

Measurement of the underlying event in pp collisions using the ATLAS detector and development of a software suite for Bayesian unfolding



Benjamin Michael Wynne

A thesis submitted in fulfilment of the requirements
for the degree of Doctor of Philosophy to the
University of Edinburgh

August 2012

Abstract

First measurements are made of the underlying event in calorimeter jet events at the LHC, using 37 pb^{-1} of pp collisions at $\sqrt{s} = 7\text{ TeV}$, recorded during 2010 by the ATLAS detector. Results are compared for an assumed di-jet topology based on a single identified jet, and an exclusive di-jet requirement. The number of charged particles in the azimuthal region transverse to the jet axis is recorded, as well as their total and average transverse momentum. The total energy carried by all particles — charged and neutral — is also calculated, using the full calorimeter acceptance $|\eta| < 4.8$. Distributions are constructed to show the variation of these quantities versus the transverse momentum of the selected jet, over the range $20 - 800\text{ GeV}$. Additional jets in the transverse region are shown to dramatically influence the measured activity.

Software is developed to perform Bayesian iterative unfolding, testing closure of the process and stability with respect to the number of iterations performed. Pseudo-experiments are used to propagate systematic errors, and the intrinsic error due to unfolding is estimated. Although the correction relies on a prior probability distribution, model-dependence is reduced to an uncertainty comparable to or smaller than experimental systematic errors.

The software is used to correct underlying event measurements for effects introduced by the ATLAS detector. Unfolded results are compared to predictions from different Monte Carlo event generators used in LHC analyses, showing general agreement in the range $|\eta| < 2.5$, but discrepancies in the forward region. Comparison with other ATLAS results shows compatible behaviour in events defined by any high-momentum charged particle, or by leptonic Z-boson decays.

Declaration

This dissertation is the result of my own work, except where explicit reference is made to the work of others, and has not been submitted for another qualification to this or any other university. This dissertation does not exceed the word limit for the respective Degree Committee.

Ben Wynne

Acknowledgements

Phil Clark, for supervising.

Andy Buckley and Peter Clarke, for help and advice.

Deepak Kar and Sebastian Wahrmund, for collaborating.

Ailsa, Alastair, Alex, Alex, Alisdair, André, Anne, Ant, Ben, Caspar, Caterina, Chrissy, Clare, Conor, Craig, Dave, Dave, Emma, Freya, Gareth, Gary, Gemma, Gemma, Gordon, Graham, Hari, Ioan, James, James, Jamie, Jody, Katie, Katy, Kyle, Kym, Lawrence, Lee, Madeline, Mark, Mark, Nathan, Nikki, Øyvind, Rachel, Rob, Ross, Sam, Sam, Samuel, Tim, Tom, and Vicky, for fun times and bright ideas.

Mum and Dad, for remarkable patience.

Preface

None of the work presented in this thesis would have been possible without the thousands of people responsible for the design, construction and operation of the ATLAS detector, the Large Hadron Collider, and its supporting accelerators. Any one analysis of ATLAS data relies on calibrations, techniques, and performance studies created by countless others. Direct contributions by the author presented in this thesis are outlined below:

Chapter 3: The author performed upgrade work on the Inner Detector control software, as described in Section 3.8. The operation of the Inner Detector cooling system was supported.

Chapter 6: For all data and MC samples used in the analysis, topocluster variables were constructed by the author. Final selections and jet corrections were applied, and the full analysis was performed on some MC samples.

Chapter 7: The Imagiro software package described in Section 7.3 was written by the author. Its use in several ATLAS analyses was supported, and all unfolding for the Underlying Event analysis was performed by the author, including studies of iteration stability and closure.

Chapter 8: All leading jet underlying event plots presented were created by the author, unless otherwise specified. The data stability study in Section 8.3 was performed by the author.

Contents

1. Introduction	2
2. The ATLAS experiment: overview	4
2.1. The Large Hadron Collider	4
2.1.1. LHC magnets	6
2.1.2. LHC beam acceleration	8
2.2. The ATLAS detector	9
2.2.1. Coordinate system	9
2.2.2. Muon tracker	11
2.2.3. Calorimetry	14
2.2.4. Hadronic calorimeter	16
2.2.5. Electromagnetic calorimeter	18
2.2.6. Inner Detector	19
2.3. ATLAS infrastructure	20
2.3.1. Detector Control Systems	21
2.3.2. Trigger	21
3. The ATLAS experiment: Inner Detector	24
3.1. Momentum measurement	25
3.2. Vertex measurement	26
3.3. Pixel detector	27
3.4. Semiconductor tracker	29
3.5. Transition Radiation Tracker	30
3.6. Radiation damage	31
3.7. Cooling	33
3.8. Inner Detector DCS	35
3.8.1. Distribution rack upgrade	35
3.8.2. TEH automatic recovery	36
3.8.3. SCT power-cut detection	37

3.8.4. SCT automatic switch-on	37
3.9. Combined Inner Detector performance	39
4. The Standard Model	42
4.1. Quantum Chromodynamics	44
4.1.1. The QCD coupling strength	46
4.2. Phenomenology	47
4.2.1. Parton density functions	49
4.2.2. Multiple parton interactions	50
4.2.3. Parton showers and hadronisation	52
4.2.4. Jets	53
5. Monte Carlo generators	57
5.1. Generators and tunes	58
5.2. Multiple parton interactions	60
5.2.1. Evolution of $p_{T\min}$	61
5.2.2. Hadronic form factor	62
5.3. Parton showers	63
5.4. Hadronisation	64
5.4.1. String hadronisation	65
5.4.2. Cluster hadronisation	67
5.4.3. Colour reconnection	68
6. The Underlying Event: analysis	69
6.1. Topological selection	70
6.2. Event selection	72
6.3. Jet selection	73
6.4. Track selection	76
6.5. Topocluster selection	78
6.6. Correction procedure	79
6.7. Monte Carlo samples	79
7. Bayesian unfolding	82
7.1. Mapping and unfolding	82
7.2. Alternatives to matrix inversion	84
7.2.1. Bin-by-bin unfolding	84
7.2.2. Hit Backspace Once More (HBOM)	85

7.2.3. Smearing matrix regularisation	85
7.2.4. Bayesian iterative unfolding	86
7.3. Imagiro	87
7.3.1. Multiple priors	87
7.3.2. Closure tests	88
7.3.3. Conditions for convergence	88
7.3.4. Unfolding correlated variables	90
7.3.5. Fake or missing events	90
7.4. Error treatment	91
7.4.1. Statistical error calculation	91
7.4.2. Fast statistical error estimation	92
7.4.3. Intrinsic systematic error	92
7.4.4. Systematic error propagation	93
7.5. Demonstration with toy models	93
8. The Underlying Event: results	97
8.1. ATLAS Underlying Event measurements	97
8.1.1. Charged particle $\sum p_T$ and multiplicity vs. p_T^{lead}	98
8.1.2. Charged and neutral particle $\sum E_T$ vs. p_T^{lead}	101
8.1.3. Charged particle $\langle p_T \rangle$ vs. p_T^{lead} and N_{ch}	104
8.2. Comparing the inclusive and exclusive selections	105
8.3. Data stability	112
8.4. Systematic errors	115
8.4.1. Background	115
8.4.2. Tracks per vertex	115
8.4.3. Merged vertices	116
8.4.4. Jet energy scale uncertainty	117
8.4.5. Jet energy resolution uncertainty	118
8.4.6. Jet reconstruction efficiency uncertainty	118
8.4.7. Track reconstruction efficiency uncertainty	118
8.4.8. Cluster energy uncertainty	118
8.4.9. Unfolding model-dependence	119
8.4.10. Results of error propagation	119
8.5. Comparison with other UE studies	120
9. Conclusion	124

A. Unfolding closure and iteration stability	127
B. Data stability plots	140
C. Propagation of systematic errors	146
D. Underlying event observable 1D plots	153
D.1. Full p_T^{lead} range	153
D.2. Restricted p_T^{lead} ranges	159
Bibliography	165
List of Figures	176
List of Tables	181

Chapter 1.

Introduction

The ATLAS detector at the Large Hadron Collider (LHC) was built to investigate the fundamental forces of nature at unprecedented energy scales. It allows physicists to test the Standard Model of particle physics, and to search for new physics beyond it. One area of investigation is the theory of Quantum Chromodynamics (QCD), which — while very successful — struggles to make predictions in energy regimes where the strength of the force grows large. Since hadron collisions at the LHC are governed by QCD, the ATLAS detector can make detailed investigations of interactions in this regime. Searches for new physics also require an understanding of the QCD-dominated background from Standard Model processes.

Chapter 2 gives an overview of the LHC and the ATLAS detector, while the ATLAS Inner Detector is described in more detail in Chapter 3. A brief introduction to the Standard Model is given in Chapter 4, with a focus on QCD and how it relates to the different processes occurring in a hadron collision. When QCD cannot predict these processes, approximate models are introduced and adjusted to reproduce observed behaviour. Some widely-used models are described in Chapter 5, along with software packages that use them to simulate LHC physics.

In any experimental procedure, care must be taken to correct the result for any effects produced by the apparatus used to make the measurements. As the apparatus grows more complex, more factors arise that must be corrected for. With an apparatus as complex as the ATLAS detector, the risk is that components will behave or interact in unexpected ways, meaning that it is impossible to explicitly correct for every effect. To address this problem, unfolding is introduced as a final corrective step. The behaviour of the detector is simulated to construct a mapping between ideal experimental results and the imperfect measurements, and then this mapping is used to correct real data.

An unfolding method based on Bayes' theorem is described in Chapter 7. The method is implemented in a new software package, and the features of this software are discussed and demonstrated.

The ‘underlying event’ is the part of a hadron-hadron collision that is not associated with the hard scattering process. It is modelled as arising primarily from additional partonic interactions between the hadrons. Measuring the underlying event allows these models to be tested, and gives better description of the expected background for new physics searches by providing input for MC model tuning.

Chapter 6 gives details of the first analysis measuring the underlying event associated with jet production at the LHC. The results are corrected using Bayesian unfolding, then compared with different QCD model predictions and experimental results in Chapter 8.

Chapter 2.

The ATLAS experiment: overview

The European Organisation for Nuclear Research (called CERN after “Conseil Européen pour la Recherche Nucléaire”) was established in 1954, and operates a number of particle accelerators at its laboratory site on the French-Swiss border northwest of Geneva. Between 1984 and 1989 a circular tunnel 26.7 km in circumference was built to house the proposed Large Electron-Positron collider (LEP) [1], with two 2.5 km tunnels linking it to the existing CERN particle accelerator chain. When planning a particle accelerator to succeed LEP, reusing this tunnel was the obvious choice given the construction costs involved. Therefore, in 2000 LEP was closed down and dismantled to make way for this new accelerator: the Large Hadron Collider (LHC) [2].

2.1. The Large Hadron Collider

Unlike LEP, the LHC is designed to collide hadrons: specifically protons or lead nuclei. Two beams of these charged particles circulate in opposite directions around the tunnel, with superconducting magnets used to direct them. At four points around the tunnel the beams are brought together and the particles collide: these are known as ‘interaction points’. Each of these points is home to a large particle detector.

Of these four experiments, ATLAS (A Toroidal LHC ApparatuS) [3] and CMS (Compact Muon Solenoid) [4] are multi-purpose detectors to search for new particles. ALICE (A Large Ion Collider Experiment) [5] is designed with more specific features to study properties of lead ion collisions. LHCb (Large Hadron Collider beauty) [6] is a specialised detector for identifying and measuring decays of B-hadrons.

There are three other experiments using the LHC. TOTEM (TOTal Elastic and diffractive cross section Measurement) [7] and LHCf (Large Hadron Collider forward) [8] are both designed to detect particles leaving collisions travelling very close to the beams, with TOTEM centred on the CMS interaction point, and LHCf on the ATLAS interaction point. MoEDAL (Monopole and Exotics Detector at the LHC) [9] is a plastic track-etch detector for highly-ionising, stable, massive particles leaving the LHCb interaction point.

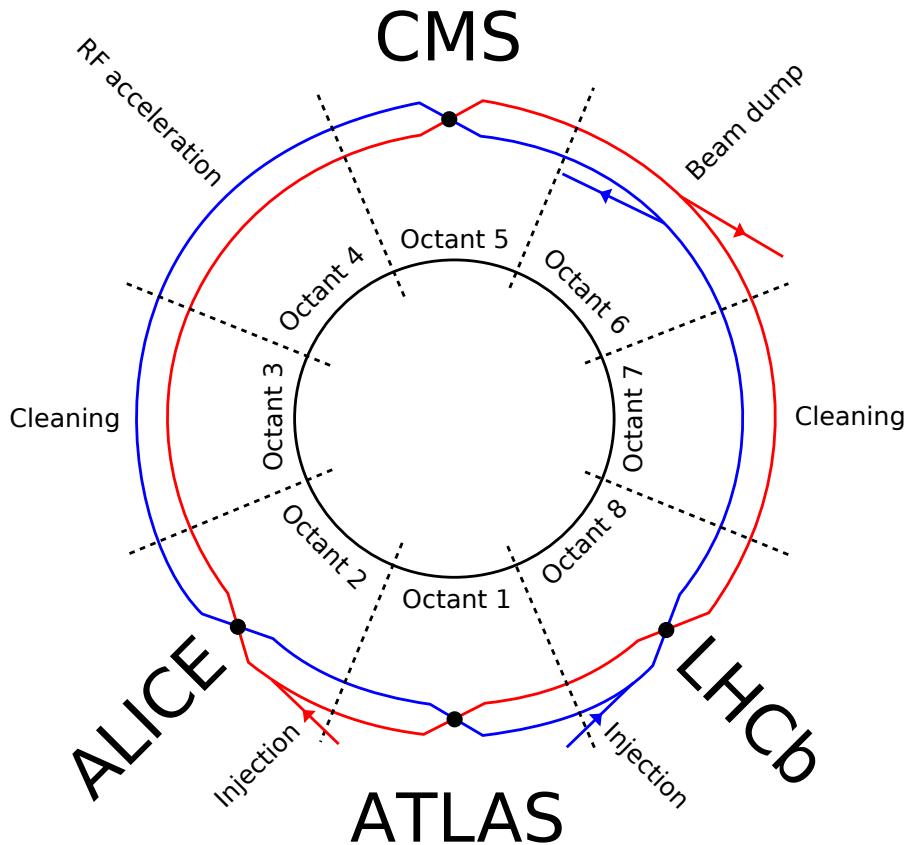


Figure 2.1.: Diagram of the LHC adapted from Reference [2], showing the four experiments at the interaction points, the two beams (beam 1 in red, beam 2 in blue), and the accelerator infrastructure.

The LHC ring is organised in eight sections, as shown in Figure 2.1. The ATLAS, ALICE, CMS and LHCb experiments are found in octants 1, 2, 5 and 8 respectively. Of the remaining octants, 3 and 7 house collimators to ‘clean’ the proton beams. Particles are scattered out of the beam in octant 3 if their momentum differs too greatly from the intended value, and in octant 7 they are scattered if they are too far from the beam axis. Octant 4 contains equipment to accelerate the particles within the beams using electric fields (see Section 2.1.2). The circulating beams can be removed from the LHC in octant 6 when they have degraded too much to be useful, or in the event of a fault.

2.1.1. LHC magnets

The design goals of the LHC were to collide protons with a centre-of-mass energy of $\sqrt{s} = 14 \text{ TeV}$ and a luminosity of $L = 10^{34} \text{ cm}^2\text{s}^{-1}$. This allows experiments at the LHC to search for new physics at previously inaccessible energy scales (compare with $\sqrt{s} = 1.96 \text{ TeV}$ for the Tevatron [10], a $p\bar{p}$ collider at Fermilab), and to detect rare processes. Protons were chosen rather than electrons to limit beam energy losses due to synchrotron radiation, which scales with particle mass m^{-4} . The rate at which a process occurs at a collider is the product of the collider luminosity and the cross-section, σ_{event} , for that process:

$$N_{\text{event}} = L\sigma_{\text{event}},$$

where N_{event} is the number of a particular event occurring per second. The relatively high design luminosity required that the LHC be a particle-particle collider — not particle-antiparticle — since presently there is not the capability to produce sufficient numbers of antiprotons. Circulating like-charged particles in opposite directions required that there be two LHC beam-pipes with opposite magnetic fields, whereas particle-antiparticle colliders only need a single beam pipe. The two beam-pipes also allow the circulation of heavier particles (like lead nuclei) that cannot practically be produced as antiparticles.

Designed to fit into the LEP tunnel, the LHC has a maximum radius of 2804 m [2] for bending the paths of the beams (the beams do not travel in perfect circles, there are straight sections). Using this and the intended 14 TeV centre-of-mass energy (i.e. 7 TeV energy for a particle in one of the two beams) the necessary magnetic field strength can be calculated. Starting with the Lorentz force on a charged particle moving in an electromagnetic field:

$$F = q(E + v \times B) \tag{2.1}$$

and the centripetal force required for circular motion

$$F = mv^2/r \tag{2.2}$$

Equate the forces and maximise the bending by putting the magnetic field perpendicular to the motion of the particle (no electric field):

$$qvB = mv^2/r \quad (2.3)$$

$$qBr = p \quad (2.4)$$

For relativistic particles the substitution $p = E/c$ can be used, giving

$$B = \frac{E}{qrc} \quad (2.5)$$

This gives a required magnetic field strength of 8.3 T. Superconducting magnets are needed to produce this field strength. Like previous particle accelerators, niobium-titanium magnets were used in the construction of the LHC. The magnets were built within cryostats that use superfluid helium to maintain a temperature below 1.9 K, as shown in Figure 2.2.

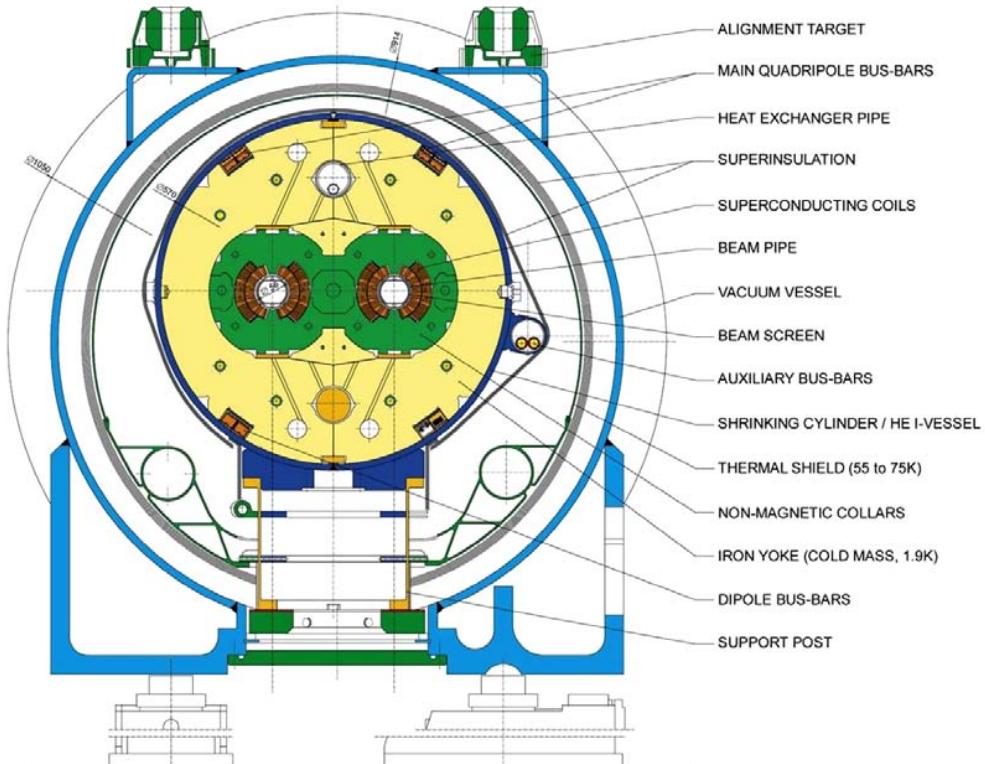


Figure 2.2.: Cross-section through an LHC dipole magnet and the structures around it, showing the two beam-pipes and the helium cryostat [2].

While superconducting magnets allow the creation of strong magnetic fields, they come with an associated problem: the risk of ‘quenching’. If a magnet loses superconductivity then resistive heating will occur, driven by the current in the magnet itself. The magnet rapidly heats up and can be damaged; the cooling system can be overloaded and the liquid helium boils. Quenches are more likely at higher magnetic field strength, and so to date the LHC has been operating at around half the design beam energy: 3.5 TeV in 2010 and 2011, and 8 TeV in 2012.

Besides the dipole magnets used to direct the beam, quadrupoles are used to focus it and a number of more specialised magnets are used in irregular areas like the beam injection points.

2.1.2. LHC beam acceleration

Particles in the LHC have already passed through a chain of smaller accelerators at CERN, as shown in Figure 2.3.

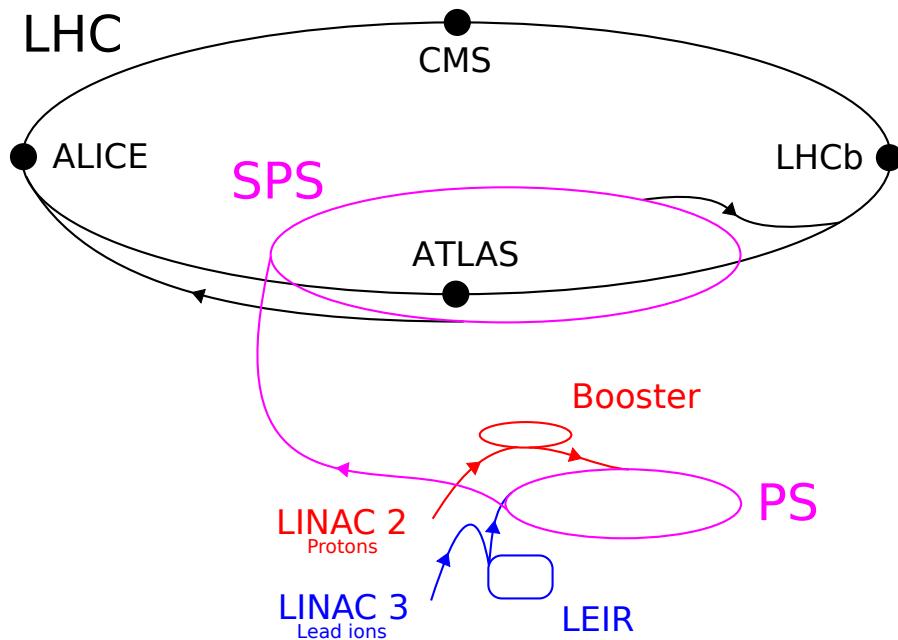


Figure 2.3.: Schematic of the CERN accelerator chain from the initial particle sources and LINear ACcelerators (LINACs) to the LHC. Adapted from Reference [11], other experiments using these accelerators not shown.

The final accelerator before the LHC is the Super Proton Synchrotron (SPS) [12] that injects particles at relativistic speeds, but still well below the design energy of the

LHC. For protons this injection energy is 450 GeV. To reach the collision energy the LHC uses superconducting Radio Frequency (RF) cavities. In these cavities the beam passes through an electric field oscillating at 400.8 MHz [13]. The beams are injected as bunches of $\sim 10^{11}$ particles, and the RF waves maintain these bunches: if the centre of a bunch coincides with the rising edge of a wave then particles at the front of the bunch will encounter a lower field and are thus accelerated less, falling back towards the centre. Conversely those particles that arrive late encounter a higher field which accelerates them more, and so they catch up with the centre of the bunch. The 400.8 MHz frequency is chosen because it is the highest multiple of the SPS frequency (200.4 MHz) that can accommodate the SPS bunch length of 1.6 ns [14]. To accelerate the particles the strength of the electric field is increased from 8 MV when the particles are injected to 16 MV at the designed collision energy [2].

2.2. The ATLAS detector

The ATLAS detector is one of the largest scientific apparatuses ever built, approximately 44 m long, 25 m tall and weighing around 7000 t [3]. The whole experiment can be thought of as a set of nested cylinders, centred on the interaction point. Each of these is a subdetector: a more-or-less independent machine with a specific measurement to make. From the outside layer in, the subdetectors are the muon tracker, the hadronic calorimeter, the electromagnetic calorimeter, and the Inner Detector. Every one can be further subdivided by the technologies used to make their measurement, and by whether they are in the ‘barrel’ region coaxial with the beam, or the ‘end-cap’ regions perpendicular to it. The overall layout of the ATLAS detector is shown in Figure 2.4.

2.2.1. Coordinate system

The ATLAS coordinate system is centred on the nominal interaction point (i.e. the point where particle collisions are expected). The z -axis is parallel with the beam pipe, positive in the anti-clockwise direction around the LHC when viewed from above. The half of the detector in the positive z region is called side A, the other side C. The $x - y$ plane is perpendicular to the beam, with the y -axis vertical and positive y directed upwards. The x -axis is positive pointing towards the centre of the LHC. Coordinate angles are defined

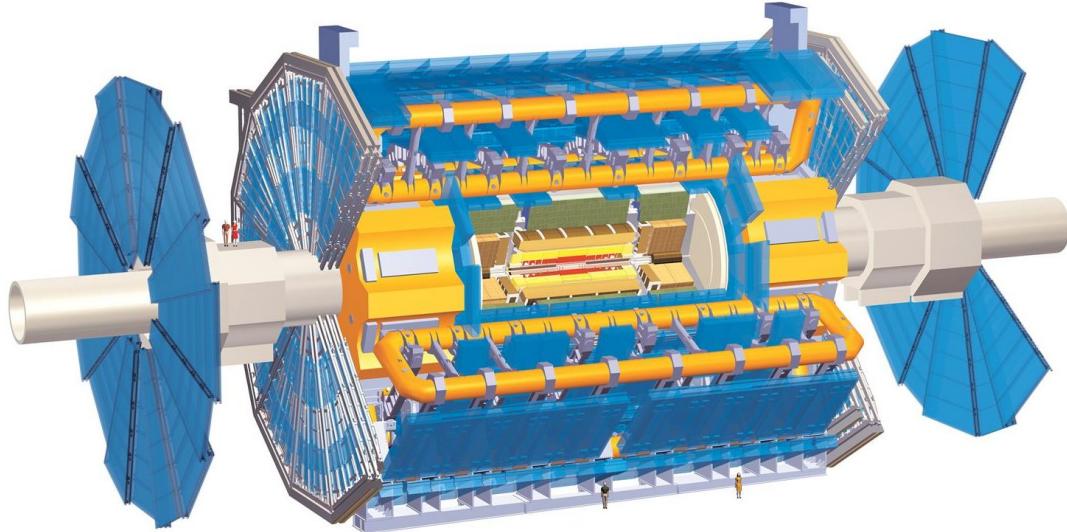


Figure 2.4.: Diagram of the ATLAS detector, showing the different subdetectors [15]. The muon tracker is in blue with its associated magnet coils in orange. The tile hadronic calorimeters are in green, and the liquid argon hadronic and electromagnetic calorimeters are brown. The Inner Detector is the multicoloured region in the very centre.

conventionally, with the polar angle θ relative to the positive z -axis, and the azimuthal angle ϕ in the $x - y$ plane relative to the positive x -axis. When ‘transverse’ quantities are mentioned, such as transverse momentum (p_T), this refers to the component of the vector in the $x - y$ plane. Radius R is defined perpendicular to the beam pipe: $R = \sqrt{x^2 + y^2}$. Figure 2.5 illustrates the coordinate system.

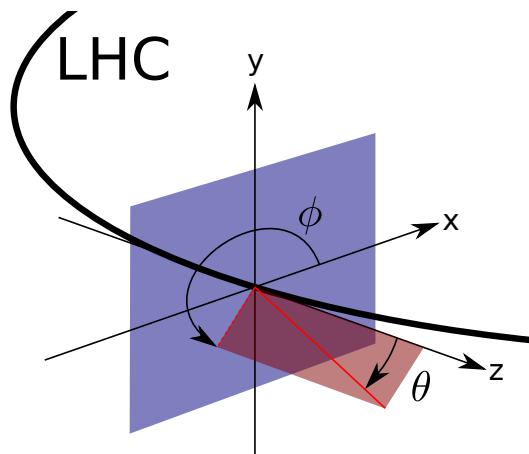


Figure 2.5.: Diagram of the ATLAS coordinate system relative to the LHC beam axis, with the $x - y$ plane in blue and a vector projected onto it in red.

In addition to these coordinates, the pseudorapidity (η) is defined as:

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

For relativistic particles, η approximates the Lorentz invariant rapidity (y):

$$y = \frac{1}{2} \ln \left(\frac{E + p_z}{E - p_z} \right), \quad (2.7)$$

where E is the energy of the object and p_z is the component of its 3-momentum parallel to the z -axis. Pseudorapidity is easier to calculate, but for higher mass particles the approximation breaks down and so rapidity is used instead.

Particle production is uniform in pseudorapidity, meaning that in higher pseudorapidity regions (close to the beam pipe) more particles can be expected in a given volume, as illustrated by Figure 2.6.

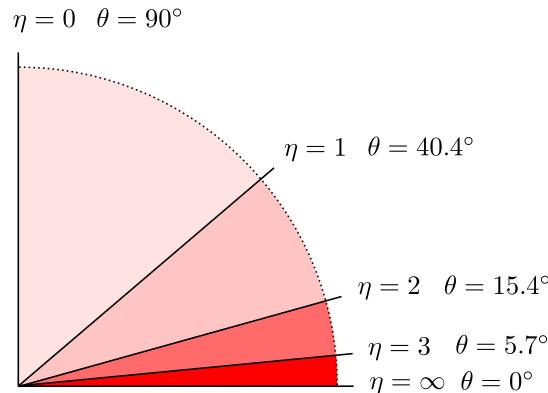


Figure 2.6.: The variation of pseudorapidity (η) with the polar angle (θ). The shading gives an idea of relative particle density.

2.2.2. Muon tracker

Muons have a long lifetime relative to their transit of the ATLAS detector, and are penetrating particles. Therefore the muon tracker forms the outer layer of ATLAS, with muons passing through all the other subdetectors first. This provides particle selection, as it is unlikely that other particles will escape the inner layers (except for neutrinos, which ATLAS cannot detect).

The muon system is built around the huge toroidal magnets that give ATLAS its name, and these produce a field which bends the paths of the passing muons. A field strength of 0.5 T in the barrel region and 1.0 T in the end-caps is achieved. Using precision tracking to measure the curvature of these paths allows the muon momentum to be calculated (see Section 3.1). The goal is a 10 % momentum resolution for muons with an energy of 1 TeV: this corresponds to a tracking uncertainty $\leq 50 \mu\text{m}$ [3] in the bending direction, z (R in the end-caps). Additionally, the detector is designed to trigger on muons: to rapidly identify the presence of a muon with transverse momentum greater than a given threshold.

To achieve these goals there are two sets of muon detectors, for tracking and triggering, each using two different technologies depending on the detector region, as described in Table 2.1. Tracking detectors are designed with a focus on good position resolution, while triggering detectors focus on fast response. Different technologies are used in higher pseudorapidity regions because of the higher occupancy (number of particles in a given detector volume).

All of these detectors use the same central principle: a closed volume of inert gas with an electric field across it produced by sets of electrodes. A passing muon ionises gas along its path, and the ions are accelerated towards the electrodes by the field, creating a pulse of current. The difference is found in the field strengths, the gas mixtures, and the design of the electrodes.

Purpose	Technology	Region	Gases	Voltage
Tracking	Monitored Drift Tube (MDT)	$ \eta < 2.7$	Ar CO ₂	3 kV
Trigger	Resistive Plate Chamber (RPC)	$ \eta < 1.05$	C ₂ H ₂ F ₄ iso-C ₄ H ₁₀ SF ₆	9.8 kV
Tracking	Cathode Strip Chamber (CSC)	$2.0 < \eta < 2.7$	Ar CO ₂	1.9 kV
Trigger	Thin Gap Chamber (TGC)	$1.05 < \eta < 2.4$	CO ₂ n-C ₅ H ₁₂	2.9 kV

Table 2.1.: The different technologies used in different pseudorapidity ranges of the ATLAS muon systems [3].

- MDTs have a long metal tube (diameter 30 mm) as a cathode, containing the gas mixture at 3 bar and a central anode wire. The tubes are all arranged parallel to the $x - y$ plane, giving the z coordinate of muons with a precision of 35 μm when combining the results from 6 – 8 layers of tubes grouped in a single chamber.
- CSCs use a chamber filled with a similar gas mixture and parallel anode wires, with layers of cathode strips in between the wire layers. Charge collected by the wires is not read out — position information comes from the cathodes alone, which are arranged in alternating layers parallel or perpendicular to the wires. Combining 4 layers of cathodes and anodes in each chamber gives an R coordinate with 40 μm precision, and ϕ with 5 mm precision.
- RPCs do not use wires at all, but two parallel flat plates of resistive material, separated by a 2 mm gap. Metallic strips detect charge on these plates by capacitative coupling, with strips for one plate perpendicular to those for the other. The plates are attached to MDT chambers in the central rapidity region, providing a ϕ coordinate measurement. At least three layers of RPCs will be encountered by muons leaving the interaction point.
- TGCs are very much like CSCs, with the distinction that the distance between cathode strips and anode wires is small relative to the wire spacing (1.4 mm versus 1.8 mm). This allows for faster charge collection, in response to the higher occupancy in the forward region. A chamber contains a single layer of anode wires with cathode strips on either side.

The muon system has the largest dimensions of any ATLAS subdetector, yet the components must be aligned with a precision better than 30 μm . To achieve this around 12000 optical sensors monitor the relative positions of the components. The magnetic field is also monitored (rather than assume uniformity over such a large volume) using 1800 Hall probes.

The muon tracker as designed is intended to measure muon p_{T} with an uncertainty of 3% at 100 GeV, and up to 10% at 1 TeV [16]. Cosmic ray muons can be used to test this resolution: the momentum of a muon passing down through the entire detector is measured separately in the upper and lower halves of the muon tracker, and the results are compared. Figure 2.7 shows the results from early commissioning, indicating that (particularly at high p_{T}) the design goals had not yet been met. Improved understanding of the detector alignment has since lead to the momentum resolution approaching its design goals [17]. However, since the analysis presented in Chapter 6 focusses on data

taken early in ATLAS operations, the initial performance of the detector is presented throughout this thesis.

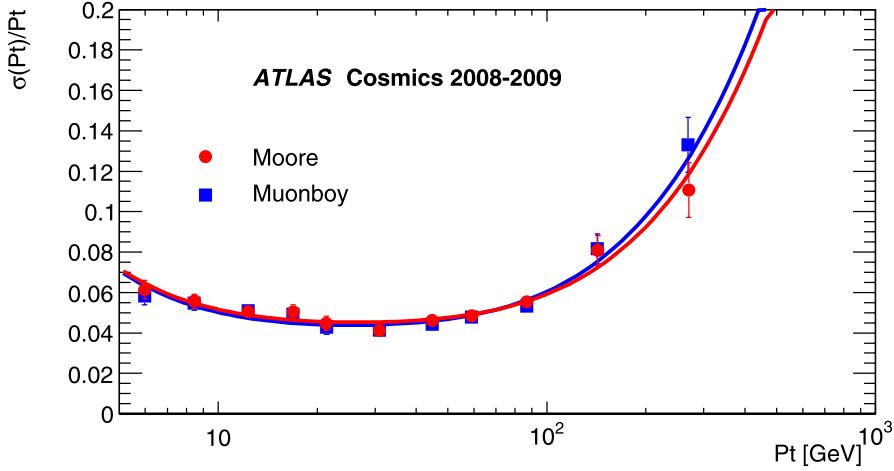


Figure 2.7.: The momentum resolution achieved with the muon tracker in commissioning with cosmic rays [16]. The two trend lines represent results from different track reconstruction algorithms.

2.2.3. Calorimetry

Within the muon system are the calorimeters, designed to absorb incident particles and measure their energy. All of the calorimeters in ATLAS are sampling calorimeters: they use different materials for absorption and energy measurement. The absorbing material is chosen to interact with the incoming particles so that a single high energy particle will produce a shower of lower energy particles. Characteristic parameters for absorbing materials are the radiation length X_0 and interaction length λ . Where a particle's interactions with the material are predominantly electromagnetic (e.g. e^\pm and γ) the relevant parameter is X_0 . Over a distance of X_0 an e^\pm will be reduced to $1/e$ of its initial energy through bremsstrahlung, and a photon has a $7/9$ chance of producing an e^+e^- pair. Strong force interactions are dominant for hadrons, with the mean free path of a hadron in a medium given by λ for that medium.

As electromagnetic and hadronic showers interact with a medium, the number of particles in the shower increases and the energy of those particles decreases. For lower energy particles the dominant mode of energy loss is through ionisation of the medium. Measuring this ionisation allows the energy of the shower — and thus the original particle — to be calculated. Layers of material sensitive to ionisation are placed between the layers

of absorbing material in a sampling calorimeter in order to make the energy measurement. Although ionisation will occur throughout the calorimeter, only ionisation in the sensitive material will be detected, hence the name ‘sampling’ calorimeter. The final energy result must be calibrated accordingly: this is discussed in Section 6.3.

As mentioned already, neutrinos cannot be detected by ATLAS given how rarely they interact with any detection medium. Their presence must be inferred by missing energy: the observation that having summed over the energies of all the particles that were detected, energy conservation is not satisfied. To make this measurement the ATLAS calorimeters enclose the interaction point as completely as possible, extending to $|\eta| < 4.9$. Separate calorimeter systems are designed to overlap, preventing particles escaping through insensitive regions. The calorimeter is also at least 9.7λ thick, ensuring the amount of energy lost by particles (other than muons) ‘punching through’ the calorimeter is small. The layout of the different calorimeter systems is shown in Figure 2.8.

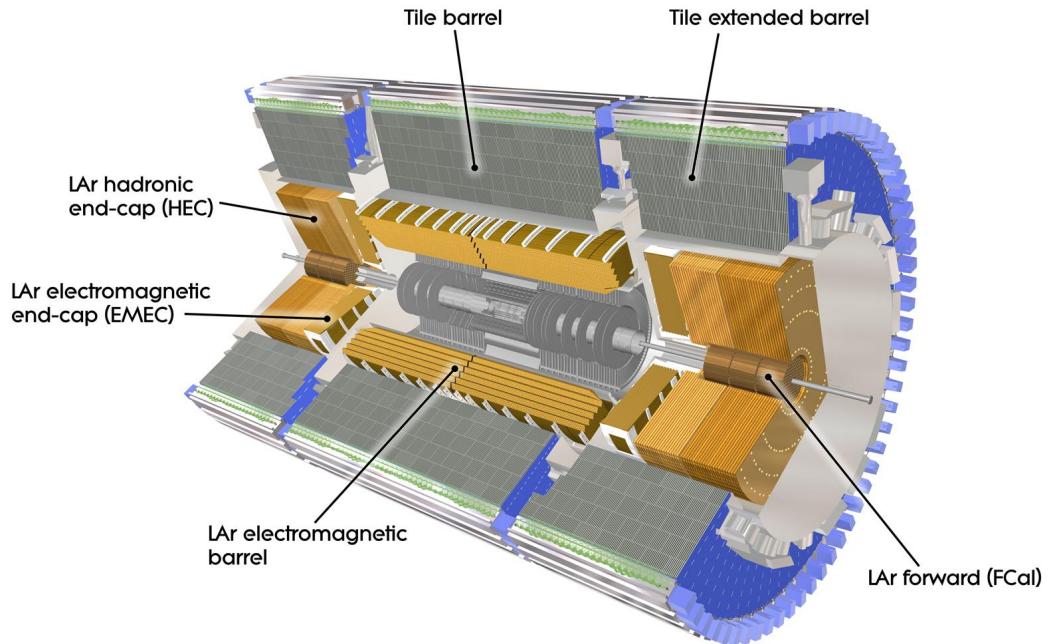


Figure 2.8.: The layout of the ATLAS electromagnetic and hadronic calorimeters [18].

2.2.4. Hadronic calorimeter

The outer calorimeter systems must capture all hadrons leaving the interaction point, and so they make up the majority of the material encountered by these particles, as shown in Figure 2.9. Given the necessarily large size and mass of these objects, practical

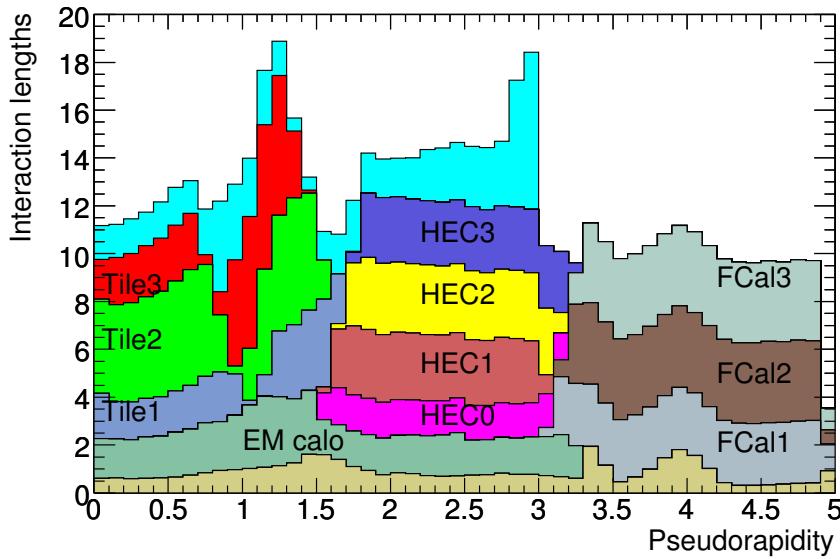


Figure 2.9.: The distribution of material (measured in interaction lengths) in the ATLAS detector up to the inner edge of the muon system. This includes all material between the interaction point and the calorimeters, such as the Inner Detector and solenoid. The majority of the material comes from the hadronic calorimetry systems, with the barrel tile layers, hadronic end-caps (HEC) and forward calorimeters (FCal) highlighted [3].

difficulties in constructing the detector had to be considered beside performance goals, leading to three different calorimeter designs.

The largest part of the hadronic calorimeter covers the region $|\eta| < 1.7$, using alternating layers of steel absorbers and polystyrene scintillating tiles. When hadronic shower particles ionise the polystyrene, recombining ion-electron pairs produce flashes of ultraviolet light. This light is guided to photomultiplier tubes through optical fibres that also act to shift the wavelength into the visible spectrum. The response of the scintillator tiles changes with radiation damage, so reference radioactive sources are used for calibration.

Larger particle flux at higher $|\eta|$ implies more rapid radiation damage, and polystyrene scintillator would degrade too rapidly to be of use. In these regions liquid argon is used

as the sensitive material instead, since it can be pumped out and replaced as needed. Ionisation of the argon is measured directly with electrodes at the edges of the argon volumes. A high voltage between these electrodes accelerates the ions towards them, creating a pulse of current.

The hadronic end-caps (HEC) cover the range $1.5 < |\eta| < 3.2$, with flat copper plates perpendicular to the beam pipe acting as absorbing material. The plates are separated by 8.5 mm gaps filled with liquid argon. Each of these gaps is split into 4 separate layers by the readout electrodes, so ions travel at most 1.8 mm.

The forward calorimeter (FCal) in the range $3.1 < |\eta| < 4.9$ experiences even higher particle flux. Smaller gaps between electrodes (0.3 – 0.5 mm) are used here to draw ions quickly out of the argon and prevent them from building up. This is achieved using coaxial copper tubes and tungsten rods as electrodes, with the rod diameter only slightly smaller than the inner diameter of the tube that encloses it. These electrodes are arranged parallel to the beam pipe in holes through the bulk absorber, also made of tungsten.

Position information for energetic particles is given by the segmentation of the readout electrodes or scintillator tiles in η and ϕ . In the central rapidity range $\Delta\eta \times \Delta\phi = 0.1 \times 0.1$ is achieved, but in the forward region $2.5 < |\eta| < 3.2$ this is reduced to 0.2×0.2 [3].

Table 2.2 compares the design goals for the hadronic calorimeter jet resolution with measurements using charged pion test beams with known energy. The resolution is parameterised with a constant term and a stochastic term that gets smaller as deposited energy increases. This relationship comes from the showering behaviour in the calorimeter: larger energy deposits mean more particles in the shower, and so more measurements contributing to the total shower energy.

	Tile barrel		Liquid argon end-caps		Forward	
	Design	Test beam	Design	Test beam	Design	Test beam
Stochastic ($\times \sqrt{E}$)	50%	56%	50%	71%	100%	70%
Constant	3%	5.5%	3%	5.8%	10%	3.0%

Table 2.2.: The design energy resolution of the ATLAS hadronic calorimeter, compared with pion test beam results for the tile barrel and hadronic end-caps [3].

Note that while the tile barrel and liquid argon end-cap modules tested did not meet their design goals in isolation, the performance of the combined system is satisfactory. Stochastic and constant terms of 52% and 3% respectively were achieved by combining tile barrel and electromagnetic calorimeter measurements.

2.2.5. Electromagnetic calorimeter

The electromagnetic calorimeter systems are the next closest layer to the interaction point, with the intention of fully absorbing electromagnetic showers before they reach the hadronic calorimeter. To that end, the total depth of these systems is at least $22X_0$, compared to only $\sim 2\lambda$.

Like the high-pseudorapidity hadronic calorimeter regions, the electromagnetic calorimeter uses liquid argon as the sensitive material. The electromagnetic FCal is of essentially identical design to the hadronic FCal, with tube and rod electrodes parallel to the beam pipe through bulk absorbing material. In this case the absorber is copper rather than tungsten (in the bulk and the rods), prioritising heat transfer rather than interaction length.

The region $|\eta| < 3.2$ uses a different calorimeter structure, with accordion-shaped lead absorbers separating the liquid argon volumes. This design was chosen to avoid creating absorbing or support structures parallel to the radial direction. Such structures would create regions in the calorimeter where particle energy is not measured (compromising the ability to detect missing energy).

Readout of the electromagnetic calorimeter is largely segmented into regions of $\Delta\eta \times \Delta\phi = 0.025 \times 0.025$ to give the positions of energetic particles. The first layer of the calorimeter for $|\eta| < 1.8$ is split into thin strips ($0.025/8 \times 0.1$) rather than square cells. At higher $|\eta|$ the segmentation is coarser, up to 0.1×0.1 for $2.5 < |\eta| < 3.2$. Multiple readouts are also grouped together into lower-resolution ‘towers,’ providing a rapid summary of energy deposition for the trigger system (see Section 2.3.2). The absorber structure and readout segmentation is shown in Figure 2.10.

Table 2.3 compares the design requirements for the electromagnetic calorimeter with the resolution measured using electron beam beams with known energy. As for the hadronic calorimeter, this resolution is parameterised with a stochastic and a constant term.

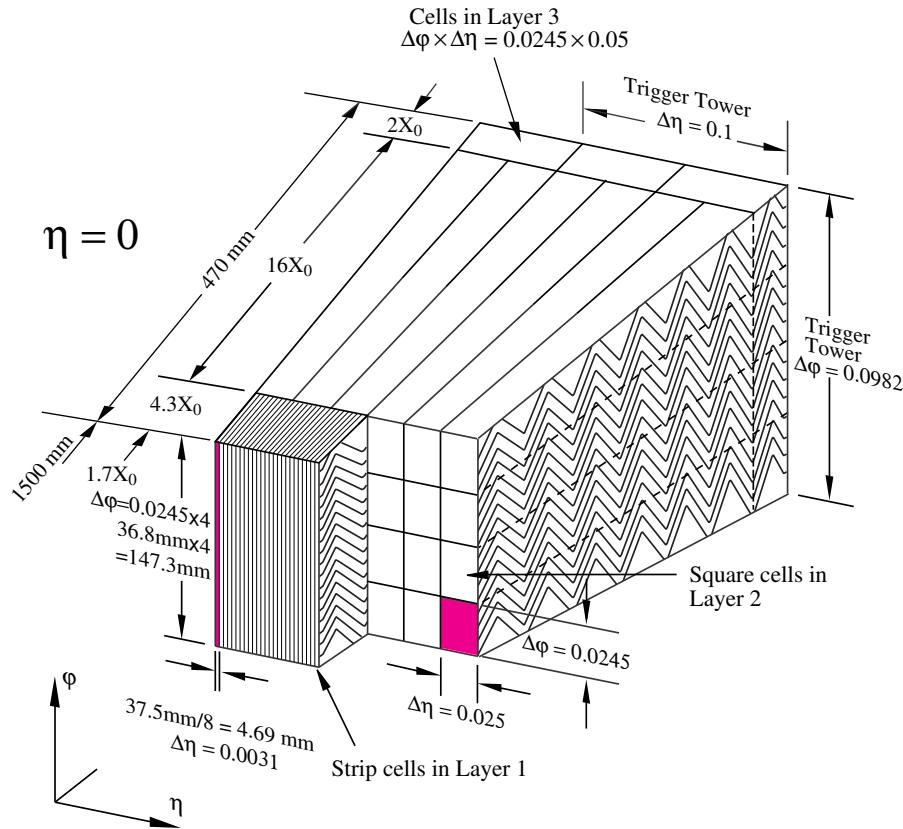


Figure 2.10.: The accordion shaped structure of the lead absorber material in the ATLAS electromagnetic calorimeter barrel. The depth of the calorimeter in radiation lengths X_0 and segmentation of readout electrodes in $\Delta\eta$ and $\Delta\phi$ are shown [3].

Liquid argon barrel		
Design	Test beam	
Stochastic ($\times \sqrt{E}$)	10%	10.1%
Constant	0.7%	0.2%

Table 2.3.: The design energy resolution of the ATLAS electromagnetic calorimeter, compared with electron test beam results [3].

2.2.6. Inner Detector

The Inner Detector is so called because it is the closest to the interaction point. It is a tracking detector, designed to measure the paths of passing particles. Since it is of greater relevance to this thesis than the other subdetectors it is worth discussing in detail — this is covered in Chapter 3. Like the muon tracker it is designed to measure the curvature of charged tracks in a magnetic field, allowing their momentum to be calculated. The Inner

Detector has the highest position resolution of all the ATLAS detectors, allowing vertices to be identified where multiple tracks originated. A magnetic field of 2 T is provided by a solenoid surrounding the Inner Detector.

