \times \Delta\phi = 0.1 \times 0.1$.

The design energy resolution of the ECAL is

$$\frac{\sigma(E)}{E} = \frac{10\%}{\sqrt{E}} \oplus 0.7\%.$$

The time response of the calorimeter is approximately 450 ns, which is significantly longer than the 25 ns bunch crossing interval at the LHC, leading to out-of-time pile-up.

2.2.2.4 Hadronic calorimeter (HCAL)

The hadronic calorimeter (HCAL) surrounds the ECAL. It is divided into different parts with different technologies to suit the different requirements and radiation environments.

The central barrel and extended barrel of the hadronic calorimeter cover the regions $|\eta| < 1$ and $0.8 < |\eta| < 1.6$, respectively. These sections use steel as absorber material and scintillating tiles as the active medium (TileCal) [73]. The TileCal is segmented in depth in three layers. The scintillator produces ultraviolet photons, which are collected by wavelength-shifting fibres that convert them into visible light, which is turned into an electric signal by photomultiplier tubes.

The end-cap hadronic calorimeters (HEC) consist of two independent wheels covering the region $1.7 < |\eta| < 3.2$, and use two copper plates as absorbers and liquid argon (LAr) as the active material. The design energy resolution for the TileCal and HEC is

$$\frac{\sigma(E)}{E} = \frac{50\%}{\sqrt{E}} \oplus 3\%.$$

The region $3.1 < |\eta| < 4.9$ is covered by the Forward Calorimeters (FCal), which is segmented longitudinally into three and uses LAr as the active material and a combination

of copper and tungsten as absorbers. The use of high-density materials ensures sufficient absorption in the high-radiation environment of the forward region. The hadronic calorimeters have coarser granularity than the electromagnetic calorimeters, ranging from $\Delta\eta \times \Delta\phi = 0.1 \times 0.1$ for the first layers of tiles to about $\Delta\eta \times \Delta\phi = 0.2 \times 0.2$ for the FCAL. The design energy resolution for the FCAL is

$$\frac{\sigma(E)}{E} = \frac{100\%}{\sqrt{E}} \oplus 10\%.$$

2.2.3 The muon spectrometer

The muons go through the ID and the calorimeter system without being absorbed. As a result, a dedicated system is required to detect them and precisely measure their momentum. The Muon Spectrometer (MS) [74] forms the outermost layer of the ATLAS detector, covering up to $|\eta| < 2.7$, as shown in Figure 2.7.

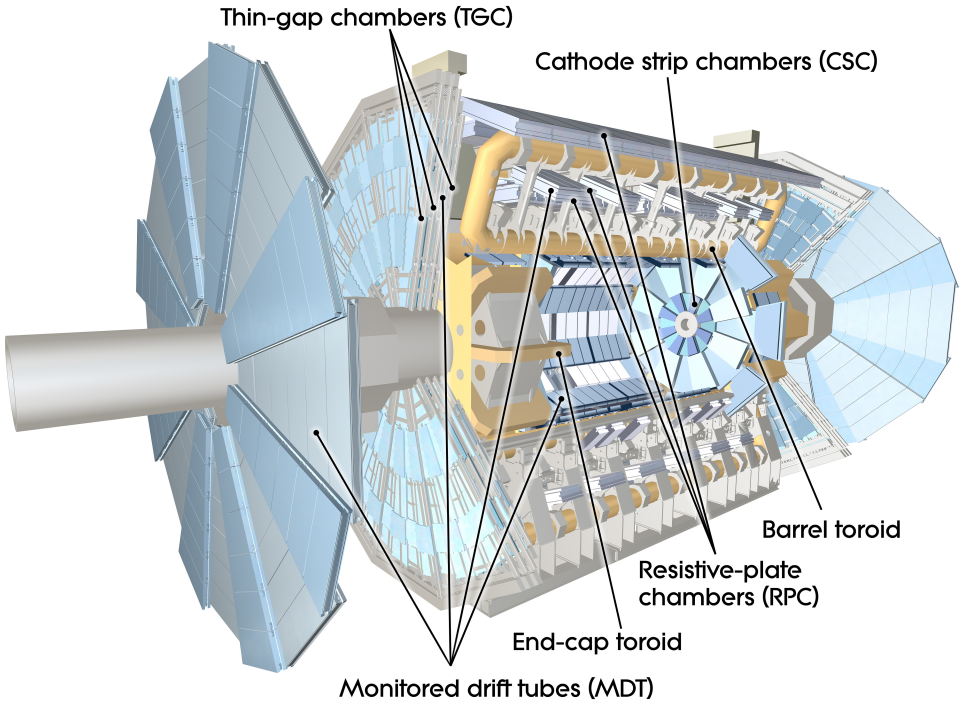


Figure 2.7: Layout of the ATLAS muon system. Taken from [75].

It is based on the magnetic deflection of muon tracks using large superconducting air-core toroid magnets. Magnetic fields are provided by a large barrel toroid covering $|\eta| < 1.4$ and by two smaller end-cap toroids covering $1.6 < |\eta| < 2.7$. The magnetic field strength is not uniform. The bending power varies from 1.5 to 5.5 T · m in the barrel region and from 1.7 to 7.5 T · m in the endcap regions. The time stability of the field strength is monitored by a set of approximately 1,800 Hall probes. The Muon Spectrometer is composed of four types of detectors: two dedicated to precision tracking, and two used for the triggering system:

- The Monitored Drift Tubes (MDTs), used for precision tracking in the $|\eta| < 2.7$ region,

- The Cathode Strip Chambers (CSCs), which are multiwire proportional chambers providing precision tracking in the $2.0 < |\eta| < 2.7$ region,
- The Resistive Plate Chambers (RPCs), which are gaseous electrode-plate detectors with coarser granularity but excellent timing resolution operating in the $|\eta| < 1.05$ region, used for triggering,
- The Thin Gap Chambers (TGCs), which are fast multiwire proportional chambers, provide a second measurement in the forward region $1.05 < |\eta| < 2.7$ for the trigger system.

In addition to triggering, the RPCs and TGCs provide bunch-crossing identification.

In the barrel region, muon tracks are measured using chambers arranged in three cylindrical layers around the beam axis. In the transition and end-cap regions, the chambers are installed in three planar layers perpendicular to the beam axis.

2.2.4 The trigger system

The collision rate in ATLAS is of the order of 40 MHz, which amounts to 10^9 collisions per second. This rate is vastly higher than the event recording capability of the detector, which is limited to approximately 1.5 kHz due to constraints on computing and storage resources. To reduce the data flow to manageable levels while preserving events of potential interest for physics analyses, a two-level trigger system [76] is employed to efficiently select relevant collisions and discard the rest.

The architecture of the ATLAS trigger system is illustrated in Figure 2.8 and described below:

- The first-level triggers (Level 1 triggers or L1) [77] are fast and hardware-based. It uses a subset of detector information to perform a fast partial reconstruction of the event and make a basic selection based on physics criteria. It is implemented in custom electronics and reduces the event rate from 40 MHz to a maximum of 100 kHz.
- The second-level triggers (High-level triggers or HLT) are seeded by the L1 triggers. Software algorithms with access to the full detector information are used to reconstruct the event thoroughly. It yields a recorded event rate of about 1.5 kHz.

Typically, those triggers search for high transverse momentum lepton, photons, jets or missing E_T. There are multiple triggers with very different criteria on all sorts of objects. They are designed to target a specific phase space such that analysers can rely on only a few of them, and the criteria are broad enough not to lose interesting events.

To accommodate the bandwidth limit, both L1 and HLT triggers are prescaled. A trigger with a prescale of w means that only $1/w$ of the events that pass the kinematic condition of this trigger are randomly chosen to be kept: those events are said to have fired the trigger. They are stored and assigned a weight w so that they represent all the events that fulfil the trigger conditions. All triggers are prescaled independently. The events are then stored and reconstructed, and further selections can be applied according to the event topology required for the analysis performed.

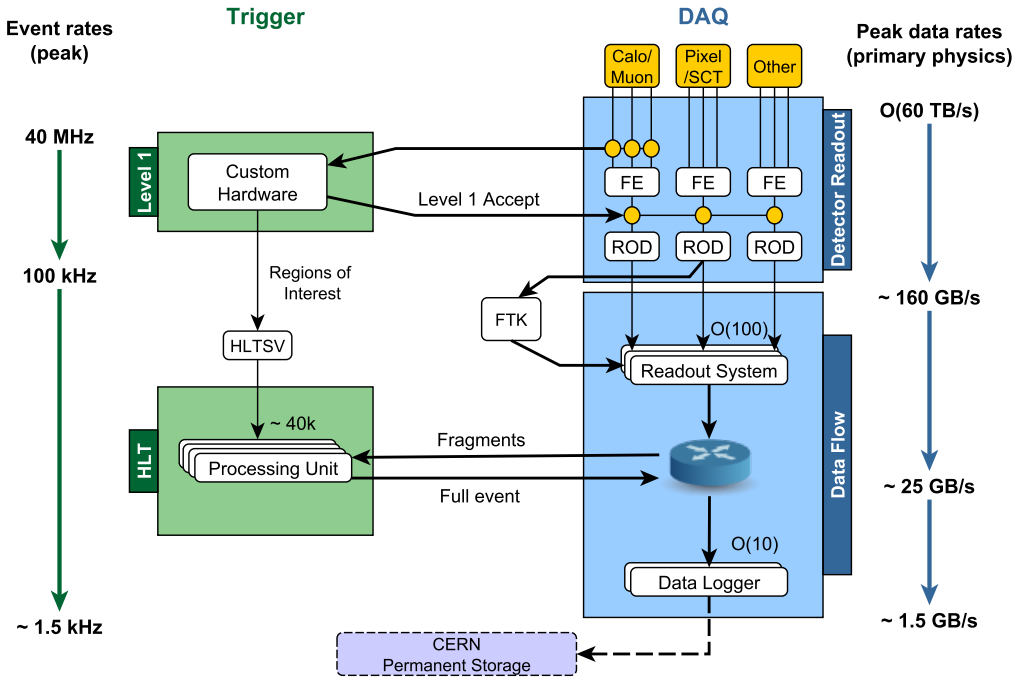


Figure 2.8: Trigger and Data Acquisition system used by ATLAS for Run 2. Taken from [78].

2.3 Reconstruction and data processing in ATLAS

At this stage, the datasets consist of hits in the tracking detectors and energy deposits in the calorimeters. These raw signals must be interpreted and associated with the particles that produced them. The information needs to be stored in a structured and efficient format to optimise memory usage and computing performance for physics analyses.

2.3.1 General object reconstruction

This section provides a brief overview of the reconstruction of tracks, vertices, electrons, photons, muons, and neutrinos. The reconstruction of jets is discussed in detail in the following section.

2.3.1.1 Tracks

Tracks are described by five parameters and a reference point, using a perigee representation [79]. The reference point is the average position of the pp interactions (beamspot position). The transverse and longitudinal impact parameters d_0 and z_0 are defined with respect to this reference point. The directional information is given by the azimuthal angle ϕ and the polar angle θ. Finally, the charge of the reconstructed track divided by the magnitude of its momentum q/p characterises the curvature of the track due to the magnetic field.

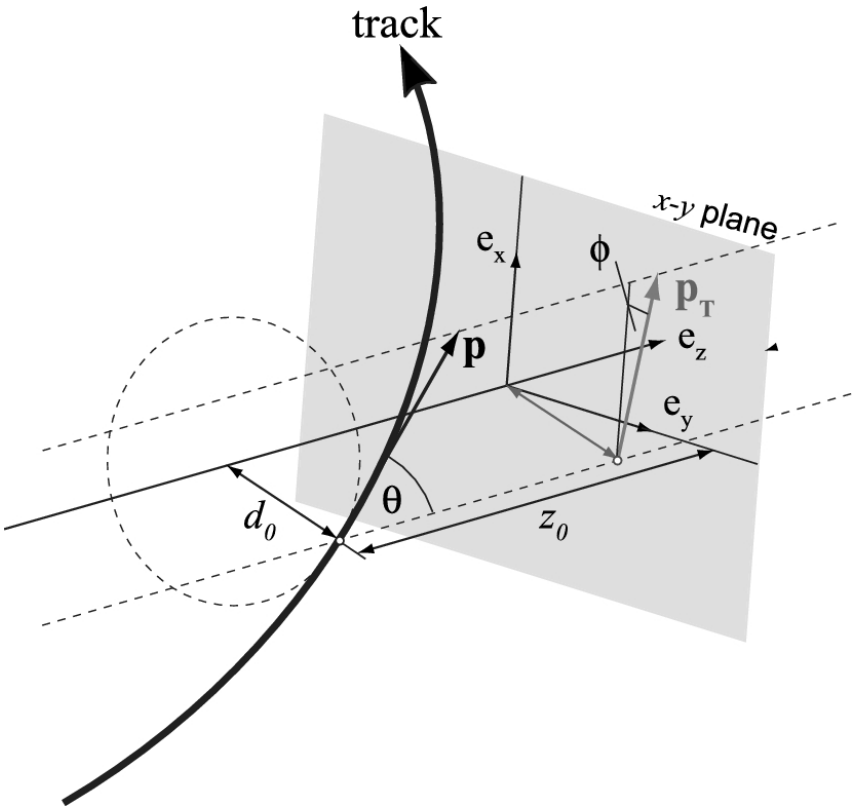


Figure 2.9: Illustration of the five global track parameters, d_0, z_0, ϕ, θ and q/p. These are defined relative to a reference point, the perigee, or point of closest approach to the beam line. Taken from [79].

The track finding process starts by identifying the seeds, which are combined signals from adjacent channels in the Pixel and SCT subdetectors. The position uncertainty of these seeds is determined by the detector geometry and sensor pitch. Then, roads are built along the seed, where the roads define where compatible clusters are searched for, to potentially extend the seed into a full track.

Track trajectories are constructed using a combinatorial Kalman filter, where at each step the trajectory is extended by clusters from consecutive detector layers. Once track candidates are formed, each is assigned a score based on a range of quality criteria. Ambiguity resolution is then performed: lower-quality candidates sharing a large number of hits with high-quality candidates are discarded. After this ambiguity resolution, the remaining track candidates are then refitted using a global χ^2 method to obtain the final, high-precision track parameter estimate.

2.3.1.2 Vertices

After all track candidates have been reconstructed, the locations of the underlying pp interactions, vertices, are identified using a dedicated vertex reconstruction procedure. This is crucial as generally only one pp interaction in a given bunch crossing is relevant for the physics analysis, and its associated particles must be isolated from those produced in pile-up interactions. Since Run 2, ATLAS uses an imaging algorithm for

vertex reconstruction [80]. A three-dimensional histogram is built to image the volume where vertex candidates are expected. The reconstructed track trajectories are linearised and back-projected onto this histogram. A Fourier transform is then applied to the histogram to reveal high-density regions, which correspond to possible vertex locations. To reduce high-frequency fluctuations and noise, two filters are applied to suppress short-range noise and eliminate artefacts from detector symmetries. The filtered image is then inverse-Fourier transformed back to position space, producing a smoothed density map of potential vertices. The resulting image is then passed to a separate clustering algorithm, where all seeds are identified from peaks in the image. Finally, a vertex fitting procedure iteratively assigns tracks to the closest vertex, and the hard scatter vertex is the one with the highest scalar sum of transverse momenta squared $\sum_{\text{tracks}} p_{\text{T}}^2$.

2.3.1.3 Electrons and photons

Electrons and photons are reconstructed using information from both the ID and the ECAL [81–83]. The reconstruction begins with the identification of energy clusters in the electromagnetic calorimeter. Cells with significant energy deposits (well above noise thresholds) are selected, and a clustering algorithm iteratively grows the cluster by considering the energy and spatial proximity of neighbouring cells.

For electrons, tracks are matched to the clusters by extrapolating them into the calorimeter and comparing the expected impact points with the cluster position. If several tracks match a single cluster, further selection is performed to refine the matching, based on track quality and consistency of the track's energy with the cluster's energy deposit.

Photons do not leave tracks in the tracker unless they convert into e^+e^- pairs; they are treated differently. Photon conversions are reconstructed by identifying pairs of tracks consistent with a common conversion vertex. These vertices are then matched to calorimeter clusters using similar spatial and energy-based criteria.

Electron and photon reconstruction is performed simultaneously, which may lead to ambiguities when a cluster can be associated with either particle type. These ambiguities are resolved using additional information, such as the presence of a track with a bremsstrahlung pattern (for electrons) or a well-identified conversion vertex (for photons). Finally, photons and electrons are selected according to isolation criteria to have a good purity, and their energies are calibrated to account for detector effects.

2.3.1.4 Others

Muons leave tracks both in the MS and the ID. In the MS, a χ^2 fit is performed to associate the hits with the track. Then, MS tracks are extrapolated inward to the ID, where they are matched to ID tracks, and their compatibility with the primary interaction point is evaluated. As for the electrons and photons, isolation criteria are used to enhance purity, and a dedicated calibration is applied.

Neutrinos interact only via the weak force and therefore leave no direct signal in the detector. As a result, they are inferred from the missing transverse momentum in an event. The missing transverse momentum is estimated as the negative vector sum of the transverse momentum of all identified hard physics objects (electrons, photons, jets and

muons). To account for additional activity not associated with these objects, tracks not matched to any of them are included in a separate component known as the soft term.

2.3.2 Jet reconstruction

In ATLAS, jets are reconstructed following the procedure detailed in Section 1.2.4, based on information collected from the detector. This algorithm requires the four-momenta of particle-like objects as input. Several approaches exist in ATLAS to construct such objects: topological clusters, optionally corrected using local hadronic calibration weights, as well as Particle Flow objects. These methods are detailed in the sections below.

2.3.2.1 Topological clusters

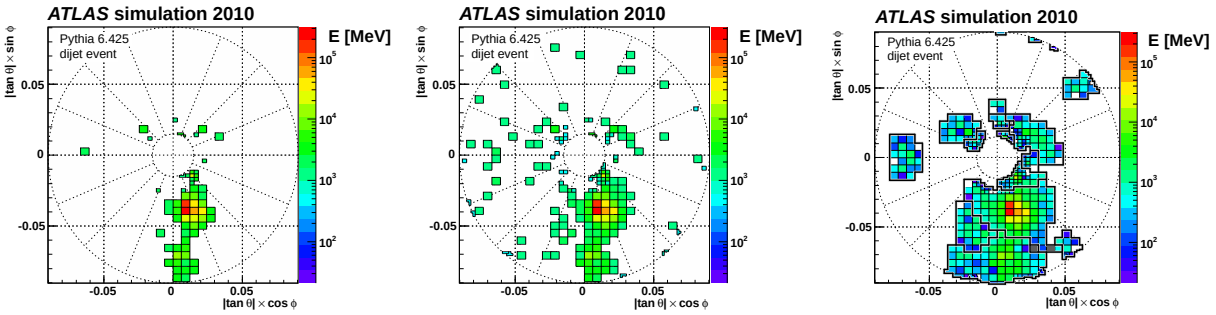
Hadrons entering the ATLAS detector deposit energy across multiple calorimeter cells. Thanks to the fine granularity of the calorimeter system, the resulting hadronic showers can be reconstructed in three dimensions. To achieve this, particle-like objects built from groups of calorimeter cells, called topo-clusters [84,85], are used. These clusters allow for the full reconstruction of the four-momentum of the particle that initiated the shower, along with detailed information about the shower shape. Additionally, the clustering algorithm suppresses contributions from electronic noise and pile-up, improving the purity of the reconstructed signal. The algorithm works as follows:

1. For each calorimeter cell, the signal significance is defined as the ratio between the measured energy and the expected average noise: $\mathcal{S}^{\mathrm{EM}}_{\mathrm{cell}} = E^{\mathrm{EM}}_{\mathrm{cell}}/\sigma^{\mathrm{EM}}_{\mathrm{noise,cell}}$. The noise includes contributions from both electronic sources and expected pile-up, the latter estimated from the average number of interactions per bunch crossing.
2. Cells with a high significance ($\mathcal{S}^{\mathrm{EM}}_{\mathrm{cell}} > 4$) are selected as seeds for the proto-clusters.
3. The cluster is expanded by including neighbouring cells which satisfy $\mathcal{S}^{\mathrm{EM}}_{\mathrm{cell}} > 2$. Two cells are considered neighbours if they are directly adjacent within a given sampling layer or, in the case of adjacent layers, if they overlap in the (η, ϕ) plane.
4. All direct neighbouring cells are included, regardless of their signal significance. Cells with negative energy are discarded.

During the cluster-building process (steps 2 and 3), proto-clusters may merge if they share common cells. To prevent proto-clusters from becoming too large to provide a good measurement of the energy flow, those with two or more local maxima, defined by $E^{\mathrm{EM}}_{\mathrm{cell}} > 500$ MeV and at least four neighbours with smaller energy, are split between the corresponding signal peaks. The splitting is performed in all three spatial dimensions. The resulting proto-cluster is characterised by a core of cells with highly significant signals, surrounded by an envelope of cells with less significant signals. The topo-cluster formation is represented in Figure 2.10.

Then, various moments of the topo-clusters are computed based on their shape and internal signal distribution:

- Geometrical moments, such as the barycentre in (η, ϕ) and the cluster direction, which can be used to reconstruct the momentum of the topo-cluster, under the



(a) (b) (c)

Figure 2.10: Representation of the different stages of topo-cluster formation. (a) Seeds (b) Neighbouring cells with significance above 2 (c) all neighbouring cells included. Taken from [84].

assumption that clusters are massless particle-like objects.

- Signal moments, which describe the distribution and quality of the signal within the cluster. These include the overall signal significance, the energy density, the cluster timing, the composition (i.e. the distribution of energy among contributing cells) and the topological isolation from nearby clusters or noise.

They provide information sensitive to the nature of the shower generating the cluster signal.

2.3.2.2 Local hadronic cell weighting

To provide a calorimeter signal suitable for physics object reconstruction without assuming the particle type, topo-clusters are calibrated using the local hadronic cell weighting (LCW) scheme. This approach consists of applying a moment-dependent, cluster-by-cluster calibration to correct several detector effects, including the non-compensating nature of ATLAS calorimeters, accidental signal losses due to the clustering strategy (in particular at the margins of the topo-clusters), and energy lost in inactive material in front, between and inside the calorimeter modules. The calibration is derived using MC simulations of single pions (charged and neutral) at various energies in the ATLAS calorimeter regions.

2.3.2.3 PFlow objects

Most particles within jets are charged and therefore leave tracks in the inner tracking detector. This tracking information can be matched to topo-clusters in the calorimeter to improve jet reconstruction. The objects resulting from this combined use of tracking and calorimeter information are referred to as particle flow objects [86], (PFlow objects or PFO).

The particle flow algorithm consists of the following steps:

1. Tracks are filtered based on quality criteria (p_{T} threshold, hit requirements).

2. Selected tracks are matched to nearby calorimeter clusters based on geometric and energy compatibility.
3. For each matched track-cluster pair, the energy-to-momentum ratio (E/p) is computed.
4. The algorithm checks if energy from one particle is spread over multiple clusters (shower splitting). If there is splitting, additional clusters in the vicinity are added to account for split showers, and E/p is recomputed, and this step is repeated.
5. A cell-based energy subtraction algorithm is used to remove overlapping contributions between the momentum measured in the inner detector and the energy measured in the calorimeters.
6. If the remaining calorimeter energy is consistent with expected shower fluctuations from a single particle, the residual topo-cluster energy is discarded.

Figure 2.11 illustrates how the particle flow algorithm deals with the different cases encountered: one or more particles (charged or not) producing one or more topo-clusters. This method helps resolve ambiguities arising when a single particle deposits energy across multiple topo-clusters or when multiple particles contribute to the same topo-cluster. It also significantly reduces the impact of pile-up by removing energy deposits associated with tracks matched to pile-up vertices. Moreover, tracking provides superior resolution at low p_{T}, where calorimeter measurements are less precise.

This procedure is applied to tracks sorted in descending p_{T} order, starting with cases where a single topo-cluster is matched to the track, and then to the remaining selected tracks. In recent implementations, tracks with transverse momentum above 100 GeV are not used, and the PFlow algorithm becomes equivalent to EMtopo, benefiting from excellent calorimeter performance at high energies.

2.3.2.4 Constituent-level pile-up rejection

The collection of PFOs consists of charged PFOs (tracks matched to clusters, with calorimeter energy subtracted) and neutral PFOs (derived from modified and unmodified topo-clusters not fully explained by charged tracks). To reduce the contribution from pile-up, the Charged Hadron Subtraction (CHS) method is employed: charged PFOs that are not associated with the primary vertex are removed, by requiring $|z_0 \sin(\theta)| < 2$ mm, where z_0 is the longitudinal impact parameter, and θ is the polar angle of the charged PFO momentum.

2.3.2.5 Jet reconstruction

Jets are reconstructed from PFlow objects ("PFlow jets") or topo-clusters ("EMTopo jets"), using the anti-k_{T} algorithm, introduced in Section 1.2.4.2. Both types of jets are referred to as "reco-jets," whether they are defined in data samples or simulations. Jets that do not satisfy $p_{\mathrm{T}} > 7$ GeV and $|\eta| < 4.5$ are discarded.

In MC samples, jets can be reconstructed from the simulated particles, using stable final-state particles (with lifetime τ satisfying $c\tau > 10mm$), excluding muons, neutrinos, and

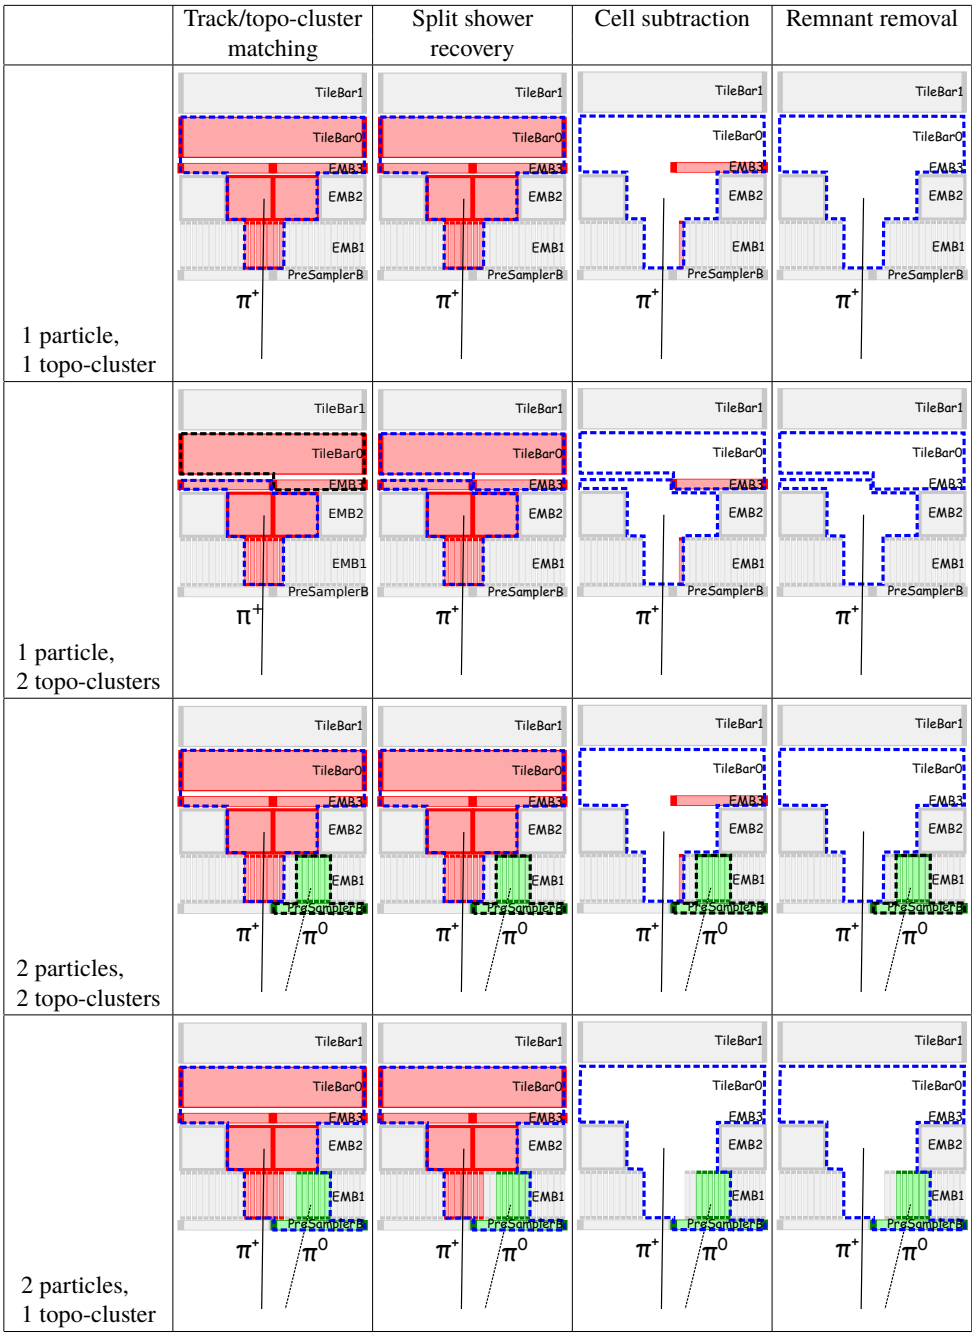


Figure 2.11: Idealised examples of how the algorithm is designed to deal with several different cases. The red cells are those which have energy from the π^+, the green cells energy from the photons from the π^0 decay, the dotted lines represent the original topo-cluster boundaries, with those outlined in blue having been matched by the algorithm to the π^+, while those in black are yet to be selected. Taken from [86].

particles from pile-up interactions. Such jets are referred to as "truth-jets". They are selected with the same $p_{\mathrm{T}} > 7$ GeV and $|\eta| < 4.5$ thresholds as EMtopo and PFlow jets, and are geometrically matched to those jets using the angular distance R with the requirement $R < 0.3$.

Afterwards, reco-jets are calibrated with the Jet Energy Scale (JES) calibration, thoroughly described in Chapter 3.1.

2.3.3 Jet Cleaning

At each stage of jet reconstruction, specific measures are taken to mitigate the effects of pile-up. For instance, CHS is applied consistently to suppress pile-up contributions in PFOs. The anti-k_{T} jet algorithm itself is inherently resilient to soft radiation. Additionally, the jet calibration procedure includes dedicated corrections to account for pile-up.

To further improve jet quality and reduce residual pile-up effects, a range of jet-level and event-level mitigation techniques has been developed.

2.3.3.1 Jet Vertex Tagger

The JVT (jet-vertex-tagger) variable [87] is a 2-dimensional likelihood that quantifies the probability for a jet of being signal-like. It is defined in the phase space: $20 < pT < 50$ GeV and $|\eta| < 2.4$. It relies on variables that capture the fraction of a jet's p_{T} originating from tracks associated with the hard-scatter vertex:

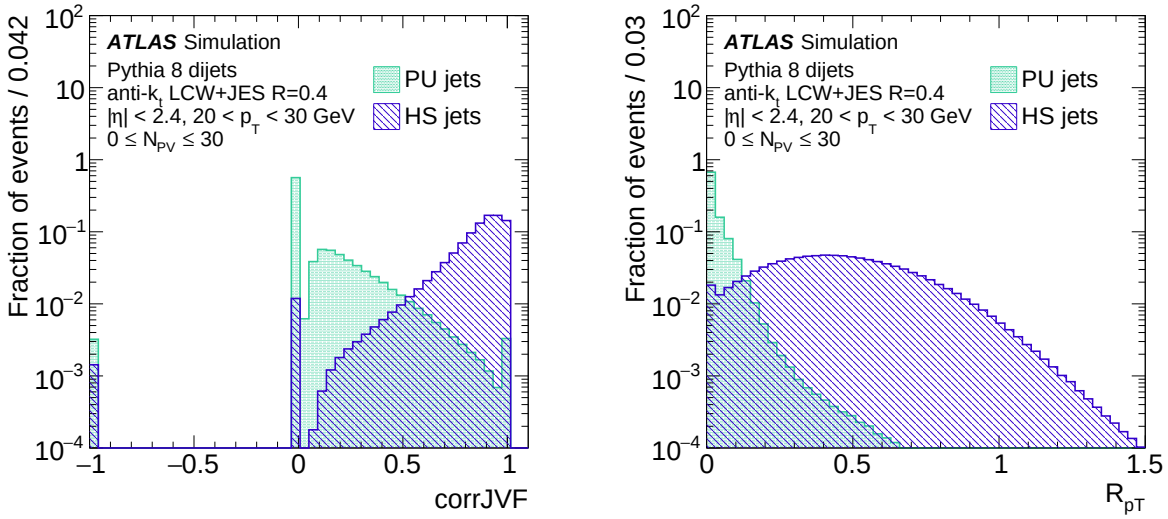
- corrJVF (corrected Jet Vertex Fraction) compares the scalar sum of p_{T} of tracks matched to the primary vertex to the total p_{T} of tracks associated with any vertex, with a correction to account for the increasing track multiplicity with the number of primary vertices N_{PV}. The Jet Vertex Fraction (JVF) was the variable used in almost all physics analyses in Run 1 for the same purpose.
- $R_{p\mathrm{T}}$ compares the scalar sum of track p_{T} associated with the primary vertex to the total calibrated jet p_{T}, which includes information from the calorimeter. $R_{p\mathrm{T}}$ can be interpreted as the fraction of the total jet p_{T} that comes from charged particles associated with the hard-scatter vertex.

The distributions of these two variables for pile-up jets and jets coming from the hard-scatter vertex are displayed in Figure 2.12. Together, these variables allow the JVT to efficiently distinguish between pile-up and hard-scatter jets, particularly in the low-to-moderate p_{T} region where pile-up contamination is most significant.

The likelihood is evaluated on a simulated $Z(\rightarrow ee)$+jets sample. In this sample, EMTopo jets geometrically matched to truth-level particles from the hard-scatter process (within $\Delta R < 0.4$) are labelled as hard-scatter jets, while unmatched jets are considered pile-up jets. The fake rate (fraction of pile-up jets misidentified as hard-scatter) versus efficiency (fraction of true hard-scatter jets retained) of the JVT, corrJVF and $R_{p\mathrm{T}}$, is shown in Figure 2.13.

2.3.3.2 Event level selection

At this point, jets from pp collisions need to be distinguished from misidentified jets of non-collision. These might arise from proton losses upstream of the interaction point. These losses can produce secondary cascades resulting in high-energy muons (with energies close to the beam energy) that may be reconstructed as fake jets. This phenomenon is referred



(a) (b)

Figure 2.12: Distribution of corrJVF (a) and $R_{p\mathrm{T}}$ (b), for pile-up jets (PU) and jets coming from the hard-scatter vertex (HS), for $20 < p_{\mathrm{T}} < 50$. Taken from [87].

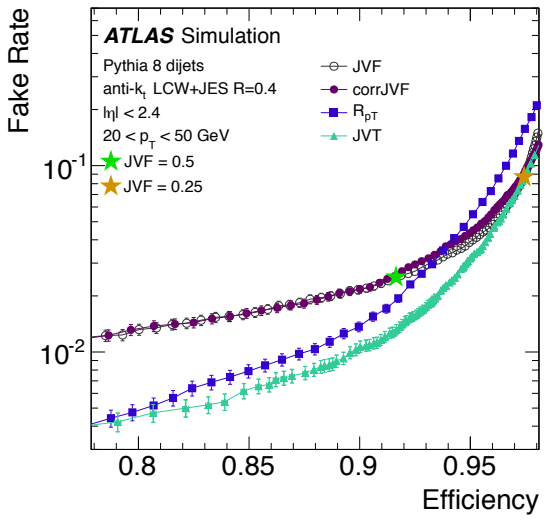


Figure 2.13: Fake rate from pile-up jets versus hard-scatter jet efficiency curves for JVF, corrJVF, $R_{p\mathrm{T}}$, and JVT, for $20 < p_{\mathrm{T}} < 50$. The widely used JVF working points with cut values 0.25 and 0.5 are indicated with gold and green stars. Taken from [87].

to as Beam Induced Background (BIB). Cosmic-ray showers produced in the atmosphere can also occasionally mimic jet-like signatures in the detector.

The selection of so-called good jets is based on several discriminating variables:

- E_{neg} the sum of all negative energy deposits in the jets,
- f_{EM} (resp. f_{HEC}) the fraction of a jet's energy deposited in the electromagnetic

(resp. HEC) calorimeter cells,

- f_{max} the maximum energy fraction in a single calorimeter layer, defined as the fraction of the jet's total energy deposited in the layer with the highest contribution;
- f_{ch} the jet charged fraction, i.e. the ratio of the scalar sum of the p_{T} of the tracks associated to the jet and the total jet p_{T}. The ratio of $f_{\text{ch}}/f_{\text{max}}$ is also used to identify good jets.

In addition, fake-jets can arise from detector-related effects, such as calorimeter noise originating from large-scale coherent noise or isolated pathological cells not flagged by the data quality monitoring system [88]. To reject such jets, the characteristic ionisation signal shape recorded in each LAr calorimeter cell is compared to the expected pulse shape from simulations of the electronic response. The quadratic difference of the two is computed as $Q_{\text{cell}}^{\text{LAr}}$, and several jet-level quantities are derived, such as:

- $f_{\text{Q}}^{\text{LAr}}$ the fraction of the energy in the LAr calorimeter cells with poor signal shape quality, as identified by a threshold on $Q_{\text{cell}}^{\text{LAr}}$;
- $f_{\text{Q}}^{\text{HEC}}$ the same variable defined for the HEC calorimeter cells.
- $\langle Q \rangle$ the average jet quality defined as the energy–squared weighted average of $Q_{\text{cell}}^{\text{LAr}}$, across all relevant cells in the jet.

The E_{neg} variable is also used to flag fake jets from calorimeter noise.

The exact selection criteria are detailed in [89]. Their efficiency has been studied both in simulation and using in situ data. Several working points are defined to offer different compromises between selection efficiency and fake jet rejection.

2.3.4 MC simulations

In ATLAS, MC simulations are used at various stages of an analysis, including for the calibration of detector response, the modelling of signal and background processes, the optimisation of selection criteria, and the unfolding of detector effects to correct reconstructed data to particle or parton level. These simulations play a central role in ensuring the robustness and accuracy of the final physics measurements. Hard scatter collisions are emulated with general-purpose Monte Carlo simulators (GPMC, described in Section 1.2.3.1). In order to accurately reflect the conditions of pp collisions at the LHC, those samples have to be overlaid with softer pp collisions, representing pile-up. This is typically achieved by superimposing zero-bias pile-up events simulated with Pythia 8, following the N_{PV} distribution. The resulting events, combining both hard scatter and pile-up, are then passed through a detailed simulation of the detection process, accurately modelling the geometry, materials, and digitisation procedures of the detector. This is performed using the Geant4 [90] simulation toolkit. The information on the particles available before the detector simulation is denoted as "truth-level information" whereas detector information available after simulation by Geant4 is referred to as "reco-level information".

During data-taking, the observed distribution of the average number of pile-up interactions may differ from the one simulated during the pile-up overlay step in simulations.

To account for this discrepancy, simulated events are reweighted during data analysis to match the pile-up distribution observed in data. This procedure is referred to as pile-up reweighting.

2.3.5 The ATLAS data flow

During Run 2, ATLAS recorded approximately 18 billion complete events. To support physics analyses, 2–3 times more simulated (Monte Carlo) events are typically produced than real detector events. This large volume of simulation is essential to reduce statistical uncertainties in predictions and to enable studies of potential new physics signals.

ATLAS computing relies on a globally distributed infrastructure composed of about one million CPU cores, 350 PB of disk storage, and nearly 450 PB of tape storage. This infrastructure includes the CERN Tier-0 site, responsible for prompt data processing, as well as the Worldwide LHC Computing Grid (WLCG), which includes 11 ATLAS Tier-1 sites and around 70 Tier-2 sites. Additional resources include high-performance computing centres, the high-level trigger processing farm located near the detector to minimise latency, and commercial cloud services. These distributed resources are collectively used to process and analyse all real and simulated data produced by the ATLAS experiment. ATLAS uses an extensive software suite [91,92] for data simulation, for the reconstruction and analysis of real and simulated data, for detector operations, and for the trigger and data acquisition systems of the experiment.

The processing of real and simulated data in ATLAS involves multiple steps. Once events pass the trigger selection, detector data are recorded in the RAW format, with each event typically around 1.6 MB. In the case of simulation, the energy deposits are passed through a digitisation step, where the detector electronics are modelled. This results in Raw Data Objects (RDO), which conceptually resemble the RAW data and are slightly larger (typically around 2 MB per event)[1].

Then, both real and simulated events undergo the same reconstruction and analysis chain. Physics objects are reconstructed, and events are stored in the Analysis Object Data (AOD) format, which typically ranges between 0.2 and 0.5 MB per event. As a result, the total size of full AOD datasets can reach about 1 PB per year, making them too large for direct use in most analyses. To address this, the data are centrally reduced into Derived Analysis Object Data (DAOD) formats, which are tailored to the needs of specific analyses or physics groups. Each DAOD contains only the relevant object selections, calibrations, and variables, with a typical event size of 50–150 KB depending on the analysis. Finally, individual analysis teams often produce Ntuples, which contain only the minimal subset of information required for their specific studies. These Ntuples are typically a few tens of gigabytes in size and allow for efficient and flexible analysis.

[1] The exact size of these data formats can vary from year to year, depending on pile-up conditions and improvements in software and data reduction.

Chapter 3

The Jet η-intercalibration

Introduction

At high-energy hadron colliders like the LHC, jets have a production cross-section that exceeds most other final-state observables by several orders of magnitude, making them a crucial element in most analyses, whether as signal or background. As such, they must be accurately reconstructed, with their energy and direction (i.e. four-vector) precisely measured at the particle level.

The goal of calibration is to ensure that, on average, the reconstructed energy of the jet matches the true energy of the corresponding jet of stable particles that entered the detector, in real data or Monte Carlo simulations. To achieve this, several effects must be accounted for and corrected, such as:

- The energy lost in the upstream material (before the calorimeter),
- The energy lost in dead material (beam pipe, structure, cryostat, electronics...),
- The non-compensating nature of ATLAS calorimeters, i.e. the fact that the response of the calorimeters is different for hadrons and electromagnetic objects (electrons or photons),
- The bending of charged particles in the magnetic field,
- The busy data-taking environment resulting from multiple proton-proton interactions (pile-up),
- The possible biases in the jet η reconstruction, caused by the transition regions between different calorimeter technologies or layouts. Pseudo-rapidity is used rather than rapidity because it directly reflects the detector geometry and segmentation. For small-R jets, the jet mass is typically small enough that the difference between pseudorapidity and rapidity is negligible.

The calibration must ensure good response linearity, energy resolution, and position reconstruction. Additionally, it must provide a uniform response across the (η, φ) plane, particle compositions and energy distributions. Since jets are complex objects, the uncertainty in

their energy measurement is the dominant experimental uncertainty for numerous physics results. Therefore, it needs to be refined to allow more precise physics measurements.

There are multiple sources of systematic uncertainties, arising from both the detector description and the physics modelling. These must be combined with statistical uncertainties, which stem from the finite amount of data.

This chapter deals with the η-intercalibration, one of the last steps of the jet calibration chain. First, an overview of the current status of the Jet Energy Scale (JES) in ATLAS is given. Then, the η-intercalibration method is detailed, and its derivation as part of the pre-recommendation for the ATLAS collaboration is presented. Finally, the various improvements made to the method are discussed.

3.1 Current status of jet calibration in ATLAS

The JES calibration involves multiple consecutive stages derived from a combination of methods based on the precise comparison between data and Monte Carlo simulations and in situ techniques. The aim is to determine correction factors to apply to the jet four-vector, depending on its pseudorapidity and transverse momentum, to match the four-momentum of the stable particles within the jet. The set of corrections is illustrated in Figure 3.1.

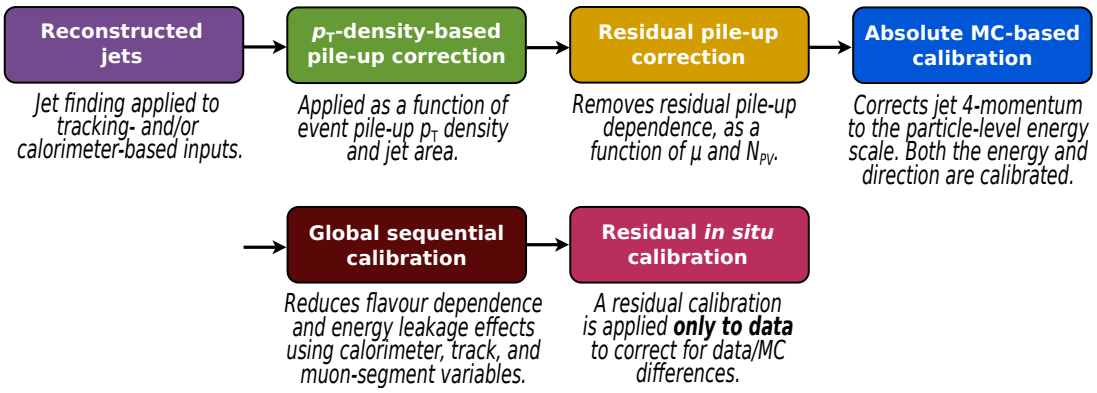


Figure 3.1: Overview of the 2015 ATLAS calibration scheme for EM-scaled jet, at $\sqrt{s} = 13$ TeV. Taken from [93]. Some calibrations, not included in this flow chart, are dedicated to specific phase space regions.

First, two sets of pile-up corrections are derived to remove the extra energy coming from additional proton-proton collisions (Section 3.1.2). Then, MC-based calibration steps correct the reconstructed jet four-momentum to match that derived from the particles within the jet (Section 3.1.3) and improve the jet p_{T} resolution and uncertainty by removing the dependence of the jet response on selected variables (Section 3.1.4). In situ techniques are then used to measure the difference in jet response between data and Monte Carlo simulation, with residual corrections applied to jets in data only (Section 3.1.5).

3.1.1 Reconstruction and origin correction

For most ATLAS analyses as well as for the derivation of the calibration, jets are reconstructed using the anti-k_{T} algorithm [43] with a radius parameter $R = 0.4$, commonly referred to as small-R jets. Details on the anti-k_{T} algorithm can be found in Section 1.2.4.2. The inputs to the algorithm are four-vector objects, which can represent various observables derived from detector information or stable particles in Monte Carlo simulations.

For example, jets can be reconstructed from topo-clusters [84], which are 3D clusters of noise-suppressed calorimeter cells, described in Section 2.3.2.1. The topo-clusters are calibrated at the electromagnetic scale (EM-scale), which is the basic signal scale accounting for the energy deposited in the calorimeter by an electromagnetic shower. Jets reconstructed from topo-clusters are called EMTopo jets. Topo-clusters can be further calibrated using the local cell weighting (LCW) method, which is designed to provide the correct energy scale for charged pions originating from the interaction point. This method reduces energy fluctuations caused by the non-compensating nature of the ATLAS calorimeter, as well as losses due to energy deposition outside the cluster and in dead material.

Another approach to jet reconstruction consists of incorporating information from the tracking system. Particle Flow Objects (PFlow) [86] are created by matching energy deposits in the calorimeter to extrapolated tracks from the tracking system. They are described in Section 2.3.2.3. The final set of constituents is a combination of tracks and energy clusters. Jets reconstructed from PFlow jets demonstrate improved performance and greater stability to pile-up at low p_{T} compared to EMTopo jets. In either case, jets follow the same calibration scheme afterwards.

Jets must meet specific quality criteria designed to reject fake jets, described in Reference [93]. In Monte Carlo simulations, the reco-jets (reconstructed from detector outputs) are geometrically matched to a truth-jet (reconstructed from the four-vector of the stable particles), such that the angular distance between the two $\Delta R = \sqrt{\Delta\varphi^2 + \Delta\eta^2}$ is smaller than 0.3. Then, a jet origin correction is performed to correct the direction of the jet to match the primary vertex instead of the centre of the detector, used in the initial jet reconstruction.

3.1.2 Pile-up corrections

There are two types of pile-up: in-time pile-up, which results from additional proton-proton collisions occurring in the same bunch-crossing as the collision of interest, and out-of-time pile-up, which comes from proton-proton collisions in nearby bunch-crossings. Pile-up events deposit additional energy in the detector, which combines with that from hard-scattering objects, leading to a bias in the total jet energy and degrading the resolution of hard-scatter jets. The main goal of the pile-up corrections is to estimate and subtract the contribution from pile-up to the object energy and filter out objects that originate from pile-up events. For that purpose, two corrections are derived.

The first correction is a jet area-based pile-up correction. For each jet, an active area A^{jet} [94] representing the jet susceptibility to contamination is computed. The calculation of A^{jet} is done by ghost-association: a large number of soft particles (called ghosts) is

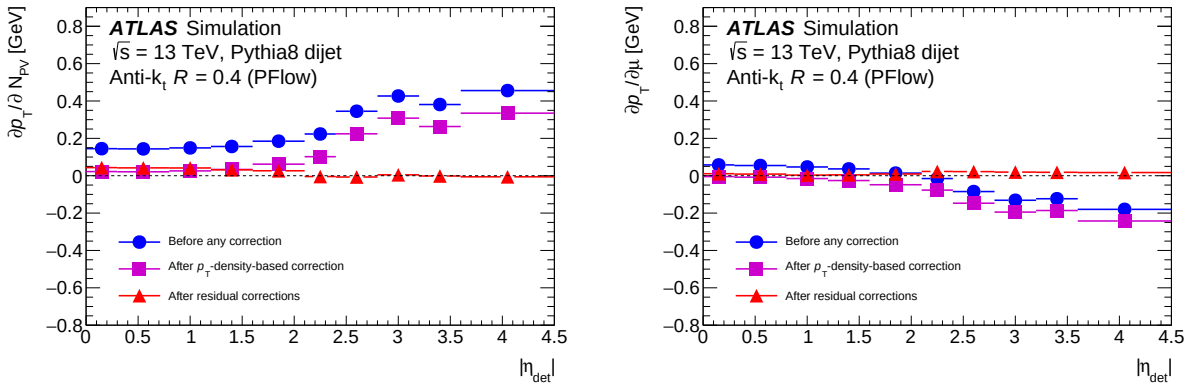
added, and the region in the (η, φ) plane where the ghosts are clustered into the jet is assessed. The pile-up contribution is estimated from the median p_{T} density of jets in the (η, φ) plane: $\rho = \mathrm{median}(p_{T,i}^{\mathrm{jet}}/A_i^{\mathrm{jet}})$. A correction proportional to $A * \rho$ is then removed from the jet p_{T}. This method is applied event-by-event, jet-by-jet and is described in Reference [95].

However, it is observed that the corrected jet p_{T} still exhibits a small residual dependence on pile-up activity. To address this, a residual correction is derived. This second correction is based on Monte Carlo simulations, where the difference between the calorimeter-jet p_{T} and the associated truth-jet p_{T} is evaluated as a function of N_{PV} and $\langle\mu\rangle$. Here, N_{PV} denotes the number of reconstructed primary vertices in the event, which is sensitive to out-of-time pile-up, while $\langle\mu\rangle$ represents the mean number of interactions per bunch crossing, which is sensitive to in-time pile-up.

The jet p_{T} after the two corrections is given by:

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

The coefficients α and β are derived with independent linear fits in bins of $p_{\mathrm{T}}^{\mathrm{true}}$ and $|\eta_{\mathrm{det}}|$, where $p_{\mathrm{T}}^{\mathrm{true}}$ is the p_{T} of the truth-jet that matches the calorimeter-jet. The effect of the various pile-up corrections is illustrated in Figure 3.2.



(a) (b)

Figure 3.2: Dependence of PFlow jet p_{T} on pile-up as a function of $|\eta_{\mathrm{det}}|$ for $p_{\mathrm{T}}^{\mathrm{true}} = 25$ GeV, for (a) in-time pile-up and (b) out-of-time pile-up. Errors are taken from the fit results and are too small to be visible on the scale of the plot. Taken from Reference [93].

Three uncertainties are derived to account for the modelling of μ, N_{PV} and ρ in MC simulations. An additional uncertainty is included to cover the residual p_{T} dependence.

3.1.3 MC absolute calibration (MCJES)

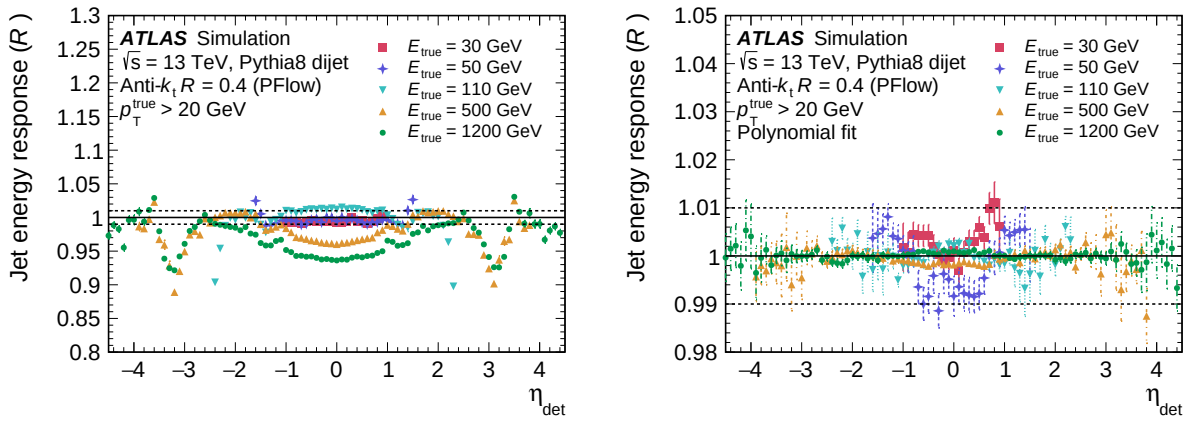
At this point, a baseline jet energy scale calibration is applied to correct the calorimeter-jet energy to that of the truth-jet, in MC-simulated dijet samples. Here, both energy and direction are calibrated.

The jet energy response $\mathcal{R}_E$, defined as

$$\mathcal{R}_E = \left\langle \frac{E_{\text{reco}}}{E_{\text{true}}} \right\rangle$$

is calculated in bins of η_{det} and E_{true}, with E_{reco} the energy of the reco-jet and E_{true} the energy of the matching truth-jet. The brackets indicate that $\mathcal{R}_E$ is defined as the mean parameter of a Gaussian fit to the distribution. Then, a calibration function is determined for each η_{det} bin by fitting a smooth function to a graph of $\mathcal{R}_E$ vs E_{reco}, as described in Reference [96].

Jet energy response before and after MC-based jet calibration is shown in Figure 3.3 as a function of E_{true} and η. Small non-closure effects (deviation from unity) after application of the correction factor are seen at low-p_{T}, implying that the kinematic observables of the calibrated calorimeter-jet are not fully restored to those of the corresponding truth-jet.



(a) (b)

Figure 3.3: The jet energy response before (a) and after (b) the MCJES calibration, at fixed energies as a function of η_{det}. Taken from Reference [97].

For the jet pseudorapidity, an additive correction is included to account for biased reconstruction due to changes in the detector technology in some regions of the detector. In the poorly instrumented regions, the topo-clusters are reconstructed with lower energy, which entails a bias in the jet direction toward the better-instrumented regions. The correction is derived as the average difference between η_{true} and η_{origin}, in E_{true} and η_{det} bins, and parameterised as a function of the calibrated jet energy $E_{\text{reco}}^{\text{EMJES}}$ and the uncorrected η_{det}. The overall correction is very small, and larger in the transition regions, as shown in Figure 3.4.