Given the proximity to the interaction point the detector must be able to cope with a much greater number of particles per unit volume than other detectors, so it is designed with high granularity and a particular focus on resistance to radiation damage. It is also intended to place minimal material in the path of particles, since any interactions with the Inner Detector will affect their energy before they reach the calorimeters.

2.3. ATLAS infrastructure

The detector hardware described above is not only large and complex in itself, it also requires a large infrastructure to support it. Rather than give a complete summary of the topic, this section lists some example infrastructure needs, and goes into more detail in two specific areas of relevance to this thesis.

Support structures: The subdetectors and magnets weigh thousands of tonnes, they must be precisely positioned with respect to the beam pipe, and must not sag over time.

Power supplies: All the electrical components require power to operate, including high voltages to accelerate ions in the muon system and liquid argon calorimeters, and high currents for the toroids and solenoids.

Cryogenics: The magnets are cooled to 4.5 K with liquid helium, and the liquid argon for the calorimeters must be maintained at 80 K. There are also tanks near the detector to store the liquid argon in case of a fault.

Detector cooling: The electrical power dissipated in the detector produces several hundred kW of heat, which must be removed to keep the subdetectors at their operational temperatures. This is discussed further in Section 3.7.

Gas supplies: The different gas mixtures required by the subdetectors (e.g. in the four muon system technologies) must be maintained and monitored.

Data readout: Collisions occurring at a rate of around 40 MHz produce outputs on $\sim 10^8$ output channels, and all these data must be read out rapidly to the computers that process it.

2.3.1. Detector Control Systems

The ATLAS detector is not a single-state machine: some of the systems have multiple modes of operation, and variables that can be adjusted within those modes. In a simplistic sense the detector can be switched on and off, to record data from particle collisions or stand idle when there are none. Since ‘turning the detector on’ is in fact a complex and time-consuming process there also exist standby states where the detector is not yet recording data, but is ready to start as soon as collisions begin. It must be possible to change the definitions of these states since the performance of the detector will change with time as the components age. The definitions of states, and transitions between them, are controlled by a logical structure called a Finite State Machine (FSM).

In parallel with this, the status of the detector must be monitored to ensure that everything is as behaving as expected, and to identify faults as efficiently as possible. Component temperatures are monitored to ensure their associated cooling systems are working, with automatic failsafe measures to shut them down if they overheat. The outputs of the many different power supplies are monitored so that engineers can be notified in the event of a blown fuse, or components shut down to protect from damaging electrical fluctuations. Automatic safety systems such as these — and many others — are combined with human monitoring to keep the detector running smoothly.

Controlling the detector and monitoring the sensors within is the task of the Detector Control Systems (DCS). A network of computers monitors the output of the sensors, and displays the results through a specially designed Graphical User Interface (GUI) in the ATLAS control room. Some of the automatic failsafe behaviour is also implemented by these computers, although the most basic protection is built into the hardware itself. The GUI allows shifters in the control room to issue commands to the detector: their commands are interpreted by the DCS software and relayed to the appropriate devices. In this way, an overall state affecting many different systems can be selected without having to command each component individually, although fine-grained controls are available if needed.

2.3.2. Trigger

The designed rate for particle collisions in ATLAS is 40 MHz, and a single event requires approximately 1.3 MB of storage [3]. Therefore storing information from every collision is completely impractical, both in terms of the amount of storage required and the

computational power needed to subsequently analyse this amount of data. Instead of recording each event immediately, they are first passed through the trigger system. The purpose of the trigger is to rapidly identify interesting features in an event, or discard it if there are none. A hierarchy of triggers — level one, level two, and the event filter — reduces the rate of events to 75 kHz, then 3.5 kHz, and finally to 200 Hz. Each level uses progressively more information about, and analysis of, the event, and increasingly strict criteria for acceptance.

The level one trigger combines three sources of information:

- The RPCs and TGCs of the muon system identify high- p_T muons.
- Calorimeter towers give a low-resolution map of energy deposited in the calorimeter.
- Attached to the inner faces of the calorimeter end-cap cryostats are scintillator tiles similar to those used in the hadronic calorimeter barrel. These cover the rapidity range $2.09 < |\eta| < 3.84$, and are separated into two units in pseudorapidity and eight in azimuth around the beam pipe. Called the Minimum Bias Trigger Scintillators (MBTS) they are intended to indicate whenever a proton-proton collision occurs in the ATLAS detector.

Regions of the detector containing a high- p_T muon or large calorimeter energy deposit are termed Regions Of Interest (ROIs), and events containing at least one are passed to the level two trigger. The level one trigger can also produce a representative sample of all events in the detector using the MBTS information: this selection has the smallest possible bias, hence the name.

The level two trigger can access the full detail of measurements from ATLAS, but to save time only does so in the ROIs identified by level one — about 2% of the total information about an event. With the greater detail it can apply more precise selections to these events, then passes those still deemed interesting to the event filter. Since the readout rate is already much reduced, the event filter uses the full detector information about a given event to decide whether it should be stored or not. This process takes around 4 s per event, compared to 40 ms for level two and 2.5 μ s for level one.

Each event is tested against a number of different criteria, since what is interesting to one analysis group may not be interesting to another. An event may pass many different trigger selections at the same time: it is stored if it passes at least one. A particular analysis will then use events that pass one or more relevant trigger selections. Since some of these selections will necessarily have less stringent criteria than others, but the overall

maximum output rate must not be exceeded, some trigger selections are ‘prescaled’. This means that even if an event passes the criteria for that selection it will have a random chance of being marked a failure anyway. For example, a trigger selection with a prescale of 10 means that there is only a 1/10 chance that an event meeting the criteria will be passed to the next level. The minimum bias selection in particular has a large prescale.

Chapter 3.

The ATLAS experiment: Inner Detector

The ATLAS Inner Detector is designed to track the paths of charged particles leaving the interaction point. As the subdetector closest to the centre of ATLAS it must have the highest resolution and radiation tolerance. It must also have the least possible effect on passing particles so that it does not bias the measurements made by the other subdetectors. The Inner Detector comprises three different systems: from the outside inwards they are the Transition Radiation Tracker (TRT), the Semi-Conductor Tracker (SCT) and the Pixel detector. Figure 3.1 shows how these systems are laid out.

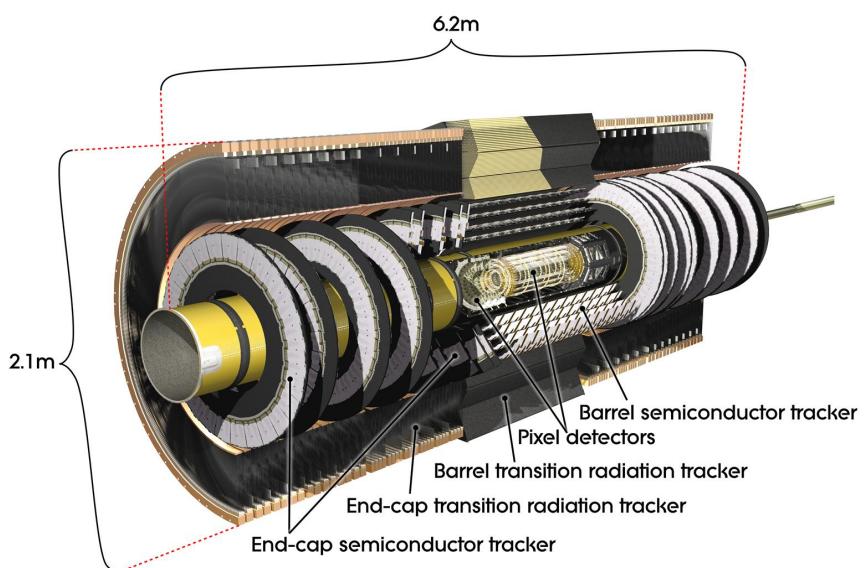


Figure 3.1.: Diagram of the ATLAS Inner Detector, showing the TRT, SCT and Pixel detector [19].

3.1. Momentum measurement

Besides simply identifying the presence of charged particles, tracking their path through the detector allows their momentum to be determined. When a magnetic field is applied the paths of the particles are bent, with the radius of curvature proportional to momentum. The Inner Detector is within a solenoid providing a 2 T field parallel to the beam pipe. The direction of curvature will indicate whether the charge is positive or negative, and it is assumed that stable charged particles in ATLAS have charge of ± 1 . Therefore the momentum can be calculated from a measurement of the radius of curvature. This requires at least three measurements of track position, i.e. three detector layers.

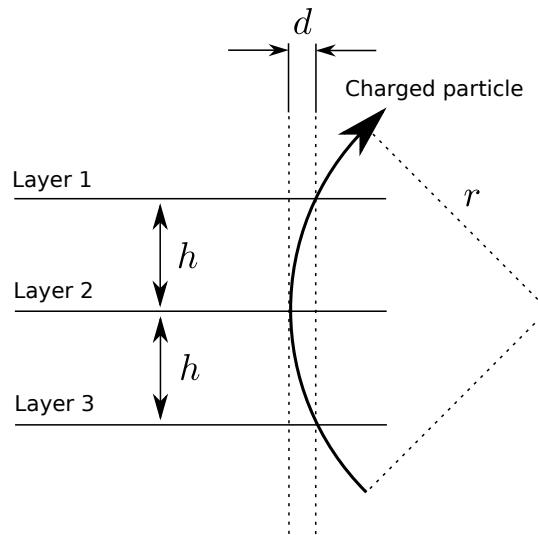


Figure 3.2.: The geometry of measuring the momentum of a charged particle, where r is the radius of curvature, h is the distance between detector layers, and d is the deviation of the track from a straight line as determined by the central position measurement — the sagitta of the arc.

For example, with a simplistic geometry as shown in Figure 3.2, the three track position measurements can be combined to find the radius of curvature as follows:

$$r^2 = (r - d)^2 + h^2$$

$$r = \frac{1}{2} \left(d + \frac{h^2}{d} \right)$$

Using Equation 2.4, the relationship between particle momentum and detector geometry is:

$$p = qB \left(\frac{1 + h^2}{2d} \right) \quad (3.1)$$

Examining this equation helps understand the limits on momentum measurement. Accurate measurement of a high momentum track requires well-separated detector layers (large h) with high position resolution (to measure small d values). Conversely a low momentum track requires detector layers to be close together (intuitively, the track will curve tightly upon itself and the detector layers must fit inside this). These requirements apply to measurements in the bending direction for the particle tracks. For the Inner Detector the solenoid field is parallel to the beam pipe, so the bending direction is ϕ .

3.2. Vertex measurement

Once charged particle tracks have been identified and measured, a second task must be undertaken: associating them with a proton-proton collision. Projecting tracks back into the beam pipe should identify a common origin, or ‘vertex,’ where (presumably) the collision occurred. The distance of closest approach of a track to its vertex is called the impact parameter, and a large impact parameter may indicate a poorly-measured or fake track.

This process is complicated by the possible existence of more than one vertex per event. As mentioned in Section 2.1.2, the proton beams consist of bunches of $\sim 10^{11}$ protons, and so there are usually multiple interactions in a single bunch-crossing. The ATLAS detector was designed with the expectation of (on average) 23 additional interactions with any event passing trigger selection [20], assuming the design luminosity $10^{34} \text{ cm}^2 \text{s}^{-1}$ and bunch spacing of 25 ns. These extra interactions are known as ‘pile-up,’ and each one can create an additional vertex in the detector. The Inner Detector must be able to associate tracks to the vertex they came from, requiring high position resolution in the z -direction. Additionally, long-lived but unstable particles such as B-mesons may travel some distance into the detector before decaying, creating a secondary vertex that must be distinguished from proton-proton collision (primary) vertices.

3.3. Pixel detector

The Pixel detector is the smallest of the ATLAS subdetectors, and is closest to the interaction point – in fact it is attached to the beam pipe. Despite the small size it has around 80 million readout channels, several times more than all of the rest of the detector combined. Each of the 80 million channels is associated with a single pixel: a $50 \times 400 \mu\text{m}^2$ section of a silicon wafer. This is because the Pixel detector is designed to provide the highest possible resolution for tracking charged particles, in order to associate tracks accurately to primary and secondary vertices. Tracks can be associated to a primary vertex with a transverse impact parameter (distance of closest approach to the vertex in the $x - y$ plane) resolution of better than $15 \mu\text{m}$, and in the z -direction better than 1 mm.

The active area of the detector is $\sim 1.7 \text{ m}^2$, split into 1744 modules. Each module has a $6.08 \times 1.64 \text{ cm}^2$ active area of silicon – comprising 47232 pixels [21] – as well as readout electronics and power connections. The modules themselves are identical, but the restricted space available for readout connections means that about 10% of pixels per module are larger than normal ($50 \times 600 \mu\text{m}^2$), or are grouped together with a single readout. This reduces the number of readout channels per module to 46080. Like the other ATLAS subdetectors, the modules are arranged into barrel and end-cap regions, and there are three pixel layers in each case. The barrel layers are at radii of 50.5, 88.5 and 122.5 mm, and the end-cap disks at $|z| = 495, 580$ and 650 mm , as shown in Figure 3.3. Pixel modules are oriented so that the $50 \mu\text{m}$ pixel edges (and hence the best resolution) are in the $R - \phi$ direction in both barrel and end-caps.

Each pixel is a reverse-biased diode: the voltage across it is in the wrong direction for current to flow. The voltage is large enough (initially 150 V) that the silicon is depleted, i.e. there are no free electrons/holes. A charged particle passing through the silicon produces electron-hole pairs, leading to a pulse of current as they move in the electric field and the silicon is depleted again. Around 2×10^4 pairs are produced by a minimum ionising particle, depending on the operating conditions [21].

As the closest subdetector to the beam pipe, the Pixel detector encounters the greatest flux of particles. At LHC design luminosity, the innermost pixel layer is expected to encounter 10^{15} neutron-equivalent particles per square centimetre in 5 years. This high occupancy is another driver of the fine resolution of the detector, since it must be possible to distinguish multiple nearby tracks. In one second of data taking, $\sim 10^8$ charged track

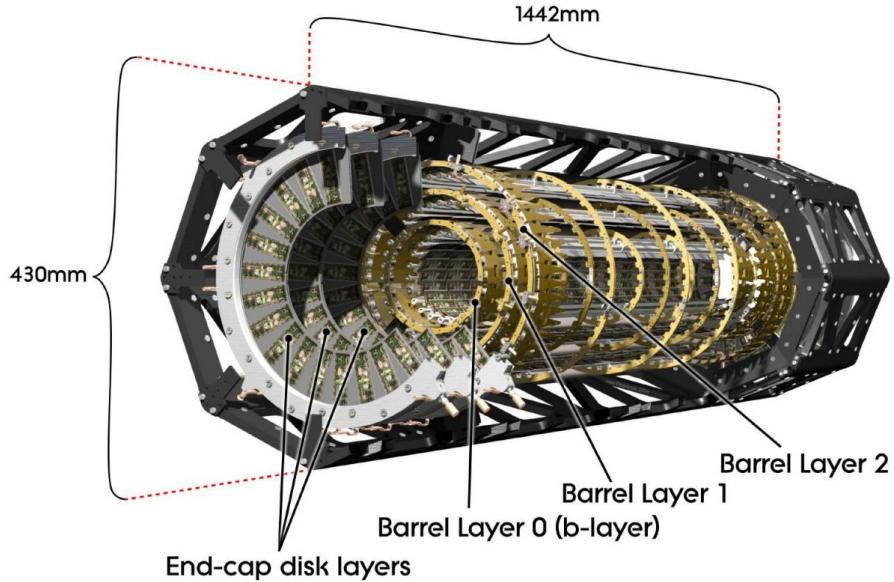


Figure 3.3.: Diagram of the ATLAS Pixel detector, showing the different layers of modules [22].

hits are expected for each square centimetre of the innermost Pixel detector layer [21]. This extreme particle flux also damages the silicon itself, as discussed in Section 3.6.

On top of all these constraints, the Pixel detector must place minimal material in the path of particles, since any interactions with the Pixel detector will affect the measurements made by outer subdetectors. To this end, the silicon wafers themselves are only $256 \pm 3 \mu\text{m}$ thick, and the attached printed circuit is only $100 \mu\text{m}$ in addition (readout chips add another $180 \mu\text{m}$ in some areas). All support structures are made of carbon fibre composite, chosen for its low density and high structural strength. The final material distribution with pseudorapidity is shown in Figure 3.4. Compare this with the hadronic calorimeter (Figure 2.9) where the intention is to absorb particles.

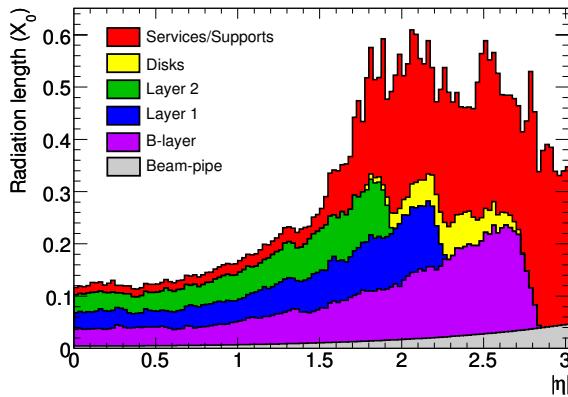


Figure 3.4.: The distribution of material in the Pixel detector, measured in radiation lengths.
Adapted from Reference [21].

3.4. Semiconductor tracker

The SCT encloses the Pixel detector, and uses similar technology: the detection of charged particles by pulses of current in reverse-biased silicon diodes. Given the greater distance from the interaction point the performance requirements of the detector are less stringent. The active area of the detector is also much larger (63 m^2 compared to $\sim 1.7\text{ m}^2$), imposing budgetary constraints. With this in mind lower resolution sensors are used, with strips $80\text{ }\mu\text{m}$ wide [3] running the length of the silicon. Sensors are paired back-to-back at a 40 mrad stereo angle to provide position information along the length of the strips. With these specifications, the nominal resolution achieved by a single module is $17\text{ }\mu\text{m}$ perpendicular to the strips, and $580\text{ }\mu\text{m}$ parallel to them. Like the Pixel detector, the strips are oriented to give the best resolution in the $R - \phi$ direction for both the barrel and end-caps.

There are 4088 SCT modules, again split between barrel and end-cap regions. The 4 barrel sensor cylinders are at radii of 299, 371, 443 and 514 mm, and all use double-layers of silicon as described above. The 9 end-caps are centred at $|z| = 853.8, 934.0, 1091.5, 1299.9, 1399.7, 1771.4, 2115.2, 2505.0$ and 2720.2 mm. Due to geometrical constraints some of the end-cap modules are irregular shapes, are not double-layered, or are displaced in z from the nominal position of the end-cap. This layout is shown in Figure 3.9. Note that the SCT end-caps extend to the extremes of the Inner Detector volume in z .

Given the large area covered by SCT modules, careful attention must be paid to their relative alignment to take advantage of the good intrinsic resolution of a single module. Using a robot to mount the modules on the carbon fibre support structure

allowed placement precision of $60\text{ }\mu\text{m}$ (measured in the z -direction) to be achieved in construction. However, the heat dissipated by the modules, and the cooling system to compensate for it (see Section 3.7) mean that significant thermal distortions of the support structure can occur during normal operation. To account for this, a system of laser interferometers constantly measures distances between reference points in the support structure. This frequency scanning interferometer system is able to measure shape changes of around $10\text{ }\mu\text{m}$, which uncorrected would degrade tracking performance by 20% [23].

3.5. Transition Radiation Tracker

Between the SCT and the Inner Detector solenoid is the TRT, which measures ionisation from charged tracks in gas volumes rather than silicon wafers. In this respect it has more in common with the ATLAS muon system than the SCT or Pixel detector. The gas is contained in kapton tubes 4 mm in diameter, with a $5 - 6\text{ }\mu\text{m}$ graphite-kapton surface layer protecting a $0.2\text{ }\mu\text{m}$ thick aluminium cathode [24]. Down the centre of each tube (or ‘straw’) is a gold-plated tungsten wire anode of diameter $31\text{ }\mu\text{m}$. The tubes are reinforced with carbon fibre, reducing their expansion with heat or humidity and allowing them to support the TRT structure.

The gas mixture in the TRT straws is 70% Xe, 27% CO_2 and 3% O_2 . With a 1530 V potential difference, the design signal gain in the gas of 2.5×10^4 is achieved [24]. To protect the gas from contamination, the pressure in the straws is maintained at $5 - 10\text{ mbar}$ above that of the CO_2 envelope that surrounds them. The gas is also recirculated through a system which measures its quality and filters out impurities.

Straws are arranged in multiple layers parallel to the z -direction in the TRT barrel, or the R -direction in the end-caps. They give no tracking information along their length, although the barrel straws are split at $z = 0$ to reduce occupancy. There are 73 layers of straws in the barrel, and 160 in the end-caps, but straws are held apart by the supporting structures. Accounting for the gaps between straws, a charged particle with $p_{\text{T}} > 0.5\text{ GeV}$ and $|\eta| < 2.0$ will encounter around 36 straws (min. 22) [25]. This is still a large number of hits compared to the silicon detectors, and so aids the identification of tracks and the rejection of fakes.

Measuring the drift time of ions in the gas mixture (and thus the distance of the charged track from the anode) gives each straw an intrinsic resolution of $130\text{ }\mu\text{m}$ in the $R - \phi$ direction. This extends precision tracking of charged particles in the magnetic field bending direction up to a radial distance of around 1 m, contributing to the precision of momentum measurement. The effect of the TRT is approximately equivalent to making a track position measurement at $R = 85\text{ cm}$, with $R - \phi$ resolution of $50\text{ }\mu\text{m}$ [26]. As described in Section 3.1, this measurement at large R (h in that notation) is particularly useful when measuring high-momentum tracks. For example, a muon with $p_{\text{T}} = 500\text{ GeV}$ and $|\eta| < 2.0$ will have its momentum measured by the Inner Detector with 16% precision, compared to 36% without the TRT. Figure 3.10 shows the effect of the TRT on various Inner Detector measurements: again, note the improvement in momentum resolution at high p_{T} .

The TRT also provides the ability to identify electrons through transition radiation. A relativistic charged particle radiates photons when it crosses a boundary between two media of different dielectric constants. More relativistic (higher γ -factor) particles are more likely to radiate, distinguishing the lighter electrons from heavier hadrons when the momentum measurement is included. The x-ray photons radiated by electrons in the TRT are absorbed by the xenon gas, creating signals with much larger amplitude than those from ionisation along charged tracks [3].

3.6. Radiation damage

The Inner Detector will be subjected to extremely high levels of radiation, with the innermost pixel layer encountering 10^{15} neutron-equivalent particles per square centimetre in 5 years of LHC operation at design luminosity. Although the flux is less for the larger components, the SCT inner layers will still encounter 2×10^{14} neutron-equivalent particles per square centimetre in 10 years [3]. The TRT is resistant to damage, but for the silicon sensors it presents a problem, as the radiation introduces defects into the silicon. These are effectively intermediate energy levels in the semiconductor band-gap, so can trap deposited charge from particle tracks and allow ‘leakage current’ to flow through the sensor.

Charge trapping in radiation-damaged silicon leads to an increase in the voltage required for depletion. Oxygenated silicon is used for the pixel sensors because it is more resistant to this effect than pure silicon, but nonetheless the initial 150 V operating

voltage will be increased as needed up to a design maximum of 600 V [21]. The effect is less severe in the SCT given the lower particle flux, but still 250 – 350 V will be required for normal operation after 10 years irradiation.

Increasing leakage current means there will be a larger background to the current pulses that indicate a hit from a charged particle. This background must be accounted for by the readout electronics. It also means that the silicon sensors will draw more power, and hence produce more heat.

The effect of radiation damage shows a temperature-dependent time evolution. In the short term, annealing reduces the density of defects in the silicon, as expected in a naive model. However, continued annealing increases their density again – behaviour referred to as ‘reverse annealing.’ This effect is well measured, but the mechanism behind it is not well understood (attempts have been made to model it, e.g. Reference [27]). This is the dominant effect at large timescales, eventually causing the defect density to surpass the initial effect of the radiation. An example of annealing in a silicon sensor is shown in Figure 3.5.

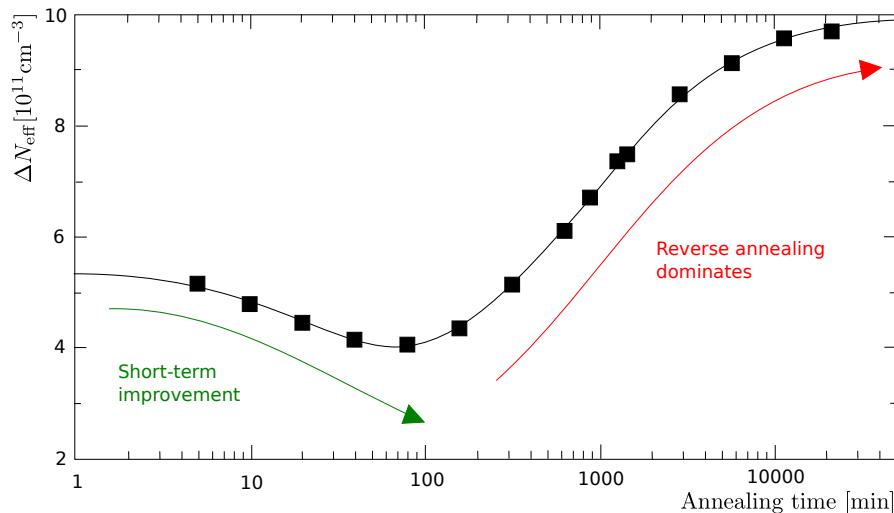


Figure 3.5.: An example of the evolution of silicon defects with time, adapted from Reference [28]. A test detector was irradiated with a neutron fluence of $1.4 \times 10^{13} \text{ cm}^{-2}$, then annealed at 60 °C. The change in effective doping concentration, ΔN_{eff} , was measured over the course of the annealing.

The density of defects in the silicon can be measured indirectly by finding the depletion voltage of the sensor, and from that calculating the effective doping concentration, N_{eff} . Any change in N_{eff} from when the sensor was first manufactured must be due to the introduction of defects in the silicon. Reverse annealing has an effect on N_{eff} that can be

parameterised [28] as:

$$\Delta N_{\text{eff}}(t) \sim \Delta N_{\text{eff}}(\infty) \left(1 - e^{-\frac{t}{\alpha}}\right) \quad \text{for large } t,$$

$$\ln(\alpha) \propto \frac{1}{T},$$

where t is the time elapsed, and T is the annealing temperature. Note the temperature dependence of the reverse annealing time constant, α , indicating that lower temperatures lead to slower annealing. On the timescale of LHC detector operations the beneficial effects of annealing can be ignored, and reverse annealing dominates. Therefore the silicon detectors are kept cold whenever possible, including outside data-taking periods, to slow the annealing process. An operational temperature of -7°C [3] was chosen, leading to reverse annealing $\sim 10^3$ times slower than at room temperature [28]. Over the lifetime of the ATLAS detector this will greatly reduce the effect of radiation damage.

3.7. Cooling

Overall $\sim 85\text{ kW}$ of heat must be removed from the Inner Detector enclosure during normal operations. The Pixel detector and SCT are expected to dissipate about 5 W of power per module, and yet these modules have operating temperatures of -7°C to protect from radiation damage. Taking into account the heat transfer through the modules to the coolant pipes, a -7°C silicon temperature requires a coolant temperature of around -25°C [29]. The requirements for the TRT are less stringent as it is operated at room temperature.

An evaporative cooling system was chosen for the SCT and Pixel detector, because using the latent heat of vapourisation of a liquid (rather than its specific heat capacity) reduces the mass-flow needed for a given cooling capacity. This means cooling pipes can be smaller, placing less material in the path of particles leaving the interaction point. Coolant is pumped into the detector as liquid at high pressure, and passes through capillaries that restrict its flow. Past the capillaries the pressure is allowed to drop below the gas phase boundary and the liquid boils, cooling the pipe and the detector modules attached to it. The temperature is regulated by controlling the pressure of the vapour as it leaves the system: lower pressure gives lower temperature. The coolant used is C_3F_8 , which is pumped in as liquid at $11 - 14\text{ bar}$ and leaves as vapour at 1.67 bar to cool the silicon to its operational temperature [29].

In ATLAS the cooling system includes a heat exchanger, using the cold vapour exhaust to cool the C_3F_8 liquid before it enters the detector. The vapour then passes through a heater, ensuring that all liquid has boiled before it enters the output pressure regulator. A schematic of this system is shown in Figure 3.6.

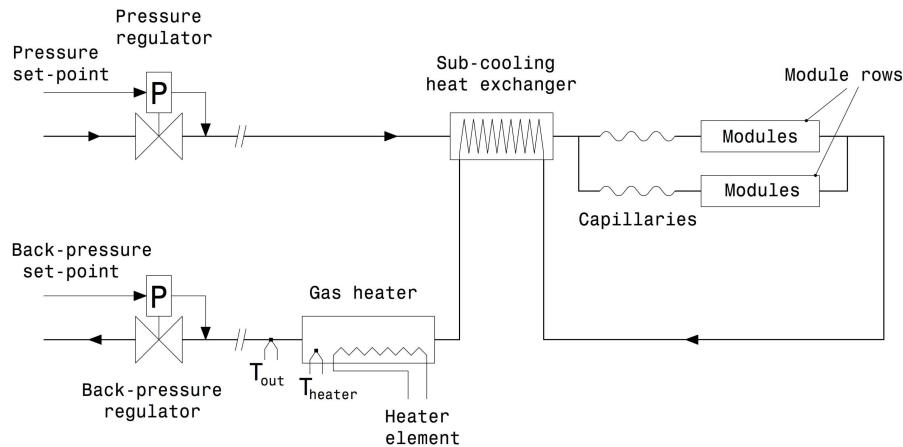


Figure 3.6.: Diagram of a cooling loop in the ATLAS Inner Detector [3]. The system is the same whether cooling the SCT or Pixel detector — simply indicated by the silicon ‘modules’ in this schematic. Liquid coolant enters through the pressure regulator (top left), boils after it passes the capillaries, and leaves through the back-pressure regulator (bottom left) which is used to control the temperature.

The TRT is cooled using a monophase system: no evaporation, just heat transported by a constantly flowing fluid with a particular specific heat capacity. In the barrel region C_6F_{14} is used, but in the end-caps temperature is regulated by forcing a higher flow rate in the CO_2 gas surrounding the straws: $50 \text{ m}^3\text{h}^{-1}$ compared to $\sim 3 \text{ m}^3\text{h}^{-1}$ in the barrel [3]. Thermal separation is needed between the room-temperature TRT and the much colder silicon detectors. Rather than use insulation — which would be too bulky for the limited space — electrical heater pads are attached to the outside of the SCT enclosure to warm the surface to 20°C . These are known as thermal enclosure heaters (TEH). Unfortunately the power supply connections to the heaters in the barrel region were damaged in early detector operations, and so the monophase cooling system is used at above design temperature to keep the TRT barrel warm. The outermost layer of the SCT barrel is also set to a higher operational temperature to protect the TRT.

3.8. Inner Detector DCS

The subsystems of the Inner Detector described in this chapter are all controlled and monitored by the DCS software first mentioned in Section 2.3.1. Temperature sensors on modules in the SCT and Pixel detector check the performance of the evaporative cooling system, and the coolant exhaust pressure can be adjusted to maintain the silicon at -7°C . The flow rate and pressure of gas in the different detector enclosures is monitored, and humidity sensors are used to ensure that the components remain dry. As the silicon accumulates radiation damage the voltage supplied to the modules can be adjusted to ensure the sensors are depleted. The TEH system is monitored to ensure that the TRT is thermally separated from the SCT.

Over the course of detector commissioning and operations in 2009 and 2010, numerous upgrades were made to the Inner Detector DCS. Drawing on the experience of running the detector, these changes were made to improve performance or aid fault identification and recovery. Besides small tweaks, improvements and bug fixes, four substantial changes of function were implemented, which are discussed below.

3.8.1. Distribution rack upgrade

Pumping liquid coolant into the detector for the evaporative cooling system is performed by a single compressor plant. However, inside the detector the system is split into 204 ‘cooling loops,’ each with its own pressure and back-pressure regulators as shown in Figure 3.6. These loops can be controlled individually, so that different regions of the detector can be switched on or off independently, or set to different temperatures. The common pipes to and from the compressor plant connect to the loops at ‘distribution racks.’ Loop pressure and back-pressure regulators are mounted on these racks, but there were not initially any sensors monitoring the system.

In 2009, pressure and temperature sensors were added to the distribution racks, allowing for the first time direct measurements of the input and output pressures rather than inferring them from temperature sensors inside the detector. Diagnostically this is useful, as problems in the coolant flow can be disentangled from the effect of detector components running unexpectedly hot or cold. The DCS GUI was modified to show the measurements from these sensors, with alerts if they deviate too far from the values fed to the pressure and back-pressure regulators.

These sensors also allowed a new diagnostic procedure called a ‘leak-down test.’ During long maintenance periods when the LHC beam is not present, the loops can be filled with coolant exhaust vapour and sealed by closing the manual valves on the distribution racks. The pressure is monitored over the course of several days to measure the leak rate of coolant into the detector, and verify that it is within acceptable bounds. Results of the leak-down test during the 2009-2010 maintenance period are shown in Figure 3.7. Leak rates are calculated separately for the pressure and back-pressure sensors, producing the pairs of values shown.

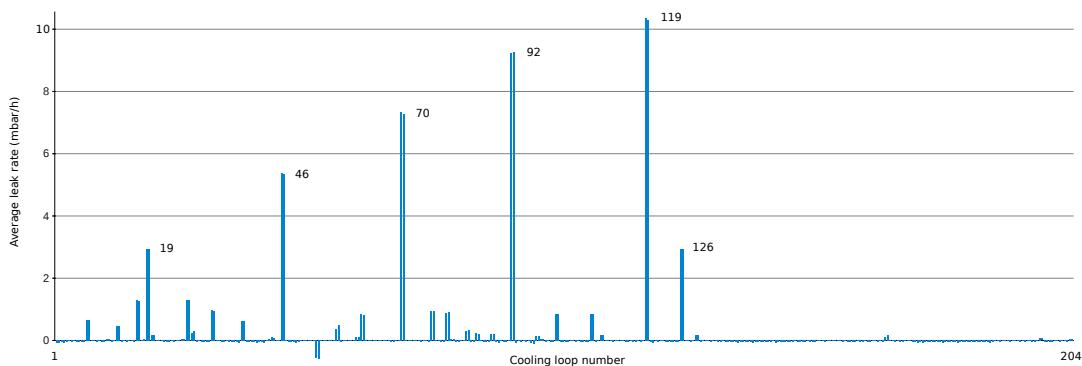


Figure 3.7.: Plot showing the average leak rates of all cooling loops for the leak-down test from January 7th – 13th 2010. Loops with leak rates over 2 mbar h⁻¹ are labelled.

3.8.2. TEH automatic recovery

The TEH pads are electric heaters, controlled by power supplies in a utility cavern away from the detector. These power supplies use information from temperature sensors to maintain a temperature of 20 °C at the SCT-TRT interface by toggling the heaters on and off. Combined with occasional fluctuations in the mains power supply or currents induced in the heaters as the ATLAS magnets ramp up and down, this varying load can cause heater power supply circuit breakers to trip.

The faster the recovery from a TEH trip the better, as this reduces temperature variation in the TRT. However, shifters were having difficulty identifying a trip due to the large number of separate warnings produced: one power supply controls multiple heaters, and each heater produces a separate warning for temperature and electrical current out of range. Further delay was introduced because the process of reactivating

the power supplies required expert intervention, and had to be performed separately for each one. New code was introduced to rapidly identify a trip, presenting a list of problems and proposed actions to the shifter, and requiring only a confirmation before sending the recovery commands automatically. This process is now routinely used, and expert intervention is no longer required.

3.8.3. SCT power-cut detection

The most crucial parts of the DCS infrastructure are protected by uninterruptible power supplies (UPS). However, most of the data-taking components of ATLAS are not: it is impractical given the huge power requirements. In particular, the power supplies to the silicon detector modules are not protected by UPS, but the cooling system is. This means that in a power cut, the SCT modules will stop producing heat, but will continue to be cooled.

In situations like this the cooling should be deactivated quickly, to prevent the Inner Detector from getting too cold (thermal contraction and expansion can damage components). To simplify this process, new software monitoring was introduced to create an alert if more than 3 SCT module power supply racks are without power. This is assumed to indicate a power cut, and so another monitoring process is started for the TEH temperature sensors. If any one of them registers a temperature below 14 °C for more than 60 s the SCT cooling is shut down automatically. This system has been successfully tested and used, both for detecting power cuts and responding to them.

3.8.4. SCT automatic switch-on

After beam is injected into the LHC it is accelerated and focussed before collisions begin. During this period of ‘unstable’ beams there can be a large background of charged particles in the ATLAS Inner Detector from the unfocused beams hitting collimators in the beam pipe. These are not useful data, and the background may be large enough to overload the readout electronics for the silicon detectors, so they are not powered. However, once the beams are stable and collisions have begun the detectors should be switched on as soon as possible.

The process of switching on the SCT manually is quite slow, and was found to be delaying the start of data-taking by several minutes. To improve this, a new method

was introduced whereby the shifter could pre-approve the decision to switch on the SCT, handing control to an automatic system. This system monitors beam background conditions, determines when collisions start, and then switches on the SCT if pre-set criteria are met. Additional inputs include the positions of collimators in the beam pipe near the ATLAS cavern, the positions of the beams within the beam pipe, and the status of the beam indicated by the LHC control room.

The decision process for declaring conditions safe to turn on the SCT is shown in Figure 3.8. At all times a simple process monitors the beam intensities, and the reported LHC beam mode, to determine if the beams are present. When the beams are present a more detailed monitoring process begins, looking for evidence of collisions and a low beam background. Readiness is indicated to the SCT when all of these tests are passed, but it can be withdrawn at any stage.

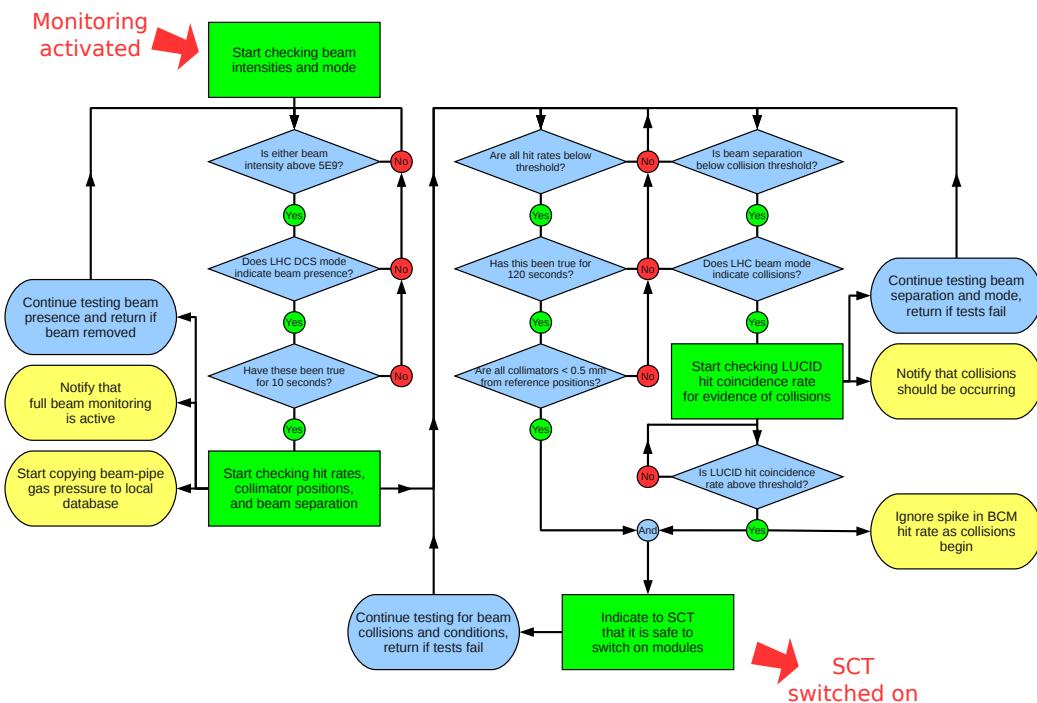


Figure 3.8.: Flowchart showing the criteria required for the DCS beam monitoring to approve automatic switch-on for the SCT.

Beam background is monitored using the ATLAS beam loss monitor (BLM), the beam conditions monitor (BCM) and “Luminosity measurement using Cherenkov Integrating Detector” (LUCID) [3]. The BCM and BLM are sets of small diamond sensors positioned close to the beam pipe (8 and 12 sensor modules respectively). They measure ionisation

from charged tracks just like the silicon sensors, but using diamond as the detector material makes them much more resistant to radiation damage. LUCID detects Cherenkov light emitted by charged particles travelling through tubes filled with C_4F_{10} – there are 20 of these tubes on each side of the detector, at $|z| = 17\text{ m}$ [3]. The SCT also produces low-precision hit rate information in its standby state, and this is used as well.

3.9. Combined Inner Detector performance

When all the subsystems of the Inner Detector are combined, they provide charged particle tracking with complete ϕ coverage extending to $|\eta| < 2.5$. The relative sizes and positions of the detector components are shown in cross-section in Figure 3.9

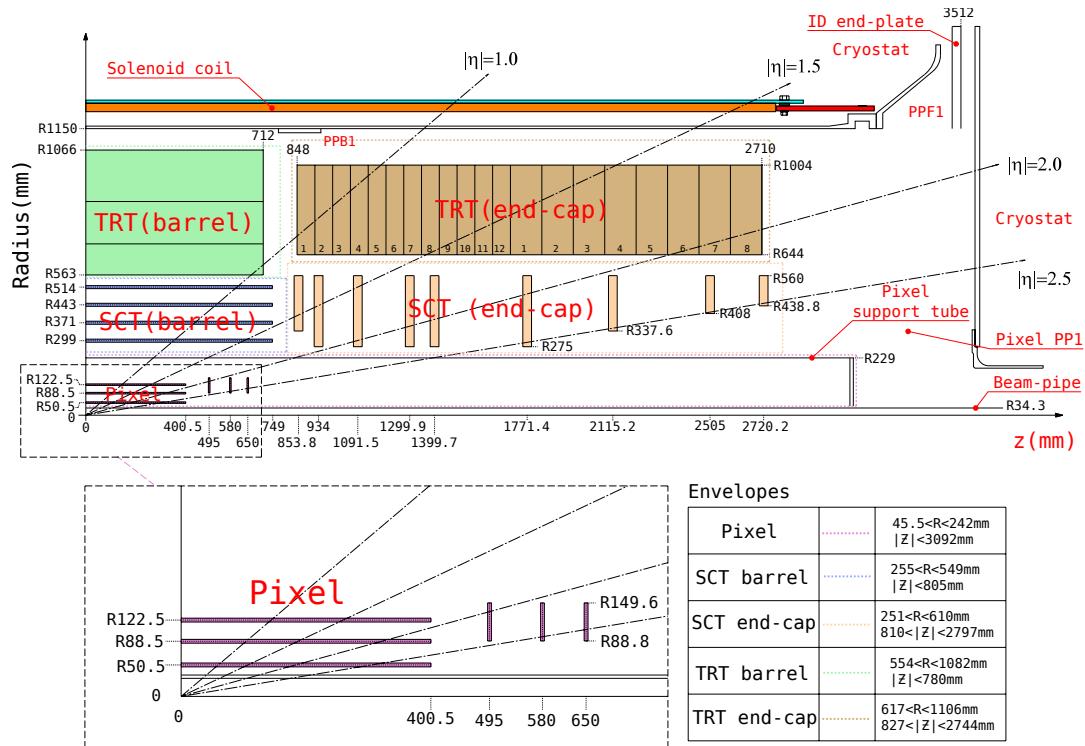


Figure 3.9.: A schematic cross-section through one quarter of the Inner Detector, showing the layout and relative sizes of the Pixel detector, SCT and TRT [3].

The positions of all of the detector components relative to each other must be accurately measured to take full advantage of their fine intrinsic resolution. This is an ongoing process of improvement and response to slow deformations of the detector – the results shown here are from the 2008 commissioning period [30], in order to best reflect the performance of the detector when collecting the data used in Chapter 6. Cosmic

ray muons can be used to measure detector alignment because they should appear as straight tracks through the detector (when the magnets are off). Since the majority of cosmic rays measured travel vertically downwards, additional data for the end-caps is provided by producing sprays of charged particles along the beam pipe by directing a beam into a collimator.

After alignment, the resolution achieved in measuring various track parameters was measured by splitting tracks as they cross the interaction point and analysing them separately. For any of the parameters under investigation, its value, T , is measured for each half of the track, and the difference in measurements ΔT is taken. The resolution of the parameter is the root mean squared value for the ΔT distribution divided by $\sqrt{2}$. Results of these measurements are shown in Table 3.1.

Parameter	Detector resolution (2008)	Ideal resolution
d_0 (μm)	22.1 ± 0.9	14.3 ± 0.2
z_0 (μm)	112 ± 4	101 ± 1
ϕ_0 (mrad)	0.147 ± 0.006	0.115 ± 0.001
θ (mrad)	0.88 ± 0.03	0.794 ± 0.006
q/p (GeV ⁻¹)	$(4.83 \pm 0.16) \times 10^{-4}$	$(3.28 \pm 0.03) \times 10^{-4}$

Table 3.1.: The resolution of the combined Inner Detector for measuring given parameters of tracks with $p_T > 30$ GeV [30]. The transverse and longitudinal impact parameters (d_0 and z_0) are the perpendicular distances of closest approach of a track to a vertex.

Track parameter resolutions can depend on the detector region and particle properties – the values in Table 3.1 are averages. Figure 3.10 shows how they vary with particle p_T . The angular and impact parameter resolutions get worse at low p_T due to multiple scattering in the beam pipe and Pixel detector.

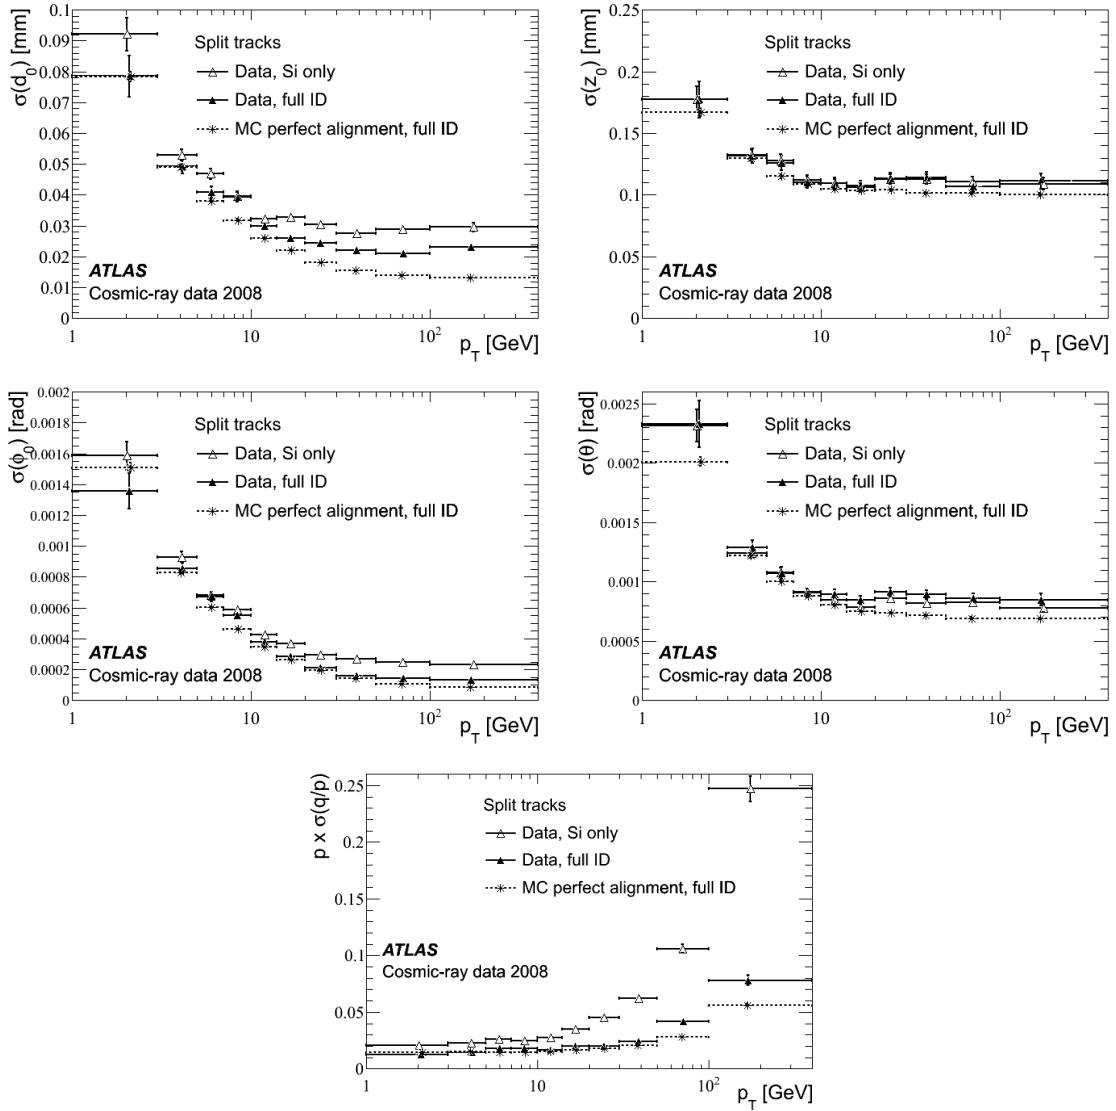


Figure 3.10.: The variation of the resolution of track parameters with particle p_T [30]. The performance of the SCT and Pixel detector are shown with and without contributions from the TRT, and both results are compared with Monte Carlo simulations using perfect detector alignment.

Chapter 4.

The Standard Model

Describing three of the four known forces and all particles that have thus-far been identified, the Standard Model of particle physics encompasses a vast swathe of experimental results and theoretical predictions. The Standard Model unites electromagnetism and weak interactions as the electroweak force [31–33], and then adds in the strong force (see Section 4.1) to model all particle interactions with the notable exception of gravity. For this reason, the Standard Model has been called “The theory of almost everything.” Figure 4.1 shows the particle content of the Standard Model, divided into the three generations of interacting fermions, and the gauge bosons (or ‘force-carriers’) that are exchanged in their interactions. The charged particles — all quarks, electrons, muons, taus and W-bosons — can interact through Quantum Electrodynamics (QED) by exchanging photons. Quarks can interact through Quantum Chromodynamics (QCD) by exchanging gluons, and all fermions can interact through the weak force by exchanging W or Z-bosons.

The Standard Model forces are described using $SU(3) \otimes SU_L(2) \otimes U_Y(1)$ symmetry¹, for the strong force, weak force, and weak hypercharge respectively. Electromagnetism arises from the breaking of the $SU_L(2) \otimes U_Y(1)$ symmetry. These groups represent degrees of freedom of the theory — transformations under which the interaction Lagrangian is invariant. Each degree of freedom gives rise to a boson: the eight gluons of QCD; the three weak bosons W^+ , W^- and Z^0 ; and the photon. Since the strong and weak force symmetries are non-Abelian (i.e., there are non-commuting members of the symmetry group), the corresponding bosons can interact with each-other.

¹Here L denotes left-handed particles, and Y denotes the weak hypercharge $Y_W = 2(Q - T_3)$

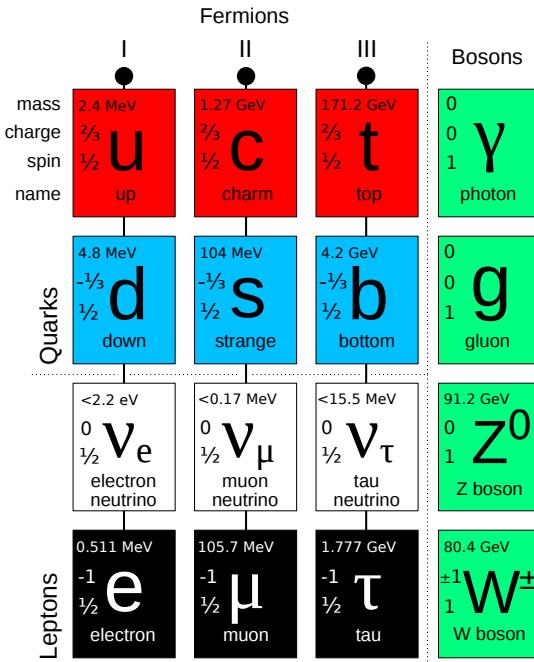


Figure 4.1.: Particle content of the Standard Model, adapted from Reference [34]. Fermions are divided into three generations, where each contains a positively charged (up-type) quark, a negatively charged (down-type) quark, a lepton and a lepton neutrino. Each fermion has an antiparticle partner, with the same properties except the opposite-sign electric charge. Bosons are the force-carrying particles for QED (photons), QCD (gluons) and the weak force (W and Z).

Combined with the mechanisms for particle interactions, the Standard Model also gives masses to the particles via the Higgs field. The fermions receive their masses by the strength of their coupling with the field, while the W and Z-bosons mix with degrees of freedom of the field itself [35–37] (the photon also mixes with this field, but remains massless). One remaining degree of freedom in the field should give rise to a new particle, the scalar Higgs boson [36, 38], which has been the subject of a prolonged experimental search. Recent observations of a new boson at LHC experiments — with mass $126.0 \pm 0.4(\text{stat.}) \pm 0.4(\text{syst.}) \text{ GeV}$ [39] or $125.3 \pm 0.4(\text{stat.}) \pm 0.5(\text{syst.}) \text{ GeV}$ [40] — suggest that this missing component of the Standard Model may have been found. Further analysis of the new boson candidate will determine whether it is compatible with Standard Model predictions.

The ATLAS detector is intended to search for other undiscovered particles besides the Higgs boson. One target for searches is a dark matter candidate: a massive, stable particle with no electric charge that could give rise to the non-luminous matter distributions

revealed on cosmological scales by gravitational effects [41]. Other possibilities include a fourth generation of fermions, or supersymmetric partners to existing particles.

New particles will most likely appear as either missing energy signatures or through their decays to known particles. In the former case it is necessary to have a detailed understanding of the particle background at the LHC, produced by pile-up or by multiple quark-quark interactions within a hadron collision (see Section 4.2.2). The latter case requires the decays of known particles to be well measured, so that additional or unusual decays can be identified. Either way, these searches rely upon Standard Model results, and in particular the theory of QCD.

4.1. Quantum Chromodynamics

With the discovery that the atomic nucleus contained many positively-charged protons came the realisation that there must be some new force to counter their like-charge repulsion. The force had to have similar strength as electromagnetism, so had to have a limited range in order to explain why it had not been observed outside the nucleus. This suggested a force-carrying particle with a significant mass: around 100 MeV [42]. Examining cosmic ray particles lead to the discovery of the pion as a potential candidate [43].

However, the discoveries did not end with the pion. Examining tracks from cosmic rays and colliders lead to the identification of a “zoo” of new particles — see Figure 4.2 for some examples. Much like the Periodic Table of Chemical Elements, attempts to classify the new particles hinted at a more fundamental structure [44]. Proton-electron inelastic scattering experiments also suggested that the proton was a composite particle [45]. These results are explained using the quark model [46], where mesons (like the pion) and baryons (like the proton) are bound states of two or three quarks respectively.

As with protons in the nucleus, a force is required to hold the quarks together. The discovery of baryons containing three quarks of the same flavour (such as the Δ^{++} and Ω^- [48]) suggested the need for a new quantum number, since at most two spin- $\frac{1}{2}$ particles can occupy the same state without violating Pauli exclusion [49]. This quantum number was given the arbitrary name ‘colour,’ although this is purely a label, and implies no connection to visible light. Eventually the model of quarks, the force that holds them together, and the behaviour of that force in inelastic proton scattering experiments, were united in the theory of Quantum Chromodynamics [50–52]. The prefix ‘chromo-’ refers

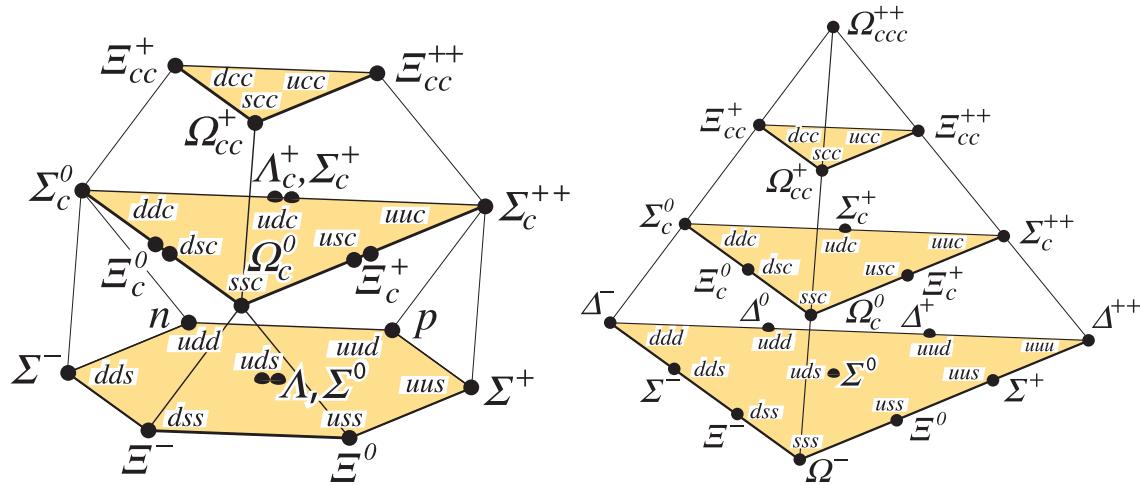


Figure 4.2.: Baryons with spin 1/2 (left) and 3/2 (right), composed of u, d, s and c-type quarks [47]. Note the baryons containing three quarks of the same flavour (e.g. $\Delta^{++}(uuu)$) which require an additional quark quantum number besides spin to avoid violating the Pauli exclusion principal.

to the colour quantum number, identified as the charge carried by particles that interact through QCD. Pion exchanges between nucleons can be interpreted as a consequence of QCD, rather than a fundamental force.

QCD is described using SU(3) group symmetry, with three equivalent colour charges (called red, green and blue), and eight force-carrying particles called gluons. The gluons each carry some combination of the colour charges — and anti-charges — meaning they can interact with each-other as well as with the quarks. This is in contrast to QED, where the photon has no charge and so cannot interact with other photons. The Feynman rules for QCD in fact allow two different gluon-only vertices, as shown in Figure 4.3. Gluons and quarks are the only particles to carry a colour charge: the other particles of the Standard Model do not interact through QCD.

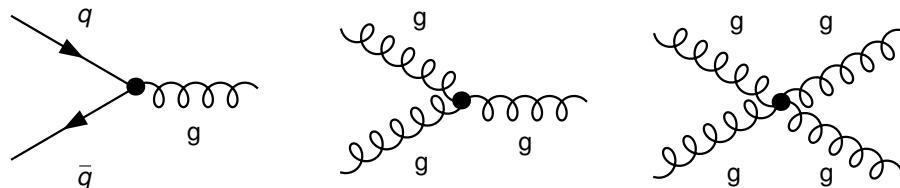


Figure 4.3.: Feynman diagrams of the vertices of QCD: qgg (left), ggg (middle) and $gggg$ (right) [53]. The vertices between three particles have a coupling of g_s , and the four gluon vertex has a coupling of g_s^2 .

It is important to note that quarks and gluons are never observed as free particles, and can only be treated as such in the limit of high energy interactions. This is discussed in more detail below.

4.1.1. The QCD coupling strength

The quantity $\alpha_s = \frac{g_s^2}{4\pi}$ is the QCD coupling constant, which governs the strength of the force or the likelihood of a particular interaction. In fact it is not constant at all (the term is historical), but is a function of Q^2 , the momentum transfer in an interaction. Figure 4.4 shows the variation of α_s , which is large at low Q^2 and decreases as the momentum transfer increases.

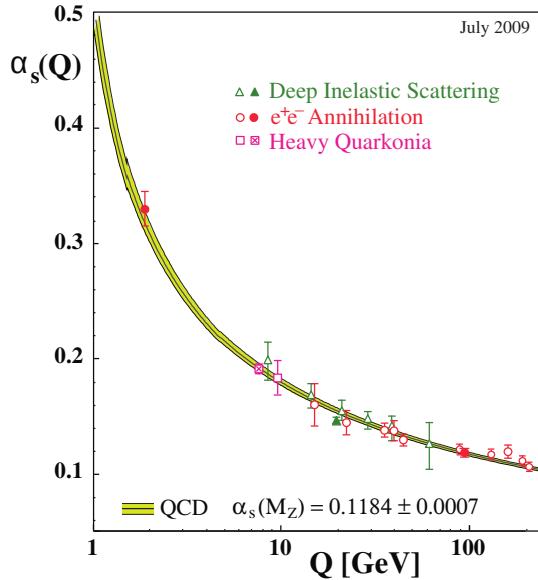


Figure 4.4.: The variation of α_s with Q , also called the running coupling of QCD [54]. Different data-point styles indicate the source of the result. Full symbols are results based on N3LO QCD, open circles are based on NNLO, open triangles and squares on NLO QCD. The cross-filled square is a computational result based on lattice QCD.

Two key features of QCD can be explained in terms of the variation of α_s . Firstly there is confinement: colour-charged quarks and gluons are never observed directly, but are always part of a bound state with no net colour. Since momentum and distance scales are inversely related, increasing the distance between colour-charged particles implies QCD interactions with lower Q^2 , and therefore higher α_s . In consequence, moving colour-charges apart increases the strength of the QCD force between them. Truly

separating the particles would require an infinite amount of energy — the effective QCD potential can be modelled as rising linearly with distance [55].

The second feature of QCD is asymptotic freedom: despite confinement, colour-charged particles can be treated as free when interacting with sufficiently high-momentum probes. This comes directly from the reduction in α_s with higher Q^2 . Without asymptotic freedom, quarks could not be considered as distinct objects — as free particles — and the whole mathematical basis of QCD would be in doubt.

Proton-proton scattering illustrates these two features well. When the transfer of momentum between the protons is small, elastic scattering occurs and the protons ‘bounce’ off each-other. Effectively, the protons are behaving as single particles: the quarks and gluons within them are confined. With higher momentum-transfer there is inelastic scattering, and the protons fragment. This is literally a breaking of the bound state, where momentum has been transferred to a quark, or gluon, within the proton, rather than the proton as a whole.

4.2. Phenomenology

The behaviour of α_s at low Q^2 creates a problem when trying to calculate the cross-section for QCD interactions. The cross-section for a particular process is proportional to the square of the matrix element for that process. The matrix element is calculated by adding all the contributions from all the Feynman diagrams that could be responsible for that process. Each diagram gives a contribution proportional to $\alpha_s^{n/2}$, where n is the number of vertices. When $\alpha_s \ll 1$, the matrix element can be evaluated using a perturbative approximation.

Perturbation theory models an interaction as a ground state (i.e., no interaction) with small corrections applied. The simplest diagrams for an interaction — called leading order (LO) — must have at least two vertices, and so give contributions proportional to α_s . More elaborate diagrams — called next-to leading order (NLO), next-to next-to leading order (NNLO) and so on — give contributions with larger powers of α_s . So, when $\alpha_s \ll 1$, the LO diagrams are a small correction, or perturbation, to the ground state, and NLO diagrams an even smaller correction. Consequently, calculating a matrix element to reasonable precision only requires that the simplest diagrams be calculated,

although greater precision can be gained by adding higher order contributions. Examples of LO and NLO diagrams for a process are shown in Figure 4.5.

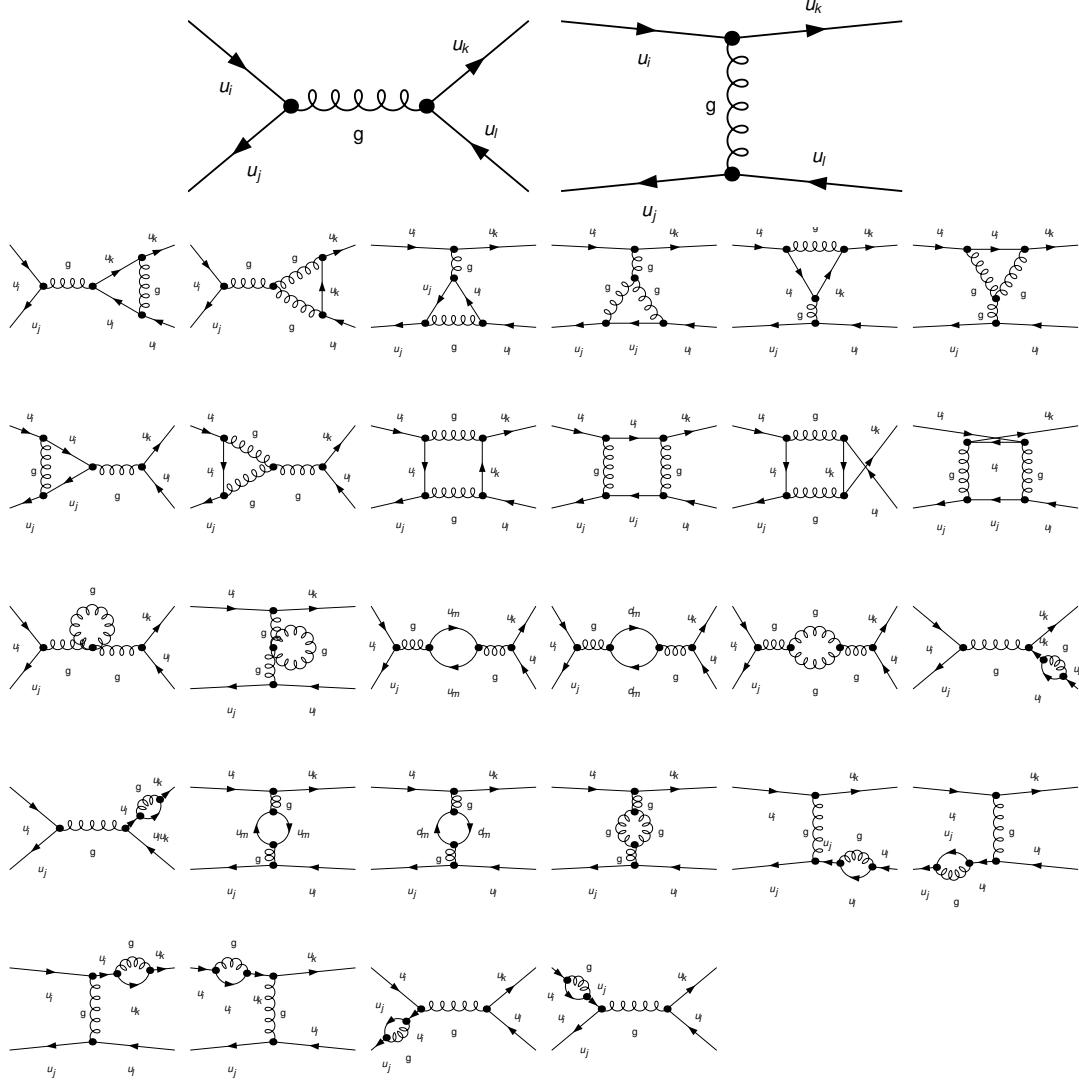


Figure 4.5.: LO and NLO diagrams for quark scattering in QCD [53]. The LO diagrams (top) are the simplest possible for this process, and so they have only two vertices and have a matrix element contribution proportional to α_s . Allowing another two $q\bar{q}g$ or ggg vertices (or one $gggg$ vertex) gives the NLO diagrams, which are proportional to α_s^2 . Note that the quarks could also interact via QED, but the coupling constant is much smaller and so the contribution is less significant.

As Q^2 reduces until $\alpha_s \simeq 1$, the perturbative approximation becomes invalid. Higher order diagrams provide substantial contributions to the matrix element, and the interaction can no-longer be thought of as a small correction to the ground state. It is not practical to calculate the matrix element as the sum over a large (or infinite, as $\alpha_s \rightarrow 1$) number of increasingly complex diagrams. There are therefore two regimes of QCD:

perturbative and non-perturbative, also called hard and soft in reference to the size of the momentum transfer. The factorisation theorem [56] implies that these regimes can be treated separately, with perturbative matrix element calculations for hard processes, and soft-QCD broken down into a series of models. This section summarises areas of soft-QCD modelling.

4.2.1. Parton density functions

The description of a proton as a bound state of two up quarks and one down quark gives the impression that the quarks are the only important feature. In fact, the QCD interactions holding the quarks together cannot be ignored when considering the behaviour of a proton at the LHC. Looking at the masses of the particles involved illustrates the size of the effect: compare the ~ 2 MeV up quark and the ~ 5 MeV down quark with the entire proton at 938 MeV [47]. Surrounding the three quarks that define a proton (the ‘valence quarks’) is a sea of gluons and virtual quarks that contribute to the total mass of the system. The general term ‘parton’ is used to describe any object bound within a proton, including the valence quarks. All hadrons can be described in the same terms — the only differences are the number and flavours of the valence quarks.

When probing the proton with another particle (with high-enough momentum to overcome confinement), it is quite possible that the parton encountered by the probe will be a gluon or virtual quark, rather than a valence quark. The probability of a particular interaction is governed by the distribution of partons within the proton, which cannot yet be calculated from first principles, largely because they interact at energies too low for perturbative QCD to be used. While computational methods such as lattice QCD [57] are making progress in this area, for the moment the parton distributions are modelled by fitting parameterised functions to experimental measurements. These Parton Density Functions (PDFs) vary with x , the fraction of a proton’s momentum carried by a given parton, and Q^2 , the momentum transferred by the particle probe.

Different methods of fitting to experimental results produce PDF sets with differing behaviour when they are extrapolated beyond the energy ranges of those results. One such PDF set is MSTW, which has 30 free parameters in the functions that define the different parton components of a proton [58]. For example, the contribution from up-type

valence quarks is given by:

$$xu_v(x, Q_0^2) = A_u x^{\eta 1} (1 - x)^{\eta 2} (1 + \epsilon_u \sqrt{x} + \gamma_u x),$$

where A_u , $\eta 1$, $\eta 2$, ϵ_u and γ_u are free parameters. The evolution of the MSTW PDF set with x , at two different values of Q^2 , is shown in Figure 4.6. Note the large contribution from gluons, which dominates at lower x values and has to be scaled by 0.1 in order to display it with the other parton contributions. At larger values of x the valence quark contributions become more significant.

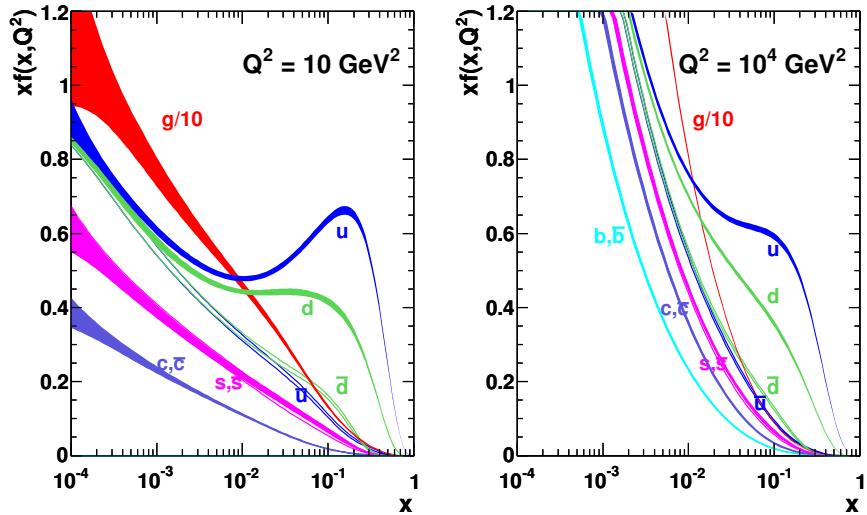


Figure 4.6.: The MSTW PDF [58] at NLO for two different values of Q^2 . Bands represent the 68% confidence level.

The NNPDF set [59] takes an alternative approach. Rather than use a set of parameterised functions fitted to experimental data, neural networks are trained to reproduce it (hence the ‘NN’ in the name). This approach avoids potential systematic bias from assuming a functional form for the PDFs.

4.2.2. Multiple parton interactions

Although the complexity of proton structure may tempt physicists to consider parton collisions in isolation, this gives unphysical results. The divergence of QCD matrix elements at low Q^2 leads to the cross-sections for these processes growing rapidly with \sqrt{s} . For example, at $\sqrt{s} \sim 1$ TeV the calculated cross-section for two partons producing two jets (described in Section 4.2.4) with $p_T > 2$ GeV becomes greater than the observed total

cross-section for proton-proton collisions [60]. The evolution of these two cross-sections is shown in Figure 4.7. This result indicates that a hadronic collision cannot be treated as

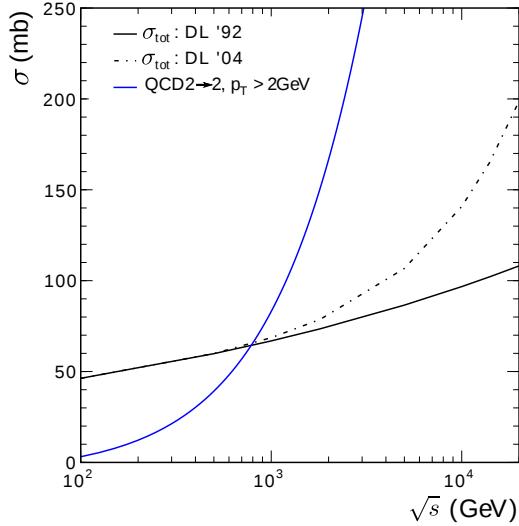


Figure 4.7.: The evolution of the cross-section for production of jets with $p_T > 2 \text{ GeV}$, in a QCD $2 \rightarrow 2$ process, compared to the Donnachie-Landshoff parameterisation of the total proton-proton cross-section [61].

a single partonic interaction.

A common approach to this problem is to introduce Multiple Parton Interactions (MPI). For example, if a single hadron collision contains two partonic interactions, it will “count” twice towards σ_{parton} but only once towards σ_{hadron} . Hadron collisions are therefore interpreted as containing a mean number of parton interactions $\bar{N} = \sigma_{\text{parton}}/\sigma_{\text{hadron}}$. Thus the unphysically increasing cross-section for a single parton interaction becomes a sensible cross-section for an increasing number of interactions. Each interaction is assumed to happen independently, and so the number of interactions in a particular hadron collision is Poisson-distributed.

This interpretation is supported by experimental evidence of pairs of jets with balanced p_T within multi-jet events at hadron colliders [62]. If all jets had come from a single interaction then they would only need to conserve p_T as an ensemble, so finding a pair of balanced jets suggests that they may have come from a separate interaction. More recent measurements examine events with a high- p_T photon and three jets [63]. Here MPI events can be identified when one of the jets balances the photon and the other two jets balance each other. MPI models give good performance simulating hadron collisions, and are commonly used in general purpose Monte Carlo generators (see Section 5.2).

4.2.3. Parton showers and hadronisation

Once a colour-charged object has left the proton-proton interaction it fragments into lower energy objects: a parton shower. Given the LO interactions in QCD, a quark might radiate a gluon, and a gluon might become a pair of gluons or a quark-antiquark pair. This showering can be described approximately with the DGLAP evolution equations, named for the co-creators Dokshitzer [64], Gribov and Lipatov [65], Altarelli and Parisi [66].

Note that there is some ambiguity between the production of quarks and gluons in parton showers and their emission from the original interaction. If (for example) a quark leaving a $2 \rightarrow 2$ interaction radiates an energetic gluon in showering, this is approximately the same as the interaction having created the gluon itself as a $2 \rightarrow 3$ process. It is more difficult to calculate the QCD matrix elements for interactions with more outgoing particles, so some models restrict themselves to $2 \rightarrow 2$ processes and allow the parton shower to approximate the other possibilities. Alternatively these more complex matrix elements can be used, and parton showers are restricted to give results consistent with the number of outgoing particles.

Like the matrix elements in perturbative calculations, there can be many different possible diagrams describing the evolution of a parton shower. There is no unique ordering of the shower — no definite way to say which particle split from which other particle. Parton shower calculations impose their own ordering, and although the details may vary, the end result is that the splittings with the largest opening angle are considered to have happened earliest in the shower. Some different methods are discussed in Section 5.3.

As the outgoing parton showers separate, eventually they cover distance scales where α_s becomes large enough for confinement to be a consideration. Colour-charged particles group together into colour-neutral mesons and baryons which can then behave as free particles. Again, there are different approaches to modelling this grouping — called hadronisation — which are discussed in Section 5.4. The transition between parton shower and hadronisation models is performed at a model-dependent cut-off, typically at an energy scale of 1 GeV.

Parton showering and hadronisation must occur for all outgoing partons, although if they are below the cut-off they will not shower before hadronisation. These partons may be direct products of the scattering process(es), or may have arisen through initial or final state radiation. Initial state radiation (ISR) is emitted from the incoming particles of a scattering process, and final state radiation (FSR) is emitted from those outgoing.

There is a distinction in the case of ISR: here parton shower models are run backwards, starting from the hard scattering [67].

4.2.4. Jets

After hadronisation, a parton shower has become an observable object: a collection of colourless particles leaving the interaction point, called a jet. In ATLAS, these particles (or their decay products) will be detected by the calorimeters (Section 2.2.3) and by the tracking system (Chapter 3) if they are charged. Ideally, the properties of the original parton can be determined from the properties of the jet it creates.

The complication arises when trying to identify jets: there is no unambiguous criterion for determining which particles should be grouped together. Any collection of particles that are close together might be called a jet, but what defines the boundary? Ultimately there are many different answers to this question, and jets are defined in a purely experimental fashion by the algorithm used to group their constituent particles.

A conceptually simple algorithm is the seeded cone. High- p_T particles are identified as potentially defining a jet, and all particles within some radius $\Delta R = \sqrt{\Delta\phi^2 + \Delta y^2}$ are added to that jet. The centre of the jet is now adjusted to reflect the average position of all the particles within it, weighted by their p_T . This may affect which particles are in the jet, and so the adjustment is iterated.

Although the simplicity of this method has made it popular in the past, it suffers from significant problems. Firstly, there is no guarantee that the high- p_T particles that seed the jets will actually be included in them after iteration. Since these particles are the most likely to have come from the hard process, missing them out will cause poor reconstruction of the partons. Secondly, the seeding process violates a logical requirement called “collinear safety:” that a jet be reconstructed the same way if one particle is replaced by two travelling in the same direction with half the momentum. This means that a large group of low- p_T particles will not be seeded as a jet. Finally, the algorithm fails a requirement called “infra-red safety:” that two nearby, separate jets should not be merged if very low p_T particles are added between them.

The infra-red safety requirement is motivated by the divergent behaviour of QCD matrix elements for soft gluon emission. These divergences cancel for observables that are insensitive to the number of partons in the final state [68, 69], allowing predictions to be made. If the radiation of a soft gluon can affect the jet algorithm, then the algorithm

is sensitive to the number of partons (as the emission of a soft gluon does not observably change the parton it has split from).

An alternative family of algorithms — sequential clustering algorithms — are inherently collinear safe and infra-red safe. Rather than jump straight to finding groups of particles, these algorithms combine objects two-at-a-time. Every particle i is given a value, $k_{T,i}^2$, and every possible pair of particles i, j is also given a value, $k_{T,(i,j)}^2$, which is smaller for particles that are close together (see Equation 4.2). All of these values — for separate objects and all pairings — are ordered by size, and the smallest is examined. If the smallest value corresponds to a pair of objects then they are replaced by a new object with the sum of their momenta, building up a jet. The values of $k_{T,i}^2$ and $k_{T,(i,j)}^2$ are now calculated for the jet, and added to the ordering. If the smallest value corresponds to a single object it is removed from the process. Either this is a finished jet (if made of multiple particles), or a single particle that is isolated from the others. The process is repeated until there are no more objects left to examine.

There are different ways of calculating the $k_{T,i}^2$ and $k_{T,(i,j)}^2$ values, which yield different results[70]. They all have the same basic form, shown in Equations 4.1 and 4.2 respectively, but with different values for the parameter n . Note the free parameter D in the definition of $k_{T,(i,j)}^2$. This distance parameter is similar to the cone jet radius parameter, and there is no physically-favoured value.

$$k_{T,i}^2 = p_{T,i}^{2n} \quad (4.1)$$

$$k_{T,(i,j)}^2 = \min(p_{T,i}^{2n}, p_{T,j}^{2n}) \times \frac{\Delta R_{i,j}^2}{D^2} \quad (4.2)$$

$$\Delta R_{i,j}^2 = (\phi_i - \phi_j)^2 + (y_i - y_j)^2 \quad (4.3)$$

The k_T algorithm [71] is perhaps the most intuitive: with $n = 1$ it starts by clustering low- p_T particles to their nearest neighbours, building up higher- p_T objects. Conceptually it is reversing the evolution of the parton shower, and so does a good job of reconstructing the original parton. However, the resulting jets can be irregularly-shaped and cover a large range in y and ϕ , meaning they are potentially difficult to calibrate in a non-homogeneous detector like ATLAS.

The anti- k_T algorithm [70] has $n = -1$, so instead starts by examining the high- p_T particles, collecting up objects that are near to them. This tends to give conical jets, but the connection to the original parton shower is less obvious.

The Cambridge/Aachen algorithm [72] uses $n = 0$, so has no p_T -dependence: it just clusters nearby objects. It tends to give similar performance to the anti- k_T algorithm, but with a greater range of jet areas. However, the algorithm may construct jets where there was no original parton, simply by finding a large number of very low- p_T particles in a small area.

Historically, sequential clustering algorithms were computationally expensive and so not widely used. With improvements to the implementation of the algorithm (i.e. the FastJet library [73]), and increasingly complex corrections required in cone algorithms, sequential clustering is now the more practical choice. It also tends to give better results, with the ATLAS default method being the anti- k_T algorithm as it appears to be the most efficient at identifying jets [74]. A representative example of using different jet algorithms on the same set of particles is shown in Figure 4.8. Note the inclusion of the Seedless, Infra-red Safe Cone algorithm (SIScone) [75], which identifies jets from clusters of particles rather than high- p_T seeds.

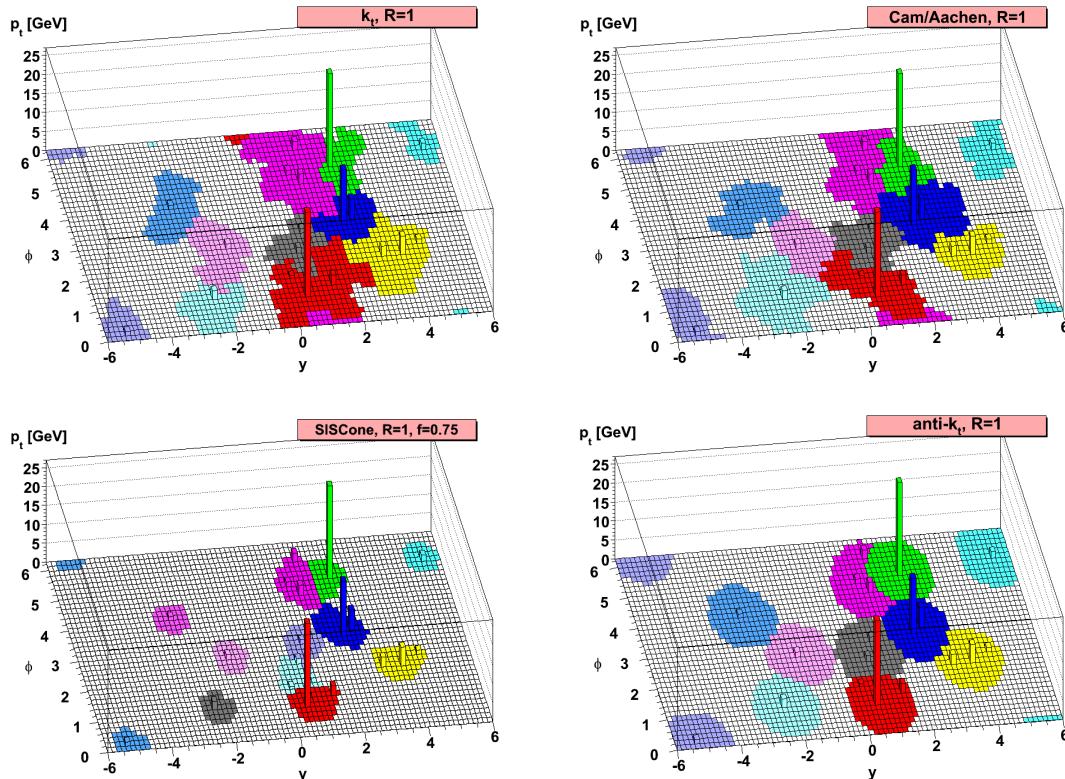


Figure 4.8.: Example of clustering particles in the same event to form jets using different algorithms [70]. Note the additional parameter f in the SIScone result, which indicates the fraction of momentum that must be shared by two overlapping cones for them to be merged into a single jet candidate.

Jet production is one of the dominant processes at a hadron collider like the LHC, as shown using predictions from the MSTW2008 PDF set in Figure 4.9. Given the high probability of jet production, jets are used to identify parton interactions for the analysis in Chapter 6.

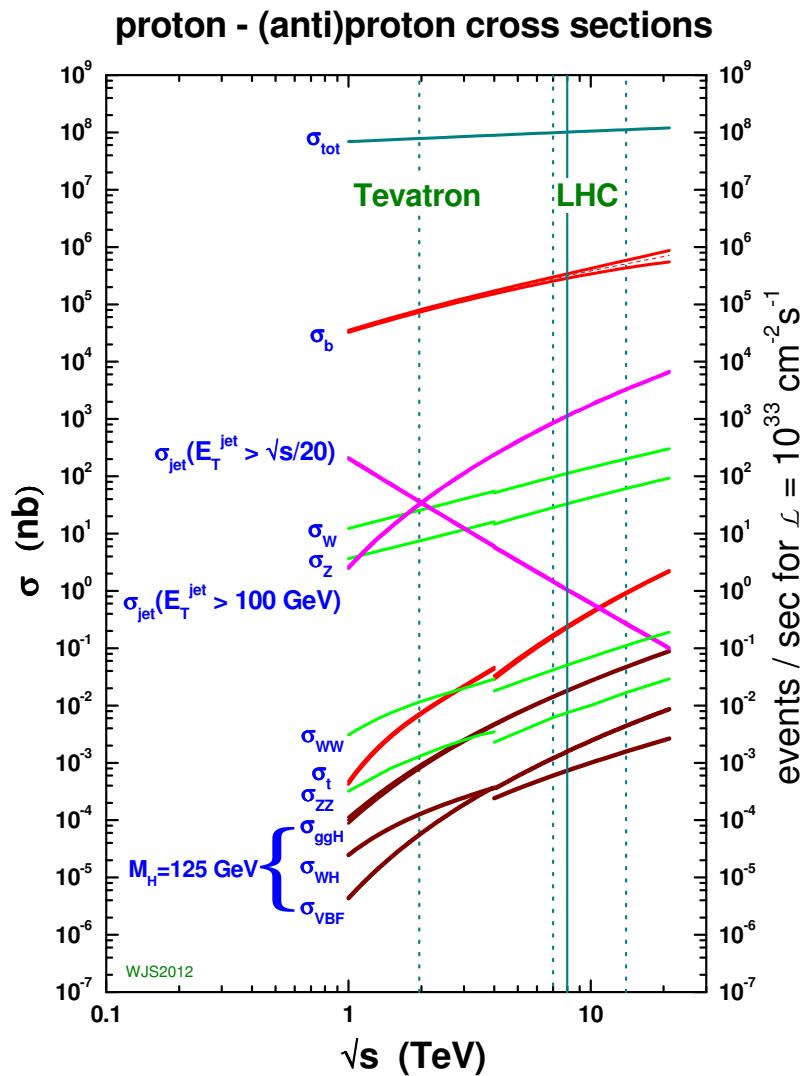


Figure 4.9.: Evolution of the cross-sections for different Standard Model processes with \sqrt{s} at hadron colliders [76]. Vertical lines correspond to particular \sqrt{s} values for the Tevatron (1.96 TeV) and LHC (7, 8, and in future 14 TeV). The discontinuity between Tevatron and LHC energies reflects the switch from proton-antiproton collisions to proton-proton. The results were calculated using the MSTW2008 (NLO) parton distributions [58].

Chapter 5.

Monte Carlo generators

The purpose of a Monte Carlo (MC) generator is to simulate the events in a particle physics experiment. At the LHC this means all aspects of a proton-proton collision, including perturbative interactions between partons, the evolution of parton showers, and the formation of stable hadrons (but not pile-up or the interaction of particles with the detector itself). Pseudo-random number generators are used to generate specific events from the probability distributions describing these processes, hence the name ‘Monte Carlo.’

Where possible, the probability distributions are calculated from first principles (or at least, use stored results of such calculations). However, most processes require some modelling assumptions that are not uniquely constrained by theory, particularly in the non-perturbative QCD regime. The models can also include free parameters, which are adjusted to best reproduce experimental measurements in a process called ‘tuning.’ A particular Monte Carlo generator can have many tunes: each is a set of values for the parameters. The generators may provide options for which models and which PDF set to use, in which case a tune will specify the choices made. Figure 5.1 shows an example of different models, and different tunes of the same model, all producing different behaviour.

Many generators are designed to test a particular model or simulate a specific physics process. General purpose Monte Carlo generators attempt to cover all possibilities, giving output as close as possible to what would be observed in a real experiment. Two distinct families of general purpose generators — PYTHIA and HERWIG — are widely used, and the different models they use are described in this section. PYTHIA generators use p_T or virtuality-ordered parton showers (see Section 5.3), and the Lund string model of hadronisation (see Section 5.4.1). In contrast to this, HERWIG generators use angular

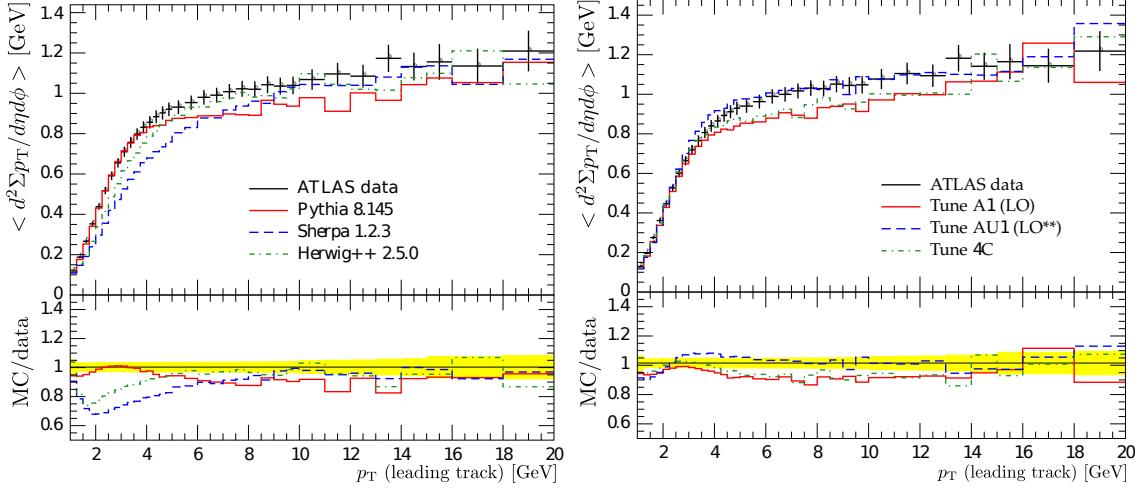


Figure 5.1.: Comparison of Monte Carlo generator output with ATLAS leading track underlying event data (see Chapter 6 for definitions of the observables). Different generators are compared in the left-hand plot [78], and in the right-hand plot three different tunes of Pythia 8 are compared [79].

ordering in parton showers, and the cluster model of hadronisation (see Section 5.4.2). These are characteristic differences, but there are many others, such as their approaches to MPI modelling (see Section 5.2) which is of particular relevance to the analysis in Chapter 6. PYTHIA typically provides more model options and free parameters for tuning than HERWIG.

Another general purpose Monte Carlo generator, Sherpa [80], is in common usage. It is not used in Chapter 6 due to unsolved technical difficulties in generating a sample of events containing jets spanning the full p_T range under investigation. Consequently it is not discussed in detail below, but it has broad similarities with HERWIG such as the use of cluster hadronisation.

5.1. Generators and tunes

The PYTHIA and HERWIG families are both split between an older and a newer version of the generator, in each case the newer is a complete re-write of the older FORTRAN code in C++, with some new features introduced as well. PYTHIA 6 [81] and Pythia 8 [82] differ in their treatment of MPI, which in the case of Pythia 8 must use the same parton shower phase space as the initial and final-state radiation from the hard scatter. In early PYTHIA 6 versions the MPI was treated entirely separately, although later an option was

introduced to share phase space between MPI and ISR. HERWIG [83] and Herwig++ [84] generators differ in their approach to MPI as well. Herwig++ has an internal MPI model, whereas for HERWIG it must be introduced through an external module, JIMMY [85].

Both PYTHIA and HERWIG generators use leading-order matrix elements to describe parton scatters, and only $2 \rightarrow 2$ processes are considered. Rather than include additional matrix elements for a larger number of outgoing particles, equivalent behaviour is produced with the parton shower models. On the other hand, the Monte Carlo generator ALPGEN [86] does include matrix elements for $2 \rightarrow n$ processes, where the maximum value of n depends on the process in question. Of interest to the analysis in Chapter 6 is its ability to produce up to six outgoing partons. ALPGEN is not a full implementation of a general purpose generator, but instead can interface with others to use their MPI, parton shower and hadronisation models.

The analysis results in Chapter 8 are compared with the following generators and tunes:

PYTHIA6 DW: The D0-Willis (DW) tune [87] of PYTHIA 6 is the oldest in this comparison, designed to describe Tevatron Run II data. It uses the CTEQ5L1 PDF set [88], the virtuality-ordered parton shower model, and the separate treatment of ISR and MPI found in the earliest versions of PYTHIA 6. Although it predates LHC data-taking, the DW tune has proven surprisingly effective at describing ATLAS results.

PYTHIA6 AUET2B: A much more recent tune of PYTHIA 6 than DW, the ATLAS underlying event tune 2B (AUET2B) [79] was tuned using ATLAS data, as well as data from the CDF and D0 experiments at the Tevatron. The PDF set has been updated to CTEQ6L1 [89], and the newer PYTHIA 6 models of p_T -ordered showers and shared MPI and ISR phase-space are used. Consequently it has little in common with the DW tune, despite the shared generator.

Pythia 8 AU2: ATLAS underlying event tune 2 (AU2) [90] uses p_T -ordered parton showers — the only choice available in Pythia 8 — and the CT10 PDF set [91]. This tune was made to describe ATLAS data only, and is the standard choice for ATLAS jet simulation.

HERWIG+JIMMY AUET2: Like PYTHIA 6 AUET2B, AUET2 is a recent tune of an old generator (HERWIG+JIMMY), using ATLAS data as well as data from LEP and the Tevatron. The MRST LO** PDF [92] was chosen, allowing the leading-

order generator to approximate higher-order behaviour. As with all HERWIG tunes, angular-ordered parton showers and cluster hadronisation are used.

Herwig++ UE7-2: This tune of Herwig++ to ATLAS underlying event results at $\sqrt{s} = 7 \text{ TeV}$ — UE7-2 [93] — was produced in response to the poor performance of Herwig++ version 2.4.2 in describing ATLAS data. It introduced a new feature in the hadronisation model: colour reconnection, described in Section 5.4.3. Like HERWIG+JIMMY AUET2, this tune uses the MRST LO** PDF, angular-ordered parton showers and cluster hadronisation (with the addition of colour reconnection).

ALPGEN: As mentioned above, ALPGEN [86] must interface to another generator for simulation of MPI, parton showers and hadronisation. In this case HERWIG+JIMMY was used, meaning angular ordered parton showers (matched to the number of outgoing particles from the ALPGEN matrix element [94]) and cluster hadronisation.

The correction process described in Chapter 7 requires input from Monte Carlo generators as well. Two are used: PYTHIA 6 with the AMBT1 tune [95], and the default tune of Herwig++ version 2.5.0 (versus 2.5.1 used with the UE7-2 tune). The details of these tunes are not important as the correction is approximately model independent; they were chosen because a large number of events from these tunes had been produced and passed through simulation of the ATLAS detector.

5.2. Multiple parton interactions

As described in Section 4.2.2, the increase of the $2 \rightarrow 2$ partonic cross-section above the total hadronic cross-section is interpreted as multiple partonic interactions occurring within the hadron collision. The earliest MPI model [96] defines a mean number of parton interactions $\bar{N} = \sigma_{\text{parton}}/\sigma_{\text{hadron}}$, with the actual number of interactions obeying a Poisson distribution:

$$P(n) = \frac{1}{n!} \bar{N}^n e^{-\bar{N}} \quad (5.1)$$

The MPI interpretation still gives divergent behaviour for \bar{N} , and it is in the treatment of this divergence that Monte Carlo generators differ. From perturbative QCD it can be shown that the divergence is driven by the low- p_{T} behaviour of the partonic cross-section,

which varies approximately as:

$$\sigma_{\text{parton}} \sim \frac{1}{p_T^2} \quad (5.2)$$

The simplest approach to this divergence — adopted by most early Monte Carlo generators and still used by HERWIG+JIMMY — is to introduce a minimum p_T value as a free parameter. By only evaluating the cross-section above $p_{T\text{min}}$ the divergence is controlled. A slightly more sophisticated approach used in PYTHIA 6 and Pythia 8 is to introduce $p_{T\text{min}}$ as a smooth variation rather than an abrupt cut-off. The partonic cross-section behaviour in Equation 5.2 is modified as follows:

$$\sigma_{\text{parton}} \sim \frac{p_T^2}{(p_T^2 + p_{T\text{min}}^2)^2} \quad (5.3)$$

The introduction of $p_{T\text{min}}$ can be interpreted in terms of QCD confinement. Since the p_T of the products of a partonic interaction indicates the magnitude of the momentum transfer Q^2 in that interaction, low p_T implies high α_s . Thus, below $p_{T\text{min}}$ the partons will not be available as free particles for interactions. Herwig++ extends this concept by modelling σ_{parton} below $p_{T\text{min}}$ using behaviour observed in proton elastic scattering.

Despite their differences, all of these models have the same qualitative behaviour: below $p_{T\text{min}}$ MPI activity is suppressed, so choosing a higher value for $p_{T\text{min}}$ leads to fewer parton interactions per hadron collision.

5.2.1. Evolution of $p_{T\text{min}}$

The PYTHIA MPI model also introduces variation in the value of $p_{T\text{min}}$ as a function of the collider centre-of-mass energy \sqrt{s} :

$$p_{T\text{min}}(\sqrt{s}) = p_{T\text{min}}(1.8 \text{ TeV}) \left(\frac{\sqrt{s}}{1.8 \text{ TeV}} \right)^{\frac{e}{2}} \quad (5.4)$$

The evolution of $p_{T\text{min}}$ is governed by a new free parameter, e , and the reference point of 1800 GeV corresponds to data from Tevatron Run I. Values for the free parameters in this model were first determined by comparing its output to this Tevatron data. HERWIG+JIMMY and early versions of Herwig++ do not include evolution of $p_{T\text{min}}$, but it has now been introduced to Herwig++.

5.2.2. Hadronic form factor

Measurements of MPI in association with a hard scattering process — the underlying event, described in Chapter 6 — show increasing activity with the scale of the hard process. The rise is initially rapid, but then reaches a plateau. This is included in MPI models by considering the ‘centrality’ of hadron collisions, i.e. whether the collision is head-on or glancing. As protons have a finite size, a head-on collision will mean a greater overlap of their matter distributions, and this should increase the probability of multiple parton interactions. Although the probability is still calculated as shown in Equation 5.1, the mean number of interactions \bar{N} is now taken to be a function of the impact parameter b , with $b = 0$ indicating perfect overlap.

The plateau behaviour arises because b is correlated with the hard process scale Q^2 , as well as with MPI. A head-on collision involves greater momentum transfer, and so lower b increases the likelihood of a high- Q^2 process. Therefore, as the hard process scale increases so does the MPI activity, assuming a reduction of the impact parameter until the collisions are fully overlapping. With $b \simeq 0$, this mechanism for varying MPI reaches a maximum, hence the plateau.