3.1.4 Global sequential calibration (GSC)

To reduce the jet-by-jet response dependence, the Global Sequential Calibration scheme (GSC) is performed. The jet topology, and therefore the detector response, fluctuates according to the parton at the origin of the jet, i.e. whether it comes from a gluon, a light quark or a heavy quark. For example, a gluon jet will typically contain a higher multiplicity of particles with softer p_{T} than quark jets, which entails a lower calorimeter response

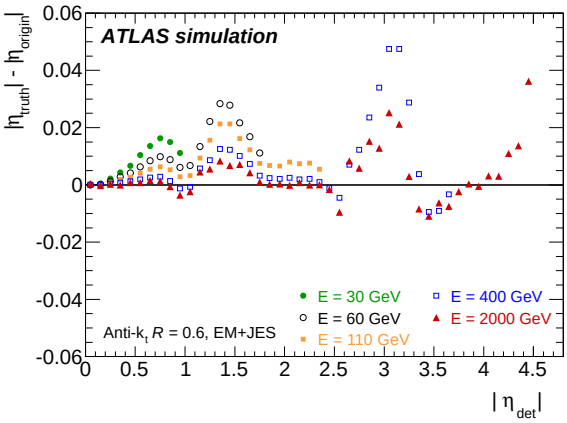


Figure 3.4: Average difference between η_{true} and η_{origin}, in bins of E^{true} and η_{det}, as a function of the uncorrected η_{det}. This view highlights the transition regions where the detector is poorly instrumented. Taken from Reference [96]

and a wider transverse profile. The GSC consists of six multiplicative corrections, exploiting observables that describe the longitudinal shower structure, its transverse structure, or the activity in the muon chambers. The observables, in the order in which they are derived and applied, are the following:

- f_{charged}: The fraction of the jet transverse momentum carried by charged particles, calculated using ghost-associated tracks with $p_{\text{T}} > 500$ MeV within a detector pseudorapidity range of $|\eta_{\text{det}}| < 2.5$.
- f_{Tile0}: The fraction of the jet energy measured in the first layer of the hadronic tile calorimeter, constrained to $|\eta_{\text{det}}| < 1.7$.
- f_{LAr3} The fraction of the jet energy measured in the third layer of the electromagnetic LAr calorimeter, considering $|\eta_{\text{det}}| < 3.5$.
- n_{trk}: The number of tracks with $p_{\text{T}} > 1$ GeV that are ghost-associated with the jet, within $|\eta_{\text{det}}| < 2.5$.
- w_{trk}: The average transverse distance, weighted by p_{T}, between the jet axis and all ghost-associated tracks with $p_{\text{T}} > 1$ GeV in the $\eta - \varphi$ plane, for $|\eta_{\text{det}}| < 2.5$.
- n_{segments}: Refers to the number of muon track segments ghost-associated with the jet, defined within $|\eta_{\text{det}}| < 2.7$. This corrects the effect of high-p_{T} jets that are not fully contained in the calorimeter, and it is also known as the "punch-through" correction.

f_{LAr3} and f_{Tile0} serve as indicators of the jet penetration depth, while n_{trk} and w_{trk} act as indicators of the jet particle multiplicity and their transverse profile. The n_{segments} correction reduces the tails of the response distribution caused by high-p_{T} jets that are not fully contained in the calorimeter. Correction factors are then derived by a procedure similar to that of the MCJES, in bins of η and p_{T}, for each variable successively. The process is further explained in Reference [93].

The mean response in the dijet sample where these correction factors are derived remains

unchanged, but the jet energy resolution (JER) improves as corrections are added and flavour dependence is reduced, as shown in Figure 3.5. The GSC is validated in situ

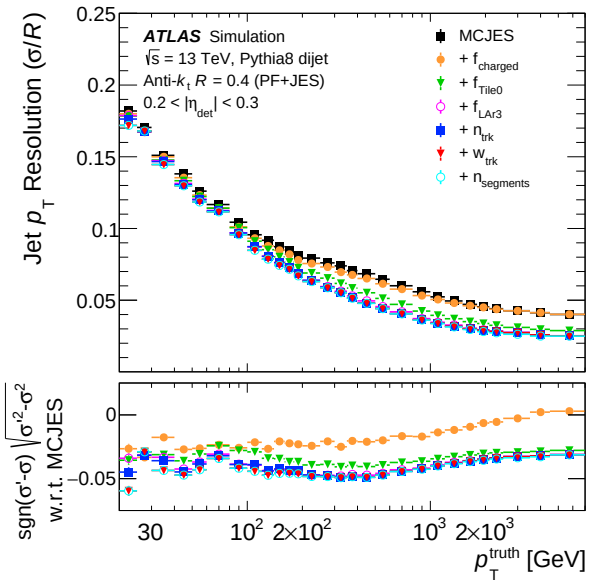


Figure 3.5: Jet Energy Resolution after MCJES and after the successive steps of the GSC. The JER is improved at each step. Taken from Reference [93].

with dijet events: the average p_{T} asymmetry between back-to-back jets is measured as a function of each observable, to check that the dependence of the jet response on each variable is reduced.

One of the recent developments in jet calibration has been to replace this sequential calibration with a Deep Neural Network (DNN) trained to predict the p_{T} response. This method, called the Global Neural Network Calibration (GNNC), exploits the correlations between the variables used and therefore allows using more of them. In particular, calorimeter-based variables, such as the fraction of energy measured in various sub-parts of the calorimeter system, and jet kinematics variables, are used in addition to the GSC variables. The GNNC is used in the latest versions of the calibration and is thoroughly described in Reference [97]. It shows an improved p_{T} resolution (about 15-25%) on average compared to baseline GSC.

Five uncertainties are derived to cover the effects of the residual dependence of the calorimeter response on jet flavour:

- Proportional sample composition of quarks and gluons,
- Response of gluon-initiated jets,
- Response of b-quark-initiated jets,
- Punch-through correction
- Single-particle response

3.1.5 Residual in situ calibration

The in situ calibrations aim to compare data and Monte Carlo with a data-driven approach and to correct for the residual differences between the two using reference events. These

differences can be caused by imperfect simulation of both the detector materials and the physics processes involved: the hard scatter and underlying event, jet formation, pile-up, and particle interactions with the detector. Moreover, modelling the response to hadrons is complicated by the wide range of physics processes involved in hadronic interactions and the non-compensating design of the ATLAS calorimeters.

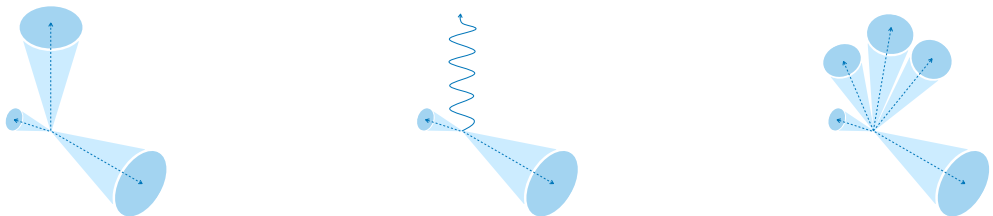
The in situ calibration relies on the "tag and probe" method, which consists of exploiting the transverse momentum balance between a well-calibrated object (tag, or reference) and the remaining jet under study (probe). The reference object may be another jet from a well-calibrated region (η-intercalibration, described in Section 3.2), a Z or a γ (Z/γ+jet method, described in Section 3.1.5.2), or multiple jets (Multi Jet Balance, or MJB, described in Section 3.1.5.3). Figure 3.6 shows the event topologies for the three methods. Calibration factors $c(p_{\mathrm{T}}, \eta_{\mathrm{det}})$ are derived from the detector response both in Monte Carlo simulations and data, then the correction factors are taken as the ratio of the two, and applied only to data:

$$\mathcal{C} = \Big[c(p_{\mathrm{T}}, \eta_{\mathrm{det}})\Big]_{\mathrm{data}} \Big/ \Big[c_(p_{\mathrm{T}}, \eta_{\mathrm{det}})\Big]_{\mathrm{MC,reco}}$$

This way, the data is corrected to match the Monte Carlo.

Another implementation of the tag and probe method consists in using the truth-level information in the Monte Carlo simulation: the calibration factor c_{MC} is calculated at truth-level and used as the denominator of the correction factor $\mathcal{C}$. It allows the deconvolution of physics and detector effects in the calibration coefficient and therefore reduces the modelling uncertainties. It is called the absolute implementation and was studied in the case of the η-intercalibration by Louis Ginabat [98].

The three in situ analyses are performed sequentially, and therefore systematic uncertainties are propagated from each to the next. The in situ chain is presented in Figure 3.7 It starts with the η-intercalibration, which aims at making the calibration homogeneous in pseudorapidity across the whole detector. Then, the Z/γ+jet and MJB calibration, which focus on the central part of the detector, are derived.



(a) (b) (c)

Figure 3.6: Topologies of events for the tag and probe techniques of the in situ calibration step: (a) dijet events for the η-intercalibration, (b) Z+jet or γ+jet events, (c) Multijet system for the multi-jet balance. In all these topologies, additional radiations can affect the balance between the probe jet and the reference object.

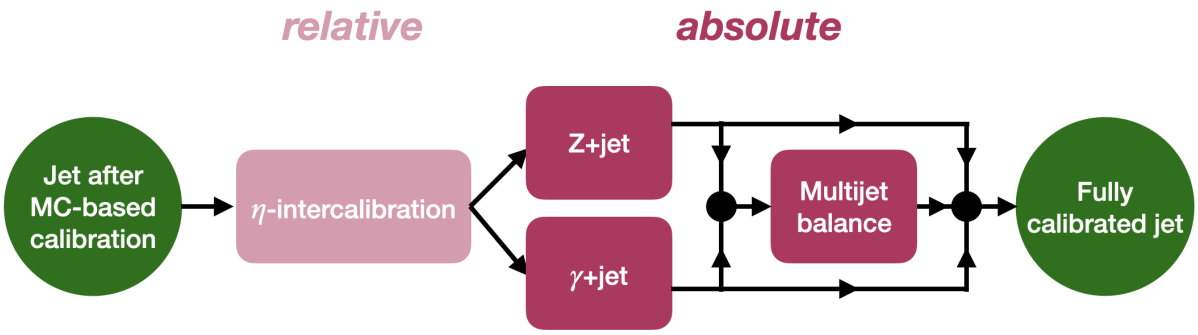


Figure 3.7: In situ calibration chain. Reproduced from [98].

3.1.5.1 η-intercalibration

The purpose of this method is to investigate the jet response across the detector in terms of pseudorapidity and remove any residual dependence, as variations in the detector structure lead to different jet responses across its various regions. This variation arises from differences in the technologies employed in different sections, gaps between instrumented regions, or the reduced tracking capabilities at high η. The η-intercalibration method is also referred to as the relative in situ correction since it calibrates regions of the detector relative to others. The other in situ methods are derived in the central region only, and the JES is propagated thanks to the η-intercalibration. The method is the subject of Section 3.2, and will therefore not be detailed here.

3.1.5.2 Z/γ+ jet method

The next step of the in situ calibration uses the p_{T} balance between a jet and a well-calibrated photon or Z boson (decaying to two electrons or two muons) to derive calibration factors in the central region of the detector ($|\eta| < 0.8$). Electrons and muons, coming from Z boson decays, and photons are well measured in the ATLAS detector, and the uncertainties in their respective energy scale are small [82] [99]. The Z+jet response measurement suffers from low statistics in the high-p_{T} region, thus, it covers the low/medium p_{T} range, from 17 GeV to 1 TeV. On the contrary, the γ+jet response measurement benefits from higher statistics at higher p_{T} and therefore can extend to 1.2 TeV, but starts at 25 GeV due to jet trigger prescales and prevalence of soft jet misidentified as photons.

Two methods can be used to derive the calibration. The Direct Balance method (DB) directly exploits the p_{T} balance between the two objects by deriving the calibration factor c in bins of $p_{\mathrm{T}}^{\mathrm{ref}}$ such that

$$\frac{1}{c} = \langle \mathcal{R}_{\mathrm{DB}} \rangle = \left\langle \frac{p_{\mathrm{T}}^{\mathrm{jet}}}{p_{\mathrm{T}}^{\mathrm{ref}}} \right\rangle$$

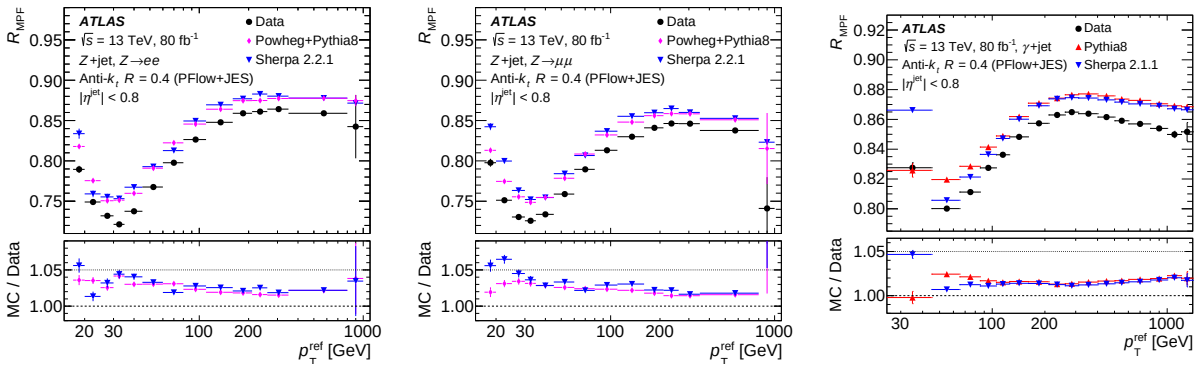
with $\langle \mathcal{R}_{\mathrm{DB}} \rangle$ being the mean of the R_{DB} distribution, evaluated using a fit of a modified

Poisson distribution. The DB method is sensitive to p_{T} imbalances from out-of-cone radiation (OOC), initial or final state radiations (ISR, FSR) and pile-up.

The Missing E_{T} Projection Fraction (MPF) method [96, 100] balances the boson against the full hadronic recoil measured with Particle Flow Objects. Assuming $\vec{p_{\mathrm{T}}}^{\mathrm{ref}} + \vec{p_{\mathrm{T}}}^{\mathrm{recoil}} + \vec{p_{\mathrm{T}}}^{\mathrm{miss}} = \vec{0}$, the calibration factors c are then defined in bins of $p_{\mathrm{T}}^{\mathrm{ref}}$ such that:

$$\frac{1}{c} = \langle \mathcal{R}_{\mathrm{MPF}} \rangle = \left\langle 1 + \frac{\hat{n}^{\mathrm{ref}} \cdot \vec{p_{\mathrm{T}}}^{\mathrm{miss}}}{p_{\mathrm{T}}^{\mathrm{ref}}} \right\rangle$$

with $\hat{n}^{\mathrm{ref}}$ being the unit vector in the direction of the reference object. Again, $\langle \mathcal{R}_{\mathrm{MPF}} \rangle$ is the mean of the R_{MPF} distribution evaluated with a fit of a modified Poisson distribution. It is displayed in Figure 3.8 as a function on the reference p_{T}, for $Z(\rightarrow ee)$ + jet, $Z(\rightarrow \mu\mu)$ + jet and γ+jet events.



(a) (b) (c)

Figure 3.8: $\langle \mathcal{R}_{MPF} \rangle$ for PFlow jets as a function of reference p_{T} for (a) $Z(\rightarrow ee)$ + jet events, (b) $Z(\rightarrow \mu\mu)$ + jet events and γ+jet events (c). Taken from Reference [93].

The MPF method is robust against both pile-up and underlying event, as their contributions each cancel out directionally on average over a large collection of events. It is also less sensitive to out-of-cone radiation since the jet definition only intervenes in the calibration application. However, applying calibration factors derived from the hadronic recoil to jets of finite radius is not a straightforward process and requires careful consideration to ensure accuracy.

Additionally, two corrections are derived from MC: the "showering correction" and the "topology correction". The showering correction quantifies the migration of energy across the jet boundaries in calorimeter jets, relative to the matching truth-jets. The topology correction accounts for the differences in the calorimeter response for different types of energy deposits, whether they be sparse or dense, like in the core of jets.

The MPF method is found to be more stable compared to the DB method for low p_{T} jets and is therefore used as a baseline.

3.1.5.3 Multijet balance

The Multijet Balance method (MJB) is used to derive a correction for jets with very high p_{T}. For this purpose, high p_{T} jets are calibrated against a system of lower p_{T} jets,

which are well-calibrated thanks to the previous in situ methods. This is done in two successive iterations in order to extend the p_{T} reach. The first iteration focuses on jets with $p_{\mathrm{T}} < 1.8$ TeV, with jets with $p_{\mathrm{T}} < 1.2$ TeV considered for the multi-jet system. In the second iteration, the p_{T} range is extended up to 2.4 TeV, enabling jets from the multi-jet system to have p_{T} values between 1.2 and 1.8 TeV. The exact threshold depends on $\sqrt{s}$ and on the luminosity of the data sample, and the values quoted here are those of Reference [93]. Those jets are calibrated with the calibration factors from the first iteration. The response is defined as:

$$\mathcal{R}_{\mathrm{MJB}} = \left\langle \frac{p_{\mathrm{T}}^{\mathrm{jet,lead}}}{p_{\mathrm{T}}^{\mathrm{multi\text{-}jet}}} \right\rangle$$

where $p_{\mathrm{T}}^{\mathrm{jet,lead}}$ is the transverse momentum of the probe jet and $p_{\mathrm{T}}^{\mathrm{multi\text{-}jet}}$ the vectorial sum of lower p_{T} jets. The mean response is obtained in bins of $p_{\mathrm{T}}^{\mathrm{multi\text{-}jet}}$ and mapped to $p_{\mathrm{T}}^{jet,lead}$. The response $\mathcal{R}_{\mathrm{MJB}}$ is illustrated on Figure 3.9.

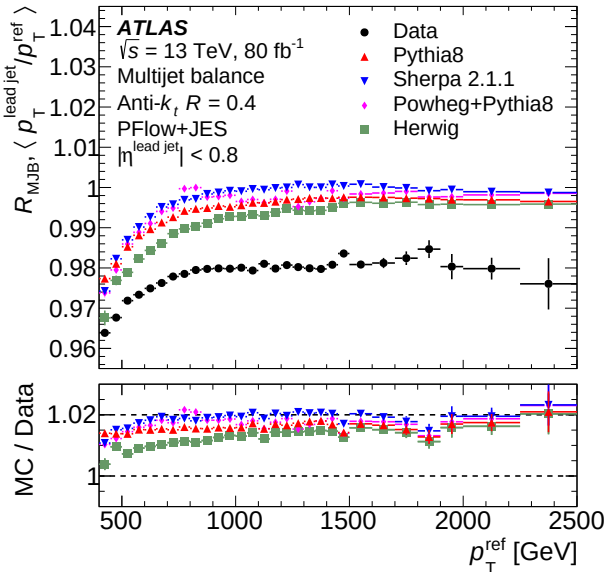


Figure 3.9: Detector response as defined in the MJB method, as a function of $p_{\mathrm{T}}^{\mathrm{ref}}$, for data and various Monte Carlo generators. Taken from Reference [93].

3.1.5.4 Systematic uncertainties

For each in situ method, uncertainties are derived as a function of η and p_{T} to account for several effects.

- Mis-modelling of physics in Monte Carlo, for which an uncertainty component is defined as the difference in calibration obtained using two different Monte Carlo generators.
- Event topology dependence on the momentum balance of events is taken into account by deriving the calibration on a dataset obtained by varying the selection cuts and comparing it to the nominal calibration.

- Uncertainties associated with the reference object (electrons, muons, and photons) are propagated by varying each object calibration by $\pm 1\sigma$ of their uncertainty, the calibration is derived and compared to the nominal calibration.
- Statistical limitations of the data and Monte Carlo datasets.
- Non-closure, for which an uncertainty can be evaluated for each method by deriving the calibration on a calibrated dataset, and taking the difference to 1.

A re-binning procedure is applied to each systematic uncertainty to enhance statistical significance by avoiding fluctuations from small event counts. This is done only where the response does not vary sharply with p_{T} to preserve real physics effects. Then, a smoothing is performed to prevent the uncertainties from being too impacted by statistical fluctuations.

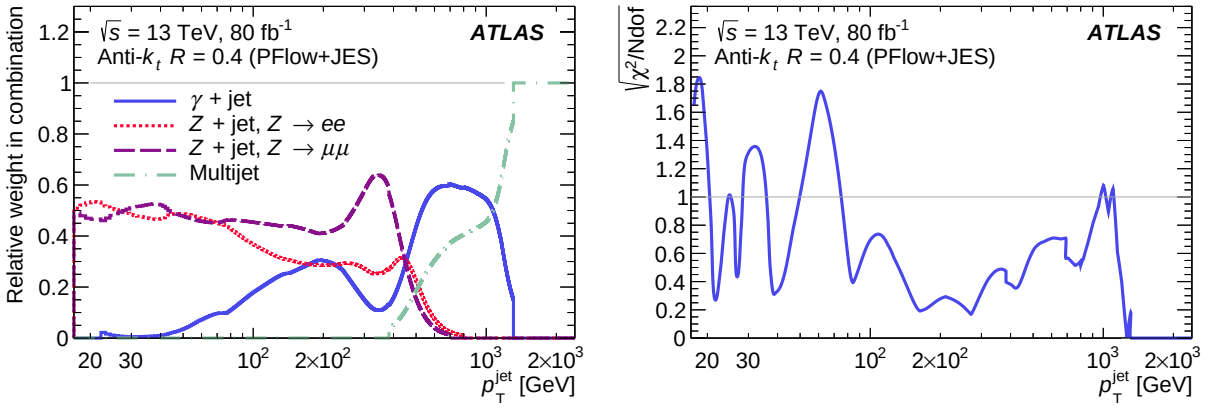
3.1.5.5 Combination of the in situ measurements

The various absolute in situ response ratios have to be combined to derive a single calibration covering the whole jet p_{T} range, from 17 GeV to 2.4 TeV. Each calibration ranges over a different p_{T} range, but those ranges can overlap, and the combination procedure must handle the possible tensions in measurement and account for their relative uncertainties. In the combination, the $Z(\rightarrow ee)$+jet and $Z(\rightarrow \mu\mu)$+jet channels are treated separately since they are affected by different systematic uncertainties.

The procedure is as follows. For each calibration, the parameterisation is adjusted to be based on the jet p_{T} rather than on the reference object p_{T}. An interpolation by a second-order spline is then used to obtain a smoother calibration in 1 GeV bins. A χ^2 minimisation is performed in each bin, using the measurements and their associated uncertainties as inputs. This minimisation acts as a weighted average, where the weights are determined by the uncertainties of the measurements. As a result, the measurement with the smallest uncertainty has the greatest influence on the estimate of the response ratio within that bin. The weights assigned to the different techniques are shown in Figure 3.10a. It shows that for p_{T} below 500 GeV, the Z+jet method dominates, then the γ+jet prevails at intermediate p_{T} up to 1000 GeV, where the MJB method dominates.

To account for tensions between the input measurements, the value of $\sqrt{\chi^2/\mathrm{NDoF}}$ is used to scale the uncertainty when it exceeds 1, according to the PDG guidelines [11]. This adjustment ensures that the uncertainties reflect the level of consistency between the measurements. The $\sqrt{\chi^2/\mathrm{NDoF}}$ metric can be seen in Figure 3.10b. It demonstrates that the different processes yield compatible results in the overlapping p_{T} regions. The final calibration curve is determined by smoothing the output from the combination with a Gaussian kernel, to minimise the influence of statistical fluctuations, ensure continuity and avoid artificial fluctuations. The result is shown in Figure 3.11. Starting from 2.4 TeV and above, the calibration factors are kept constant.

Since the three in situ analyses are performed sequentially, systematic uncertainties are propagated from each to the next. For each method, the measured response is varied by 1σ in the uncertainty component, and the propagation procedure is repeated using the weight of the nominal combination to prevent the varied uncertainty from altering the



(a) (b)

Figure 3.10: (a) Weights assigned to the different absolute in situ methods in the combination process. (b) $\sqrt{\chi^2/\mathrm{NDoF}}$ metric obtained during the combination. It represents the tension between the different measurements and is used to rescale the uncertainties. Taken from Reference [93].

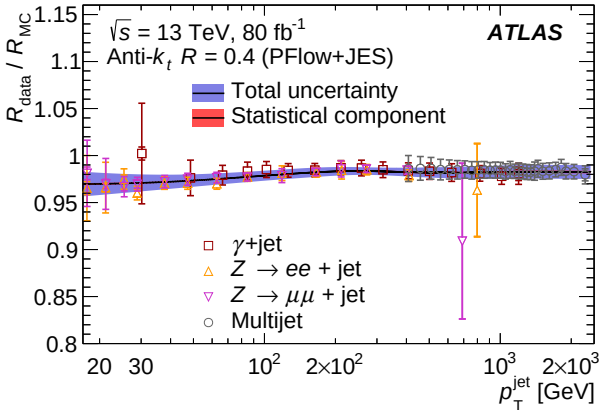


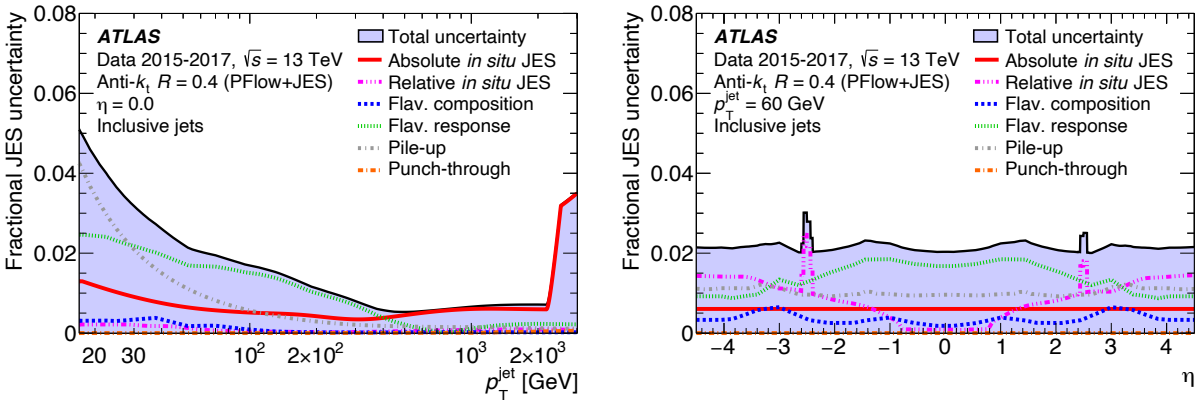
Figure 3.11: Combination of the corrections of the in situ methods, along with their combined uncertainties. Taken from Reference [93].

contribution to the measurement. The uncertainties are treated as fully correlated over the p_{T} and η range but fully uncorrelated with all other uncertainty sources in most cases. The uncertainties in $(Z \to ee)$+jet and γ+jet arising from the e/γ calibration are taken as fully correlated. Similarly, uncertainties common to Z+jet, γ+jet, and MJB are also treated as fully correlated, since MJB uses the outputs of the Z+jet and γ+jet analyses as inputs.

The sum in quadrature of the uncertainties is shown in Figure 3.11 as an uncertainty band around the calibration. The fractional uncertainties are shown as a function of p_{T} and η in Figure 3.12.

3.1.5.6 Single-particle response calibration

The in situ p_{T}-balanced based corrections (Z + jet and γ + jet) suffer from limited kinematic reach in the high-p_{T} regime due to low statistics, as well as limitations due to MC



(a) (b)

Figure 3.12: Fractional uncertainties as a function of p_{T} (a) and η (b). Taken from Reference [93].

modelling of additional radiations altering the p_{T} balance. Previously, the MJB method was used to derive a correction for jets with very high p_{T}, by calibrating them against a system of lower-p_{T} jets, which were well-calibrated thanks to the previous in situ methods. A more general in situ approach, using jet substructure information, can be used to obtain a calibration over the full p_{T} range. Since jets are composed of collections of single particles, the single-particle response measurements can be used to derive an in situ calibration and its corresponding uncertainties.

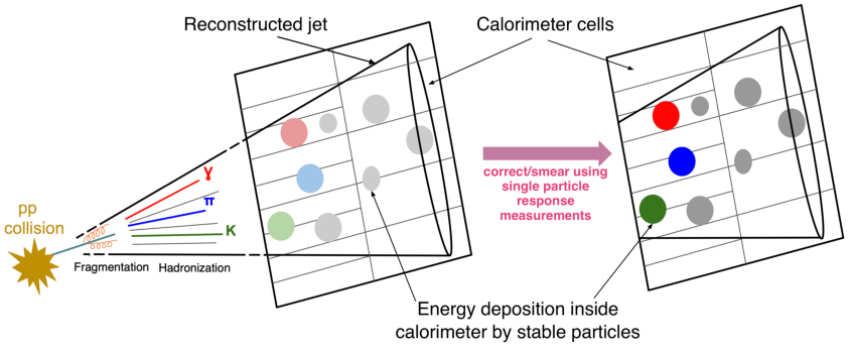


Figure 3.13: Sketch of the particles associated with a jet and the corresponding energy depositions in the calorimeter. The colours represent the energy depositions of the various particle types. In the deconvolution method, the energy depositions are corrected and smeared to obtain the data-to-MC jet energy correction from calorimeter response measurements to single particles. Taken from [101].

The deconvolution method [101] treats reconstructed jets as a superposition of individual

particle energy depositions. It relies on MC simulations but includes data-driven single-particle calorimeter response measurements for different particle types to compute the JES correction. For each calorimeter cell associated with a jet, truth particle energy deposits are identified and corrected using the measured data-to-MC response ratio for isolated particles of the same type and kinematics. The ratio of total energy depositions before and after correction is computed in data and in MC, and the mean of the data-to-MC ratio is used for calibration.

Single-particle response measurements depend on particle type and kinematics. The most abundant particle types in jets are charged and neutral pions, but some other particle types contribute as well. Charged hadron responses are measured from pion decays in $W \to \tau\nu$ events and low-p_{T} minimum-bias collisions, while electron and photon calibrations use $Z \to e^+e^-$ events. The same data-to-MC correction is used for all types of hadrons, and additional uncertainties are derived to account for the difference in the calorimeter response to hadrons of different types.

The uncertainty in the JES is derived from single-particle response uncertainties through MC pseudo-experiments. In each pseudo-experiment, the single-particle response is randomly varied within its uncertainty range, the JES correction is recalculated, and the final JES uncertainty is taken as the spread of corrections across pseudo-experiments. In addition, a jet fragmentation modelling uncertainty is derived by comparing Pythia and Sherpa MC generators.

The single-particle method and the p_{T}-balance methods yield consistent results, but the uncertainty from the single-particle response method is lower than that from the p_{T}-balance methods for $p_{\mathrm{T}} > 200$ GeV, providing better precision at high-p_{T}. Since the two types of methods are complementary in terms of p_{T} range, they can be combined to reduce systematic uncertainties and ensure consistency across the full p_{T} range. The combination procedure is the same as described in Section 3.1.5.5, where MJB is replaced by the single-particle method.

The output of the combination is shown in Figure 3.14. The combined JES uncertainty reaches 0.3% at $p_{\mathrm{T}} = 300$ GeV and 0.6% at 4 TeV, the most precise jet energy calibration achieved in ATLAS. For this reason, the single-particle response calibration will replace the MJB method in future derivations of the JES.

3.1.6 JES uncertainties application and reduction scheme

To evaluate the JES uncertainties in physics analyses, the jet p_{T} is scaled up and down by 1σ for each uncertainty component, reflecting the impact of that component on the measurement. The full jet energy scale uncertainty is described by 125 independent components, presented in Appendix A.1, which is excessive for most analyses. To address this, reduced uncertainty schemes are used, offering simplified descriptions that retain the essential correlation structure. This is done by constructing a covariance matrix of the full set and identifying the largest uncorrelated contributions, which are retained as individual components. Two reduction schemes are available: the global reduction, which combines all p_{T}-dependent in situ uncertainties regardless of origin into 8 components, and the category reduction, which groups uncertainties based on their source (detector, statistical,

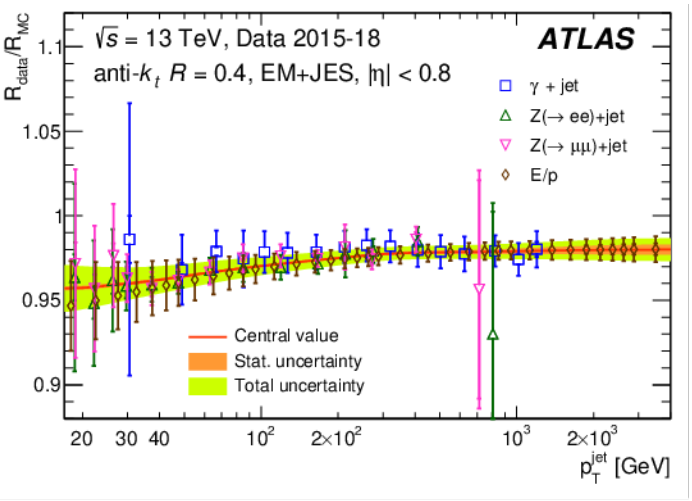


Figure 3.14: Data to MC jet response ratio as a function of the jet p_{T}, for the combined in situ E/p and p_{T}-balance methods ($Z+$jet and $\gamma+$jet). The points indicate the various measurements with their statistical (inner bars) and systematic uncertainties added in quadrature (outer bars). The line indicates the nominal combination result, while bands indicate its statistical (inner bands) and total uncertainty (outer band), respectively. Taken from [101].

modelling, or mixed), resulting in 15 components. A stronger reduction strategy can be employed for analyses requiring even fewer parameters.

3.1.7 The Jet Energy Resolution (JER)

Once the JES is applied, the transverse momentum of jets can be determined with high accuracy, but the measurements still exhibit some spread due to the finite resolution of the detector. This spread is called the Jet Energy Resolution (JER), and it is generally assumed to have a Gaussian shape of width $\sigma(p_{\mathrm{T}})$. The relative JER can be parametrised using a functional form expected for calorimeter-based resolutions:

$$\frac{\sigma(p_{\mathrm{T}})}{p_{\mathrm{T}}} = \frac{N}{p_{\mathrm{T}}} \oplus \frac{S}{\sqrt{p_{\mathrm{T}}}} \oplus C \tag{3.1}$$

where $\oplus$ is the sum in quadrature ($a \oplus b = \sqrt{a^2 + b^2}$). The precise measurement of the JER in data and MC allows for validating the simulation of the detector, as well as correcting for the data/MC difference by smearing the distributions of MC to match those of data.

3.1.7.1 JER measurement using dijet events

For the most recent measurement [93], a tag-and-probe method relying on dijet events is used to measure the JER in situ. A perfect dijet system (after application of the JES) should be perfectly balanced in p_{T}. However, this is not observed in practice because of a combination of experimental resolution, presence of additional radiation in the event, and biases due to the event selection used in the measurement. This imbalance is measured through the p_{T} asymmetry variable $\mathcal{A} = (p_{\mathrm{T}}^{\mathrm{probe}} - p_{\mathrm{T}}^{\mathrm{ref}})/p_{\mathrm{T}}^{\mathrm{avg}}$. Here, $p_{\mathrm{T}}^{\mathrm{ref}}$ is the p_{T} of a

reference jet in the central region of the calorimeter, while $p_{\mathrm{T}}^{\mathrm{probe}}$ is the p_{T} of the jet for which the resolution is measured, located anywhere in the detector. The average p_{T} is defined as $p_{\mathrm{T}}^{\mathrm{avg}} = (p_{\mathrm{T}}^{\mathrm{probe}} + p_{\mathrm{T}}^{\mathrm{ref}})/2$.

The width of the $\mathcal{A}$ distribution at reco level for a particular $(p_{\mathrm{T}}^{\mathrm{avg}}, \eta^{\mathrm{probe}})$, $(\sigma_{\mathcal{A}})_{\mathrm{meas}}$ is extracted using an iterative fitting procedure, assuming a Gaussian convolution of detector effects with the particle-level balance spectrum. To isolate detector effects, the contribution from physics phenomena such as additional radiation (e.g. a third jet or out-of-cone radiation) is estimated from the width of the true $\mathcal{A}$ distribution, denoted $(\sigma_{\mathcal{A}})_{\mathrm{true}}$, and subtracted from $(\sigma_{\mathcal{A}})_{\mathrm{meas}}$ to get $\sigma_{\mathcal{A}}$, which represent the intrinsic detector resolution. Finally, the JER for the probe jet is obtained by quadratically subtracting the relative resolution of the reference jet: $\left\langle \frac{\sigma_{p_{\mathrm{T}}}}{p_{\mathrm{T}}} \right\rangle^{\mathrm{probe}} = \sigma_{\mathcal{A}}^{\mathrm{probe}} \ominus \left\langle \frac{\sigma_{p_{\mathrm{T}}}}{p_{\mathrm{T}}} \right\rangle^{\mathrm{ref}}$.

The main uncertainties in the determination of the JER are:

- The non-closure uncertainty, evaluated as the difference between the resolution measured in MC simulation and the particle-level resolution, $\sigma(R)/R$, where $R = p_{\mathrm{T}}^{\mathrm{reco}}/p_{\mathrm{T}}^{\mathrm{true}}$. This uncertainty is around 0.4% and increases with p_{T}.
- The uncertainties propagated from the JES, dominant at low-p_{T}.
- The uncertainties due to the event selection, derived by varying the selection cut and comparing with the nominal result
- Physics modelling uncertainty, obtained by comparing the nominal result with that derived with an alternative MC generator.

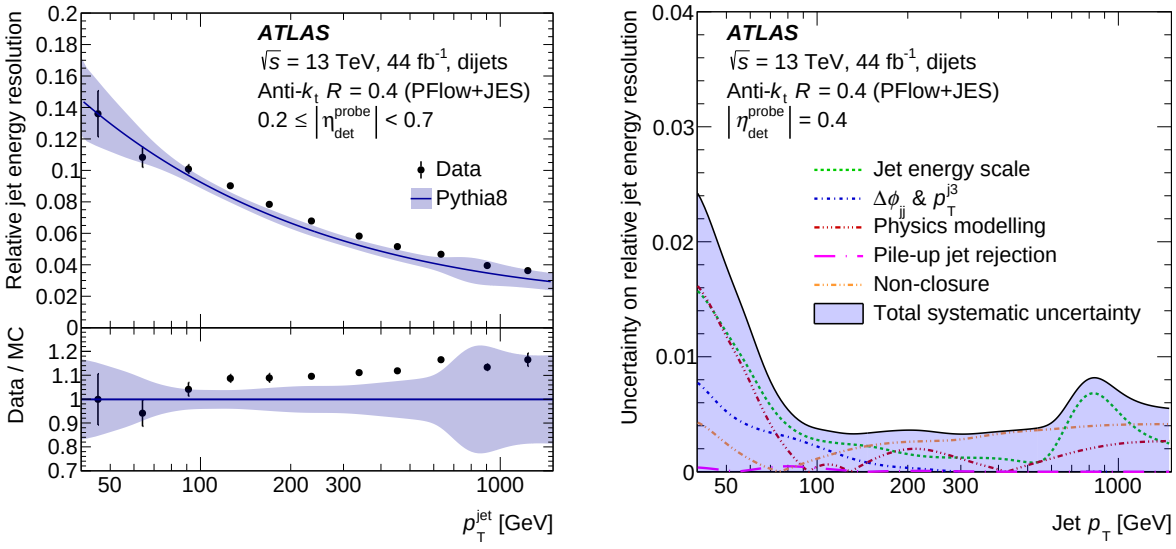
The relative JER and its absolute uncertainty are displayed as a function of p_{T} in Figure 3.15.

In previous publications, measurements based on γ+jet and Z+jet events have also been employed [102].

3.1.7.2 Noise measurement

The noise term of Equation 3.1 can be directly estimated by measuring the fluctuations in energy deposits due to pile-up, using data samples. These measurements are performed using the random cones method, in which energy deposits in the calorimeter are summed at the energy scale of the constituents, within circular regions that mimic the typical jet area. Two cones are defined at random ϕ values and in opposite $\pm\eta$ regions, and the difference of their transverse momentum sums, $\Delta p_{\mathrm{T}}^{\mathrm{RC}}$, is computed. The width of that distribution σ_{RC}, defined by the 68% confidence interval, provides a measurement of the random fluctuations of deposited energy. The estimated pile-up noise is estimated in multiple $|\eta|$ bins as $N^{\mathrm{PU}} = \sigma_{\mathrm{RC}}/(2\sqrt{2})$. This method captures the non-Gaussian behaviour of noise contributions to the JER. Since the energy deposits are defined at the constituent scale, a conversion factor is required to obtain the measurement of the noise term at the fully calibrated scale. The nominal JES calibration factor is used to perform this conversion. The result is then fitter with $N/p_{\mathrm{T}} \oplus S/\sqrt{p_{\mathrm{T}}}$, to account for non-negligible stochastic contributions

A closure test is performed by comparing the in situ measurement of the calibrated N^{PU}



(a) (b)

Figure 3.15: (a) Relative jet energy resolution and (b) absolute uncertainty in the relative resolution as a function of p_{T} in the central region of the detector, measured using the dijet balance method. The resolution in data is shown in black points with error bars indicating statistical uncertainties; the resolution in detector-level simulated events is shown by the blue curve, with total systematic uncertainty given by the blue band. The systematic uncertainty is dominated by terms propagated from the JES uncertainty. Taken from Reference [93].

with the expectation in MC simulations. To isolate the contribution to the JER from pile-up in MC simulations, the noise term is calculated using samples with and without pile-up. The contribution from pile-up noise is then extracted as the quadratic difference of the two noise terms obtained. The result is then fitted with N/p_{T}, and the non-closure is the difference between the N values obtained.

Finally, to estimate the electronic noise, the JER measured in MC simulations without pile-up is fitted with the function of Equation 3.1, where the electronic noise term is the N term. The total noise term used in the JER combination is taken as the quadratic sum of the calibrated N^{PU} and the electronic noise term. The estimation of the noise term with the random cone method is affected by the following sources of uncertainty:

- An uncertainty due to the non-closure of the method, which is dominant,
- uncertainties in the JES,
- Statistical uncertainties,
- The variations of the quantile of the confidence interval for the measurement of σ_{RC},
- An uncertainty due to the fit parameterisation and stability of the electronic noise term,
- The data/MC differences in the electronic noise component are conservatively esti-

mated to be 20%.

The total noise term and its systematic uncertainties are presented in Figure 3.16.

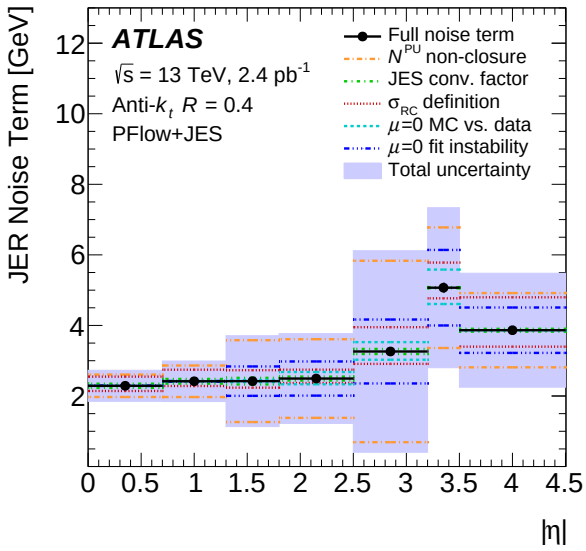


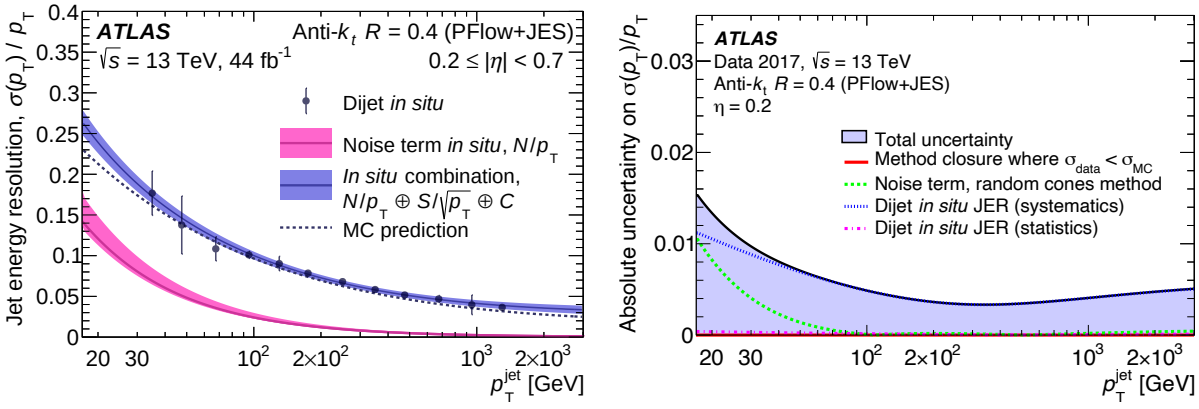
Figure 3.16: Noise term due to pileup estimated using the random cone method and its uncertainties as a function of $|\eta|$. $\mu = 0$ in the legend refers to the electronic noise term, measured in no-pile-up conditions. Taken from [93].

3.1.7.3 Combination and application

A combined measurement of the JER is obtained by fitting the dijet balance measurement using the function form of Equation 3.1, with the N term constrained by the total noise term measured with the method presented in Section 3.1.7.2. The uncertainties in the noise term are propagated by repeating the fit with different values of N. The uncertainties from the dijet measurement are propagated in the same way as for the JES, using a similar eigenvalue decomposition to reduce the final number of nuisance parameters. The resulting combined measurement is shown with its uncertainties in Figure 3.17.

For data analyses, a smearing procedure is recommended to ensure that the resolution of the jet energy scale in simulation matches that in data, wherever possible. Specifically, in regions of jet p_{T} where the resolution in data exceeds that in simulation, the simulated jets are smeared to match the broader resolution observed in data. On the contrary, no smearing is performed when the resolution is smaller than in the simulation.

The uncertainties are propagated through analyses by smearing the jets according to a Gaussian function of width σ_{smear} such that $\sigma^2_{\mathrm{smear}} = (\sigma_{\mathrm{nominal}} + |\sigma_{\mathrm{varied}}|)^2 - \sigma^2_{\mathrm{nominal}}$, with $\sigma_{\mathrm{nominal}}$ the nominal JER (after MC smearing if necessary) and σ_{varied} the one-standard-deviation variation in the uncertainty component evaluated.



(a) (b)

Figure 3.17: Relative jet energy resolution (a) and its absolute uncertainty (b) as a function of p_{T}. Taken from Reference. [93].

3.2 Detailed description of the η-intercalibration

The η-intercalibration is one of the main focuses of this thesis. It is the first in situ method applied in the calibration sequence. Its goal is to remove any dependence in η_{det} of the jet p_{T} response, defined by:

$$\mathcal{R} = \frac{p_{\mathrm{T}}^{\mathrm{probe}}}{p_{\mathrm{T}}^{\mathrm{ref}}} \tag{3.2}$$

where $p_{\mathrm{T}}^{\mathrm{probe}}$ is the transverse momentum of the jet under investigation and $p_{\mathrm{T}}^{\mathrm{ref}}$ is the transverse momentum of the reference jet. By conservation of the transverse momentum in the centre-of-mass frame, the transverse momenta of the two jets of a dijet event should be balanced. However, in practice, $\mathcal{R}$ is sensitive to effects such as the presence of additional radiative jets or the detector resolution. Therefore, the intercalibration factors c are defined using dijet events to measure the average imbalance between the two leading jets of a multi-jet event:

$$c(p_{\mathrm{T}}, \eta_{\mathrm{det}}) * \mathcal{R}(p_{\mathrm{T}}, \eta_{\mathrm{det}}) = 1$$

Then, calibration factors $\mathcal{C}$ are derived in bins of p_{T} and η_{det}, following the detector structure, as the double ratio of the intercalibration factors c:

$$\mathcal{C} = \Big[c(p_{\mathrm{T}}, \eta_{\mathrm{det}})\Big]_{\mathrm{data}} \Big/ \Big[c(p_{\mathrm{T}}, \eta_{\mathrm{det}})\Big]_{\mathrm{MC}}$$

3.2.1 Event selection

3.2.1.1 Datasets

The method makes use of both experimental data and Monte Carlo samples. For this study, I used experimental data recorded during the Run 2 (2015-2018) and the early Run 3 (2022-2023) data-taking period of the LHC, with collisions of protons at an energy in the centre of mass of 13 TeV (Run 2) and 13.6 TeV (Run 3), at high pile-up conditions. The simulated data uses Powheg+Pythia 8 [21,22] as the nominal Monte Carlo generator to simulate proton-proton collisions under the same conditions as for data. The simulation

provides information on the "truth" level, meaning it contains detailed information on the final particles themselves. Then, the interactions of stable particles in the ATLAS detector are fully simulated using GEANT4 [4] [90], which generates outputs at the reconstruction ("reco") level, meaning that objects are reconstructed from the detector information in the same way as in the data.

In each type of dataset, jets are reconstructed with the anti-k_{T} algorithm [43] with a radius of $R = 0.4$ and calibrated with the first calibration steps described in Section 3.1. No in situ correction is applied since the η-intercalibration correction is always the first in situ calibration to be derived and applied.

3.2.1.2 Basic event selection and triggers

First of all, events that may be affected by data-taking issues, whether in LHC or ATLAS, are discarded. This is achieved using a "Good Runs List" (GRL), which lists periods during which the data-taking conditions are satisfactory. Moreover, events that involve parts of the detector affected by local problems are removed. To ensure good quality jets, a primary vertex is required, which must have at least two tracks in the tracker and two track particles associated with it. Then, an event-level cleaning is performed to reject jets of non-collision origins, which can stem from the beam-induced background due to proton loss upstream of the interaction point, cosmic ray showers or calorimeter noise. It is described in Section 2.3.3.2.

Events are required to pass a combination of two single-jet triggers made of one central jet trigger (working in the $|\eta| < 3.2$ range) and one forward jet trigger (working in the $3.2 < |\eta| < 4.9$ range). Each jet trigger is made of an L1 jet trigger and an HLT jet trigger, as explained in Section 2.2.4, which both impose a criterion on the p_{T} of the jet. Each combination is considered in the p_{T} region where it is 99% efficient, as the method requires the sample to cover the whole pseudorapidity range covered by the detector. For the η-intercalibration method, events are only required to pass one or the other trigger of the combination to be added to the dataset. Each selected event is then weighted by the appropriate amount (depending on whether it passes a central trigger, a forward trigger or both), following the inclusion method described in Reference [103]. The "trigger menu" summarises the list of triggers used for the analysis. It is specified in Table 3.1 for Run 2. I performed the trigger efficiency study for Run 3, see Section 3.4.2.1.

3.2.1.3 Dijet selection

To ensure a clean back-to-back dijet topology, the following event selection cuts are applied to both data and MC:

- $\Delta\phi^{\mathrm{ref,probe}} > 2.5$ to ensure a large azimuthal separation between the two leading jets,
- $p_{\mathrm{T}}^{j3}/p_{\mathrm{T}}^{\mathrm{avg}} < 0.25$ to suppress contributions from any additional third jet,
- JVT < 0.25 for jets having $p_{\mathrm{T}} < 50$ GeV and $|\eta| < 2.4$, to minimise the contribution from pile-up (see Section 2.3.3.1). This criterion is only applied to the two leading jets.

Central trigger	**Forward trigger**	**L1 p_{T} criterion**	**HLT p_{T} criterion**
HLT_j15	HLT_j15_320eta490	0	15
HLT_j25	HLT_j25_320eta490	0	25
HLT_j35	HLT_j35_320eta490	0	35
HLT_j45	HLT_j45_320eta490	15	45
HLT_j60	HLT_j60_320eta490	20	60
HLT_j85	HLT_j85_320eta490	20	85
HLT_j110	HLT_j110_320eta490	30	110
HLT_j175	HLT_j175_320eta490	50	175
HLT_j260	HLT_j260_320eta490	75	260
HLT_j360	HLT_j360_320eta490	100	360
HLT_j380	-	100	380
HLT_j400	-	100	400
HLT_j420	-	100	420

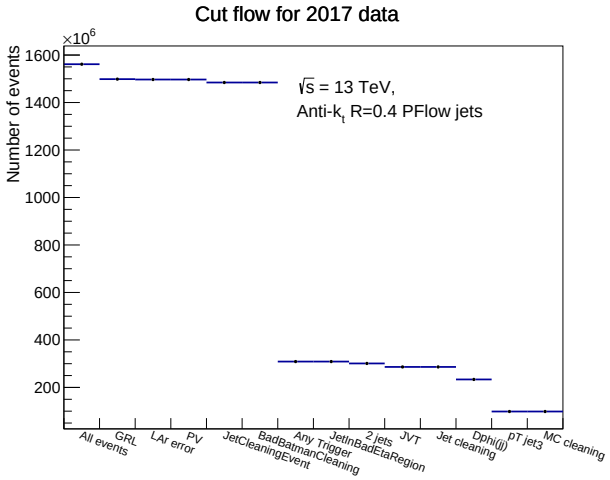
Table 3.1: Trigger table with p_{T} criteria for different triggers.

These selection criteria, with threshold values optimised as in Reference [104], aim to eliminate systematic bias in the sample. In particular, they prevent physics effects like additional radiation (third jet, out-of-cone radiation) from becoming dominant. Those effects could bias the p_{T} balance between the two jets of the dijet system. When a variation of one of the physics selection criteria is applied, a different set of events is selected, resulting in a distinct calibration. This alternative calibration is used to derive a systematic uncertainty in the nominal calibration, as explained in Section 3.2.4.1. These cuts are applied in data and Monte Carlo simulations, both at reco and truth levels, except for the JVT cut since there is no pile-up and thus no use of the JVT selection criteria at truth level. In addition, an "MC cleaning" is performed in Monte Carlo simulations at the reco level to mitigate the mismodelling of the pile-up. It consists in comparing the truth and the reco p_{T} and rejecting an event if the two leading jets at reco level are too far from the truth leading jet: $p_{\mathrm{T}}^{\text{avg, reco}}/p_{\mathrm{T}}^{\text{leading jet, truth}} < 1.4$. Rejected events can happen when pile-up jets can be overlaid to lower p_{T} events at high p_{T}. Figure 3.18 shows the effect of the successive cuts in data, Monte Carlo truth and Monte Carlo reco.

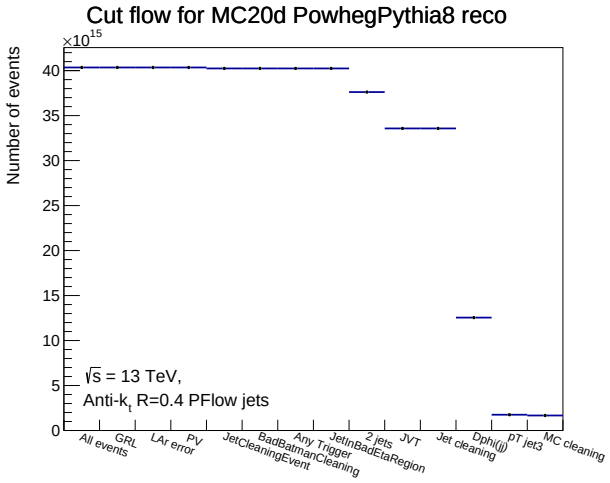
3.2.1.4 Binning

After this selection, the events are separated into $p_{\mathrm{T}}^{\text{avg}} = \frac{p_{\mathrm{T}}^{\text{ref}} + p_{\mathrm{T}}^{\text{probe}}}{2}$ bins, which are then treated independently in the first steps of the η-intercalibration. The binning is chosen to be nearly equally spaced in $\log(p_{\mathrm{T}}^{\text{avg}})$ and to contain a satisfying amount of statistics. One trigger combination is used for one $p_{\mathrm{T}}^{\text{avg}}$ bin. The $p_{\mathrm{T}}^{\text{avg}}$ binning used here had been chosen up to 1200 GeV, as explained in Reference [96], and bins up to 2000 GeV have been added, see Reference [104]. The η-intercalibration is empirically known to be rather smooth in p_{T}, hence the $p_{\mathrm{T}}^{\text{avg}}$ bins can be quite broad. The events that do not enter any $p_{\mathrm{T}}^{\text{avg}}$ bin are discarded.