The variation of \bar{N} with b is determined by modelling the distribution of partons within a hadron. PYTHIA generators provide a choice of models: single or double Gaussian, with the widths as free parameters; or a general overlap function of the form

$$O(b) \propto e^{-b^\alpha}, \quad (5.5)$$

where α is a free parameter. HERWIG generators provide only one model, using the Fourier transform of the proton electromagnetic form factor [60]:

$$G(\mathbf{b}) = \int \frac{d^2\mathbf{k}}{2\pi} \frac{e^{i\mathbf{k} \cdot \mathbf{b}}}{(1 + \mathbf{k}^2/\mu^2)^2}, \quad (5.6)$$

where μ is taken as a free parameter.

A further refinement has been introduced in Pythia 8 and recent Herwig++ versions. The ‘hot spot’ model modifies the distribution according to the fraction x of the hadron momentum carried by a parton. Harder partons are modelled as being concentrated in smaller regions than the overall distribution allows [61].

5.3. Parton showers

A parton shower is the evolution of a single parton far off the mass-shell into a group of partons that are on-shell or nearly so. The mass-shell is simply the set of solutions to the relation $E^2 - \vec{p}^2 = m^2$ from special relativity, and particles that do not satisfy this relationship are called off-shell or virtual. Shower evolution is accomplished by repeated splitting of the partons: a quark (or antiquark) radiates a gluon; a gluon splits into a pair of gluons; or a gluon splits into quark-antiquark pair.

The phase space for parton showers is dominated by diagrams where the splittings are strongly ordered in the scale t at which they occur. Thus, when generating a shower from a parton i with scale t_1 , what is needed is the probability distribution that it will split at a particular lower scale t_2 . This distribution is calculated using a Sudakov form factor [97]:

$$\Delta_i(t_2, t_1) = \exp \left(- \int_{t_2}^{t_1} dt' \sum_{jk} \frac{\alpha_{ijk}(t')}{2\pi} \int dz P_{i \rightarrow jk}(z) \right), \quad (5.7)$$

where $\alpha_{ijk}(t')$ is the coupling constant for the splitting at scale t' , and j, k are possible species of the two partons produced. Parton j carries a fraction z of the available momentum. This form factor can be interpreted as the probability that no splittings of any kind occurred between t_1 and t_2 , hence the sum over all product species and momentum fractions.

To use this distribution a random number $0 \leq r \leq 1$ is generated, and a value for t_2 is sought such that $r = \Delta_i(t_2, t_1)$. No splitting occurs if the value for t_2 produced is below the scale at which hadronisation occurs. If splitting does occur, its properties are given by another probability distribution, proportional to the DGLAP [64–66] splitting kernels $P_{i \rightarrow jk}(z)$. This gives the species of partons j and k , and the fraction z of the available energy that is taken by parton j . Flavour and momentum are conserved at

each splitting. The forms of the splitting kernels used in PYTHIA 6 are as follows [81]:

$$P_{q \rightarrow qg}(z) = \frac{4}{3} \left(\frac{1+z^2}{1-z} \right) \quad (5.8)$$

$$P_{g \rightarrow gg}(z) = 3 \frac{(1-z(1-z))^2}{z(1-z)} \quad (5.9)$$

$$P_{g \rightarrow q\bar{q}}(z) = \frac{n_f}{2} (z^2 + (1-z)^2) \quad (5.10)$$

$$P_{q \rightarrow q\gamma}(z) = e_q^2 \left(\frac{1+z^2}{1-z} \right) \quad (5.11)$$

where n_f is the number of quark flavours allowed in the parton shower, and e_q is the electric charge of the quark. Note the inclusion of a kernel for a quark to radiate a photon (a lepton may also do so, with a kernel of the same form). In PYTHIA 6 the photon takes no part in the showering, but Pythia 8 includes kernels for $P_{\gamma \rightarrow l\bar{l}}$ and $P_{\gamma \rightarrow q\bar{q}}$. Such splittings are rare, hence their omission in PYTHIA 6. The emission of a photon is itself relatively unlikely, given the small coupling constant compared to α_s .

The scale t has not been defined thus far because it has multiple possible definitions. In early PYTHIA models $E^2 - \vec{p}^2$ was used (referred to a ‘virtuality-ordered’), but the option has been added for a p_T -ordered shower, with $p_T^2 = z(1-z)(E^2 - \vec{p}^2)$. Soft gluon emission is only treated correctly in parton showers if splittings are ordered by decreasing splitting angle, a result referred to as colour coherence [78]. This angular ordering occurs automatically in p_T -ordered showers, but when using virtuality ordering the angular condition must be applied separately. HERWIG generators approach angular ordering directly, using

$$t = \frac{p_j \cdot p_k}{E_j E_k} \simeq \frac{1}{2} \theta_{jk}^2 \quad \text{for small } \theta. \quad (5.12)$$

5.4. Hadronisation

As parton showers evolve they reach an energy scale where QCD confinement becomes significant. Colour-charged particles must be grouped into colour-neutral bound states, in a process called hadronisation — the formation of hadrons. There are two major models of hadronisation in use in Monte Carlo generators: string and cluster.

5.4.1. String hadronisation

String models of hadronisation take as a starting point the linearly-rising effective potential of QCD at large distance scales [55]. If the QCD potential is modelled simply as $V(r) = \kappa r$ then the coefficient $\kappa \sim 1 \text{ GeV/fm}$ [47]. This potential can be treated as a string with tension κ , hence the name of the model. Separating a $q\bar{q}$ pair is equivalent to stretching this colour potential string, and the energy required to do so quickly becomes larger than the energy required to create a new $q\bar{q}$ pair. The string can therefore be broken by creating a $q\bar{q}$ pair at some point along its length, as shown in Figure 5.2. Hadrons can then be made by grouping these new quarks.

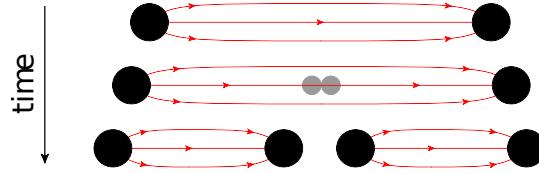


Figure 5.2.: The QCD potential between two colour-charged particles rises linearly with their separation. Eventually the potential energy is sufficient to allow creation of $q\bar{q}$ pairs. Adapted from Reference [47].

The Lund string model [98] is a widely-used version of this approach. Starting at one end, the string is broken by inserting a new $q\bar{q}$ pair and splitting off the outermost quarks as a hadron bound state. This process continues towards the middle of the string, randomly switching which end is under consideration. There is no causal connection between string breaks, so they can be considered in any order, but this approach makes it easy to ensure that only valid hadronic states are produced. A simple example to illustrate the method is shown in Figure 5.3.

Quantum tunnelling is used to model the distribution of quark masses (m) produced, and their momentum perpendicular to the string (p_\perp), giving the following functional form:

$$P(m^2, p_\perp^2) \propto e^{-Am^2} e^{-Bp_\perp^2}. \quad (5.13)$$

The mass dependence essentially eliminates the production of charm quarks or heavier. Strange quarks are produced, but this is suppressed in an additional step in order to best reproduce experimental results. The suppression can be considered a response to the uncertainties on the masses of the light quarks.

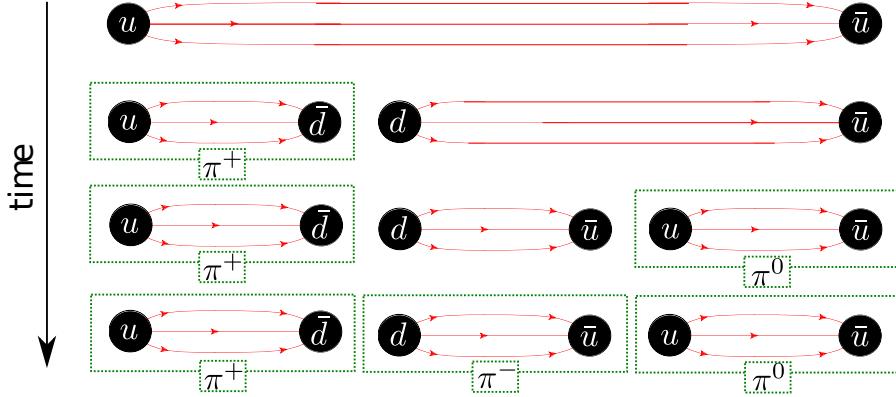


Figure 5.3.: An example of Lund string hadronisation, starting from a $u\bar{u}$ pair. First the string is broken with the creation of a $d\bar{d}$ pair, allowing the quarks at the end to be split off as a positive pion $u\bar{d}$. Then a $u\bar{u}$ pair is created at the other end, splitting off a neutral pion $u\bar{u}$. The remaining quarks can form a negative pion $d\bar{u}$. Adapted from Reference [47].

The Lund string model method of creating a $q\bar{q}$ pair and then immediately splitting off a quark bound state would suggest that only mesons could ever be created. Baryon production is included by also allowing the creation of *diquark* pairs, i.e. four quarks in total, split into two ‘loosely bound’ states. Equation 5.13 again governs the distribution of the masses of the diquark pairs produced. Like strange quarks, uncertainties on the masses of the diquarks require that an additional step is applied to ensure that the ratio of meson to baryon production best matches experimental data.

Once the hadrons have been created and given momentum perpendicular to the string, the remaining task is to assign their longitudinal momenta. This is expressed as a fraction z of the longitudinal momentum of the quarks at the ends of the string, and the probability distribution of z arises from the symmetry of the string model itself:

$$f(z) \propto \frac{1}{z} (1-z)^a \exp\left(-\frac{1}{z} b(m_h^2 + p_{\perp}^2)\right) \quad (5.14)$$

The parameters a and b are included to regulate the behaviour of $f(z)$ as $z \rightarrow 1$ and $z \rightarrow 0$ respectively. Note the mass-dependence, which gives a larger expectation value of the longitudinal momentum for heavier hadrons.

Gluons exist in the Lund string model only as modifications to the momentum of the string. Soft gluons from the parton shower are simply absorbed, while harder ones bend the string. The transverse and longitudinal momenta of the hadrons produced must be adjusted accordingly.

5.4.2. Cluster hadronisation

The cluster model of hadronisation is based on the concept of QCD preconfinement [99]. Even in the showering process — before confinement is a concern — partons tend to be arranged in systems that are colour neutral, with total mass that obeys a universal distribution. In the cluster model [100], all gluons are forced to split into $q\bar{q}$ pairs, and then these colour neutral clusters are identified.

Each cluster is treated as a resonance that can decay into hadrons. Two-body decay is assumed, either to two mesons, or to a baryon and an antibaryon. Either a $q\bar{q}$ pair or a diquark pair is introduced to each cluster, with flavour randomly chosen from the light quarks, and a list of the possible decays corresponding to the quark flavours is created. The relative probabilities of each decay are calculated using the same kinematic phase-space approach as for a first-principles cross-section calculation. A random number is compared to each probability to determine if that decay occurs. If no decay occurs at all, an alternative set of quarks is tried and the process is repeated. Note that the kinematics of the decay are entirely determined by the original quarks from the parton shower: the newly introduced (di)quarks are given no physical properties other than flavour.

In about 10% of cases [100], the cluster is too massive (> 4 GeV) to decay immediately to hadrons. Instead, the cluster is split by the introduction of a $u\bar{u}$, $d\bar{d}$ or $s\bar{s}$ pair with equal probability, forming two new clusters each with a fraction of the original's momentum. This splitting can occur multiple times until all clusters are below the threshold. Clusters containing heavy quarks are also split by forcing the quarks to decay (via the weak force), producing either two new clusters, or one new cluster and leptons. In the case of a decay $b \rightarrow c + X$, the charm quark is forced to undergo a subsequent decay. This means that only the well-understood light quark hadrons need to be considered when determining cluster decay products.

Some clusters do not undergo two-body decay, but are allowed to decay directly to a single hadron. This is included to better describe experimental data, where a single hadron is observed carrying a large fraction of the total momentum of a jet. The probability of clusters decaying to a single hadron is a free parameter, set to reproduce the observation.

5.4.3. Colour reconnection

The idea of colour reconnection is that in a collision with MPI, the strings or clusters that decay to hadrons may be formed from partons leaving different interactions. An ‘annealing’ step is included, which may produce a more physically-favoured set of objects for hadronisation. In the Lund string model this means allowing strings to be formed between any partons with compatible colour charges, and then choosing the set of strings with the smallest total potential energy [81]. The approach in Herwig++ cluster hadronisation is similar, this time attempting to minimise the sum of the cluster invariant masses [93].

Both models include a free parameter to regulate how much reconnection is allowed to occur. In PYTHIA 6 this gives a probability for each string not to take part in the reconnection:

$$P = (1 - \alpha)^n, \quad (5.15)$$

where α is the free parameter and n is the total number of partonic interactions, thus reconnection is more likely in events with more MPI. In Herwig++ the free parameter is the probability that a particular pair of clusters will be replaced by another pair with lower total invariant mass.

In the PYTHIA generators there is an additional free parameter included to make strings with high p_T less likely to participate in colour reconnection. This is motivated by the idea that such strings will have less time to participate in annealing.

Chapter 6.

The Underlying Event: analysis

The underlying event (UE) encompasses all particles produced in a hadron collision that are not associated with the hard scattering process. This is primarily MPI activity (as described in Section 4.2.2), but measurements may unavoidably include contributions from other effects. Particles in the underlying event come from soft-QCD processes and the energy scales involved are too low for perturbative methods to be used.

Instead, these processes are described by approximate models (see Chapter 5), with free parameters adjusted to match experimental results. Underlying event measurements are a direct test of these models, allowing their relative merits to be examined and their subsequent performance improved. As well as potentially improving understanding of soft-QCD, if the underlying event can be accurately modelled then it can be removed from the analyses of hard scattering processes where it is an inconvenient background.

Underlying event observables are expressed as the collective properties of all the selected particles. Some typical observables (which will be used in this analysis) are shown in Table 6.1.

Since particles in the underlying event are produced in association with a hard scatter, the observables are measured with reference to the scale of that hard scatter. Underlying event measurements have already been made at the CDF experiment at the Tevatron, at centre-of-mass energies $\sqrt{s} = 1.8$ TeV [101] and $\sqrt{s} = 1.96$ TeV [102], but the LHC allows the investigation of higher energy scales. Note that results from the two colliders may not be directly comparable, as the $p\bar{p}$ collisions at the Tevatron will give rise to more high-momentum $q\bar{q}$ interactions than at the LHC (where there are no valence anti-quarks).

A similar family of measurements examine particles from soft-QCD interactions without requiring a hard scattering process. Typically these just select events as inclusively as possible, hence their name: ‘minimum bias’ measurements. Although similar particle production mechanisms are at work, these measurements are distinct from the underlying event because they do not examine how particle production varies with the scale of the interaction. Minimum bias measurements have been made at the Tevatron [103] and LHC [104].

Observable	Interpretation
$\frac{d^2 N_{\text{ch}}}{d\eta d\phi}$	Number of stable charged particles (N_{ch}) per unit $\eta - \phi$
$\frac{d^2 \sum p_{\text{T}}}{d\eta d\phi}$	Scalar p_{T} sum of stable charged particles per unit $\eta - \phi$
$\langle p_{\text{T}} \rangle$	Mean p_{T} of stable charged particles (requires ≥ 1 particle)
$\frac{d^2 \sum E_{\text{T}}}{d\eta d\phi}$	E_{T} sum of stable charged and neutral particles per unit $\eta - \phi$

Table 6.1.: Typical underlying event observables, and their symbolic representation.

6.1. Topological selection

A topological selection process is used to find the activity in an event that is not associated with the hard process. For each event a particular object is identified as coming from the hard process, and its ϕ -coordinate is measured, defining the orientation of the event. Every other object can now be assigned a relative $\Delta\phi$ value. The region $|\Delta\phi| < 60^\circ$ is called the ‘toward’ region, and is associated with the identified object from the hard process. From momentum conservation it is reasonable to assume that the hard process will also have produced an object travelling in the opposite direction. Therefore the region $|\Delta\phi| > 120^\circ$ is also associated with the hard process, and is referred to as the ‘away’ region.

Assuming the back-to-back topology is valid then the region $60^\circ < |\Delta\phi| < 120^\circ$ should not contain significant contributions from the hard process. Consequently the particles in this ‘transverse’ region should be associated with the underlying event. A diagram of

the different regions is shown in Figure 6.1. The 60° angle used to define the different regions is purely conventional [105], but is adhered to in order to simplify comparing results between different analyses.

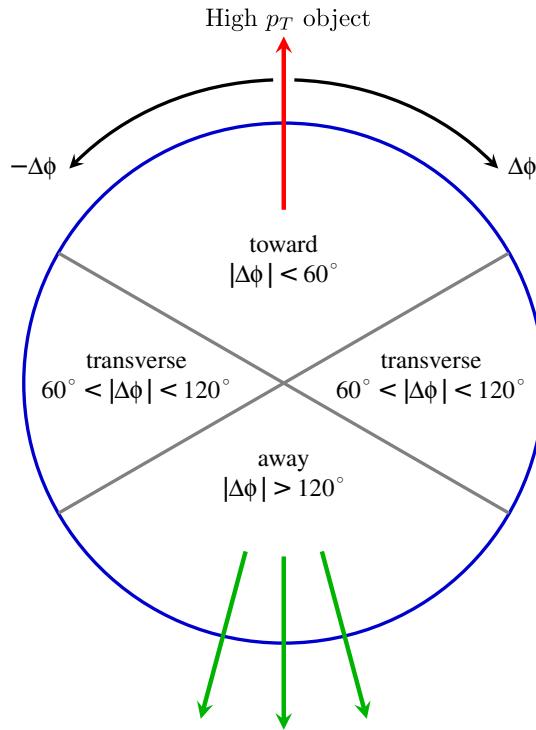


Figure 6.1.: The regions in ϕ around a hard process, defining the underlying event.

The transverse region can be split into two by the sign of the $\Delta\phi$ value. These subdivisions are called transverse maximum and minimum (trans-max and trans-min), depending on which has the most activity in the underlying event observable being measured. Note that trans-max and min are undefined for the $\langle p_T \rangle$ observable.

The object defining the event orientation varies depending on the analysis. Possibly the simplest object to use is the highest p_T track, where a track is any charged particle in a tracking detector. This ‘leading track’ is easy to identify, and is the most inclusive selection for different hard processes. Leading track underlying event measurements have been made at $\sqrt{s} = 900$ GeV and 7 TeV at the LHC with the ATLAS detector [106], and the ALICE detector [107].

In hadron collisions at the LHC, the dominant cross-section for hard scattering processes is jet production, as shown in Figure 4.9. Although it is natural to assume that the highest p_T jet will contain the highest p_T track, this is not guaranteed. Using the leading track may also underestimate the energy scale of the hard process since a jet may

contain many low- p_{T} or neutral particles, depending on the details of the algorithm used (see Section 4.2.4). Therefore, the ‘leading jet’ can be used to define the orientation for underlying event studies instead. Analyses using jets constructed from charged particles alone (‘track-jets’) have been performed with the ATLAS detector [108], and the CMS detector [109]. For the analysis in this chapter, jets are constructed from both charged and neutral particles.

More specific selections may be used to study the underlying event associated with a particular process, for example Drell-Yan [110] production of muon pairs with the CMS detector [111], or electrons and muons with ATLAS [112].

Since underlying event measurements rely on the identification of a particular object, care must be taken to address the potential *misidentification* of this object. As a simple example, if there are two tracks in the detector with similar (high) p_{T} values, a small uncertainty in momentum measurement may lead to the wrong track being identified as the lead. Orienting an event relative to the misidentified lead track means that the definition of the transverse region — and hence which particles are considered part of the underlying event — may be completely incorrect. A correction procedure is needed that can account for this: see Chapter 7.

6.2. Event selection

As the LHC data-taking programme has continued, the intensity of the beams has been increased towards the design value. This has lead to the expected number of proton-proton interactions per bunch crossing exceeding 20 (see Section 3.2). These pile-up interactions are unlikely to produce additional hard scattering processes, but will create particles through soft-QCD activity. The underlying event is a property of a single hadron collision, and so pile-up effects should not be included. To that end, only the early data (collected in 2010) is used in this analysis, since at this time the beam intensity was lower, producing fewer pile-up events.

To reject events due to cosmic ray muons and other non-collision backgrounds, events are required to have at least one primary vertex, constrained by the reconstruction to be consistent with the beam-spot position. This vertex must have at least five associated tracks. The efficiency for collision events to pass these vertex requirements —

as measured in a sample of events passing all other event and jet selection requirements — is well over 99%.

To reduce the contributions from pile-up interactions, events with more than one vertex are removed, here requiring only two tracks to define a vertex. By the end of 2010 data-taking there was sufficient pile-up that 80% of events fail this selection. The only remaining pile-up effect comes from events where the primary and pile-up vertices are too close in the z -direction to be resolved. This contribution is small, and is treated as a systematic error (see Section 8.4.3).

All data events considered in this analysis were required to have good detector status flags for the L1 central trigger processor, solenoid magnet, Inner Detector, calorimeters and luminosity detectors, as well as good tracking, jet, and missing energy reconstruction performance flags. In addition, good data quality was required for the high-level trigger during the periods when this device was used for rejection (see Table 6.2).

6.3. Jet selection

All jets in this analysis were constructed using the anti- k_T algorithm, as described in Section 4.2.4. The jet distance parameter value chosen was $D = 0.4$, the narrower of the two ATLAS default values, although results using $D = 0.6$ were also examined and found to be similar. Jets were constructed from detected energy deposits called topoclusters (see Section 6.5), adding their 4-momenta to produce the uncalibrated jet momentum. A three-step calibration process is then applied:

Pile-up correction: Particles from pile-up interactions may incorrectly be included in jets from the hard process. An average correction to subtract the additional energy due to pile-up interactions is applied. The correction is measured as a function of η and number of pile-up vertices, so in this analysis it is limited to the effects of merged vertices.

Jet origin correction: The position of the jet is corrected such that the jet direction points to the primary vertex of the interaction instead of the geometrical centre of ATLAS. The kinematics of each topocluster are recalculated using the direction from the primary vertex to the cluster centroid. The raw jet four-momentum is then redefined as the four-vector sum of the clusters. This correction improves the angular resolution while the jet energy is unaffected.

Final jet energy scale: The final part of the jet calibration is the “electromagnetic and jet energy scale” (EM+JES) calibration. This corrects for the sampling behaviour of the calorimeters, energy losses in inactive regions, out-of-cone showering effects (where the jet algorithm does not manage to collect all the components of a jet) and inefficiencies in the calorimeter clustering and jet reconstruction. This calibration is primarily dependent on the jet energy and on the region of the detector in which the jet is found, due to the variations in calorimeter technology and amounts of material in front of the calorimeters (see Section 2.2.3).

The EM+JES calibration is calculated by simulating the detector. The correction factors are ratios of the energies of isolated jets in events simulated by an MC generator (PYTHIA 6 AMBT1), compared to these same events reconstructed after the behaviour of the calorimeter has been simulated. An isolated jet is defined as a jet that has no other jet within $\Delta R = 2.5D$, where D is the distance parameter of the jet algorithm.

Following this, a small η -dependent correction is applied to remove a bias in the reconstructed η of jets, which arises due to reduced response from less well-instrumented regions of the calorimeter. Clusters in these regions contribute a smaller weight than those in more active regions when their four-vectors are combined to build the jet. This η -correction is parameterised as a function of jet energy and η , and is small ($\Delta\eta < 0.01$) for most regions of the calorimeter.

Any event in this analysis is required to have at least one jet with a calibrated $p_T > 20\text{ GeV}$. ATLAS triggers (see Section 2.3.2) are chosen to efficiently select events with jets above this p_T threshold, with the specific triggers used shown in Table 6.2.

The MBTS trigger is fully efficient for selecting events with hard jets because it detects the additional soft radiation of the underlying event. However, the rate of events passing this trigger is too high for all of them to be recorded, and so the trigger was heavily prescaled in all but the earliest data-taking. This necessitates the use of the jet triggers — denoted JX — which have smaller prescales, and so provide more events in the higher jet p_T ranges. These jet triggers cover the central rapidity range $|y| < 2.8$.

The jet triggers were brought into full operation gradually over data-taking in 2010, with only the level one triggers used for event selection before period G. Before run 152777 (within period A), mis-timings in the level one central jet trigger hardware meant that all jet triggers were inefficient. Therefore, L1_MBTS_1 was used for all p_T ranges before

p_T range [GeV]	Period A-F	Period G-I
	30/03/10 – 30/08/10	22/09/10 – 29/10/10
20 – 60	L1_MBTS_1	L1_MBTS_1
60 – 110	L1_J5	EF_J20_jetNoEF
110 – 160	L1_J15	EF_J35_jetNoEF
160 – 210	L1_J30	EF_J50_jetNoEF
210 – 260	L1_J55	EF_J75_jetNoEF
260 – 310	L1_J75	EF_J95_jetNoEF
310 – 400	L1_J95	L1_J95
> 400	L1_J95	L1_J115

Table 6.2.: The triggers used for this analysis, chosen to efficiently select events with the leading jet in the p_T range shown [113].

this time — the collision rate was much lower at this very early stage and so the MBTS trigger was not prescaled.

The jet triggers select events containing at least one jet with energy above the threshold indicated. Different thresholds are used in the level one and high level triggers in order to give them similar efficiency. Figure 6.2 shows the performance of the different level one triggers.

There are additional jet quality selections, designed to identify fake jets reconstructed from cosmic rays, LHC beam background or erroneous calorimeter behaviour [113]. An event is rejected if it contains any jets identified as problematic in this way.

Underlying event measurements are presented using this basic jet selection, and also with a di-jet selection. The topology in Section 6.1 *assumes* a di-jet structure in all events, whereas this extra selection enforces it. Besides the leading jet, a second jet must be identified with $p_T > 20$ GeV and $|y| < 2.8$, and there must be no other jets in the event that fulfil these criteria. Furthermore, the second ('sub-leading') jet must balance the leading jet, satisfying $p_{T,sub} > \frac{1}{2}p_{T,lead}$ and $|\Delta\phi| > 2.5$. This is referred to as the exclusive di-jet selection, compared to the inclusive leading jet selection without this requirement.

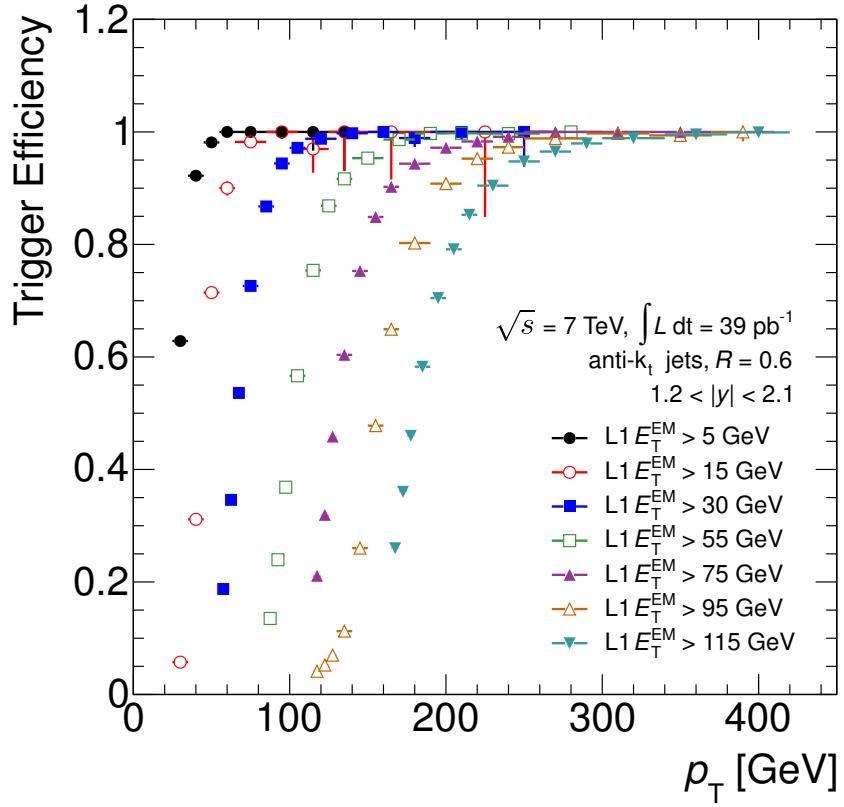


Figure 6.2.: The efficiencies of the level one jet triggers as a function of the lead jet p_T [114].

6.4. Track selection

The observables N_{ch} , $\sum p_T$ and $\langle p_T \rangle$ are all constructed using charged particle tracks from the ATLAS Inner Detector, described in Chapter 3. Although underlying event observables should be constructed as inclusively as possible, some quality selections must be applied to ensure that only properly reconstructed tracks are used. The following selections were chosen to remove badly reconstructed (i.e. potentially fake) tracks [104]:

- $|\eta| < 2.5$. This restricts the track to the fully-instrumented region of the Inner Detector.
- $p_T > 500$ MeV. Below this value the track reconstruction efficiency drops dramatically, as shown in Figure 6.3.
- ≥ 1 Pixel detector hit, and ≥ 6 SCT hits. This requires that the charged particle was detected several times, so the track is less likely to be fake.

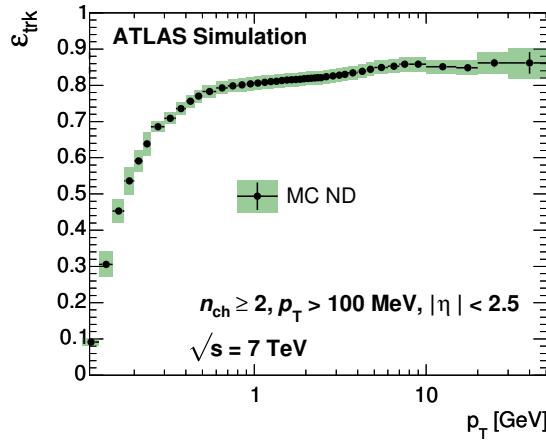


Figure 6.3.: Efficiency of track reconstruction in the ATLAS detector, as a function of track p_T [104]. Result obtained from MC samples of non-diffractive (ND) events, comparing tracks at truth-level with those reconstructed.

- A hit in the innermost pixel layer (the b-layer), if the corresponding pixel module was active. This reduces the background of tracks from secondary vertices.
- Require transverse and longitudinal impact parameters with respect to the primary vertex to be $|d_0| < 1.5 \text{ mm}$ and $|z_0| \sin \theta < 1.5 \text{ mm}$. This requires that tracks be associated with one of the measured vertices, and so they are more likely to originate from a proton-proton collision and not background in the detector. The inclusion of a factor of $\sin \theta$ gives a better measure of the closest approach of a track to the primary vertex than z_0 alone, as shown in Figure 6.4.

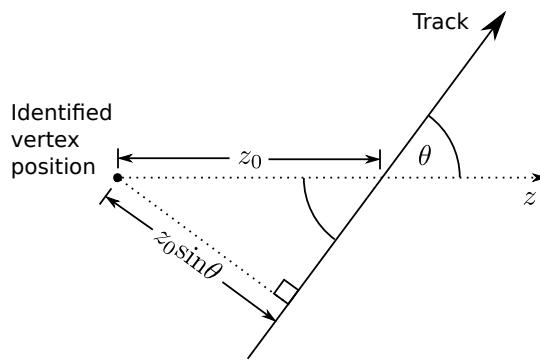


Figure 6.4.: The longitudinal impact parameter (z_0) does not give a good indication of the track-vertex distance unless a factor of $\sin \theta$ is included.

- A long non-Gaussian tail in the track momentum resolution, combined with the steeply falling p_T spectrum, means that tracks with higher reconstructed p_T are increasingly likely to have arisen from incorrect measurement of low- p_T particles.

These are referred to as mis-measured tracks, and to remove them the track fitting χ^2 probability was required to be > 0.01 for tracks with $p_T > 10 \text{ GeV}$.

6.5. Topocluster selection

Topoclusters are energy deposits in the ATLAS calorimeters (see Section 2.2.3). Each topocluster is constructed from a seed calorimeter cell with $|E_{\text{cell}}| > 4\sigma$, where σ is the RMS of the measured cell noise. Neighbouring cells are iteratively added to the topocluster if they have $|E_{\text{cell}}| > 2\sigma$. Finally, an outer layer of surrounding cells is added [115]. In reconstruction, each calorimeter cluster is considered as a massless particle with energy $E = \sum E_{\text{cell}}$, originating from the geometrical centre of the ATLAS detector.

Besides being used as input objects for jet reconstruction, topoclusters are also used to construct the underlying event observable $\sum E_{\text{T}}$. Although this is similar to the $\sum p_{\text{T}}$ observable, the use of topoclusters allows the contribution to the underlying event from neutral particles to be measured as well.

Two selections are applied to the topoclusters used in this analysis:

- $|\eta| < 4.8$. Although the calorimeter acceptance extends to $|\eta| = 4.9$, this slightly tighter selection avoids effects from topoclusters that deposit some portion of their energy outside the sensitive region. A separate ‘central’ selection is made with $|\eta| < 2.5$ to allow comparison with the charged track observables.
- In the range $1.3 < |\eta| < 1.32$ a disproportionately large number of clusters are recorded with the majority of their energy deposited in the hadronic calorimeter [116]. The exact cause of this is unknown, but the η range is consistent with the transition from barrel to end-cap calorimeters. Therefore, in that $|\eta|$ range clusters are discarded if their hadronic energy fraction $f_{\text{had}} > 0.4$.

Note that there is no minimum p_{T} threshold for topocluster selection. This is deliberate, as it gives intrinsic cancellation of noise in the calorimeters. At the EM-scale the noise distribution is symmetric about zero, so including noise clusters with negative energy in $\sum E_{\text{T}}$ will approximately cancel the positive energy noise.

6.6. Correction procedure

After all selections have been made, the data must be corrected for detector effects, i.e. imperfect measurement. Firstly, the tracking performance of the ATLAS Inner Detector is corrected for:

Track reconstruction efficiency: To obtain the true number of charged particles, it is necessary to take inefficiencies of the track measurement and reconstruction into account. The efficiency, $\epsilon_{\text{trk}}(p_{\text{T}}, \eta)$, of the ATLAS detector to reconstruct a charged particle track was measured as a function of the particle p_{T} and η [104]. These efficiencies are corrected for by weighting each track with the following factor:

$$w_{\text{trk}} = \frac{1}{\epsilon_{\text{trk}}(p_{\text{T}}, \eta)}.$$

Secondaries, fakes, and acceptance: The fractions of tracks that are reconstructed from secondary particles ($f_{\text{sec}}(p_{\text{T}}, \eta)$), fake tracks ($f_{\text{fake}}(p_{\text{T}}, \eta)$), or particles migrating in from outside of the kinematic range ($f_{\text{OKR}}(p_{\text{T}}, \eta)$), have been measured [104]. These effects are corrected for by weighting each track with another factor:

$$w_{\text{sec,fake,OKR}} = (1 - f_{\text{sec}}(p_{\text{T}}, \eta)) \cdot (1 - f_{\text{fake}}(p_{\text{T}}, \eta)) \cdot (1 - f_{\text{OKR}}(p_{\text{T}}, \eta)),$$

as a product with the weight for the track reconstruction efficiency. For tracks with $p_{\text{T}} > 500 \text{ MeV}$, fakes and migrations from outside the kinematic region are negligible.

The underlying event observables are constructed after the tracking performance correction — note that each observable is a property of a whole event, rather than an individual particle. The distributions of these observables are corrected using a procedure based on MC, described in Chapter 7.

6.7. Monte Carlo samples

In this analysis, two sets of MC samples are used (see Section 5.1 for details). The first set is passed through the ATLAS detector simulation [117] to provide input for the MC-based correction. Particle interactions with the detector — and the creation of secondary particles — are simulated using a detailed model of the detector hardware within the

GEANT4 [118] software package. Having simulated the behaviour of particles within the detector, the detector’s response is also simulated in a process called ‘digitisation.’ If GEANT4 calculates that a particle has encountered detector hardware then the properties of that particle are passed to the corresponding digitisation simulator for that hardware. While GEANT4 is software for general use, digitisation simulators are written and tuned specifically to reproduce the electronic response of a particular sub-detector component. These simulated detector readouts are reconstructed in exactly the same way as real data from the detector, and the same underlying event observables are produced. The AMBT1 tune of PYTHIA 6 is used for this, as well as the default tune of Herwig++ version 2.5.0.

The simulation of the ATLAS detector does not reproduce topocluster behaviour entirely accurately. Since the MC-based correction relies on accurate simulation, MC topocluster energies have an η -dependent scaling applied to better match the real detector performance. This scaling was calculated by examining the decays of π^0 mesons [119]. The variation of this scaling with η is shown in Figure 6.5.

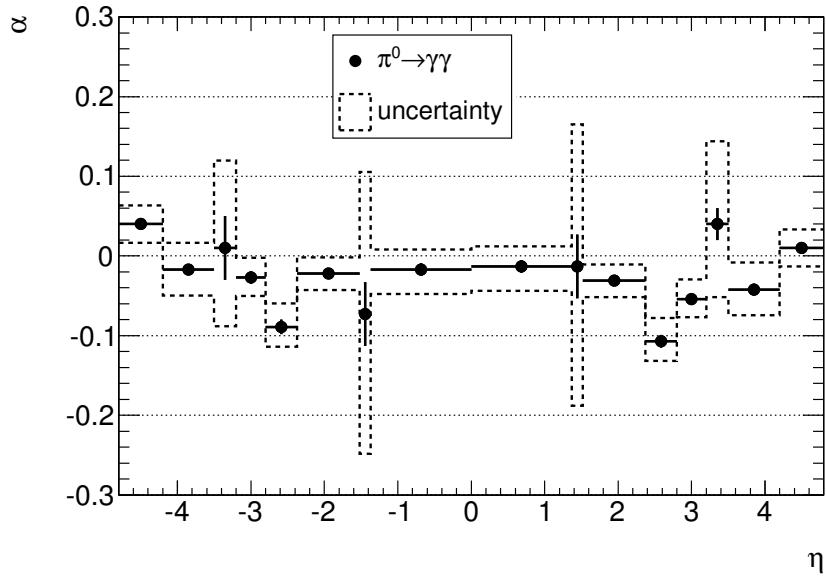


Figure 6.5.: The η -dependent correction to simulated topocluster energies, which are scaled by a factor of $1 + \alpha$ [119].

The recorded data — once corrected — are compared with a second set of MC samples to determine how well different MC models perform. These samples are used only at generator level: there is no simulation of detector effects. The AUET2B and DW tunes of PYTHIA 6 are used, as well as Pythia8 AU2. From the HERWIG generator family

there is HERWIG+JIMMY AUET2 and Herwig++ version 2.5.1 with the UE7-2 tune. The multi-legged MC generator ALPGEN is used as well.

Different objects are used when constructing underlying event observables at truth level, since tracks and topoclusters are defined with respect to the detector. Jets are constructed by applying the anti- k_T algorithm to all generated particles — rather than topoclusters — although the selection cuts are the same. In place of tracks, any charged particle is used, provided that it is marked as stable by the generator. The $|\eta| < 2.5$ and $p_T > 500$ MeV selections remain to match the acceptance for the data, but none of the reconstruction quality cuts are needed.

To construct the $\sum E_T$ observable, all stable particles are used. The $|\eta| < 4.8$ or $|\eta| < 2.5$ selections are applied to correspond to the data, and a p_T minimum threshold is introduced (calorimeter noise is not present at truth level). Charged particles must have $p_T > 200$ MeV, and neutral particles must have $p_T > 500$ MeV. Particles below these thresholds are unlikely to reach the ATLAS calorimeters due to the other material they will encounter first [116].

Chapter 7.

Bayesian unfolding

The complex design of the ATLAS detector (and other detectors in high energy physics) means that the properties of a particle may be measured in different ways depending on which region of the detector it travels through. Particles will also interact with the detector and its support structures, potentially losing energy, changing direction, or creating additional particles. Consequently, the measurements made by the detector may differ substantially from the true properties of the collisions that occurred within it, with a complex relationship between the two.

Insofar as possible, these effects are corrected by detailed first-principles treatments of individually understood detector effects such as track-finding efficiencies and jet scale calibration, as described in Chapter 6. However, attempting to correct for every effect this way may be impractical, and runs the risk that some unexpected effect is neglected. Unfolding can be used as the final step to produce a fully faithful correction of collider signatures to the particle level. The principle is that it is easier to calculate the output from the detector given the true properties of an event than to do the reverse. By simulating physics events in the detector, then simulating the detector itself, it is possible to map the true properties of an event onto those reported by the detector. Unfolding applies this mapping in reverse to measurements made by the detector in order to recover the true values of observables.

7.1. Mapping and unfolding

The mapping between true and measured (‘reconstructed’) properties of events is made using events from Monte Carlo generators, and the corresponding events after they have

been passed through the detector simulation. For a given variable the truth (as reported by the Monte Carlo generator) is plotted against the simulated reconstructed value. Repeating for many truth-reconstructed pairs builds up a two-dimensional histogram of the frequency that a particular MC truth gives a particular reconstructed value: see Figure 7.1 for an example. This is only valid for observables with a known connection between the true and reconstructed values, such as the properties of entire events where each event has a unique identifier. In cases where the connection is more subjective (e.g. associating energy in the calorimeter with the energy of specific truth particles) this unfolding technique is not appropriate.

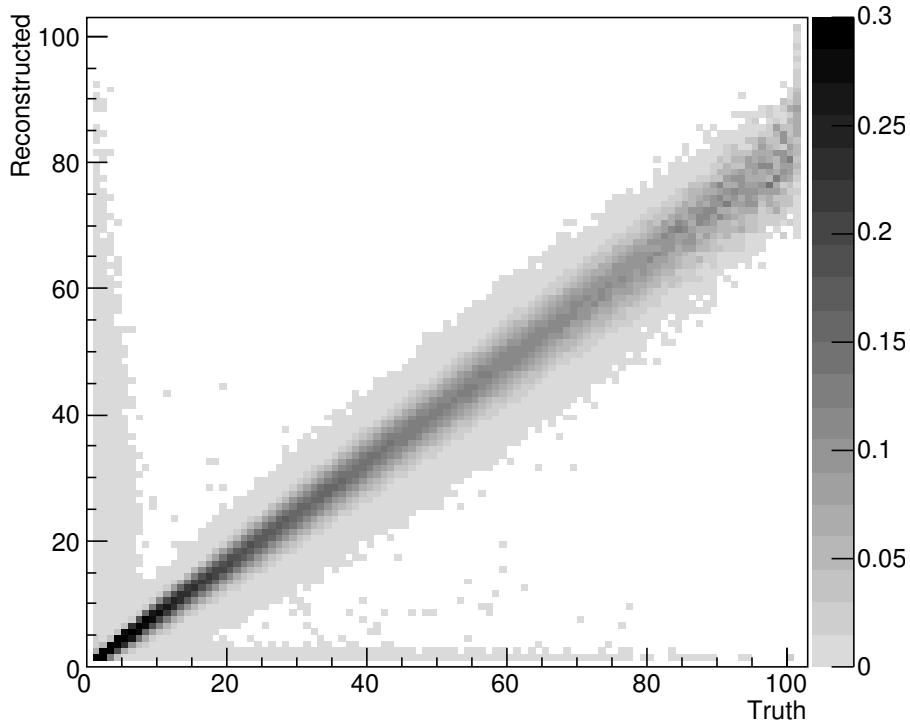


Figure 7.1.: An example of the smearing matrix – the mapping between MC truth and reconstructed values of a variable.

Elements on the diagonal of the truth-to-reconstructed mapping histogram give the probability of a particular MC truth value being correctly reconstructed, while the off-diagonal elements give the probability that a value belonging in bin i is mistakenly measured as belonging in bin j . This mapping is called the ‘smearing matrix’¹ S_{ij} , and

¹Often also referred to as the migration or response matrix.

the effect of a detector on the measurement of an observable can be thought of as:

$$\sum_i S_{ij} T_i = R_j, \quad (7.1)$$

where T_i is the true distribution of the observable and R_j is the reconstructed distribution.

Naively, it should be possible to recover the true distribution of a variable by inverting the smearing matrix and applying it to the measured distribution:

$$\sum_j S_{ij}^{-1} R_j = T_i. \quad (7.2)$$

Unfortunately there is no guarantee that S_{ij} is invertible or that the inversion has a unique solution. Attempting this method tends to give unphysical results, such as regions of negative probability density in the result distribution, or small changes in input causing disproportionate fluctuations in the output. Such problems are called ‘ill-posed’ or ‘improper,’ as discussed in References [120–122].

7.2. Alternatives to matrix inversion

Since exact inversion of the smearing matrix is not a viable approach, several approximate unfolding methods have been developed. Some notable methods are summarised below.

7.2.1. Bin-by-bin unfolding

This method is extremely simplistic: it takes the ratio of truth and reconstructed MC distributions, then multiplies the measured data distribution by that ratio. The ratio is calculated and applied separately for each bin of the histograms for each distribution — hence the name — so naturally they all must have the same binning. For the sake of clarity, the correction can be expressed as:

$$n(C_i^{\text{data}}) = n(R_i^{\text{data}}) \frac{n(T_i^{\text{MC}})}{n(R_i^{\text{MC}})}, \quad (7.3)$$

where $n(C_i^{\text{data}})$, $n(R_i^{\text{data}})$, $n(T_i^{\text{MC}})$, and $n(R_i^{\text{MC}})$ are the numbers of events in bin i of the corrected data, uncorrected data, truth MC and reconstructed MC histograms respectively.

While this method is easy to implement and is (as a consequence) widely-used, it suffers from two major problems. Firstly, it can be model-dependent: the ratio of the truth and reconstructed distributions may be affected by the shape of the truth distribution itself. Secondly, it does not explicitly account for migrations of events between histogram bins. Their effect will contribute to the truth-reconstructed MC ratio, but if the data and MC distributions are significantly different then migrations should manifest differently in each case.

7.2.2. Hit Backspace Once More (HBOM)

A substantial extension to the bin-by-bin unfolding technique, with far greater sophistication, the HBOM method [123] is relatively new but shows great potential. Detector effects for a particular distribution are parameterised in any suitable way, such as constructing a smearing matrix. Starting from the uncorrected data distribution (i.e. one application of detector effects), additional distributions are constructed where the detector effects have been applied two, three, four or more times. The value of a particular histogram bin varies with each application of the detector effects, and a polynomial function is fitted to these values. Extrapolating the function back to *no* applications of the detector effects should give the corrected data distribution.

At first examination this method seems extremely powerful, however, problems may appear as it is studied in greater detail. The method was developed too recently to have been used in this thesis, but is included for reference.

7.2.3. Smearing matrix regularisation

The problems with matrix inversion can be attributed to statistical fluctuations of the distributions involved. Regularisation is based on the assumption that the true distributions of observables are smoothly-varying, and that any high-frequency fluctuations are spurious. Many versions of regularised unfolding exist, but a relatively well-known one is presented in Reference [124]. Here Singular Value Decomposition (SVD) of the smearing matrix is used to remove the high-frequency fluctuations.

While approaches like this are theoretically well-motivated, they tend not to be widely used due to the offputting complexity of the mathematics they rely on. In the particular case of SVD unfolding the result may vary significantly with the value of the

regularisation parameter (effectively how much smoothing to perform), which casts some doubt on the reliability of the method.

7.2.4. Bayesian iterative unfolding

The rest of this chapter focusses on the method proposed by D'Agostini [125]. Here the full information of the smearing matrix is used (and so bin-to-bin migrations are accounted for), but there is no direct regularisation of the smearing matrix and so no need to choose a value for an abstract regularisation parameter. This method uses Bayes' theorem, which is usually stated as:

$$P(\text{hypothesis} \mid \text{result}) = \frac{P(\text{result} \mid \text{hypothesis}) P(\text{hypothesis})}{P(\text{result})}, \quad (7.4)$$

i.e. it relates the probability of a hypothesis being correct given an experimental result to the probability of that hypothesis producing the result.

Restating Equation 7.4 in the terminology of unfolding, each hypothesis is a bin T_i^{MC} in the MC truth distribution histogram and each result a bin R_j^{MC} in the reconstructed MC histogram:

$$P(T_i^{\text{MC}} \mid R_j^{\text{MC}}) = \frac{P(R_j^{\text{MC}} \mid T_i^{\text{MC}}) P(T_i^{\text{MC}})}{P(R_j^{\text{MC}})}. \quad (7.5)$$

Now $P(R_j^{\text{MC}} \mid T_i^{\text{MC}})$ corresponds to the smearing matrix S_{ij} – a mapping of truth to reconstructed values – and $P(R_j^{\text{MC}})$ is given by

$$P(R_j^{\text{MC}}) = \sum_i P(R_j^{\text{MC}} \mid T_i^{\text{MC}}) P(T_i^{\text{MC}}), \quad (7.6)$$

which is equivalent to Equation 7.1. We can use the unfolding matrix $P(T_i^{\text{MC}} \mid R_j^{\text{MC}})$ from Equation 7.5 to correct a measured distribution as follows:

$$n(C_i^{\text{data}}) = \sum_j P(T_i^{\text{MC}} \mid R_j^{\text{MC}}) n(R_j^{\text{data}}), \quad (7.7)$$

where $n(R_j^{\text{data}})$ indicates the number of events in bin j of the measured data histogram, and $n(C_i^{\text{data}})$ the number in bin i of the corrected histogram.

This method not only requires the smearing matrix as an input, but also a prior distribution $P(T_i^{\text{MC}})$. Using the Monte Carlo truth distribution for the prior seems natural, but raises the question of model-dependence. Clearly the truth distribution depends on the Monte Carlo generator that created it, and using different priors for the unfolding will give different results. Model-dependence is addressed with *iterative* unfolding: take each unfolded distribution and use it as the prior for the next round of unfolding. Repeating this process should cause the output to converge with the true data distribution. The convergence of the procedure is discussed in more detail in Section 7.3.3.

7.3. Imagiro

Imagiro [126] is a software package – designed for the analysis in Chapter 6, but available for general use – that provides Bayesian iterative unfolding with robust tests and safeguards. It attempts to achieve the best unfolding performance automatically, without requiring manual fine-tuning. The Monte Carlo and data events are loaded from ROOT format [127] files, and the unfolded distributions stored as ROOT histograms. Besides the unfolding algorithm there are a number of additional features, described in this section.