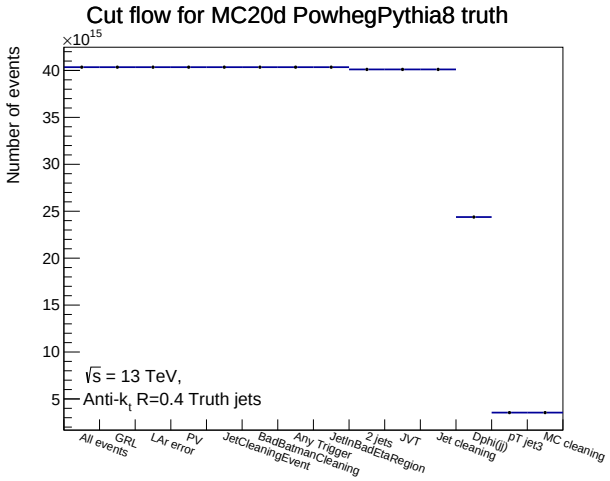
Each $p_{\mathrm{T}}^{\text{avg}}$ is further split into rapidity bins $(\eta^{\text{ref}}, \eta^{\text{probe}})$. The η binning is chosen independently for each $p_{\mathrm{T}}^{\text{avg}}$ bin. In each case, the η binning is designed to reflect the detector



(a)



(b)



(c)

Figure 3.18: Unweighted number of events after each selection step for (a) 2017 data, (b) MC20d reco and (c) MC20d truth levels.

structure. For example, a finer binning is used in the transition region around $|\eta| = 2.5$, which corresponds to the end of the tracker coverage. Since the η-intercalibration is the only in situ calibration that corrects in η, the binning is chosen to be fairly thin while keeping a sufficient amount of statistics in each $(p_{\mathrm{T}}^{\mathrm{avg}}, \eta^{\mathrm{ref}}, \eta^{\mathrm{probe}})$ bin. The exact binning used is displayed in Figure 3.19.

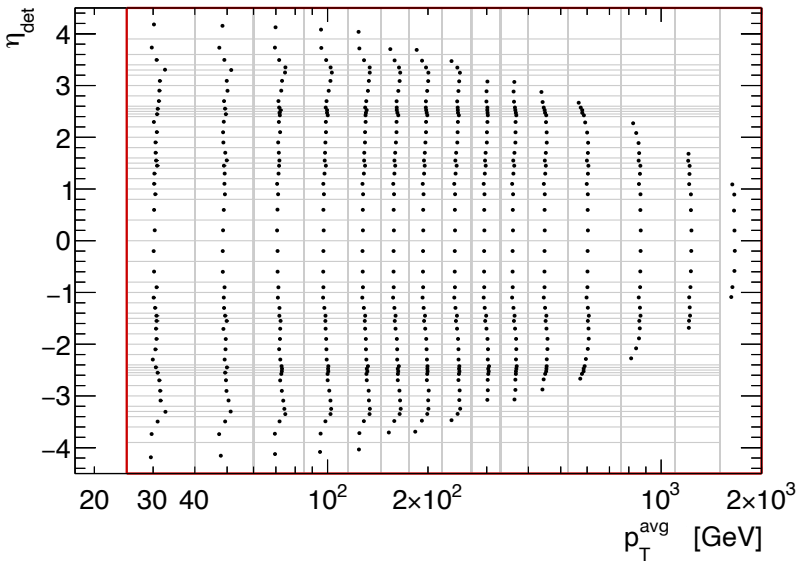


Figure 3.19: $(p_{\mathrm{T}}^{\mathrm{avg}}, \eta)$ binning. The red box delimits the phase space region in which calibration points are obtained, and the dot shows the barycentre of the used bin.

3.2.2 Derivation of the intercalibration factors

Two strategies exist to derive the correction. The simplest strategy is the "central-reference method", or "standard method", which uses dijet systems in which the reference jet is in the central region of the detector ($|\eta| < 0.8$) and the probe jet is in the forward region ($|\eta| > 0.8$). The second strategy, called the "matrix method", uses all dijet systems of the sample. For each method, the mean of the p_{T} asymmetry of the dijet system is computed as a first step. In the next section, the assessment of the mean of the asymmetry is explained in the case of the central-reference method, but the reasoning is similar to the matrix method.

3.2.2.1 Determining the mean asymmetry

The $p_{\mathrm{T}}^{\mathrm{probe}}/p_{\mathrm{T}}^{\mathrm{ref}}$ distribution is not symmetrical around its peak, which makes it difficult to retrieve the intercalibration factor c by a simple numerical inversion. Therefore, the dijet asymmetry variable is used. It is defined by:

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

Here, $p_{\mathrm{T}}^{\mathrm{probe}}$ is the transverse momentum of the jet investigated, $p_{\mathrm{T}}^{\mathrm{ref}}$ is the transverse momentum of the reference jet, and $p_{\mathrm{T}}^{\mathrm{avg}} = \frac{p_{\mathrm{T}}^{\mathrm{probe}} + p_{\mathrm{T}}^{\mathrm{ref}}}{2}$ is the average of the two. The asymmetry quantifies the deviation from a perfectly balanced p_{T} of the system, which would infer $p_{\mathrm{T}}^{\mathrm{probe}} - p_{\mathrm{T}}^{\mathrm{ref}} = 0$. Its distribution is symmetrical around its peak due to the fluctuations of $p_{\mathrm{T}}^{\mathrm{probe}}$ and $p_{\mathrm{T}}^{\mathrm{ref}}$ compensating, and the tails fall rapidly. This makes it easier to determine the average of the distribution as the position of the peak.

For every $p_{\mathrm{T}}^{\mathrm{avg}}$ bin, the asymmetry distribution $\mathcal{A}$ is evaluated in bins of η_{probe} (with η_{ref} represented by a single bin), resulting in a distribution that can be approximated by a Gaussian. The asymmetry distributions can look very different depending on the phase space region considered. For example, the forward η bins suffer from reduced statistics compared to bins with two central jets (η close to zero) for kinematic reasons, and the bins in the higher $p_{\mathrm{T}}^{\mathrm{avg}}$ region benefit from a large number of statistics since the triggers used here are unprescaled. For this reason, the asymmetry binning is chosen relatively small (0.06), and a re-binning procedure, explained in the next paragraph, is implemented for each asymmetry distribution to ensure an appropriate representation of the distribution.

The mean value of the distribution $\langle\mathcal{A}\rangle$ is defined as the central value of a Gaussian fit to the $\mathcal{A}$ distribution, if the statistics allow, or else as an arithmetic mean of the distribution. In both cases, an iterative procedure is used to extract the mean value. In the following, $\langle\mathcal{A}\rangle_n$ defines the mean asymmetry computed at the n-th step of the iterative procedure, and $\langle\mathcal{A}\rangle$ is the final value. The process to obtain the mean asymmetry through the arithmetic mean is the following:

1. $\langle\mathcal{A}\rangle_1$ is computed as the arithmetic mean in the $[-0.7, 0.7]$ interval.
2. $\langle\mathcal{A}\rangle$ is computed as the arithmetic mean in the $[\langle\mathcal{A}\rangle_1 - 0.5, \langle\mathcal{A}\rangle_1 + 0.5]$ interval, and the uncertainty $\Delta\langle\mathcal{A}\rangle$ is taken as the RMS of the distribution divided by the effective number of entries.

The process to obtain the mean asymmetry through the Gaussian fit is as follows:

1. The asymmetry histogram is re-binned to facilitate an accurate Gaussian description. This is done by successively merging pairs of bins until the bin width exceeds the optimal value, defined as $w_{optimal} = 3.5 * \mathrm{RMS} * N_{\mathrm{eff}}^{-1/3}$. Here, N_{eff} is the effective number of entries defined by $N_{\mathrm{eff}} = \left(\sum_{i=1}^{n} w_i\right)^2 / \left(\sum_{i=1}^{n} w_i^2\right)$, where w_i is the weight of the i-th event entering into the histogram[1]. This optimal bin formula is taken from Reference [105].
2. $\langle\mathcal{A}\rangle_1$ and $\Delta\langle\mathcal{A}\rangle_1$ is taken as the central value and associated uncertainty of a Gaussian fit to the distribution in the $[-0.8, 0.8]$ interval.
3. $\langle\mathcal{A}\rangle_2$ and $\Delta\langle\mathcal{A}\rangle_2$ is taken as the central value and associated uncertainty of a Gaussian fit to the distribution in the $[\langle\mathcal{A}\rangle_1 - \alpha * \Delta\langle\mathcal{A}\rangle_1, \langle\mathcal{A}\rangle_1 + \alpha * \Delta\langle\mathcal{A}\rangle_1]$ interval, with $\alpha = 1.5$.
4. $\langle\mathcal{A}\rangle$ and $\Delta\langle\mathcal{A}\rangle$ are taken as the central value and associated uncertainty of a Gaussian fit to the distribution in the $[\langle\mathcal{A}\rangle_2 - \alpha * \Delta\langle\mathcal{A}\rangle_2, \langle\mathcal{A}\rangle_2 + \alpha * \Delta\langle\mathcal{A}\rangle_2]$ interval, with

[1] For an unweighted histogram, N_{eff} is simply the number of entries. For a weighted histogram, it is the hypothetical number of entries an unweighted histogram with the same statistical power would have.

$\alpha = 1.5$.

An example of a distribution and its associated fit is presented in Figure 3.20.

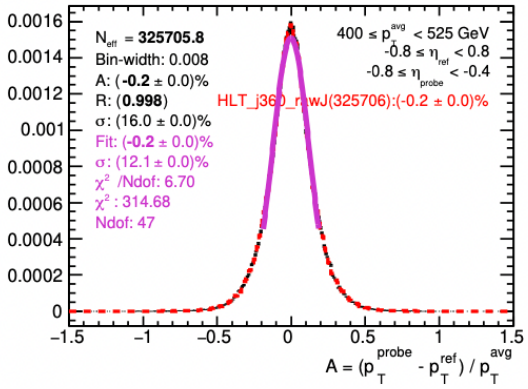


Figure 3.20: Asymmetry distribution for a $400 < p_{\mathrm{T}}^{\mathrm{avg}} < 525$ GeV, $-0.8 < \eta_{probe} < -0.4$ bin and associated Gaussian fit (in purple), for 2017 data. The black dots represent the total distribution, while the coloured dashed line represents one trigger combination.

The entire procedure relies heavily on the accurate estimation of the mean asymmetry. Ensuring that the mean asymmetry is determined with minimal bias is crucial for the success of the η-intercalibration.

3.2.2.2 The central reference method

At first order, within a relatively narrow bin of $p_{\mathrm{T}}^{\mathrm{avg}}$, the average value of the asymmetry is given by $\langle \mathcal{A} \rangle = 2 \times \frac{\langle p_{\mathrm{T}}^{\mathrm{probe}} \rangle - \langle p_{\mathrm{T}}^{\mathrm{ref}} \rangle}{\langle p_{\mathrm{T}}^{\mathrm{probe}} \rangle + \langle p_{\mathrm{T}}^{\mathrm{ref}} \rangle}$. The intercalibration factors c_i, which is the inverse of the jet p_{T} response defined by $\langle \mathcal{R} \rangle = \langle p_{\mathrm{T}}^{\mathrm{probe}} \rangle / \langle p_{\mathrm{T}}^{\mathrm{ref}} \rangle$ in one $p_{\mathrm{T}}^{\mathrm{avg}}$ bin and the η^{probe} bin i, can then be written:

$$\left(\frac{1}{c}\right)_i = \frac{2 + \langle \mathcal{A} \rangle_i}{2 - \langle \mathcal{A} \rangle_i}$$

and the associated uncertainty is defined as

$$\Delta\left(\frac{1}{c}\right)_i = \frac{4\Delta\langle \mathcal{A} \rangle_i}{(2 - \langle \mathcal{A} \rangle_i)^2}$$

Jets in the central region ($|\eta| < 0.8$) receive an average correction of 1 at this step of the calibration. The result for one $p_{\mathrm{T}}^{\mathrm{avg}}$ bin is shown in Figure 3.21.

Due to the constraints imposed on the dijet system, this method suffers from reduced statistics. The Matrix Method was created to circumvent this issue by using all dijet events.

3.2.2.3 The matrix method

The principle of the Matrix Method is to use all dijet events to calibrate every region relative to the others. To maximise statistics, neither jet is required to be in the central

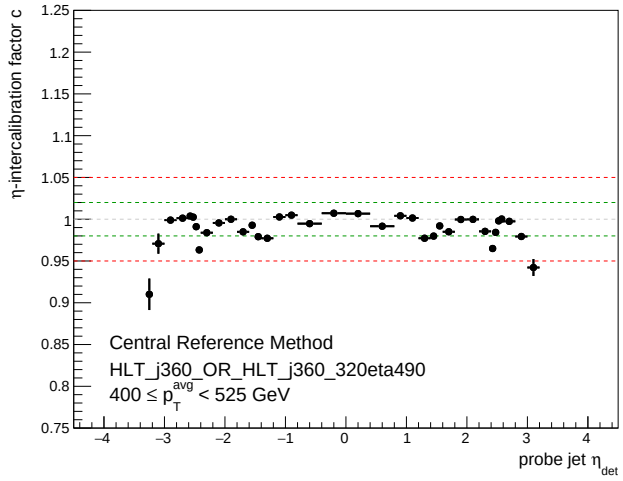


Figure 3.21: Calibration factors derived with the central reference method for $400 < p_{\mathrm{T}}^{\mathrm{avg}} < 525$ GeV, for 2017 data. The uncertainty in the calibration factors is propagated linearly from the uncertainty in the asymmetry fit, provided by the MINUIT minimiser [106].

region: instead, each jet of the dijet system serves as a reference jet for the other one. This way, the intercalibration factors for one η bin are calculated using measurements from all the other η bins, contrary to the central reference method, where only the reference region is used. If one measurement fails, is inaccurate or imprecise, its impact is reduced by the combination with all the other measurements. This provides strong robustness to the method, along with significantly enhanced statistics.

The two leading jets are labelled "left" and "right", such that $\eta_{\mathrm{left}} < \eta_{\mathrm{right}}$, and the asymmetry variable is defined as:

$$\mathcal{A} = \frac{p_{\mathrm{T}}^{\mathrm{left}} - p_{\mathrm{T}}^{\mathrm{right}}}{p_{\mathrm{T}}^{\mathrm{avg}}}$$

Similarly to the central reference method, for each $p_{\mathrm{T}}^{\mathrm{avg}}$ bin, the jet response is defined in bins of η_{left} (i), η_{right} (j) as:

$$\langle \mathcal{R}_{ij} \rangle = \frac{2 + \langle \mathcal{A} \rangle_{ij}}{2 - \langle \mathcal{A} \rangle_{ij}}$$

and the associated statistical uncertainty is

$$\Delta \langle \mathcal{R}_{ij} \rangle = \frac{4 \Delta \langle \mathcal{A} \rangle_{ij}}{(2 - \langle \mathcal{A} \rangle_{ij})^2}$$

For each p_{T} bin, the response matrix is computed as the 2D distribution of $\mathcal{R}_{ij}$ in $p_{\mathrm{T}}^{\mathrm{left}}$, $p_{\mathrm{T}}^{\mathrm{right}}$ bins. An example of a response matrix is presented in Figure 3.22.

It is an upper triangular matrix (hence the name of the method), and its components act like a set of $N^2 - N$ constraints on the evaluation of the N intercalibration factors of the considered $p_{\mathrm{T}}^{\mathrm{avg}}$ bin. The intercalibration factors are chosen to fulfil $c_{\mathrm{right}} \langle \mathcal{R} \rangle - c_{\mathrm{left}} = 0$, which results in $\langle \mathcal{A}_{\mathrm{calibrated}} \rangle = 0$. This results in an over-constrained system, and for

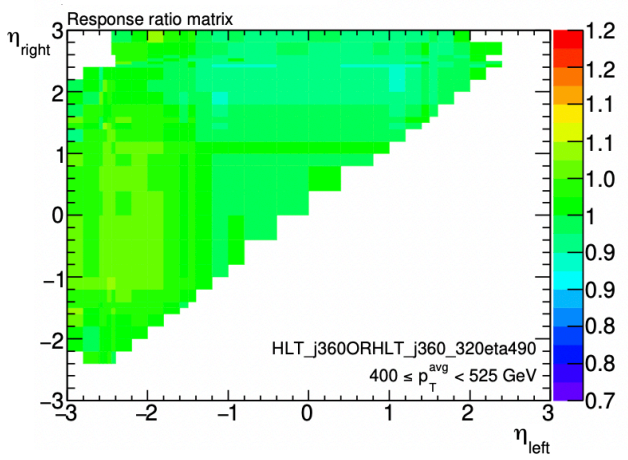


Figure 3.22: Response matrix for $400 < p_{\mathrm{T}}^{\mathrm{avg}} < 525$ GeV, in 2017 data. This matrix is used to constrain the intercalibration factors.

each $p_{\mathrm{T}}^{\mathrm{avg}}$ bin the intercalibration factors c_i are derived simultaneously from an analytical minimising procedure of the response matrix (thoroughly described in Reference [98]), which consists in minimising this function:

$$S(c_1, ..., c_N) = \sum_{j=1}^{N} \sum_{i=1}^{j-1} \left(\frac{1}{\Delta \langle \mathcal{R}_{ij} \rangle} \left(c_i \langle \mathcal{R}_{ij} \rangle - c_j \right) \right)^2 + X(c_1, ..., c_N) \tag{3.3}$$

Finding the minimum of this function ensures that the c_i is chosen to correct the jet response to be closer to 1. The system is over-constrained, so more weight can be given to the constraints for which the uncertainty is lower. In addition, in some cases where the determination of $\langle \mathcal{R}_{ij} \rangle$ is not satisfactory (bad asymmetry fit, lack of statistics, central value too far from the expected value of 0), the contribution from the corresponding bin can be dropped with little impact on the result. However, the more constraints there are from the η combinations, the more robust the method is. This is why, during my PhD, I focused on recovering the maximum number of constraints, as detailed in Section 3.3.1. Moreover, the compatibility between those constraints can be assessed by evaluating the contribution of each combination to a global χ^2 of the global minimisation. This is a novelty developed in the context of my PhD. It is further explained in Section 3.3.2.

Since the normalisation of the c_i factors is not constrained by the minimisation of the first term of the S function, the X function, defined by

$$X(c_1, ..., c_N) = \lambda \left(\frac{1}{N} \sum_{i=1}^{N} c_i - 1 \right)^2$$

is added to fix the value of the average of c_i to unity. It also prevents the minimising algorithm from choosing the trivial solution, which is all $c_i = 0$, thanks to the Lagrange multiplier λ, which is arbitrarily chosen to be large ($\lambda \sim N$), for numerical stability. The value of λ does not affect the solution. A recent development of the η-intercalibration has been the implementation of an analytical minimisation of the S function instead of the

standard MINUIT minimiser [106], by Robert Hankache [104]. It made the minimisation approximately a thousand times faster, enabling the use of time-consuming techniques such as the Bootstrap Method, described in Section 3.3.3, as well as finer η binning to better scrutinise detector features. The exact definition of the analytical minimisation function is provided in Reference [93].

The correction being relative, the resulting intercalibration factors c_i are defined up to an overall multiplicative constant. To fix this arbitrariness, the central region is chosen as the reference. The calibration factors are scaled afterwards such that the simple average of the factors in the central region is equal to 1. The result is shown in Figure 3.23. The statistical uncertainties in the calibration factors are smaller than in the case of the central reference method (see Figure 3.21).

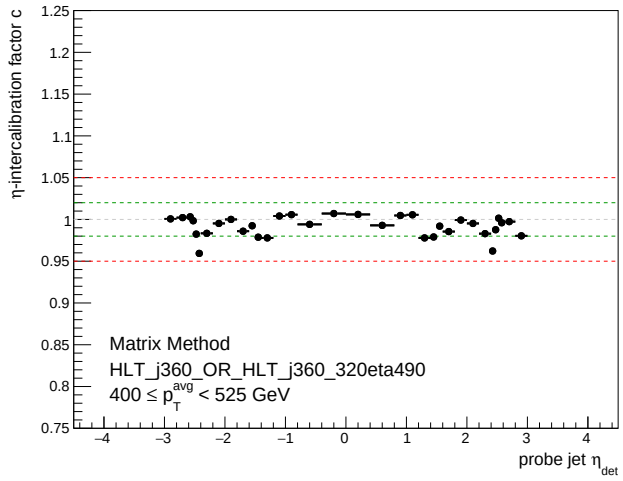


Figure 3.23: Calibration factors derived with the matrix method for $400 < p_{\mathrm{T}}^{\mathrm{avg}} < 525$ GeV, for 2017 data. The statistical uncertainties in the calibration factors are propagated from the uncertainty in the asymmetry fit using pseudo-experiments (see Section 3.2.4.2). Dashed lines are drawn to guide the eye: green and red lines correspond to corrections of 2.5% and 5%, respectively.

3.2.3 Derivation of the calibration factors

The residual calibration is derived in bins of $(p_{\mathrm{T}}^{\mathrm{avg}}, \eta)$ as the ratio of the intercalibration factors, in data and in MC.

$$\mathcal{C} = \Big[c(p_{\mathrm{T}}, \eta)\Big]_{\mathrm{data}} \Big/ \Big[c_(p_{\mathrm{T}}, \eta)\Big]_{\mathrm{MC,reco}}$$

However, the calibration is meant to be applied to any jet regardless of the topology of the event, so the calibration factors need to be derived in bins of (p_{T}, η). With the Central Reference method's notations, the relation between $p_{\mathrm{T}}^{\mathrm{avg}}$ and p_{T} is defined considering a jet with transverse momentum p_{T} as part of a dijet system along with a jet in the reference region. Thus, the reference jet has $p_{\mathrm{T}}^{\mathrm{ref}} = p_{\mathrm{T}} * c = p_{\mathrm{T}}/\mathcal{R}$ (with $\mathcal{R}$ defined in Equation 3.2) and the $p_{\mathrm{T}}^{\mathrm{avg}}$ would be:

$$p_{\mathrm{T}}^{\mathrm{avg}} = \frac{p_{\mathrm{T}} + p_{\mathrm{T}}^{\mathrm{ref}}}{2} = \frac{1+\mathcal{R}}{2\mathcal{R}} * p_{\mathrm{T}}$$

Then, the calibration is expressed at the barycentre of the (p_{T}, η) bin and a Gaussian kernel smoothing is performed to smooth statistical fluctuations and ensure a continuous calibration, removing any unphysical discontinuity. The result is a calibration function $F(p_{\mathrm{T}}, \eta)$ which is expressed in fine bins: the Gaussian kernel also performs the interpolation and extrapolation over the whole (p_{T}, η) range. The calibration is frozen above the highest p_{T} calibration point $p_{\mathrm{T}}^{\mathrm{freeze}} = 1651$ GeV. An example of the final calibration for 2017 data is given in Figure 3.24. The calibration is smooth with respect to p_{T}, and displays fast variations around $|\eta| \approx 1.4$, $|\eta| \approx 2.5$, $|\eta| \approx 3.1$, as expected given the detector structure.

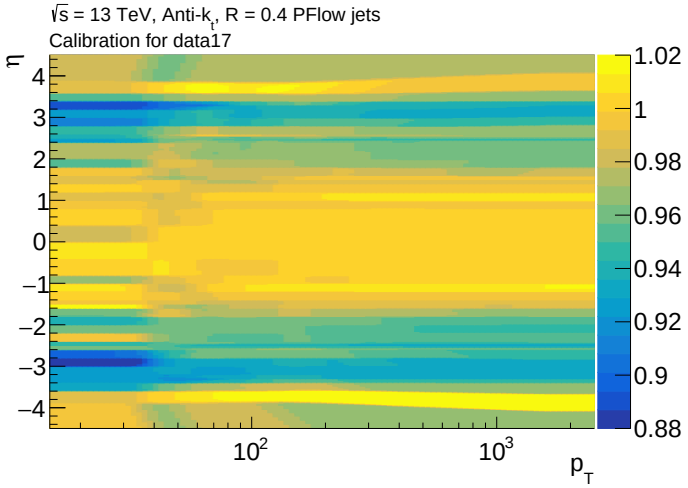


Figure 3.24: Full calibration after smoothing and freezing, for 2017 data.

3.2.4 Uncertainties

3.2.4.1 Systematic uncertainties

The η-intercalibration aims to correct the mismodelling of the detector simulation. Still, the calibration factors are sensitive to other effects, such as the mismodelling of physics processes, and the choice of criteria for the event selection. Systematic uncertainties cover these effects.

Event selection Selection cuts are applied to target the dijet topologies as described in Section 3.2.1.3. The threshold values have been optimised in Reference [104] to remove unwanted events while keeping a reasonable amount of statistics: they result from a compromise and do not fully remove physics effects, arising from additional radiation in the event (third jet, out-of-cone radiation) that would bias the p_{T} balance of the dijet system. As a result, systematic uncertainties have to be estimated to account for the residual physics effects. Since the choice of the exact cut-off values is arbitrary, the systematic uncertainties are evaluated by deriving the final result after varying these selection criteria. Each selection criterion is varied up and down, with the values shown in Table 3.2, the calibration is derived for both variations, and the result is compared to the nominal calibration. The relative systematic uncertainty is taken as the maximum difference between the varied and the nominal calibration.

	Cut		
	Loose	Nominal	Tight
JVT	-	0.2	0.5
$\Delta\phi^{j1,j2}$	2.3	2.5	2.8
$p_{\mathrm{T}}^{\mathrm{jet\ 3}}/p_{\mathrm{T}}^{\mathrm{avg}}$	0.35	0.25	0.15

(a) Selections for reco events

	Cut		
	Loose	Nominal	Tight
$\Delta\phi^{j1,j2}$	2.3	2.5	2.8
$p_{\mathrm{T}}^{\mathrm{jet\ 3}}/p_{\mathrm{T}}^{\mathrm{avg}}$	0.35	0.25	0.15

(b) Selections for truth events

Table 3.2: Variation of the selection for reco (a) and truth (b) events.

Modelling uncertainties In Monte Carlo simulations, different models can be used to describe the showering and the hadronisation processes. Uncertainties related to the modelling of the parton shower, jet fragmentation, and multiple parton interactions, which impact the p_{T} balance between the photon and the jet, can be estimated by using different Monte Carlo generators. The calibration is derived using the alternative Monte Carlo generators, and the relative systematic uncertainty is taken as the difference from the nominal calibration. Here, the nominal generator is POWHEG+PYTHIA 8 [21,22], and the alternative generator is SHERPA [24]. They use different models for parton showering and different PDFs, and SHERPA works at Leading Order (LO) while POWHEG+PYTHIA 8 works at Next to Leading Order (NLO). The modelling uncertainties can be reduced using the absolute implementation method explained in Section 3.1.5 and presented in Reference [97].

Combination of the uncertainties The systematic uncertainties are then smoothed using the same method as for the nominal case, but with stronger smoothing, achieved by using a wider Gaussian in the Gaussian kernel. This is because no fine structure is anticipated in the systematic uncertainties since they are designed to cover physics effects which are expected to be smooth. The residual structure that might be seen would then be due to detector effects, which are not meant to be covered at that stage, so they must be smoothed out.

Usually, the systematic uncertainties evaluated are not symmetrical relative to $\eta = 0$. This is expected to come from statistical fluctuations since the physics effects are symmetrical. This is corrected conservatively by choosing $\max(|\Delta C(p_{\mathrm{T}}, \eta)|, |\Delta C(p_{\mathrm{T}}, -\eta)|)$ both for η and $-\eta$. Then, for the azimuthal separation and third jet radiation, the uncertainties obtained with the up and down variations are combined by taking the maximum absolute value of the two. Finally, the four remaining uncertainties are summed in quadrature to provide a single uncertainty to physics analyses.

An example of systematic and statistical uncertainties, derived for PFlow jets at $\sqrt{s} =$ 13 TeV, is presented in Figure 3.25. At all p_{T}, the MC modelling term dominates the uncertainty, except in the central region of the detector ($|\eta_{\mathrm{det}}| < 0.8$). The derivation of the statistical uncertainties is explained in the next section.

Closure test The calibration closure is tested to check whether the calibration has successfully corrected all the differences in the jet response. For that, the calibration is derived once, then applied to a sample and derived again on the corrected sample. If the

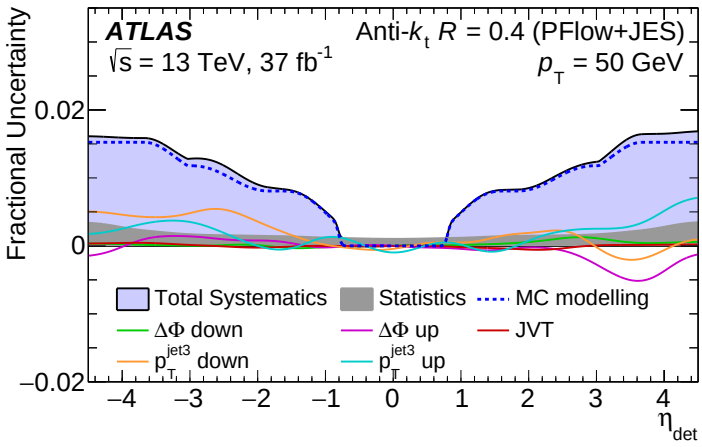


Figure 3.25: Statistical and systematic uncertainties associated to the η-intercalibration as a function of η for PFlow+JES jets of $p_{\mathrm{T}} = 50$ GeV. The solid purple band shows the total systematic uncertainty, while the grey band shows the statistical uncertainty alone. Coloured lines mark individual sources of uncertainty. Taken from [93].

method is ideal, the correction factors from the second derivation are compatible with unity, and the closure is good. The non-closure of the method, i.e. the difference of those factors to unity, is taken as a methodological systematic uncertainty.

3.2.4.2 Statistical uncertainties

For the central reference method, the statistical uncertainty is simply linearly propagated from the uncertainty in the mean value of the asymmetry $\Delta\mathcal{A}$, taken from the fit and provided by the MINUIT minimiser [106]. For the Matrix Method, the statistical uncertainty is derived as follows. The minimisation gives the central values of the c_i for each p_{T} bin k. To assess the statistical uncertainty, 1000 toy models are generated by randomly fluctuating the nominal $\langle\mathcal{R}\rangle$ distribution, following a central Gaussian of width $\Delta\langle\mathcal{R}\rangle = 4 * \Delta\langle\mathcal{A}\rangle/(2 - \langle\mathcal{A}\rangle)^2$, $\Delta\langle\mathcal{A}\rangle$ being the uncertainty in the central value of the asymmetry fit:

$$\langle\mathcal{R}\rangle^{\mathrm{toy}} = \langle\mathcal{R}\rangle + \Delta\langle\mathcal{R}\rangle \times G(0,1)$$

Then, the new S function, as defined in 3.3, is minimised again and the c_i^{toy} are derived for the 1000 toys. The uncertainty of each c_i is then defined as $\Delta c_i = \mathrm{RMS}(c_i^{\mathrm{toy}})$. Figure 3.26 shows an example of the calibration factors derived with the matrix method for the nominal sample and 1000 toys.

This method relies on the approximation that the distribution of the central value of the Gaussian fit to the asymmetry distributions is itself Gaussian. However, it has been found that this is not always accurate. I have studied and implemented the so-called Bootstrap method to evaluate the statistical uncertainties. The Bootstrap Method does not rely on approximation and generates statistical fluctuations at the event level to provide additional information on the correlations between the measurements. It is further explained in Section 3.3.3.

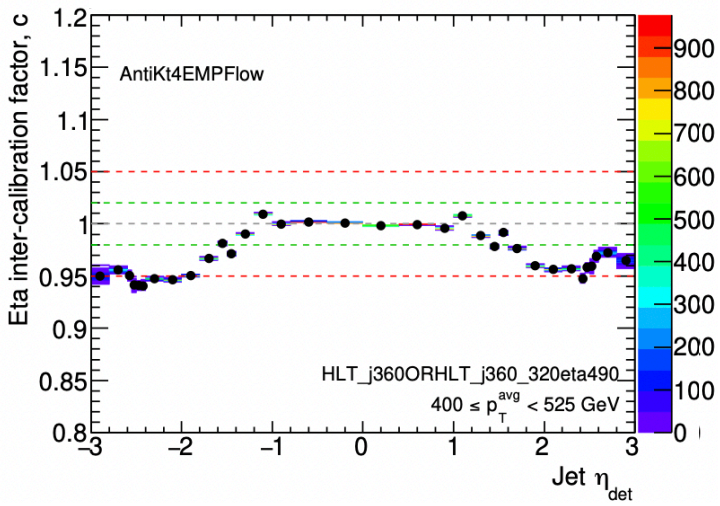


Figure 3.26: Calibration factors derived with the matrix method for $400 < p_{\mathrm{T}}^{\mathrm{avg}} < 525$ GeV, for 2017 data (black dots), along with the results for the 1000 toys, which each enter on the Z axis. The statistical uncertainties are taken as the RMS of the distribution of these toy results.

3.3 Improvements to the η-intercalibration method

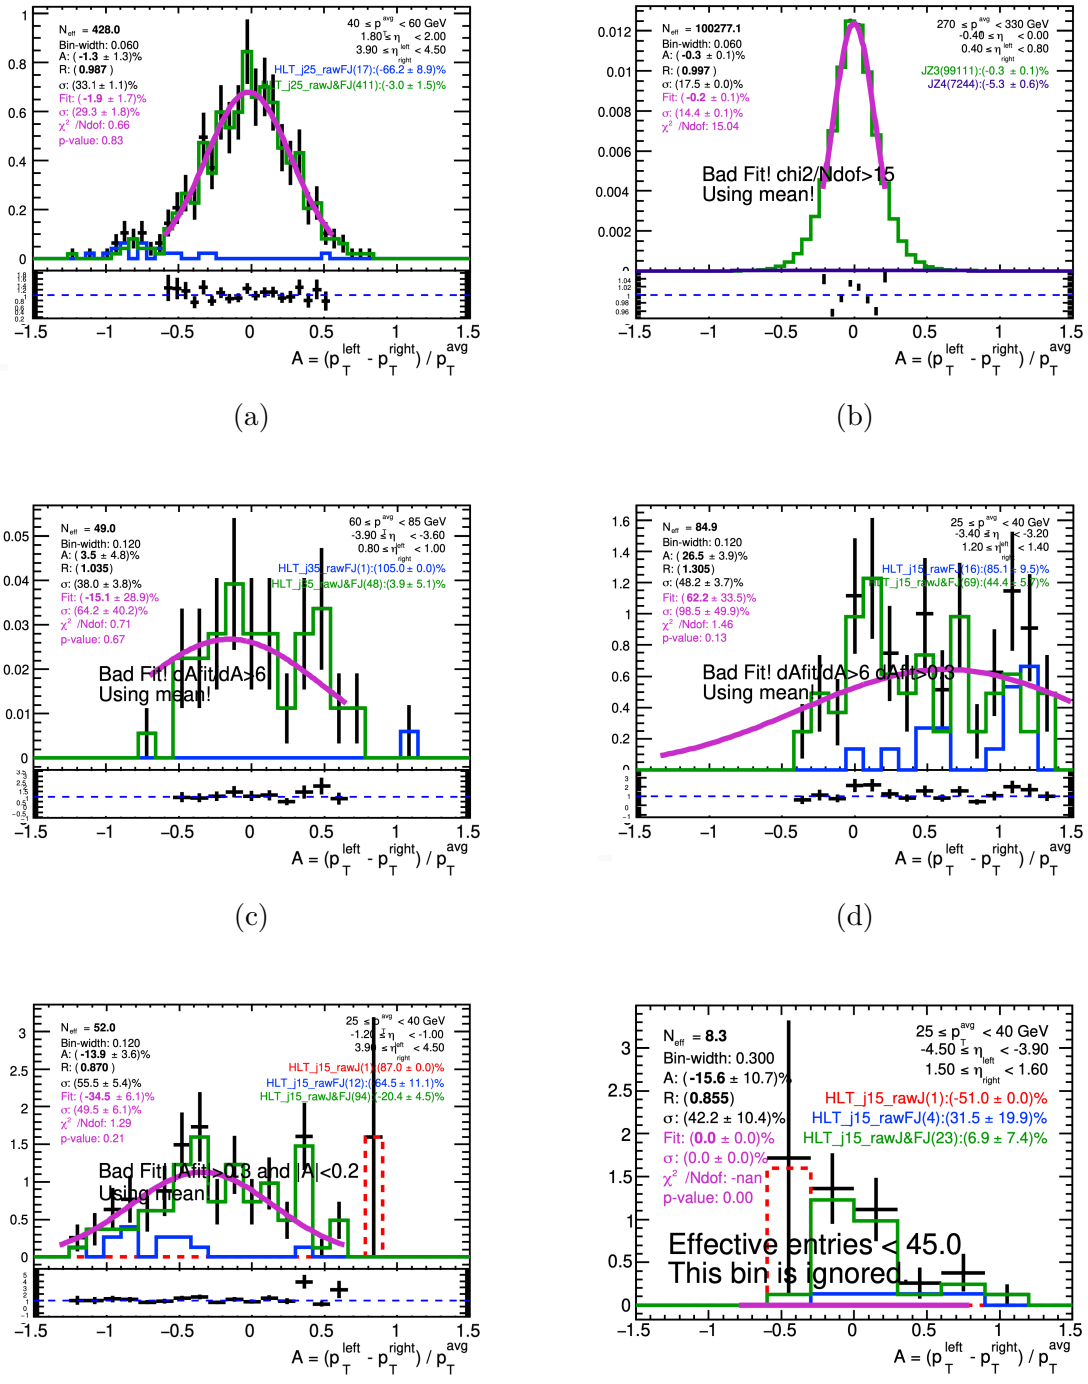
3.3.1 Recovering constraints

3.3.1.1 Discarded constraints

The matrix method is based on the fact that the intercalibration factors are constrained by many contributions corresponding to all the possible combinations of η_{left} and η_{right}. In practice, not all those combinations are kept: several criteria on the quality of the asymmetry distribution need to be met.

- If the effective number of entries N_{eff} is too low (< 45), the bin can either be dumped or merged with its neighbour.
- If the asymmetry fit is not deemed good enough, e.g. if its χ^2/NDF is too large (> 15) or if the uncertainty in the central value is too large, the bin is not dumped but $\langle \mathcal{A} \rangle$ is taken as the arithmetic mean of the distribution, which is a weaker indicator compared to the mean value of the Gaussian fit.
- If the central value of the Gaussian fit is not reasonable, i.e. too far from 0 or too far from the arithmetic mean, the arithmetic mean is taken instead.
- The arithmetic mean is taken if the bin width after the rebinning procedure is too coarse.

Examples of asymmetry bins discarded in the matrix method are shown in Figure 3.27. For a more robust method, the number of constraints coming from the mean value of the Gaussian fit needs to be maximised, without introducing any biased contributions.



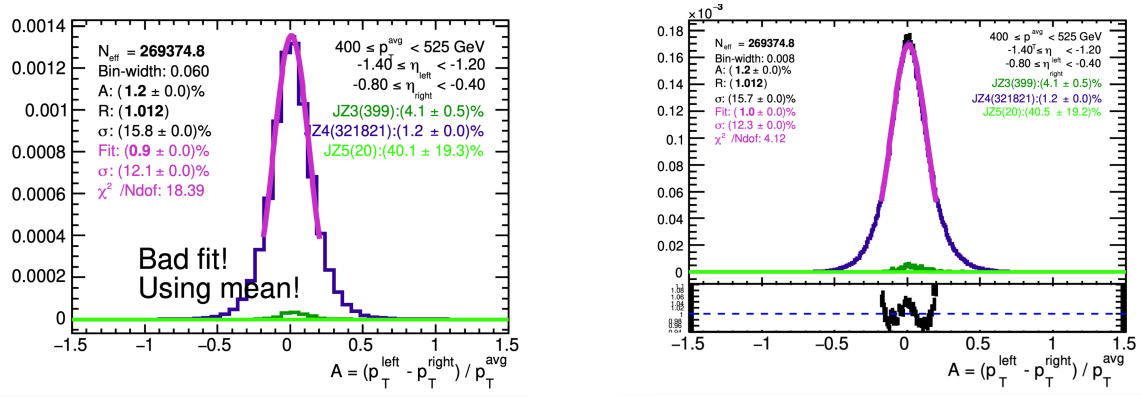
(a) (b) (c) (d) (e) (f)

Figure 3.27: Determination of the mean asymmetry for different ($p_{\mathrm{T}}^{\mathrm{avg}}$, η_{left}, η_{right} in data15, by a Gaussian fit (a) or by the arithmetic mean (b, c, d, e). (a) Example where the fit works (b) Example where the χ^2/NDF of the fit is too large (c,d) Examples where the fit is rejected because the uncertainty in the fit central value is too large (e) Example where the central value of the fit is too different from the arithmetic mean of the distribution (f) Example where the bin is dropped because there are not enough effective entries.

3.3.1.2 Asymmetry binning

One way to recover constraints is to tune the asymmetry binning as a function of the number of effective entries in the bin: when the statistics are high (e.g. at high p_{T}, low $|\eta|$), the χ^2/NDF of the Gaussian fit to the asymmetry distribution becomes very large, even though the fit looks correct. This happens when the bins are too broad: they contain many entries, meaning that the bin represents a wide range of data values. The fitting function is evaluated at the centre of each bin, but since the bin covers a large range, the actual data within the bin may vary significantly from the Gaussian function at the centre. As a result, the difference between the bin content and the value of the Gaussian at the bin's centre becomes larger, leading to a poorer fit. In other words, broad bins reduce the accuracy of the fit by failing to capture the finer structure of the data, which, in the context of the Matrix Method, can result in the exclusion of valid contributions.

Therefore, a finer binning is implemented in this region. The re-binning strategy is also adapted to take advantage of this: the optimal bin width defined in Section 3.2.2.1 is optimised empirically. Both the initial bin width and the optimal bin width have been studied to ensure an optimal constraint recovery. An initial bin width of 0.008 and an optimal bin width of $w_{\mathrm{optimal}} = 1.5 * \mathrm{RMS} * N_{\mathrm{eff}}^{-1/3}$ are found to give the best results. In addition to these optimisations, a ratio plot was added below the asymmetry distribution plots to check for potential biases in the fit. The ratio plot compares the data to the fit, and if a noticeable pattern or deviation appears in this ratio, it indicates that the fit may be biased or inaccurate. An example of a constraint recovered after this optimisation is shown in Figure 3.28. A pattern can be seen in the ratio plot. It is symmetrical around the peak, suggesting that the bias originates from the tails of the distribution, which are not well-represented by the Gaussian fit. These non-Gaussian tails also enhance the χ^2/NDF of the fit. This can be addressed by tuning the fit range.



(a) (b)

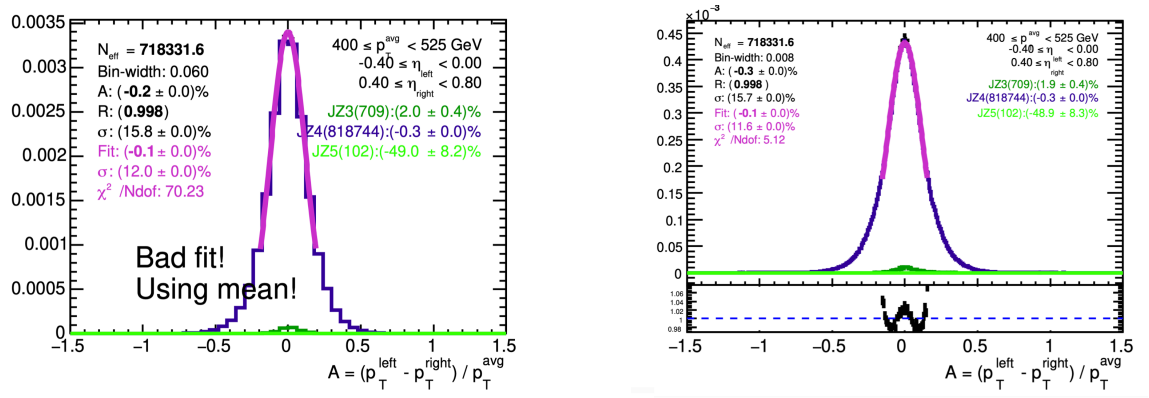
Figure 3.28: Asymmetry distribution for a given $p_{\mathrm{T}}^{\mathrm{avg}}$, η_{left}, η_{right} bin, before (a) and after (b) optimising the binning, as explained in Sec 3.3.1.2.

3.3.1.3 Fit range

The non-Gaussian behaviour of the asymmetry distribution, particularly in the tails, tends to inflate the χ^2/NDF in high-statistics bins. However, since the tails are not crucial for determining the mean value, the fit range can safely be reduced to focus on the peak of the distribution. Therefore, for the last two steps of the iterative procedure, the fit range is narrowed for bins where the effective number of entries exceeds a certain threshold. As a reminder, the fit range was defined as $[\langle\mathcal{A}_i\rangle - \alpha * \Delta\langle\mathcal{A}_i\rangle, \langle\mathcal{A}_i\rangle + \alpha * \Delta\langle\mathcal{A}_i\rangle]$, with $\langle\mathcal{A}_i\rangle$ and $\Delta\langle\mathcal{A}_i\rangle$ the value extracted from the previous fit, and $\alpha = 1.5$. Now, the multiplicative factor α depends on the effective statistic in the bin. The optimal threshold and final fit range have been carefully studied to ensure accurate constraint recovery. The best values are found to be the following:

- If $N_{\text{eff}} > 100,000$; α=0.7,
- If $N_{\text{eff}} > 20,000$; α=1.0,
- else, the previous fit range is kept: α=1.5.

An example of a bin recovered by this process is illustrated in Figure 3.29.

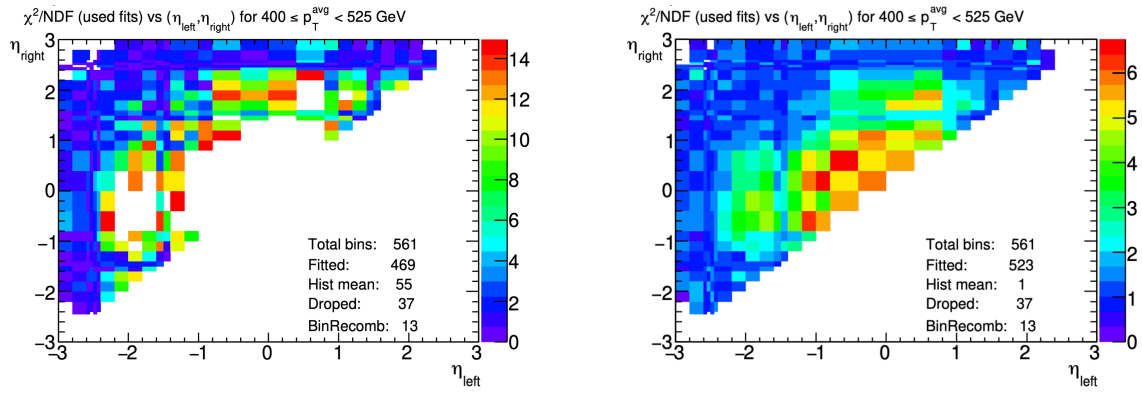


(a) (b)

Figure 3.29: Asymmetry distribution for a given $p_{\text{T}}^{\text{avg}}$, η_{left}, η_{right} bin, before (a) and after (b) optimising the binning and the fit range, as explained in 3.3.1.3.

3.3.1.4 Results

With the implementation of these two improvements, many contributions are recovered, mainly in the central region of the response matrices (high p_{T}, low $|\eta|$) where the statistic is the highest. The result of these two improvements is best illustrated in Figure 3.30, which represents the χ^2/NDF of the asymmetry fit for each η_{left} and η_{right} combination, before and after the constraint recovery, for one $p_{\text{T}}^{\text{avg}}$ bin. An empty bin means that the arithmetic mean was used or that the bin was dumped. In this example, in addition to recovering 54 constraints, the χ^2/NDF are globally reduced, meaning that the method is more reliable.



(a) (b)

Figure 3.30: Distribution of the χ^2/NDF of the Gaussian fit for a given $p_{\mathrm{T}}^{\mathrm{avg}}$ bin, before (a) and after (b) optimisations.

3.3.2 The global χ^2/NDF test

3.3.2.1 Principle

As described in the previous sections, calibration of all detector regions relative to one another is achieved by solving an over-constrained system of linear equations through a χ^2 minimisation. For corresponding calibration coefficients, a global reduced χ^2 can be evaluated for each $p_{\mathrm{T}}^{\mathrm{avg}}$ bin, using the following expression (based on the first term of Equation 3.3):

$$\chi^2 = \sum_{j=1}^{N}\sum_{i=1}^{j-1}\left(\frac{1}{\Delta\langle\mathcal{R}_{ij}\rangle}\left(c_i\langle\mathcal{R}_{ij}\rangle - c_j\right)\right)^2$$

The reduced χ^2 is given by χ^2/n_{dof}, where the number of degrees of freedom n_{dof} is given by the number of used relative response terms minus the number of calibration factors. Since there is one term per bin in the upper half of the $N \times N$ response matrix and N calibration factors, n_{dof} corresponds to $n_{\mathrm{dof}} = (N^2 - N)/2 - N - N_{\mathrm{drop}}$, where N_{drop} is the number of contributions excluded due to insufficient statistics. The global reduced χ^2 provides insight into the consistency of the calibration obtained from the input $\langle\mathcal{R}_{ij}\rangle$ terms and their uncertainties.

To study potential tension within the minimisation, a pull statistic is used, defined by:

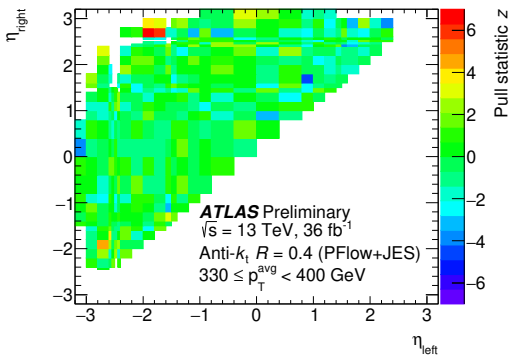
$$z_{i,j} = \frac{1}{\Delta\langle\mathcal{R}_{ij}\rangle}\left(c_i\langle\mathcal{R}_{ij}\rangle - c_j\right)$$

The χ^2 of Equation 3.3 is hence a sum of pulls squared. The pull can reveal inconsistencies between the obtained calibration and the $\langle\mathcal{R}_{ij}\rangle$ terms, and no such test has been performed in the ATLAS η-intercalibration before.

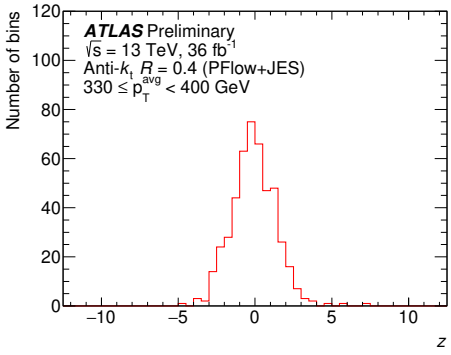
3.3.2.2 Results

The global reduced χ^2 test was performed on Run 2 data and their associated Monte Carlo samples. The results for the combined 2015 and 2016 data sample (called data1516

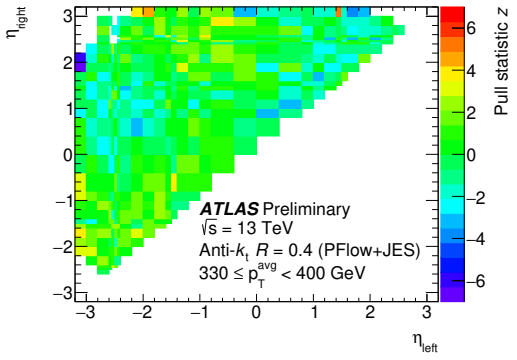
hereafter) for one particular $p_{\mathrm{T}}^{\mathrm{avg}}$ bin are presented in Figure 3.31. Overall, there does not appear to be any significant bias originating from a specific region. Minor tensions are noticeable in some isolated bins (e.g. red bins in Figure 3.31a and 3.31c) and can be traced back to bad evaluations of the average of the asymmetry distribution they originate from. This often occurs due to poor asymmetry fits that still pass the quality criteria, such as double-peaked structures where only one peak is properly fitted, or very asymmetrical distributions. The minor tensions can be explored further to optimise the computation of the mean asymmetry values. A study was conducted where bins contributing significantly to the global reduced χ^2 were excluded, and the minimisation was repeated. This study confirmed that the calibration factors are stable relative to such a change of inputs.



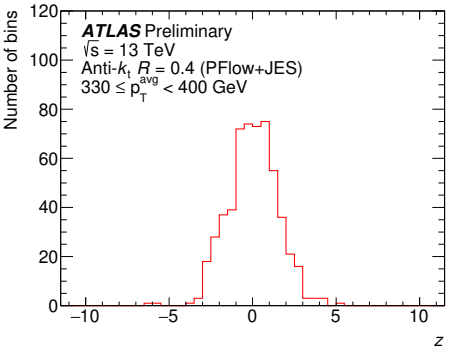
(a) data1516



(b) data1516



(c) Simulations



(d) Simulations

Figure 3.31: Pull statistic for each $(\eta_{\mathrm{left}}, \eta_{\mathrm{right}})$ for $330 < p_{\mathrm{T}}^{\mathrm{avg}} < 400$ GeV for data collected in 2015–2016 (a) and MC (c), and distribution of the pull (b) and (d).

3.3.2.3 Perspective

To reduce inconsistencies in the system, a second minimisation can be performed, omitting the bins that contribute the most to the tension. This second minimisation approach was tested, but it resulted in only marginal improvements in the calibration and, therefore, was dropped. The application of this test also raises the question of whether the uncertainties should be rescaled using the global reduced χ^2 to account for the tensions between con-

tributions, following similar techniques of the Particle Data Group [11]. However, such a quantitative use of the χ^2/NDF would require including systematic uncertainties in the global χ^2 evaluation, along with detailed studies of the corresponding (uncertainties in) the correlations.

3.3.3 The bootstrap method for the evaluation of statistical uncertainties

3.3.3.1 Principle

The Bootstrap method [107] is a technique used for evaluating and propagating statistical uncertainties. It is especially helpful when the analytic linear propagation of statistical uncertainties is challenging or impossible. The principle of the Bootstrap method is to perform the analysis with alternative datasets, each representing a possible dataset that could have been collected under the same conditions. Such alternative datasets are referred to as replicas and lead to corresponding replica measurements.

A set of N replicas that are statistically equivalent to the nominal dataset is created by introducing Poisson perturbations that simulate statistical fluctuations: to each event is assigned a unique and deterministically-generated weight, following a Poisson distribution with a mean of unity. Then, each replica is analysed in the same way as the nominal dataset, resulting in a distribution for each observable. These distributions contain N entries, corresponding to the number of replicas, and they can be treated as the probability distribution function for this observable. An uncertainty in this observable can be retrieved e.g. from the standard deviation.

In addition, the bootstrap method allows the evaluation of the statistical covariance between any two observables a and b as:

$$\mathrm{Cov}_{a,b} = \frac{1}{N-1}\sum_{r=0}^{N}(a_r - \bar{a})(b_r - \bar{b})$$

where N is the number of bootstrap replicas, a_r, b_r are the measured values of the observable for the replica r, and $\bar{a} = \frac{1}{N}\sum_{r=0}^{N} a^r$. The correlation between the measurement is then given by

$$\mathrm{Corr}_{a,b} = \frac{\mathrm{Cov}_{a,b}}{\sqrt{\mathrm{Cov}_{a,a}\mathrm{Cov}_{b,b}}}$$

This correlation matrix encodes information on the correlations between two observables: a value close to 1 means that the observables are 100% correlated, and a value close to -1 means that the observables are 100% anti-correlated.

Moreover, the Bootstrap method is deterministic. Specifically, the seed for the random number generator is based on the event number, ensuring that the results are fully correlated with those from different analyses using the Bootstrap method as well. This feature enables the evaluation of statistical correlations between measurements that use either fully or partially overlapping input data.

3.3.3.2 For the η-intercalibration

In the context of the η-intercalibration, the goal is to use this method to evaluate the statistical uncertainties, taking into account the correlations between different η regions. Correlations between the calibration factors are expected, as each selected event contains a dijet pair that impacts two η bins. The Bootstrap method shall replace the current process, described in Section 3.2.4.2, which does not account for the correlations.