An alternative implementation of Bayesian iterative unfolding can be found in the RooUnfold package [128]. While the central algorithm is the same, the user experience is quite different in each case. Imagiro attempts to minimise the requirements on the user by providing robust solutions to anticipated tasks such as data access, plotting, systematic error propagation and so on. By contrast, RooUnfold provides the bare mathematical tools for unfolding as a library to be included in code written by the user. Thus, Imagiro might be simpler to use but RooUnfold more flexible, giving two complimentary approaches to the same problem.

7.3.1. Multiple priors

Iterating the unfolding procedure should remove the dependence of the result on the prior distribution. However, this assertion must be tested to gauge the success of any attempted unfolding. Imagiro is designed so that the whole process can be repeated for multiple different priors (i.e., Monte Carlo samples), and the results combined. The final value for each histogram bin is the mean of the results of unfolding with each prior

distribution:

$$n(F_i^{\text{data}}) = \frac{1}{N_{\text{priors}}} \sum_{d=1}^{N_{\text{priors}}} n(C_{i,d}^{\text{data}}), \quad (7.8)$$

where $n(F_i^{\text{data}})$ is the number of events in bin i of the final result histogram, $n(C_{i,d}^{\text{data}})$ is the number in bin i of the histogram unfolded using prior distribution d , and N_{priors} is the total number of prior distributions. Any model dependence is considered a systematic error, as discussed in Section 7.4.3.

Ideally, the smearing matrix should depend only on the detector simulation used, and should not depend on the Monte Carlo model. Therefore if the MC samples use the same detector simulation they can share a single smearing matrix, constructed using all events from all samples. This reduces the statistical uncertainty, but a separate matrix for each sample can be made instead in order to test the model-independence of the matrix itself.

7.3.2. Closure tests

Closure tests are used to demonstrate that the software is working correctly. Given that the smearing matrix maps the true MC distribution onto reconstructed MC, the software should be able to unfold the reconstructed MC distribution to reproduce MC truth. By comparing the result from unfolding with the actual MC truth distribution, Imagiro demonstrates that it is performing correctly. When multiple MC samples are provided, a closure test is performed with each one. The criterion for success is that $\chi^2/N_{\text{df}} < 1$ in the comparison of the corrected and true distributions, where N_{df} is the number of degrees of freedom.

When a single MC sample is used to produce the reconstructed distribution, prior distribution, and smearing matrix used for a closure test, the closure will be perfect. Therefore, this test is most useful when MC samples have been combined into a single smearing matrix, in order to compensate for a lack of statistics in the separate samples.

7.3.3. Conditions for convergence

Iterating the unfolding process causes the unfolded distribution to converge with the true data distribution. However, it also compounds the effects of statistical uncertainties in the smearing matrix. Therefore a larger number of iterations does not guarantee a

better result: eventually the true distribution will be obscured by random fluctuations. There is hence an optimum point when the unfolded distribution most closely describes the truth. Figure 7.2 shows an example of this behaviour, by unfolding a reconstructed MC distribution and each iteration comparing the result to the truth.

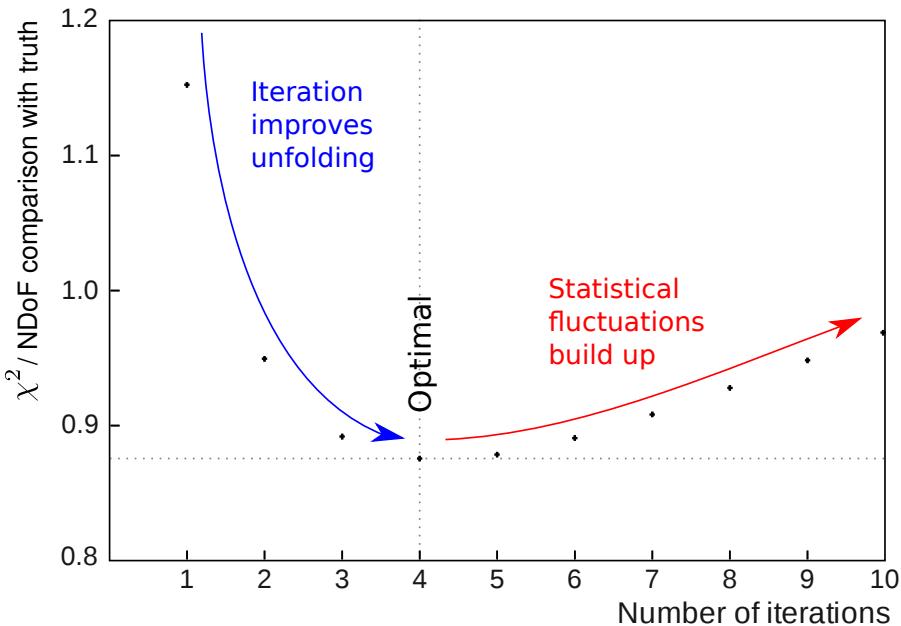


Figure 7.2.: While the first iterations of unfolding bring the result closer to the truth, there is an optimum point past which the accumulation of statistical errors becomes more significant. This is a representative example of results from Imagiro.

When multiple MC samples are provided, Imagiro attempts to identify the optimum point automatically. MC reconstructed distribution A is unfolded with MC truth distribution B as a prior. Each iteration the output is compared with MC truth A , which it should reproduce exactly in an ideal situation. This process is repeated for all combinations of true and reconstructed distributions from different MC samples, since different priors will give different unfolding performance. For each pair of MC samples 10 iterations are performed, and the iteration that gives either the minimum value of χ^2/N_{df} or the maximum result of Kolmogorov-Smirnov comparison – whichever comes first – is taken to be the optimum. Taking the mean of the optimal iteration numbers from each of these tests gives the number of iterations to use when unfolding the data. Note that unfolding a Monte Carlo reconstructed distribution with its own truth as a prior is not attempted because this is equivalent to a closure test. The unfolding may perform differently for different observables, so the optimum iteration number is calculated separately for each one.

In the ATLAS leading jet underlying event analysis two iterations are performed, minimising the statistical errors while still accounting for prior dependence. The effect of a higher and lower number of iterations is shown using MC closure tests in Appendix A.

7.3.4. Unfolding correlated variables

Unfolding corrects distributions, not individual events. When investigating the correlation between two observables (i.e. making a profile histogram) it is not possible to unfold each variable separately, as the correlations are not preserved in the process. To preserve these correlations Imagiro first makes a combined observable. As an example, if one observable has bins A , B and C , and a correlated observable has bins x , y , and z , then Imagiro would make a combined observable with bins Ax , Ay , Az , Bx , By , Bz , Cx , Cy and Cz . Unfolding this combined observable preserves the correlations.

Converting to and from this combined observable may introduce errors through poor choice of binning. Good binning of the y -axis observable is particularly important, since to construct a profile histogram conventionally only the x -axis is binned. Imagiro checks for mistakes here by constructing two profile histograms from the prior distribution: one without binning the y -axis; and one using the same binning as the combined observable. These two profiles are compared, and if any bin value differs by more than 1% the problem is reported.

Potentially this process could be abstracted to higher dimensions, but it is unlikely to be practical. Already for two correlated observables the smearing matrix histogram must have $(N_{\text{xbins}} \times N_{\text{ybins}})^2$ total bins, which may require a large number of MC events to populate. Adding more correlated observables will make the smearing matrix even larger, requiring even more MC events. There are alternative approaches that are not as demanding, at a cost of restricting the potential for event migrations.

7.3.5. Fake or missing events

It is possible for an MC truth event to have no corresponding partner in the list of reconstructed MC events, or vice versa. This is likely when events must pass a selection criterion based on uncorrected reconstructed observables. If a truth event has no reconstructed partner it is called ‘missed,’ and if a reconstructed event has no partner in truth it is called ‘fake.’ D’Agostini provides an explicit treatment for missed events by

introducing an efficiency ϵ_i [125], which modifies Equation 7.7 as follows:

$$n(C_i^{\text{data}}) = \frac{1}{\epsilon_i} \sum_j P(T_i^{\text{MC}} | R_j^{\text{MC}}) n(R_j^{\text{data}}), \quad (7.9)$$

The efficiency is defined as the probability that an event in bin i of the true distribution has a partner in *any* bin of the reconstructed distribution.

Imagiro can also correct for fake events by assigning them a prior probability. Each prior distribution has an extra bin added, and the smearing matrix contains an extra column giving the probability that a particular reconstructed event was fake.

7.4. Error treatment

Imagiro provides full statistical and systematic error calculations. Statistical errors on the MC distributions affect the final result, not just errors from the measured distribution. Model-dependence of the unfolded results gives a potential systematic error which may not be entirely removed by iteration. There is also the capability to propagate experimental systematic errors through the unfolding process. The different error calculations are discussed in this section.

7.4.1. Statistical error calculation

The calculation of statistical errors on the unfolded distribution is complex, since each bin of the measured distribution and prior distribution has some effect on each bin of the corrected distribution. Additionally, there is a contribution to the statistical error from the smearing matrix, since that is constructed with a finite number of MC events. The errors are compounded each iteration as the smearing matrix is re-used. Imagiro can exactly calculate the full covariance matrix [128], but the calculation may be impractically slow (execution time scales linearly with the number of histogram bins, and combined observables can require many bins). Therefore the option is provided to only calculate the diagonal elements of the covariance matrix.

7.4.2. Fast statistical error estimation

As an alternative to the full statistical error calculation, a near-instantaneous method is provided. The statistical error for each bin in the input distribution is scaled proportionately to the correction made to that bin value:

$$\sigma(C_i^{\text{data}}) = \sigma(R_i^{\text{data}}) \frac{n(C_i^{\text{data}})}{n(R_i^{\text{data}})}, \quad (7.10)$$

where $\sigma(C_i^{\text{data}})$ is the statistical uncertainty of bin i in the corrected histogram, and $\sigma(R_i^{\text{data}})$ is the corresponding error for the measured data. This method does not account for migrations of events between bins in the unfolding, or for the statistical uncertainty of the smearing matrix, so should be treated with caution. In particular, the uncertainty on the smearing matrix should be assessed. The covariance matrix is not calculated.

7.4.3. Intrinsic systematic error

Unfolding the measured data with different priors gives different corrected distributions. Iterating the unfolding reduces the difference between corrected distributions, but it is never eliminated. This remaining model-dependence is treated as a systematic error intrinsic to the unfolding process. Each bin of the final result histogram will have N_{priors} different unfolded values associated with it. The central value for that bin is given by the mean of all these unfolded values, as described in Section 7.3.1. The systematic error is calculated by ordering all the unfolded values and taking the range of the central 68% of them. If the ordered values for bin i are indexed by A running from smallest to largest, the systematic error is given by:

$$A_{\text{upper}} = \lceil N_{\text{priors}} \times 0.84 \rceil, \\ \text{sys}_{\text{upper}}(F_i^{\text{data}}) = n(C_{i, A_{\text{upper}}}^{\text{data}}) - n(F_i^{\text{data}}),$$

$$A_{\text{lower}} = \lfloor N_{\text{priors}} \times 0.16 \rfloor, \\ \text{sys}_{\text{lower}}(F_i^{\text{data}}) = n(F_i^{\text{data}}) - n(C_{i, A_{\text{lower}}}^{\text{data}}),$$

where $\text{sys}_{\text{upper}}(F_i^{\text{data}})$ and $\text{sys}_{\text{lower}}(F_i^{\text{data}})$ are the asymmetric systematic error ranges of bin i in the final result histogram, $n(F_i^{\text{data}})$ is the number of events in that bin, and

$n(C_{i,A}^{\text{data}})$ is the A^{th} largest number of events from bin i of each of the different corrected histograms.

7.4.4. Systematic error propagation

Systematic errors on the measured data distributions can be propagated through the unfolding in Imagiro using pseudo-experiments. When constructing a data distribution for unfolding, additional pseudo-experiment distributions are made by offsetting the input data values. The offset for each value is either sampled per-event from a Gaussian with a user-defined width, or is a specific value requested by the user. Absolute or fractional errors can be specified, and the user can use the value of one event-level observable as the error for another.

Each pseudo-experiment is unfolded with each available prior distribution, just like the measured data. Therefore each bin of the final result histogram will have $N_{\text{priors}} \times N_{\text{pseudo}}$ extra unfolded values associated with it, in addition to those from unfolding the measured data without offsets. The systematic error for that bin is calculated just as in Section 7.4.3, except that A_{upper} and A_{lower} are now given by:

$$\begin{aligned} N_{\text{unfolded}} &= N_{\text{priors}}(N_{\text{pseudo}} + 1), \\ A_{\text{upper}} &= \lceil N_{\text{unfolded}} \times 0.84 \rceil, \\ A_{\text{lower}} &= \lfloor N_{\text{unfolded}} \times 0.16 \rfloor. \end{aligned}$$

The final value for bin i is now given by the mean of all N_{unfolded} values for that bin, modifying Equation 7.8 as follows:

$$n(F_i^{\text{data}}) = \frac{1}{N_{\text{unfolded}}} \sum_{A=1}^{N_{\text{unfolded}}} n(C_{i,A}^{\text{data}}). \quad (7.11)$$

7.5. Demonstration with toy models

Two examples of results from unfolding with Imagiro are provided. The first is for a single toy model, shown in Figure 7.3. Events are given uniform distributions in the arbitrary observables X and Y , with X ranging from 0 to 10000, and Y from 0 to 1000. The values of X and Y are correlated such that $X = 10Y$. To create toy reconstructed distributions, a smearing is applied to each true value such that $X_{\text{reco}} = X_{\text{true}}(1 + q)$, and

the same for Y , where q is sampled each event from a Gaussian distribution of width 0.5 and mean 0.0. Imagiro is then used to recover the true distributions of X and Y . The unfolding is perfect because the true distributions of X and Y are used as the prior distributions – a closure test. The correlation between X and Y is recovered using the method described in Section 7.3.4.

In the second example, the observable X has three different possible distributions. Each is a Gaussian with width 500, with a mean of either 4000, 5000, or 6000. These are referred to as ‘Gauss 4,’ ‘Gauss 5,’ and ‘Gauss 6’ respectively. To create toy reconstructed distributions a smearing is applied to each true value such that $X_{\text{reco}} = X_{\text{true}} + q$, where q is sampled each event from a Gaussian distribution of width 1000 and mean 0.0. The reconstructed events from Gauss 5 are then unfolded with Imagiro to recover their true distribution. Each of the different possible distributions for X is used as a prior for the unfolding, as shown in Figure 7.4. The original distribution for X is recovered, with a systematic uncertainty given by the different results from unfolding with the different priors. The unfolding is repeated using only Gauss 4 and Gauss 6 as priors (i.e. without using knowledge of the correct result) with almost identical results, although the statistical errors are larger because there are fewer events in the smearing matrix.

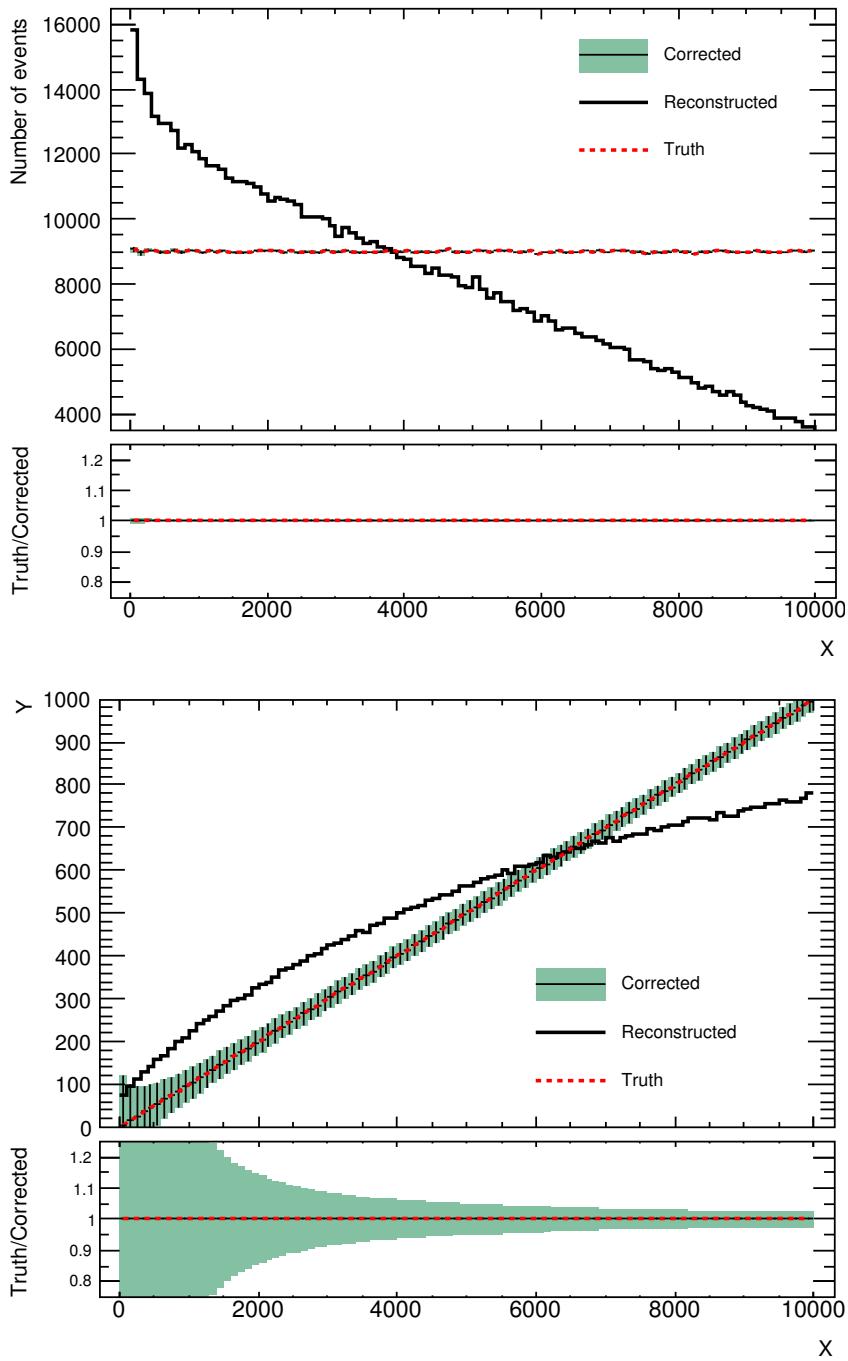


Figure 7.3.: Events have uniformly distributed values for the observables X and Y , and $X = 10Y$. A Gaussian smearing is applied, then Imagiro used to correct for this and recover the true variable distributions. In the upper plot, the points showing the corrected X distribution fall exactly on the horizontal line of the true distribution. In the lower plot, the points showing the corrected X vs Y distribution fall exactly on the line $X = 10Y$ of the true distribution. Both plots show perfect closure, despite each reconstructed distribution differing substantially from the truth. Corrected distributions show statistical errors as narrow bars, with a shaded background showing the statistical and systematic errors combined in quadrature.

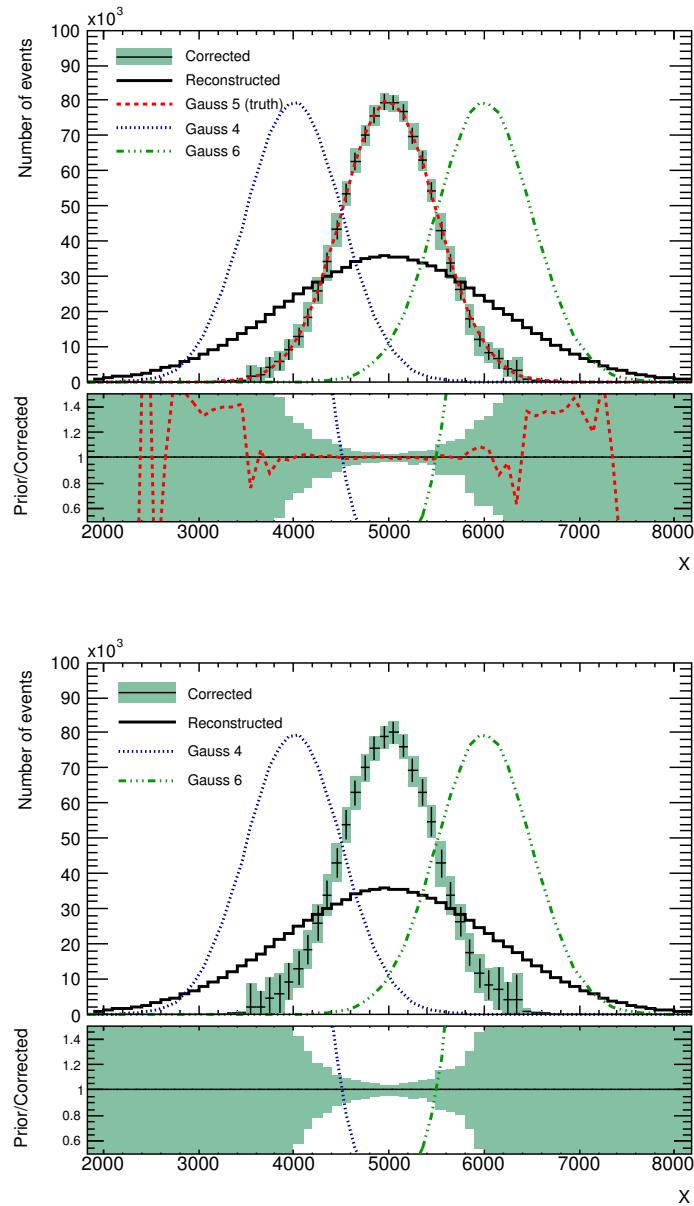


Figure 7.4.: In the upper plot, reconstructed events from Gauss 5 are unfolded using Gauss 4, Gauss 5 and Gauss 6 as priors. The corrected distribution closely matches Gauss 5. In the lower plot, reconstructed events from Gauss 5 are unfolded using only Gauss 4 and Gauss 6 as priors. Again, the corrected distribution closely matches Gauss 5, despite the fact that it was not used as a prior. Corrected distributions show statistical errors as narrow bars, with a shaded background showing the statistical and systematic errors combined in quadrature.

Chapter 8.

The Underlying Event: results

The underlying event observables defined in Table 6.1 are plotted as profile histograms, showing their dependence on another observable. Primarily their variation with the scale of the hard process — characterised by p_T^{lead} , the leading jet p_T — is shown, but $\langle p_T \rangle$ is also plotted with respect to the charged particle multiplicity N_{ch} to investigate soft processes. In these plots, all observables have been computed for both the inclusive leading jet and exclusive di-jet event selections. The corrected data are compared to the Monte Carlo samples described in Section 5.1.

To allow direct comparison between the transverse region, the trans-min/max regions (defined in Section 6.1), the different η ranges, and with other experiments with different angular acceptances, the raw quantities are divided by angular area in η - ϕ space to produce densities, $d^2 N_{\text{ch}} / d\eta d\phi$, $d^2 \sum p_T / d\eta d\phi$, and $d^2 \sum E_T / d\eta d\phi$. The central transverse observables are normalised by $\Delta\phi \Delta\eta = (2 \times \pi/3) \times (2 \times 2.5) = 10\pi/3$. The topocluster $\sum E_T$ is also shown for the full calorimeter η range, so its area normalisation is $\Delta\phi \Delta\eta = (2 \times \pi/3) \times (2 \times 4.8) = 19.6\pi/3$. The trans-max and min regions have only half the angular area since they only consider one side in $\Delta\phi$.

The distribution of each separate quantity is also shown, both for the full p_T^{lead} range (see Appendix D.1) and for restricted p_T^{lead} ranges (see Appendix D.2), as the profiles only provide information about the mean behaviours of underlying event observables.

8.1. ATLAS Underlying Event measurements

Here the profile plots for fully-corrected underlying event observables are shown, in comparison to six different MC models. The ratios of the MC predictions to the data are

shown at the bottom of the plots. In all cases the error bars show the statistical uncertainty while the shaded area shows the combined statistical and systematic uncertainties.

8.1.1. Charged particle $\sum p_T$ and multiplicity vs. p_T^{lead}

In Figures 8.1 and 8.2 the $\sum p_T$ and N_{ch} profiles (respectively) show similar behaviour. For the inclusive leading jet results, the total activity in the transverse region increases with the p_T^{lead} , while for the exclusive di-jet selection the activity decreases.

The inclusive results can be broken down further by examining the trans-max and min regions. The trans-max activity (for both $\sum p_T$ and N_{ch}) grows with p_T^{lead} , similar to the total transverse region trend, but the trans-min activity is almost constant over the whole range of p_T^{lead} . The slow rise of the $\sum p_T$ trans-min profile with p_T^{lead} may either indicate the contributions of multi-jet topologies (which have a larger cross-section for higher jet p_T) or higher- p_T MPI products as the hard process scale increases.

The relative insensitivity of the trans-min region to changes in p_T^{lead} suggests that pure MPI activity can indeed be modelled as plateauing as a function of hard process scale. Once the collisions are central the proton matter distributions overlap fully, reaching a maximum in this mechanism for the evolution of MPI (see Section 5.2.2). The MC models all reproduce the qualitative features of the data, but the best description (particularly of N_{ch}) is given by HERWIG+JIMMY, which even outperforms Herwig++.

The exclusive di-jet topology provides an alternative view of the same observables, where multi-jet events are explicitly excluded — including extra jets produced by the UE itself. There is little difference between the regions: all profiles fall with increasing p_T^{lead} , although the fall in trans-min is steepest. This behaviour, as opposed to p_T^{lead} -independence, implies that an important effect of the exclusive di-jet selection has been to exclude events where jets with $p_T > 20 \text{ GeV}$ were produced by MPI activity. Excluding MPI jets means that only events with low transverse activity pass the selection, leading to the profiles falling despite more MPI particle production.

This interpretation is verified by the 1D distributions of $\sum p_T$ and N_{ch} for $p_T^{\text{lead}} > 210 \text{ GeV}$ (see Figures D.6 and D.7): the peak (modal) positions of the distributions are relatively unaffected by the addition of the exclusive di-jet selection, but the lengths of the high activity tails are restricted, reducing the mean.

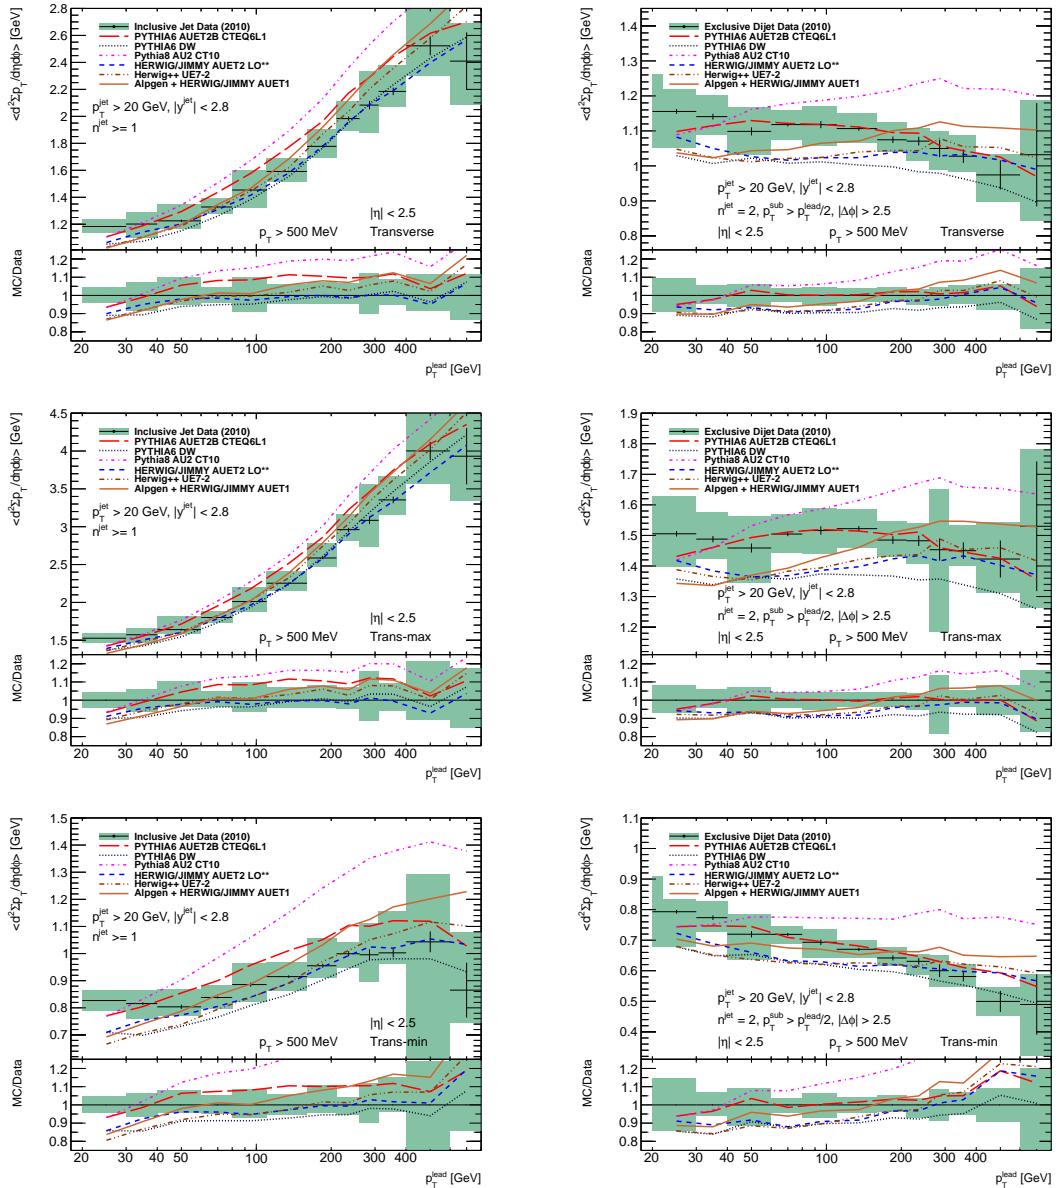


Figure 8.1.: Charged particle $\sum p_T$ vs. p_T^{lead} profiles, shown from top to bottom for the transverse, trans-max and trans-min regions. Inclusive leading jet and exclusive di-jet selections in the left and right columns respectively.

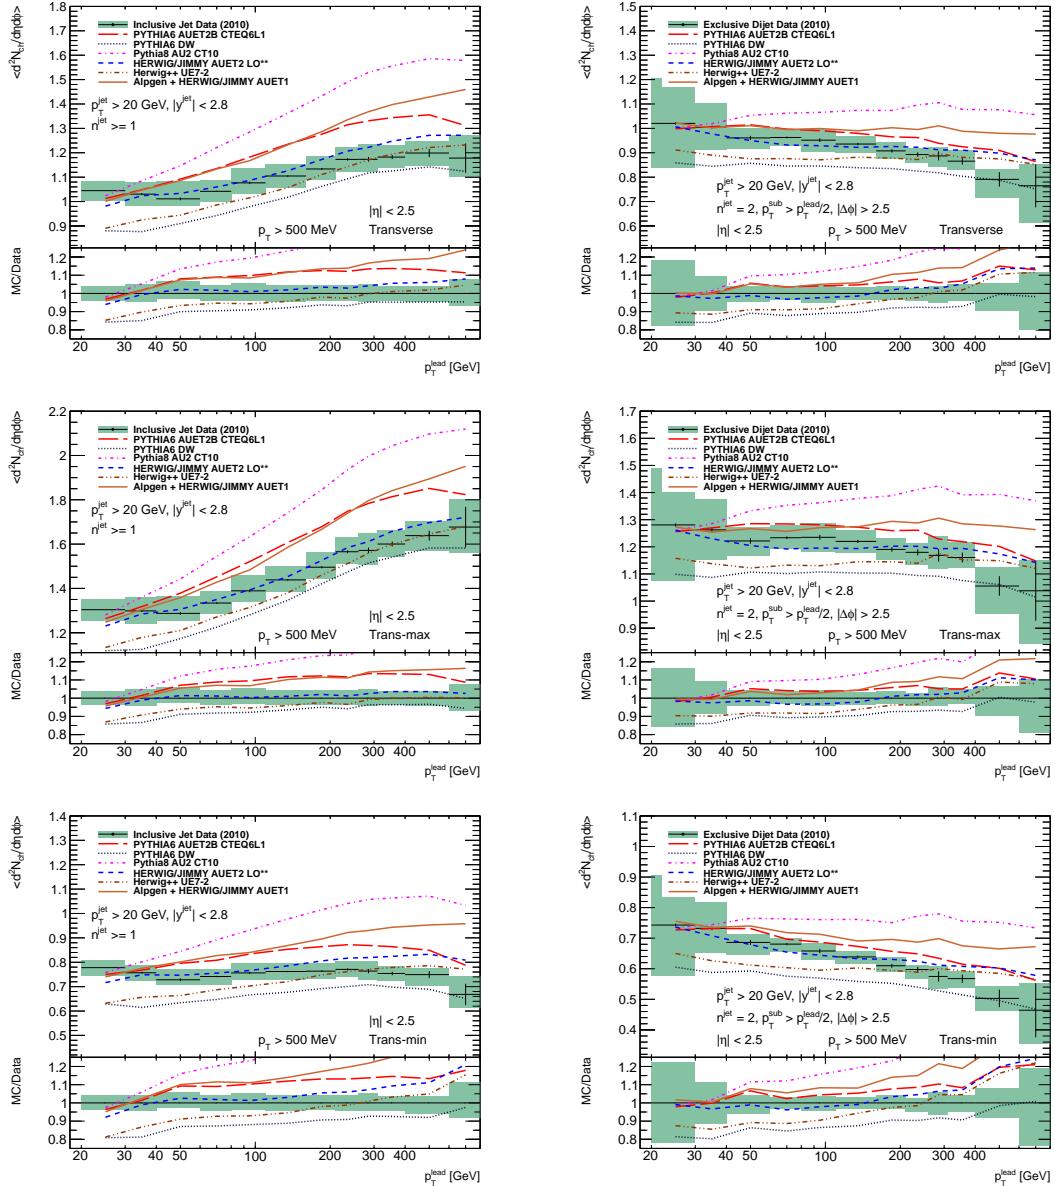


Figure 8.2.: Charged particle N_{ch} vs. p_T^{lead} profiles, shown from top to bottom for the transverse, trans-max and trans-min regions. Inclusive leading jet and exclusive di-jet selections in the left and right columns respectively.

Note that while the `HERWIG+JIMMY` description of the N_{ch} profiles is still good after the exclusive di-jet selection, `PYTHIA 6 AUET2B` has dramatically improved. In the exclusive $\sum p_{\text{T}}$ profiles AUET2B performs best of all the tunes.

8.1.2. Charged and neutral particle $\sum E_{\text{T}}$ vs. $p_{\text{T}}^{\text{lead}}$

The corrected topocluster $\sum E_{\text{T}}$ is shown in both the full η acceptance range (Figure 8.3), and in the central region (Figure 8.4), for the inclusive leading jet and exclusive di-jet topologies. For the central $|\eta|$ range the results are similar to the track-based $\sum p_{\text{T}}$ observable, with the inclusive profiles rising with $p_{\text{T}}^{\text{lead}}$, and the exclusive profiles falling. The MC models show qualitative agreement, again with `HERWIG+JIMMY` and `PYTHIA 6 AUET2B` performing best. However, the full $|\eta|$ range plots show more substantial disagreement between MC and data: the MC models undershoot the observed level of activity at low $p_{\text{T}}^{\text{lead}}$ values in both the inclusive and exclusive selections. Perhaps this discrepancy is not surprising, as the models have only been tuned to measurements from the central $|\eta|$ range.

The ratio of the charged particle $\sum p_{\text{T}}$ to central topocluster $\sum E_{\text{T}}$ approximately indicates the fraction of charged particles in underlying event activity. This ratio appears to be independent of $p_{\text{T}}^{\text{lead}}$, and is described within uncertainties by all MC models, as shown in Figure 8.5. It is notable that `HERWIG+JIMMY` and `ALPGEN` (which uses `HERWIG+JIMMY` for soft-QCD modelling) have similar behaviour, distinct from the other MC models.

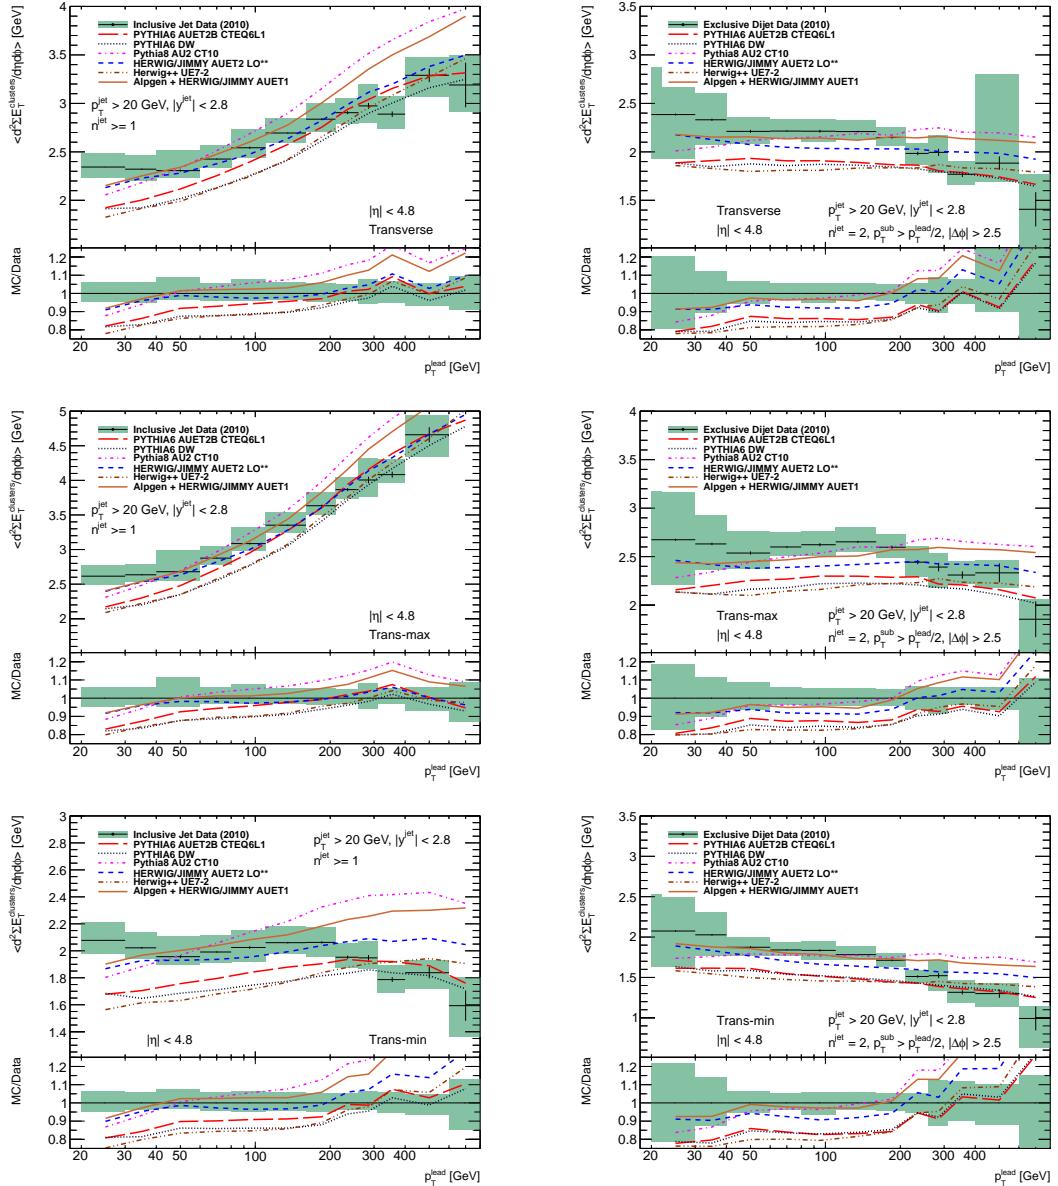


Figure 8.3.: Topocluster $\sum E_T$ vs. p_T^{lead} profiles ($|\eta| < 4.8$), shown from top to bottom for the transverse, trans-max and trans-min regions. Inclusive leading jet and exclusive di-jet selections in the left and right columns respectively.

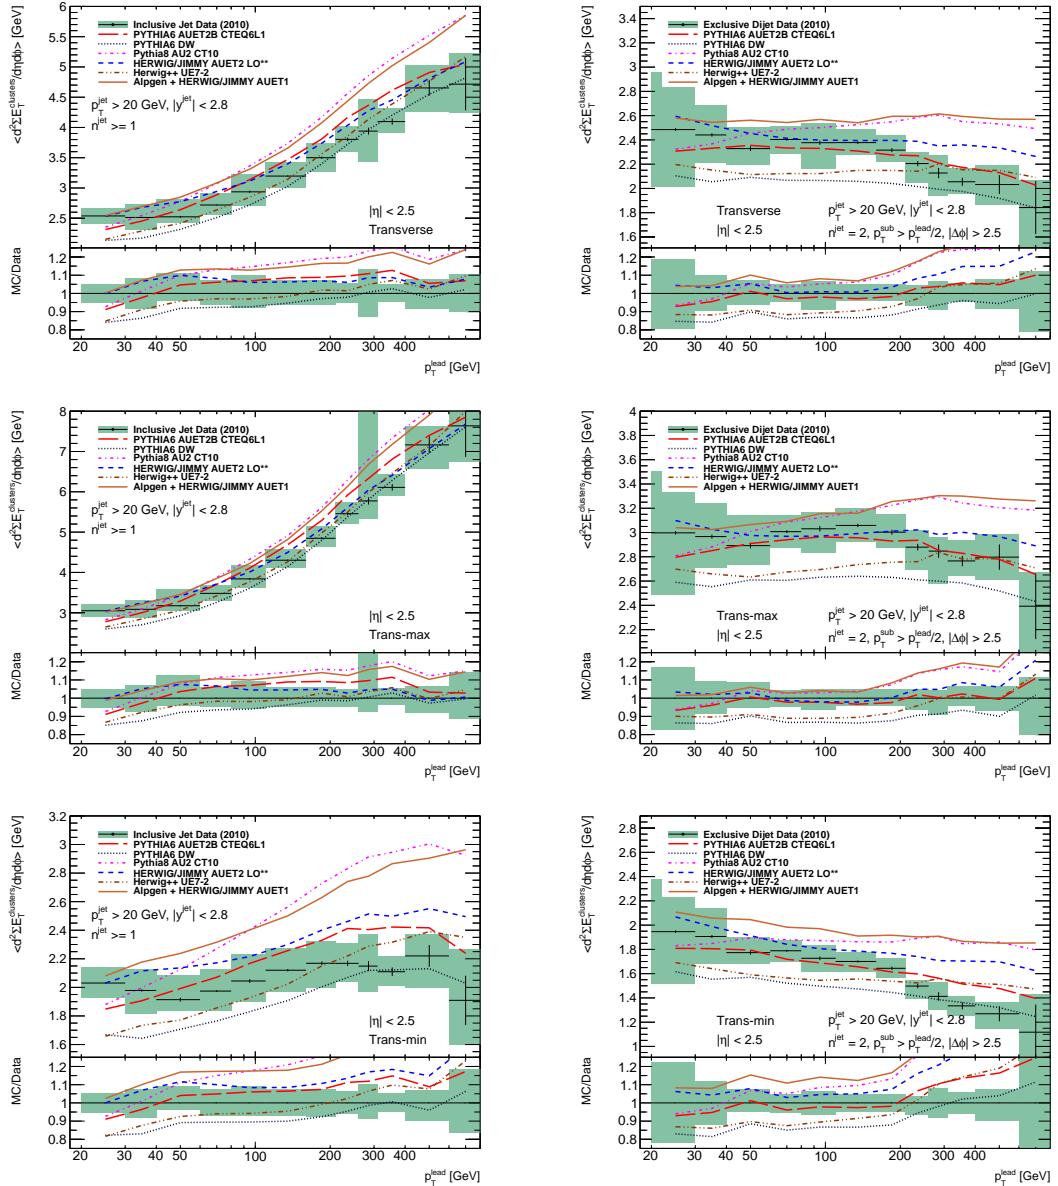


Figure 8.4.: Topocluster $\sum E_T$ vs. p_T^{lead} profiles ($|\eta| < 2.5$), shown from top to bottom for the transverse, trans-max and trans-min regions. Inclusive leading jet and exclusive di-jet selections in the left and right columns respectively.

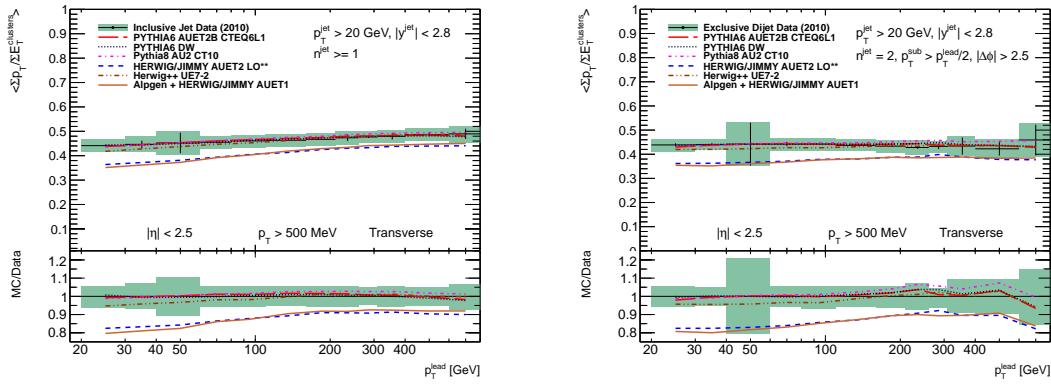


Figure 8.5.: Charged particle $\sum p_T$ and central topocluster $\sum E_T$ ratio vs. p_T^{lead} . Inclusive leading jet and exclusive di-jet selections on the left and right respectively.

8.1.3. Charged particle $\langle p_T \rangle$ vs. p_T^{lead} and N_{ch}

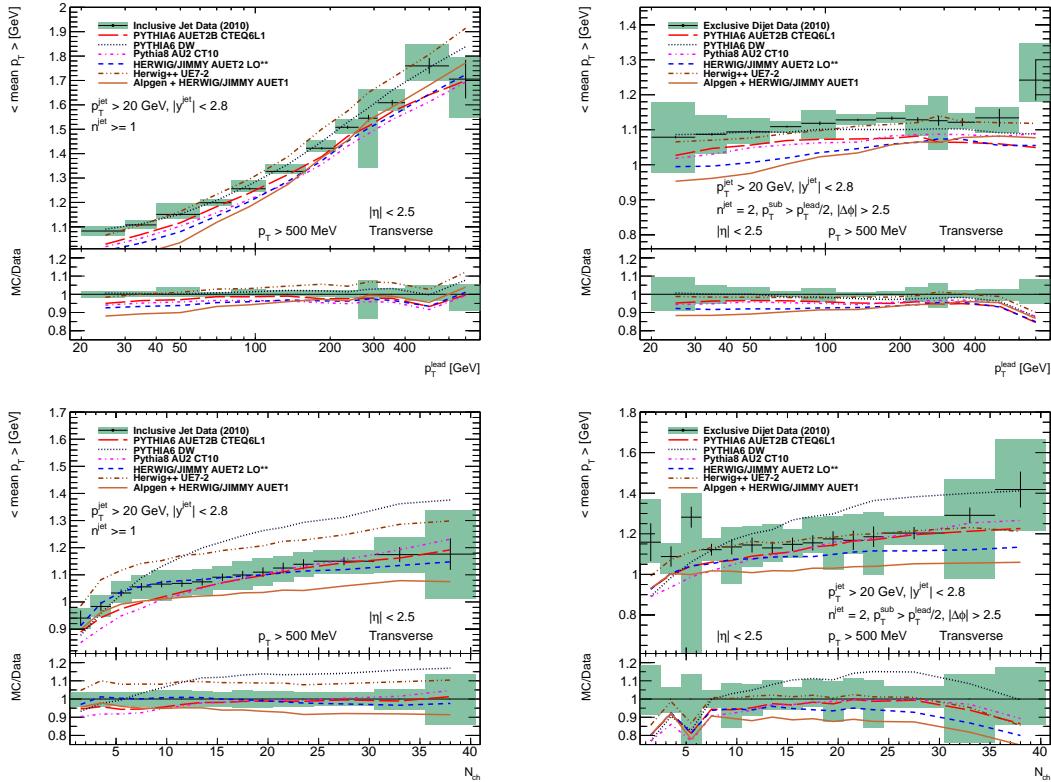


Figure 8.6.: Charged particle $\langle p_T \rangle$ profiles vs. p_T^{lead} (top) and N_{ch} (bottom). Results for the transverse region, inclusive leading jet and exclusive di-jet selections in the left and right columns respectively.

The $\langle p_T \rangle$ vs. p_T^{lead} profiles in Figure 8.6 have very different behaviour between the inclusive leading jet and exclusive di-jet event selections. In the inclusive case $\langle p_T \rangle$ rises

with the increasing let jet p_T , but the exclusive result is independent of p_T^{lead} within uncertainties. Again, this can be explained by MPI jets or multi-jet hard processes contributing to the inclusive result. Herwig++ describes the data well, but so does the old DW tune of PYTHIA 6, surprisingly outperforming the much more recent AUET2B tune.

Examining $\langle p_T \rangle$ vs. N_{ch} reveals more about parton showers and hadronisation than MPI, and so the results for the two selections are similar (although the more limited statistics for the exclusive selection lead to larger errors). In both cases $\langle p_T \rangle$ rises with N_{ch} , but in the inclusive leading jet profile the gradient changes at $N_{\text{ch}} \simeq 5$. It is possible that this behaviour is also present in the exclusive selection, but is just concealed by the uncertainties — certainly the MC predicts it. The inclusive result is very well described by HERWIG+JIMMY, reflecting its overall good performance.

8.2. Comparing the inclusive and exclusive selections

In general, the inclusive leading jet selection profiles show increasing underlying event activity with p_T^{lead} , whereas the activity is constant or falling in the exclusive di-jet profiles. Using the $\sum p_T$ observable as an example, the origin of this behaviour can be shown. Figure 8.7 gives a side-by-side comparison of the $\sum p_T$ profiles for these two selections, using the same axis ranges for each.