For this purpose, 1000 bootstrap replicas of the asymmetry distributions are computed and the mean asymmetry is retrieved as explained in Section 3.2.2.1. The response matrix is evaluated in bins of $p_{\mathrm{T}}^{\mathrm{avg}}$, and correction factors are deduced with the χ^2 minimisation. As a result, for one $(p_{\mathrm{T}}^{\mathrm{avg}}, \eta_{\mathrm{det}})$ bin, there are $N = 1000$ corrections factors C_i^r for $r \in [\![1, N]\!]$. The statistical uncertainty in the calibration factor for this bin is then determined as the RMS of the C_i distribution.

The covariance matrix between the calibration factors is then defined as:

$$\mathrm{Cov}_{i,j} = \frac{1}{N-1} \sum_{r=0}^{N} (C_i^r - \bar{C}_i)(C_j^r - \bar{C}_j)$$

with $\bar{C}_i = \frac{1}{N} \sum_{r=0}^{N} C_i^r$. Then the correlation matrix is computed as follows:

$$\mathrm{Corr}_{i,j} = \frac{\mathrm{Cov}_{i,j}}{\sqrt{\mathrm{Cov}_{i,i}\mathrm{Cov}_{j,j}}}$$

3.3.3.3 Results

I implemented the Bootstrap method in the η-intercalibration code using the THBootstrap class from ROOT [108–110], to produce statistically correlated distributions of asymmetry, response matrices, intercalibration and calibration factors. Those distributions, their mean and standard deviation, are studied at each intermediate step of the process, and compared with the results of the previous method. The previous method is applied to the nominal distribution and is referred to as such in the following. An example of each is given in Figure 3.32. The mean values of those distributions are compatible with the nominal values, and the standard deviations are about two times larger than the uncertainty propagated from the fit to the nominal asymmetry distributions, as indicated in Table 3.3. This is the case for most $(\eta_{\mathrm{left}}, \eta_{\mathrm{right}})$ bins, as shown in Figure 3.33a and Figure 3.33b, where the ratio is generally below unity, indicating that the uncertainty from the fit to the asymmetry distribution underestimates the statistical uncertainty of the mean asymmetry.

The distributions of the mean asymmetries obtained from the asymmetry fits exhibit non-Gaussian tails (as illustrated in Figure 3.32), and their widths tend to be larger than the uncertainty obtained from the fit to the nominal distribution. However, the previous method assumed these distributions were Gaussian, with widths matching the uncertainties in the central values derived from the asymmetry fits. As a result, the statistical uncertainties in the calibration factors yielded by the bootstrap method are found to be up to two times larger than with the previous method, the enhancement varying with both $p_{\mathrm{T}}^{\mathrm{avg}}$ and η. This indicates that the previous method underestimated

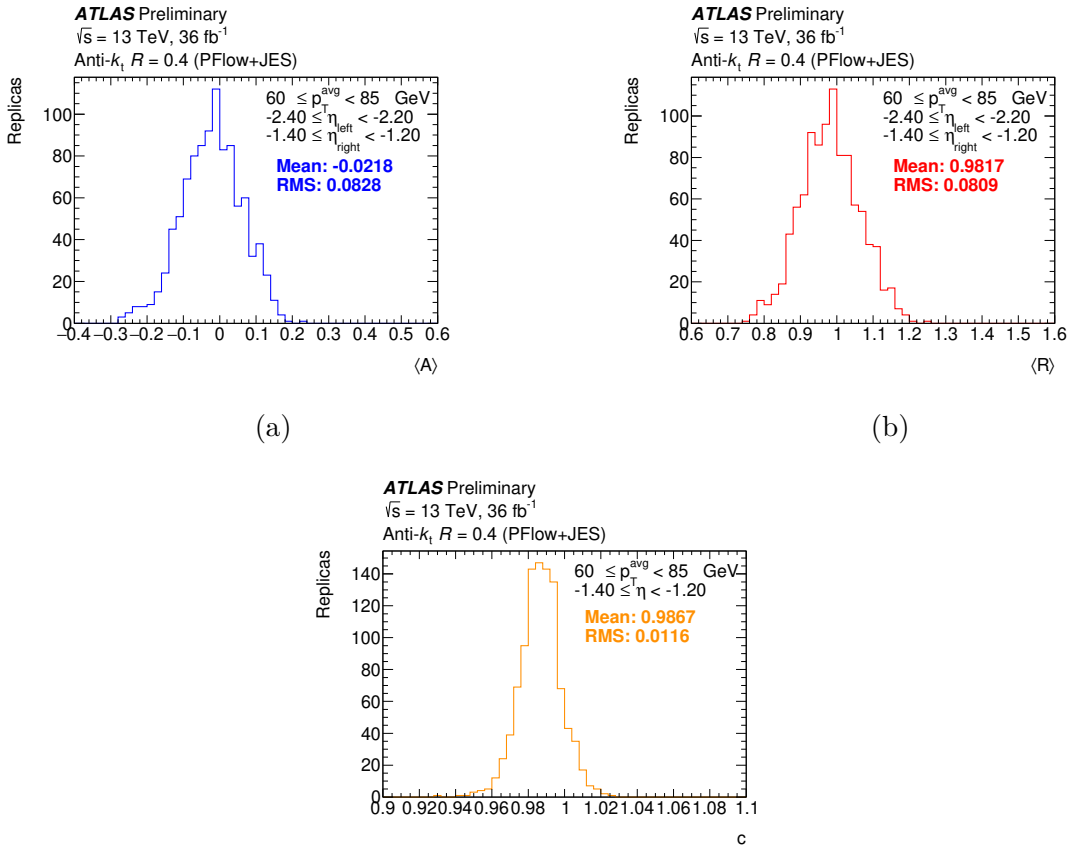
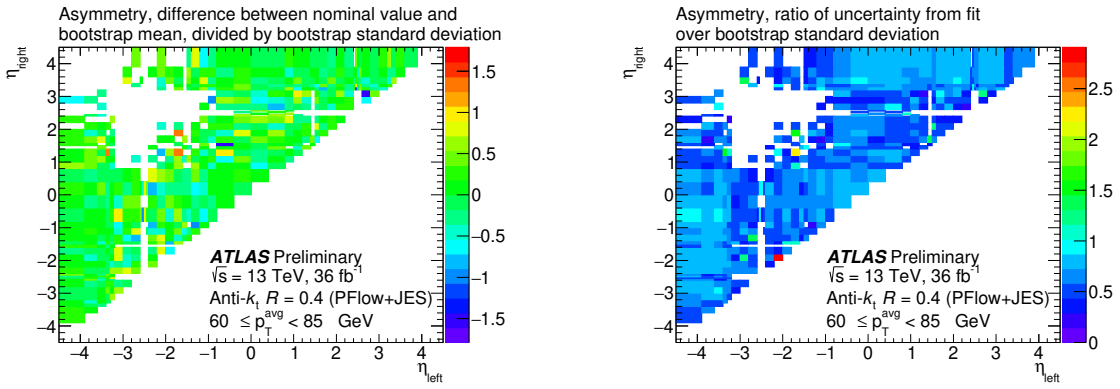


Figure 3.32: Distribution of 1000 Bootstrap replicas at intermediate steps of the η-intercalibration process, for $60 < p_{\mathrm{T}}^{\mathrm{avg}} < 85$ GeV. Mean asymmetry (a) and response (b) distributions are shown for $-2.40 < \eta_{\mathrm{left}} < -2.20$, $-1.40 < \eta_{\mathrm{right}} < -1.20$. (c) Intercalibration factor distribution is shown for $-1.40 < \eta_{\mathrm{det}} < -1.20$.

	Value $\pm$ uncertainty from fit	Bootstrap mean $\pm$ std. dev.
Mean asymmetry	0.0108 ± 0.0460	-0.0218 ± 0.0828
Relative response	1.0108 ± 0.0465	0.9817 ± 0.0809
η-intercalibration factor	0.9898 ± 0.0041	0.9867 ± 0.0116

Table 3.3: Comparison of the mean asymmetry $\langle\mathcal{A}\rangle$, relative response $\langle\mathcal{R}\rangle$, and η-intercalibration factor c obtained from the fit to the asymmetry distributions, and the means and standard deviations of the associated bootstrap distributions. The results correspond to the bins illustrated in Figure 3.32.

the statistical uncertainties in comparison with the bootstrap method. Examples for three $p_{\mathrm{T}}^{\mathrm{avg}}$ bins are shown in Figure 3.34. Given that the Bootstrap method does not rely on the strong assumptions of the previous method, and that the study demonstrates that those assumptions were incorrect, it can be confidently concluded that the Bootstrap method is more accurate and should be used from now on.



(a) (b)

Figure 3.33: (a) The difference between the mean asymmetry derived from the nominal dataset with the average of the mean asymmetry of the 1000 replicas, divided by the Bootstrap RMS, for every $(\eta_{\text{left}}, \eta_{\text{right}})$ combination. The difference is close to zero, indicating no bias in the Bootstrap method. (b) The ratio of the uncertainties obtained with the two methods (previous/Bootstrap), for every $(\eta_{\text{left}}, \eta_{\text{right}})$ combination. The empty bins correspond to cases where the lack of statistics prevents a reliable evaluation of the mean asymmetry. Figure for $60 < p_{\text{T}}^{\text{avg}} < 85$ GeV.

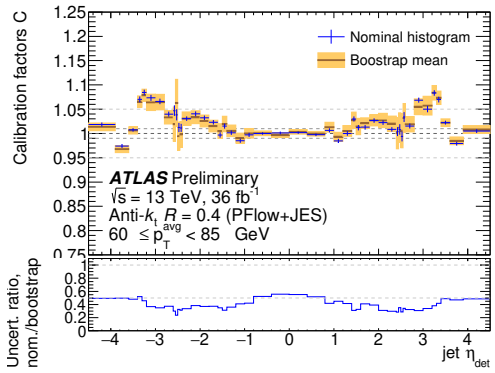
The Bootstrap method also allows the derivation of correlation matrices for each $p_{\text{T}}^{\text{avg}}$ bin. An example is shown in Figure 3.35. Correlation between the bins of high η and anti-correlations between the bins of low η can be seen, especially at low $p_{\text{T}}^{\text{avg}}$. Several effects can explain this:

- One mean asymmetry value $\langle \mathcal{A} \rangle$ constrains two different η bins,
- The normalisation tends to shift matrix entries to positive values,
- Some bins have higher statistical significance than others,
- The strong anti-correlations in the 4 central bins are due to the normalisation of the central part such that their mean is equal to 1.

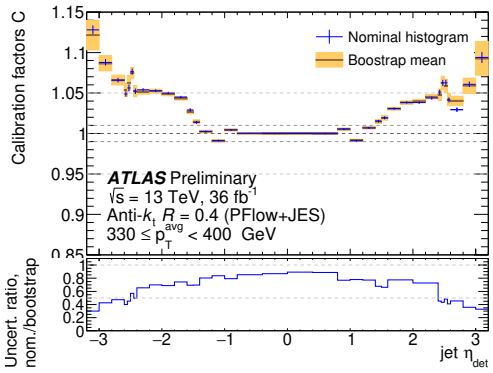
Until now, the correlation matrices had never been evaluated, and the correlations had not been considered in the computation of the statistical uncertainties of the η-intercalibration method. These (anti-) correlations between the η-intercalibration factors affect the statistical uncertainties associated with the method and should be taken into account to prevent double-counting of uncertainties. Furthermore, this method allows the correlation of the statistical uncertainties of the η-intercalibration with those of any other analysis that uses overlapping datasets.

3.3.3.4 Application in physics analysis

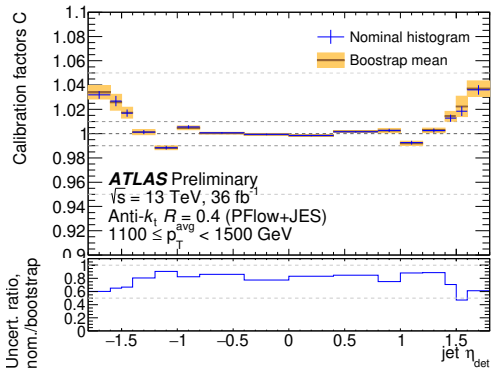
In most physics analyses using jets, the events used to perform the study overlap with those used for the η-intercalibration. Namely, the same jet triggers are used. For this reason, the statistical fluctuations in the sample impact the outcome at two levels: event counting and calibration. Using the bootstrap replicas generated at the event level during the calibration and propagating the statistical uncertainties from there in the analysis



(a) $60 \leq p_{\mathrm{T}}^{\mathrm{avg}} < 85$ GeV



(b) $330 \leq p_{\mathrm{T}}^{\mathrm{avg}} < 400$ GeV



(c) $1100 \leq p_{\mathrm{T}}^{\mathrm{avg}} < 1500$ GeV

Figure 3.34: Calibration factors for various $p_{\mathrm{T}}^{\mathrm{avg}}$ ranges. The statistical uncertainty propagated with the previous approach is in dark blue, while the statistical uncertainty propagated with the Bootstrap Method is in yellow. The bottom panel displays the ratio of the statistical uncertainties evaluated with the two approaches.

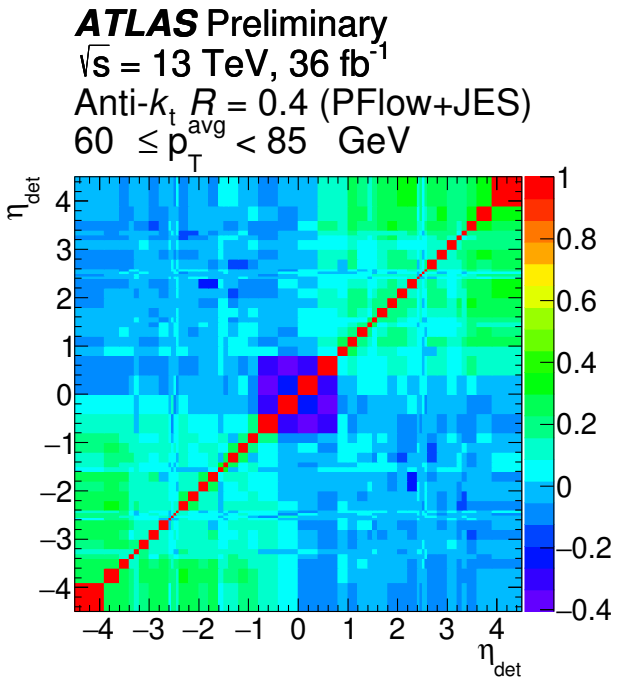


Figure 3.35: Correlation matrix for $60 < p_{\mathrm{T}}^{\mathrm{avg}} < 85$ GeV. Correlation between the bins of high η and anti-correlations between the bins of low η can be seen.

allows for properly accounting for these correlations, i.e. to evaluate and propagate them.

The following procedure is proposed for analysers who wish to use correlated uncertainties for the η-intercalibration factors.

1. For each data taking period (2015-2016, 2017, 2018), for each $p_{\mathrm{T}}^{\mathrm{avg}}$ bin, I provide a Bootstrap histogram (i.e. an object that stores the nominal histogram and the 1000 replicas), each replica providing a set of correction factors derived by doing

$$\mathcal{C}_{\mathrm{replica}_i}/\mathcal{C}_{\mathrm{nominal}} = \Big(\frac{c_{\mathrm{MC}}}{c_{\mathrm{data}}}\Big)_{\mathrm{replica}_i} / \Big(\frac{c_{\mathrm{MC}}}{c_{\mathrm{data}}}\Big)_{\mathrm{nominal}}$$

2. The analyser creates a Bootstrap histogram with 1000 replicas for the variable of interest.

3. The analyser fills each replica with the jets p_{T} divided by the appropriate correction factors $\mathcal{C}_{\mathrm{replica}_i}(p_{\mathrm{T}}^{\mathrm{avg}}, \eta)$ (from the right data taking period), from the i^{th} replica of my Bootstrap histogram.

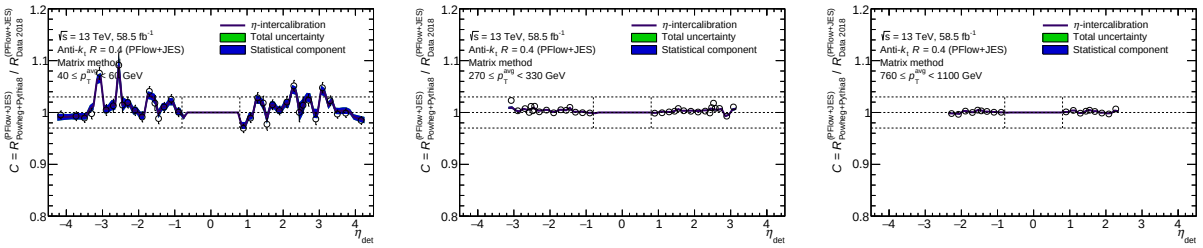
The results on correlations are considered robust to minor differences that may be found between the analysis and the calibration, such as jet collection, recommendation, version of the software, etc. These have little impact on the event selection, and any difference in the result on the correlation can be considered as an uncertainty in the statistical uncertainties. This procedure was used in the context of the inclusive-jet cross-section measurement and is described in Chapter 4.

3.4 Deriving the official η-intercalibration corrections and uncertainties

To become an ATLAS author, one needs to complete a "qualification task" (QT), which consists of performing a task that is beneficial to the whole ATLAS community, such as performance studies. In the context of my qualification task, I worked on providing pre-recommendations, which are early results of the calibration derived using the latest version of the ATLAS software, to be used by analyses using jets. In particular, my goal was to derive an uncertainty in the calibration factors from the η-intercalibration, linked to an update of the global analysis framework of the ATLAS collaboration, Analysis Base, which was updated from release 21 (R21) to release 22 (R22). When I produced the pre-recommendation, the improvements to the method described in the last sections were implemented, except for the Bootstrap Method. All samples and calibrations were provided by the Jet/ETMiss group.

3.4.1 Pre-recommendations phase-I

In the context of the pre-recommendations, early results are provided using the new version of the software. Major changes are not anticipated, so deriving the calibration from scratch is not necessary, and a non-closure uncertainty is derived to cover any difference between the calibration derived with R21 and R22. For this, the η-intercalibration code is ported to R22 and the calibration is derived on samples calibrated with the R22 MCJES (Section 3.1.3) and GSC (Section 3.1.4) and with the R21 in situ calibration. Examples of calibration factors obtained for a low-p_{T} bin, a medium-p_{T} bin and a high-p_{T} bin of 2018 data are shown in Figure 3.36. Non-closure (i.e. deviation from unity) is observed at low-p_{T} while good closure is achieved at higher p_{T}. The difference to 1 of the output



(a) $40 < p_{\mathrm{T}}^{\mathrm{avg}} < 60$ GeV (b) $270 < p_{\mathrm{T}}^{\mathrm{avg}} < 330$ GeV (c) $760 < p_{\mathrm{T}}^{\mathrm{avg}} < 1100$ GeV

Figure 3.36: Calibration factors at (a) low-p_{T} bin, (b) medium-p_{T} bin, (c) high-p_{T} bin, derived from a sample calibrated with the R21 in situ calibration, for 2018 data and associated MC16e. It shows the non-closure of the R22 η-intercalibration relative to the R21 η-intercalibration. Non-closure is observed at low-p_{T} while good closure is achieved at higher p_{T}.

calibration factors is taken as the non-closure uncertainty for the pre-recommendation phase-I: $\Delta\mathcal{C}(p_{\mathrm{T}}, \eta) = \mathcal{C}(p_{\mathrm{T}}, \eta) - 1$.

It was decided that a unique uncertainty would be provided for all data years. Thus, the global non-closure uncertainty is the sum of the non-closure uncertainty for each year,

weighted by their relative luminosity:

$$\Delta\mathcal{C}_{\mathrm{TOT}} = \frac{l_{15}\cdot|\Delta\mathcal{C}_{15}| + l_{16}\cdot|\Delta\mathcal{C}_{16}| + l_{17}\cdot|\Delta\mathcal{C}_{17}| + l_{18}\cdot|\Delta\mathcal{C}_{18}|}{l_{15}+l_{16}+l_{17}+l_{18}}$$

with l_{15} resp. l_{16}, l_{17} and l_{18} the luminosity for the 2015 resp. 2016, 2017, and 2018 years of data taking. An additional Gaussian kernel smoothing was applied to this result in a ($p_{\mathrm{T}}^{\mathrm{avg}}$, η) plane, with a strong smoothing effect to provide users with highly smooth uncertainties. This is needed because sharp transitions in uncertainties can affect analyses, potentially causing convergence or fit stability issues. The result as a function of $p_{\mathrm{T}}^{\mathrm{avg}}$ and η is shown in Figure 3.37. It demonstrates small non-closure uncertainties across the plane, with larger uncertainties at low p_{T} and high η.

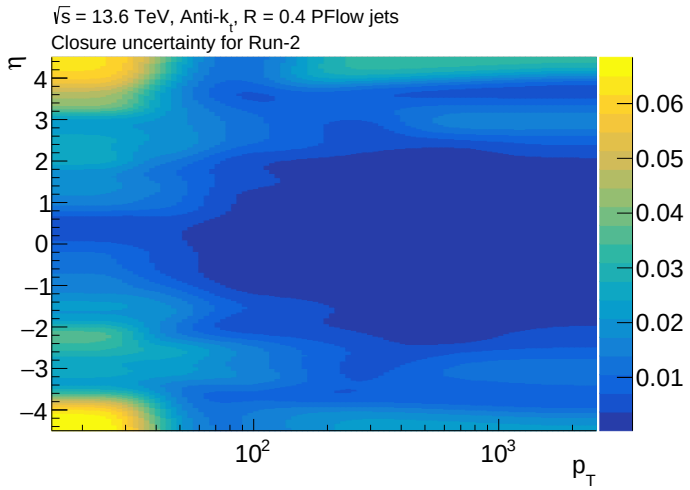


Figure 3.37: Smoothed non-closure uncertainty for the full Run 2.

After this, a non-closure uncertainty was also derived for the other two in situ methods, and the pre-recommendation phase-I result was propagated to the collaboration.

3.4.2 Pre-recommendations phase-II

Phase-II of the pre-recommendations provides a preliminary calibration for analysers to work with early Run 3 data (2022 and 2023 campaigns). A new MC-JES is derived using the latest Monte Carlo samples (MC23) that emulate the Run 3 conditions. For the in situ part, similarly to phase-I, a non-closure uncertainty linked to the application of the Run 2 R21 in situ calibration to Run 3 data is derived. To achieve this, Run 3 samples (2021 and 2022 years of data-taking and MC23) are calibrated with the new MC-JES and the R21 in situ calibration, and the calibration factors are derived. The difference to one is taken as a non-closure uncertainty.

To work with Run 3 data, a trigger efficiency study must be conducted, since several significant improvements were introduced during the long shutdown between Run 2 and Run 3. In particular, the central HLT triggers now use PFlow HLT jets for the partial reconstruction. The Run 3 trigger menu was also updated (see Table 3.4 requiring the development of a new trigger strategy for the η-intercalibration.

3.4.2.1 Trigger efficiency study

Principle In order to avoid biases when selecting events using triggers, it is essential to ensure that each trigger is used only within the region where it is fully efficient. For this purpose, a trigger efficiency study has to be performed. The efficiency of a given trigger (referred to as the probe trigger) is determined using an unbiased reference dataset relative to the trigger being evaluated. The efficiency is calculated by measuring the proportion of events from this reference dataset that pass the probe trigger, as a function of the relevant kinematic variable in the analysis:

$$\text{Eff} = \frac{N_{\text{probe}} \cap N_{\text{reference}}}{N_{\text{reference}}}$$

The representation of the efficiency as a function of the kinematic variable of the analysis is called a turn-on curve. In the case of the η-intercalibration, it is evaluated as a function of $p_{\text{T}}^{\text{avg}}$. A trigger is considered fully efficient when its efficiency is at least 99%. The associated $p_{\text{T}}^{\text{avg}}$ value is called the turn-on p_{T}.

To construct an unbiased dataset, a reference trigger can be used. It should be independent of the signal or probe trigger, provide a sufficiently large dataset to ensure good statistical precision, and have significant overlap with the dataset selected by the signal trigger. Using a minimum bias trigger as the reference trigger is the most unbiased approach because it triggers on random events. In particular, it operates independently of the calorimeter trigger system, making it completely separate from the probe trigger. However, only a limited number of minimum bias events are recorded, which restricts the p_{T} reach of this trigger. Additionally, the probability of a low-p_{T} trigger and a minimum bias trigger firing simultaneously is quite low, as both triggers have high prescales.

Another way to construct a reference dataset for a trigger efficiency study is with the so-called Bootstrap method (NB: despite the similar name, this is different from the bootstrap method for the evaluation of the statistical uncertainties, described in Section 3.3.3). In this method, the efficiency of the trigger with the lowest threshold is evaluated with an independent trigger, such as a minimum bias trigger. That trigger can then be used as a reference trigger to evaluate the efficiency of a trigger with a kinematic threshold higher than the turn-on p_{T} of the first trigger. This process is done iteratively until the jet trigger with the highest threshold is reached. This method maximises the available statistics, since the prescales decrease with increasing p_{T} thresholds of the triggers.

To further enhance the statistical precision of the trigger efficiency study, the emulation method is used. The aim of this method is to recover events where a jet meets the kinematic threshold of the probe trigger but does not activate the trigger due to prescaling. This can happen when the event has fired another trigger, but still fulfils the conditions of the trigger under study. Since the HLT triggers are prescaled independently, only a limited number of events fire both the probe and the reference trigger. Thus, instead of the trigger information, the information on the L1 trigger before prescale, together with the information on the HLT jets (both available in the event record), can be used. The HLT jets are reconstructed from partial information and partially calibrated to assess the trigger decision during a run. They enable the emulation of the HLT decision by requiring that at least one HLT jet pass the trigger threshold. This method is only limited by the

number of events firing the reference trigger and is independent of the probe jet trigger prescale. With this method, the efficiency is defined by:

$$\text{Eff} = \frac{N(\text{fire reference trigger}) \cap N(\text{pass L1 trigger before prescale \& pass probe emulation})}{N(\text{fire reference trigger})}$$

Results For the η-intercalibration, dijet events with jets that span the whole η range of the detector are needed, so the jets are required to fire a forward jet trigger or a central trigger, or both (as described in Section 3.2.1.2. Therefore, for the trigger efficiency study, the probe trigger combines a central and a forward trigger, with comparable p_{T} thresholds.

The trigger efficiency study was initially conducted for the central triggers, followed by the forward triggers, to identify a combination of triggers with a similar turn-on p_{T}. Then, the trigger efficiency study was done for the combination. An issue with the low-p_{T} triggers during the early Run 3 prevented the use of the first two triggers, so the study starts with the third lowest p_{T} trigger. The resulting turn-on curves for data22 are presented in Figure 3.38. The turn-on p_{T} is estimated as the $p_{\text{T}}^{\text{avg}}$ for which the efficiency crosses

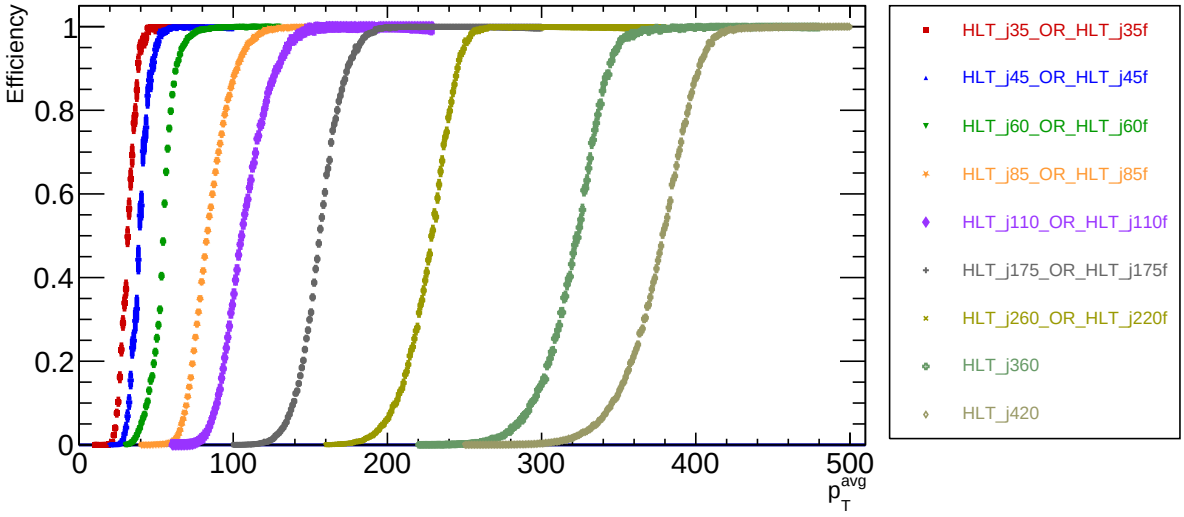


Figure 3.38: Trigger efficiencies as a function of $p_{\text{T}}^{\text{avg}}$ for year 2022 of data-taking. Each trigger combination is shown in a different colour and with a different symbol. The turn-on p_{T} is defined as the $p_{\text{T}}^{\text{avg}}$ for which the efficiency crosses 0.99. The full name and characteristics of the triggers are shown in Table 3.4.

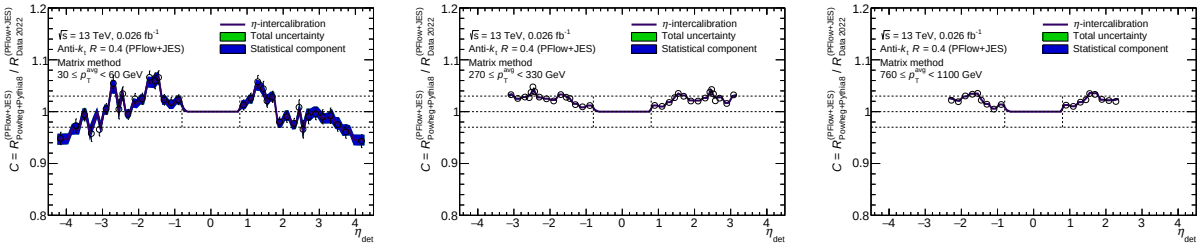
0.99. For safety, a trigger is considered efficient for a p_{T} that is around 10GeV above the turn-on p_{T}, and turn-on curves are produced in a $(\eta, p_{\text{T}}^{\text{avg}})$ plane to ensure that the trigger is considered on its efficiency plateau in the whole η range.

3.4.2.2 Phase-II pre-recommendations results

A new $p_{\text{T}}^{\text{avg}}$ binning strategy is adopted to match the Run 3 trigger menu, following the recommendations presented in Section 3.2.1.4. To circumvent the issue with the low-p_{T}

triggers, a zero-bias sample is used for the two lowest $p_{\mathrm{T}}^{\mathrm{avg}}$ bins. It corresponds to a sample triggered randomly with a very high prescale. It contains just enough data to perform the η-intercalibration at low-p_{T} as a temporary solution for the pre-recommendations. Since the p_{T} spectra fall rapidly, this sample contains very few events with p_{T} high enough for the second p_{T} bin: the binning is adapted to ensure sufficient statistics in those two bins.

Similarly to what was done for the phase-I of pre-recommendation, the η-intercalibration is derived from the Run 3 sample. Example of calibration factors derived for the non-closure are shown in Figure 3.39. Non-closure is observed at low-p_{T} while good closure is achieved at higher p_{T}. The difference to 1 is taken as a non-closure uncertainty for the



(a) $30 < p_{\mathrm{T}}^{\mathrm{avg}} < 60$ GeV (b) $270 < p_{\mathrm{T}}^{\mathrm{avg}} < 330$ GeV (c) $760 < p_{\mathrm{T}}^{\mathrm{avg}} < 1100$ GeV

Figure 3.39: Calibration factors at (a) low-p_{T} bin, (b) medium-p_{T} bin, (c) high-p_{T} bin, derived from a sample calibrated with the R21 in situ calibration, for 2022 data and associated MC23a. It shows the non-closure of the Run 3 η-intercalibration relative to the R21 Run 2 η-intercalibration. The open symbols show the discrete η-intercalibration factors from which the smoothed calibration is deduced.

phase-II of pre-recommendations. The 2022 and 2023 years of data-taking are combined to provide a single uncertainty, with a similar approach to phase-I:

$$\Delta\mathcal{C}_{\mathrm{TOT}} = \frac{l_{22} \cdot |\Delta\mathcal{C}_{22}| + l_{23} \cdot |\Delta\mathcal{C}_{23}|}{l_{22} + l_{23}}$$

The resulting uncertainty is smoothed in the $(\eta, p_{\mathrm{T}}^{\mathrm{avg}})$ plane with a Gaussian Kernel smoothing, to avoid sharp transitions in uncertainties, which can be detrimental to analysis. The non-closure uncertainty for the early Run 3 is shown in Figure 3.40. It shows larger non-closure at high η and in the transition regions, especially at low $p_{\mathrm{T}}^{\mathrm{avg}}$.

3.5 Conclusion

Jets are ubiquitous in the hadron collider environment and often represent either signal or background in most analyses. As such, accurate jet calibration is crucial and must be continuously monitored and improved. This chapter presented several studies aimed at improving the reliability of the jet η-intercalibration. In particular, the overconstrained system used to derive the calibration factors has been strengthened by incorporating additional terms and introducing a compatibility test between them. Furthermore, a new, more robust method for estimating the statistical uncertainties in the calibration factors has been implemented. This approach results in larger, more realistic uncertainties

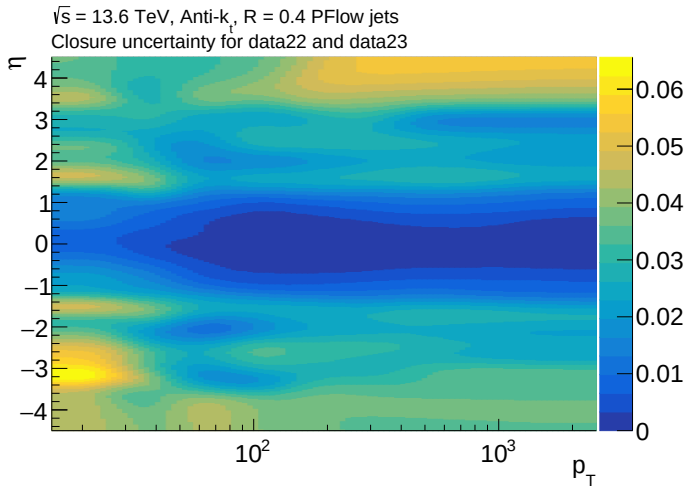


Figure 3.40: Smoothed non-closure uncertainty for the 2022 and 2023 datasets. At low-$p_{\mathrm{T}}^{\mathrm{avg}}$, a zero-bias sample with lower statistics is used.

and enables the study of their correlations among different phase-space regions and with other measurements. Finally, the η-intercalibration step of the JES has been derived for direct application by analysts working with Run 2 and Run 3 data. This required the development of a new trigger strategy for the application of the calibration to Run 3 data.

Central Trigger name	Preselection	L1 criterion	HLT criterion
HLT_j0_perf_L1RD0_FILLED	0	0	0
HLT_j15_pf_ftf_L1RD0_FILLED	0	0	15
HLT_j25_pf_ftf_L1RD0_FILLED	0	0	25
HLT_j35_pf_ftf_L1RD0_FILLED	0	0	35
HLT_j45_pf_ftf_preselj20_L1RD0_FILLED	20	0	45
HLT_j60_pf_ftf_preselj50_L1J20	50	20	60
HLT_j85_pf_ftf_preselj50_L1J20	50	20	85
HLT_j110_pf_ftf_preselj80_L1J30	80	30	110
HLT_j175_pf_ftf_preselj140_L1J50	140	50	175
HLT_j260_pf_ftf_preselj200_L1J75	200	75	260
HLT_j360_pf_ftf_preselj225_L1J100	225	100	360
HLT_j400_pf_ftf_preselj225_L1J100 (data23)	225	100	400
HLT_j420_pf_ftf_preselj225_L1J100 (data22)	225	100	420
Forward Trigger name	**Preselection**	**L1 criterion**	**HLT criterion**
HLT_j15f_L1RD0_FILLED	0	0	15
HLT_j25f_L1RD0_FILLED	0	0	25
HLT_j35f_L1RD0_FILLED	0	0	35
HLT_j45f_L1J15p31ETA49	0	15	45
HLT_j60f_L1J20p31ETA49	0	20	60
HLT_j85f_L1J20p31ETA49	0	20	85
HLT_j110f_L1J30p31ETA49	0	30	110
HLT_j175f_L1J50p31ETA49	0	50	175
HLT_j220f_L1J75p31ETA49	0	75	175

Table 3.4: Run 3 jet triggers table with p_{T} criteria for different triggers. The preselection refers to a cut on calorimeter cells before the HLT jet reconstruction to reduce the timing cost of tracking when using the PFlow algorithm. The central triggers are defined in the $|\eta| < 3.2$ range while the forward triggers are defined in the $3.2 < |\eta| < 4.9$ range. L1RD0_FILLED triggers on bunch crossing, and HLT_j0 requests the presence of a jet with no condition on its p_{T}. For 2022, the first unprescaled trigger is HLT_j420_pf_ftf_preselj225_L1J100, while for 2023 it is HLT_j400_pf_ftf_preselj225_L1J100.

Trigger	**Turn-on $p_{\mathrm{T}}^{\mathrm{avg}}$ (GeV)**	**$p_{\mathrm{T}}^{\mathrm{avg}}$ bin (GeV)**
HLT_noalg_L1ZB	-	20-25
HLT_noalg_L1ZB	-	35-60
HLT_j35_OR_HLT_j35f	46.5	60-85
HLT_j45_OR_HLT_j45f	56.5	85-125
HLT_j60_OR_HLT_j60f	84.5	125-150
HLT_j85_OR_HLT_j85f	120.5	150-175
HLT_j110_OR_HLT_j110f	147.5	175-220
HLT_j175_OR_HLT_j175f	192.5	220-270
HLT_j260_OR_HLT_j260f	274.5	270-330
HLT_j360	377.5	330-400
HLT_j420 (resp HLT_j400)	419.5 (resp 418.5)	400-525
HLT_j420 (resp HLT_j400)	419.5 (resp 418.5)	525-760
HLT_j420 (resp HLT_j400)	419.5 (resp 418.5)	760-1100
HLT_j420 (resp HLT_j400)	419.5 (resp 418.5)	1100-1500
HLT_j420 (resp HLT_j400)	419.5 (resp 418.5)	1500-2000

Table 3.5: Turn-on p_{T} for the Run 3 jet triggers and $p_{\mathrm{T}}^{\mathrm{avg}}$ binning strategy for the η-intercalibration. Due to an issue with low-p_{T} triggers, a zero-bias sample is used for the first two bins. For the highest $p_{\mathrm{T}}^{\mathrm{avg}}$ bins, HLT_j420 is used in 2022, while HLT_j420 is used in 2023.

Chapter 4

The Inclusive Jet Cross-Section Analysis

Introduction

Jets are ubiquitous in high-energy proton-proton collision environments, making them one of the main observables at the LHC. They enable a wide range of physics studies, spanning from low-energy non-perturbative QCD to high-energy QCD measurements, while also playing a role in searches for physics beyond the Standard Model. The jet production cross-section measurement is one of the high-energy QCD measurements. It represents the probability of emission of a jet in proton-proton collisions at a given centre-of-mass energy. Experimentally, this translates to counting jets and normalising the measurement by luminosity. These measurements have multiple uses. First, they enable the precise evaluation of Standard Model parameters, such as the strong coupling constant α_s or the determination of Parton Density Functions (PDFs). Additionally, they can be used to optimise the tuning of free parameters in Monte Carlo generators. These measurements can also be compared with theoretical predictions to test the Standard Model and indirectly probe physics beyond it. Any discrepancy between the measurement and theory predictions could hint at new physics effects in virtual loops.

4.1 The inclusive jet double-differential cross-section measurement

Various observables can be used to study the jet production cross-section, including the inclusive jet production double-differential cross-section. This chapter focuses on small-R jets ($R = 0.4$) reconstructed using the anti-k_T algorithm, regardless of their flavour. In this measurement, each jet in an event is considered individually, as a function of jet p_T in eight equal-size bins of the absolute jet rapidity $|y|$. Only jets in the kinematic range $p_T > 55$ GeV and $|y| < 4.5$ are considered. The inclusive jet production cross-section is defined as the ratio of the number of jets in data N_{jets} within a given range of transverse momentum Δp_T and rapidity Δy, to the integrated luminosity $\mathcal{L}$:

$$\frac{\mathrm{d}^2\sigma}{\mathrm{d}p_T \mathrm{d}y} = \frac{N_{\text{jets}}}{\mathcal{L}\Delta p_T \Delta y}.$$

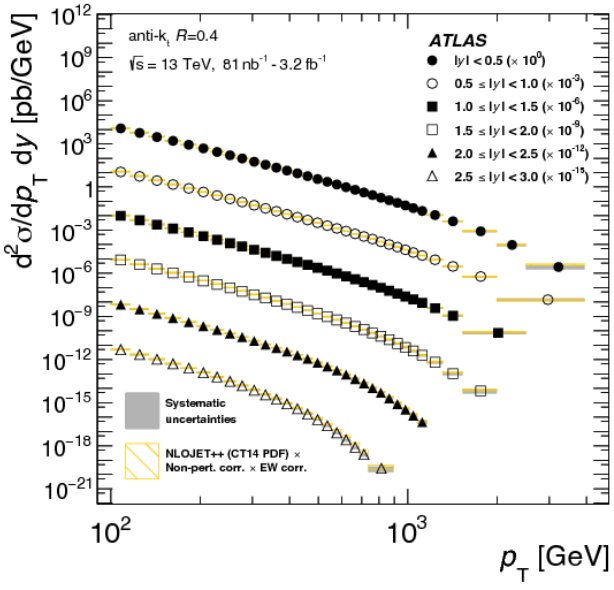


Figure 4.1: Inclusive jet cross-section as a function of p_{T} and $|y|$, for anti-k_{T} jets with $R = 0.4$. The statistical uncertainties are smaller than the size of the symbols used to plot the cross-section values. The shaded areas indicate the experimental systematic uncertainties. Taken from [109].

It has been thoroughly studied in proton-proton collisions, both in ATLAS [109–111] and CMS [112, 113]. Figure 4.1 shows the latest results of the ATLAS collaboration. Historically, this measurement has exhibited significant tensions in the quantitative comparisons between experimental results and theoretical predictions, both in ATLAS and CMS. In particular, while there is good agreement within individual $(p_{\mathrm{T}}, |y|)$ bins, the global χ^2/NDOF, which accounts for the full bin range and their correlations, is remarkably high. This points to the importance of accounting for the correlations, and multiple ideas have been tested to reduce this tension from the experimental side. For example, even if the amplitude of this uncertainty is at the few-permil level, having hundreds of local statistical uncertainties for the η-intercalibration instead of one single fully correlated uncertainty over the full (p_{T}, η) range has reduced the χ^2 by more than 200 units. This gives hope that the new treatment of the η-intercalibration statistical uncertainties, presented in Chapter 3 Section 3.3.3, could have a significant impact on the inclusive jet production cross-section analysis.

4.2 Dataset

4.2.1 Data

The measurement uses proton-proton collision data at a centre-of-mass energy of $\sqrt{s} = 13$ TeV recorded by the ATLAS detector during the Run 2 data-taking period, in high pile-up conditions. The 2015 and 2016 samples are combined and treated as a whole, alongside the 2017 and 2018 samples. The total integrated collected luminosity is presented in

Table 4.1.

2015	2016	2017	2018	Total
3.24 ± 0.04	33.40 ± 0.30	44.63 ± 0.50	58.79 ± 0.64	140.07 ± 0.64

Table 4.1: Per-year and total integrated luminosities of the Run 2 of LHC, recorded by the ATLAS detector, in fb^{-1}. Values are from Reference [58]. The relative uncertainty in the total integrated luminosity amounts to 0.46%.

Jets are reconstructed from PFlow objects [86] using the anti-k_{T} algorithm [43] with parameter $R = 0.4$ and calibrated following the most recent recommendation of [93], with single-particle calibration (Section 3.1.5.6) replacing the MJB in situ method (Section 3.1.5.3), offering better precision relative to earlier derivations. For the η-intercalibration part, the calibration of 2017 is applied to the 2018 sample and a non-closure uncertainty is added for the 2018 sample to account for any discrepancy. Events are required to pass multiple selections to ensure the good quality of the sample:

- Events that may be affected by data-taking issues, whether in LHC or ATLAS, are discarded using the ATLAS Good Run List,
- Corrupted events with jets landing in parts of the detector affected by local problems are removed with an acceptance cut,
- To ensure good quality jets, a primary vertex is required, and must have at least two tracks in the tracker and two track particles associated with it,
- The event must pass at least one trigger from the trigger menu, which consists of single-jet triggers with thresholds varying from 55 GeV to 360 GeV, considered in p_{T} regions where their efficiency is above 99.9% (see Section 3.4.2.1),
- The jet must be considered in-time relative to the collision that produced the event. This is achieved via a simple cut $|t_{\text{jet}}| < 10$ ns where t_{jet} is computed relative to the event time recorded by the trigger (which is synchronised with the bunch-crossing time),
- Jet cleaning, to reject fake jets in the high η and high p_{T} region. It relies on a simple set of selection criteria based on noise-sensitive variables. The jet cleaning process has been tailored to this analysis and is thoroughly presented in [114],
- The transverse momentum of the jet is required to be above 100 GeV and the absolute rapidity is required to be below 4.5. For the unfolding procedure, an additional p_{T} bin extending down to 70 GeV is introduced to account for possible migrations of jets into or out of the region below 100 GeV.

These selection criteria can introduce inefficiencies. The inefficiencies from the jet timing selection and the trigger selection are found to be minimal and consistent between data and Monte Carlo. The jet cleaning induces a larger inefficiency ($\sim 2\%$ over the entire jet p_{T} range) and therefore requires a correction. This is done at the unfolding step, described in Section 4.3. No additional uncertainty due to this correction is assumed [89].

The dataset is then split into nine equally spaced absolute rapidity bins $|y|$ of width 0.5.

4.2.2 MC simulations

Simulated data is used for the correction of detector effects. It is produced using Pythia 8.230 [21, 22], including matrix elements for the leading order processes of dijet production interfaced to a p_{T}-ordered parton shower. The interaction of the particles in the ATLAS detector is fully simulated using Geant4 [90]. As described in Section 2.3.4, the distribution of the average number of pile-up interactions in simulations is re-weighted during data analysis to match that observed in the Run 2 data.

At the reconstructed level, events are required to fulfil the same conditions as for the data sample, except for the GRL, the corrupted events and the trigger selection. An MC cleaning criterion is added: the event must have at least one truth jet and $p_{\mathrm{T}}^{\mathrm{avg,reco}}/p_{\mathrm{T}_{\mathrm{jet}}}^{1,\mathrm{truth}} \leq 1.4$. At the truth level, only the timing and kinematic selection are required.

4.3 Unfolding

4.3.1 Motivation

In high-energy physics, experimental distributions are modified by detector effects such as limited acceptance, finite resolution, and other systematic factors that entail a transfer of events between different regions of the spectra. For example, the non-zero JER (see Section 3.1.7) causes some jets to be reconstructed with a p_{T} value that is either higher or lower than their true value, resulting in their migration to adjacent analysis bins. When these effects are well understood and controlled, they can be included in MC simulations of the detector response, which can then be used to correct the data accordingly. This correction procedure is known as unfolding.

In most applications, unfolding is used to estimate the truth-level distribution of an observable, such as the transverse momentum of a jet. Due to detector inefficiencies and resolution effects, the truth-level and reconstructed-level distributions may indeed differ. Unfolding methods aim to provide an estimator that corrects for these effects statistically, providing a more accurate representation of the underlying physical process with minimal bias and minimal model dependence. Unfolding data enables a direct comparison of measurements from different experiments or with theory, as well as data preservation and re-interpretation. It also allows for simplifying phenomenological studies.

4.3.2 General unfolding strategy

The following notations are used: d represents the data distribution, while r and t correspond to the reconstructed and truth-level distributions in the MC simulation. The unfolded data distribution is denoted by u.

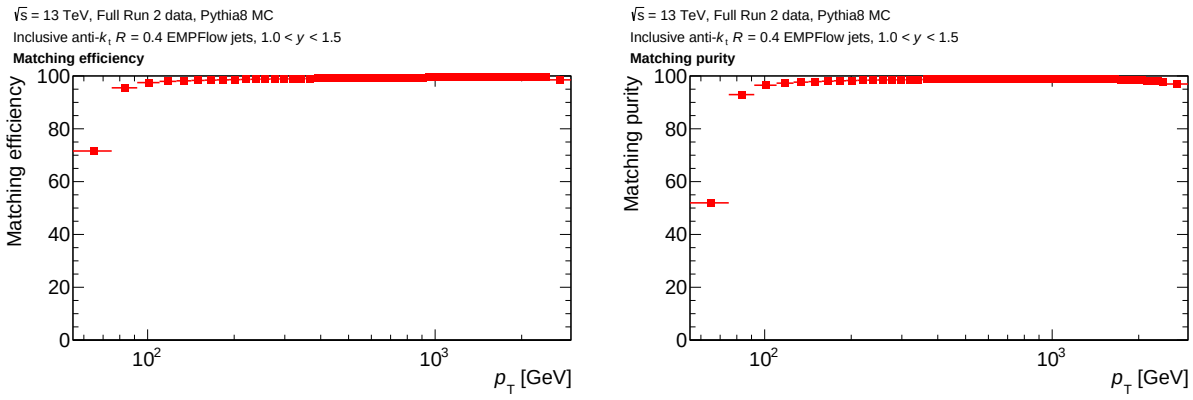
For binned observables, the unfolding method relies on a transfer matrix $T_{i,j}$ which represents the number of events originating from the truth bin j that is reconstructed in the reco bin i. The folding matrix describes detector effects and can only be determined using simulations. An unfolding matrix $U_{i,j}$ is constructed from the transfer matrix such that each element of the unfolding matrix represents the probability that an event in a particular reconstructed-level bin originates from a specific particle-level bin, given the

assumed input particle-level spectrum:

$$U_{ij} = p\left(r \in (\Delta r)_i,\, t \in (\Delta t)_j\right)$$

In practice, the unfolding matrix can be refined using various methods. In the case of the inclusive jet cross-section measurement, an iterative, dynamically stabilised method (IDS) [115, 116], described in Section 4.3.3, is used.

Selected events, chosen to minimise potential biases, are used to build the histograms of the cross-section spectra and the transfer matrix. However, some events may not contribute to the transfer matrix due to detector inefficiencies, such as limited acceptance or finite resolution. To address this, the matching purity and matching efficiency are defined as: $\epsilon_i^{\text{purity}} = \frac{\sum_j T_{i,j}}{r_i}$ and $\epsilon_j^{\text{efficiency}} = \frac{\sum_i T_{i,j}}{t_j}$, respectively. The first step of the unfolding procedure is applying the impurity correction to the data spectrum, ensuring consistency between the number of jets used to build the transfer matrix and the number of jets in the corrected data spectrum. Specifically, events in the reconstructed bin i are weighted by the matching purity at the reconstruction level. In the second step, the unfolded data spectrum is obtained using the unfolding matrix. The third step is the application of a correction to account for matching efficiency losses: events in bin j are weighted by the inverse of the matching efficiency at the particle level. An example of efficiency and purity for the inclusive jet cross-section measurement is shown in Figure 4.2.



(a) (b)

Figure 4.2: Matching efficiency (a) and purity (b) for the inclusive jet cross-section measurement, for $1.0 < |y| < 1.5$.

4.3.3 The IDS method

In practice, several methods are used to evaluate the unfolding matrix, one of which is the Iterative Dynamically Stabilised (IDS) [115, 116] method. The method relies on the transfer matrix giving an accurate description of the detector effects and on the quality of the model used for the truth-level MC spectrum. The latter may not be perfect, as the truth-level spectrum (and thus the transfer matrix) is iteratively improved throughout the unfolding process, by comparing the truth-level MC and unfolded distributions.

The transfer matrix may also be imperfect, and this limitation is accounted for through the JER systematics. The number of iterations and the binning serve as regularisation parameters.

An initial folding (resp. unfolding) matrix is constructed from the transfer matrix by normalising row-by-row (resp. column-by-column) so that the sum of the elements corresponding to a given truth-level (resp. reconstructed-level) bin is unity.

$$F_{ij} = \frac{T_{ij}}{\sum_k T_{kj}}, \qquad U_{ij} = \frac{T_{ij}}{\sum_k T_{ik}}$$

F_{ij} gives the probability for an event in the truth bin j to arrive in the reco bin i, and U_{ij} provides the probability for an event in the reconstructed bin i to come from the truth bin j. An example of a folding matrix for the inclusive jet cross-section measurement is shown in Figure 4.3. At each iteration, the unfolded spectrum is given by:

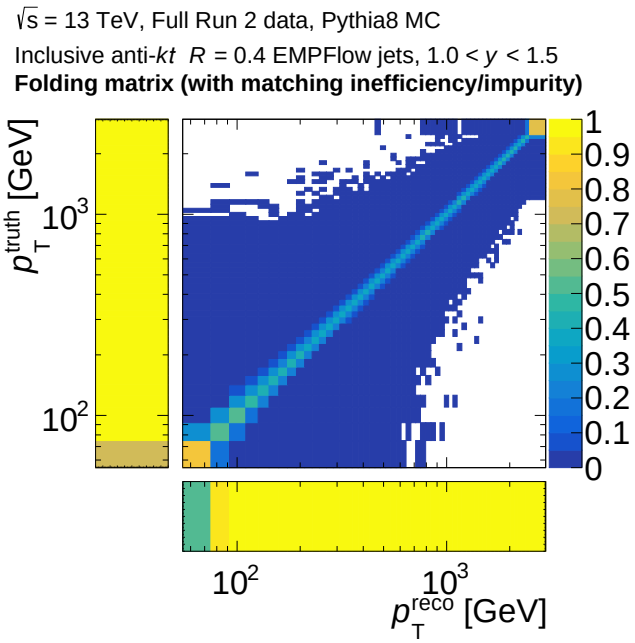


Figure 4.3: Example of a folding matrix for the inclusive jet cross-section measurement, for $1.0 < |y| < 1.5$. The side bands are the projection along the $p_{\mathrm{T}}^{\mathrm{reco}}$ and $p_{\mathrm{T}}^{\mathrm{truth}}$ axis. They represent the matching purity and efficiency, respectively.

$$u_j = t_j \frac{N_{\mathrm{data}}}{N_{\mathrm{MC}}} + \sum_k f_k \Delta d_k U_{kj} + (1 - f_j)\Delta d_j,$$

where $\frac{N_{\mathrm{data}}}{N_{\mathrm{MC}}}$ is a normalisation factor, and $\Delta d_k = d_k - \frac{N_{\mathrm{data}}}{N_{\mathrm{MC}}} r_k$ quantifies the difference between the data and scaled MC reconstructed distribution. The term f_k acts as a regularisation term that dictates the fraction of events that migrate from the bins k, based on the associated uncertainty, with σ_k the statistical uncertainty in d_k:

$$f_k = 1 - e^{-(\frac{\Delta d_k}{\lambda \sigma_k})^2}.$$

In this equation, λ is a positive regularisation parameter that can, in principle, be different at each step of the procedure. The setup used here employs, however, very little regularisation, other than the one induced by the choice of the number of iterations.

Between iterations, the truth-level spectrum is refined through a reweighting procedure that accounts for the discrepancy between the truth-level MC and the unfolded data at the current step. This process leads to an update of the transfer matrix as follows:

$$T'_{ij} = T_{ij} + \Delta u_j F_{ij} \frac{N_{\mathrm{MC}}}{N_{\mathrm{data}}}$$

with $\Delta u_j = u_j - \frac{N_{\mathrm{MC}}}{N_{\mathrm{data}}} t_j$. Each iteration reduces the bias caused by an imperfect MC model and shape differences between data and MC, but leads to increased statistical uncertainties. The number of unfolding iterations must be chosen carefully to balance bias reduction and the control of statistical uncertainties.