Figure 8.8 compares the distribution of $\sum p_T$ values directly, showing that the inclusive selection has a long tail of events with high $\sum p_T$, while in the exclusive selection the distribution rapidly falls to zero. The long tail in the inclusive selection is dominated by events with high p_T^{lead} , as shown in Figure 8.9. Each distribution has a maximum at around 8–10 GeV (like the distribution from the exclusive selection); it is just the tail that becomes more significant.

Since the exclusive selection removes events with a third jet with $p_T > 20$ GeV, it seems reasonable to conclude that the long tail of $\sum p_T$ values in the inclusive selection is due to contributions from these additional jets. That is consistent with the observation that the $\sum p_T$ distribution broadens with p_T^{lead} , as multijet events are more likely in higher momentum-transfer collisions.

The contribution of jets to the measured underlying event activity can be illustrated by varying the third jet veto on the exclusive di-jet selection. By varying the maximum

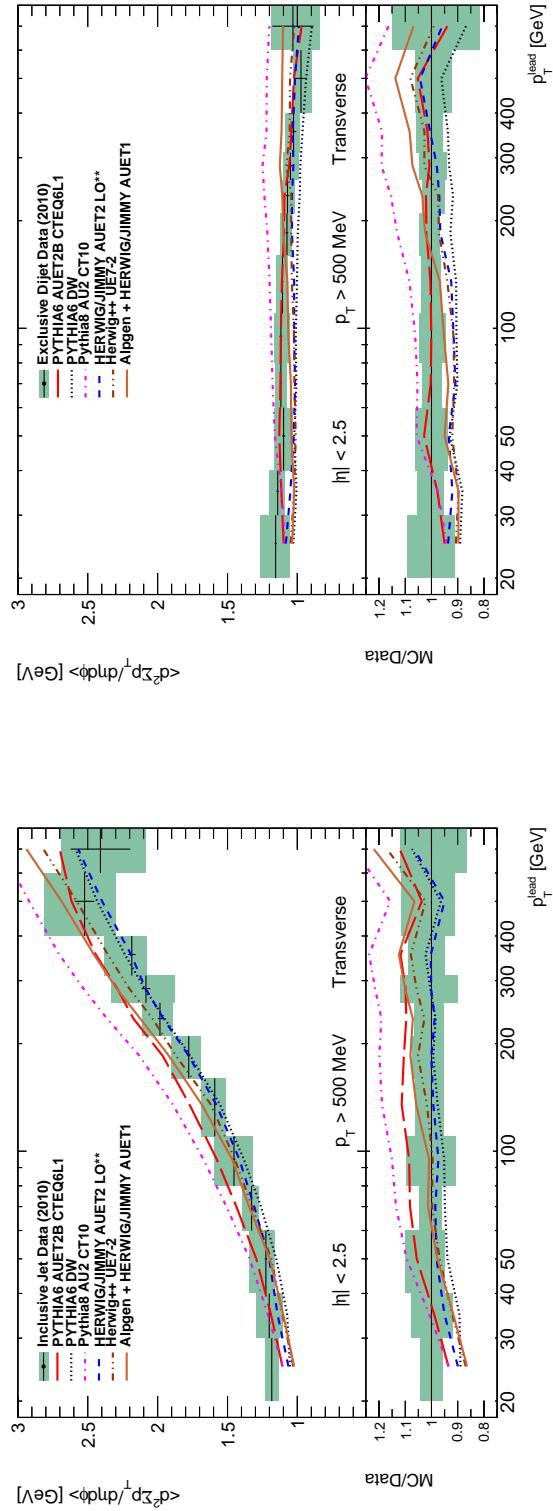


Figure 8.7.: $\sum p_T$ vs p_T^{lead} profiles in the transverse region for inclusive leading jet (left) and exclusive di-jet (right) selections.

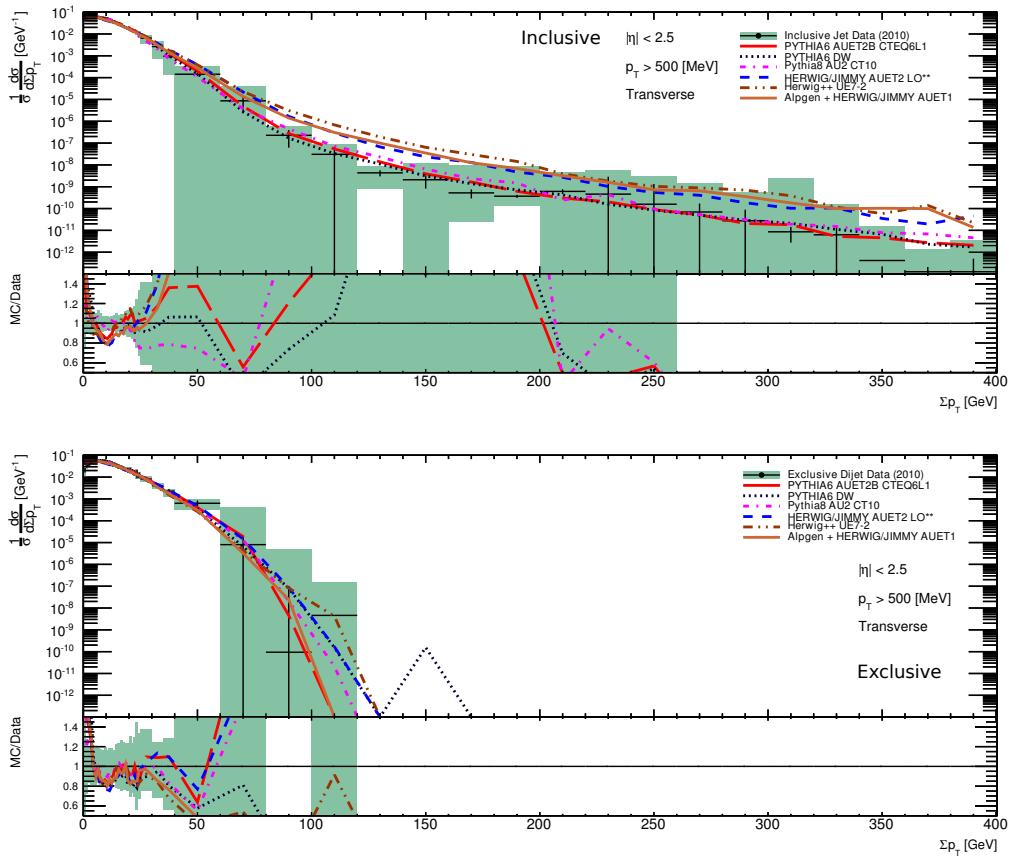


Figure 8.8.: The $\sum p_T$ distribution in the transverse region for inclusive leading jet (top) and exclusive di-jet (bottom) selections.

third jet p_T upwards from the initial value of 20 GeV, events with more jet activity will be included in the selection. As shown in Figure 8.10, this leads to a corresponding rise in underlying event activity. The effect of each third jet p_T cut appears near the corresponding value of p_T^{lead} , implying that some of the jets that affect the transverse region have comparable momentum to the leading jet itself. This is confirmed in Figure 8.11, which shows the jet p_T and multiplicity distributions for the transverse region of the inclusive leading jet selection. Around 1% of events contain one or more jets in the transverse region, and jet $p_T > 100$ GeV is quite possible.

The motivation for studying the trans-min region is that it should be less susceptible to contributions from multijet topologies. This hypothesis can be tested by examining the $\sum p_T$ distribution in the trans-min region, for inclusive and exclusive selections, as shown in Figure 8.12. In the inclusive selection the trans-min region still shows a long tail in $\sum p_T$, although it is much less significant than in the combined transverse region. Comparison with the exclusive selection again suggests that this long tail is due to jet

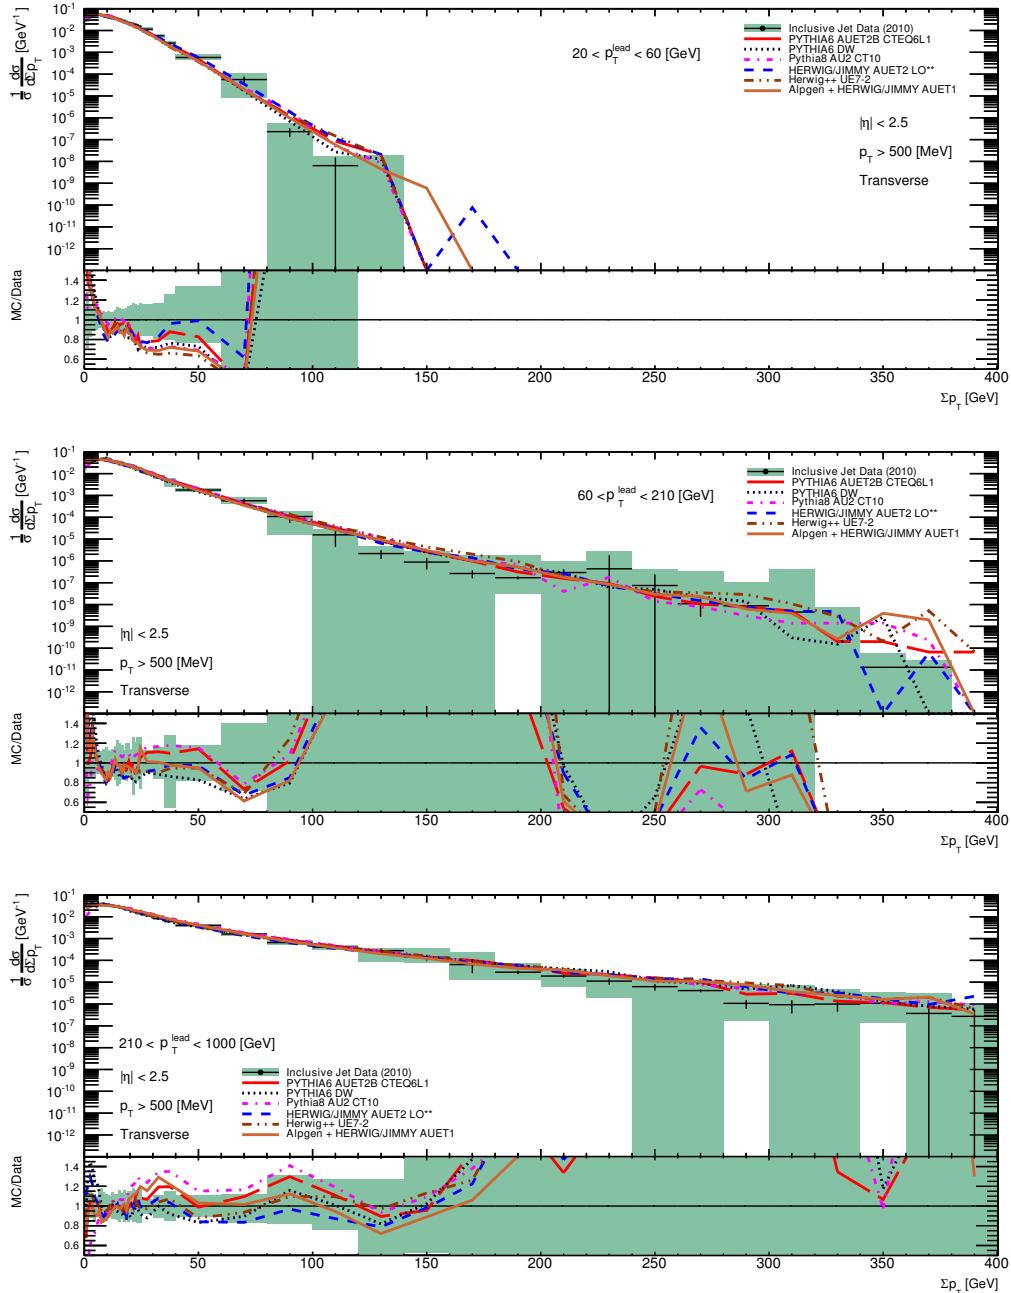


Figure 8.9.: The $\sum p_T$ distribution in the transverse region for the inclusive leading jet selection, shown here in three different ranges of p_T^{lead} : 20–60 GeV (top), 60–210 GeV (middle), and 210–1000 GeV (bottom).

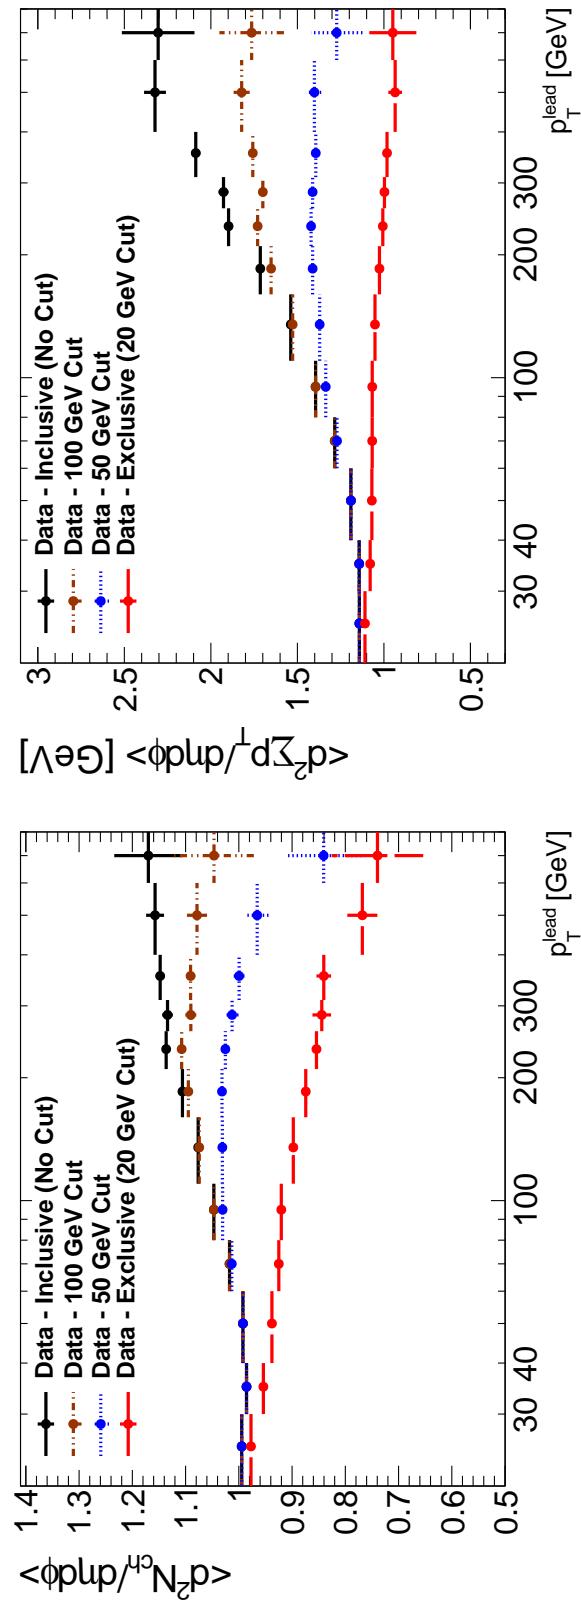


Figure 8.10: The transition from inclusive leading jet to exclusive di-jet selections at detector level, for transverse region N_{ch} (left) and $\sum p_T$ (right) profiles by varying the maximum allowed p_T for a third jet in the event [129]. The error bars represent the statistical error only.

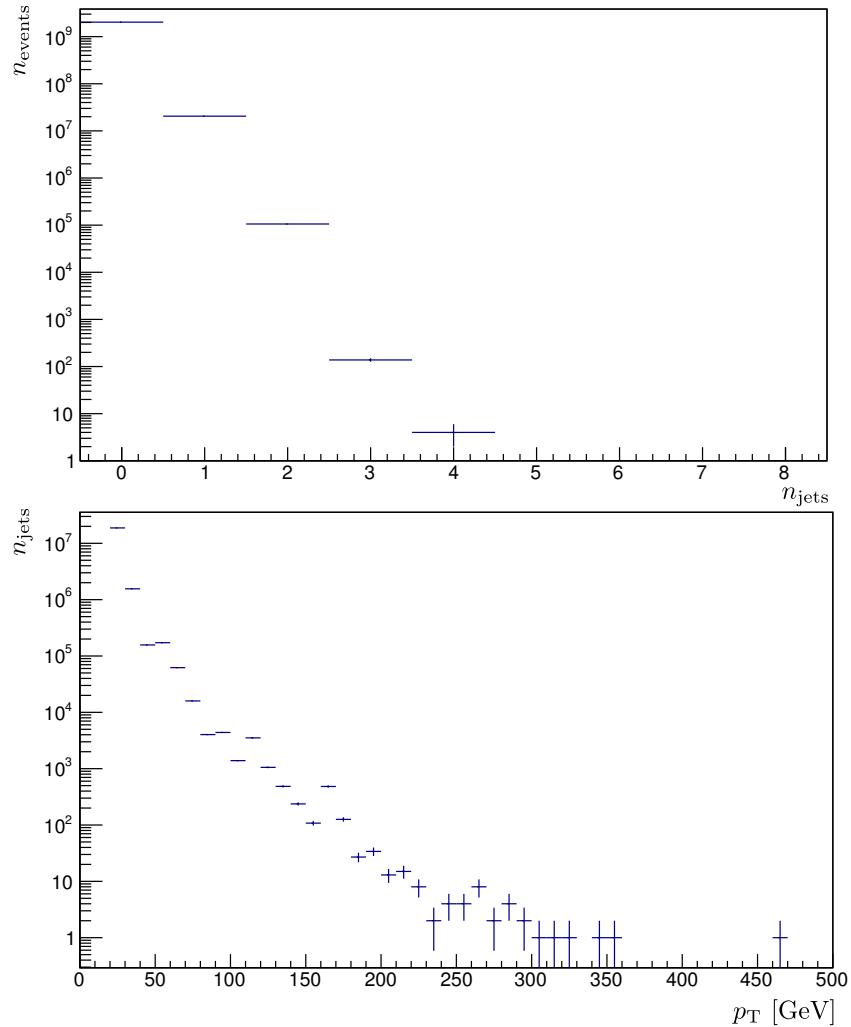


Figure 8.11.: The multiplicity (top) and p_{T} (bottom) distributions of jets found in the transverse region of events passing the inclusive leading jet selection. These results are presented at detector level.

activity, and so the trans-min region can indeed be said to be less affected by jets. This explains the relatively similar results in the trans-min region for the two selections.

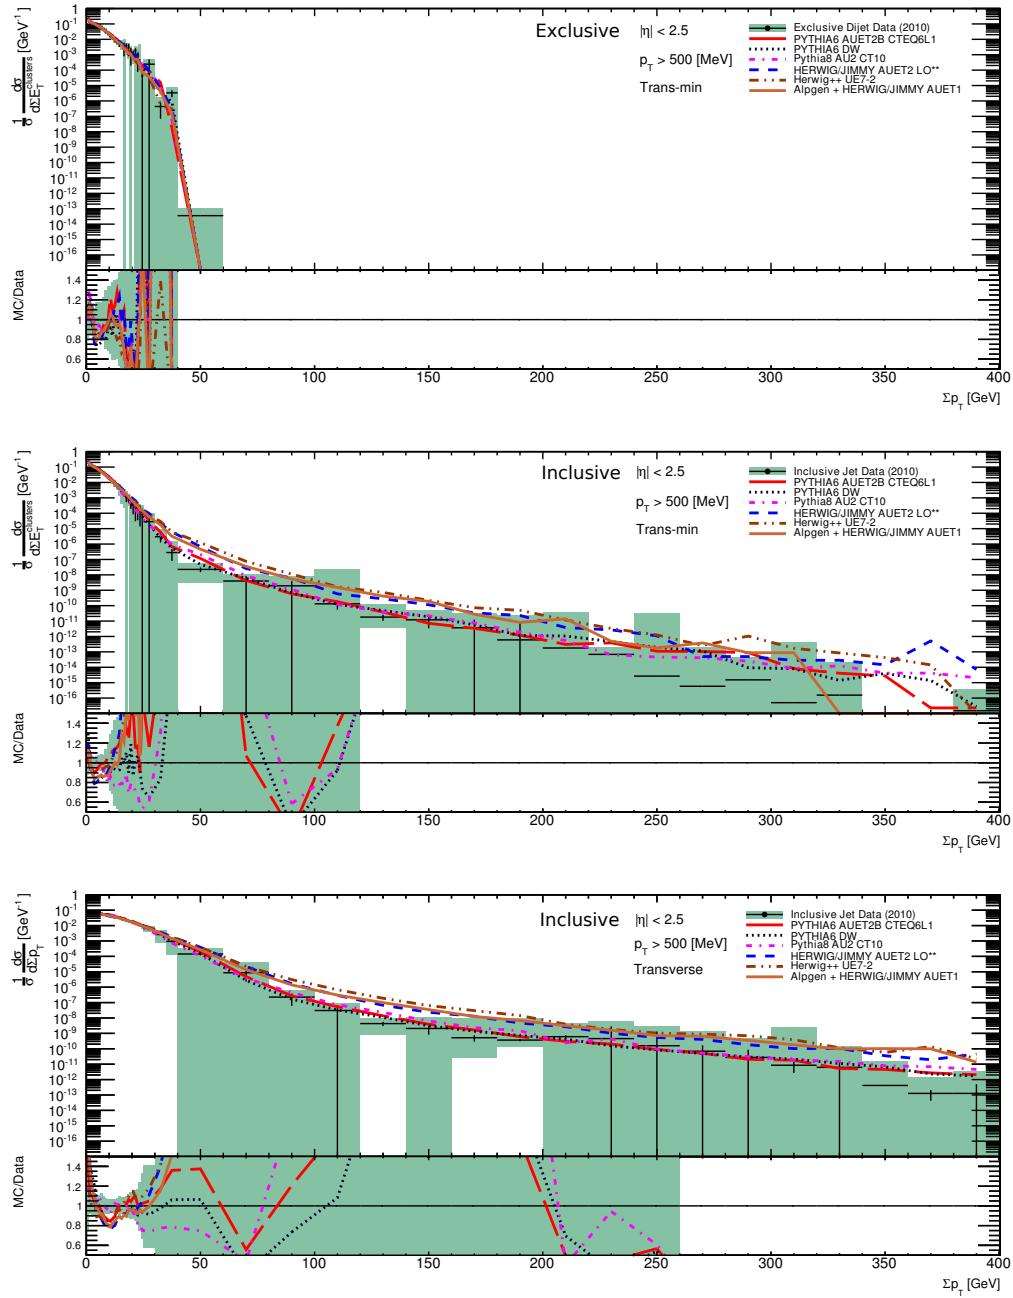


Figure 8.12.: Comparing the $\sum p_T$ distribution in the trans-min region for exclusive di-jet (top) and inclusive leading jet (middle) selections shows there is still a long tail of high $\sum p_T$ events in the latter case. The tail in the trans-min region is several orders of magnitude less significant than in the combined transverse region (bottom).

8.3. Data stability

The use of different triggers in different ranges of p_T^{lead} (see Table 6.2) could potentially create a bias in the results. In early 2010 the collision rate was low enough that there were no trigger prescales, and so the more inclusive triggers (L1_MBTS_1 and the lower JX triggers) recorded a large number of events. Higher collision rates later in the year required larger prescales for these triggers — so they recorded fewer events — but also meant there were potentially more events that could pass the higher JX trigger thresholds. Overall, this means that events selected for this analysis with higher p_T^{lead} are more likely to have been recorded later in 2010. This effect is shown in Figure 8.13.

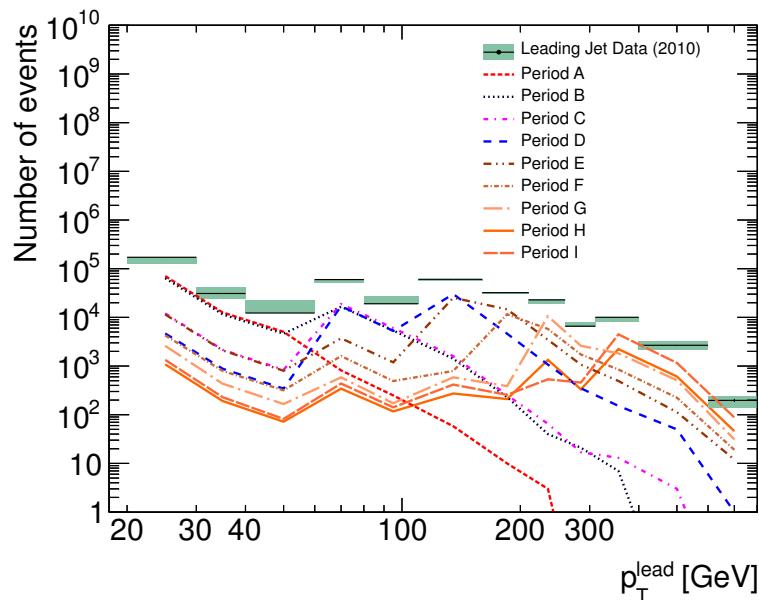


Figure 8.13.: The number of events from each data period is shown as a function of the p_T^{lead} . Later periods contribute more events to higher p_T ranges.

The potential bias arises as the beam conditions were changing throughout 2010, particularly the expected amount of pile-up. If beam background or pile-up does effect the measurement then the effect would be concentrated in the higher p_T^{lead} ranges. To test this, the observables measured in different data periods were compared in the p_T ranges where those periods provide the most events, as follows:

- Periods A and B in p_T^{lead} 20 – 60 GeV.
- Periods B, C and D in p_T^{lead} 60 – 110 GeV.

- Periods D and E in p_T^{lead} $110 - 160 \text{ GeV}$.
- Periods E and F in p_T^{lead} $160 - 260 \text{ GeV}$.
- Periods F and G in p_T^{lead} $210 - 310 \text{ GeV}$.
- Periods G, H and I in $p_T^{\text{lead}} > 310 \text{ GeV}$.

Within statistical uncertainty, no variation between data periods is observed. An example of this comparison is given in Figure 8.14, and comparisons are made for all underlying event observables in Appendix B.

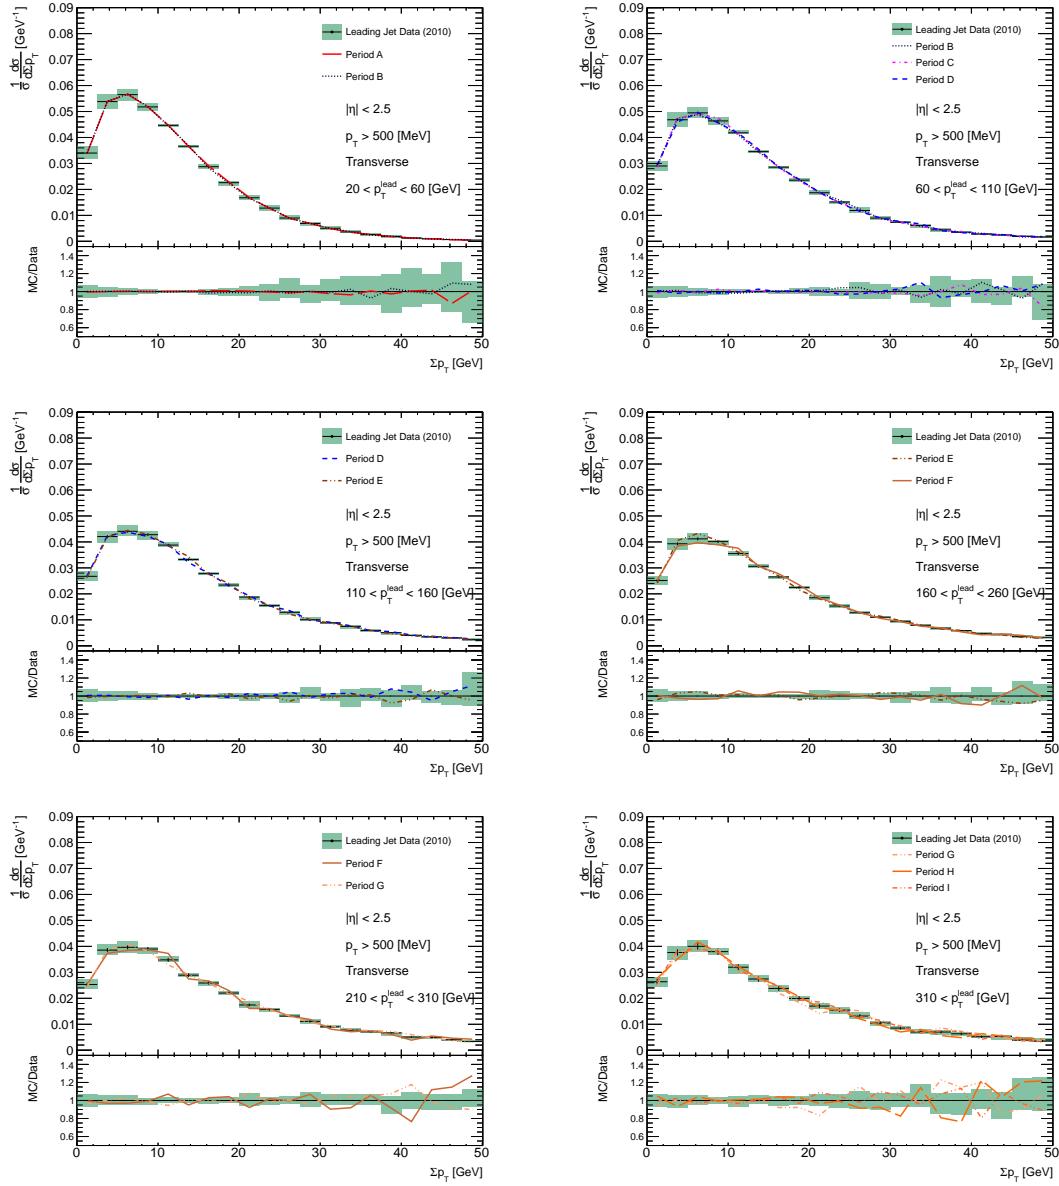


Figure 8.14.: Comparison of the charged particle $\sum p_T$ distribution for different data periods, in the p_T^{lead} ranges indicated.

8.4. Systematic errors

Many different potential sources of uncertainty in the leading jet underlying event measurements were investigated, as described below. Some were found to be insignificant, but others are treated as systematic errors on the final result. The errors were propagated through the unfolding process using pseudo-experiment distributions, as described in Section 7.4.4. Most error values varied from event to event, and so only two pseudo-experiments (plus and minus one standard deviation) were used for each source of error: multiple sets of error values would require an impractical amount of storage. Uncorrelated sources of error were combined in quadrature.

8.4.1. Background

Background contributions from sources other than proton–proton collisions were evaluated using events from trigger streams in which no real collision candidates are expected. The L1_J5_Uncalibrated trigger selects events with only a single proton bunch passing through the detector, and L1_J5_Empty selects events with no bunches in the detector at all. Cosmic ray events were also selected using the CosmicCalo trigger. In a sample of events from these triggers, no jets that satisfy the selection criteria were found. Given the observed stability of the data with changing beam conditions (see Section 8.3), the background rates across the entire data period are considered negligible after the jet quality cuts [130].

8.4.2. Tracks per vertex

The possible bias due to requiring a minimum of 5 tracks per vertex (compared to 2 tracks for pile-up vertex rejection) was studied and found to be negligible. This cut reduces the overall number of events by only $\approx 0.025\%$, as shown in Figure 8.15. The effect on the underlying event observables is of order 10^{-3} . Although a great many vertices are expected with < 5 associated tracks, the 20 GeV jet requirement eliminates almost all of these independently of the vertex track number cut.

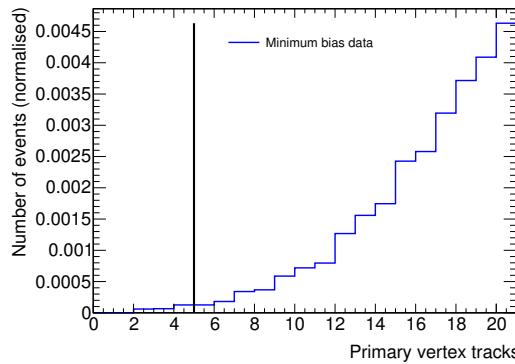


Figure 8.15.: The number of tracks associated with the primary vertex of an event that passes all other selection criteria [129].

8.4.3. Merged vertices

The z -coordinate distribution of pile-up vertex positions with respect to the primary collision vertex is shown in Figure 8.16. If the distance between the primary and pile-up vertex is less than 10 mm, the reconstruction algorithm may be unable to resolve the two. Consequently the two vertices are merged into one, giving rise to the dip in the centre of the vertex distribution. The fraction of vertices which are merged with the primary vertex is calculated by dividing the integral of the vertex distribution by the integral of a fitted Gaussian distribution:

$$f = 1 - \frac{\text{Integral over vertex distribution}}{\text{Integral over Gaussian}}.$$

The merged vertex fraction is calculated in this way for the data sample, and for a PYTHIA 6 AMBT1 MC sample, with pile-up added in simulation. In the simulated pile-up, the fraction of merged vertices is about 0.015, compared to 0.0095 for the combined 2010 data set. Therefore, comparing MC events with and without pile-up should give an overestimate of its impact on the underlying event measurement. A systematic error of 1% is introduced based on this comparison.

Figure 8.16 also shows that in the early data periods the distribution of pile-up vertices is not well-fitted by a single Gaussian. This is probably because of the rapid changes in the beam conditions over these periods. Since the early periods have very low pileup this should not be a cause for concern — the effect of merged vertices will be small regardless.

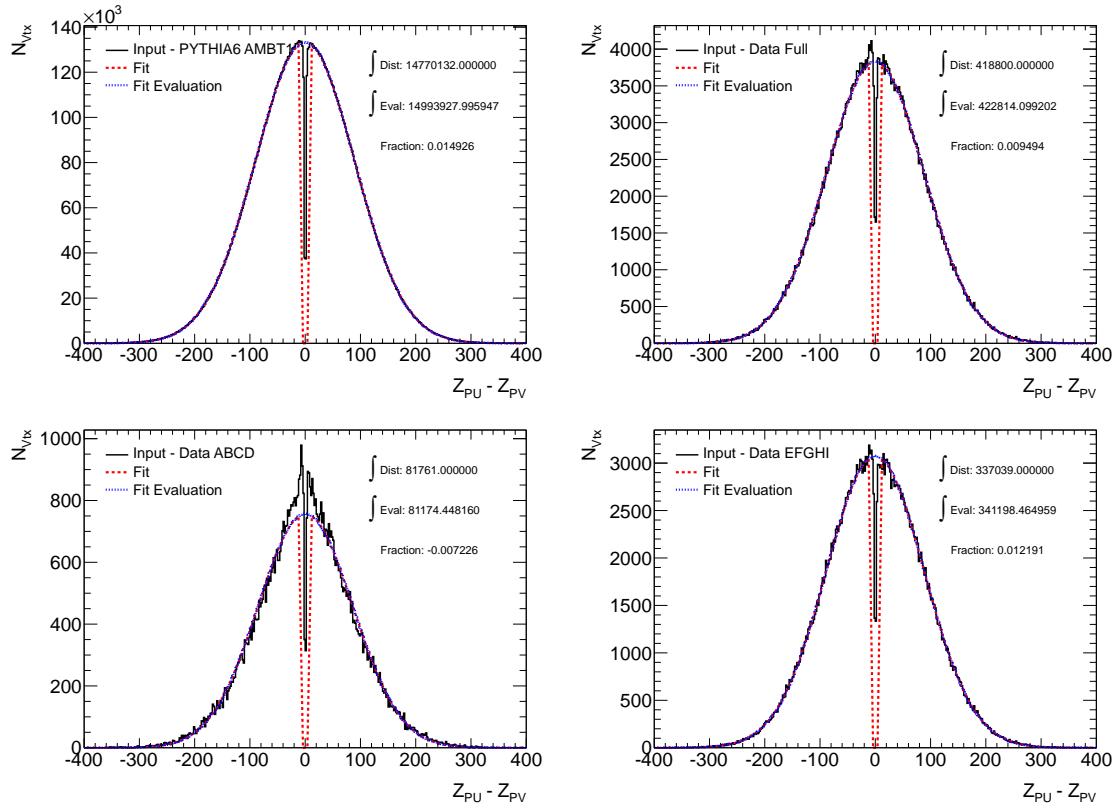


Figure 8.16.: Figures showing the contribution of merged pile-up vertices in an AMBT1 Monte Carlo sample and in data [129].

8.4.4. Jet energy scale uncertainty

The jet energy calibration described in Section 6.3 has an associated uncertainty [131]. The size of this uncertainty was primarily established by the measurement of the single hadron response using test beam data, but was also verified during 2010 data-taking by exploiting momentum conservation in di-jet events. In the central region ($|\eta| < 0.8$), the uncertainty is lower than 4.6% for all jets with $p_T > 20$ GeV, and this decreases to less than 2.5% uncertainty for jet transverse momenta between 60 and 800 GeV. The impact of pile-up on the jet energy is accounted for by the energy scale corrections and uncertainty, and becomes negligible above $p_T > 250$ GeV. The uncertainty is taken as a systematic error.

8.4.5. Jet energy resolution uncertainty

The jet energy resolution is calculated using both data and Monte Carlo [132]. In both cases the resolution was estimated by exploiting momentum conservation in di-jet systems, and then the results were compared between data and MC. The two resolutions agree within 14%, and the difference is taken as a systematic error, varying with jet p_T and rapidity.

8.4.6. Jet reconstruction efficiency uncertainty

The uncertainty on the jet reconstruction efficiency (within the tracking acceptance) is evaluated using track jets. Efficiency is given by the likelihood of constructing a calorimeter jet if there is a nearby track jet. The disagreement in efficiency between data and MC is found to be 2% for calorimeter jets with $p_T > 20$ GeV and less than 1% for those with $p_T > 30$ GeV. These values are taken as a conservative estimate of this uncertainty.

8.4.7. Track reconstruction efficiency uncertainty

Systematic uncertainties due to the tracking efficiency were studied [104], and the largest were found to be due to the material in the Inner Detector and the χ^2 probability cut to remove incorrectly reconstructed tracks. Material budget uncertainty in the Inner Detector was determined to affect the efficiency by around 2% in the barrel region, rising to over 7% for $2.3 < |\eta| < 2.5$, for tracks with $p_T > 500$ MeV. The maximum difference between the fraction of events in data and MC which passed the χ^2 probability cut was found to be 10%. This value was taken as a conservative estimate of the systematic uncertainty, applied to tracks with $p_T > 10$ GeV.

8.4.8. Cluster energy uncertainty

The correction to the energy of simulated topoclusters, described in Section 6.7, has an associated uncertainty. This systematic error comes from uncertainties on the fits to the $M_{\gamma\gamma}$ distributions in $\pi^0 \rightarrow \gamma\gamma$ candidates. The total uncertainty depends on the $|\eta|$ region and is typically 2 – 4%, but increases up to 15% in the regions where different calorimeter subsystems overlap [119].

8.4.9. Unfolding model-dependence

The uncertainty due to model-dependence of the unfolding procedure is taken to be the difference between the results of unfolding with each of the two MC samples, PYTHIA 6 AMBT1 and Herwig++. In contrast to the default procedure described in Section 7.3.1, the distribution unfolded with AMBT1 is used as the final corrected result, and the difference from the Herwig++ unfolding is used to define a symmetrical error band. For each sample, two different priors are used: the unmodified truth distribution, and the truth after reweighting so that the reconstructed distribution matches the data.

8.4.10. Results of error propagation

Tables 8.1 and 8.2 summarise the results of propagating the systematic error pseudo-experiments through the unfolding process for the inclusive leading jet and exclusive di-jet selections respectively. These values are for the plots of how the underlying event observables vary with p_T^{lead} . Plots are also made of the variation of $\langle p_T \rangle$ with N_{ch} , and error propagation gives different results in this case, as shown in Table 8.3. The full pseudo-experiment results are shown in Appendix C.

Contribution of pile-up and merged vertices			
All observables	– 1% –		
Charged tracks	Unfolding	Efficiency	
$\sum p_T$	2 – 3%	3 – 7%	
N_{ch}	1 – 2%	2 – 5%	
$\langle p_T \rangle$	1 – 2%	1%	
Topoclusters	Unfolding	Simulation	Reco
$\sum E_T, \eta < 4.8$	5 – 6%	2%	8 – 12%
$\sum E_T, \eta < 2.5$	8 – 9%	2 – 3%	8 – 12%
Jets	Energy resolution	JES uncertainty	Efficiency
p_T^{lead}	0.2 – 1%	1 – 3%	0.2 – 1.2%

Table 8.1.: Table of systematic uncertainties for inclusive leading jet profiles vs. p_T^{lead} .

Contribution of pile-up and merged vertices			
All observables	– 1% –		
Charged tracks	Unfolding	Efficiency	
$\sum p_T$	1 – 2%	2 – 5%	
N_{ch}	1 – 2%	2 – 4%	
$\langle p_T \rangle$	1 – 2%	0.2%	
Topoclusters	Unfolding	Simulation	Reco
$\sum E_T, \eta < 4.8$	5 – 9%	2 – 8%	7 – 12%
$\sum E_T, \eta < 2.5$	8 – 12%	3 – 6%	8 – 12%
Jets	Energy resolution	JES uncertainty	Efficiency
p_T^{lead}	0.2 – 1%	0.2 – 2%	0.2 – 1%

Table 8.2.: Table of systematic uncertainties for exclusive di-jet profiles vs. p_T^{lead} .

	Unfolding	Reconstruction efficiency	Pile-up
Leading jet	2%	1%	1 – 2%
Exclusive di-jet	5 – 7%	1 – 4%	4%

Table 8.3.: Table of systematic uncertainties for $\langle p_T \rangle$ vs N_{ch} profiles

8.5. Comparison with other UE studies

As described in Section 6.1, the event axis for the topological UE selection can be defined relative to any identified object. Comparing the leading jet results in this chapter with those from other analyses allows tests for consistency, and comparison of different mechanisms. Since jet production is a dominant process at the LHC, and since the highest p_T track is likely to be found in the highest p_T jet, the ATLAS leading track analysis [106] provides results that should be consistent with this leading jet analysis. The leading track results are based on minimum-bias data selections and so cover a lower p_T^{lead} range than leading jet, but as shown in Figure 8.17 the two results are consistent at the p_T^{lead} boundary.

Figure 8.17 also shows results from the ATLAS analysis of the underlying event in leptonic Z-boson decays [112]. The leptons produced have no QCD colour charge, and so are unlikely to produce parton showers through FSR. This may explain the

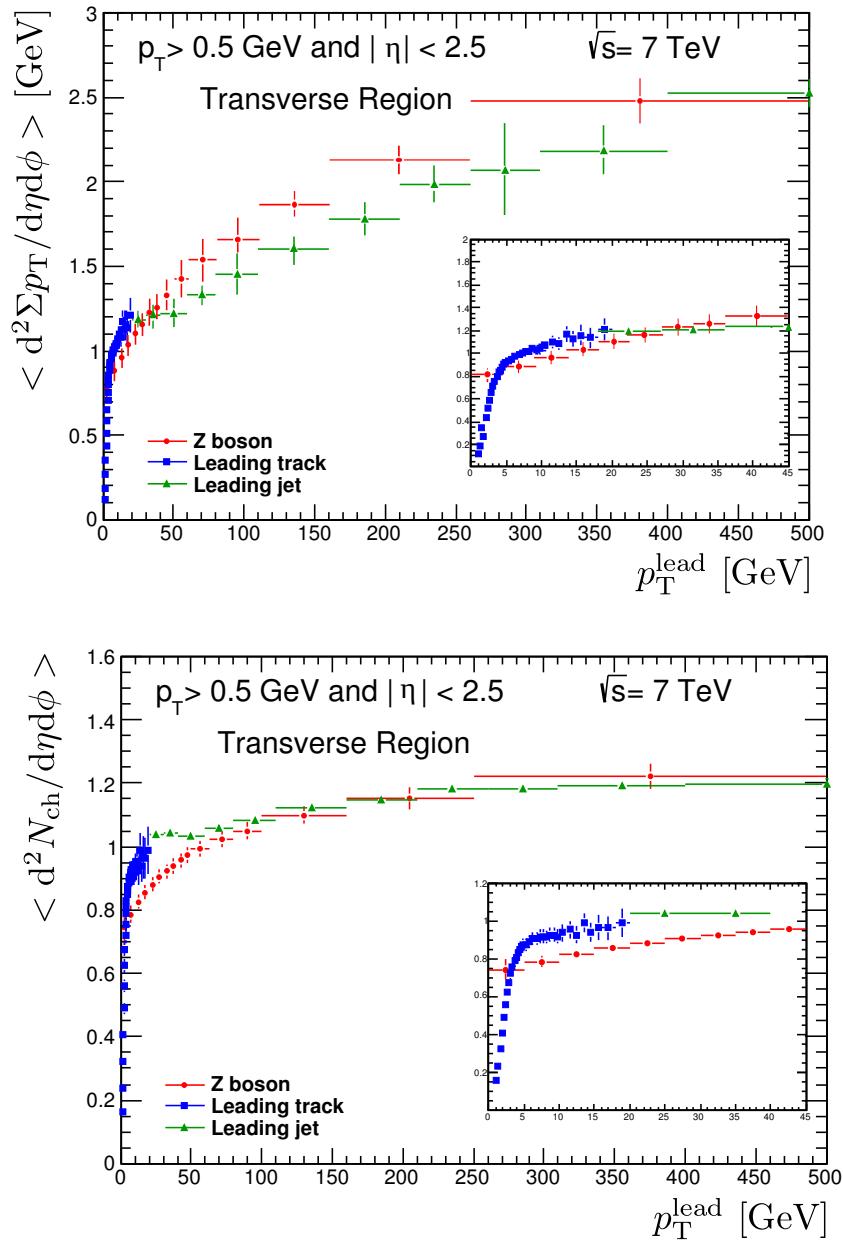


Figure 8.17.: Charged particle $\sum p_T$ and N_{ch} vs the p_T of the identified object, for three different UE analyses [133]. The results from this chapter are compared with the ATLAS leading track [106] and leptonic Z-boson [112] underlying event analyses.

relatively small amount of UE activity in Z-boson events with low p_T^Z . Performing the same comparison at the Tevatron for $p_T^{\text{lead}} < 100 \text{ GeV}$ shows similar behaviour [102].

At higher p_T^Z , the $\sum p_T$ profiles show less activity in jet events than in Z decays. This effect is not expected, but may well just be a difference in the topologies between the two analyses. In the Z-boson underlying event analysis it is quite possible for there to be a jet with greater p_T than the Z itself, without affecting the event orientation. Naturally, in the leading jet analysis this object would define the event orientation, and so there is a potential conflict. Adding a cut to remove jets with greater p_T than the Z-boson gives behaviour similar to the leading jet analysis, as shown in Figure 8.18. Therefore, to the extent that these analyses can be compared they seem to give compatible results, despite the differences in event selection.

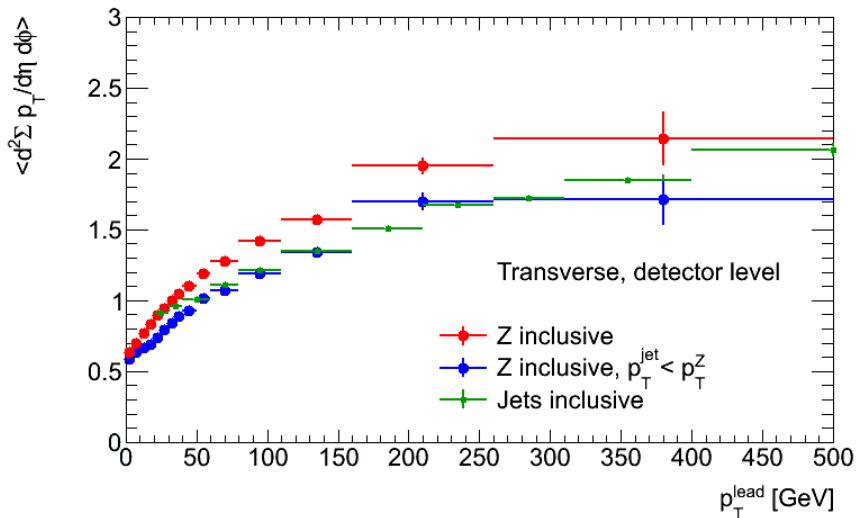


Figure 8.18.: Charged particle $\sum p_T$ vs the p_T of the identified object, for the leading jet and Z-boson UE analyses [134]. Two versions of the Z-boson analysis are presented, one unmodified, one with an additional selection which rejects any events containing a jet with higher p_T than the Z-boson. The additional selection precludes the use of the analysis correction procedure, so these results are presented at detector level.

Chapter 9.

Conclusion

The underlying event has been measured in jet events at the LHC for the first time, using 37 pb^{-1} of pp collisions at $\sqrt{s} = 7 \text{ TeV}$, recorded during 2010 by the ATLAS detector. Observables N_{ch} , $\sum p_{\text{T}}$, $\langle p_{\text{T}} \rangle$, and $\sum E_{\text{T}}$ have been calculated for events with leading jet p_{T} in the range $20 - 800 \text{ GeV}$. In the inclusive leading jet selection the profiles of these observables rise with $p_{\text{T}}^{\text{lead}}$, although in the trans-min region this rise is small. Using an exclusive di-jet selection gives rise to behaviour that has not been observed before: the observables fall with $p_{\text{T}}^{\text{lead}}$.

These different behaviours highlight the ambiguity over what exactly constitutes the underlying event. The long tails in the $\sum p_{\text{T}}$ distributions for the inclusive selection suggest that jets are contributing to this measurement. However, these are not necessarily from the hard process: additional parton interactions may give rise to jets too. Whether jets arising from MPI should be considered part of the underlying event or not, modified selections could be tested in future to study particular contributions in detail:

- Complete jet subtraction: all particles within the ΔR cone of an identified jet could be removed from the selection, with a corresponding reduction in the $d\eta d\phi$ normalisation for that event. However, this may miss particles arising from jets that are irregularly-shaped, and the normalisation would be complicated by the potential for overlapping jet cones.
- Balanced di-jets: in order to allow some jet contributions from MPI, the exclusive di-jet selection could be relaxed to allow an additional balanced di-jet. Since a hard scattering process producing four jets need only conserve momentum as an ensemble, hypothetically an event containing two balanced di-jets is more likely to have arisen from two separate partonic interactions.

Bayesian iterative unfolding was used to correct the underlying event profiles for the effects of the ATLAS detector. Although the correction relies on a prior probability distribution, the unfolding technique has reduced model-dependence to an uncertainty comparable to or smaller than experimental systematic errors.

Unfolding the data has allowed comparisons to be made to several different Monte Carlo models, and to results from other experiments. The leading jet underlying event results are consistent with those from the leading track analysis — they give overlapping values at their shared p_T^{lead} limit. Comparing with measurements from Z-boson production shows two discrepancies, although both can be explained:

- Relatively low activity in both the charged particle multiplicity and $\sum p_T$ profiles at low p_T^Z . This is probably due to the lack of QCD final state radiation from the leptonic Z-boson decays.
- A more rapidly rising $\sum p_T$ profile as p_T^Z increases. This appears to be the result of contributions from particularly high- p_T jets in the Z analysis. In the leading jet analysis the highest p_T jet would define the orientation of an event, whereas in the Z analysis jets do not affect the orientation. Removing events containing a jet with higher p_T than the Z-boson addresses this effect, giving compatible results for both analyses.

The three ATLAS analyses compared do seem to be compatible with these considerations in mind, although opportunities for further study suggest themselves. The lack of FSR in leptonic Z-boson decays could either be addressed by attempting an FSR correction in the leading jet analysis, or by studying the underlying event in *hadronic* Z-boson decays (where the decay products have colour, and so FSR can occur through QCD processes). It may also be possible to extend the p_T^{lead} range of the leading track analysis upwards, giving some overlap with the leading jet result.

When comparing the leading jet results to MC predictions, HERWIG+JIMMY AUET2 performs well in the inclusive selection for N_{ch} , $\sum p_T$ and $\sum E_T$ vs p_T^{lead} , and PYTHIA 6 AUET2B does better in the exclusive selection. In both the inclusive and exclusive selections, PYTHIA 6 DW gives the best reproduction of $\langle p_T \rangle$ vs p_T^{lead} . The inclusive profile for $\langle p_T \rangle$ vs N_{ch} also favours HERWIG+JIMMY AUET2, but in the exclusive case the uncertainties are too large to distinguish between models. No model is able to reproduce the $\sum E_T$ result over the full pseudorapidity range $|\eta| < 4.8$, reflecting the fact that tuning has not been attempted for results from the forward region. These results give the opportunity to improve the performance of Monte Carlo generators by tuning them for

the new energy and pseudorapidity ranges investigated. With that in mind, the results of the leading jet analysis will be uploaded into the HepData [135] repository for easy use in future tuning efforts.

Appendix A.

Unfolding closure and iteration stability

In this section, MC reconstructed distributions are unfolded rather than data. From one to five iterations of the unfolding are shown, shading from black to red as the iteration number increases. These unfolded distributions are compared with the true distribution shown by the black points with green error bars. The difference between the unfolded reconstructed distribution and the truth distribution gives an indication of the quality of the unfolding. As expected, unfolding a reconstructed distribution using a prior and smearing matrix made from the same MC sample gives perfect closure. In general the choice of two unfolding iterations is validated where closure is not perfect, as fluctuations in the unfolded distribution are reinforced by additional iterations. Apart from this reinforcing of fluctuations, the unfolding seems independent of iteration number (once past the prior-dependent result of a single unfolding iteration). Figures A.7 and A.8 do seem to show some trend with iteration number, but these profiles have large unfolding uncertainties, and two iterations still gives the best closure.

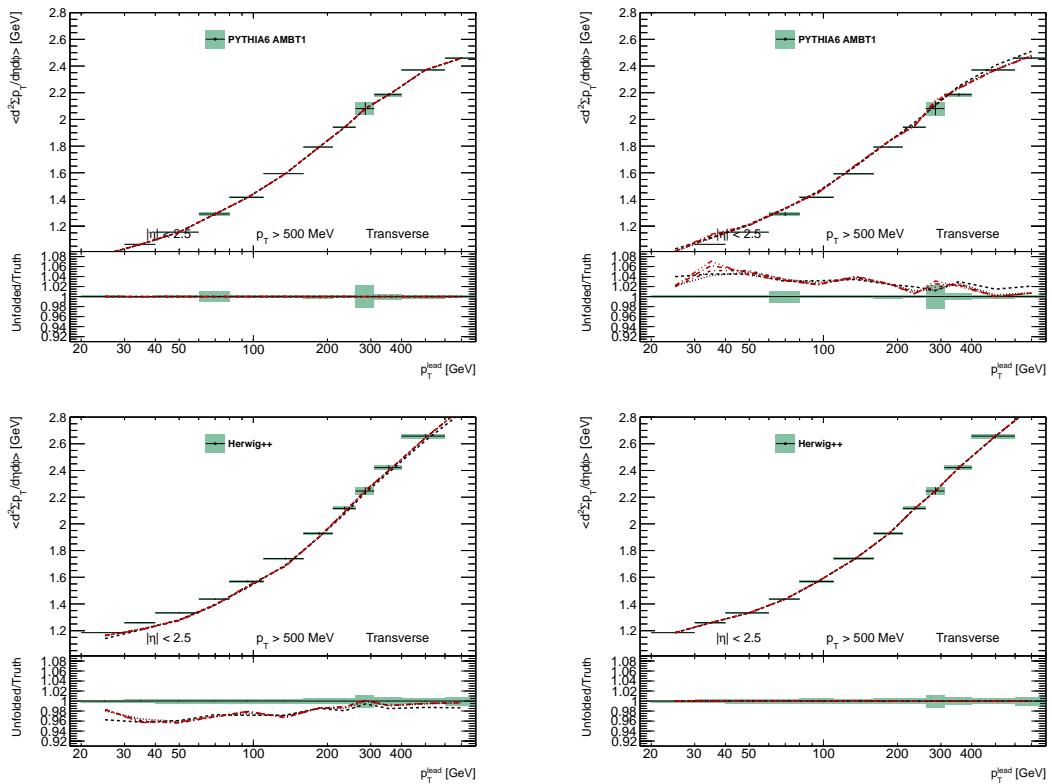


Figure A.1.: Results shown for the transverse region – leading jet selection – charged particle $\sum p_T$ vs. p_T^{lead} . In the top row, the AMBT1 reco distribution is unfolded using prior and smearing matrix from AMBT1 (left) and Herwig++ (right). In the bottom row, the Herwig++ reco distribution is unfolded using prior and smearing matrix from AMBT1 (left) and Herwig++ (right).

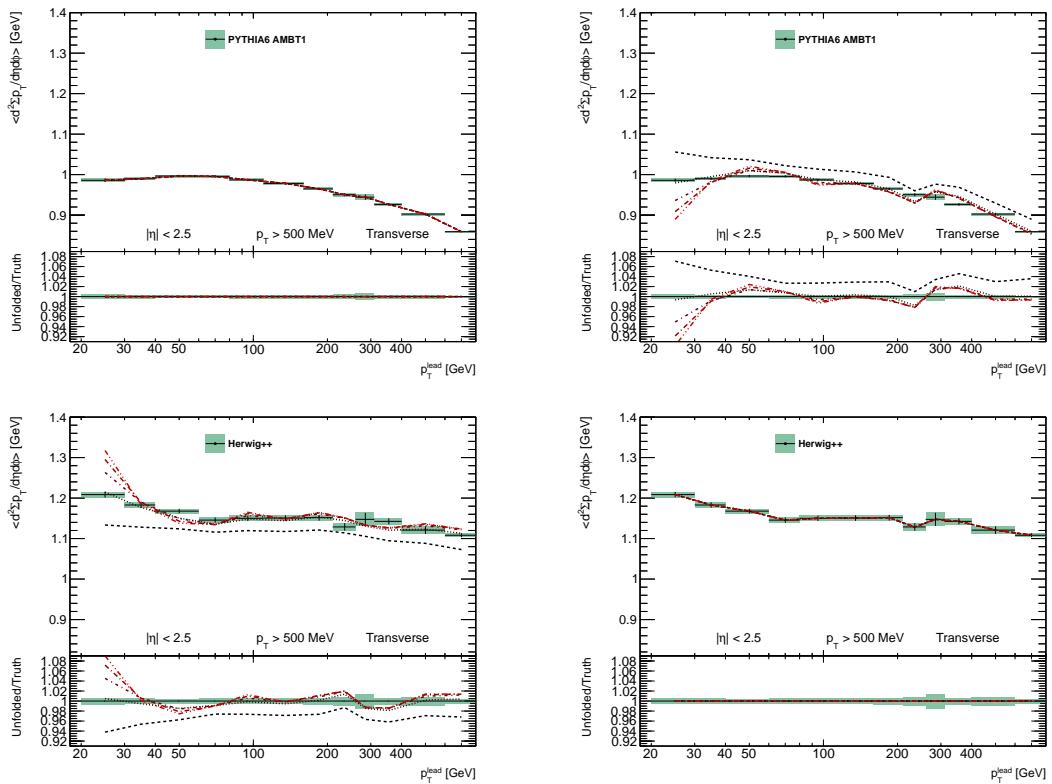


Figure A.2.: Results shown for the transverse region – exclusive di-jet selection – charged particle $\sum p_T$ vs. p_T^{lead} . In the top row, the AMBT1 reco distribution is unfolded using prior and smearing matrix from AMBT1 (left) and Herwig++ (right). In the bottom row, the Herwig++ reco distribution is unfolded using prior and smearing matrix from AMBT1 (left) and Herwig++ (right).

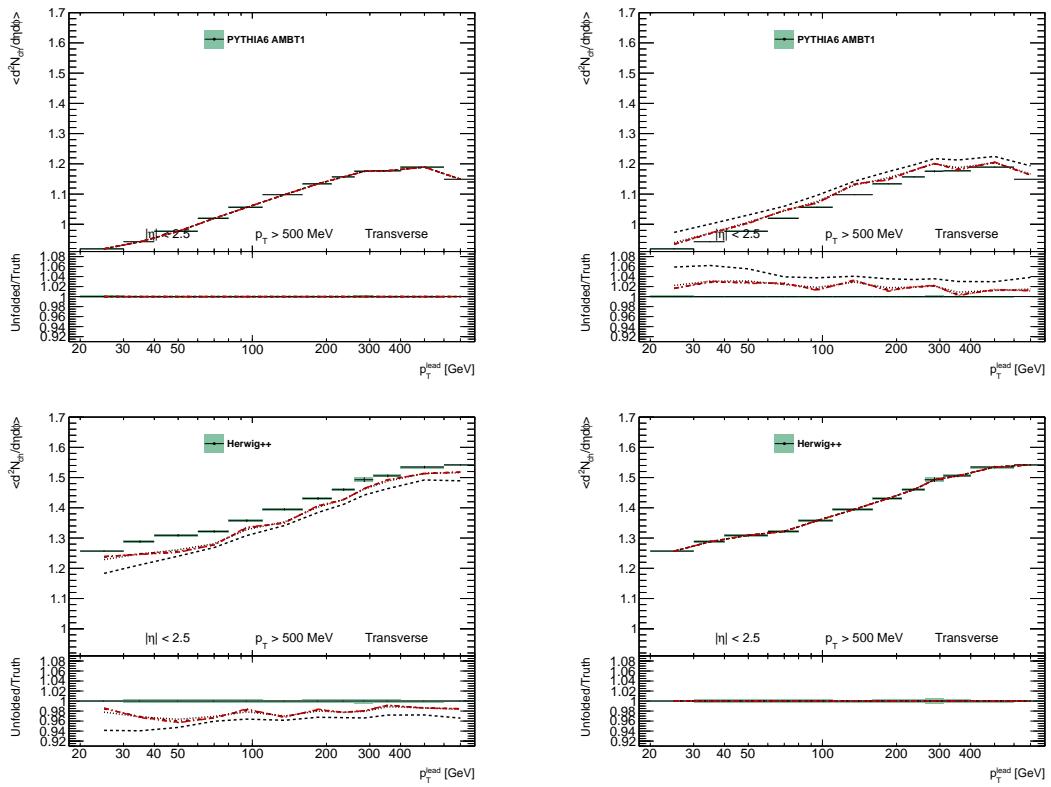


Figure A.3.: Results shown for the transverse region – leading jet selection – charged particle N_{ch} vs. p_T^{lead} . In the top row, the AMBT1 reco distribution is unfolded using prior and smearing matrix from AMBT1 (left) and Herwig++ (right). In the bottom row, the Herwig++ reco distribution is unfolded using prior and smearing matrix from AMBT1 (left) and Herwig++ (right).

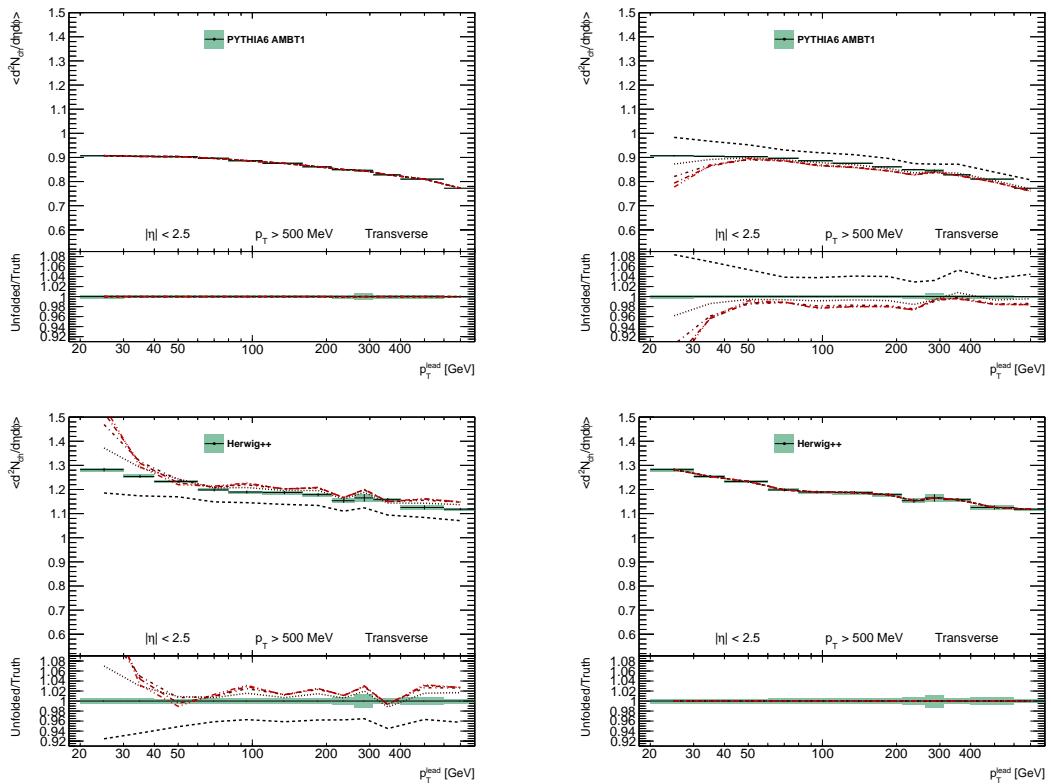


Figure A.4.: Results shown for the transverse region – exclusive di-jet selection – charged particle N_{ch} vs. p_T^{lead} . In the top row, the AMBT1 reco distribution is unfolded using prior and smearing matrix from AMBT1 (left) and Herwig++ (right). In the bottom row, the Herwig++ reco distribution is unfolded using prior and smearing matrix from AMBT1 (left) and Herwig++ (right).

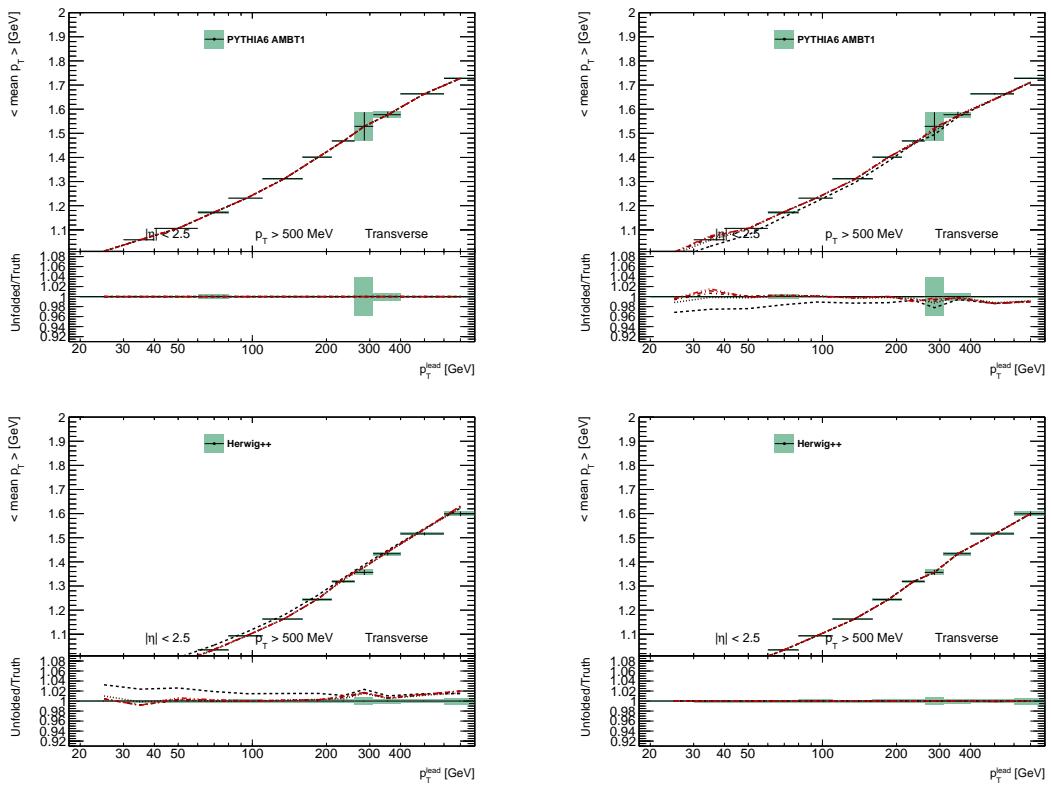


Figure A.5.: Results shown for the transverse region – leading jet selection – charged particle $\langle p_T \rangle$ vs. p_T^{lead} . In the top row, the AMBT1 reco distribution is unfolded using prior and smearing matrix from AMBT1 (left) and Herwig++ (right). In the bottom row, the Herwig++ reco distribution is unfolded using prior and smearing matrix from AMBT1 (left) and Herwig++ (right).

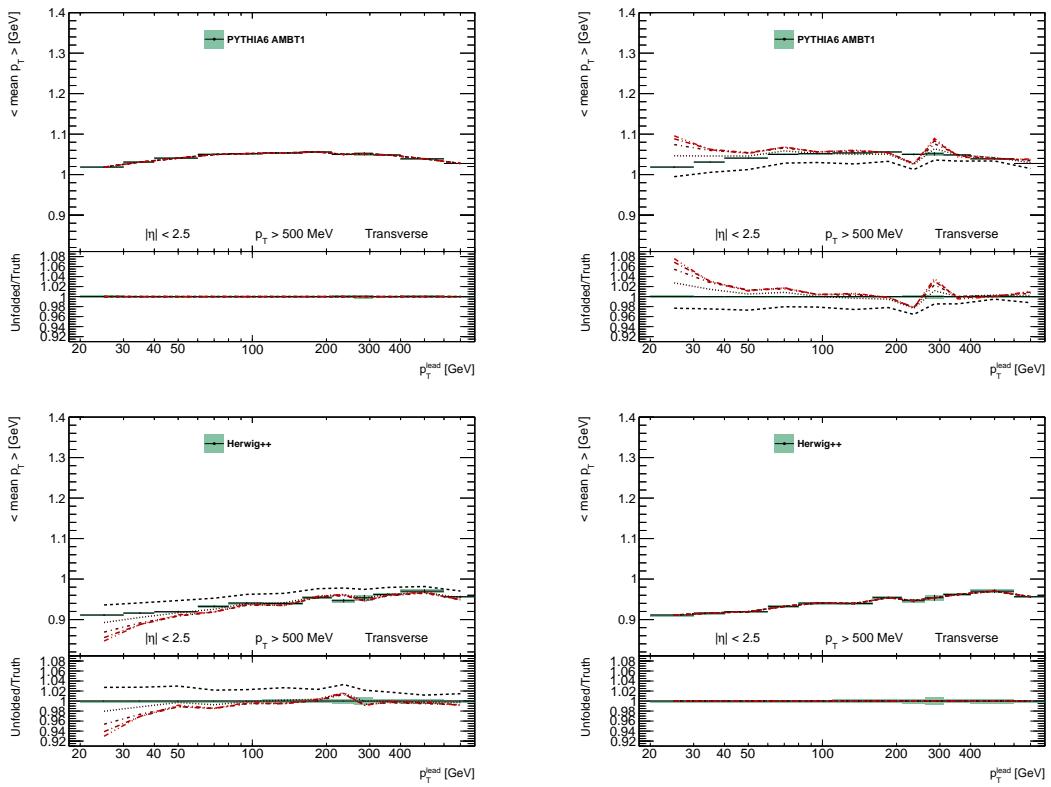


Figure A.6.: Results shown for the transverse region – exclusive di-jet selection – charged particle $\langle p_T \rangle$ vs. p_T^{lead} . In the top row, the AMBT1 reco distribution is unfolded using prior and smearing matrix from AMBT1 (left) and Herwig++ (right). In the bottom row, the Herwig++ reco distribution is unfolded using prior and smearing matrix from AMBT1 (left) and Herwig++ (right).

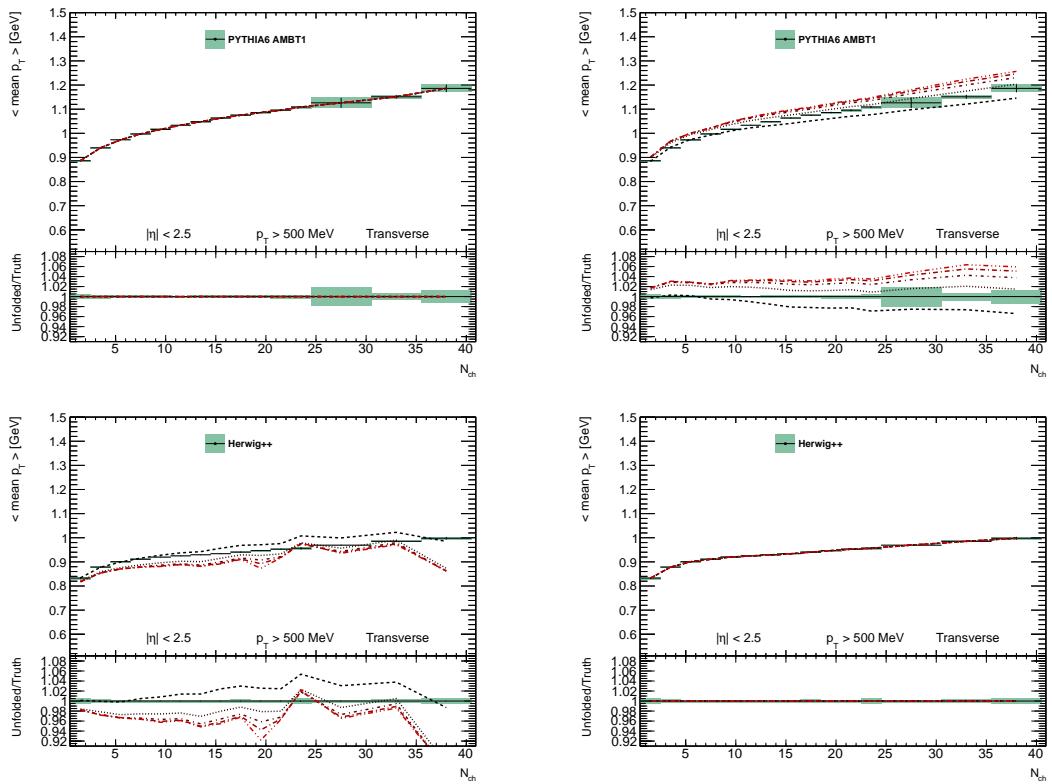


Figure A.7.: Results shown for the transverse region – leading jet selection – charged particle $\langle p_T \rangle$ vs. N_{ch} . In the top row, the AMBT1 reco distribution is unfolded using prior and smearing matrix from AMBT1 (left) and Herwig++ (right). In the bottom row, the Herwig++ reco distribution is unfolded using prior and smearing matrix from AMBT1 (left) and Herwig++ (right).

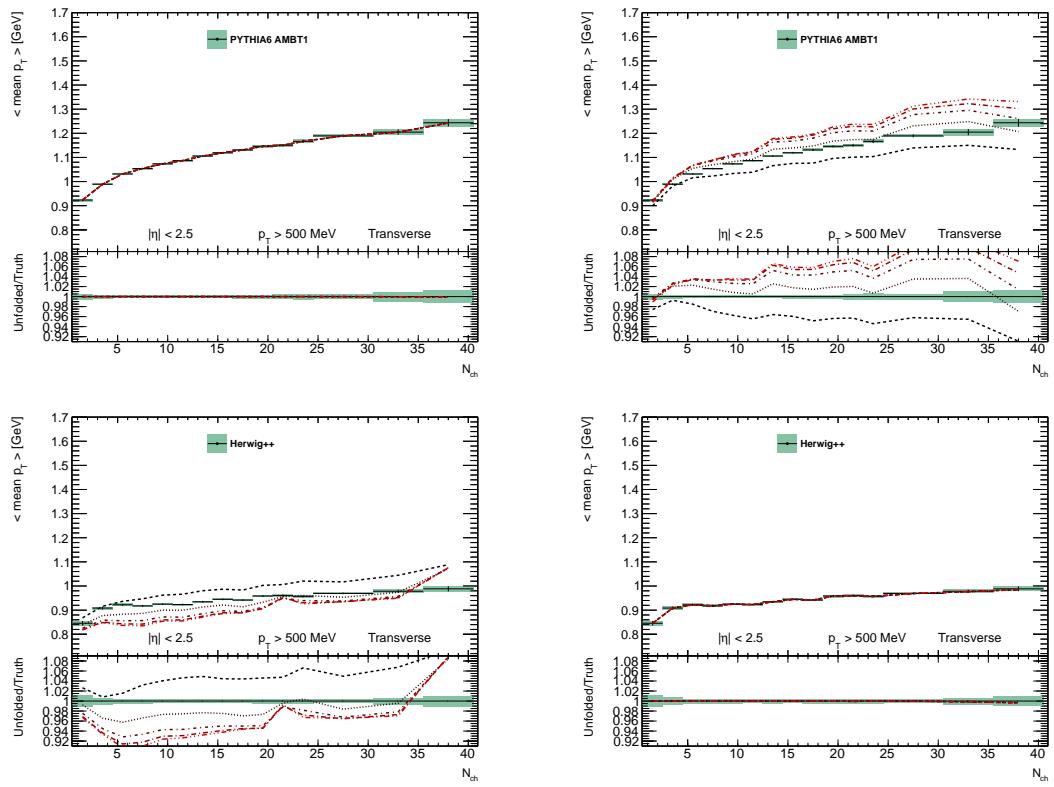


Figure A.8.: Results shown for the transverse region – exclusive di-jet selection – charged particle $\langle p_T \rangle$ vs. N_{ch} . In the top row, the AMBT1 reco distribution is unfolded using prior and smearing matrix from AMBT1 (left) and Herwig++ (right). In the bottom row, the Herwig++ reco distribution is unfolded using prior and smearing matrix from AMBT1 (left) and Herwig++ (right).

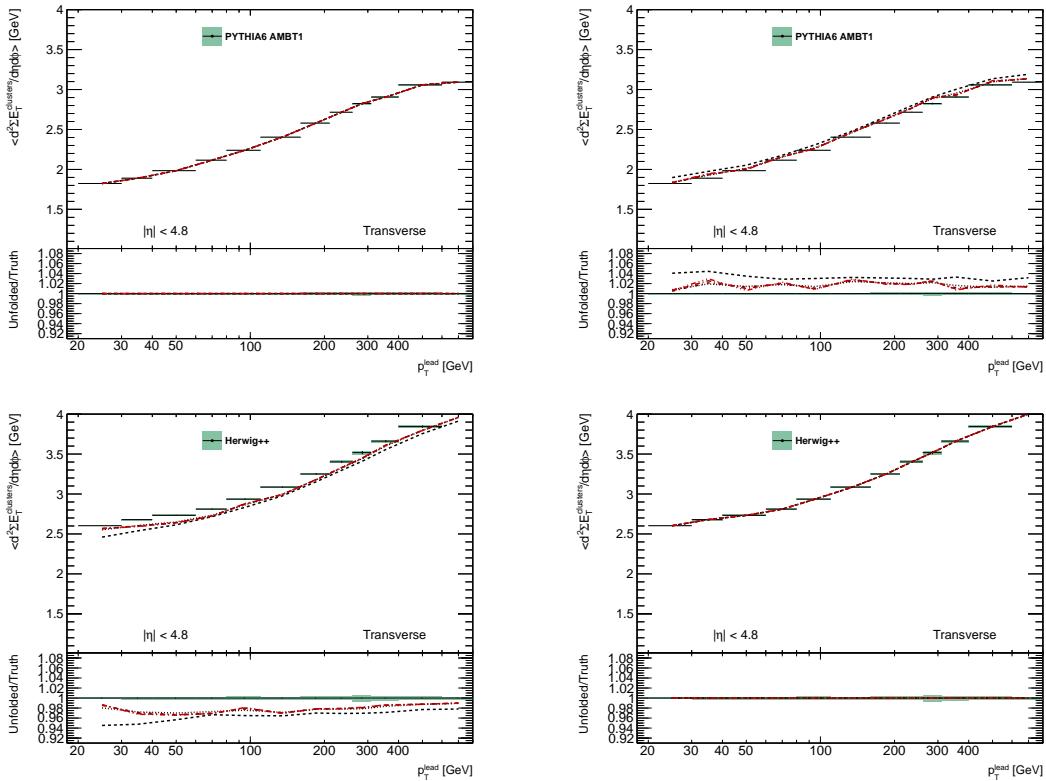


Figure A.9.: Results shown for the transverse region – leading jet selection – topocluster $\sum E_T$ for $|\eta| < 4.8$ vs. p_T^{lead} . In the top row, the AMBT1 reco distribution is unfolded using prior and smearing matrix from AMBT1 (left) and Herwig++ (right). In the bottom row, the Herwig++ reco distribution is unfolded using prior and smearing matrix from AMBT1 (left) and Herwig++ (right).

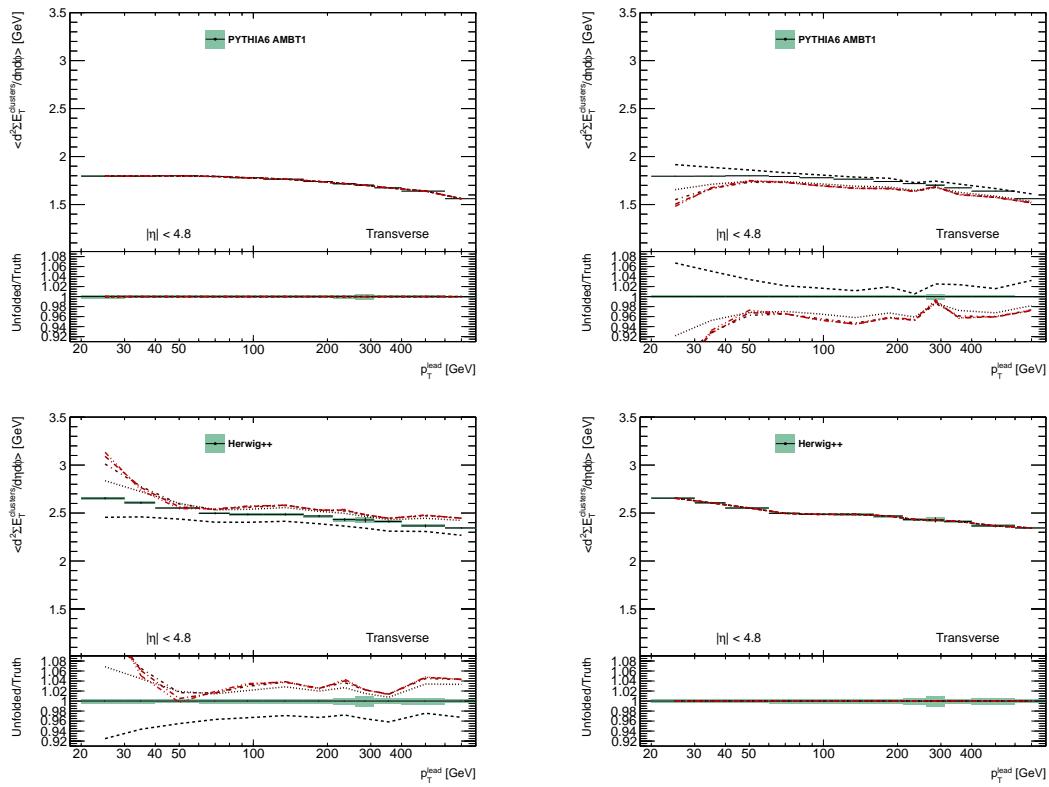


Figure A.10.: Results shown for the transverse region – exclusive di-jet selection – topocluster $\sum E_T$ for $|\eta| < 4.8$ vs. p_T^{lead} . In the top row, the AMBT1 reco distribution is unfolded using prior and smearing matrix from AMBT1 (left) and Herwig++ (right). In the bottom row, the Herwig++ reco distribution is unfolded using prior and smearing matrix from AMBT1 (left) and Herwig++ (right).

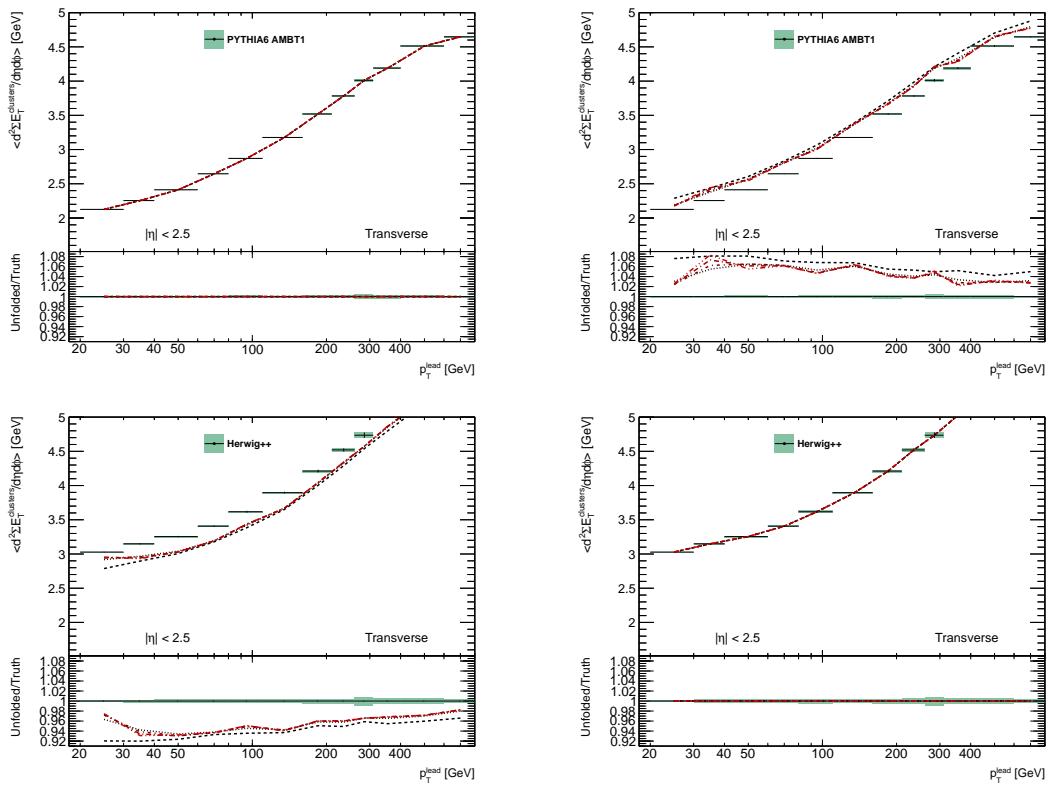


Figure A.11.: Results shown for the transverse region – leading jet selection – cluster $\sum E_T$ for $|\eta| < 2.5$ vs. p_T^{lead} . In the top row, the AMBT1 reco distribution is unfolded using prior and smearing matrix from AMBT1 (left) and Herwig++ (right). In the bottom row, the Herwig++ reco distribution is unfolded using prior and smearing matrix from AMBT1 (left) and Herwig++ (right).

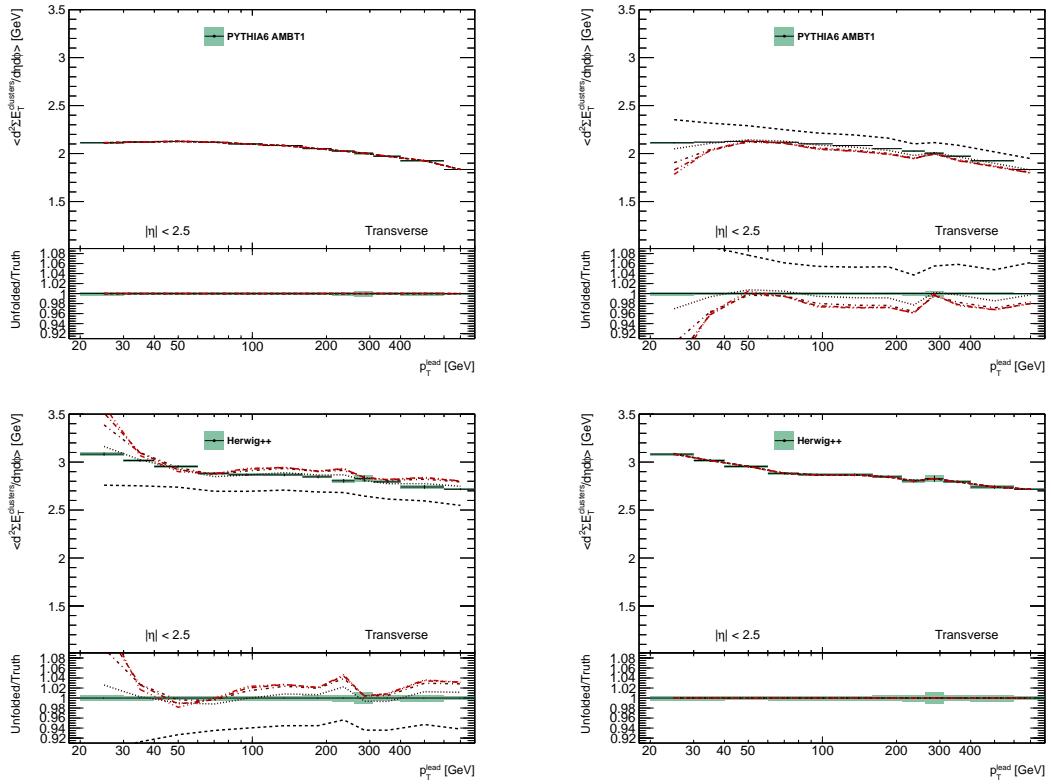


Figure A.12.: Results shown for the transverse region – exclusive di-jet selection – cluster $\sum E_T$ for $|\eta| < 2.5$ vs. p_T^{lead} . In the top row, the AMBT1 reco distribution is unfolded using prior and smearing matrix from AMBT1 (left) and Herwig++ (right). In the bottom row, the Herwig++ reco distribution is unfolded using prior and smearing matrix from AMBT1 (left) and Herwig++ (right).

Appendix B.

Data stability plots

The following figures compare 1D distributions between data periods for the ranges in p_T^{lead} where the periods' occupancies most comparably overlap. Results are shown for the transverse region.

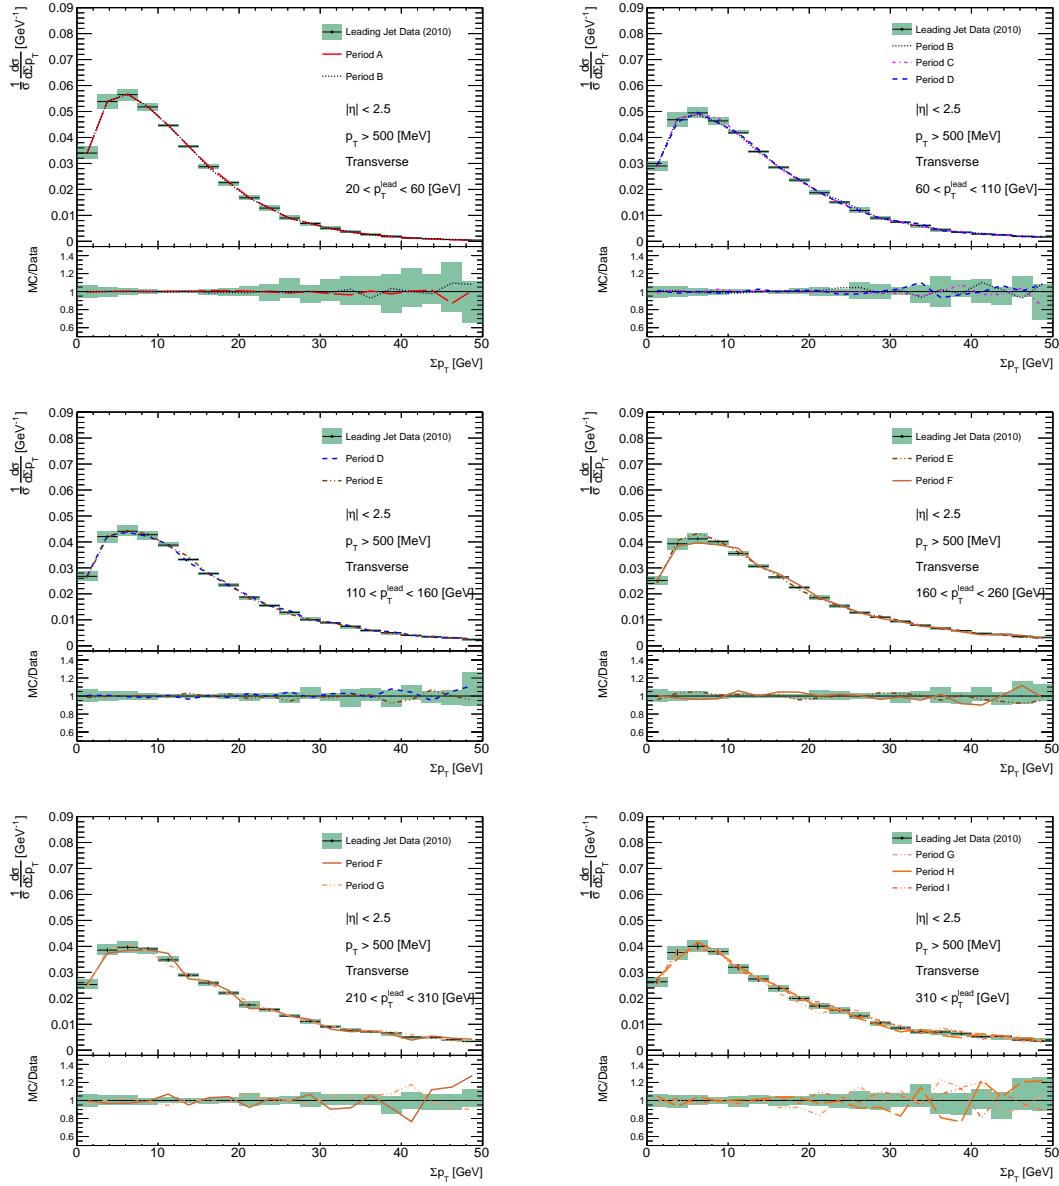


Figure B.1.: Comparison of the charged particle $\sum p_T$ distributions in the transverse region for the data periods and p_T^{lead} ranges indicated.

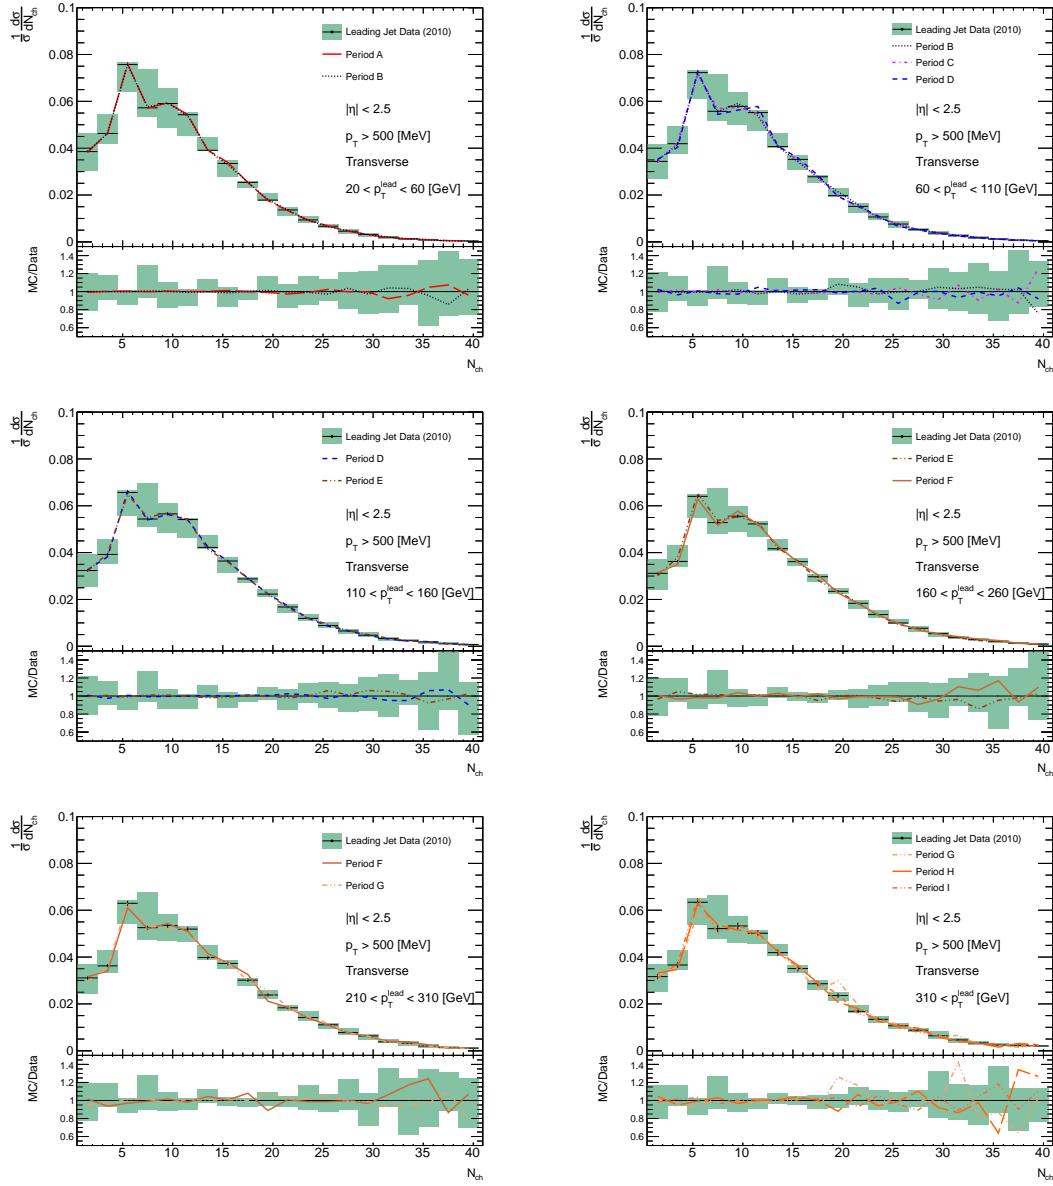


Figure B.2.: Comparison of the charged particle N_{ch} distributions in the transverse region for the data periods and $p_{\text{T}}^{\text{lead}}$ ranges indicated.

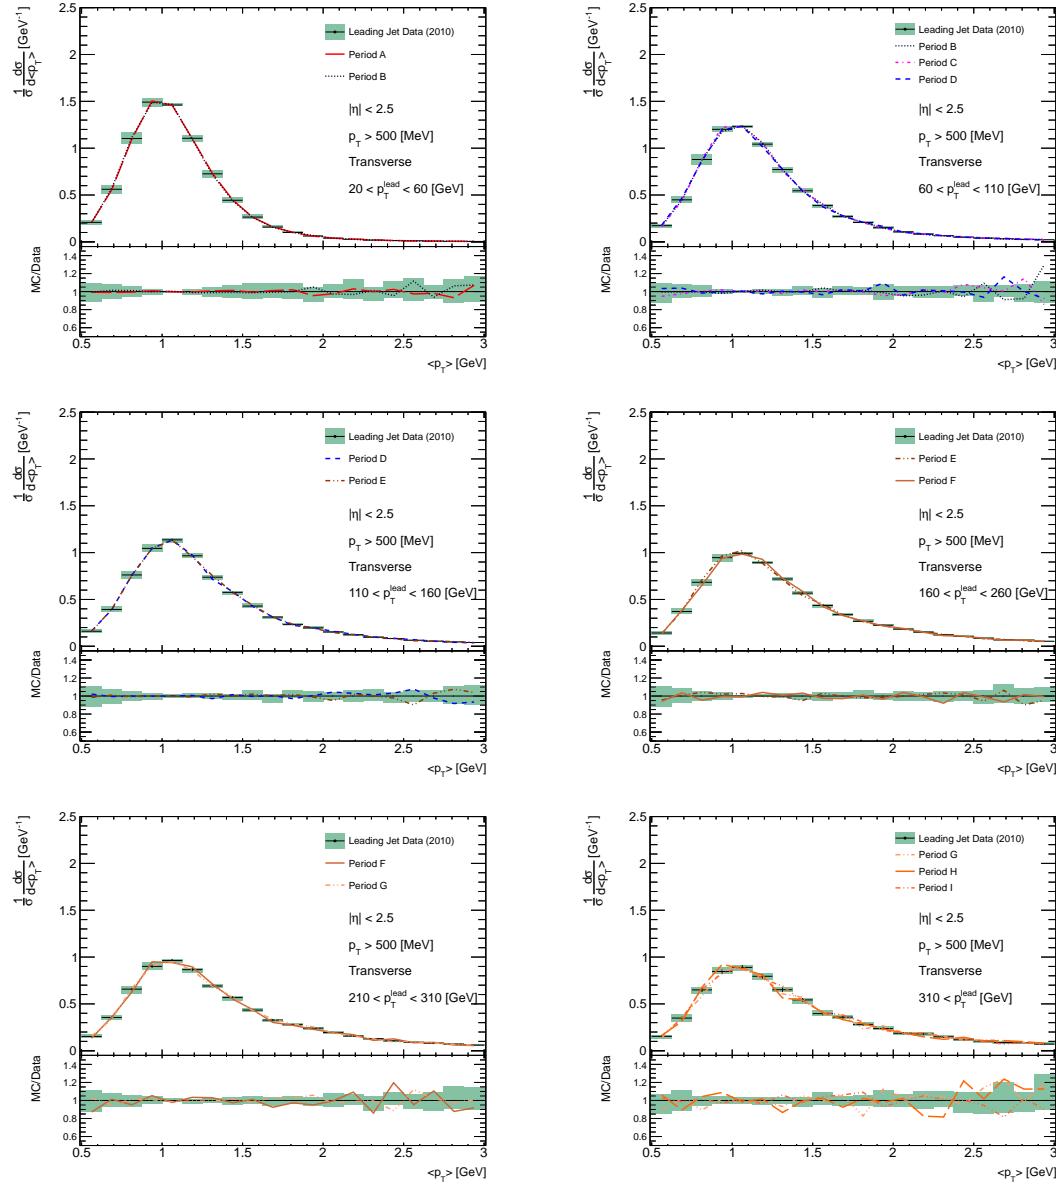


Figure B.3.: Comparison of the charged particle $\langle p_T \rangle$ distributions in the transverse region for the data periods and p_T^{lead} ranges indicated.