For the inclusive jet cross-section measurement, the number of unfolding iterations is taken as $N = 1$ (optimised in Reference [98]), meaning that the unfolding is performed twice, once with the original transfer matrix and with the reweighted one. The comparison of data before and after unfolding is provided in Figure 4.4. The differences are more pronounced at low p_{T}, where the p_{T} resolution is larger, and at high p_{T} (above 1 TeV), where the jet production cross-section starts being more affected by statistical fluctuations.

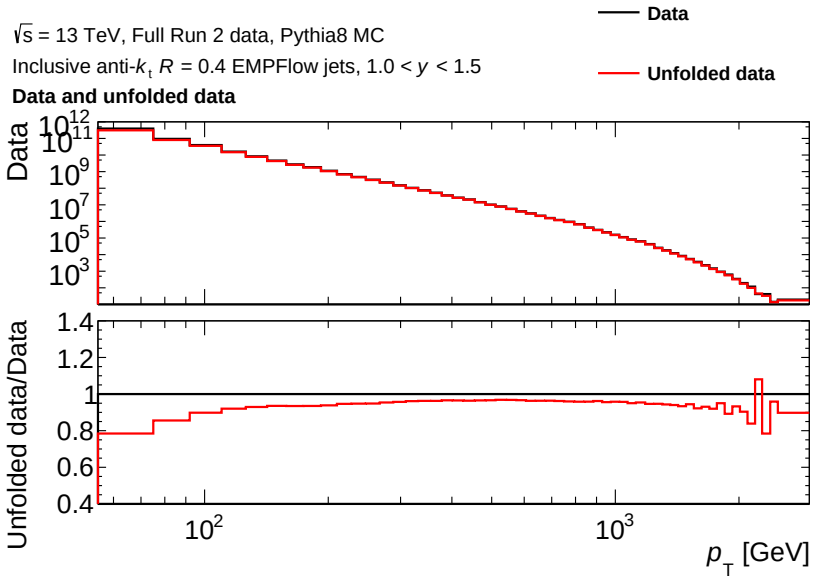


Figure 4.4: Inclusive jet cross-section spectrum for $1.0 < |y| < 1.5$, before (in black) and after (in red) unfolding with 1 iteration. The bottom panel shows the ratio of the unfolded spectra to the non-unfolded spectra.

4.3.4 Data-driven closure test

The transfer matrix $T_{i,j}$ is built from simulated data, making it dependent on the choice of the MC model. However, this model dependence decreases as the number of unfolding

iterations increases. Any residual bias can be estimated through a data-driven closure test, presented in Reference [115]. The test consists of applying the unfolding procedure to a pseudo-data distribution r' that shares the same overall characteristics as the data distribution and has a known corresponding truth distribution t'. The associated unfolded distribution u' is then compared to t': a non-closure bias is evaluated as $b = (u' - t')/t'$.

To obtain the pseudo-data distribution r', the MC events are reweighted as a function of the truth-level variable to better match the data at the reconstructed level. These weights are chosen to improve the agreement between the data (after applying the matching efficiency at the reconstructed level) and the reconstructed MC (obtained by projecting the modified matrix onto the X axis, corresponding to the p_{T} of the reconstructed jet). Then, r' and t' are normalised to have the same number of events as d, and the statistical uncertainties of r' are set to be those of d. The reweighting can be done in several ways. In the inclusive jet analysis, a smooth function of p_{T} is applied, whereas for the Lund jet plane analysis, a simple binned ratio is used.

4.3.5 Application to the inclusive jet cross-section analysis

4.3.5.1 Object definition

For the inclusive jet cross-section analysis, the detector effects are mostly bin migrations due to the large p_{T} resolution of jets. The transfer matrix is populated on a jet-by-jet basis by matching a truth-level jet with a reconstructed-level jet when the two satisfy specific criteria. Namely, they must:

- Be closer to each other than to any other jet,
- Lie within a radius of $R = 0.3$,
- Have a transverse momentum of at least $p_{\mathrm{T}} > 75$ GeV,
- Belong to the same rapidity bin.

A truth jet with an assigned reconstructed jet is considered matched, and similarly, a reconstructed jet with an assigned truth jet is also considered matched.

4.3.5.2 Data-driven closure test

For the data-driven closure test, the reweighting is performed using a smooth function $f(p_{\mathrm{T}})$. It is modelled as a polynomial with the degree N depending on the number of bins in the distribution:

$$f(p_T) = \sum_{i=0}^{N} c_i \left(\log p_{\mathrm{T}} - \log p_{\mathrm{T}}^0\right)^i$$

The f function is optimised such that r' is as close to d as possible, by minimising the following function:

$$\tilde{\chi}^2 = \sum_i \frac{(d_i - r'_i)^2}{\Delta d_i^2 + \Delta r_i'^2}$$

The shape differences between data and MC before and after reweighting of the truth distribution by f, along with the local $\tilde{\chi}^2$ before and after minimisation, are displayed in Figure 4.5.

The number of unfolding iterations has been optimised in Reference [98] by minimising the quadratic sum of the unfolding bias and statistical uncertainties. The optimal number of iterations is found to be one.

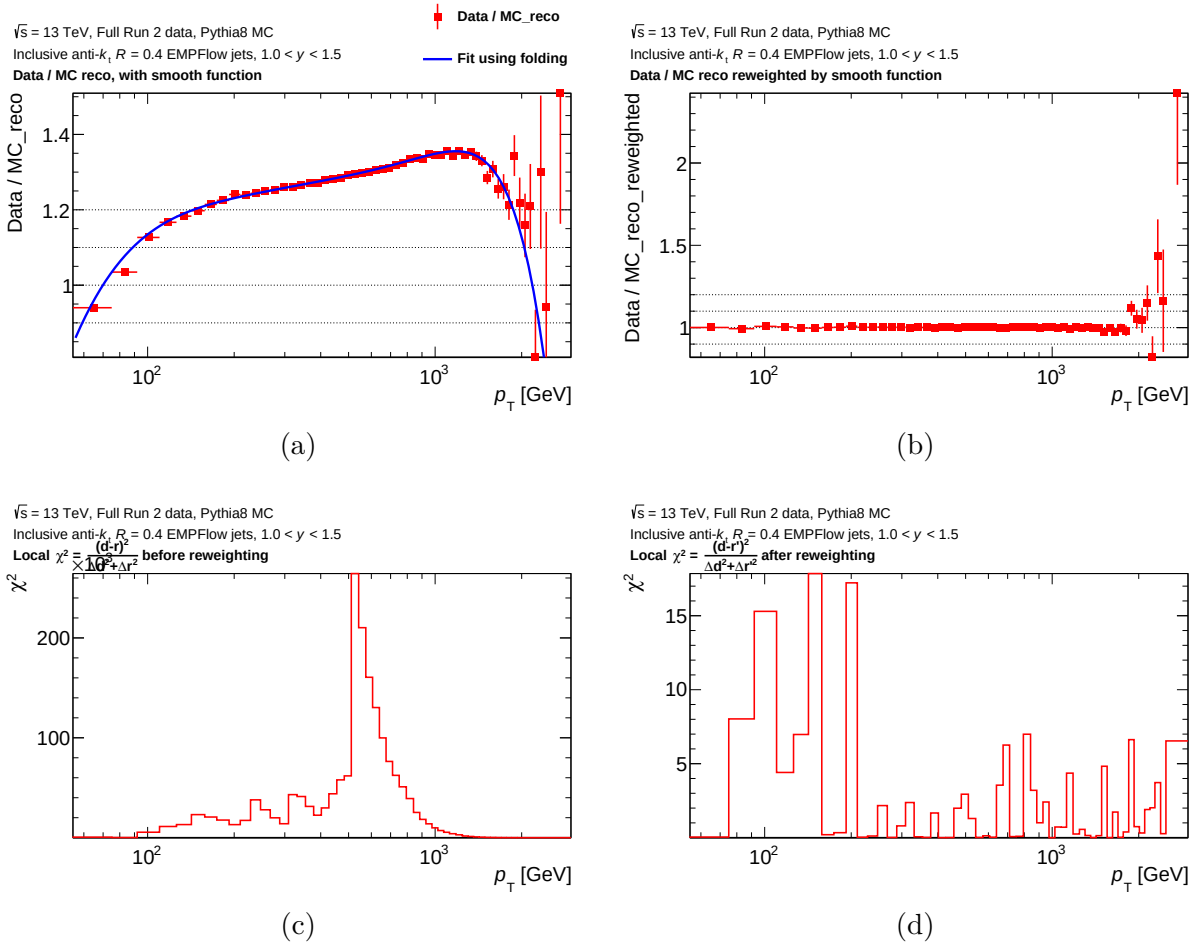


Figure 4.5: Illustration of the reweighting process: the d/r distribution is shown before (a) and after (b) reweighting. The smooth function f used for the reweighting process is shown alongside the initial d/r distribution. The local $\tilde{\chi}^2$ before (c) and after (d) reweighting is shown as well.

4.4 Experimental uncertainties

Multiple sources of uncertainty affect the cross-section measurements:

- Statistical uncertainties linked to the finite amount of data,
- Systematic uncertainties related to the missing tile modules,
- Systematic uncertainties related to the application of the JES and the JER,
- Uncertainty linked to the non-gaussianity of the JER,
- Bias related to the sensitivity of the IDS unfolding to the data/MC shape difference,
- uncertainty in the luminosity measurement,

For the quantitative comparison of the data and the theory prediction, individual system-

atic uncertainty sources are treated as uncorrelated with each other, as they correspond to independent sources.

4.4.1 Statistical uncertainties

4.4.1.1 Evaluation through the bootstrap method

The statistical uncertainties of the inclusive jet cross-section measurement are evaluated using the bootstrap method, thoroughly explained in Section 3.3.3.1. One thousand bootstrap replicas are generated for every distribution (data, MC transfer matrix) by assigning a deterministically generated weight drawn from a Poisson distribution with a mean of unity. As a convention, the random number used to select the Poisson weight is seeded by the event number, allowing for coherent statistical fluctuations across different measurements that use the bootstrap method. Each replica dataset is treated in the same way as the nominal dataset.

The total statistical uncertainty and the bin-to-bin correlations are obtained from the covariance matrix, which encodes on the diagonal the statistical uncertainty as the standard deviation of the bootstrap, and the bin-to-bin correlations in the off-diagonal elements. Such a correlation matrix is displayed in Figure 4.6. Correlations and anti-correlations

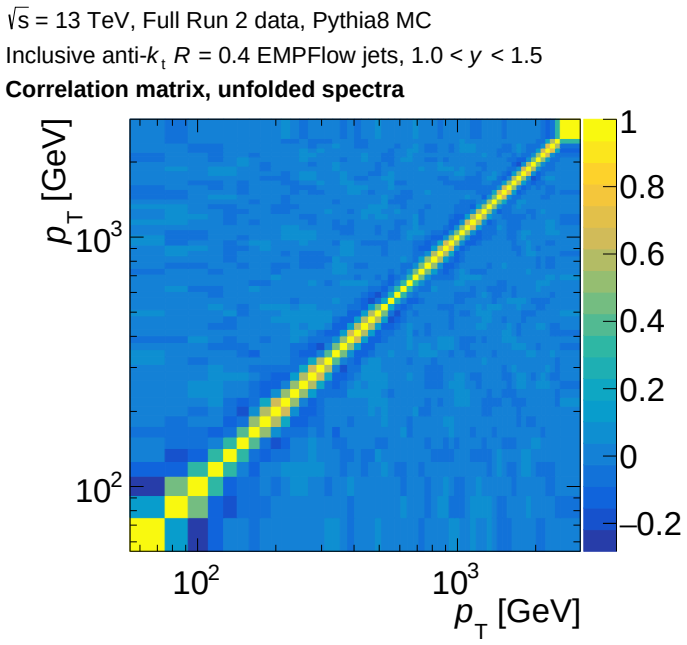


Figure 4.6: Statistical correlations obtained with the bootstrap method for the inclusive jet cross-section analysis in the $1.0 < |y| < 1.5$ bin, after unfolding.

are observed near the diagonal, as expected, since the unfolding causes jets to move to the adjacent bins. Outside the diagonal, a rather uniform background of ± 20% reflects the statistical noise in the determination of the correlations.

4.4.1.2 Correlation with the η-intercalibration factors

In addition to the evaluation of the statistical uncertainties, the bootstrap method also enables the assessment of statistical correlations between measurements that share overlapping input data, thanks to the use of a deterministic seed. The inclusive jet cross-section measurement uses an input sample that largely overlaps with that used by the jet η-intercalibration. In particular, the η-intercalibration relies on dijet events selected using the same single-jet triggers as the inclusive jet cross-section measurement. Studying these correlations has been made possible with the recent development in the η-intercalibration. It is essential for avoiding the double-counting of information and accurately estimating the total uncertainty.

The correlations between those two measurements are evaluated here for the first time. Correlation matrices are computed, where the lower-left block represents the inclusive jet cross-section correlation matrix (before unfolding), and the upper-right plot represents the correlation matrix of the η-intercalibration factors. The potential correlations between the two measurements would be displayed in the off-diagonal blocks. This is done for all combinations of the $|y|$ bins of the inclusive jet cross-section analysis and the $p_{\mathrm{T}}^{\mathrm{avg}}$ bins of the η-intercalibration, for every year of LHC data-taking.

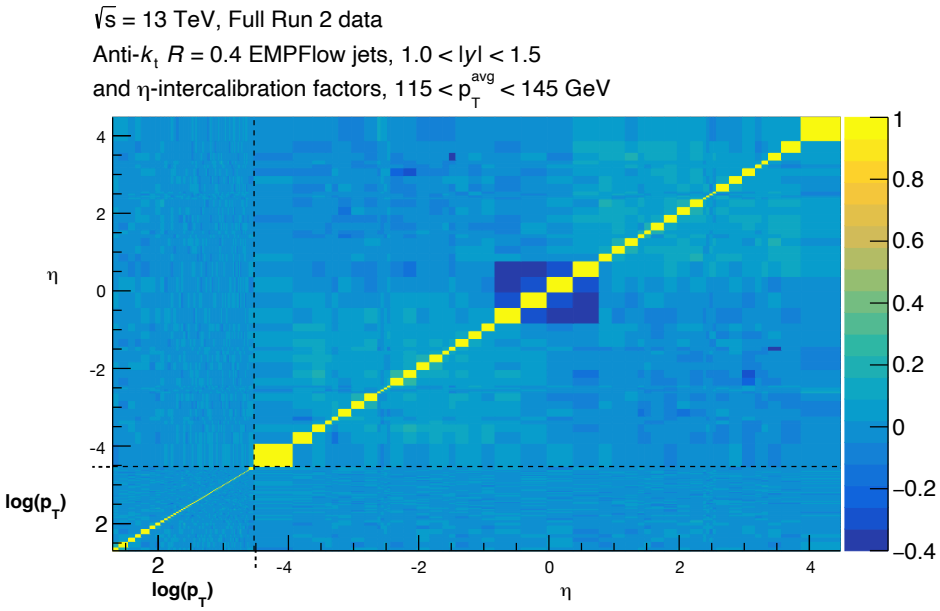


Figure 4.7: Correlation matrix for the inclusive jet cross-section measurement and the η-intercalibration, for the $1.0 < |y| < 1.5$, $60 < p_{\mathrm{T}}^{\mathrm{avg}} < 85$ GeV bin combination, for 2017 data. The lower-left block is the inclusive jet cross-section correlation matrix before unfolding, in $\log(p_{\mathrm{T}}[\mathrm{GeV}])$. The upper-right plot is the correlation matrix of the η-intercalibration factors, in η_{det}.

As shown in Figure 4.7, no significant correlations are observed on the off-diagonal, indicating no correlation between the two analyses. This is the case for all the others $|y|$,

$p_{\mathrm{T}}^{\mathrm{avg}}$ combinations, and all datasets tested. The dilution of the expected correlations can be attributed to several factors:

- The narrower p_{T} binning of the inclusive jet cross-section measurement.
- The different use of the events: the inclusive cross-section counts individual jets, whereas the η-intercalibration relies on the asymmetry of dijet systems to constrain calibration factors.
- The tighter selection applied for the calibration.

4.4.2 Systematic uncertainties

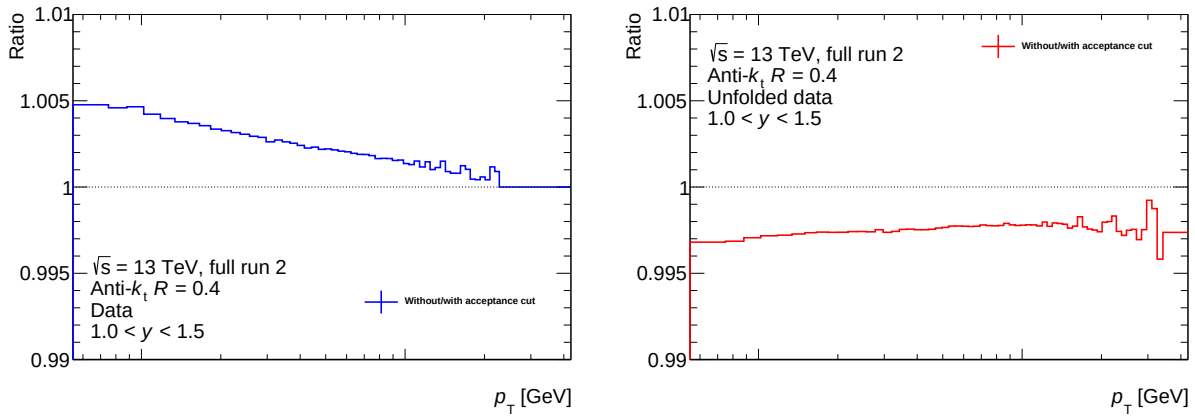
4.4.2.1 Missing Tile Modules

It occasionally occurs that some modules temporarily fail during data-taking. The modules and period affected are recorded. The missing Tile calorimeter modules induce a systematic bias in the p_{T} to be reconstructed at lower values. This bias becomes more pronounced at higher p_{T}, since the Tile calorimeter has a higher energy fraction in this regime. The previous recommended procedure was to derive an uncertainty based on the ratio of the jet cross-section with and without the missing Tile calorimeter modules, but the uncertainty was largely overestimated as jets in the regions with missing Tile calorimeter modules are either mis-measured or absent.

A more effective approach that had never been implemented before is to exclude the (η, ϕ) regions corresponding to missing Tile calorimeter modules from the reconstructed jet cross-section for each data period, in data and MC. Since these regions are removed from the reconstructed spectra but not from the truth spectra, any jet falling within these areas would remain unmatched and thus be accounted for in the matching efficiency correction of the unfolding process. An uncertainty can then be assessed to cover for the inefficiency observed. The impact of this treatment is illustrated in Figure 4.8, which shows the ratio of the inclusive jet cross-section spectra with and without the acceptance cut, before and after unfolding. Figure 4.8a corresponds to the previous treatment of dead Tile modules, where the largest effect reaches about 0.5%. In contrast, Figure 4.8a includes the acceptance correction, reflecting the new approach, and shows that the difference between spectra with and without the acceptance cut is about $0.3 - 0.3\%$. Since it is really small, no systematic uncertainty is assigned for this effect. This study also shows that in the old approach, the main effect is due to the acceptance loss and not to the bad jet measurement.

4.4.2.2 JES uncertainties

To evaluate the JES uncertainties, the jet p_{T} is scaled up and down by 1σ for each uncertainty component, the p_{T} spectra are refilled and unfolded using the nominal unfolding matrix, and the systematic uncertainty is determined as $\Delta u = u_{\mathrm{varied}}/u_{\mathrm{nominal}} - 1$. Since these uncertainties evaluate a relative effect of a change in the calibration, they can be evaluated on an MC sample acting as pseudo-data, allowing for reduced statistical fluctuations in the propagation of the systematics, especially in the high-p_{T} region. In this case, u_{nominal} corresponds to MC reco unfolded as if it were data. The observed changes in the uncertainties due to the unfolding are at the level of a few percent of the uncertainties



(a) (b)

Figure 4.8: Ratio of the inclusive jet cross-section spectra with and without the acceptance cut, before (a) and after (b) unfolding with the modified MC samples. The inefficiencies due to the missing tile modules are taken into account during the unfolding procedure, as described in Section 4.3.2.

themselves.

While the JES uncertainties are initially symmetric, the propagation of the uncertainties from the jet level to the level of the reconstructed cross-section spectra introduces asymmetries, mostly due to the non-linearly falling shape of the cross-section. Each variation is applied consistently across the entire $(p_T, |y|)$ range and is treated as fully correlated within the phase space, as it arises from a single source (whether it be of physical or statistical origin). However, it remains independent of other systematic uncertainties. The origins of the JES systematic uncertainties are shown in Table 4.2. A statistical uncertainty in the systematic uncertainties estimates is derived using the bootstrap method with 100 replicas. To prevent local statistical fluctuations from impacting the propagation of the systematic uncertainties and hence the global χ^2/NDF, the systematic uncertainties are rebinned until they become statistically significant and smoothed using a Gaussian

Category	**Statistical uncertainty**	**Systematic uncertainty**
η-intercalibration	1000	4×2 $+ 1 \times 2$ (2018 only)
Single particle at high-p_{T}	2×2	27×2
Flavour		6×2
γ+jet	16×2	8×2
Z+jet	28×2	12×2
Pile-up		4×2
Punch through		1×2
Total: 1155	1092	63

Table 4.2: Origin of the JES systematic uncertainties. The η-intercalibration statistical uncertainties are obtained from the new approach using the bootstrap method. A non-closure systematic uncertainty accounts for the application of the 2017 calibration on 2018 data. The $\times 2$ accounts for the up and down variations.

kernel.

The standard treatment of the statistical uncertainties of the η-intercalibration is the use of one nuisance parameter, fully correlated between the jet η and p_{T} bins for which the η-intercalibration was performed. Previously in the analysis, the η-intercalibration statistical uncertainties were treated as 532×2 nuisance parameters, added in quadrature, for each data period considered (2015–16 and 2017–18), corresponding to individual η-intercalibration statistical uncertainties local in $(p_{\mathrm{T}}, |y|)$. Using this scheme instead of one fully correlated uncertainty helped to reduce the χ^2 by 200 units in Reference [110]. In this scheme, the uncertainties are not smoothed, allowing for a further reduction of the χ^2. The smoothing procedure causes local uncertainties to "leak" into adjacent bins, artificially inflating correlations between them.

This motivates another approach, taking advantage of the recent improvement in the determination of the η-intercalibration statistical uncertainties using the bootstrap method (Section 3.3.3). The 1000 bootstrap replicas of the calibration can be used to provide as many alternative datasets, each calibrated using the factors derived from the statistically fluctuated samples, following the procedure described in Section 3.3.3.4. Each of the 1000 alternative datasets is unfolded, and the resulting relative uncertainty Δu is scaled by a factor of $\sqrt{1/(N_{\mathrm{replica}} - 1)}$ and handled in the same way as the other systematic uncertainties.

Using the statistical uncertainties from the bootstrap method enables the use of the accurate, enhanced statistical uncertainties (the evaluation of which is described in Section 3.3.3.3), as well as evaluating and propagating the correlations due to the large overlap between the samples if the η-intercalibration and the inclusive jet cross-section analysis. A complete study of the correlations between the two datasets is detailed in Section 4.4.1.2. The η-intercalibration statistical uncertainties derived with the older method (using 532×2 nuisance parameters) and the new method (using the bootstrap replicas) are shown in Figure 4.9 for the full Run 2. It can be observed that the amplitude of

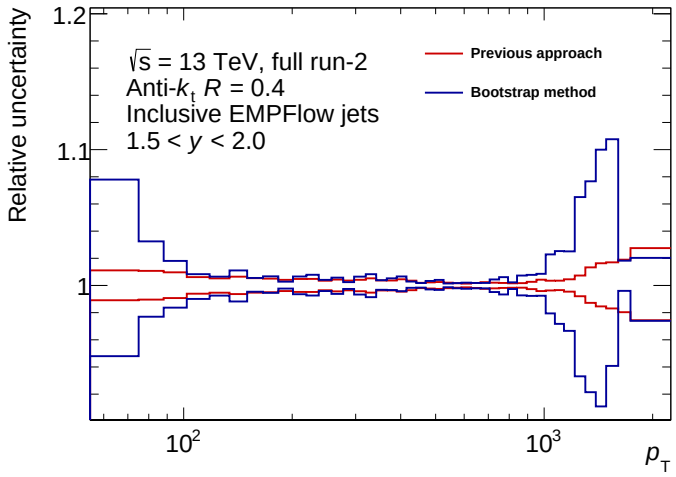
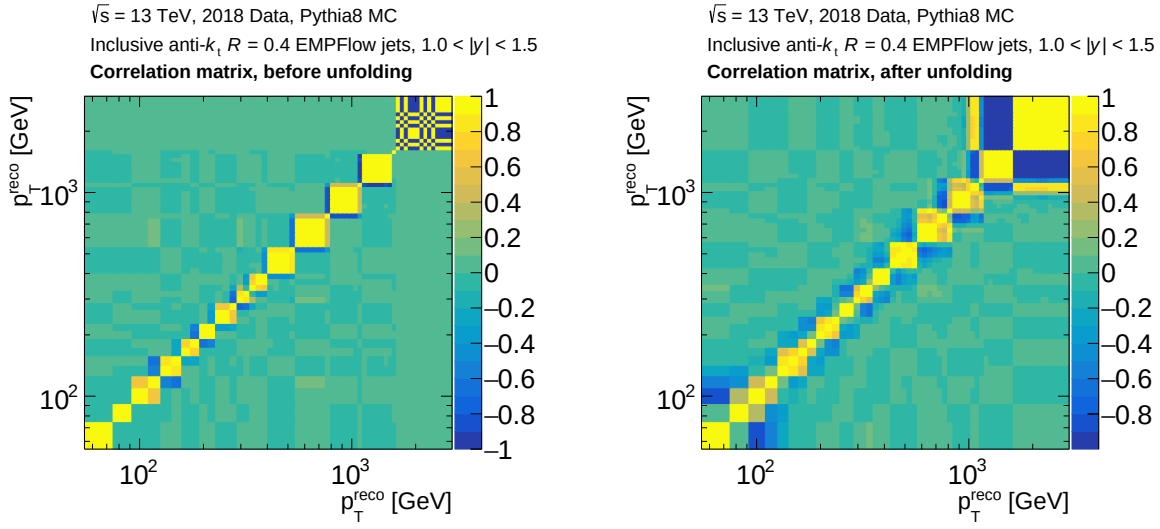


Figure 4.9: Scaled uncertainties are added in quadrature, equivalent to the RMS of the Bootstrap distribution.

the uncertainty enhancement differs between the spectra and the calibration factor levels, with the expected factor of 2 not being recovered. A further study reveals that this discrepancy arises from dilutions caused by the different p_{T} and rapidity binning used in the determination of calibration factors and in the analysis, as well as the modified correlations in rapidity.



(a) (b)

Figure 4.10: Correlation matrix for the 1000 η-intercalibration statistical uncertainty components, before (a) and after (b) unfolding. The checkboard pattern before the unfolding is caused by numerical instability in this low-statistic region.

In practice, the components are only added in quadrature for visualisation purposes, because the quadratic sum doesn't account for the correlations between them. These can be displayed in a correlation matrix, as represented in Figure 4.10. Fully correlated blocks can be seen along the diagonal. This effect arises from the η-intercalibration bins being larger than those used in the inclusive jet cross-section analysis. Each block corresponds to a $p_{\mathrm{T}}^{\mathrm{avg}}$ bin of the η-intercalibration: in one block, all jets are calibrated with the same factors, hence the 100% correlation among the bins of each block. Anti-correlations are visible at the higher-p_{T} border of the block. This effect arises because the last bin of a block straddles two η-intercalibration regions. Due to the steeply falling cross-section spectrum, most jets in a straddling bin originate from the lower-p_{T} side and are therefore calibrated with the corresponding lower-p_{T} η-calibration factors. When these factors fluctuate downward, jets in the affected bin are over-corrected and tend to migrate toward higher-p_{T} neighbour bins, which are calibrated with a different set of η-intercalibration factors. This induces anti-correlations between the straddling bin and the bins above it. When the lower-p_{T} calibration factors fluctuate upward, jets migrate toward lower-p_{T} neighbouring bins, which are calibrated with the same set of factors. This yields positive correlations between the straddling bin and the bins below it.

The checkboard pattern in the last bin before the unfolding is caused by low-statistics migrations induced by variations of the systematic uncertainties in this low-statistic re-

gion, and disappears when rebinned. During the unfolding procedure, some jets migrate into other bins, causing correlations and anti-correlations around the diagonal. This phenomenon is seen in the correlation matrices of the nominal dataset as well, as seen in Section 4.4.1. In the case of the η-intercalibration statistical uncertainties, the unfolding smears the correlation scheme around the diagonal.

4.4.2.3 JER uncertainties

The uncertainties in the jet energy resolution arise from three sources. The first source consists of systematic uncertainties in the measurement of the JER (detailed in Section 3.1.7). They are evaluated by smearing the p_{T} of the reconstructed jets in the MC simulation such that the resolution is degraded by the size of each uncertainty component. Namely, the jet p_{T} is smeared by a Gaussian function of width $\sigma^2_{\mathrm{smear}} = (\sigma_{\mathrm{nominal}} - \sigma_{\mathrm{varied}})^2 - \sigma^2_{\mathrm{nominal}}$, with $\sigma_{\mathrm{nominal}}$ the nominal JER and σ_{varied} the one-standard-deviation variation in the uncertainty component evaluated. When $\sigma_{\mathrm{varied}} > 0$, the resolution is degraded and the p_{T} of the jets entering the transfer matrix is smeared. When $\sigma_{\mathrm{varied}} < 0$, the resolution is improved. Since the transfer matrix cannot be un-smeared, the p_{T} of the jets entering the pseudo-data spectra is smeared, and the transfer matrix is left unchanged. Then, the transfer matrix is used to unfold the pseudo-data spectra. As for the JES, the uncertainty is derived as $\Delta u = u_{\mathrm{varied}}/u_{\mathrm{nominal}} - 1$.

The second source is the difference between the JER measured in data and MC simulations. Indeed, during the calibration procedure, when the resolution is better in MC than in data, the resolution of reconstructed jets in MC is smeared to match the resolution in data. When the resolution is better in the data than in MC, the difference between the two is taken as an additional JER systematic uncertainty and propagated similarly to the others.

The third source is the non-gaussianity of the JER. As described in 3.1.7, the JER is evaluated in situ as the width of the Gaussian fit to asymmetry distributions, with the assumption that those distributions are Gaussian. However, those distributions have non-Gaussian tails. An uncertainty to cover for this effect has been derived for the first time in Reference [98]. The effect of the non-gaussianity of the distributions manifests in the off-diagonal terms of the transfer matrix. Thus, the approach is to compare the unfolding result using a "Gaussian transfer matrix" with the nominal result, which includes the non-Gaussian effects, and treat the difference as an uncertainty. The Gaussian transfer matrix is constructed by replacing the reconstructed jet p_{T} in the transfer matrix with a convolution of the truth distribution and a Gaussian of width equal to the JER value in the given bin, using a very fine binning. The associated systematic uncertainty is found to be non-dominant but non-negligible.

4.4.2.4 Other uncertainties

An uncertainty is derived to account for the bias related to the sensitivity of the IDS unfolding to the data/MC shape difference. It is evaluated through the data-driven closure test, as discussed in Section 4.3.4.

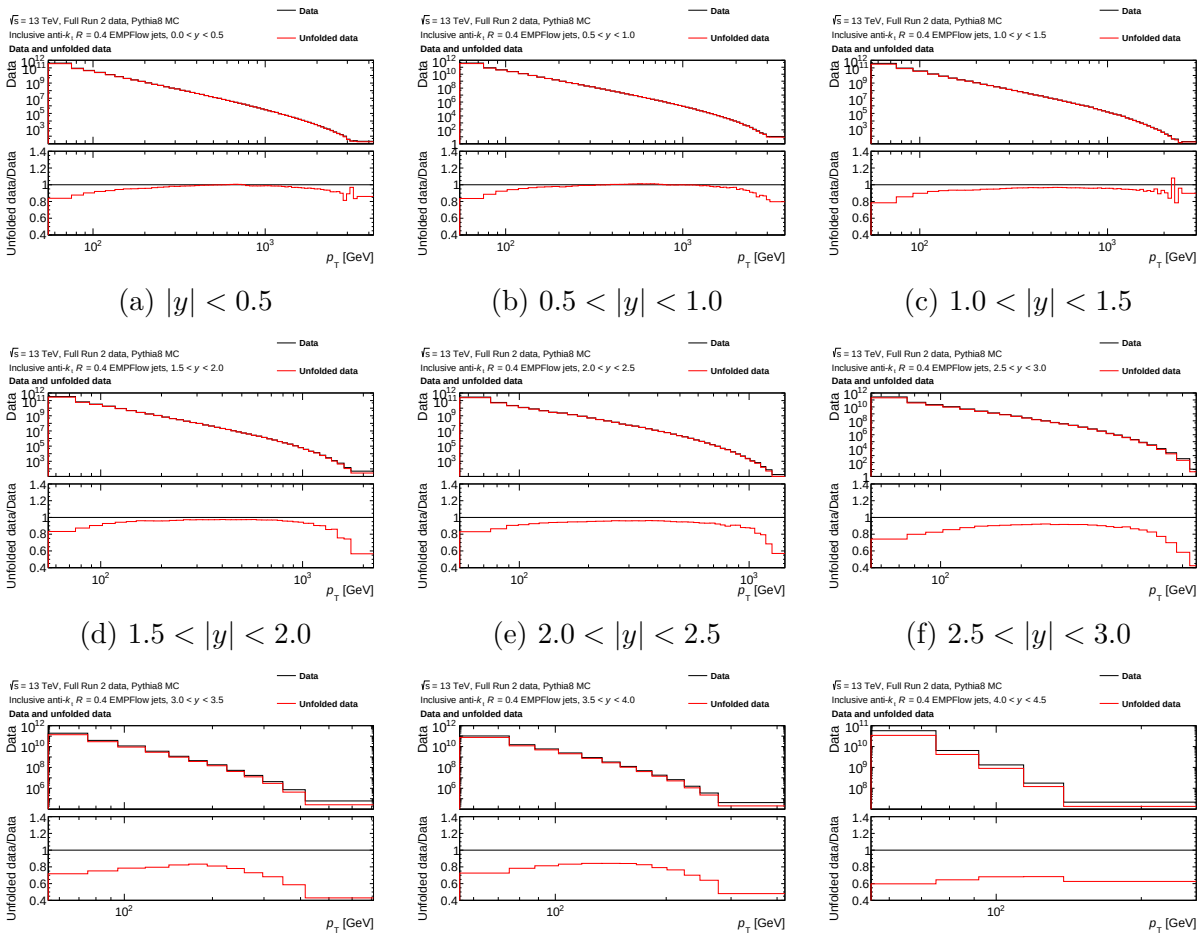
The relative uncertainty in the total luminosity recorded by ATLAS during Run 2 [58]

translates into a fully correlated uncertainty of the same magnitude across the entire phase-space in the inclusive jet cross section analysis.

4.5 Results

4.5.1 Unfolded data and uncertainties

The inclusive jet cross-section spectra are obtained in the nine $|y|$ bins. Figure 4.11 shows the unfolded distribution alongside the non-unfolded one.



(a) $|y| < 0.5$ (b) $0.5 < |y| < 1.0$ (c) $1.0 < |y| < 1.5$

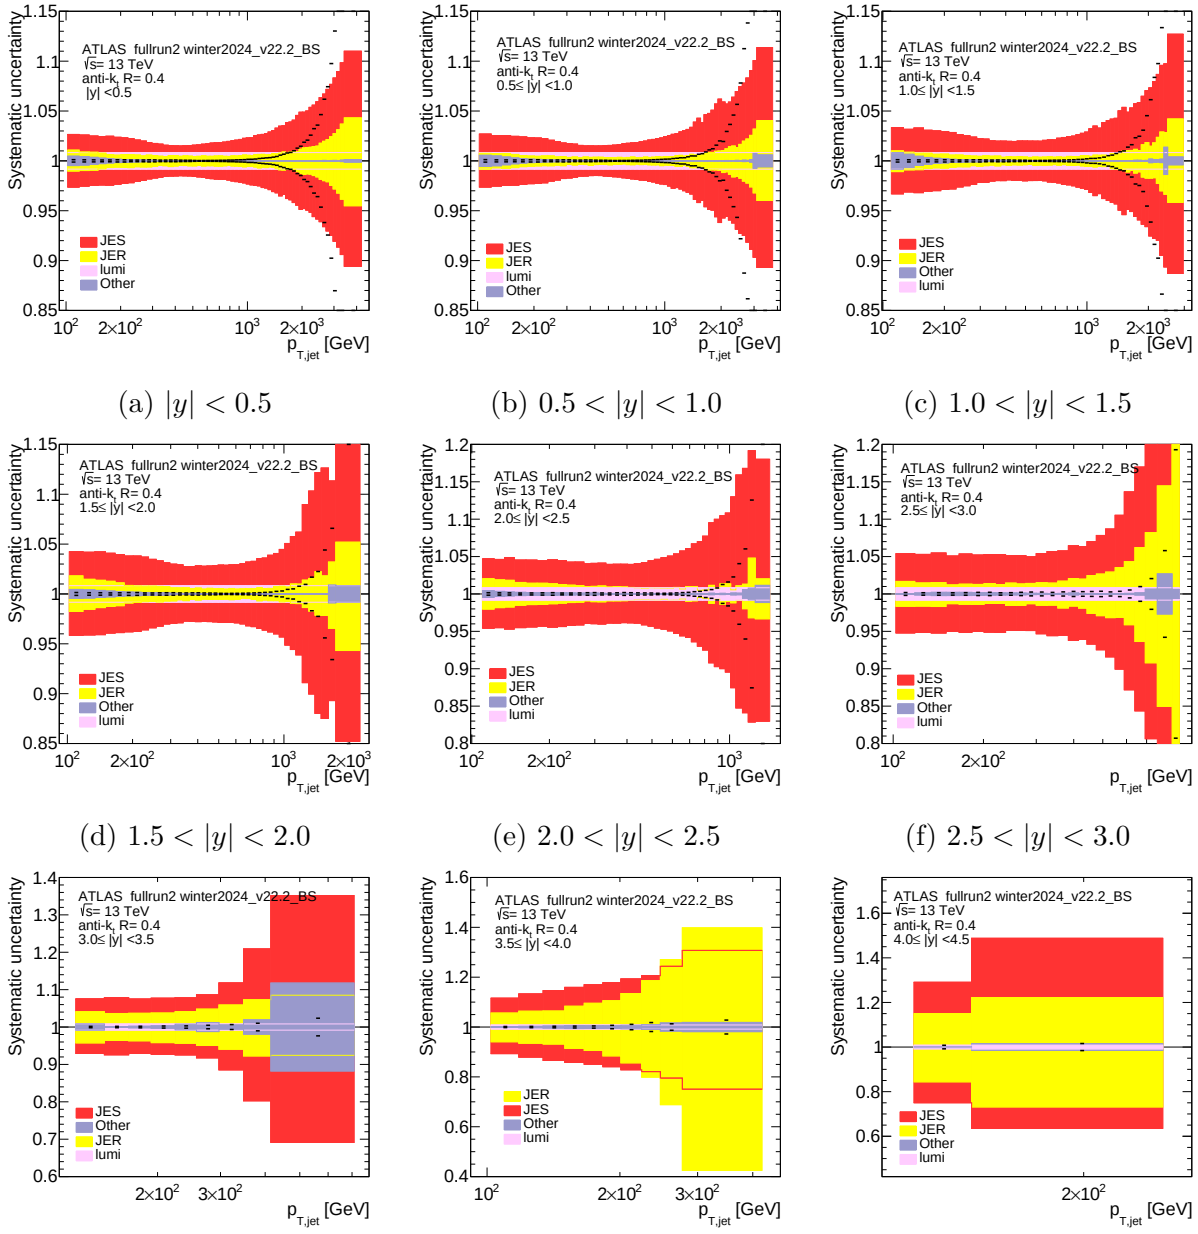
(d) $1.5 < |y| < 2.0$ (e) $2.0 < |y| < 2.5$ (f) $2.5 < |y| < 3.0$

(g) $3.0 < |y| < 3.5$ (h) $3.5 < |y| < 4.0$ (i) $4.0 < |y| < 4.5$

Figure 4.11: Inclusive jet cross-section spectra as a function of p_{T}, in rapidity bins. The unfolded data spectrum is shown alongside the non-unfolded data spectrum.

The systematic uncertainties are shown in Figure 4.12. In this figure, the individual sources of uncertainty are added in quadrature in three subcategories: JES, JER, and other (luminosity and unfolding uncertainty). The JES uncertainty is the largest among all sources, amounting to approximately 2–3% in the central region and 5–10% in the forward regions. It is smallest around $p_{\mathrm{T}} = 300$ GeV and generally increases toward higher p_{T}. In the highest-p_{T} region, it reaches about 10% in the central region and can

rise up to 30% in the most forward region. The unfolding uncertainty and the luminosity are relatively small.



(a) $|y| < 0.5$ (b) $0.5 < |y| < 1.0$ (c) $1.0 < |y| < 1.5$

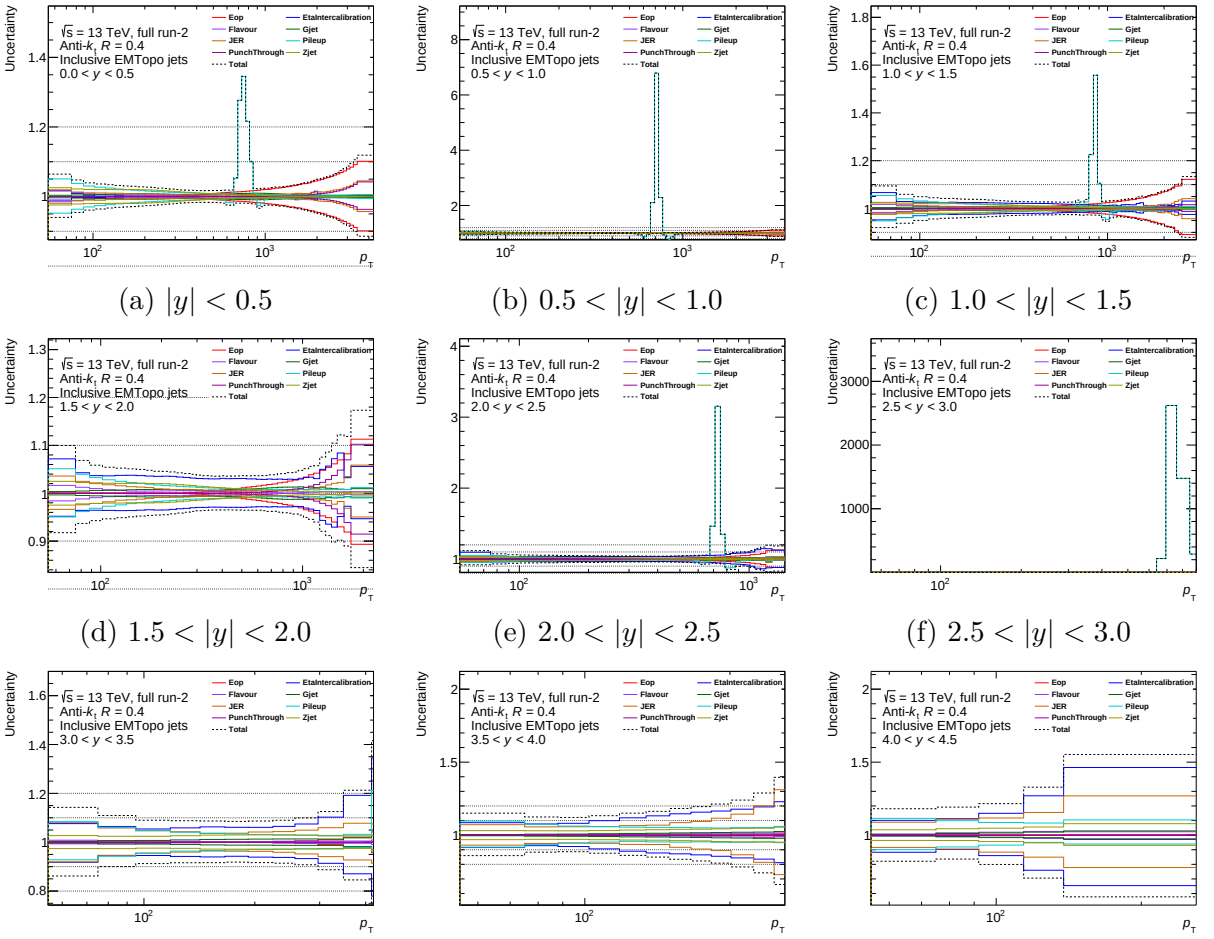
(d) $1.5 < |y| < 2.0$ (e) $2.0 < |y| < 2.5$ (f) $2.5 < |y| < 3.0$

(g) $3.0 < |y| < 3.5$ (h) $3.5 < |y| < 4.0$ (i) $4.0 < |y| < 4.5$

Figure 4.12: Systematic uncertainties components on the unfolded inclusive jet cross-section as a function of p_{T} in rapidity bins. The statistical uncertainties are shown as small black horizontal lines. Plots kindly provided by Tancredi Carli.

The breakdown of the JES and JER uncertainties propagated during the unfolding procedure is shown in Figure 4.13. The JES and JER uncertainties show little variation at the jet level. For example, the JES uncertainty is about 0.4% at $p_{\mathrm{T}} = 300$ GeV and increases to 0.6% at $p_{\mathrm{T}} = 3000$ GeV in the central region [93]. In the forward region, the uncertainties are larger due to η-intercalibration modelling and non-closure effects.

However, at the cross-section level, differences emerge due to the varying slopes of the jet p_{T} distributions. These slopes become steeper at high p_{T}, especially in the forward region where the cross-sections fall rapidly. This steep behaviour leads to larger propagated JES and JER uncertainties.



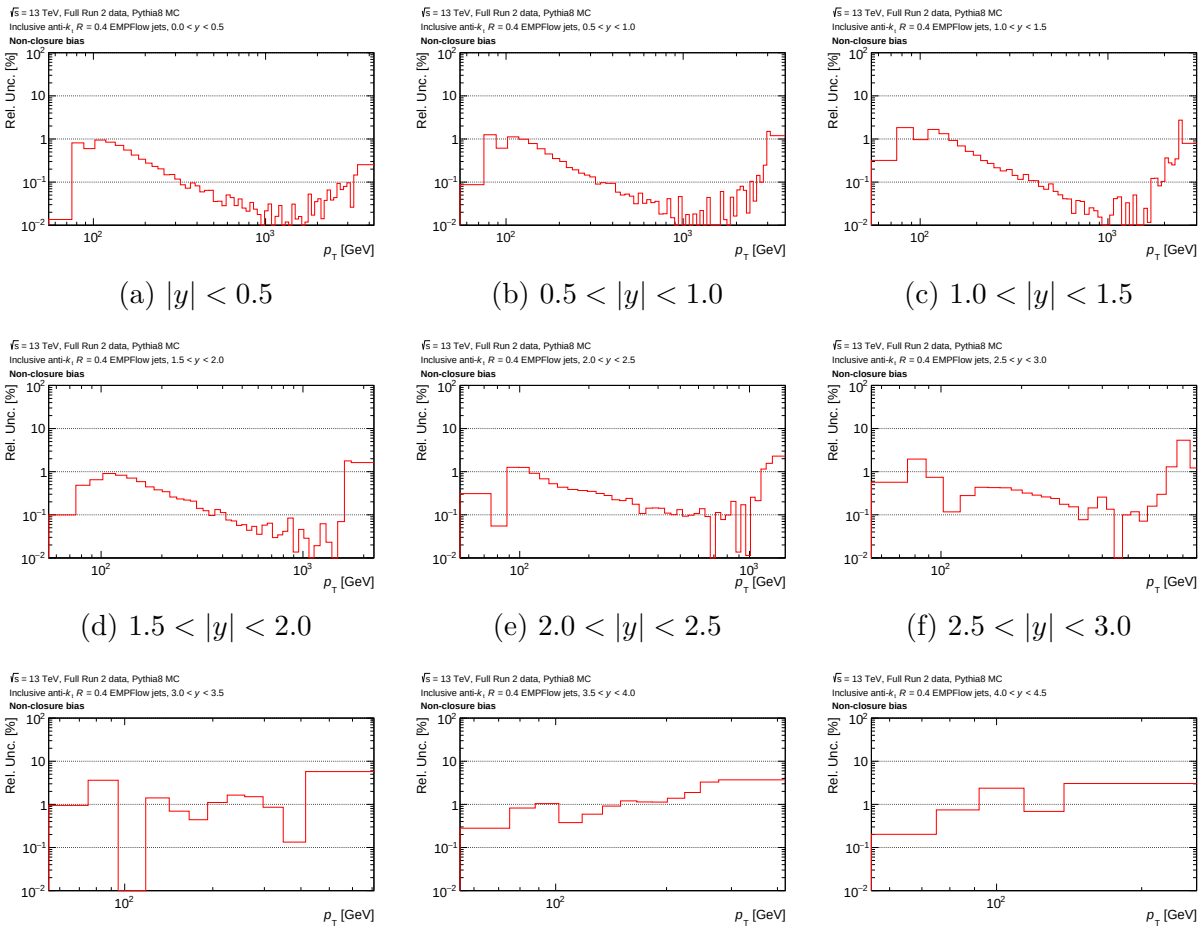
(a) $|y| < 0.5$ (b) $0.5 < |y| < 1.0$ (c) $1.0 < |y| < 1.5$

(d) $1.5 < |y| < 2.0$ (e) $2.0 < |y| < 2.5$ (f) $2.5 < |y| < 3.0$

(g) $3.0 < |y| < 3.5$ (h) $3.5 < |y| < 4.0$ (i) $4.0 < |y| < 4.5$

Figure 4.13: Systematic uncertainties arising from the propagation of JES and JER uncertainties as a function of p_{T} in rapidity bins. The uncertainties are added in quadrature within each subcategory. The peaks are due to issues related to the pile-up uncertainty, and they are corrected at a later stage.

The systematic uncertainty related to the sensitivity of the unfolding method to the shape difference between reconstructed MC and data, determined through the data-driven non-closure test (as described in Section 5.3.2.1), is shown in Figure 4.14. Finally, the statistical uncertainties in the cross-section, propagated through the unfolding using the bootstrap method, are shown in Figure 4.15. The number of iterations has been optimised in Reference [98] to optimise the quadratic sum of the two: it is set to $N = 1$.

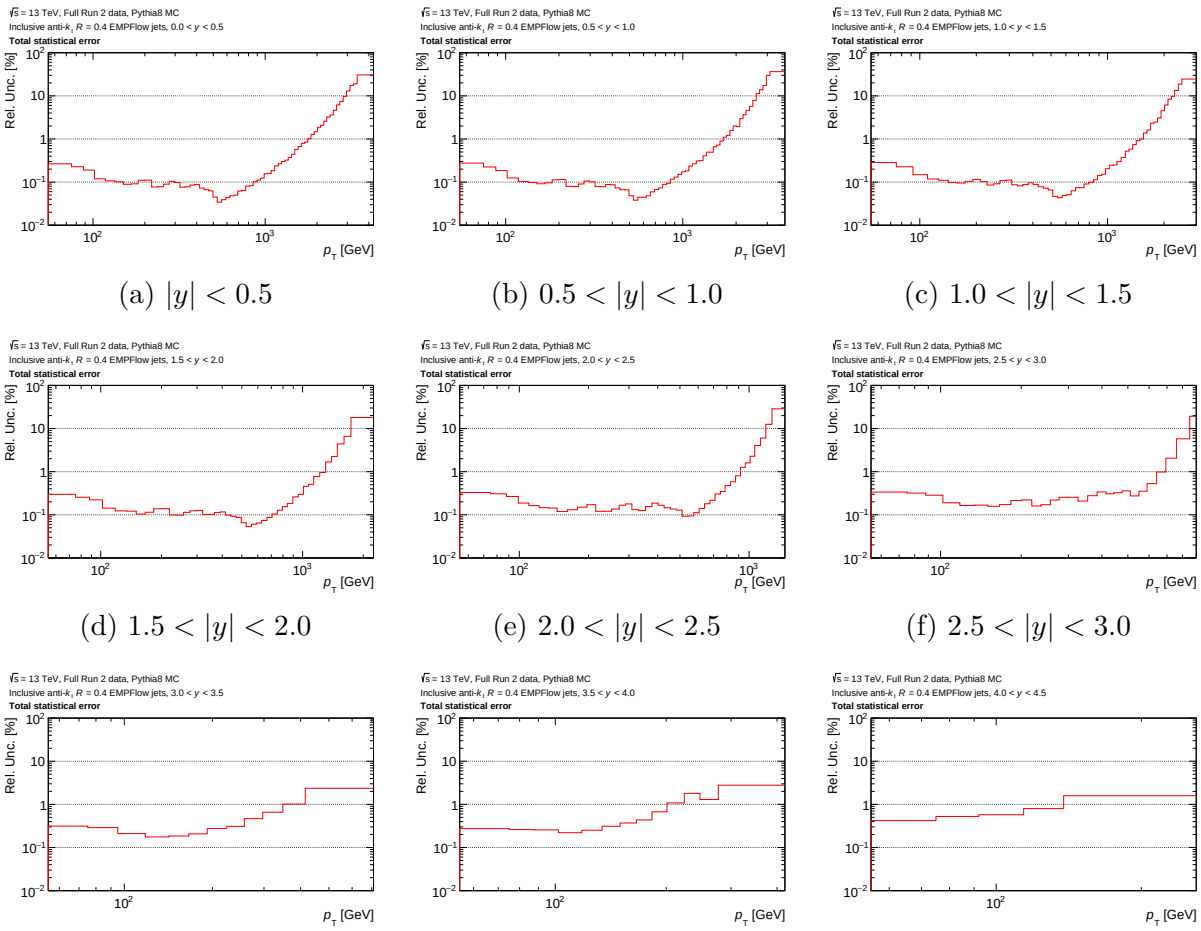


(a) $|y| < 0.5$ (b) $0.5 < |y| < 1.0$ (c) $1.0 < |y| < 1.5$

(d) $1.5 < |y| < 2.0$ (e) $2.0 < |y| < 2.5$ (f) $2.5 < |y| < 3.0$

(g) $3.0 < |y| < 3.5$ (h) $3.5 < |y| < 4.0$ (i) $4.0 < |y| < 4.5$

Figure 4.14: Uncertainty related to the sensitivity of the unfolding method to the data/MC shape difference, as a function of p_{T} in rapidity bins.



(a) $|y| < 0.5$ (b) $0.5 < |y| < 1.0$ (c) $1.0 < |y| < 1.5$

(d) $1.5 < |y| < 2.0$ (e) $2.0 < |y| < 2.5$ (f) $2.5 < |y| < 3.0$

(g) $3.0 < |y| < 3.5$ (h) $3.5 < |y| < 4.0$ (i) $4.0 < |y| < 4.5$

Figure 4.15: Statistical uncertainties in the cross-section propagated through the unfolding via the bootstrap method as a function of p_{T} in rapidity bins.

4.5.2 Comparison to theory predictions

4.5.2.1 Theory predictions

The theoretical predictions for the inclusive jet cross-section are obtained at fixed order using fixed-order perturbative QCD calculation tools (see Section 1.2.3.2). The NNLOJET++ [35, 36] framework provides predictions for the hard-scattering, on an event-by-event basis. The LHAPDF6 library [117] is used to access the PDFs, and several sets of PDFs are tested. The predictions are calculated using the scalar sum of all partons as the renormalisation scale and the factorisation scale ($\mu_{\mathrm{R}} = \mu_{\mathrm{F}} = H_{\mathrm{T,\,parton}}$). This is justified by the fact that scale choices based on the kinematics of the reconstructed jets inherit a dependence on the jet cuts and jet algorithm employed in the analysis [118].

Non-perturbative effects, such as hadronisation or the underlying event, are not simulated on an event-by-event basis; they are accounted for in the prediction through multiplicative correction factors. Those multiplicative factors are derived using Pythia 8, by turning the MPI and hadronisation models on and off and computing the bin-by-bin ratio of the inclusive jet cross-section with and without these effects. The non-perturbative correction factors obtained are close to 1, and increase up to 1.04 at low-p_{T}. Non-perturbative corrections are also derived with various Pythia 8 and Herwig 7 tunes in order to derive an uncertainty in this correction, conservatively obtained as the envelope for all studied MC tunes.

Given the high precision of both the NNLO calculations and the measurements, electroweak corrections become non-negligible and must be accounted for. They include virtual electroweak loop corrections and real emissions of weak bosons, as well as electroweak Sudakov logarithms, which take the form $\alpha \log^2(p_T^2/M_W^2)$, where α is the electroweak coupling and M_W is the W boson mass. These logarithmic terms become sizable at high energy and typically lead to a suppression of the jet production cross-section. The electroweak corrections to the inclusive jet cross-section have been calculated at fixed order by Dittmaier et al [119]. They are generally small for large $|y|$ bins and for low jet transverse momenta. However, they reach up to 12% at the highest p_{T} of 2 TeV for the central $|y|$ bin, corresponding to the phase-space region where the effects of Sudakov logarithms and virtual weak boson exchange become increasingly important.

At this stage, several choices must be made in the theoretical prediction chain, including the selection of parton distribution functions (PDFs), the values of the renormalisation and factorisation scales, the strong coupling constant α_{s}, etc. In addition, uncertainties associated with these inputs must be properly propagated through the calculations. However, fixed-order QCD predictions are computationally intensive, and recalculating them for every possible combination of parameters would be highly impractical. This is addressed more efficiently using the tools from the APPLGRID [38] (at NLO) and APPLfast [37] (at NNLO) projects, described in Section 1.2.3.2, which enables fast a posteriori convolutions of the perturbative coefficients with different choices of PDFs and scale settings.

4.5.2.2 Uncertainties

There are three main types of systematic uncertainties in the theoretical prediction:

- The scale uncertainty, estimated by varying the renormalisation and factorisation scale by a factor of two, considering all possible permutations except the opposite directions variation. Since the dependence of the cross-section on μ_R and μ_F is an artefact of the truncation of the perturbative series, this uncertainty serves as an estimate of the theoretical uncertainty due to missing higher orders.

- The strong coupling uncertainty, estimated using two additional PDF sets that assume different values of α_s, varied up or down by one standard deviation from the world average [11].

- The PDF uncertainties, provided by the PDF set authors as a collection of equally likely PDF sets (called replicas) and propagated through the calculations. The total uncertainty for each individual PDF set is the RMS of the predictions from the replicas.

In addition to the systematic uncertainties, the number of MC events generated introduces a statistical uncertainty, which amounts to 0.5% for most p_T and rapidity regions.

For NLO QCD, the renormalisation and factorisation scale uncertainties are the dominant source of uncertainty. They amount to about 5% at low p_T and increase to around 15% at high p_T. The uncertainty associated with the strong coupling constant is approximately 1–2% across the full p_T range. The PDF uncertainties are a few percent at low p_T and reach up to 10% at high p_T.

At NNLO QCD, the dependence on renormalisation and factorisation scales is significantly reduced compared to NLO. As a result, the PDF uncertainty becomes the dominant contribution to the total theoretical uncertainty.

4.5.2.3 Smoothing of the predicted jet cross-section spectra

Since the NNLO QCD calculation requires a huge amount of computing resources, it is practically impossible to further reduce the statistical uncertainty in the calculation. To mitigate statistical fluctuations in the NNLO component of cross-section predictions, a Gaussian smoothing technique is applied. This smoothing is performed on the NNLO prediction after the convolution step, using a precomputed smoothing matrix. It is then combined with the LO and NLO components to obtain the final prediction.

The smoothing process begins by normalising the NNLO cross section by dividing it by the corresponding LO cross section vector, which is much smoother. This normalisation reduces the biases arising from smoothing the steeply falling NNLO distribution directly. After normalisation, the points are convolved with a Gaussian smoothing kernel, and the LO cross-section factor is reapplied to obtain the smoothed NNLO cross-section. For computational efficiency, these steps are combined into a single Gaussian smoothing matrix transformed by the LO cross section to produce the final smoothed NNLO cross section. The smoothing can be achieved by a simple single matrix multiplication on the output results of the fast convolution.

4.5.2.4 Quantitative comparison with data

The inclusive jet cross-section measured with ATLAS Run 2 data is compared to NNLO QCD theory predictions with various NNLO PDFs, from CT18 [120], ATLASPDF [121], MHST [122] and NNPDF [123]. The individual systematic uncertainty sources are treated as uncorrelated with respect to each other. The result for $1.0 < |y| < 1.5$ is presented in Figure 4.16. The agreement between data and theory is quantified with a χ^2/NDF, in the legend of the figure. The χ^2 is calculated from data and the full covariance matrix of statistical and systematic uncertainties, including experimental and theory uncertainty. The NNLO prediction with all tested PDFs provides a similar description of the data shape, with comparable values of the associated χ^2/ndf. At lower p_{T}, they overestimated data by 10 to 20%, and they tend to underestimate at very high p_{T}. In each rapidity bin, the χ^2/NDF is generally good. The small χ^2/ndf values, in contrast with the by-eye comparison, arise from long-range correlations between the p_{T} bins.

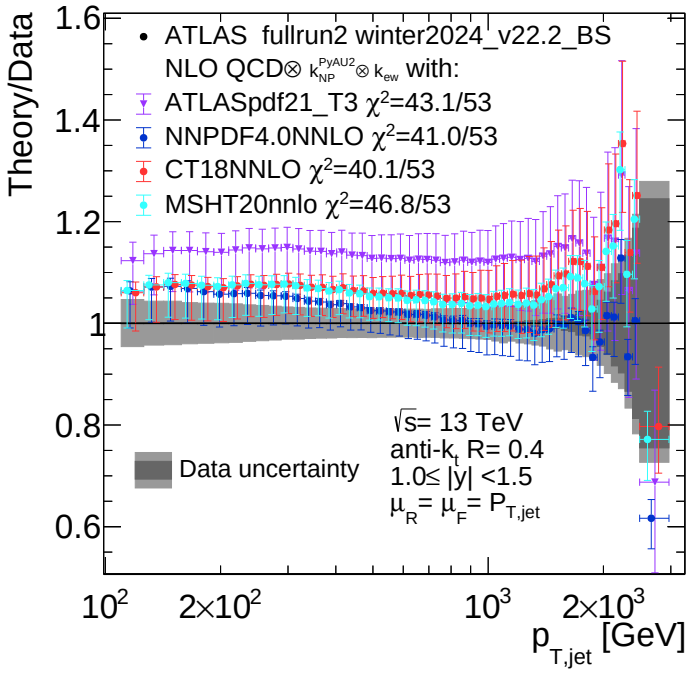


Figure 4.16: Ratio of NNLO QCD prediction to the data for various PDFs, as a function of p_{T} for $1.0 < |y| < 1.5$. The NNLO predictions are used unsmoothed in this figure. The χ^2/NDF values are indicated for each bin and PDF. The error bar indicates the PDF and scale uncertainties added in quadrature, the α_{s} uncertainty is not considered here, since it can not be computed for the ATLAS PDF. The grey bands give the statistical (dark) and systematic (light) experimental uncertainties. This plot was kindly provided by Tancredi Carli.

When taking all the rapidity bins into account, the global χ^2/NDF is very high. This tension between theory and data has already been observed in the previous measurements [109, 110].

The global χ^2/NDF is mainly sensitive to global trends, local fluctuations and correlations. Moreover, in the past, using individual η-intercalibration statistical uncertainties in (p_{T}, η)

bins instead of a single fully correlated uncertainty helped reduce the χ^2 by 200 units. Similarly, using unsmoothed η-intercalibration uncertainties also helped, yielding a further reduction of the χ^2 by 150 units.

With the updated η-intercalibration statistical uncertainties including correlations between the calibration factors, it was expected that their propagation to the inclusive jet cross section would impact the global χ^2/NDF. In addition, this propagation resulted in an overall increase in the associated uncertainty, although smaller than initially anticipated, as discussed in Section 4.4.2.2. The new χ^2/NDF, calculated after the propagation of the updated η-intercalibration statistical uncertainties, is found to be increased by a few units compared to the previous method, particularly in the forward region, starting from $2.0 < |y| < 2.5$. When considering all rapidity bins combined, the total χ^2 increases by approximately 20 units. This is likely due to the inclusion of correlations, which modify the description of the uncertainty.

4.5.2.5 QCD interpretation

The strong coupling constant is currently the least precisely known among all the fundamental couplings. A precise determination of α_s is of critical importance, as it enters the calculation of many key processes in particle physics. The current world average value of α_s is obtained from a combination of lattice QCD computations and experimental measurements, including τ-lepton decay, global PDF fits, electroweak boson decays, and hadronic final states in e^+e^- collisions [11, 124]. Jet production at LHC provides direct constraints on α_s at a very high momentum transfer. At leading order, the inclusive jet cross section is sensitive to α_s^2, and reaches up to 3 TeV.

The extraction of α_s from the inclusive jet cross-section data is currently under study. It is performed by comparing experimental measurements to NNLO QCD predictions using PDF sets fitted with different values of α_s. For each theoretical prediction, the χ^2 is evaluated, and a parabola is fitted to the χ^2 values to analytically determine the minimum. In addition, a joint PDF and α_s fit is being considered, to consistently account for correlations between the strong coupling and the parton distribution functions, potentially improving the precision and robustness of the extraction.

4.6 Conclusion

The inclusive jet production cross-section has been measured with the full Run 2 dataset in pp collisions at a centre-of-mass energy of 13 TeV. The measurement benefits from a larger dataset and extended rapidity coverage with respect to the previous similar measurement in ATLAS [109], as well as an improvement in the JES, which constitutes the dominant statistical uncertainty.

Together with the recent improvement in QCD calculation at NNLO, this measurement represents a stringent test of QCD at very high momentum transfer. It also has the potential to be included in future PDF fits. A highly competitive determination of the strong coupling constant is expected from this analysis.

In this work, I performed the unfolding of the nominal spectra and associated uncertain-

ties, as well as the evaluation of the uncertainty linked to the sensitivity of the unfolding procedure to differences in the data and MC shapes. In addition, I propagated the corrected version of the statistical uncertainty of the η-intercalibration factors, which I had previously derived, to this measurement. This represents an important contribution given the known sensitivity of the result to their modelling.

Chapter 5

The Lund Jet Plane Analysis

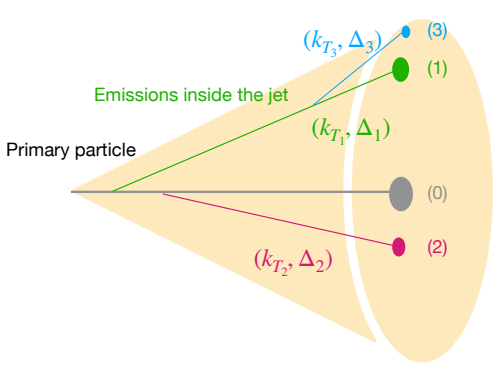
Introduction

In high-energy collisions, quarks and gluons produced in hard-scattering processes undergo fragmentation and hadronisation, ultimately forming jets. The production of jets is governed by both perturbative and non-perturbative dynamics, each influencing the precision of measurements based on jets. In addition, jets contain a wide range of particle energies spanning several orders of magnitude. This complexity makes them a rich testing ground for QCD. In recent years, the Lund Jet Plane (LJP) [51] has proven to be a powerful framework for probing jet substructure. At LHC, it has been measured by ATLAS [125, 126], CMS [127] and ALICE [128]. The LJP maps emissions within a jet in a two-dimensional space defined by their transverse momentum and angular separation. Therefore, it offers a clear distinction between perturbative and non-perturbative QCD effects, allowing the theory to be probed in different regimes. In addition, it enables tests of the running of the strong coupling constant over several orders of magnitude, benefiting from coherent systematic uncertainties. The LJP can also be used to test analytic prediction by comparing MC simulations to data, as demonstrated in Reference [129]. It also serves as a valuable tool for constraining and tuning MC event generators in the perturbative and non-perturbative regions. In addition, it has proven effective in distinguishing ("tagging") between jets originating from hadronic decay of a boosted boson (W, Z or H) or top quarks, from those arising from the fragmentation of a quark or gluon. Ongoing studies are further exploring the potential of the LJP to improve flavour tagging.

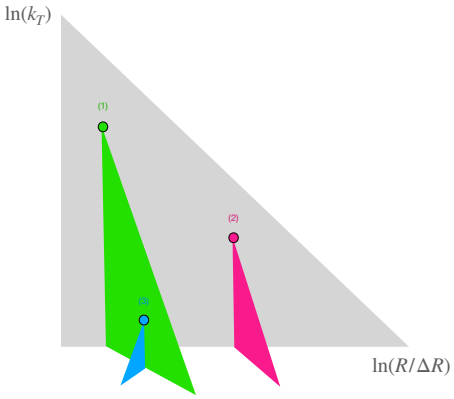
5.1 The Lund Jet Plane

5.1.1 Principle

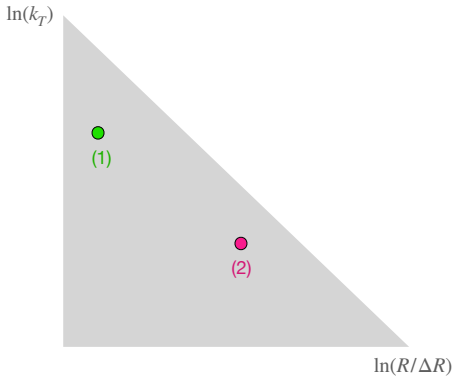
Lund diagrams are a theoretical representation of the phase space within jets. In particular, the Lund plane [130] displays the momentum fraction and angular distribution of emissions inside a jet. A graphical representation of a jet, the Lund jet diagram and primary and secondary Lund Jet Planes (LJP) are given in Figure 5.1.



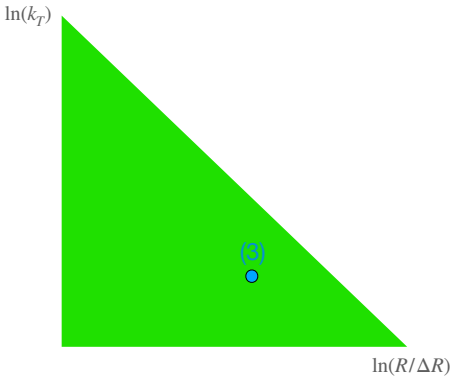
(a) Schematic representation of a jet.



(b) Lund Jet Diagram.



(c) Primary Lund Jet Plane.



(d) Secondary Lund Jet Plane.

Figure 5.1: Different representation for a jet. The figure is colour coded so that emissions visible in the schematic representation (a) are represented by triangles of the same colour in the Lund Jet Diagram (b), and points of the same colour in the primary (c) and secondary (d) Lund Jet Plane. Reproduced from [51].

The Lund plane is a powerful representation for providing information on jet substructure. However, the Lund plane is not directly observable, as it is defined in terms of quarks and gluons, which cannot be individually detected due to colour confinement. Nevertheless, a proxy for their kinematic properties can be constructed from jets, allowing the radiation patterns inside the jet to be interpreted similarly to a parton shower. An observable analogue of the Lund plane, built using jets and capturing its relevant features, has been proposed and studied in Reference [51].

The tree-like structure of k_{T}-type jet algorithms can be loosely interpreted as mimicking the parton shower evolution in reverse: the final jet corresponds to the highly energetic particle that initiated the shower, while the jet inputs represent the resulting hadrons. Following this analogy, the LJP is constructed from the transverse momenta and angles of proto-jets at the intermediate steps of the algorithm. In the Lund Jet Diagram, the phase-space of each emission inside the jet is represented as a triangle in a $(\ln \mathrm{R}/\Delta\mathrm{R}, \ln k_T)$ plane, where ΔR and k_{T} are respectively the angle and the relative transverse momentum of an emission relative to its emitter. R is a parameter relative to the jet radius, a parameter of the jet algorithm. These coordinates are dimensionless and invariant under Lorentz boosts along the z-axis. Each emission inside the jet is represented in the appropriate plane, stemming from its emitter.

The primary LJP corresponds to the emissions relative to the particle that initiated the jet. It provides a simpler and widely used representation of the jet's internal structure. Since the first splitting inside a jet is most often initiated by a gluon, the secondary LJP is less sensitive to the quark/gluon fraction in the event topology.

5.1.2 Construction

The construction of the primary LJP begins by reclustering the jets using the Cambridge-Aachen (C/A) algorithm [41, 42]. This sequential clustering algorithm groups entities based on their proximity in rapidity y and azimuth ϕ, without taking their relative energies into account. It is further described in Section 1.2.4.2. The resulting clustering structure reflects the angular ordering characteristic of gluon emissions in a parton shower.

For the reclustering, only inner-detector tracks associated with the jet are considered. This approach takes advantage of the high angular resolution of the tracker to reduce the uncertainty of the ΔR coordinate in the LJP, while being very similar to the all-particle LJP because of the approximate isospin symmetry of QCD. Each jet of the sample is then declustered along the hardest branch, following the outline below:

1. Undo the last clustering step and define two subjets j_1 and j_2, such that $p_{\mathrm{T}1} > p_{\mathrm{T}2}$.
2. Compute the following variables associated with this declustering step:

$$\Delta R \equiv (y_1 - y_2)^2 + (\phi_1 - \phi_2)^2, \quad k_t \equiv p_{\mathrm{T}2}\Delta R$$

3. Fill the LJP with the $(\ln(\mathrm{R}/\Delta\mathrm{R}), \ln(k_T))$ point,
4. Define $j = j_1$ and iterate until j is a single particle.

The advantage of the C/A algorithm is that it creates an angular-order clustering sequence that reflects the emission hierarchy, contrary to the other popular jet algorithms.

For the secondary LJP, the same procedure is followed for the lower p_{T} branch of the first declustering.

5.1.3 Analytical properties and applications

The main observable in this analysis is the average density of points per jet and per unit area of the LJP, defined by

$$\rho(k_{\mathrm{T}}, \Delta R) = \frac{1}{\mathrm{N_{jets}}} \frac{d^2(\#\text{emissions})}{d\ \ln(k_{\mathrm{T}}/\mathrm{GeV})\ d\ \ln(\mathrm{R}/\Delta\mathrm{R})}$$

To leading order in perturbative QCD, and for $\Delta R \ll R$ and $k_{\mathrm{T}} \ll p_{\mathrm{T}}^{\mathrm{jet}} \Delta R$, $\rho(k_{\mathrm{T}}, \Delta R)$ is proportional to α_{s}:

$$\rho(k_{\mathrm{T}}, \Delta R) \approx \frac{2\mathrm{C_F}\alpha_{\mathrm{s}}(k_{\mathrm{T}})}{\pi}$$

with $\mathrm{C_F} = 4/3$ the colour factor.

The emission density is expected to remain approximately uniform at high k_{T}, while increasing rapidly at low k_{T} due to the running of the strong coupling. At high R/ΔR ($\Delta R \sim R$), initial state radiation (ISR) and multi-parton interactions (MPI, UE, at lower-k_{T}) contribute to the LJP density. Hard-collinear emissions (when $p_{\mathrm{T}2}$ gets close to $p_{\mathrm{T}1}$) lie along the diagonal, corresponding to the kinematic limit. Soft-collinear emissions are emitted uniformly in the LJP. The different kinematic regimes as represented in the LJP are displayed in Figure 5.2.

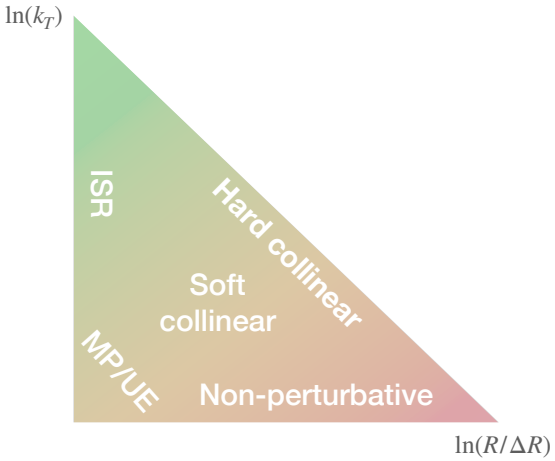


Figure 5.2: Schematic representation of the factorisation of the QCD effects in the Lund Jet Plane. Reproduced from [51].

Identifying and separating perturbative and non-perturbative regions provides a powerful advantage relative to other jet substructure variables that tend to mix perturbative and non-perturbative regimes.

5.2 Dataset

5.2.1 Data

For this measurement, proton-proton collision data are used at a centre-of-mass energy of $\sqrt{s} = 13$ TeV recorded by the ATLAS detector during the Run 2 data-taking period. It amounts to a luminosity of 140 fb^{-1}.

Jets are reconstructed from PFlow objects [86] using the anti-k_{T} algorithm [43] with parameter $R = 0.4$ and calibrated following the recommendation of [93] described in Section 3.1. Only the two leading jets of an event are considered.

Similarly to the inclusive jet cross-section measurement, events are required to pass multiple selections to ensure the good quality of the sample:

- Events that may be affected by data-taking issues, whether in LHC or ATLAS, are discarded using the ATLAS Good Run List,
- Events are required to have at least one vertex reconstructed from two or more inner-detector tracks with $p_{\mathrm{T}} > 500$ MeV,
- The event must pass at least one trigger from the trigger menu, which consists of single-jet triggers with thresholds varying from 55 GeV to 360 GeV, considered in p_{T} regions where their efficiency is above 99.9% (see Section 3.4.2.1),
- The jet must be considered in-time relative to the collision that produced the event. This is achieved via a simple cut $|t_{\mathrm{jet}}| < 16$ ns where t_{jet} is computed relative to the event time recorded by the trigger,
- To reject events contaminated by non-collision backgrounds, calorimeter noise bursts, readout errors and other pathologies, a jet cleaning procedure is applied (presented in Section 2.3.3.2),
- The transverse momentum of the jet is required to be above 60 GeV and the absolute rapidity is required to be below 2.0.
- To target dijet topologies in the sample, a kinematic requirement is placed on the p_{T} of the leading jet pair in the event: $p_{\mathrm{T}}^{\mathrm{leading}}/p_{\mathrm{T}}^{\mathrm{subleading}} < 1.5$.

The LJP is then defined differentially in $\ln(k_{\mathrm{T}})$ and $\ln(\mathrm{R}/\Delta\mathrm{R})$ in bins of the jet p_{T}. The p_{T} and $\ln(k_{\mathrm{T}})$ binning are chosen to ensure good statistics in each bin, and the borders of the eight p_{T} bins are set to coincide with those of the inclusive jet measurement: the p_{T} binning is then set to [183, 277, 415, 657, 1034, 2011, 4256] GeV. The $\ln(k_{\mathrm{T}})$ binning is optimised separately for the primary and secondary Lund Jet Plane. Additional bins at both high and low values are included to account for migrations into and out of the edge bins in the unfolding procedure. The $\ln(\mathrm{R}/\Delta\mathrm{R})$ binning goes from 0 to 5 with a step of 0.5 for both the LJPs. The eight 2-dimensional LJPs are flattened into a one-dimensional $\ln(k_{\mathrm{T}}) - \ln(\mathrm{R}/\Delta\mathrm{R}) - p_{\mathrm{T}}$ variable for the unfolding, and a rebinning procedure is applied to the one-dimensional variable to ensure that both migration efficiency and purity exceed 30%. The final distributions (after unfolding) are normalised to the number of jets entering the measurement.

5.2.2 Monte Carlo simulations

For the nominal sample, MC simulations of dijet and multijet events [131] generated using Pythia 8 [21] are used. Simulated minimum-bias interactions generated using Pythia 8 are superimposed with all simulations. The interactions of stable particles in the ATLAS detector are fully simulated using Geant4 [90]. The distribution of the average number of pile-up interactions in simulations is re-weighted during data analysis to match that observed in the Run 2 data.

Except for the GRL, MC reconstructed events are required to fulfil the same conditions as for the data sample. Truth-level jets are reconstructed from all detector-stable particles from the hard-scattering process with a lifetime τ such that $c\tau > 10$ mm, except for neutrinos and muons, which leave negligible energy in the calorimeter. In addition, truth-level jets are required to have a $p_{\mathrm{T}} > 60$ GeV and a rapidity y satisfying $|y| < 2.0$ to be selected for this analysis.

5.3 Unfolding

The goal of unfolding is to correct the measured data for effects arising from the finite efficiency, acceptance, and resolution of the detector. The motivation and general principle of unfolding, as well as the details of the IDS method, have been thoroughly explained in Section 4.3. This section focuses on the contributions of the present work: comparing two regularisation methods for the unfolding and optimising the number of unfolding iterations for this analysis. Since in this analysis the LJP is a three-dimensional observable (k_{T}, ΔR, jet p_{T}), the data and MC spectra are flattened into a one-dimensional observable by concatenating the LJP columns for each p_{T} bin, and then concatenating the results of each p_{T} bin. In the following, the spectra will be given along the $k_{\mathrm{T}} - \Delta R - p_{\mathrm{T}}$ axis. The folding matrix is constructed by matching truth-level and reconstructed-level emissions that have a transverse momentum above 0.5 GeV and lie within a radius of $R = 0.1$.

5.3.1 IDS vs IBU regularisation

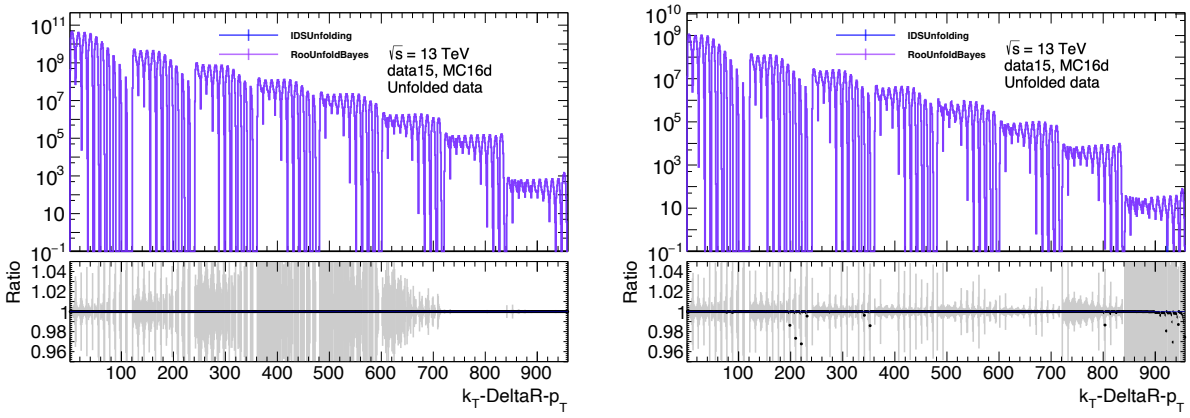
5.3.1.1 IBU regularisation principle

The Iterative Bayesian Unfolding (IBU) is based on the same principle as the IDS method, presented in Section 4.3.2: it uses Bayes' theorem to iteratively improve the estimate of the true distribution. The difference lies in the choice of the prior and regularisation scheme. The IBU method uses the regularisation proposed by d'Agostini in Reference [132]. In the IBU method, the prior used to compute the transfer matrix is flat. In other words, equal probabilities in all bins are assumed for the truth variable. Then, at each iteration, the unfolded spectrum is smoothed to suppress statistical fluctuations by a model-dependent fitting (e.g. a polynomial fit) and used as a prior for the next iteration. More recent implementations of the IBU method allow for the choice of prior by the user [133]. In the IDS method, the MC truth spectrum is used as the starting point, and the unfolding correction is dynamically adjusted in each bin depending on the statistical significance of the difference between data and reconstructed MC. The IDS method was developed to improve the stability and reliability of unfolding results, particularly in situations with

limited statistics or large migrations between bins.

5.3.1.2 Comparison with IDS

The two methods are applied to the same samples and compared. For this comparison, the number of unfolding iterations is set to one. The results and their ratio are shown in Figure 5.3. As anticipated, for the Monte Carlo simulations, both methods yield identical results: with any method, unfolding the MC reconstructed spectra must result in the MC true spectra. When unfolding the data, close agreement is found within the statistical uncertainties. These findings underscore the robustness and reliability of both unfolding techniques for accurately reconstructing underlying distributions from experimental measurements.



(a) MC (b) Data

Figure 5.3: Comparison of IDS and IBU method, for the unfolding of the data (a) and MC (b) spectra. The bottom panel displays the ratio of the unfolded spectra obtained using the two methods, with statistical uncertainties shown in grey.

The IDS method was ultimately chosen, as it is also used in the inclusive jet cross-section measurement, enabling a consistent correlation between the two measurements. This study was performed prior to the decision to exclude the first two p_{T} bins, which correspond to the first 300 bins of the one-dimensional LJP variable. In the remainder of the analysis, these bins are excluded.

5.3.2 Choice of the number of unfolding iterations

Each unfolding iteration reduces the bias caused by an imperfect MC model and shape differences between data and MC, but leads to increased statistical uncertainties. The number of iterations must be chosen carefully to balance bias reduction and the control of statistical uncertainties.

5.3.2.1 Data-driven non-closure test

The data-driven non-closure test provides an estimate of the bias arising from discrepancies between data and MC. It is described in detail in Section 4.3.4. The main difference

relative to the procedure used in the inclusive jet cross-section analysis is how the MC is reweighted to match the shape of the data spectrum at the reconstructed level better. Unlike the inclusive jet spectra, the LJP one-dimensional spectra are not easily described by an analytic function. Therefore, reweighting is performed by scaling each column of the transfer matrix by the ratio of data to reconstructed MC corresponding to that column. Afterwards, the reweighted MC reconstructed spectrum is scaled to match the total number of events observed in the data. To improve the agreement between data and reconstructed MC, this procedure can be iterated, using the projection along the Y-axis of the scaled transfer matrix as the updated MC reco.

Figure 5.4 shows the data to MC ratio after one to five reweighting iterations, and the initial ratio is restricted to the first p_{T} bin for readability. The uncertainties displayed on the Figure are the statistical uncertainties inly. They are estimated on a bin-by-bin basis by propagating the square root of the event count to the final ratio using standard error propagation techniques. While the initial ratio reveals significant discrepancies between data and MC, the first reweighting step already leads to a substantial reduction in these differences. The second reweighting further suppresses the systematic bias observed in the previous ratio. With each additional iteration, the differences continue to decrease, although the improvement becomes progressively smaller. Two reweighting iterations are found to be satisfactory for the data-driven non-closure test, as they already recover most of the data-MC shape differences.

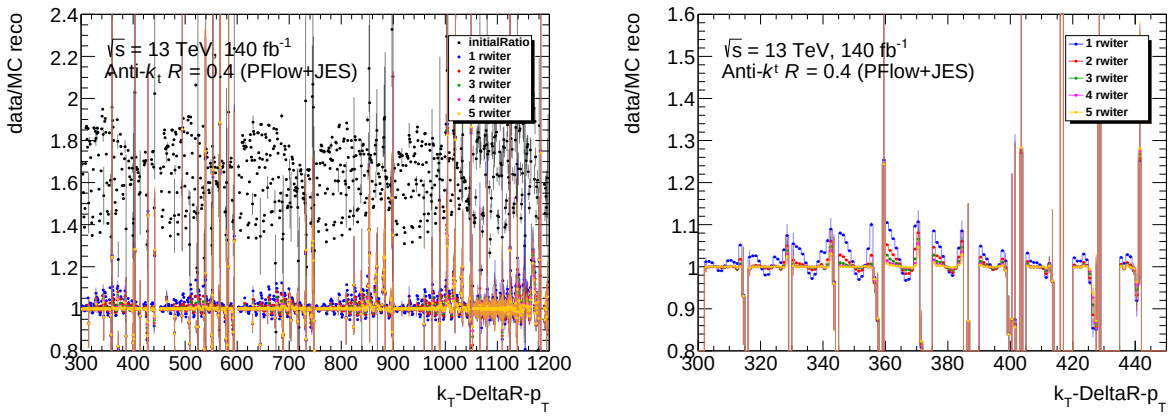


Figure 5.4: Ratio of data to reconstructed MC, without reweighting and after several instances of reweighting. The initial ratio for the whole plane is shown in (a), while (b) is a zoom on the first p_{T} bin, and without the initial ratio for readability. The error bars represent the statistical uncertainties inly.

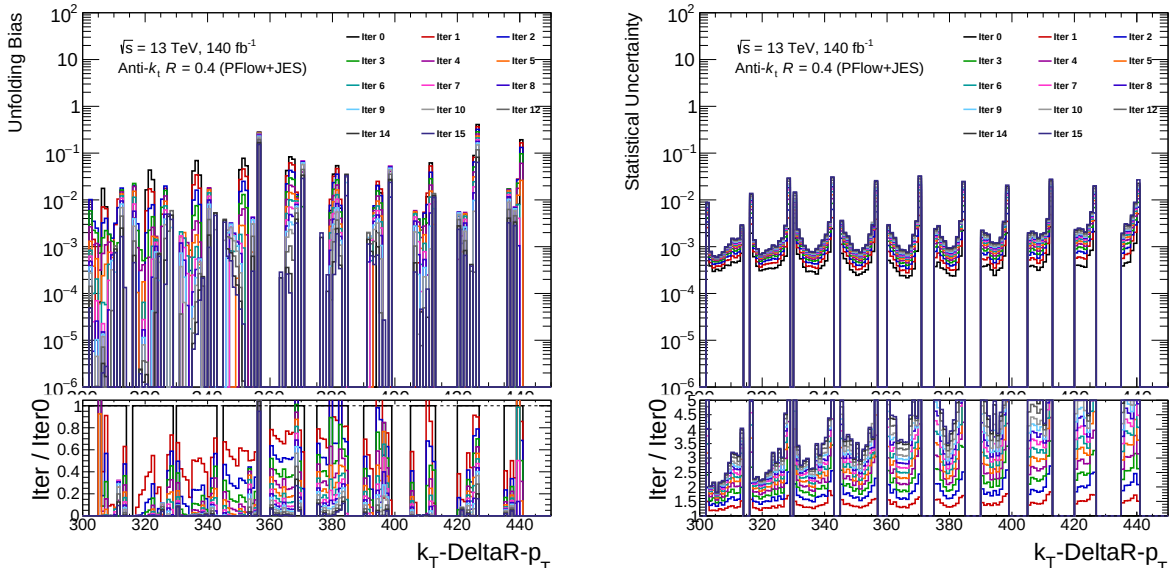
The non-closure test is then performed as described in Section 4.3.4: the reweighted MC spectrum is unfolded, and the resulting unfolded spectrum is compared to the true spectrum. An uncertainty in the unfolding bias is estimated using the bootstrap method, which is thoroughly explained in Section 3.3.3.1: one hundred bootstrap replicas are generated by assigning a deterministically generated weight drawn from a Poisson distribution with a mean of unity, and the test is repeated for each replica. The RMS of the resulting bootstrap distribution is taken as the statistical uncertainty.

5.3.2.2 Statistical uncertainties

The bias related to the discrepancies between data and MC is to be compared with the statistical uncertainties of the unfolded spectrum. Indeed, they evolve in opposite directions with the number of iterations: the statistical uncertainties grow with each iteration, while the bias decreases. The statistical uncertainties are estimated using the bootstrap method, and the statistical uncertainties in each bin are taken as the standard deviation of the corresponding bootstrap distribution.

5.3.2.3 Final choice

The non-closure bias is displayed alongside the statistical uncertainties for a given p_{T} bin, in Figure 5.5. Although the p_{T} and k_{T} binning have been optimised, the LJP suffers from quite large statistical uncertainties in some areas, namely at high k_{T}, but overall the statistical uncertainties are relatively small ($< 1\%$ in the bulk of the distributions). At high k_{T}, the non-closure bias is significant, even after many iterations, due to the large statistical uncertainties causing instabilities in the unfolding procedure.



(a) (b)

Figure 5.5: (a) Non-closure bias and (b) statistical uncertainties, for a given p_{T} bin, shown for several numbers of unfolding iterations.

Figure 5.6 shows the quadratic sum of the statistical uncertainties and the unfolding bias for a given p_{T} bin, evaluated for different numbers of unfolding iterations. In principle, this sum is expected to decrease initially and then rise again if too many iterations are used, as statistical fluctuations begin to dominate. However, within the range of reasonably tested iterations, this rise was not observed: the total uncertainty continued to decrease. Beyond four iterations, the reduction becomes marginal in the bulk of the distribution. Therefore, four unfolding iterations were chosen for the analysis. The unfolding bias arising from discrepancies between data and MC shapes for four unfolding iterations is shown in Figure 5.7. It is reasonably low in most of the distribution, but gets high where the statistical precision gets low, i.e. at high $\ln(k_{\mathrm{T}})$.

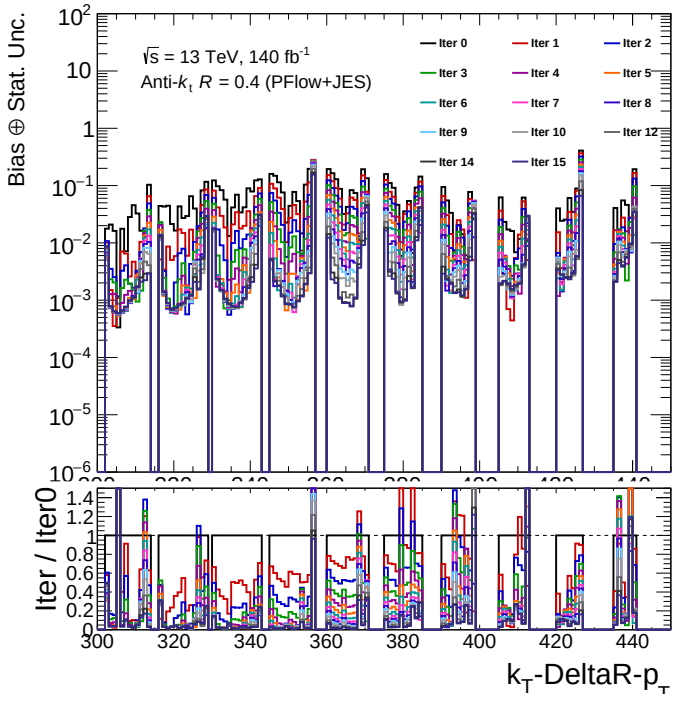


Figure 5.6: Quadratic sum of the statistical uncertainties and unfolding bias for a given p_{T} bin, shown for several numbers of unfolding iterations. The bottom panel shows the ratio of this quantity to the corresponding value obtained after a single unfolding step (0 iterations).

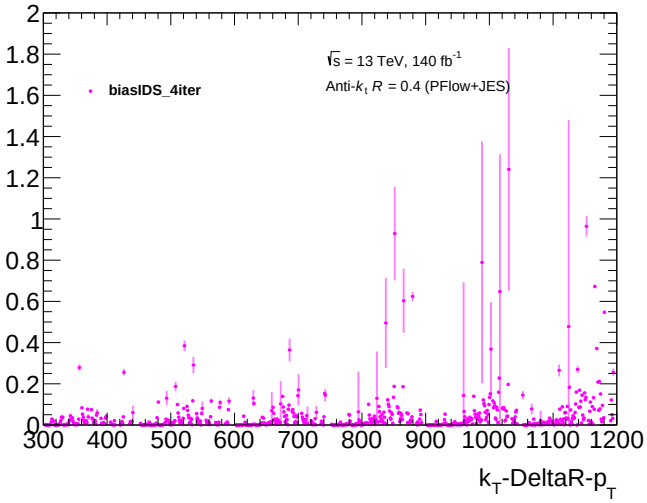


Figure 5.7: Non-closure bias estimated through the data-driven non-closure test, with four unfolding iterations. The error bars display the statistical uncertainties, evaluated through the bootstrap method.

5.4 Experimental uncertainties

Multiple sources of uncertainty affect the LJP measurements:

- Statistical uncertainties linked to the finite amount of data,
- Systematic uncertainties related to the application of the JES and the JER,
- Systematic uncertainties related to the tracking,
- Systematic uncertainties related to the pile-up reweighting,
- Bias related to the sensitivity of the IDS unfolding to the data/MC shape difference,
- Systematic uncertainties due to the choice of MC model for the unfolding procedure,
- uncertainty in the luminosity measurement.

5.4.1 Statistical uncertainties

Statistical uncertainties are estimated using the bootstrap method, thoroughly explained in Section 3.3.3.1: one hundred bootstrap replicas for every distribution are generated by assigning a deterministically generated weight drawn from a Poisson distribution with a mean of unity. For the Monte Carlo, the response matrix truth and reco spectra are all allowed to vary during this process, while for the data, the input data distributions to the unfolding procedure are varied. The unfolding of each replica dataset is performed with the associated replica transfer matrix. The statistical uncertainties are retrieved as the RMS of the bootstrap distribution in each bin. They are shown in Figure 5.8. The statistical uncertainties remain under 1% across the whole plane, and are larger along the diagonal for most p_{T} bins. For the highest p_{T} values, the statistical uncertainty reaches very high values along the diagonal.

5.4.2 Systematic uncertainties

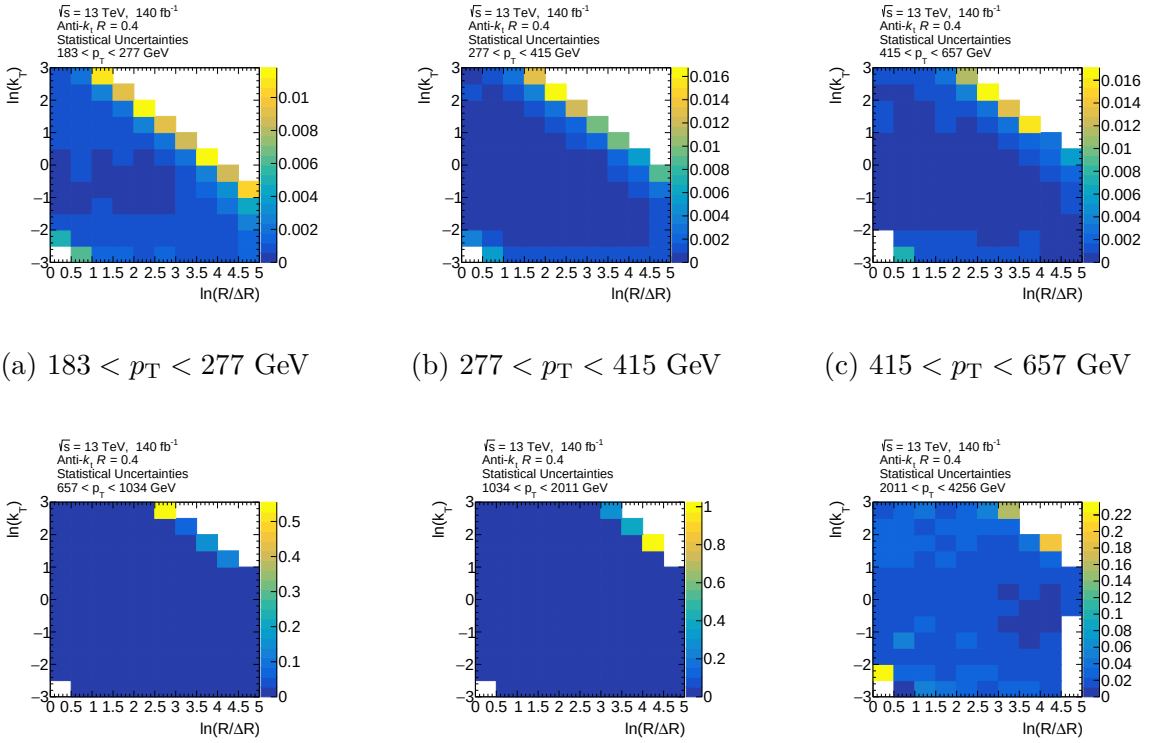
5.4.2.1 JES and JER uncertainties

Uncertainties related to the application of the JES and the JER are derived, similarly to the inclusive jet cross-section analysis. Namely, for the JES, the jet p_{T} is scaled and the unfolding is performed with the nominal matrix, while for the JER, the jet p_{T} is smeared in a pseudo-data spectrum and the transfer matrix and the unfolding are performed. The JES/JER uncertainties are expected to be one of the largest uncertainties in many bins of this measurement. They have not been derived yet at this stage of the analysis.

5.4.2.2 Tracking related uncertainties

As tracks are the primary objects from which emissions within the C/A clustering history are reconstructed, it is necessary to account for systematic uncertainties to prevent potential biases from affecting the track-based measurement. The track reconstruction process is explained in Section 2.3.1.1. In this analysis, systematic uncertainties are evaluated for track efficiencies (inclusive and inside jets), fake rate, and reconstructed parameters of all tracks.

The inclusive tracking inefficiency is due to interactions in the passive material of the detector. It is well studied using photon conversions and hadronic interactions [134]. In practice, the uncertainty is propagated by randomly dropping tracks, depending on their p_{T} and $|\eta|$. The tracking inefficiency inside jets occurs because of charge deposits from multiple particles in the same sensor. This happens mostly in dense environments, i.e. in the core of jets. It is studied by comparing the rate of pixel clusters assigned to single



(a) $183 < p_{\mathrm{T}} < 277$ GeV (b) $277 < p_{\mathrm{T}} < 415$ GeV (c) $415 < p_{\mathrm{T}} < 657$ GeV

(d) $657 < p_{\mathrm{T}} < 1034$ GeV (e) $1034 < p_{\mathrm{T}} < 2011$ GeV (f) $2011 < p_{\mathrm{T}} < 4256$ GeV

Figure 5.8: Relative statistical uncertainties propagated through the unfolding via the bootstrap method, in p_{T} bins.

tracks in the core of jets in data and simulations [135]. It is implemented by randomly dropping tracks inside of jets, depending on their p_{T} and $|\eta|$.

Fake tracks arise when a random combination of hits happens to overlap in space, or in the case of misallocation of clusters from other particles. It is highly dependent on the pile-up, and it also increases inside jets due to the high density of charge particles. Fake tracks are investigated by defining a fake-enriched control region based on Monte Carlo truth information, and performing a template fit to extract an estimate of the fraction of real and fake tracks that can be compared between data and Monte Carlo simulations [136]. From this study, the rate of fake tracks is estimated, and the uncertainty is implemented by randomly dropping tracks depending on their p_{T} and $|\eta|$.

The sagitta bias alters the measured q/p for charged particle tracks. This bias depends on both the charge of the particle and environmental conditions that can vary across data-taking periods (e.g. temperature changes that cause detector deformations). To correct for this effect, a small correction is applied. This correction is derived from resonance decays to muon pairs and from E/p measurements of isolated electrons [137]. The associated uncertainty is small and is accounted for in the analysis by smearing the track kinematics accordingly.

In all three cases, alternative samples are created with the corrections implemented, and

the uncertainties are propagated through the unfolding.

The tracking uncertainties for the first p_{T} bin are shown in Figure 5.9, alongside their statistical uncertainties and the breakdown into their subcomponents, while results for the remaining p_{T} bins are presented in the Appendix. In most regions of the plane, the uncertainties remain below 5%, but they increase along the diagonal, reaching up to 30%. The dominant contribution comes from the inclusive tracking efficiency.

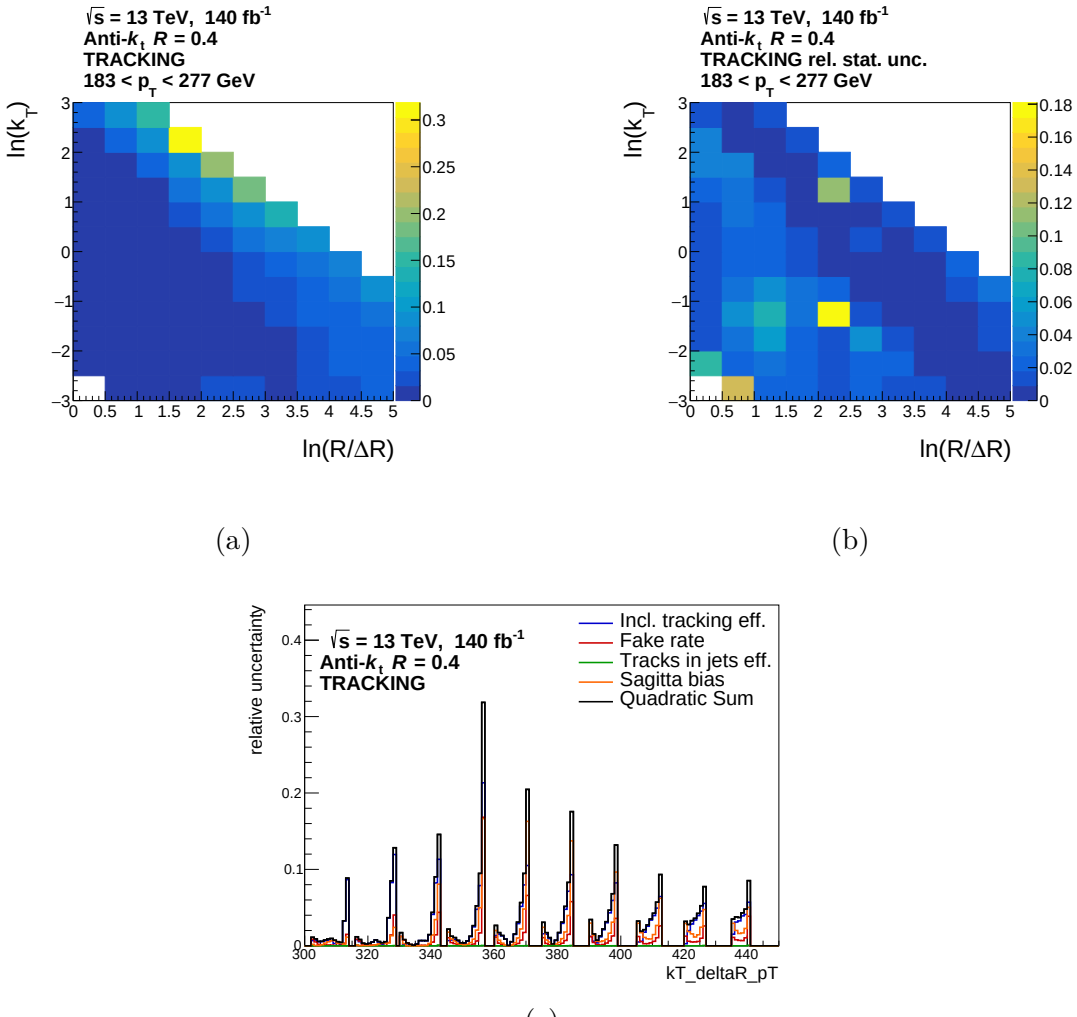


Figure 5.9: Total uncertainties due to tracking (a) and associated relative statistical uncertainties (b) and breakdown of the uncertainties into its subcomponents (c), for $183 < p_{\mathrm{T}} < 277$ GeV

5.4.2.3 Modelling uncertainty

An uncertainty is derived to account for the model dependence of the unfolding. Indeed, Pythia is chosen as the nominal MC generator to perform the unfolding, but using a different generator can lead to variations in the unfolded result. To evaluate this uncertainty,

the unfolding is repeated using the standard PYTHIA prior, but with the transfer matrix, efficiency and purity factors from an alternative MC simulation. The alternative MC generators used to evaluate this uncertainty are SHERPA 2.2.11, POWHEG+PYTHIA 8, and HERWIG 7.1. Both PYTHIA and POWHEG+PYTHIA 8 employ a dipole-style, p_{T}-ordered parton shower and the Lund string hadronisation model. In contrast, HERWIG uses an angular-ordered parton shower and a cluster hadronisation model. SHERPA uses a dipole parton shower combined with cluster hadronisation.

Uncertainties related to prior variations are accounted for through the data-driven non-closure uncertainty (see Section 5.3.2.1), and are therefore not evaluated here to avoid double-counting.

Modelling uncertainties are shown in Figure 5.10 for the first p_{T} bin, and for the SHERPA and POWHEG+PYTHIA samples only, due to a technical issue affecting the HERWIG sample. These uncertainties are among the dominant sources in the LJP measurement, as expected given the sensitivity of this observable to parton shower and hadronisation effects. They can reach up to 40% in the lower region of the plane and decrease along the diagonal, where statistical uncertainties become more significant.

To identify which component (efficiency, purity, or transfer matrix) has the largest impact, the uncertainty propagation is repeated while varying only one of these inputs at a time. The result is presented for the POWHEG+PYTHIA generator and for the first p_{T} bin in Figure 5.11. The dominant contribution arises from the modelling of the transfer matrix. This is the case for both generators tested and for all p_{T} bins.

5.4.2.4 Pile-up reweighting

As described in Section 2.3.4, the distribution of the average number of pile-up interactions in simulations is re-weighted during data analysis to match that observed in the Run 2 data. To assess the uncertainty associated with the pile-up reweighting procedure, the scale factor applied to the pile-up distribution is varied up and down from its nominal value, following ATLAS recommendations. In addition, another dataset is derived with the pile-up reweighting turned off. The resulting alternative datasets are then unfolded to propagate this uncertainty to the final measurement.

The pile-up reweighting uncertainties are shown in Figure 5.12, for the first p_{T} bin, and in the Appendix for the rest. They remain below 3% across most of the Lund plane, with a moderate increase along the diagonal.

5.4.2.5 Other uncertainties

An uncertainty is derived to account for the bias related to the sensitivity of the IDS unfolding to the data/MC shape difference. It is evaluated through the data-driven closure test, explained in Section 4.3.4. The bias evaluated is simply taken as an uncertainty in the measurement. The bias is shown in Figure 5.13 for all p_{T} bins. It remains below 5% across most of the plane but can become significantly larger along the diagonal, where statistical precision is limited.

Similarly to the jet cross-section analysis, the relative uncertainty in the total luminosity

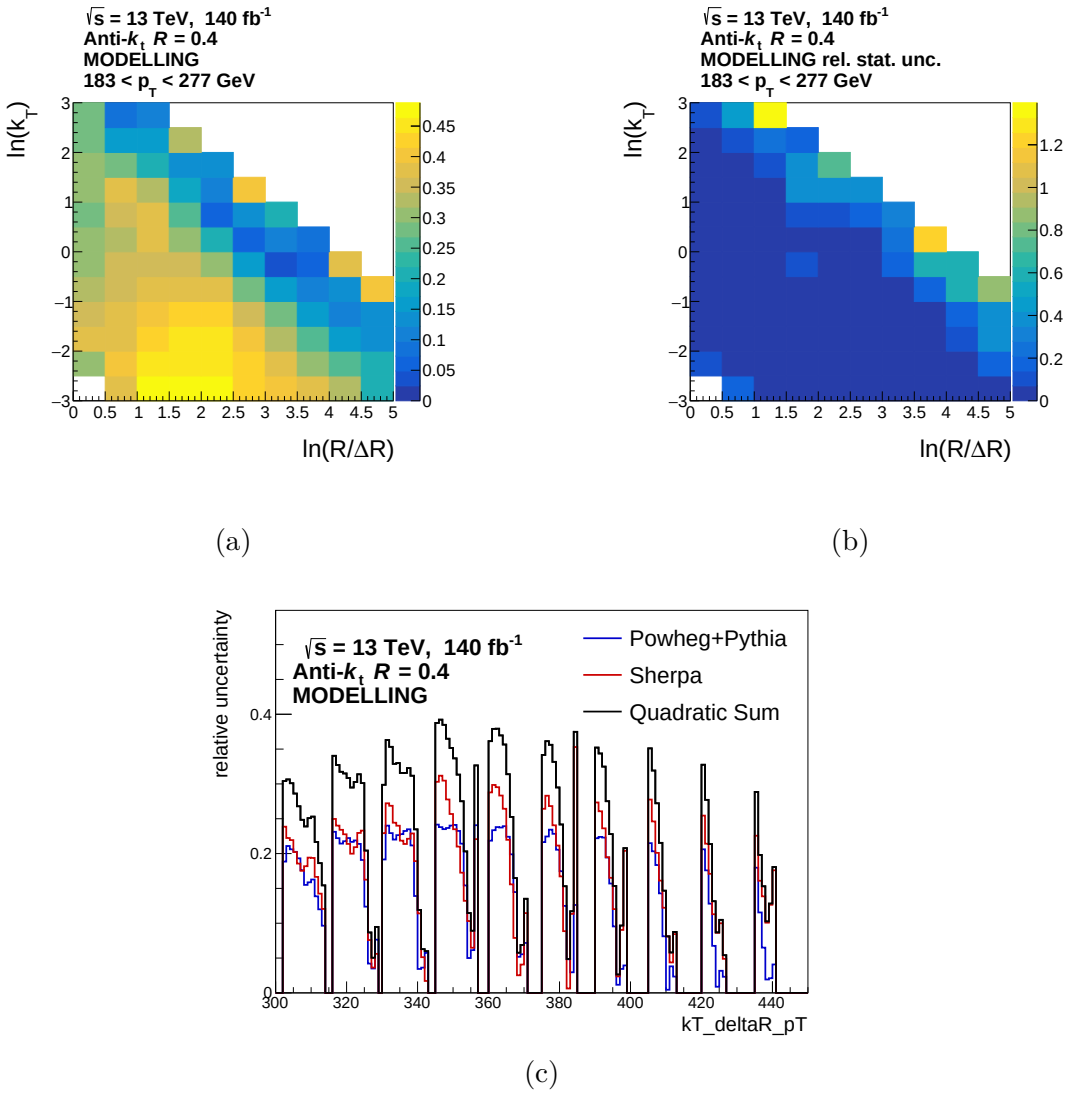


Figure 5.10: Total modelling uncertainties (a) and associated relative statistical uncertainties (b) and breakdown of the uncertainties into its subcomponents (c), for $183 < p_{\mathrm{T}} < 277$ GeV

is applied as a variation of the normalisation of the nominal MC sample. It translates into a fully correlated uncertainty of the same magnitude across the entire phase-space.

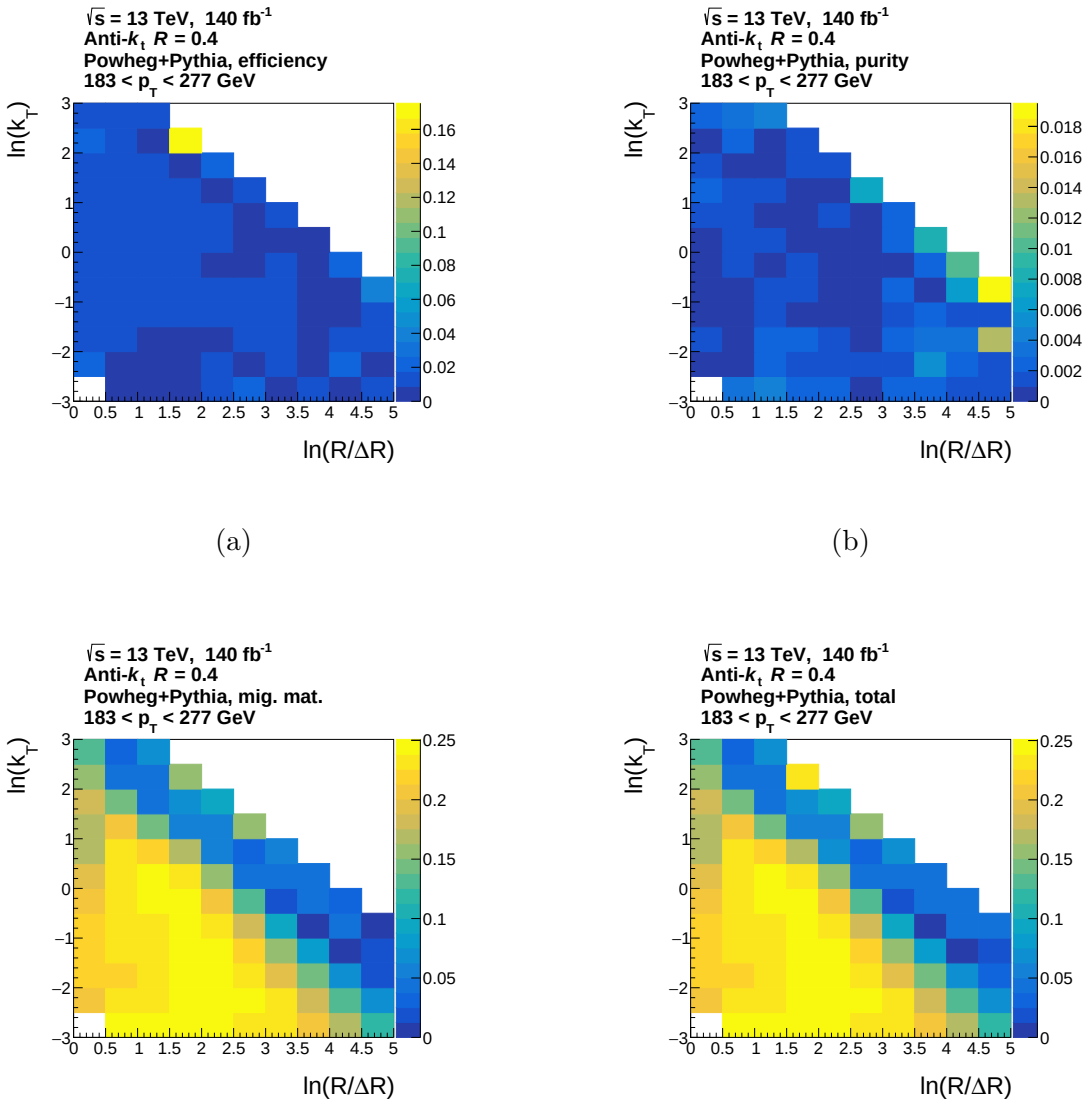
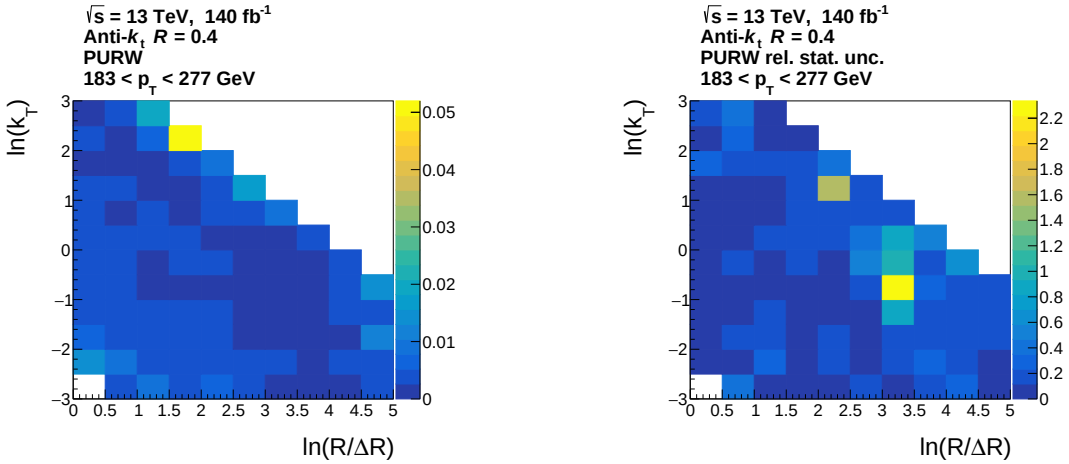
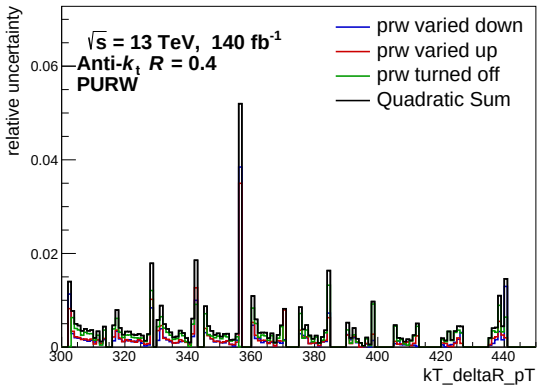


Figure 5.11: Modelling uncertainties, using POWHEG+PYTHIA generator for efficiency correction (a) purity correction (b) transfer matrix (c), or all (d) for $183 < p_{\mathrm{T}} < 277$ GeV

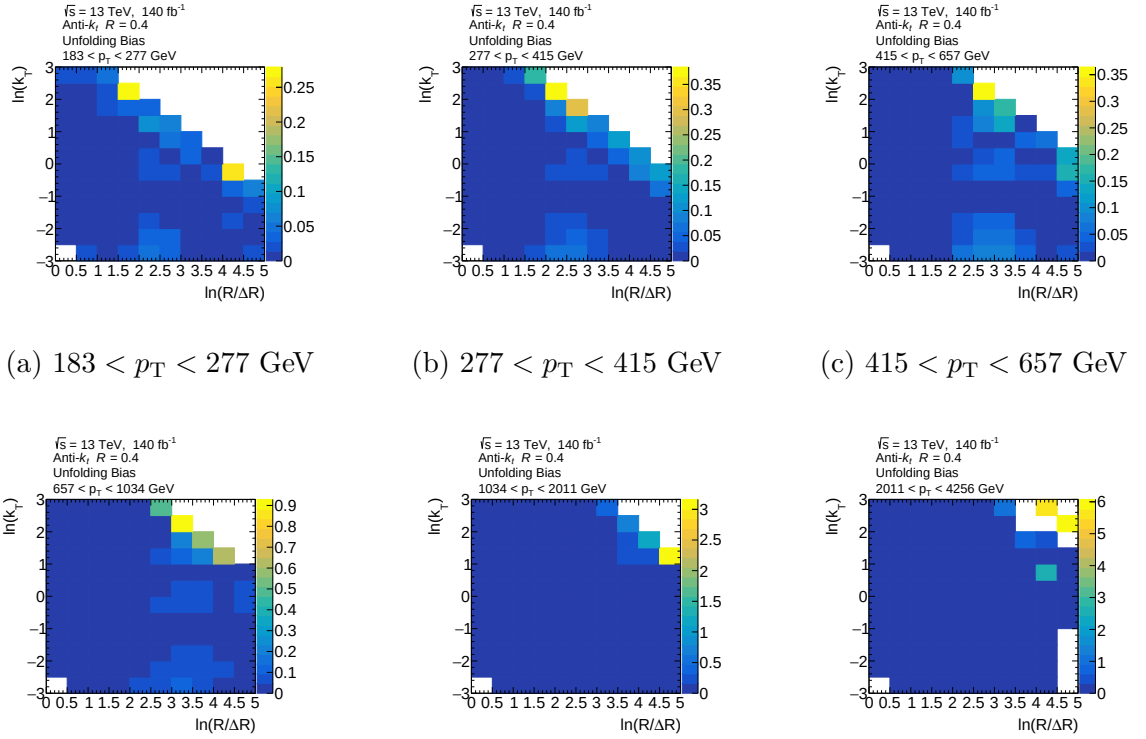


(a) (b)



(c)

Figure 5.12: Total uncertainties due to the pile-up reweighting procedure (a) and associated relative statistical uncertainties (b) and breakdown of the uncertainties into its subcomponents (c), for $183 < p_{\mathrm{T}} < 277$ GeV



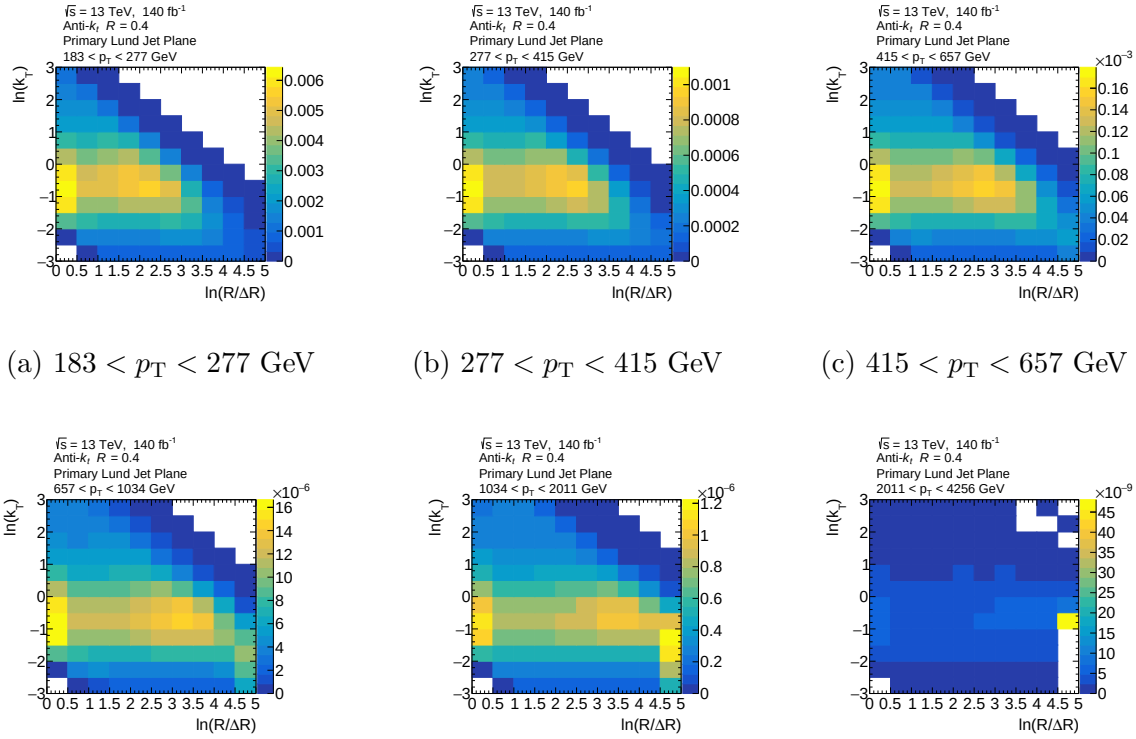
(a) $183 < p_{\mathrm{T}} < 277$ GeV (b) $277 < p_{\mathrm{T}} < 415$ GeV (c) $415 < p_{\mathrm{T}} < 657$ GeV

(d) $657 < p_{\mathrm{T}} < 1034$ GeV (e) $1034 < p_{\mathrm{T}} < 2011$ GeV (f) $2011 < p_{\mathrm{T}} < 4256$ GeV

Figure 5.13: Uncertainty related to the sensitivity of the unfolding method to the difference in shape between data and MC, in p_{T} bins.

5.5 Results

The unfolded primary LJP is shown for various p_{T} bins in Figure 5.14. A triangular region of the LJP is populated nearly uniformly by perturbative emissions. The density of emissions in this region is shaped by the running of the strong coupling $\alpha_{\mathrm{s}}(k_{\mathrm{T}})$. At low transverse scales, $\alpha_{\mathrm{s}}(k_{\mathrm{T}})$ becomes large, giving rise to the yellow band around $\ln(k_{\mathrm{T}}) \sim -1$. In this regime, the LJP density can no longer be reliably described with perturbation theory. For smaller ln(R/ΔR), additional contributions from initial-state radiation and multi-parton interactions enhance the density. The total uncertainty is not yet available at this stage of the analysis.



(a) $183 < p_{\mathrm{T}} < 277$ GeV (b) $277 < p_{\mathrm{T}} < 415$ GeV (c) $415 < p_{\mathrm{T}} < 657$ GeV

(d) $657 < p_{\mathrm{T}} < 1034$ GeV (e) $1034 < p_{\mathrm{T}} < 2011$ GeV (f) $2011 < p_{\mathrm{T}} < 4256$ GeV

Figure 5.14: Primary Lund Jet Plane normalised to the number of jets entering the measurement, in p_{T} bins.

5.6 Conclusion and perspective

The LJP analysis is currently ongoing, with several components still under development. My contributions so far have focused on the unfolding of the nominal spectra, including the selection of the unfolding method, determination of the optimal number of iterations, and validation through the technical closure test. I also implemented a data-driven closure test to assess the sensitivity of the unfolding method to differences in shape between data and MC. I am currently contributing to the propagation of systematic uncertainties through the unfolding process, ensuring their correct treatment and impact on the final

measurement. Key systematic uncertainties, such as the JES and JER, have yet to be fully evaluated. In parallel, studies focusing on the secondary LJP are actively being pursued. Once the full set of experimental results is finalised and the associated uncertainties are thoroughly validated, the next step is to compare the measurements with state-of-the-art theoretical predictions. This comparison will enable the extraction of the strong coupling constant, α_s. To facilitate this comparison, the LJP analysis has been designed to be closely aligned with the jet cross-section measurement strategy. Several deliberate choices have been made to ensure consistency between the two, including the adoption of coherent p_T binning schemes, event selection criteria, and unfolding techniques. This will allow for a meaningful correlation between the LJP and jet cross-section results in the final interpretation.

Chapter 6

Jet studies in an FCC-ee environment

Introduction

The discovery of the Higgs boson at the LHC (consistent with predictions from earlier precision measurements at LEP and other colliders) has completed the list of elementary particles and interactions described by the Standard Model. This achievement has set the stage for a new era of exploration in particle physics. However, many open and fundamental questions remain unanswered by the SM as it stands. Among them are the origin of the matter-antimatter asymmetry in the universe, the nature of dark matter, the mechanism behind the non-zero mass of neutrinos in particular, the large discrepancy in the strength of the fundamental forces, and many more. These open issues strongly suggest that the SM is not a complete theory of nature. Furthermore, no direct evidence for particles beyond the standard model (BSM) has yet been observed. This lack of guidance has led to the development of a wide range of new physics scenarios, many of which involve particles which are light and very weakly coupled to the SM. The absence of an indication on the mass scale or coupling strength of such new physics calls for a versatile and broad exploration tool.

In this context, the natural successor to an exploratory machine such as the LHC is a high-precision machine capable of delivering large data samples at energy scales corresponding to the mass of the Higgs boson, electroweak gauge bosons, and other key processes. Such a machine would be able to uncover indirect signs of BSM physics through precise measurements that could reveal deviations from SM predictions arising in quantum loop effects.

6.1 The Future Circular Collider

6.1.1 Project

The Future Circular Collider (FCC) project suggests a 91 km-circumference circular collider at CERN. In its first phase, the FCC would operate as an electron-positron collider (FCC-ee), delivering collisions at centre-of-mass energies from 91 GeV (corresponding to

the mass of the Z boson) to a maximum of 365 GeV (just above the top-quark pair production threshold). In a second phase, the same tunnel is planned to host a proton-proton collider (FCC-hh) capable of reaching energies of at least 100 TeV, significantly extending the energy frontier.

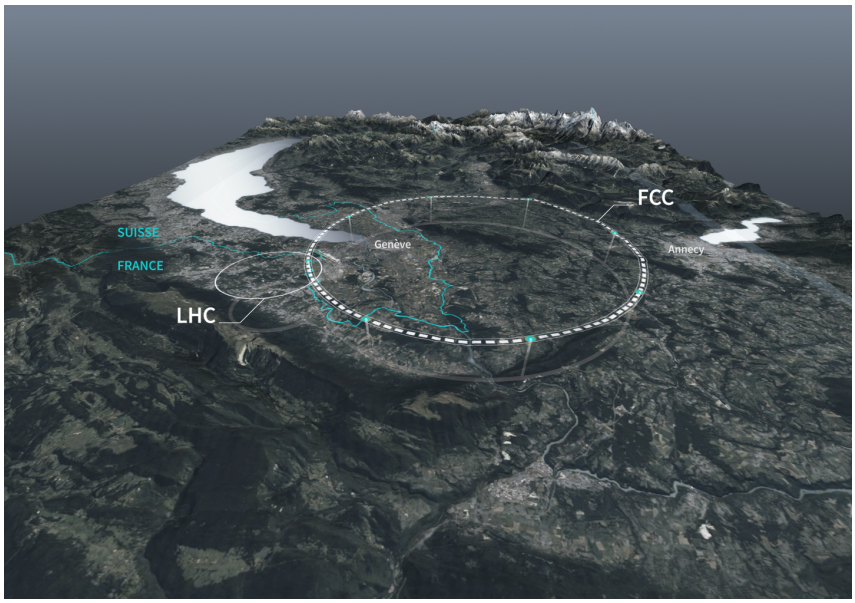


Figure 6.1: Map showing the preferred placement for the FCC. Image from [138].

The FCC-ee will enable the collection of extremely large datasets of Higgs bosons, electroweak gauge bosons, and top quarks. This will support a broad and ambitious programme of precision measurements across a wide range of physics sectors, including electroweak interactions, quantum chromodynamics (QCD), flavour physics, tau lepton properties, Higgs boson couplings, and top-quark observables. These measurements would be sensitive to tiny deviations from the predicted SM behaviour, and would probe energy scales far beyond the direct kinematic reach of the FCC-ee.

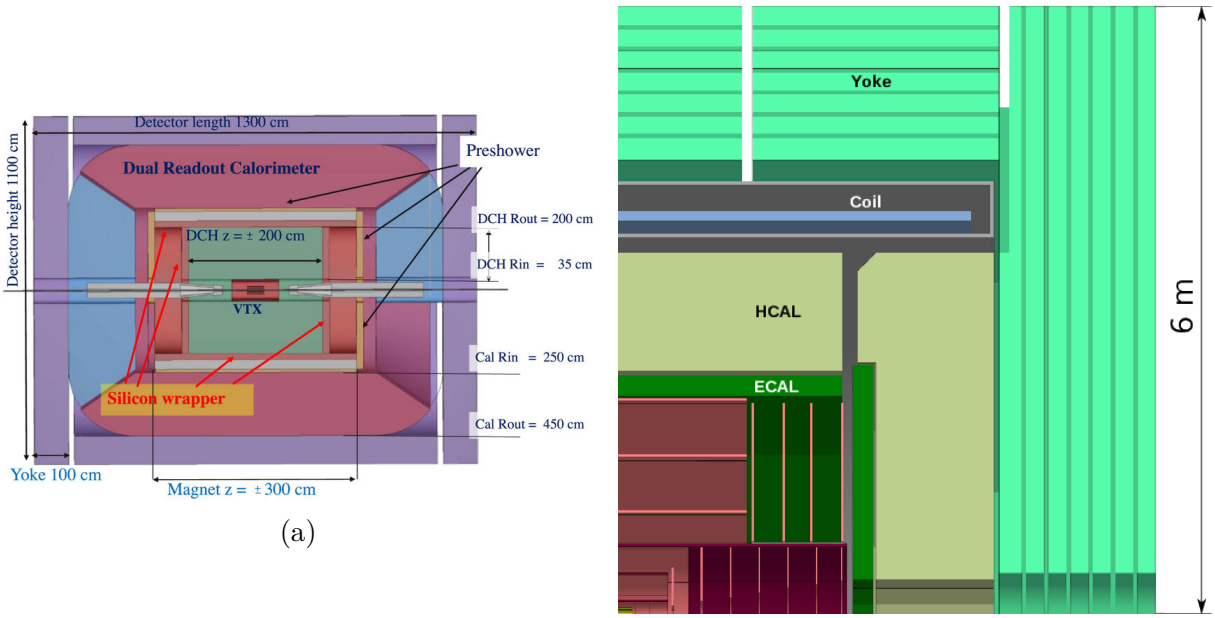
The technical design and financial viability of the FCC programme are currently being investigated as part of an extensive feasibility study coordinated under the framework of the European Strategy for Particle Physics [139–141].

6.1.2 Detector concepts

The FCC-ee is designed with four interaction points, each hosting a dedicated experiment. The FCC-ee experimental programme imposes stringent demands on detector performance [142], particularly in terms of vertex resolution, momentum and energy resolution, and flavour tagging capabilities. Several detector concepts have already been studied to meet the unique experimental requirements of the FCC-ee. In particular, two reference detector designs have been developed and studied in detail [142]: IDEA (International Detector for Electron-positron Accelerator) and CLD (CLIC-Like Detector). They are represented in Figure 6.2. Both designs build on experience from previous and proposed lepton collider experiments, such as ALEPH [143], ILD [144], and CLIC [145].

The IDEA detector [146] features a lightweight and highly granular tracking system based on a large-radius drift chamber immersed in a 2T magnetic field, complemented by a silicon vertex detector and a dual-readout calorimeter system, with crystal calorimetry for the ECAL and fibres for the HCAL. A muon system surrounds everything. This design prioritises excellent tracking resolution and minimal material budget to optimise performance for low-momentum particle reconstruction and precision measurements.

The CLD detector [147] has a more compact design derived from the CLIC detector, using a full silicon tracker combined with a high-granularity electromagnetic and hadronic calorimeter system optimised for particle flow reconstruction. Operating in a 4T magnetic field, it aims to achieve excellent jet energy resolution, crucial for precise Higgs and top-quark measurements.



(a) (b)

Figure 6.2: Concept for the IDEA (a) and CLD (b) detectors. For CLD, only the top quarter is shown. Taken from [142].

6.1.3 Working in an e^+e^- environment

6.1.3.1 Advantages

A major advantage of the e^+e^- environment is the exceptionally clean conditions it offers for precision measurements. In hadron colliders, many proton–proton interactions can occur within a single bunch crossing, leading to pile-up. At the FCC-ee, high instantaneous luminosity will instead be achieved through a very large collision rate ($\sim$ 33 MHz) using smaller bunches, allowing the interaction rates per bunch crossing to be greatly reduced. In particular, at the Z pole, 0.2% of the events are expected to contain more than one hard interaction. In addition, electron–positron collisions produce colour-singlet hadronic systems of well-defined quantum numbers and centre-of-mass energy. In other words, they are free from contamination from the underlying event and initial-state QCD radiation.

Altogether, this contributes to a low event rate and a clean experimental environment, enabling the FCC-ee to potentially operate in trigger-less mode and record all events, further simplifying the analysis.

Another critical advantage is the precisely known nature of the initial state. Since electrons and positrons are elementary particles, their momenta and energies are known with great accuracy from the beam parameters (energy and polarisation). This removes the need to account for PDFs, which introduce a major source of theoretical uncertainty in proton–proton collisions. Moreover, the well-defined initial state allows for powerful analysis techniques, such as the recoil mass method. In processes such as Higgsstrahlung ($e^+e^- \to ZH$), the mass of the Higgs boson can be determined from the measured kinematics of the recoiling Z boson, typically identified through its leptonic decay. This method is fully independent of the Higgs decay mode, making it valuable for model-independent measurement of the Higgs production cross section, for example.

The electron–positron collider design also allows for great flexibility in the centre-of-mass energy. The beam energy can be tuned with high precision to target specific physics processes, making it possible to perform threshold scans and dedicated campaigns for specific processes. These campaigns will be a key feature of the FCC-ee physics program.

In addition, the luminosity at FCC-ee can be measured with exceptional accuracy. It can be determined in situ using processes with very well-known cross sections, such as $e^+e^- \to \gamma\gamma$.

Finally, the radiation levels in the FCC-ee detector environment are significantly lower than those in hadron colliders. This results in reduced radiation damage to the detectors, particularly the innermost components, thereby enhancing detector longevity and reducing the need for frequent replacements or complex shielding strategies. In addition, it also improves the stability and performance of subdetectors over the lifetime of the experiment.

6.1.3.2 Challenges

To fully exploit the physics potential of the FCC-ee, several technical and experimental challenges must be addressed. A key challenge lies in reducing systematic uncertainties to match the level of statistical uncertainties, which will be considerably low thanks to the very high luminosity. Achieving this level of precision requires exceptional control over all aspects of the experiment, from detector design and calibration to theoretical inputs.

Furthermore, the physics event rate at the FCC-ee could reach up to 100 kHz at peak luminosity. This places strong demands on the sub-detectors response and the data acquisition (DAQ) systems, which must be capable of handling high throughput without dead time or loss of efficiency. A large and well-defined detector acceptance is also required, especially in the inner layers and forward region, to be able to reach a very precise luminosity measurement [148].

Finally, the accelerator itself must be designed to meet tight beam energy and stability specifications, particularly for precision measurements that rely on threshold scans or beam energy calibration.

6.1.4 Planned schedule

The FCC-ee is planned to run for 15 years, targeting a total integrated luminosity delivered of 144×10^{34} cm$^{-2}s^{-1}$. This corresponds to 10^5 times the full LEP dataset in total.

The FCC-ee physics program is structured around several runs at various centre-of-mass energies, each targeting key sectors in the SM. Starting at the Z pole, the collider will produce an unprecedented number of events (6×10^{12} Z bosons are expected), enabling precise measurement of the Z properties, along with QCD effects, flavour observables, and rare or invisible decay. Afterwards, it will run at the WW threshold, producing approximately 240 Million W bosons, allowing for precise determination of the W mass and width. This large number of WW events will also serve for studies of electroweak couplings and neutrino physics. The collider will then focus on Higgs physics through the Higgsstrahlung, providing around 3 million Higgs bosons for precise coupling measurements. Finally, a run above the top-pair production threshold will allow the measurement of top quark electroweak couplings and indirect sensitivity to the Higgs self-coupling. This broad physics program can be complemented with other runs dedicated to low-energies QCD, electron Yukawa or Higgs mass studies.

This running strategy is illustrated in Figure 6.3, which summarises the target energies, integrated luminosities, and key physics goals of each phase of the FCC-ee programme.

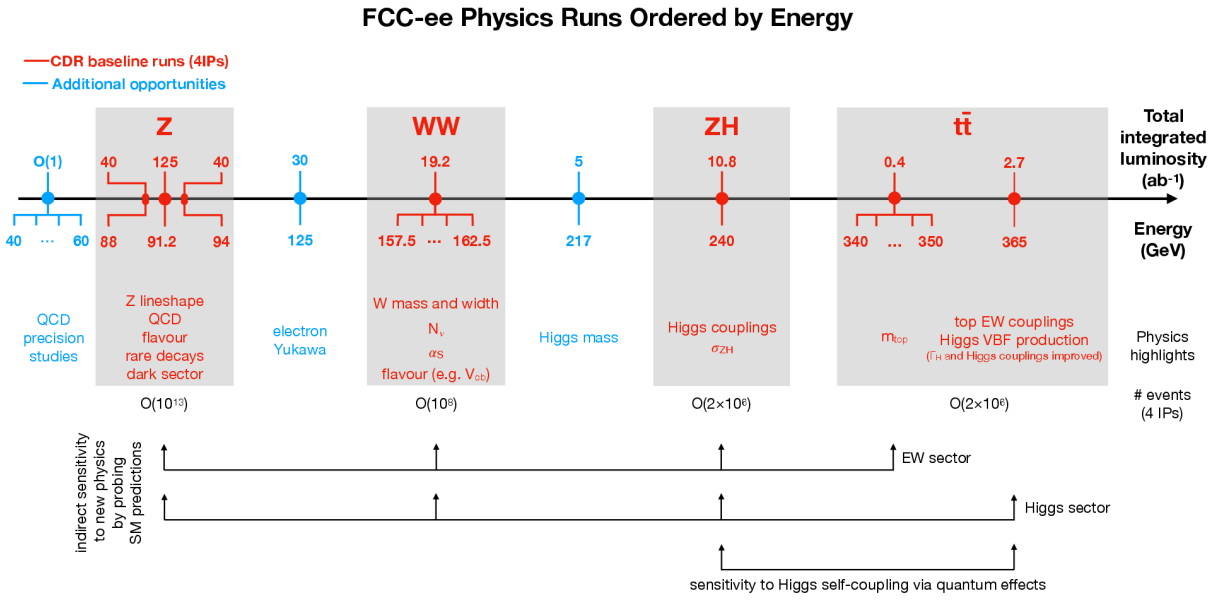


Figure 6.3: Taken from [149].

6.2 Sensitivity to α_s at the FCC-ee

6.2.1 Motivations

Theoretical predictions for many SM processes rely heavily on QCD calculations, which are expressed as a power series in the strong coupling constant. The current world average value of α_s at the Z boson mass is $\alpha_s(m_Z^2) = 0.1180 \pm 0.0009$, corresponding to an uncertainty of approximately 1% [11]. This makes α_s one of the least precisely known fundamental parameters of the SM. This uncertainty propagates into theoretical predictions for a wide range of processes, limiting the precision of SM tests and reducing sensitivity to potential new physics. Reducing the uncertainty in α_s is therefore necessary to fully exploit the capabilities of future colliders.

In parallel, the FCC-ee provides an exceptionally clean collision environment, ideal for precision QCD studies. In contrast to hadron colliders, where QCD backgrounds are ubiquitous, electron–positron collisions produce colour-singlet hadronic systems, resulting in clean dijet and multijet events with minimal contamination from underlying processes. This will contribute to a significant reduction in experimental systematic uncertainties.

At all energy points considered, a very large number of $q\bar{q}$ events will be produced from the decay of a Z boson or a photon. The large data samples expected at the FCC-ee will enable high-statistics measurements of key observables that will enable the extraction of α_s. For example, the hadronic decays of the Z and W bosons will allow for determinations of α_s with permil-level precision. Similarly, studies of hadronic decays of the τ lepton will provide sensitivity to α_s at very low energy scales. In addition, measurements of jet multiplicities and jet substructure will offer complementary probes of QCD dynamics. Together, these studies have the potential to significantly improve the current precision on α_s and its running as a function of energy.

In this section, two observables that are sensitive to α_s are studied in the context of FCC-ee, using MC simulations. The first one is the three-to-two jet cross-section ratio $R_{3/2}$, and the second one is the Lund Jet Plane, thoroughly described in Chapter 5.

6.2.2 Dataset and event selection

The analyses use 10 million events from $e^+e^- \to Z \to q\bar{q}$ process generated at next-to-leading order with MADGRAPH [31] at parton level, at $\sqrt{s} = 91$ GeV. The quarks are required to be light quarks (up or down). These events are further processed with PYTHIA 8 [21] for fragmentation and hadronisation, and with DELPHES [150] for fast simulation of the response of the detector. In this analysis, the IDEA detector concept is used.

Jets are clustered using the e^+e^- generalised k_t algorithm [46] with a radius R, where $R = 0.5$ for the LJP analysis and several $R \in [0.5, 0.7, 0.8, 1.0]$ for the $R_{3/2}$ analysis. Several other jet algorithms have been investigated for the LJP analysis. The constituents are required to be within the $[0.3, \pi - 0.3]$ θ-region to match the angular acceptance close to the beam. Jets are further required to have an energy above $E_{\text{min}} = 5$ GeV for the LJP, $E_{\text{min}} \in [5, 10, 15, 20, 25]$ GeV for the $R_{3/2}$ study. In both cases, jets are required to be in the angular acceptance (jet-θ within $[0.3 + R, \pi - 0.3 - R]$).

Five datasets are created by setting using five different α_s-values in MADGRAPH: 0.110, 0.115, 0.118, 0.120 and 0.125. At parton level, the gluon multiplicity is expected to increase with α_s. This α_s scan enables probing the sensitivity of those variables to α_s.

6.2.3 $R_{3/2}$

6.2.3.1 Definition

The observable $R_{3/2}$ is defined as the ratio of the cross sections for events with at least 3 jets to those with at least 2 jets:

$$R_{3/2} = \frac{\sigma_{\geq 3j}}{\sigma_{\geq 2j}}.$$

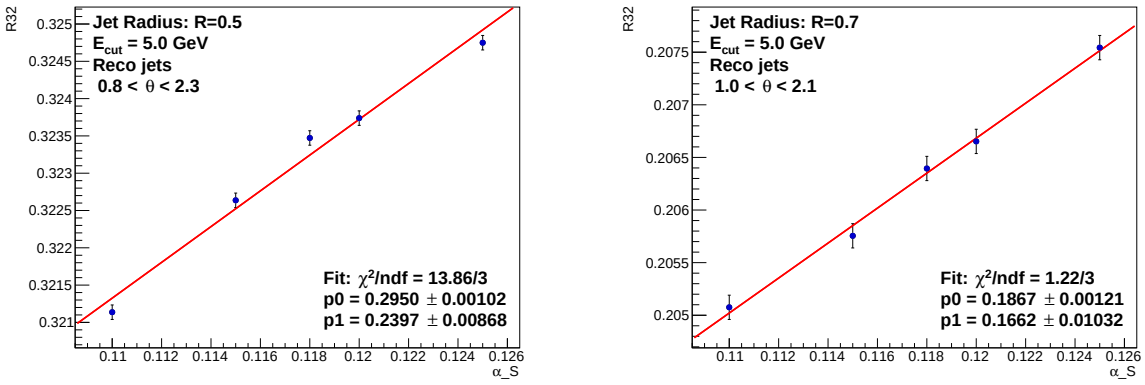
Ratios of inclusive cross-sections are a preferred method for extracting the value of the strong coupling constant because they retain sensitivity to α_s while minimising effects related to non-perturbative QCD and PDFs. At leading order, $R_{3/2}$ is directly proportional to α_s since, at the matrix-element level, the emission of an additional parton introduces an extra power of α_s. Therefore, this observable is a powerful tool for precision QCD studies at FCC-ee.

6.2.3.2 Dependence on jet algorithm variables

The dependence of $R_{3/2}$ on α_s is studied as a function of several variables: R, the radius of the jet, and E_{min}, the minimum jet energy. Both parameters influence the reconstructed jet multiplicity. A larger jet radius can cause two nearby small-R jets to be merged into a single large-R jet, while a lower energy threshold allows more low-energy jets to be included in the event selection. The result of $R_{3/2}$ as a function of α_s is then fitted with a linear function, with p_1 the slope reflecting the sensitivity of $R_{3/2}$ to α_s, and p_0 the intercept. The R-scan is presented in Figure 6.4, and is performed with an energy cut $E_{\text{min}} = 5$ GeV. The sensitivity to α_s decreases with the jet radius. This effect is primarily due to acceptance losses: larger jets are more likely to be rejected if they extend too close to the beam pipe. Additionally, larger jets tend to capture more surrounding radiation, causing nearby small jets to merge into a single large jet. As a result, the overall number of events with three or more jets is reduced, which in turn lowers the sensitivity to α_s.

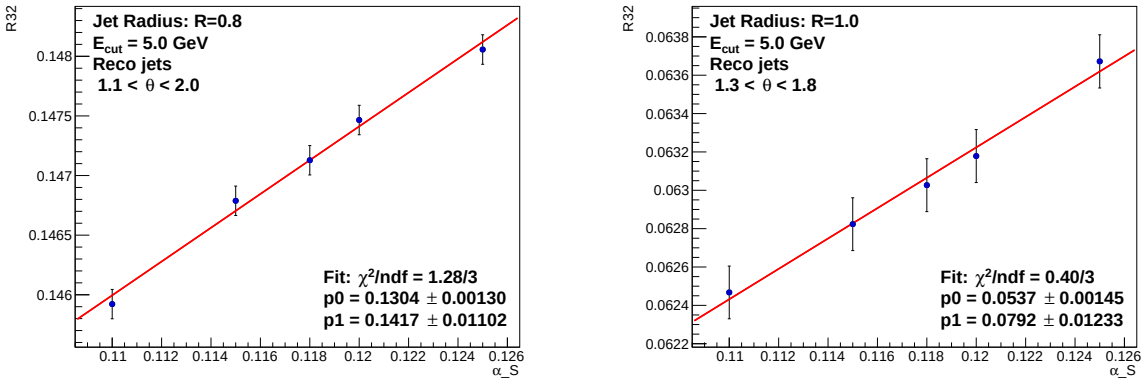
The E-scan is presented in Figure 6.5, and is done with jets of radius $R = 0.5$. The sensitivity to α_s decreases with the jet energy threshold. In this case, the limiting factor is phase space acceptance: higher energy thresholds significantly reduce the available phase space for the production of a third jet, thereby lowering the number of multi-jet events and diminishing the observable's sensitivity to α_s.

The sensitivity to α_s is maximal for $R = 0.5$ and $E_{\text{min}} = 5$ GeV. For $\alpha_s = 0.118$, the relative statistical uncertainty in the $R_{3/2}$ value is of the order of $\Delta R_{3/2} = 0.03\%$. For these parameters, and for a sample of $N^{\text{simulated}}_{Z\to u\bar{u},d\bar{d}} = 10^7$ events, a 1% change in α_s results in a $(0.240 \pm 0.009)\%$ change in the observable $R_{3/2}$.



(a) $R = 0.5$

(b) $R = 0.7$



(c) $R = 0.8$

(d) $R = 1.0$

Figure 6.4: Dependency of $R_{3/2}$ on α_s for jets clustered with various R values. The error bars represent the statistical uncertainty inly. A linear fit of the data, together with its parameters and χ^2/NDF, is also displayed.

6.2.3.3 Perspective at FCC-ee

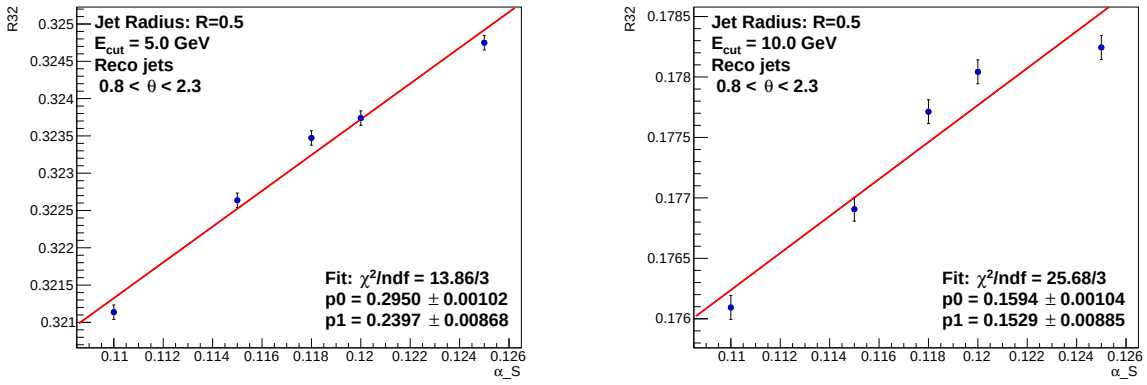
The number of Z bosons expected at FCC-ee at the Z pole is of the order of 10^{13}. Among them, only a fraction decays into light quarks ($u\bar{u}$ or $d\bar{d}$):

$$N^{\text{FCC}}_{Z\to u\bar{u},d\bar{d}} = N^{\text{FCC}}_{Z} \times \frac{\Gamma_{u\bar{u},d\bar{d}}}{\Gamma_{\text{Z}}},$$

where $\frac{\Gamma_{u\bar{u},d\bar{d}}}{\Gamma_{\text{Z}}}$ denotes the partial width of $Z \to u\bar{u}, d\bar{d}$. This partial width can be estimated from the sum of the partial widths of $Z \to (u\bar{u} + c\bar{c})/2$ and $Z \to (d\bar{d} + s\bar{s} + b\bar{b})/3$, using values reported by the Particle Data Group [11]. This yields an approximate fraction of $\frac{\Gamma_{u\bar{u},d\bar{d}}}{\Gamma_{\text{Z}}} \sim 27\%$, corresponding to

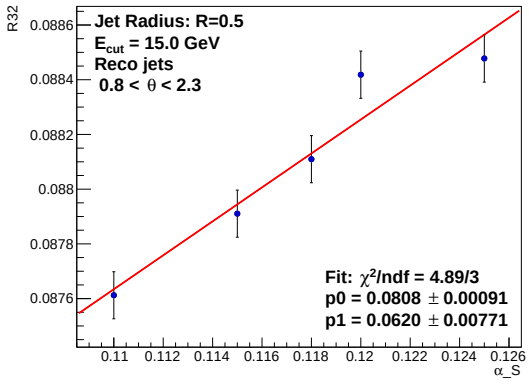
$$N^{\text{FCC}}_{Z\to u\bar{u},d\bar{d}} \sim 3 \times 10^{12}.$$

As a result, the relative statistical uncertainties in the $R_{3/2}$ observable would be reduced by a factor of $\sqrt{N^{\text{FCC}}_{Z\to u\bar{u},d\bar{d}}/N^{\text{simulated}}_{Z\to u\bar{u},d\bar{d}}} \sim 550$ with respect to this study, corresponding to

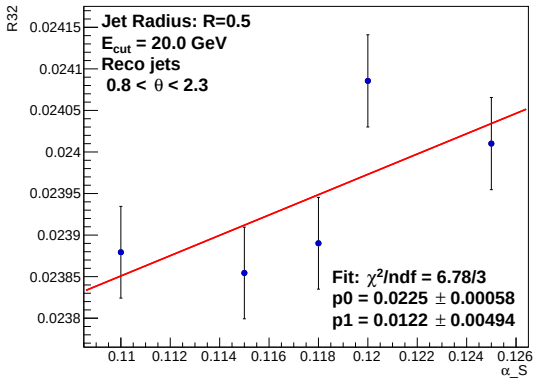


(a) $E_{\text{min}} = 5$ GeV

(b) $E_{\text{min}} = 10$ GeV



(c) $E_{\text{min}} = 15$ GeV



(d) $E_{\text{min}} = 20$ GeV

Figure 6.5: Dependency of $R_{3/2}$ on α_s for jets clustered with various E_{min} values. The error bars represent the statistical uncertainty inly. A linear fit of the data, together with its parameters and χ^2/NDF, is also displayed.

$\Delta R_{3/2\text{FCC}} = 0.00005\%$. The statistical uncertainty in the evaluation of α_s at FCC-ee using the $R_{3/2}$ variable would therefore be $\Delta\alpha_s = 0.0002\%$.

The 4% uncertainty in p_1, the slope of the function quantifying the sensitivity of $R_{3/2}$ to α_s, directly impacts the determination of α_s. To ensure that this uncertainty becomes comparable in magnitude to the statistical uncertainty in data, previously determined as $\Delta\alpha_s = 0.0002\%$, the statistical uncertainty in $R_{3/2}$ from MC should be reduced accordingly. Since the uncertainty in p_1 scales as $1/\sqrt{N}$, reducing this uncertainty by a factor of $20,000$ requires increasing the MC statistics by a factor of $400,000,000$ with respect to the current study, reaching a total number of events of the order of 10^{15}. The same reasoning with a more standard target of 0.1% yields a total number of MC events needed of the order of 10^{11}. This is an approximation, as in practice, MC predictions include loop-induced cancellations and non-perturbative effects, which will require producing more statistics to achieve the desired level of statistical uncertainties.

6.2.3.4 Discussions

The strong coupling constant was extracted at NNLO in QCD using $R_{3/2}$ from ALEPH data at LEP [151]. The analysis quotes a total experimental uncertainty of ± 0.0020 ($\sim 1.7\%$), of which the statistical component is only ± 0.0004 ($\sim 0.3\%$), with the remaining $\sim 1.4\%$ arising from detector, hadronisation, and quark-mass systematics. The theory (scale) uncertainty is ± 0.0015 ($\sim 1.3\%$) at NNLO. To fully exploit FCC-ee's exquisite statistical precision, it will be crucial to drastically reduce experimental systematics (via detector design and calibration) and hadronisation uncertainties (through modern tuning of MC models), and to have N^{3}LO QCD predictions to further shrink theory scale errors.

6.2.4 Lund Jet Plane

As presented in Section 5.1.3, the LJP organises the emissions within a jet in terms of their transverse momentum and angular separation. This approach not only distinguishes between perturbative and non-perturbative QCD regimes but also allows the extraction of α_s from the density of the LJP:

$$\rho(k_T, \Delta R) = \frac{1}{N_{jets}} \frac{d^2(\#\text{emissions})}{d\ \ln(k_T/\text{GeV})\ d\ \ln(R/\Delta R)} \approx \frac{2C_F \alpha_s(k_T)}{\pi}.$$

In this context, k_t and ΔR represent, respectively, the transverse momentum and angular separation of a soft emission relative to a hard emission, during a jet declustering step. C_F denotes the colour factor.

Several jet algorithms have been investigated for the jet finding step, such as anti-k_T, k_T, Cambridge-Aachen and Valencia. A modified version of the Cambridge-Aachen [44] algorithm is used for declustering the emissions from jets in angular order to construct the LJP. In this version, ΔR is parametrised as a function of θ rather than y, to reflect the e^+e^- environment where no direction is preferred. In this study, large-R jets are constructed ($R = 1.5$) to ensure that the LJP captures the characteristics of additional gluon radiations in the initial state, which is enhanced at higher values of α_s.

The leading jet of all events that contains at least two jets is used to construct the primary LJP. The primary LJP for a value of $\alpha_s(M_Z) = 0.118$ is given in Figure 6.6. It exhibits the same structure as the ATLAS LJP shown in Chapter 5, with a plateau at high k_T and a rapid increase at low k_T, which is consistent with the expected dependence of the emission density on α_s.

The projection along the X and Y axis of the LJP is shown for various values of $\alpha_s(M_Z)$ in Figure 6.7. Sensitivity to $\alpha_s(M_Z)$ in the shape of the LJP is hinted at, with the number of hard emissions increasing with $\alpha_s(M_Z)$. This precision is expected to improve with the large statistics available at the Z pole of the FCC-ee. However, fully exploiting this statistical power places stringent demands on tracking and reconstruction algorithms, which should be further studied using a full simulation of the IDEA detector. It also requires improvement in theoretical modelling, which has historically struggled to describe the structure of the LJP accurately. Furthermore, this study does not account for the dependence of the LJP on the running of α_s, as only variations in $\alpha_s(M_Z)$ for the hard process are considered.

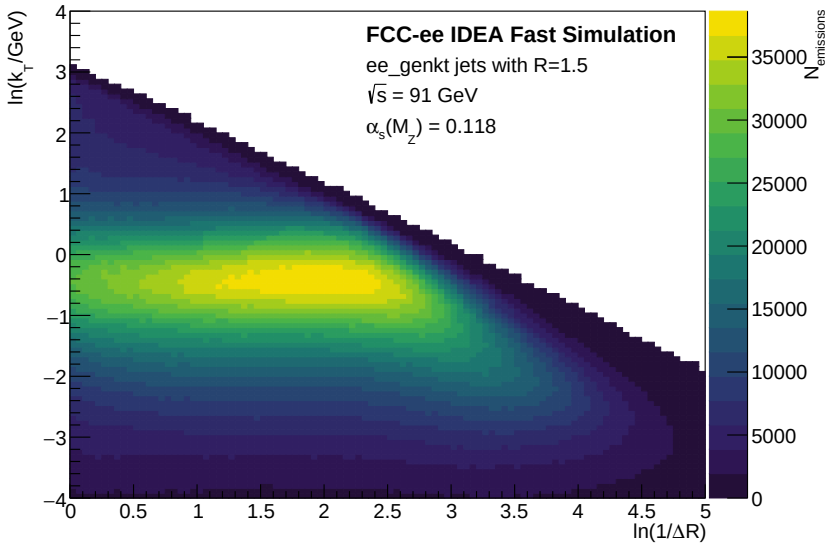
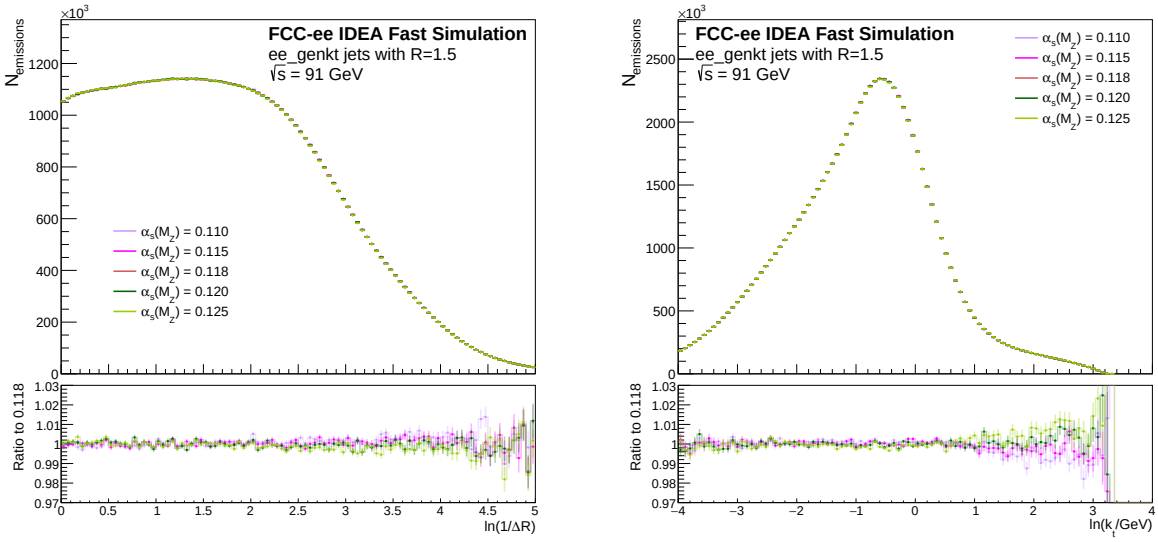


Figure 6.6: Primary Lund Jet Plane, for a value of $\alpha_s(M_Z) = 0.118$.



(a) (b)

Figure 6.7: Projection of the Primary Lund Jet Plane along the ln(R/ΔR) axis (a) and ln(k_T) axis, for various values of $\alpha_s(M_Z)$

6.3 Conclusion and perspective

Jet studies at FCC-ee offer a unique and fertile ground for advancing our understanding of QCD. The clean and well-controlled experimental environment at FCC-ee enables unprecedented precision in probing the mechanisms of jet production and evolution. However, this clean collision environment and large expected datasets bring new challenges. Among them is the need for exceptionally tight control of systematic uncertainties, which place stringent demands on detector acceptance, construction quality, and long-term sta-

bility. Jet-related measurements, in particular, impose strong requirements on the energy and angular resolution of the calorimeters, as well as on the granularity of the tracking detectors. While a full GEANT-based simulation of FCC-ee events has only recently become available, a state-of-the-art reconstruction framework based on the particle flow approach is still under development. Given that the reconstruction algorithm plays a crucial role in determining the resolution of hadronic observables, it remains difficult to set definitive performance benchmarks for the hadronic calorimeter based on current studies.

Conclusion

This thesis focuses on testing the QCD sector of the SM using jets collected in the ATLAS detector. Jets provide access to the dynamics of quarks and gluons at the smallest distance scales accessible in the laboratory, and they are among the most prominent observables at the LHC. Their production and evolution are directly governed by α_s, the strong coupling constant of QCD. High-precision measurements of jet-related observables enable stringent tests of QCD and offer sensitivity to potential signs of physics beyond the Standard Model. The SM and QCD have been described in Chapter 1, while Chapter 2 described the ATLAS experiment.

To the required precision, jets must be accurately calibrated, correcting for detector effects to ensure reliable comparisons between data and theoretical predictions. In particular, jet calibration must be continuously derived, monitored and improved. Jet calibration is the focus of Chapter 3, with a focus on the jet η-intercalibration, one of the in situ steps of jet calibration. The JES procedure is thoroughly described, and several studies aimed at improving the reliability of the jet η-intercalibration are presented. In particular, the overconstrained system used to derive the calibration factors has been strengthened by incorporating additional terms and introducing a compatibility test between them. Furthermore, a new, more robust method for estimating the statistical uncertainties on the calibration factors has been implemented. This approach results in larger, more realistic uncertainties and enables the study of their correlations among different phase-space regions and with other measurements. Finally, the η-intercalibration has been derived for direct application in data analyses using Run 2 and Run 3 data. This required the development of a new trigger strategy for the application of the calibration to Run 3 data.

Chapter 4 presents the inclusive jet production cross-section analysis, based on the full Run 2 dataset from proton–proton collisions at a centre-of-mass energy of 13 TeV. This is a double-differential measurement, counting every jet above a given p_T threshold in bins of p_T and $|y|$. Combined with recent advances in QCD calculations at NNLO, this measurement provides a stringent test of QCD at very high momentum transfer and has the potential to contribute to future global PDF fits. The chapter details the analysis strategy, including comparisons with theoretical predictions and a thorough description of the unfolding procedure. Unfolding statistically corrects for detector effects in order to recover the true particle-level distributions, enabling accurate comparisons with the-

ory. There are several method for the unfolding, this chapter focuses on the Iterative Dynamically Stabilized (IDS) method It is based on Bayes' theorem, iteratively updating the unfolded distribution by comparing data to simulation, while applying regularisation to minimise the impact of statistical fluctuations and stabilise the result. The evaluation of unfolding bias, arising from shape differences between data and simulation, is also presented. In addition, the chapter discusses the propagation of systematic uncertainties through the unfolding process, with particular emphasis on the corrected treatment of the statistical uncertainty on the η-intercalibration factors, which I previously derived. This represents a significant contribution, given the known sensitivity of the measurement to the modelling of these factors. Overall, this analysis is expected to enable a highly competitive determination of the strong coupling constant α_s.

Chapter 5 deals with the Lund Jet Plane, a jet substructure observable that maps the phase space of emissions inside a jet in terms of their momentum fraction and opening angle. This observable factorises the effects of QCD along the plane, and is sensitive to α_s in its normalisation and density. The chapter details the analysis strategy, and explains how the unfolding bias is evaluated and used to optimise the parameters of the IDS unfolding for this analysis, namely the number of unfolding iterations. The propagation of uncertainties is detailed, and prospects for theoretical interpretation are discussed.

In Chapter 6, simulated data from the IDEA detector concept of FCC-ee is used to study jet in an e^+e^- environment. The well-controlled experimental environment, as well as the large expected datasets, enables unprecedented precision in probing the mechanisms of jet production and evolution. Two jet variables are studied: $R_{3/2}$ the ratio of cross sections for events with at least three jets to those with at least two jets, and the Lund Jet Plane. The jet algorithm parameters are optimised to maximise the sensitivity of those variables to α_s. The sensitivity of $R_{3/2}$ to α_s at FCC-ee is evaluated, considering statistical uncertainties only. The LJP is produced and studied.

For the ATLAS analyses, the inclusive jet cross-section measurement is completed, and the LJP measurement is nearing completion. For the FCC-ee studies, prospects for these measurements can be explored using different detector parameter configurations. Such studies would offer valuable insight into detector design, aiming to maximise the physics potential for jet-related observables.

Bibliography

[1] E. Noether, "Invariante Variationsprobleme," *Nachrichten von der Gesellschaft der Wissenschaften zu Göttingen, Mathematisch-Physikalische Klasse*, vol. 1918, pp. 235–257, 1918.

[2] J. Goldstone, A. Salam, and S. Weinberg, "Broken Symmetries," *Phys. Rev.*, vol. 127, pp. 965–970, Aug 1962.

[3] G. S. Guralnik, C. R. Hagen, and T. W. B. Kibble, "Global Conservation Laws and Massless Particles," *Phys. Rev. Lett.*, vol. 13, pp. 585–587, Nov 1964.

[4] F. Englert and R. Brout, "Broken Symmetry and the Mass of Gauge Vector Mesons," *Phys. Rev. Lett.*, vol. 13, pp. 321–323, Aug 1964.

[5] P. W. Higgs, "Broken Symmetries and the Masses of Gauge Bosons," *Phys. Rev. Lett.*, vol. 13, pp. 508–509, Oct 1964.

[6] A. Purcell, "Go on a particle quest at the first CERN webfest. Le premier webfest du CERN se lance à la conquête des particules," no. 35/2012, p. 10, 2012.

[7] C. N. Yang and R. L. Mills, "Conservation of Isotopic Spin and Isotopic Gauge Invariance," *Phys. Rev.*, vol. 96, pp. 191–195, Oct 1954.

[8] Gross *et al.*, "50 Years of quantum chromodynamics: Introduction and Review," *The European Physical Journal C*, vol. 83, Dec. 2023.

[9] D. J. Gross and F. Wilczek, "Ultraviolet Behavior of Non-Abelian Gauge Theories," *Phys. Rev. Lett.*, vol. 30, pp. 1343–1346, Jun 1973.

[10] H. D. Politzer, "Reliable Perturbative Results for Strong Interactions?," *Phys. Rev. Lett.*, vol. 30, pp. 1346–1349, Jun 1973.

[11] S. Navas *et al.*, "Review of particle physics," *Phys. Rev. D*, vol. 110, no. 3, p. 030001, 2024.

[12] F. Bloch and A. Nordsieck, "Note on the Radiation Field of the Electron," *Phys. Rev.*, vol. 52, pp. 54–59, Jul 1937.

[13] T. Kinoshita, "Mass Singularities of Feynman Amplitudes," *Journal of Mathematical Physics*, vol. 3, pp. 650–677, 07 1962.

[14] T. D. Lee and M. Nauenberg, "Degenerate systems and mass singularities," *Phys. Rev.*, vol. B133, pp. 1549–62, 1964.

[15] 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.

[16] V. N. Gribov and L. N. Lipatov, "Deep inelastic e p scattering in perturbation theory," *Sov. J. Nucl. Phys.*, vol. 15, pp. 438–450, 1972.

[17] G. Altarelli and G. Parisi, "Asymptotic Freedom in Parton Language," *Nucl. Phys. B*, vol. 126, pp. 298–318, 1977.

[18] B. Andersson, G. Gustafson, G. Ingelman, and T. Sjostrand, "Parton Fragmentation and String Dynamics," *Phys. Rept.*, vol. 97, pp. 31–145, 1983.

[19] B. R. Webber, "A QCD Model for Jet Fragmentation Including Soft Gluon Interference," *Nucl. Phys. B*, vol. 238, pp. 492–528, 1984.

[20] B. Webber, "Parton shower Monte Carlo event generators," *Scholarpedia*, vol. 6, no. 12, p. 10662, 2011. revision #128236.

[21] C. Bierlich *et al.*, "A comprehensive guide to the physics and usage of PYTHIA 8.3," 2022.

[22] T. Sjöstrand, S. Mrenna, and P. Skands, "A brief introduction to PYTHIA 8.1," *Computer Physics Communications*, vol. 178, p. 852–867, June 2008.

[23] T. Sjöstrand, S. Ask, J. R. Christiansen, R. Corke, N. Desai, P. Ilten, S. Mrenna, S. Prestel, C. O. Rasmussen, and P. Z. Skands, "An introduction to PYTHIA 8.2," *Computer Physics Communications*, vol. 191, p. 159–177, June 2015.

[24] E. Bothmann, G. Singh Chahal, S. Höche, J. Krause, F. Krauss, S. Kuttimalai, S. Liebschner, D. Napoletano, M. Schönherr, H. Schulz, S. Schumann, and F. Siegert, "Event generation with Sherpa 2.2," *SciPost Physics*, vol. 7, Sept. 2019.

[25] M. Bähr *et al.*, "Herwig++ physics and manual," *Eur. Phys. J. C*, vol. 58, p. 639, 2008.

[26] J. Bellm *et al.*, "Herwig 7.0/Herwig++ 3.0 release note," *Eur. Phys. J. C*, vol. 76, no. 4, p. 196, 2016.

[27] J. Bellm *et al.*, "Herwig 7.1 Release Note," 2017.

[28] P. Nason, "A new method for combining NLO QCD with shower Monte Carlo algorithms," *JHEP*, vol. 11, p. 040, 2004.

[29] S. Frixione, P. Nason, and C. Oleari, "Matching NLO QCD computations with parton shower simulations: the POWHEG method," *JHEP*, vol. 11, p. 070, 2007.

[30] S. Alioli, P. Nason, C. Oleari, and E. Re, "A general framework for implementing NLO calculations in shower Monte Carlo programs: the POWHEG BOX," *JHEP*, vol. 06, p. 043, 2010.

[31] J. Alwall, R. Frederix, S. Frixione, V. Hirschi, F. Maltoni, O. Mattelaer, H. S. Shao, T. Stelzer, P. Torrielli, and M. Zaro, "The automated computation of tree-level and next-to-leading order differential cross sections, and their matching to parton shower simulations," *JHEP*, vol. 07, p. 079, 2014.

[32] Z. Nagy, "Next-to-leading order calculation of three-jet observables in hadron-hadron collisions," *Physical Review D*, vol. 68, Nov. 2003.

[33] Z. Nagy, "Three-Jet Cross Sections in Hadron-Hadron Collisions at Next-To-Leading Order," *Physical Review Letters*, vol. 88, Mar. 2002.

[34] Z. Nagy, "Next-to-leading order calculation of three-jet observables in hadron-hadron collisions," *Physical Review D*, vol. 68, Nov. 2003.

[35] A. Gehrmann-De Ridder, T. Gehrmann, E. Glover, A. Huss, and T. Morgan, "Precise QCD Predictions for the Production of a Z Boson in Association with a Hadronic Jet," *Physical Review Letters*, vol. 117, July 2016.

[36] J. Currie, E. Glover, T. Gehrmann, A. Gehrmann-de Ridder, A. Huss, and J. Pires, "Single Jet Inclusive Production for the Individual Jet p_{T} Scale Choice at the LHC," *Acta Physica Polonica B*, vol. 48, no. 6, p. 955, 2017.

[37] D. Britzger, A. G.-D. Ridder, T. Gehrmann, E. W. N. Glover, C. Gwenlan, A. Huss, J. Pires, K. Rabbertz, D. Savoiu, M. R. Sutton, and J. Stark, "Nnlo interpolation grids for jet production at the lhc," 2022.

[38] T. Carli, G. P. Salam, and F. Siegert, "A posteriori inclusion of PDFs in NLO QCD final-state calculations," 2005.

[39] S. D. Ellis and D. E. Soper, "Successive combination jet algorithm for hadron collisions," *Physical Review D*, vol. 48, p. 3160–3166, Oct. 1993.

[40] S. Catani, Y. Dokshitzer, M. Seymour, and B. Webber, "Longitudinally-invariant k-clustering algorithms for hadron-hadron collisions," *Nuclear Physics B*, vol. 406, no. 1, pp. 187–224, 1993.

[41] Y. Dokshitzer, G. Leder, S. Moretti, and B. Webber, "Better jet clustering algorithms," *Journal of High Energy Physics*, vol. 1997, p. 001–001, Aug. 1997.

[42] M. Wobisch and T. Wengler, "Hadronization Corrections to Jet Cross Sections in Deep-Inelastic Scattering," 1999.

[43] M. Cacciari, G. P. Salam, and G. Soyez, "The anti-k_t jet clustering algorithm," *JHEP*, vol. 04, p. 063, 2008.

[44] G. P. Salam, "Towards jetography," *The European Physical Journal C*, vol. 67, p. 637–686, May 2010.

[45] S. Catani, Y. L. Dokshitzer, M. Olsson, G. Turnock, and B. R. Webber, "New clustering algorithm for multi - jet cross-sections in e+ e- annihilation," *Phys. Lett. B*, vol. 269, pp. 432–438, 1991.

[46] M. Cacciari, G. P. Salam, and G. Soyez, "FastJet user manual: (for version 3.0.2)," *The European Physical Journal C*, vol. 72, Mar. 2012.

[47] M. Boronat, J. Fuster, I. García, E. Ros, and M. Vos, "A robust jet reconstruction algorithm for high-energy lepton colliders," *Physics Letters B*, vol. 750, p. 95–99, Nov. 2015.

[48] "Standard Model Summary Plots June 2024," tech. rep., CERN, Geneva, 2024. All figures including auxiliary figures are available at https://atlas.web.cern.ch/Atlas/GROUPS/PHYSICS/PUBNOTES/ATL-PHYS-PUB-2024-011.

[49] J. Thaler and K. Van Tilburg, "Identifying boosted objects with N-subjettiness," *Journal of High Energy Physics*, vol. 2011, Mar. 2011.

[50] A. J. Larkoski, G. P. Salam, and J. Thaler, "Energy correlation functions for jet substructure," *Journal of High Energy Physics*, vol. 2013, June 2013.

[51] F. A. Dreyer, G. P. Salam, and G. Soyez, "The Lund jet plane," *Journal of High Energy Physics*, vol. 2018, Dec. 2018.

[52] L. Evans and P. Bryant, "LHC Machine," *Journal of Instrumentation*, vol. 3, aug 2008.

[53] "CERN Website." https://home.cern. Accessed: 2025-06-02.

[54] The ATLAS Collaboration, "The ATLAS Experiment at the CERN Large Hadron Collider," *Journal of Instrumentation*, vol. 3, p. S08003, aug 2008.

[55] The CMS Collaboration, "The CMS experiment at the CERN LHC," *Journal of Instrumentation*, vol. 3, p. S08004, aug 2008.

[56] The LHCb Collaboration, "The LHCb Detector at the LHC," *Journal of Instrumentation*, vol. 3, p. S08005, aug 2008.

[57] The ALICE Collaboration, "The ALICE experiment at the CERN LHC," *Journal of Instrumentation*, vol. 3, p. S08002, aug 2008.

[58] ATLAS Collaboration, "Luminosity determination in pp collisions at $\sqrt{s} = 13$ TeV using the ATLAS detector at the LHC," *The European Physical Journal C*, vol. 83, oct 2023.

[59] O. Aberle *et al.*, *High-Luminosity Large Hadron Collider (HL-LHC): Technical design report*. CERN Yellow Reports: Monographs, Geneva: CERN, 2020.

[60] "ATLAS Luminosity Public Results for Run 3." https://twiki.cern.ch/twiki/bin/view/AtlasPublic/LuminosityPublicResultsRun3. Accessed: 2025-06-03.

[61] J. Pequenao, "Computer generated image of the whole ATLAS detector." https://cds.cern.ch/record/1095924, 2008.

[62] I. Neutelings. https://tikz.net/. Accessed: 2025-06-04.

[63] G. Avoni *et al.*, "The new LUCID-2 detector for luminosity measurement and monitoring in ATLAS," *JINST*, vol. 13, p. P07017, 2018.

[64] *ATLAS inner detector: Technical Design Report, 1*. Technical design report. ATLAS, Geneva: CERN, 1997.

[65] S. Haywood, L. Rossi, R. Nickerson, and A. Romaniouk, *ATLAS inner detector: Technical Design Report, 2*. Technical design report. ATLAS, Geneva: CERN, 1997.

[66] J. Pequenao, "Computer generated image of the ATLAS inner detector." 2008.

[67] N. Wermes and G. Hallewel, *ATLAS pixel detector: Technical Design Report*. Technical design report. ATLAS, Geneva: CERN, 1998.

[68] ATLAS Collaboration, "ATLAS Insertable B-Layer Technical Design Report," tech. rep., 2010.

[69] *ATLAS calorimeter performance: Technical Design Report*. Technical design report. ATLAS, Geneva: CERN, 1996.

[70] J. Pequenao, "Computer Generated image of the ATLAS calorimeter." 2008.

[71] *ATLAS liquid-argon calorimeter: Technical Design Report*. Technical design report. ATLAS, Geneva: CERN, 1996.

[72] ATLAS Collaboration, "ATLAS Liquid Argon Calorimeter Phase-I Upgrade: Technical Design Report," tech. rep., 2013. Final version presented to December 2013 LHCC.

[73] *ATLAS tile calorimeter: Technical Design Report*. Technical design report. ATLAS, Geneva: CERN, 1996.

[74] *ATLAS muon spectrometer: Technical Design Report*. Technical design report. ATLAS, Geneva: CERN, 1997.

[75] J. Pequenao, "Computer generated image of the ATLAS Muons subsystem." 2008.

[76] ATLAS Collaboration, "Performance of the ATLAS trigger system in 2015," *Eur. Phys. J. C*, vol. 77, p. 317, 2017.

[77] *ATLAS level-1 trigger: Technical Design Report*. Technical design report. ATLAS, Geneva: CERN, 1998.

[78] "Approved Plots Data Acquisition." https://twiki.cern.ch/twiki/bin/view/AtlasPublic/ApprovedPlotsDAQ. Accessed: 2025-06-06.

[79] T. G. Cornelissen, N. Van Eldik, M. Elsing, W. Liebig, E. Moyse, N. Piacquadio, K. Prokofiev, A. Salzburger, and A. Wildauer, "Updates of the ATLAS Tracking Event Data Model (Release 13)," tech. rep., CERN, Geneva, 2007. All figures including auxiliary figures are available at https://atlas.web.cern.ch/Atlas/GROUPS/PHYSICS/PUBNOTES/ATL-SOFT-PUB-2007-003.

[80] G. Borissov, D. W. Casper, K. Grimm, S. Pagan Griso, L. E. Pedersen, K. Prokofiev, M. S. Rudolph, and A. M. Wharton, "ATLAS strategy for primary vertex reconstruction during Run-2 of the LHC," Tech. Rep. 7, CERN, Geneva, 2015.

[81] ATLAS Collaboration, "Electron reconstruction and identification in the ATLAS experiment using the 2015 and 2016 LHC proton–proton collision data at $\sqrt{s}$ = 13 TeV," *Eur. Phys. J. C*, vol. 79, p. 639, 2019.

[82] ATLAS Collaboration, "Electron and photon performance measurements with the ATLAS detector using the 2015–2017 LHC proton–proton collision data," *JINST*, vol. 14, p. P12006, 2019.

[83] ATLAS Collaboration, "Electron and photon reconstruction and performance in ATLAS using a dynamical, topological cell clustering-based approach." ATL-PHYS-PUB-2017-022, 2017.

[84] ATLAS Collaboration, "Topological cell clustering in the ATLAS calorimeters and its performance in LHC Run 1," *The European Physical Journal C*, vol. 77, July 2017.

[85] ATLAS Collaboration, "Improving topological cluster reconstruction using calorimeter cell timing in ATLAS," *Eur. Phys. J. C*, vol. 84, p. 455, 2024.

[86] ATLAS Collaboration, "Jet reconstruction and performance using particle flow with the ATLAS Detector," *The European Physical Journal C*, vol. 77, July 2017.

[87] "Tagging and suppression of pileup jets with the ATLAS detector," tech. rep., CERN, Geneva, 2014. All figures including auxiliary figures are available at https://atlas.web.cern.ch/Atlas/GROUPS/PHYSICS/CONFNOTES/ATLAS-CONF-2014-018.

[88] ATLAS Collaboration, "Monitoring and data quality assessment of the ATLAS liquid argon calorimeter," *Journal of Instrumentation*, vol. 9, p. P07024–P07024, July 2014.

[89] "Selection of jets produced in 13TeV proton-proton collisions with the ATLAS detector," tech. rep., CERN, Geneva, 2015. All figures including auxiliary figures are available at https://atlas.web.cern.ch/Atlas/GROUPS/PHYSICS/CONFNOTES/ATLAS-CONF-2015-029.

[90] ATLAS Collaboration, "The ATLAS Simulation Infrastructure," *The European Physical Journal C*, vol. 70, no. 3, 2010.

[91] ATLAS Collaboration, "The ATLAS Collaboration Software and Firmware." ATL-SOFT-PUB-2021-001, 2021.

[92] ATLAS Collaboration, "Software and computing for Run 3 of the ATLAS experiment at the LHC," *Eur. Phys. J. C*, vol. 85, p. 234, 2025.

[93] 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.

[94] M. Cacciari, G. P. Salam, and G. Soyez, "The catchment area of jets," *Journal of High Energy Physics*, vol. 2008, p. 005–005, Apr. 2008.

[95] ATLAS Collaboration, "Performance of pile-up mitigation techniques for jets in *pp* collisions at $\sqrt{s} = 8$TeV using the ATLAS detector," *The European Physical Journal C*, vol. 76, Oct. 2016.

[96] ATLAS Collaboration, "Jet energy measurement with the ATLAS detector in proton-proton collisions at $\sqrt{s} = 7$ TeV," *Eur. Phys. J. C*, vol. 73, no. 3, p. 2304, 2013.

[97] ATLAS Collaboration, "New techniques for jet calibration with the ATLAS detector," *The European Physical Journal C*, vol. 83, Aug. 2023.

[98] L. Ginabat, "Étalonnage des jets et mesures précises de section efficaces de production de jets avec les données de l'expérience ATLAS," 2023. Presented 25 Sep 2023.

[99] ATLAS Collaboration, "Muon reconstruction performance of the ATLAS detector in proton–proton collision data at $\sqrt{s}$=13TeV," *The European Physical Journal C*, vol. 76, May 2016.

[100] B. Abbott *et al.*, "Determination of the absolute jet energy scale in the DØ calorimeters," *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 424, p. 352–394, Mar. 1999.

[101] ATLAS Collaboration, "A precise measurement of the jet energy scale derived from single-particle measurements and in situ techniques in proton-proton collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector," 2024.

[102] ATLAS Collaboration, "Determination of jet calibration and energy resolution in proton-proton collisions at $\sqrt{s} = 8$ TeV using the ATLAS detector," *Eur. Phys. J. C*, vol. 80, no. 12, p. 1104, 2020.

[103] V. Lendermann, J. Haller, M. Herbst, K. Kruger, H.-C. Schultz-Coulon, and R. Stamen, "Combining Triggers in HEP Data Analysis," *Nucl. Instrum. Meth. A*, vol. 604, pp. 707–718, 2009.

[104] R. Hankache, "Jet calibration, cross-section measurements and New Physics searches with the ATLAS experiment within the Run 2 data," 2019. Presented 24 Sep 2019.

[105] K. H. Knuth, "Optimal Data-Based Binning for Histograms," 2013.

[106] F. a. James, "MINUIT: Function Minimization and Error Analysis Reference Manual," 1998. CERN Program Library Long Writeups.

[107] "Evaluating statistical uncertainties and correlations using the bootstrap method," tech. rep., CERN, Geneva, 2021. All figures including auxiliary figures are available at https://atlas.web.cern.ch/Atlas/GROUPS/PHYSICS/PUBNOTES/ATL-PHYS-PUB-2021-011.

[108] ATLAS Collaboration, "Measurement of dijet cross sections in *pp* collisions at 7 TeV centre-of-mass energy using the ATLAS detector," *JHEP*, vol. 05, p. 059, 2014.

[109] ATLAS Collaboration, "Measurement of inclusive jet and dijet cross-sections in proton-proton collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector," *Journal of High Energy Physics*, vol. 2018, May 2018.

[110] ATLAS Collaboration, " Measurement of the inclusive jet cross - sections in proton -proton collisions at $\sqrt{s} = 8\,\text{TeV}$ with the ATLAS detector," *JHEP*, vol. 09, p. 020, 2017.

[111] ATLAS Collaboration, "Measurement of inclusive jet and dijet production in pp collisions at $\sqrt{s} = 7$ TeV using the ATLAS detector," *Physical Review D*, vol. 86, July 2012.

[112] CMS Collaboration, "Measurement of the double-differential inclusive jet cross section in proton-proton collisions at $\sqrt{s} = 5.02$ TeV," *Journal of High Energy Physics*, vol. 2025, Jan. 2025.

[113] CMS Collaboration, "Measurement and QCD analysis of double-differential inclusive jet cross sections in proton-proton collisions at $\sqrt{s} = 13$ TeV," *JHEP*, vol. 02, p. 142, 2022. [Addendum: JHEP 12, 035 (2022)].

[114] T. Carli, "Jet cleaning in full run-2 data," tech. rep., CERN, Geneva, 2022.

[115] B. Malaescu, "An iterative, dynamically stabilized method of data unfolding," 2009.

[116] B. Malaescu, "An Iterative, Dynamically Stabilized (IDS) Method of Data Unfolding," 2011.

[117] Buckley, Andy and Ferrando, James and Lloyd, Stephen and Nordström, Karl and Page, Ben and Rüfenacht, Martin and Schönherr, Marek and Watt, Graeme, "LHAPDF6: parton density access in the LHC precision era," *The European Physical Journal C*, vol. 75, Mar. 2015.

[118] J. Currie, A. Gehrmann-De Ridder, T. Gehrmann, N. Glover, A. Huss, and J. Pires, "Infrared sensitivity of single jet inclusive production at hadron colliders," *Journal of High Energy Physics*, vol. 2018, Oct. 2018.

[119] S. Dittmaier, A. Huss, and C. Speckner, "Weak radiative corrections to dijet production at hadron colliders," *Journal of High Energy Physics*, vol. 2012, Nov. 2012.

[120] M. Yan, T.-J. Hou, P. Nadolsky, and C.-P. Yuan, "CT18 global PDF fit at leading order in QCD," *Physical Review D*, vol. 107, June 2023.

[121] A. Collaboration, "Determination of the parton distribution functions of the proton using diverse atlas data from pp collisions at

$$\sqrt{s} = 7$$

, 8 and 13 tev," *The European Physical Journal C*, vol. 82, May 2022.

[122] S. Bailey, T. Cridge, L. A. Harland-Lang, A. D. Martin, and R. S. Thorne, "Parton distributions from lhc, hera, tevatron and fixed target data: Msht20 pdfs," *The European Physical Journal C*, vol. 81, Apr. 2021.

[123] R. D. Ball, S. Carrazza, J. Cruz-Martinez, L. Del Debbio, S. Forte, T. Giani, S. Iranipour, Z. Kassabov, J. I. Latorre, E. R. Nocera, R. L. Pearson, J. Rojo, R. Stegeman, C. Schwan, M. Ubiali, C. Voisey, and M. Wilson, "The path to proton structure at 1

[124] d'Enterria *et al.*, "The strong coupling constant: state of the art and the decade ahead," *Journal of Physics G: Nuclear and Particle Physics*, vol. 51, p. 090501, Sept. 2024.

[125] ATLAS Collaboration, "Measurement of the Lund Jet Plane Using Charged Particles in 13 TeV Proton-Proton Collisions with the ATLAS Detector," *Physical Review Letters*, vol. 124, June 2020.

[126] ATLAS Collaboration, "Measurement of the Lund jet plane in hadronic decays of top quarks and W bosons with the ATLAS detector," *The European Physical Journal C*, vol. 85, Apr. 2025.

[127] CMS Collaboration, "Measurement of the primary Lund jet plane density in proton-proton collisions at $\sqrt{\mathrm{s}} = 13$ TeV," *Journal of High Energy Physics*, vol. 2024, May 2024.

[128] ALICE Collaboration, "Measurement of the primary Lund jet plane density in pp collisions at 13 TeV with ALICE," 2021.

[129] A. Lifson, G. P. Salam, and G. Soyez, "Calculating the primary Lund Jet Plane density," *Journal of High Energy Physics*, vol. 2020, Oct. 2020.

[130] B. Andersson, G. Gustafson, L. Lonnblad, and U. Pettersson, "Coherence Effects in Deep Inelastic Scattering," *Z. Phys. C*, vol. 43, p. 625, 1989.

[131] "Multijet simulation for 13 TeV ATLAS Analyses," tech. rep., CERN, Geneva, 2019. All figures including auxiliary figures are available at https://atlas.web.cern.ch/Atlas/GROUPS/PHYSICS/PUBNOTES/ATL-PHYS-PUB-2019-017.

[132] G. D'Agostini, "A multidimensional unfolding method based on Bayes' Theorem," tech. rep., DESY, Hamburg, 1994.

[133] G. D'Agostini, "Improved iterative Bayesian unfolding," 2010.

[134] ATLAS Collaboration, "Study of the material of the ATLAS inner detector for Run 2 of the LHC," *Journal of Instrumentation*, vol. 12, p. P12009–P12009, Dec. 2017.

[135] ATLAS Collaboration, "Performance of the ATLAS track reconstruction algorithms in dense environments in LHC Run 2," *The European Physical Journal C*, vol. 77, Oct. 2017.

[136] B. P. Nachman, M. Shapiro, and F. Rubbo, "Properties of Jet Fragmentation with the ATLAS detector at $\sqrt{s} = 13$ TeV," tech. rep., CERN, Geneva, 2017.

[137] ATLAS Collaboration, "Alignment of the ATLAS Inner Detector in Run-2," *Eur. Phys. J. C*, vol. 80, p. 1194, 2020.

[138] O. Press, "CERN Future Circular Collider (FCC) photo collection for media," 2024. General Photo.

[139] W. Bartmann, J.-P. Burnet, C. Carli, A. Chance, P. Craievich, M. Giovannozzi, C. Grojean, J. Gutleber, K. Hanke, A. Henriques, P. Janot, C. Lourenco, M. Mangano, T. Otto, J. H. Poole, S. Rajagopalan, T. Raubenheimer, E. Todesco, T. P. Watson, and G. Wilkinson, "Future Circular Collider Feasibility Study Report Volume 1: Physics and Experiments," 2025.

[140] M. Benedikt, W. Bartmann, J.-P. Burnet, C. Carli, A. Chance, P. Craievich, M. Giovannozzi, C. Grojean, J. Gutleber, K. Hanke, A. Henriques, P. Janot, C. Lourenco, M. Mangano, T. Otto, J. H. Poole, S. Rajagopalan, T. Raubenheimer, E. Todesco, L. Ulrici, T. P. Watson, G. Wilkinson, and F. Zimmermann, "Future Circular Collider Feasibility Study Report Volume 2: Accelerators, technical infrastructure and safety," 2025.

[141] J. Gutleber, L. Ulrici, and T. P. Watson, "Future Circular Collider Feasibility Study Report Volume 3: Civil Engineering, Implementation and Sustainability," 2025.

[142] FCC Collaboration, "Fcc-ee: The lepton collider – volume 2: The detector," *Eur. Phys. J. ST*, vol. 228, pp. 261–623, 2019.

[143] D. Décamp *et al.*, "ALEPH: A detector for electron-positron annihilations at LEP," *Nuclear Instruments and Methods in Physics Research Section A Accelerators Spectrometers Detectors and Associated Equipment*, vol. 294, 09 1990.

[144] T. I. Collaboration, "International Large Detector: Interim Design Report," 2020.

[145] M. Aicheler *et al.*, eds., *Physics and Detectors at CLIC: CLIC Conceptual Design Report*. No. CERN-2012-003 in CERN Yellow Reports: Monographs, CERN, 2012.

[146] The IDEA Study Group, "The IDEA detector concept for FCC-ee," 2025.

[147] N. Bacchetta *et al.*, "CLD – A Detector Concept for the FCC-ee," 2019.

[148] A. Blondel and M. Dam, "Fcc-ee detector requirements: geometric acceptance requirements for dilepton and diphoton events at z pole energies," Sept. 2023.

[149] A. Blondel, C. Grojean, P. Janot, S. Rajagopalan, and G. Wilkinson, "The FCC integrated programme: a physics manifesto," 2025.

[150] J. de Favereau, C. Delaere, P. Demin, A. Giammanco, V. Lemaître, A. Mertens, and M. Selvaggi, "DELPHES 3: a modular framework for fast simulation of a generic collider experiment," *Journal of High Energy Physics*, vol. 2014, Feb. 2014.

[151] G. Dissertori, A. Gehrmann-DeRidder, T. Gehrmann, E. W. N. Glover, G. Heinrich, and H. Stenzel, "Precise Determination of the Strong Coupling Constant at NNLO in QCD from the Three-Jet Rate in Electron–Positron Annihilation at LEP," *Physical Review Letters*, vol. 104, Feb. 2010.

Appendices

Appendix A

Appendix for the jet η-intercalibration

A.1 Sources of uncertainties in the Jet Energy Scale

Table A.1: Sources of uncertainty in the JES

Component	**Description**
η-intercalibration	
Systematic mis-modelling	Envelope of generator, pile-up, event topology variations
Statistical component	Statistical uncertainty
Non-closure	Three components describing non-closure at high energy and at $\eta \sim \pm 2.4$
Non-closure, 2018 only	Single component describing non-closure at $\eta \sim \pm 1.5$
Z + jet	
Electron scale	Uncertainty in the electron energy scale
Electron resolution	Uncertainty in the electron energy resolution
Muon scale	Uncertainty in the muon momentum scale
Muon resolution (ID)	Uncertainty in muon momentum resolution in the ID
Muon resolution (MS)	Uncertainty in muon momentum resolution in the MS
MC generator	Difference between MC event generators
JVT cut	Jet vertex tagger uncertainty
$\Delta\phi$ cut	Variation of $\Delta\phi$ between the jet and Z boson
Subleading jet veto	Radiation suppression through second-jet veto
Showering & topology	Modelling energy flow and distribution in and around a jet
Statistical	Statistical uncertainty in 28 discrete p_T terms
γ + jet	
Photon scale	Uncertainty in the photon energy scale
Photon resolution	Uncertainty in the photon energy resolution
MC generator	Difference between MC event generators
JVT cut	Jet vertex tagger uncertainty
$\Delta\phi$ cut	Variation of $\Delta\phi$ between the jet and photon
Subleading jet veto	Radiation suppression through second-jet veto
Showering & topology	Modelling energy flow and distribution in and around a jet
Photon purity	Purity of sample used for γ + jet balance
Statistical	Statistical uncertainty in 16 discrete p_T terms
Multijet balance	
$\Delta\phi$ (lead, recoil system)	Angle between leading jet and recoil system
$\Delta\phi$ (lead, any sublead)	Angle between leading jet and closest subleading jet
MC generator	Difference between MC event generators
p_T^{sub} selection	Second jet's p_T contribution to the recoil system
Jet p_T	Jet p_T threshold
Statistical	Statistical uncertainty in 28 discrete p_T terms
Pile-up	
μ offset	Uncertainty in the μ modelling in MC simulation
N_{PV} offset	Uncertainty in the N_{PV} modelling in MC simulation
ρ topology	Uncertainty in the per-event ρ density modelling in MC simulation
p_T dependence	Uncertainty in the residual p_T dependence
Jet flavour	
Flavour composition	Uncertainty in the proportional sample composition of quarks and gluons
Flavour response	Uncertainty in the response of gluon-initiated jets
b-jets	Uncertainty in the response of b-quark-initiated jets
Punch-through	Uncertainty in GSC punch-through correction
Single-particle response	High-p_T jet uncertainty from single-particle and test-beam measurements

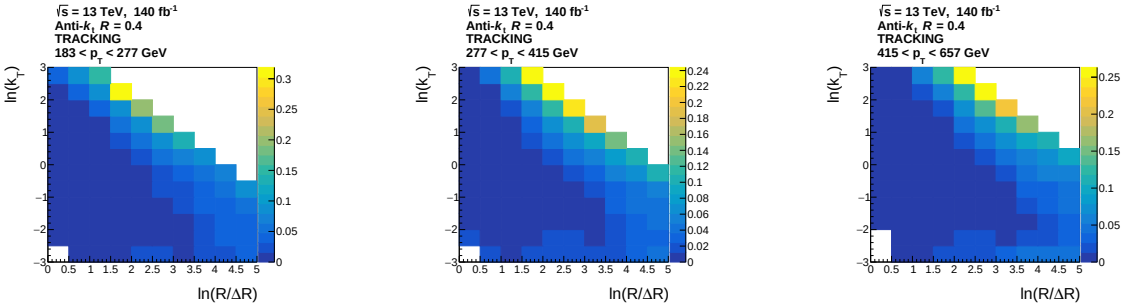
A.2 Trigger Strategy for data 2022

$p_{\mathrm{T}}^{\mathrm{avg}}$ **bin (GeV)**	**Trigger name**
25-30	HLT_noalg_L1ZB
30-60	HLT_noalg_L1ZB
60-85	HLT_j35_L1RDO_FILLED _OR_HLT_j35f_L1RD0_FILLED
85-125	HLT_j45_L1RD0_FILLED _OR_HLT_j45f_L1J15p31ETA49
125-150	HLT_j60_L1J20 _OR_HLT_j60f_L1J20p31ETA49
150-175	HLT_j85_L1J20 _OR_HLT_j85f_L1J20p31ETA49
175-220	HLT_j110_preselj80_L1J30 _OR_HLT_j110f_L1J30p31ETA49
220-270	HLT_j175_preselj140_L1J50 _OR_HLT_j175f_L1J50p31ETA49
270-330	HLT_j175_preselj140_L1J50 _OR_HLT_j175f_L1J50p31ETA49
330-400	HLT_j220_preselj150_L1J75 _OR_HLT_j220f_L1J75p31ETA49
400-525	HLT_j360_L1J100
525-760	HLT_j420_L1J100
760-1100	HLT_j420_L1J100
1100-1500	HLT_j420_L1J100
1500-2000	HLT_j420_L1J100

Table A.2: Trigger strategy table for data22. For the first two $p_{\mathrm{T}}^{\mathrm{avg}}$ bins, a zero bias trigger is used to circumvent issues with low-p_{T} triggers in the early Run-3.

Appendix B

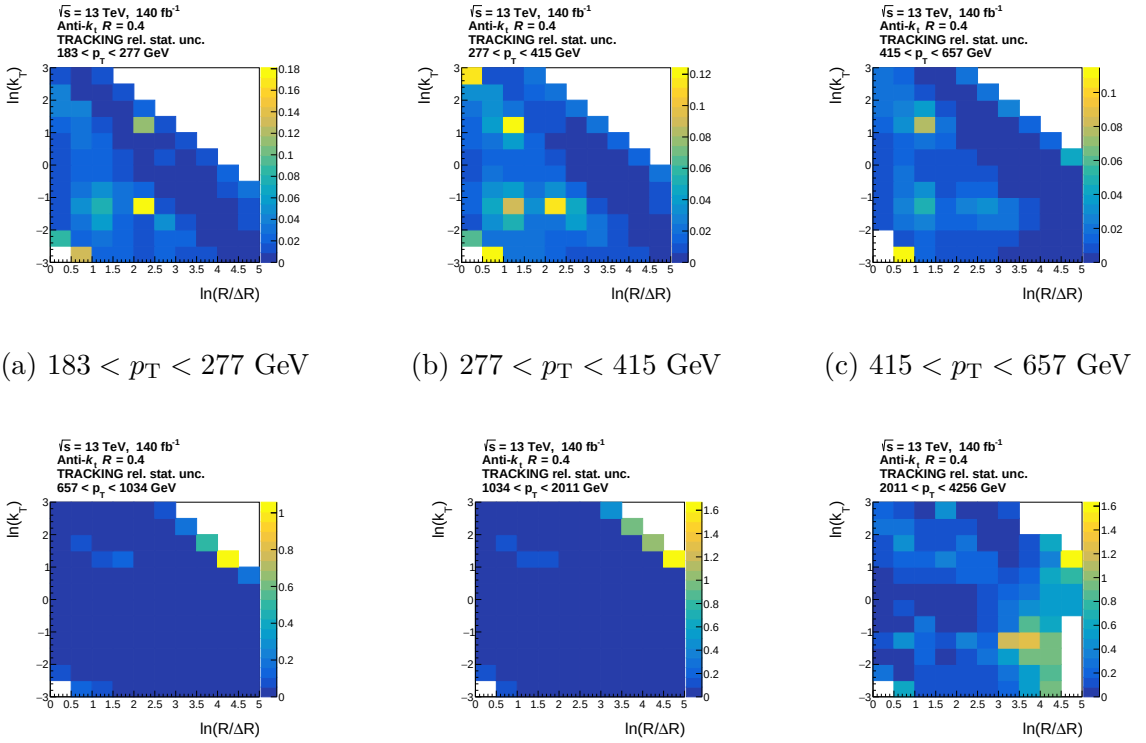
Appendix for the Lund jet plane measurement



(a) $183 < p_{\mathrm{T}} < 277$ GeV (b) $277 < p_{\mathrm{T}} < 415$ GeV (c) $415 < p_{\mathrm{T}} < 657$ GeV

(d) $657 < p_{\mathrm{T}} < 1034$ GeV (e) $1034 < p_{\mathrm{T}} < 2011$ GeV (f) $2011 < p_{\mathrm{T}} < 4256$ GeV

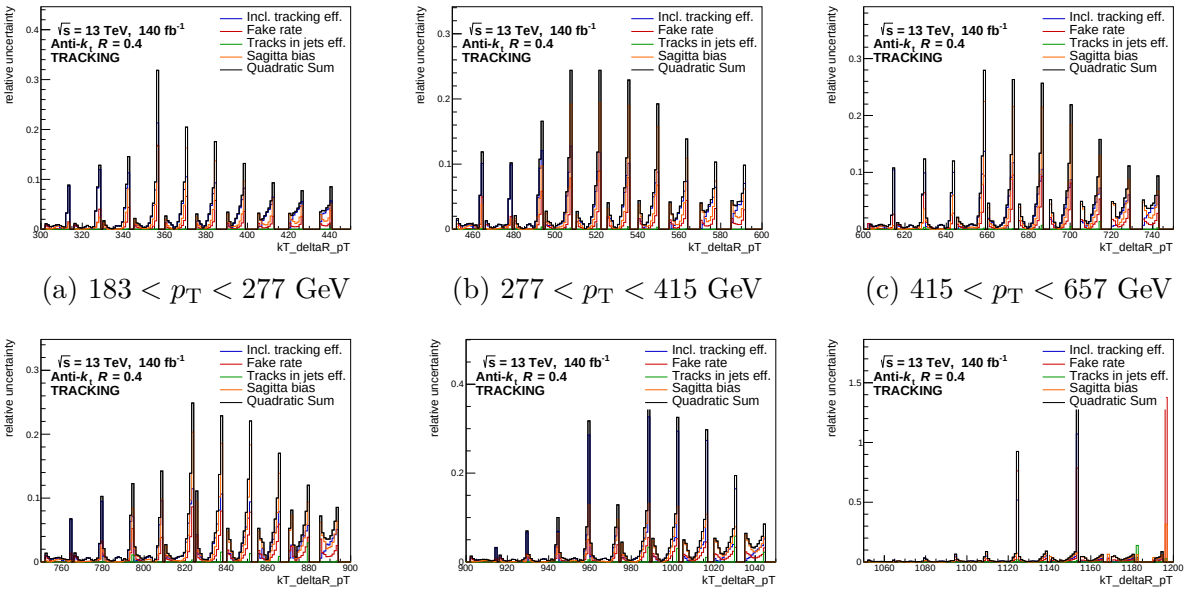
Figure B.1: Uncertainties due to tracking, in p_{T} bins.



(a) $183 < p_{\mathrm{T}} < 277$ GeV (b) $277 < p_{\mathrm{T}} < 415$ GeV (c) $415 < p_{\mathrm{T}} < 657$ GeV

(d) $657 < p_{\mathrm{T}} < 1034$ GeV (e) $1034 < p_{\mathrm{T}} < 2011$ GeV (f) $2011 < p_{\mathrm{T}} < 4256$ GeV

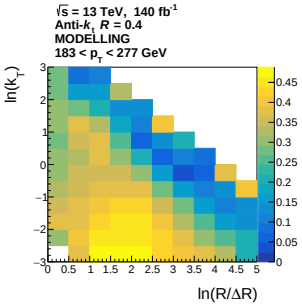
Figure B.2: Relative statistical uncertainty on tracking uncertainties, in p_{T} bins.



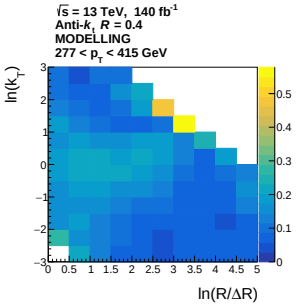
(a) $183 < p_{\mathrm{T}} < 277$ GeV (b) $277 < p_{\mathrm{T}} < 415$ GeV (c) $415 < p_{\mathrm{T}} < 657$ GeV

(d) $657 < p_{\mathrm{T}} < 1034$ GeV (e) $1034 < p_{\mathrm{T}} < 2011$ GeV (f) $2011 < p_{\mathrm{T}} < 4256$ GeV

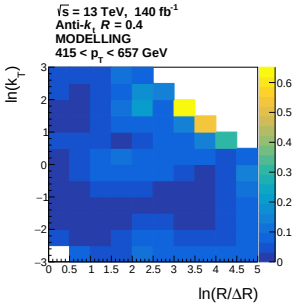
Figure B.3: Breakdown of tracking uncertainties, in p_{T} bins.



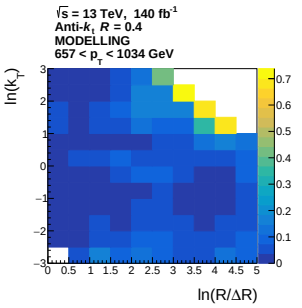
(a) $183 < p_{\mathrm{T}} < 277$ GeV



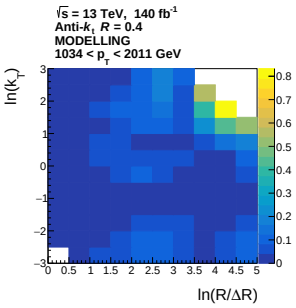
(b) $277 < p_{\mathrm{T}} < 415$ GeV



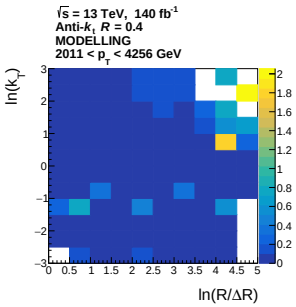
(c) $415 < p_{\mathrm{T}} < 657$ GeV



(d) $657 < p_{\mathrm{T}} < 1034$ GeV

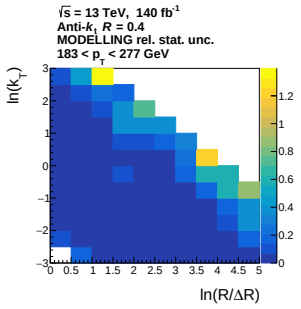


(e) $1034 < p_{\mathrm{T}} < 2011$ GeV

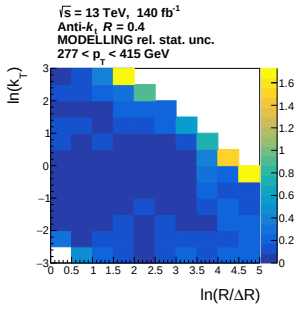


(f) $2011 < p_{\mathrm{T}} < 4256$ GeV

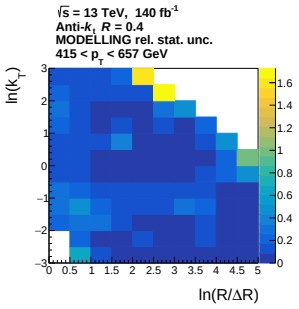
Figure B.4: Modelling uncertainty, in p_{T} bins.



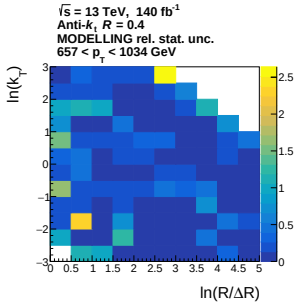
(a) $183 < p_{\mathrm{T}} < 277$ GeV



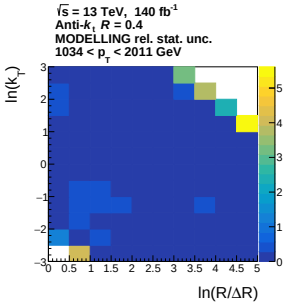
(b) $277 < p_{\mathrm{T}} < 415$ GeV



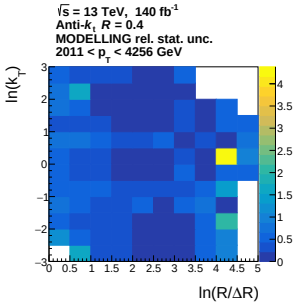
(c) $415 < p_{\mathrm{T}} < 657$ GeV



(d) $657 < p_{\mathrm{T}} < 1034$ GeV

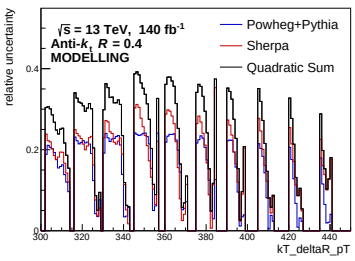


(e) $1034 < p_{\mathrm{T}} < 2011$ GeV

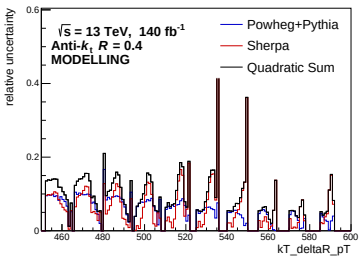


(f) $2011 < p_{\mathrm{T}} < 4256$ GeV

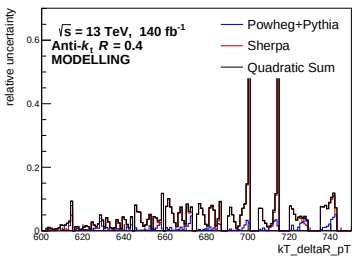
Figure B.5: Relative statistical uncertainty on modelling uncertainties, in p_{T} bins.



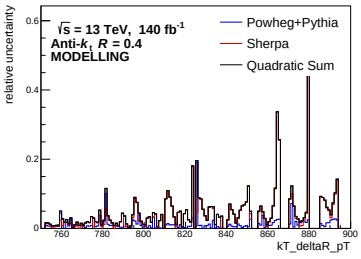
(a) $183 < p_{\mathrm{T}} < 277$ GeV



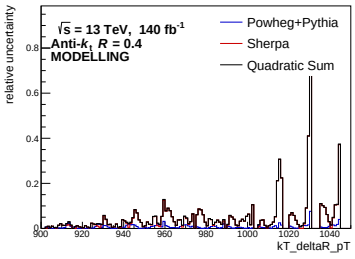
(b) $277 < p_{\mathrm{T}} < 415$ GeV



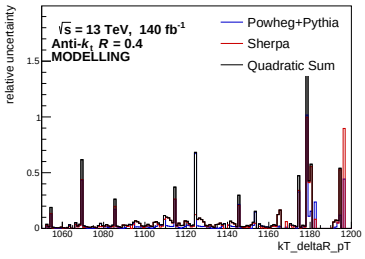
(c) $415 < p_{\mathrm{T}} < 657$ GeV



(d) $657 < p_{\mathrm{T}} < 1034$ GeV

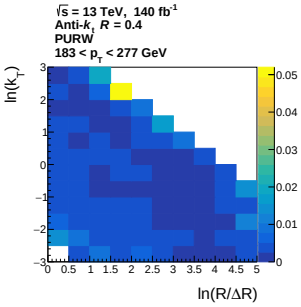


(e) $1034 < p_{\mathrm{T}} < 2011$ GeV

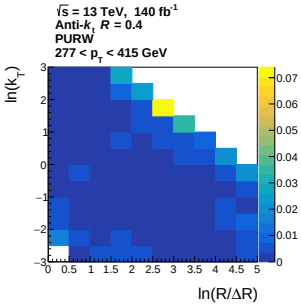


(f) $2011 < p_{\mathrm{T}} < 4256$ GeV

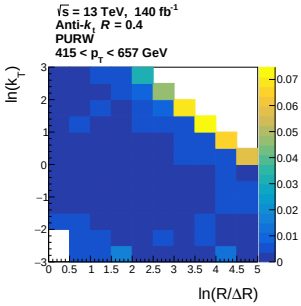
Figure B.6: Breakdown of modelling uncertainties, in p_{T} bins.



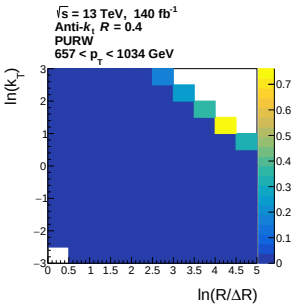
(a) $183 < p_{\mathrm{T}} < 277$ GeV



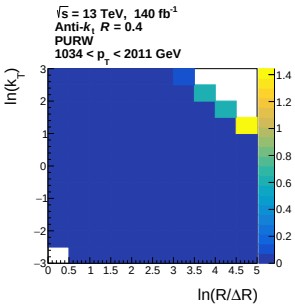
(b) $277 < p_{\mathrm{T}} < 415$ GeV



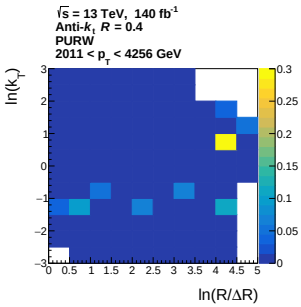
(c) $415 < p_{\mathrm{T}} < 657$ GeV



(d) $657 < p_{\mathrm{T}} < 1034$ GeV

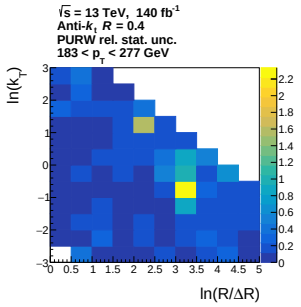


(e) $1034 < p_{\mathrm{T}} < 2011$ GeV

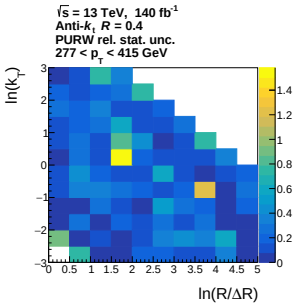


(f) $2011 < p_{\mathrm{T}} < 4256$ GeV

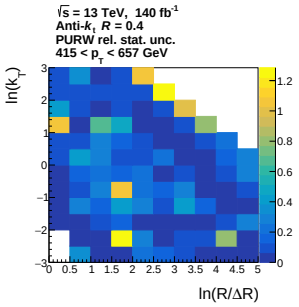
Figure B.7: Uncertainties due to pile-up reweighting, in p_{T} bins.



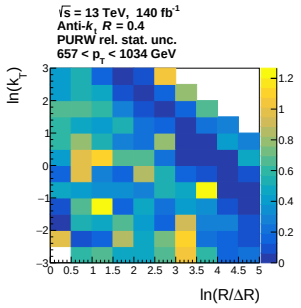
(a) $183 < p_{\mathrm{T}} < 277$ GeV



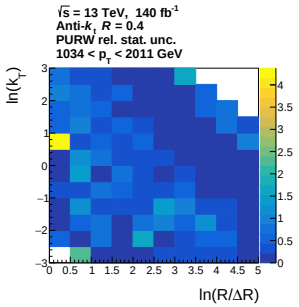
(b) $277 < p_{\mathrm{T}} < 415$ GeV



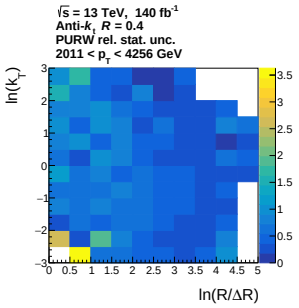
(c) $415 < p_{\mathrm{T}} < 657$ GeV



(d) $657 < p_{\mathrm{T}} < 1034$ GeV

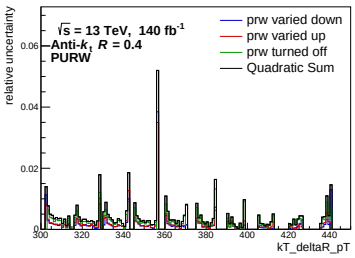


(e) $1034 < p_{\mathrm{T}} < 2011$ GeV

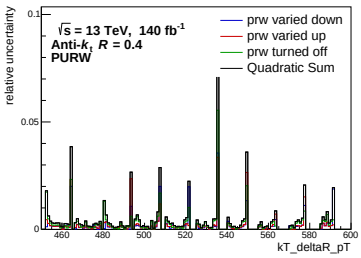


(f) $2011 < p_{\mathrm{T}} < 4256$ GeV

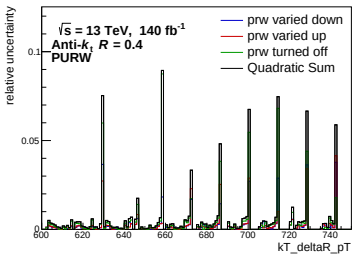
Figure B.8: Relative statistical uncertainty on pile-up reweighting uncertainties, in p_{T} bins.



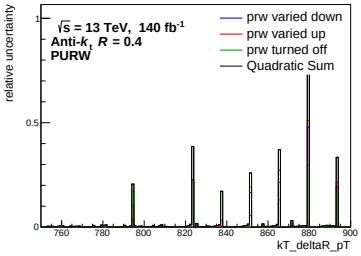
(a) $183 < p_{\mathrm{T}} < 277$ GeV



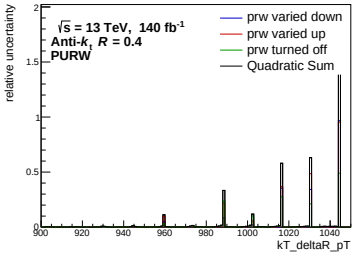
(b) $277 < p_{\mathrm{T}} < 415$ GeV



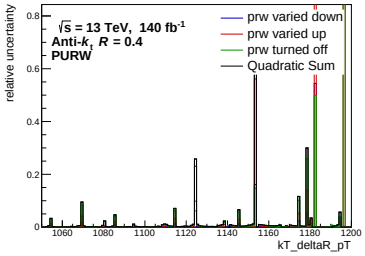
(c) $415 < p_{\mathrm{T}} < 657$ GeV



(d) $657 < p_{\mathrm{T}} < 1034$ GeV



(e) $1034 < p_{\mathrm{T}} < 2011$ GeV



(f) $2011 < p_{\mathrm{T}} < 4256$ GeV

Figure B.9: Breakdown of pile-up reweighting uncertainties, in p_{T} bins.