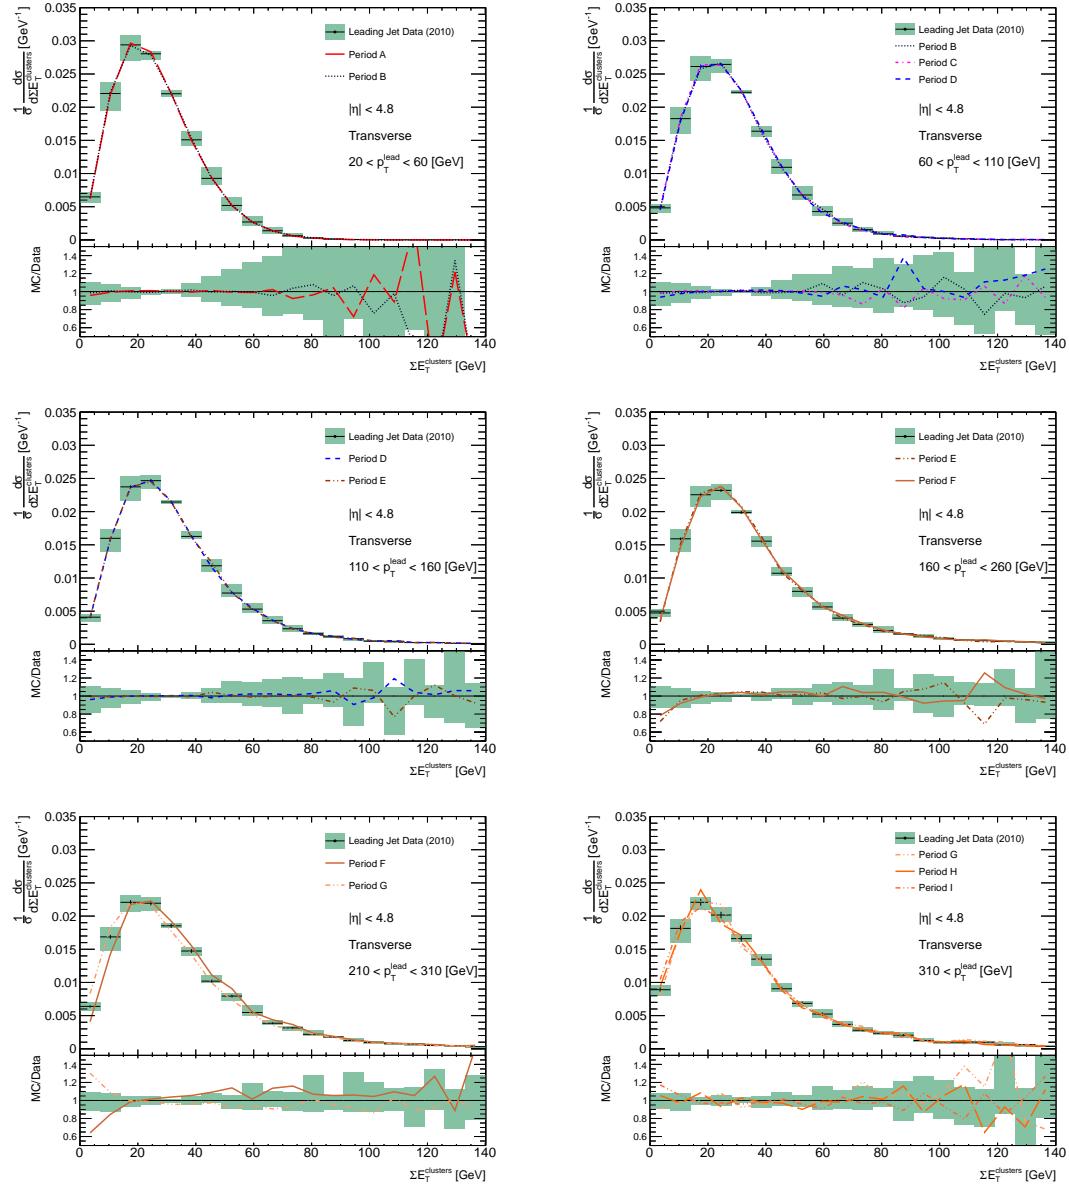


Figure B.4.: Comparison of the topocluster $\sum E_T$ distributions ($|\eta| < 4.8$) in the transverse region for the data periods and p_T^{lead} ranges indicated.

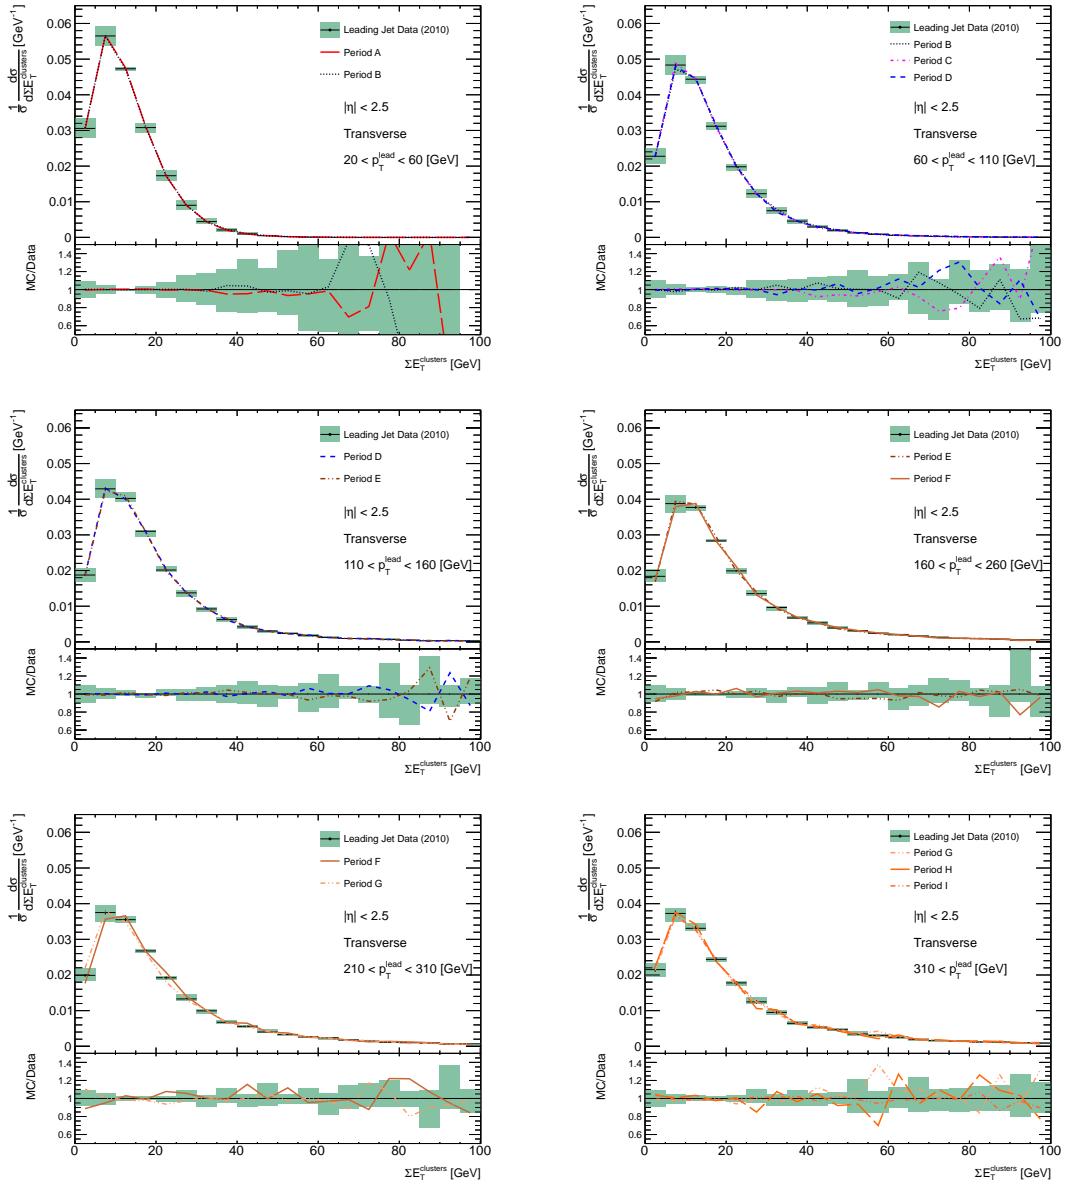


Figure B.5.: Comparison of the topocluster $\sum E_T$ distributions ($|\eta| < 2.5$) in the transverse region for the data periods and p_T^{lead} ranges indicated.

Appendix C.

Propagation of systematic errors

This section shows the results of unfolding each of the different systematic error pseudo-experiments described in Section 8.4. The unfolded, unmodified data distribution is shown in each plot by the black points, with the green error bars indicating statistical errors only. Each pseudo-experiment is labelled with the source of error, and whether it corresponds to $\pm 1\sigma$.

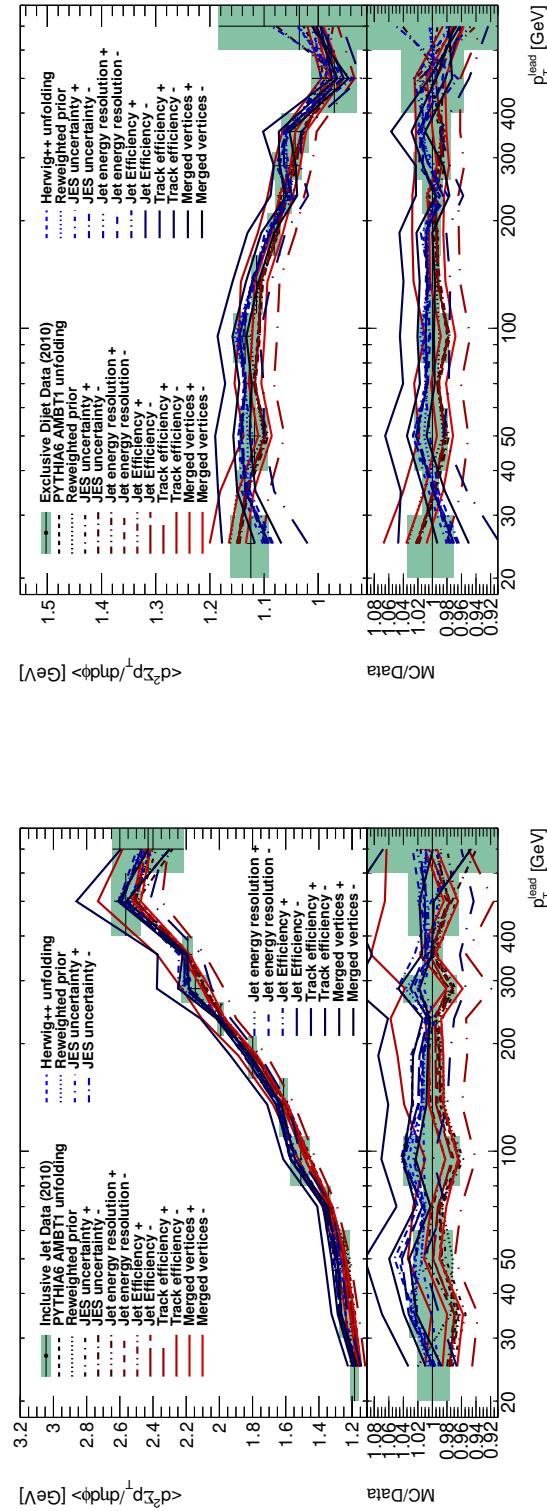


Figure C.1.: Charged particle $\sum p_T$ vs. p_T^{lead} systematic error pseudo-experiment profiles for the transverse region. Inclusive leading jet and exclusive di-jet selections on the left and right respectively.

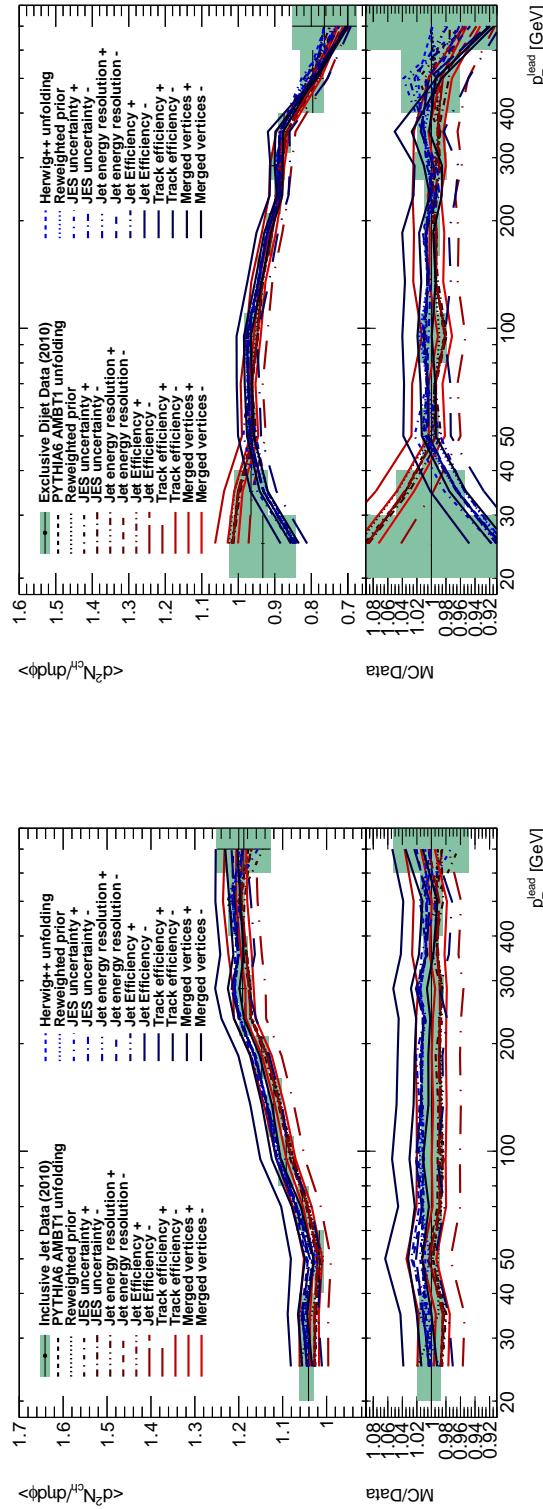


Figure C.2: Charged particle N_{ch} vs. p_T^{lead} systematic error pseudo-experiment profiles for the transverse region. Inclusive leading jet and exclusive di-jet selections on the left and right respectively.

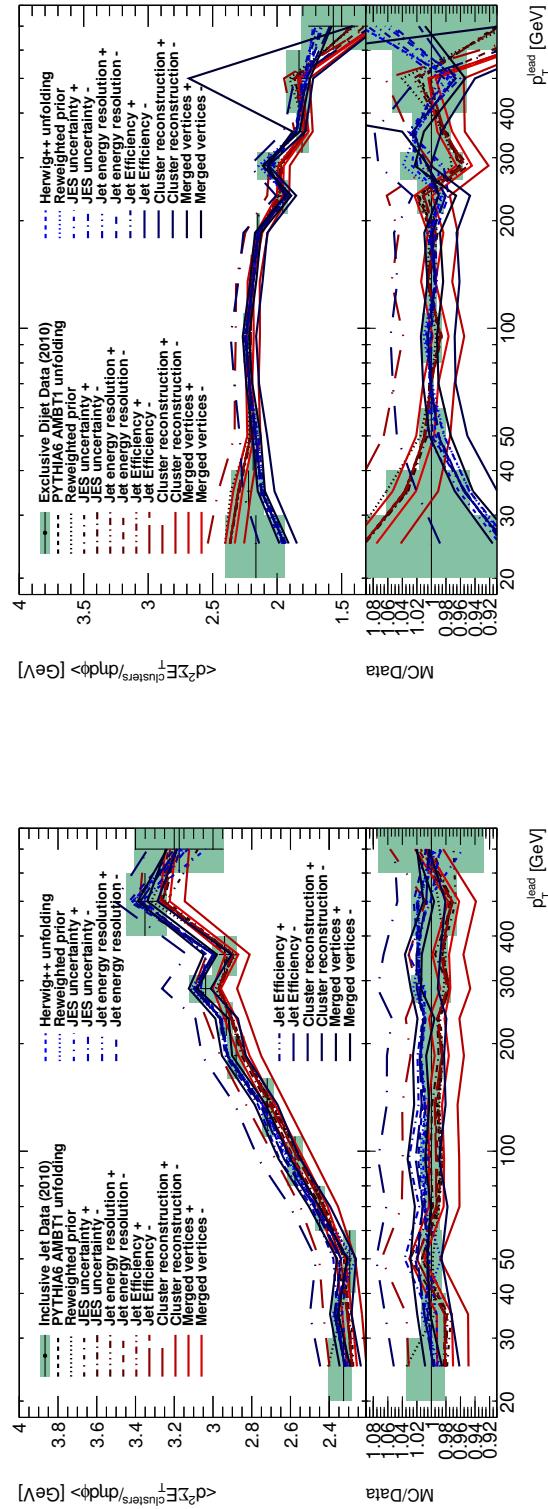


Figure C.3: Topocluster $\sum E_T$ vs. p_T^{lead} systematic error pseudo-experiment profiles ($|\eta| < 4.8$) for the transverse region. Inclusive leading jet and exclusive di-jet selections on the left and right respectively.

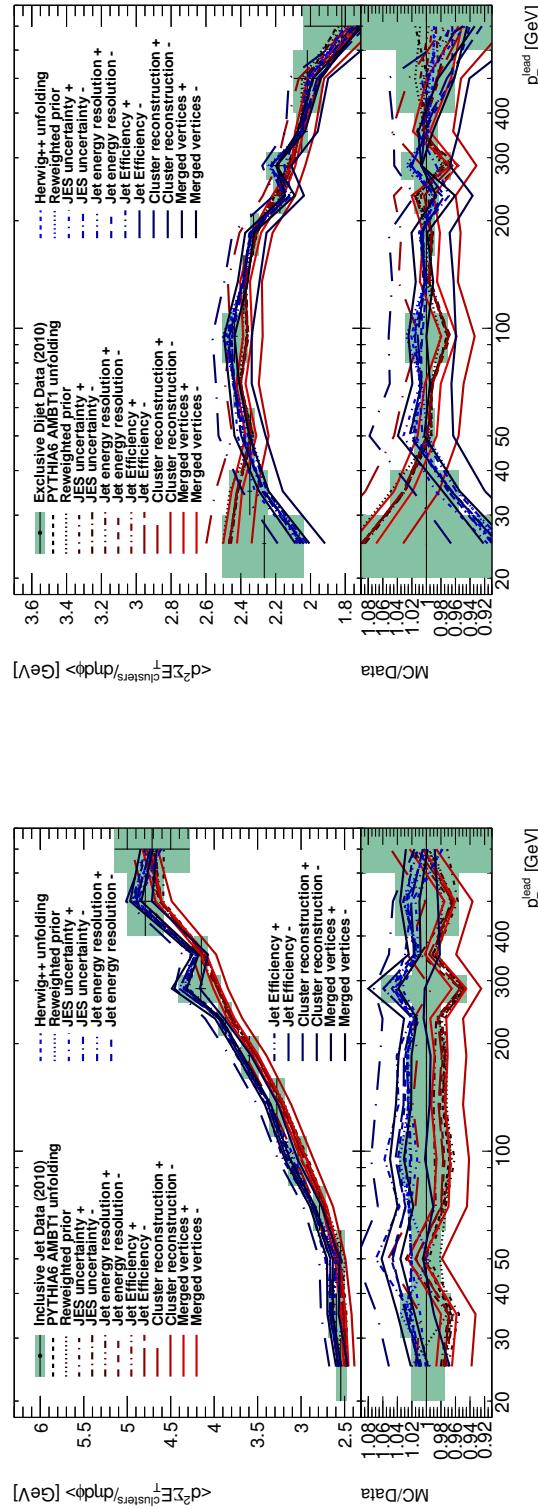


Figure C.4: Topocluster $\sum E_T$ vs. p_T^{lead} systematic error pseudo-experiment profiles ($|\eta| < 2.5$) for the transverse region. Inclusive leading jet and exclusive di-jet selections on the left and right respectively.

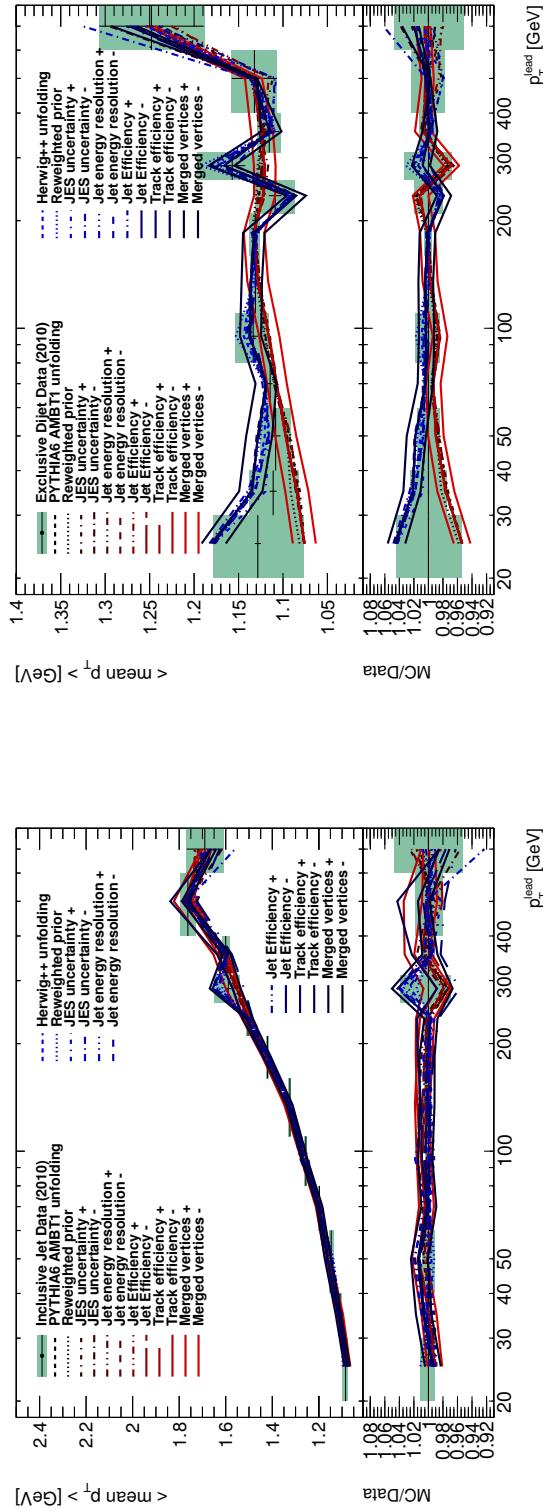


Figure C.5: Charged particle $\langle p_T \rangle$ vs. p_T^{lead} systematic error pseudo-experiment profiles for the transverse region. Inclusive leading jet and exclusive di-jet selections on the left and right respectively.

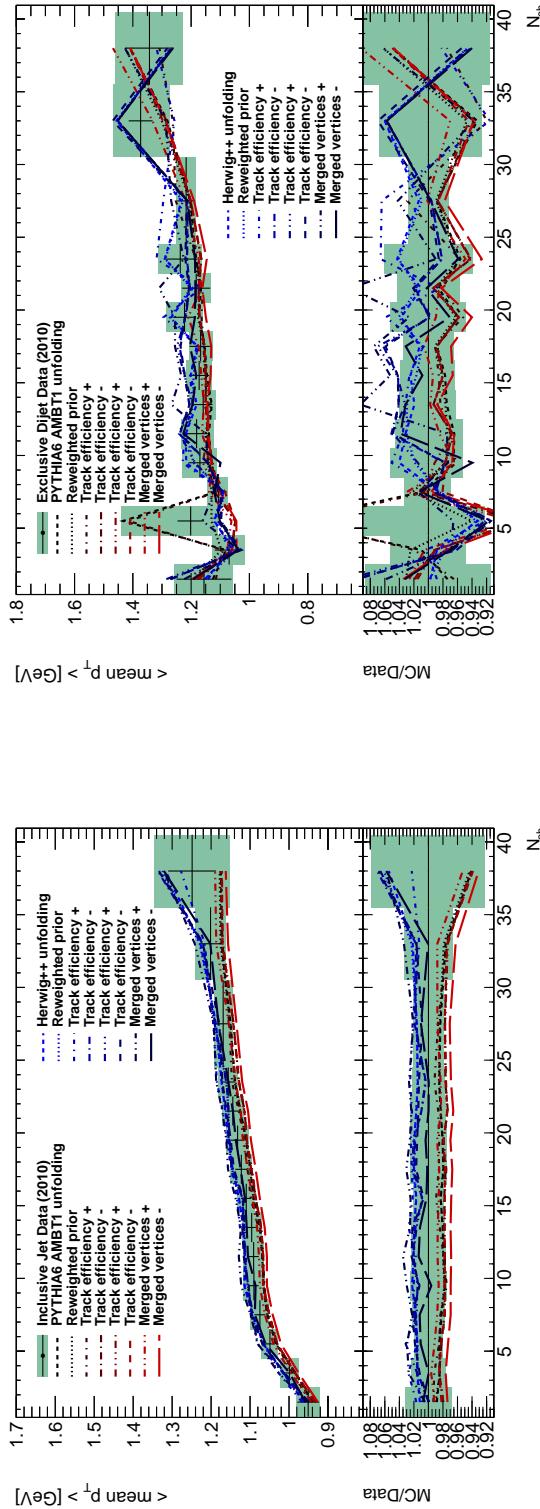


Figure C.6.: Charged particle $\langle p_T \rangle$ vs. N_{ch} systematic error pseudo-experiment profiles for the transverse region. Inclusive leading jet and exclusive di-jet selections on the left and right respectively.

Appendix D.

Underlying event observable 1D plots

D.1. Full p_T^{lead} range

The distributions of underlying event observable values, measured in the transverse, trans-min and trans-max regions and plotted for the full p_T^{lead} range.

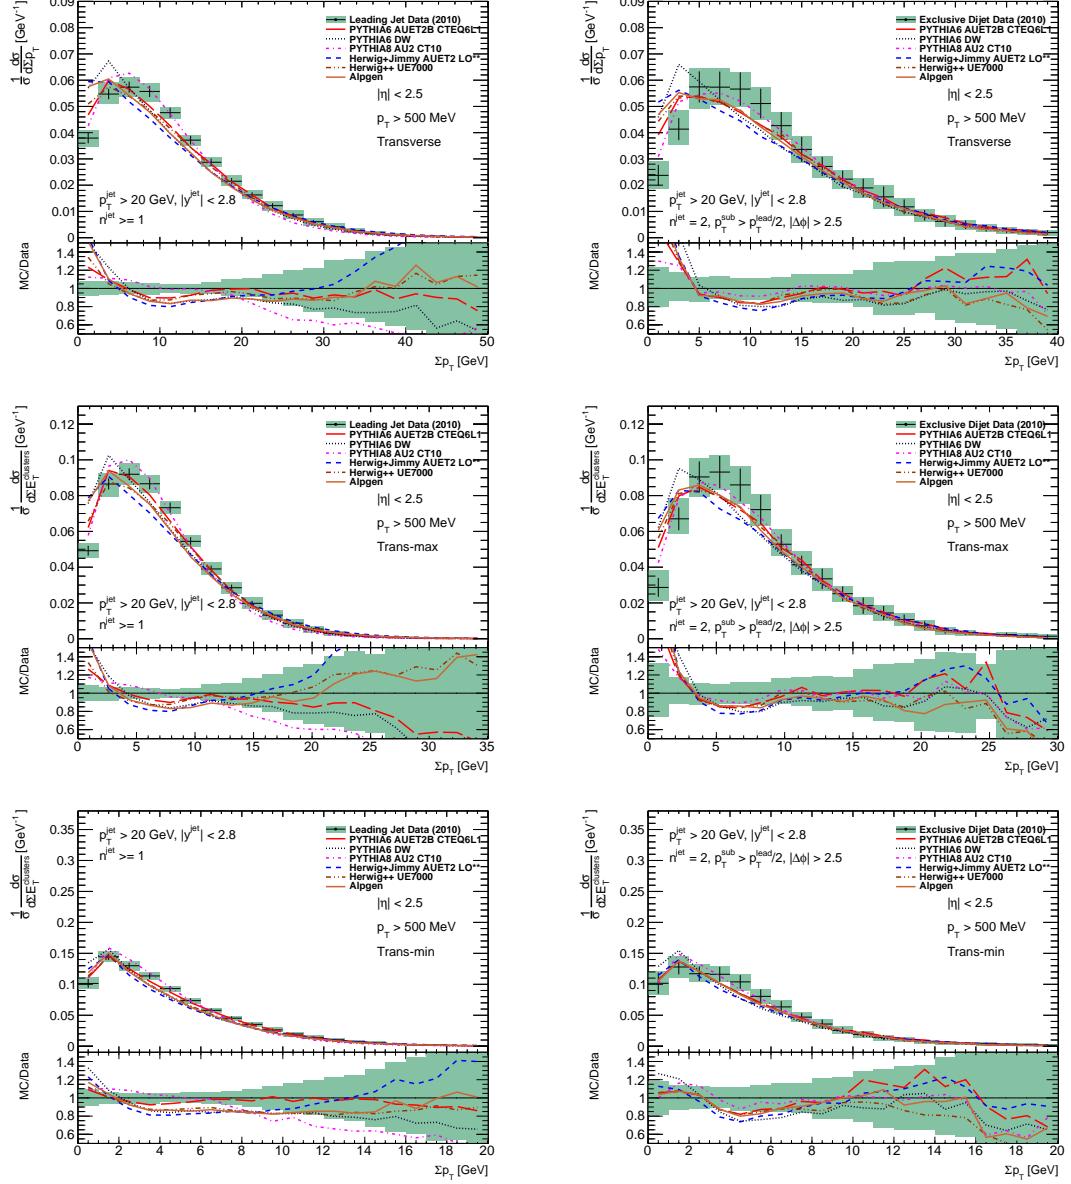


Figure D.1.: Charged particle $\sum p_T$ distributions shown from top to bottom for the transverse, trans-max and trans-min regions. Inclusive leading jet and exclusive di-jet selections in the left and right columns respectively.

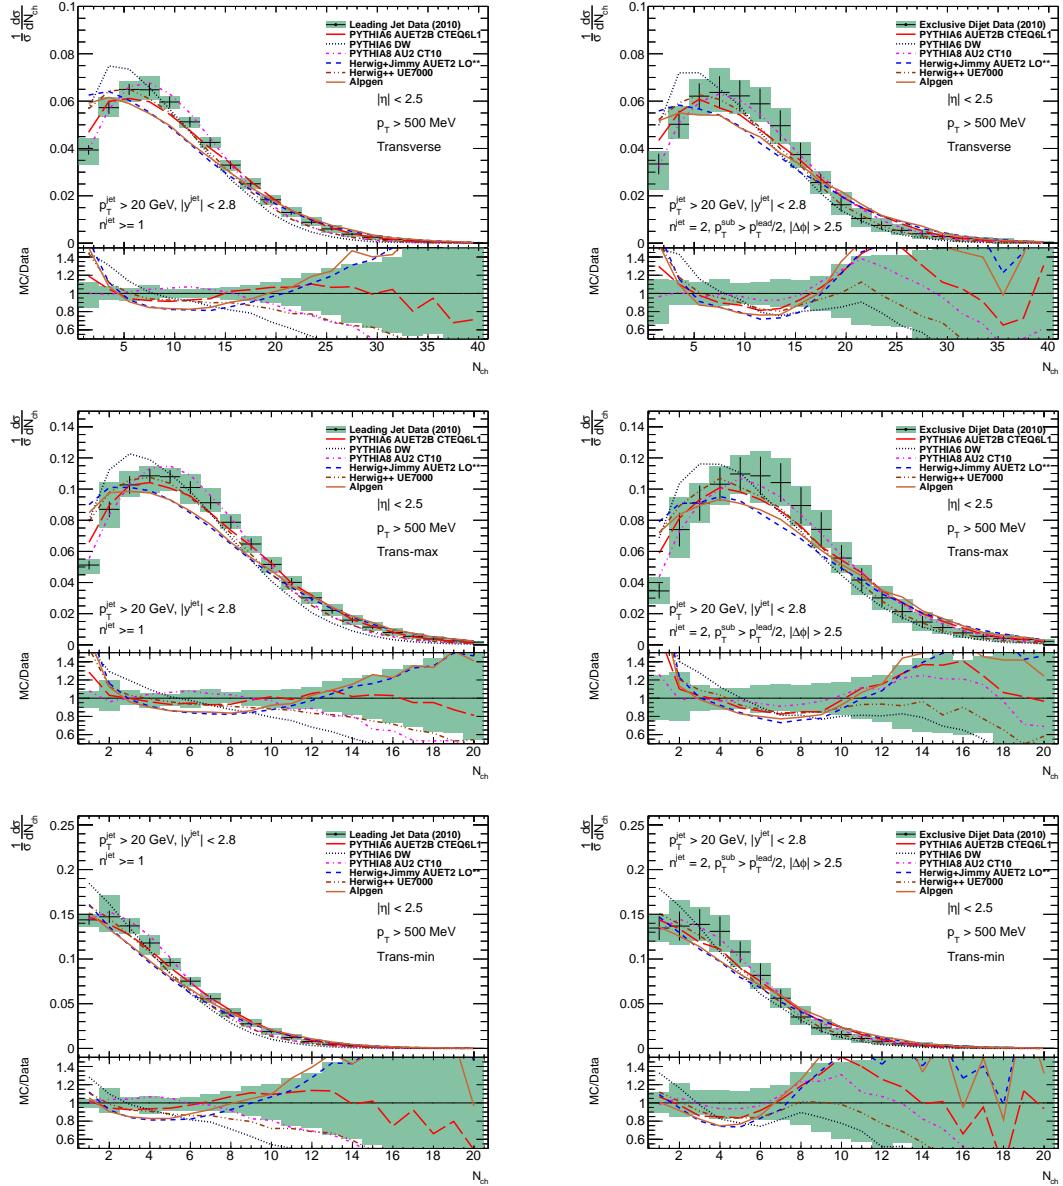


Figure D.2.: Charged particle N_{ch} distributions shown from top to bottom for the transverse, trans-max and trans-min regions. Inclusive leading jet and exclusive di-jet selections in the left and right columns respectively.

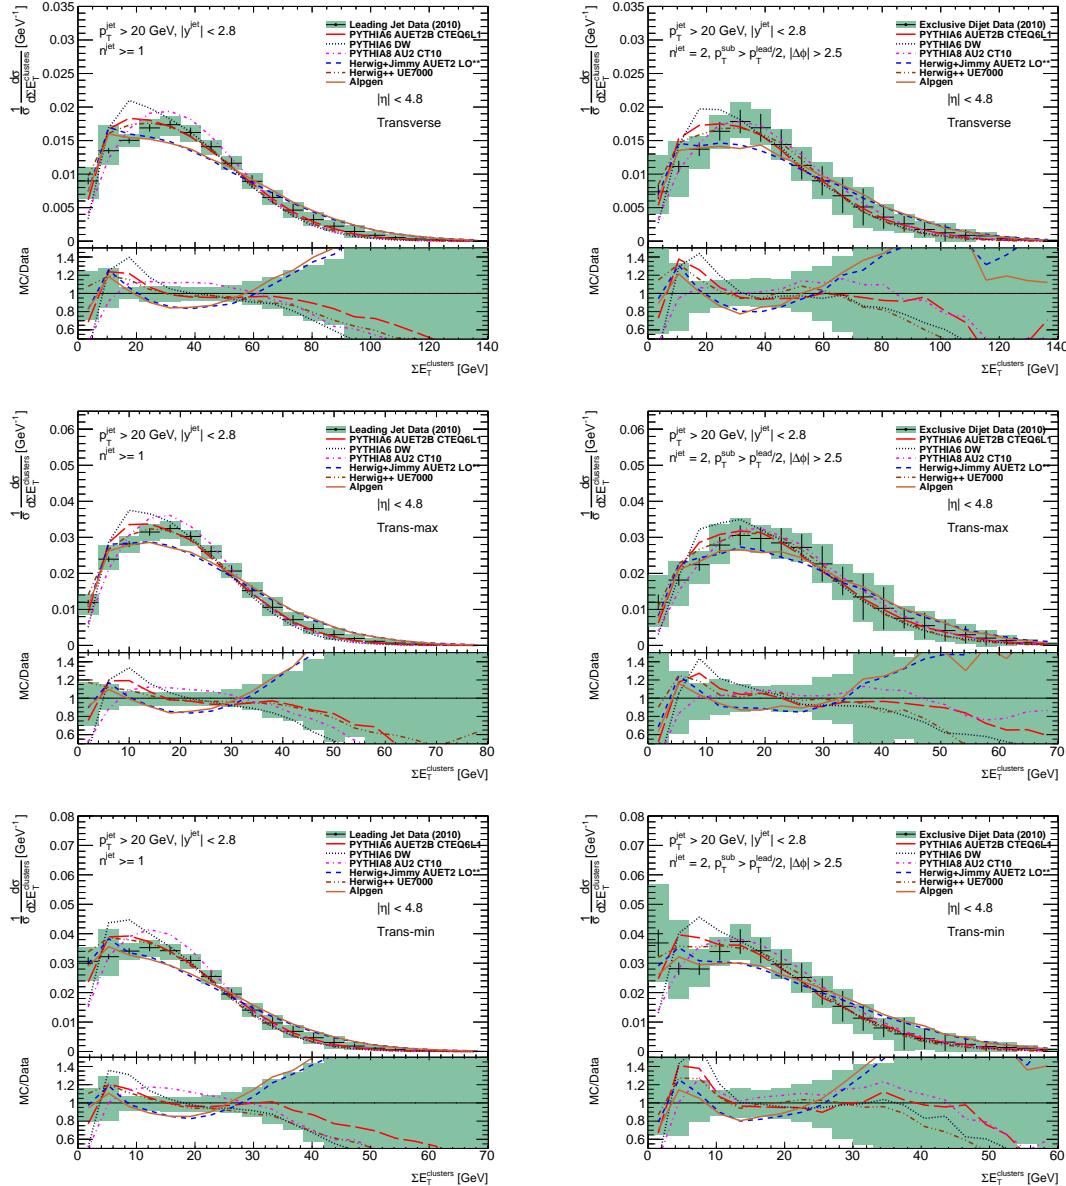


Figure D.3.: Topocluster $\sum E_T$ distributions ($|\eta| < 4.8$) shown from top to bottom for the transverse, trans-max and trans-min regions. Inclusive leading jet and exclusive di-jet selections in the left and right columns respectively.

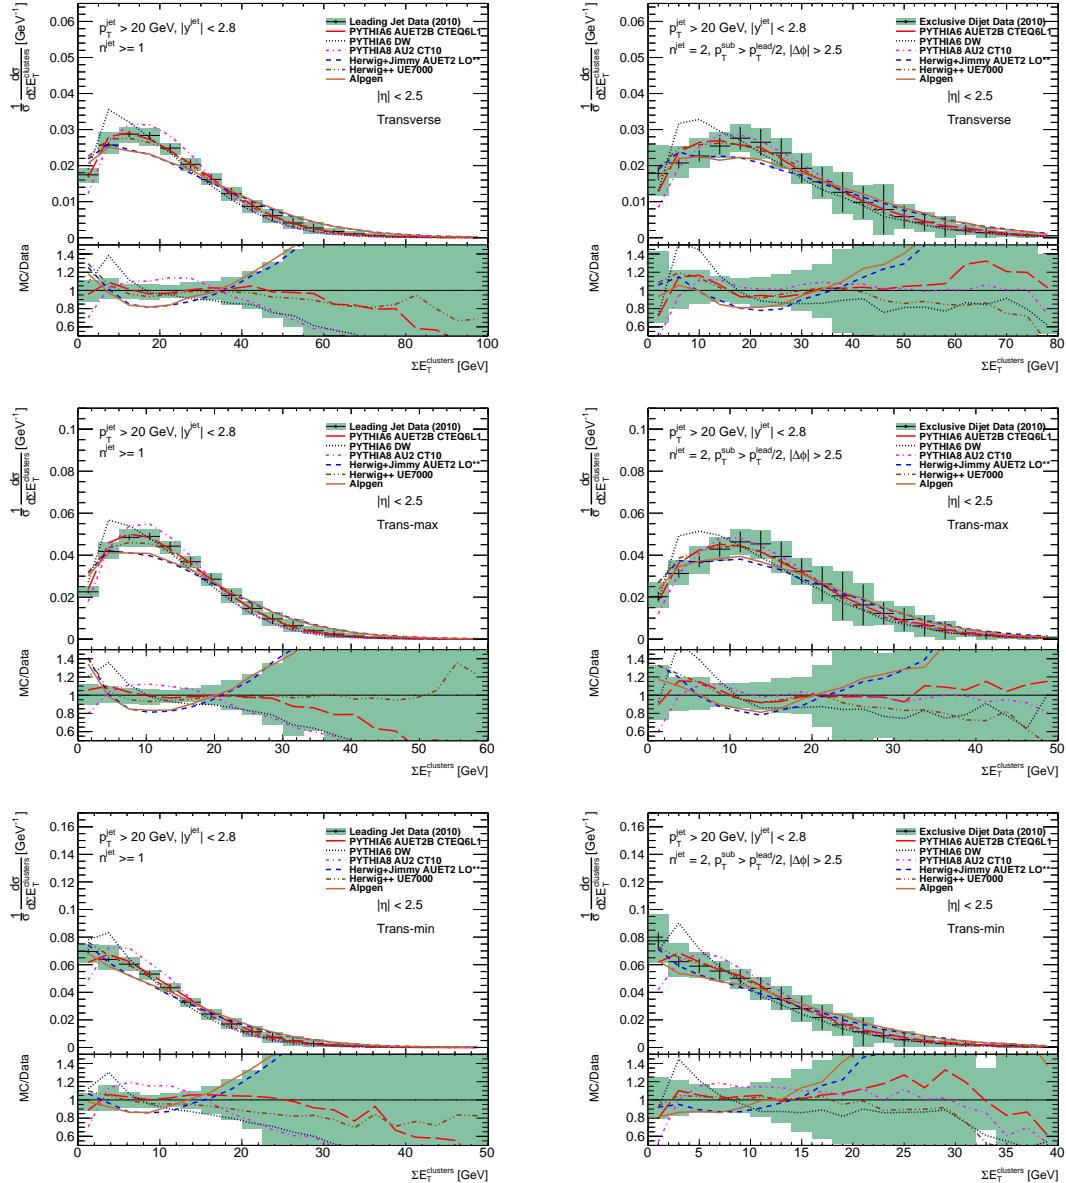


Figure D.4.: Topocluster $\sum E_T$ distributions ($|\eta| < 2.5$) shown from top to bottom for the transverse, trans-max and trans-min regions. Inclusive leading jet and exclusive di-jet selections in the left and right columns respectively.

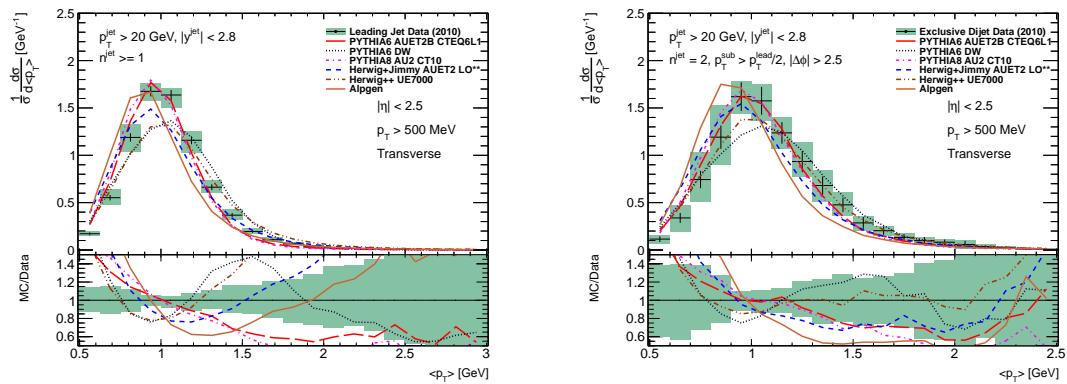


Figure D.5.: Charged particle $\langle p_T \rangle$ distributions for the transverse region. Inclusive leading jet and exclusive di-jet selections in the left and right columns respectively.

D.2. Restricted p_T^{lead} ranges

The distributions of underlying event observable values, measured in the transverse region and plotted for three different p_T^{lead} ranges.

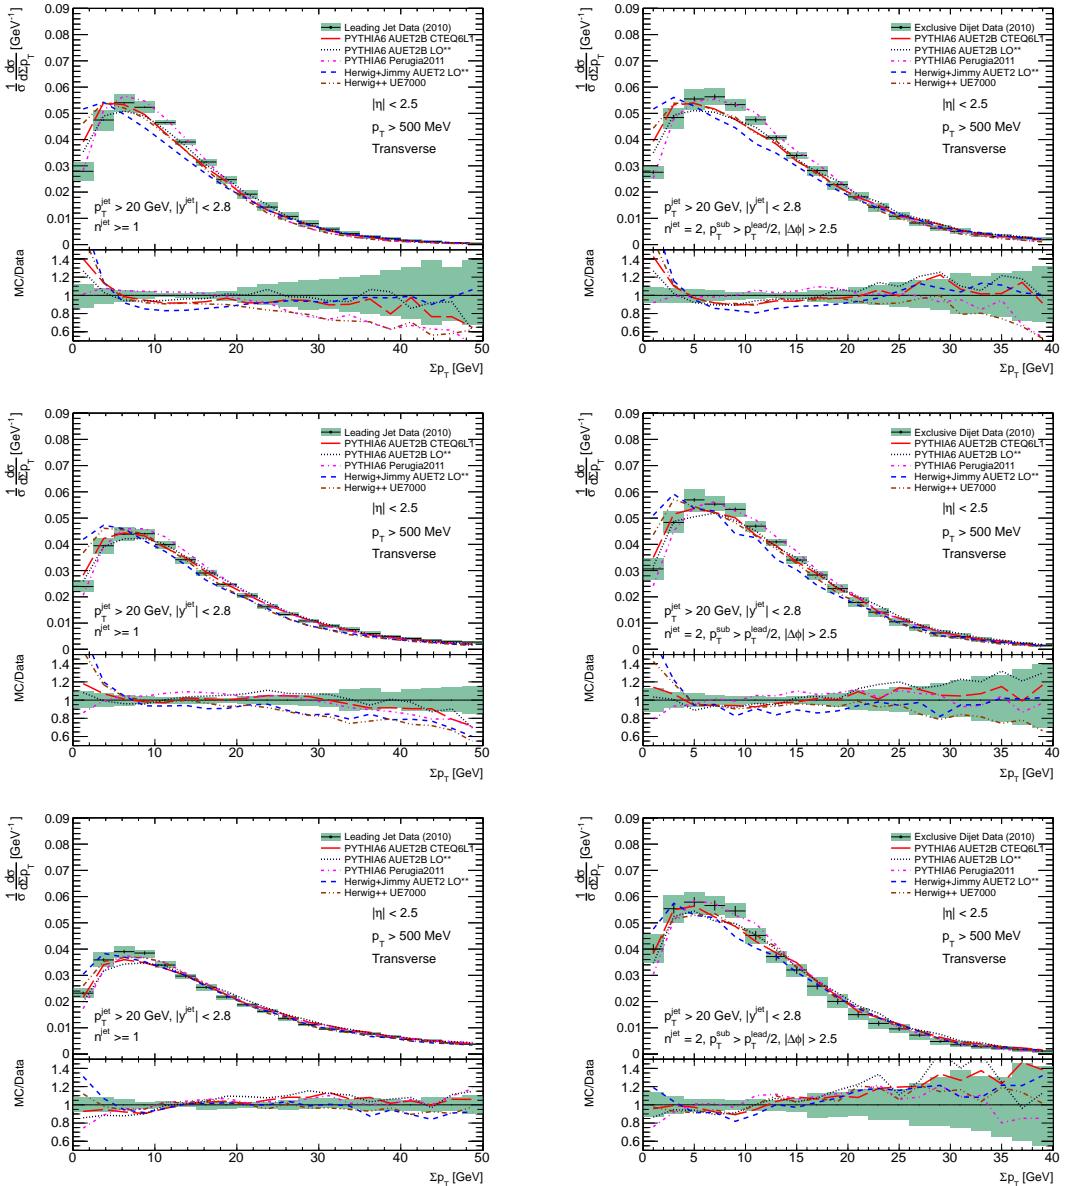


Figure D.6.: Charged particle $\sum p_T$ distributions for the transverse region shown from top to bottom for p_T^{lead} from 20–60 GeV, 60–210 GeV, and > 210 GeV. Inclusive leading jet and exclusive di-jet selections in the left and right columns respectively.

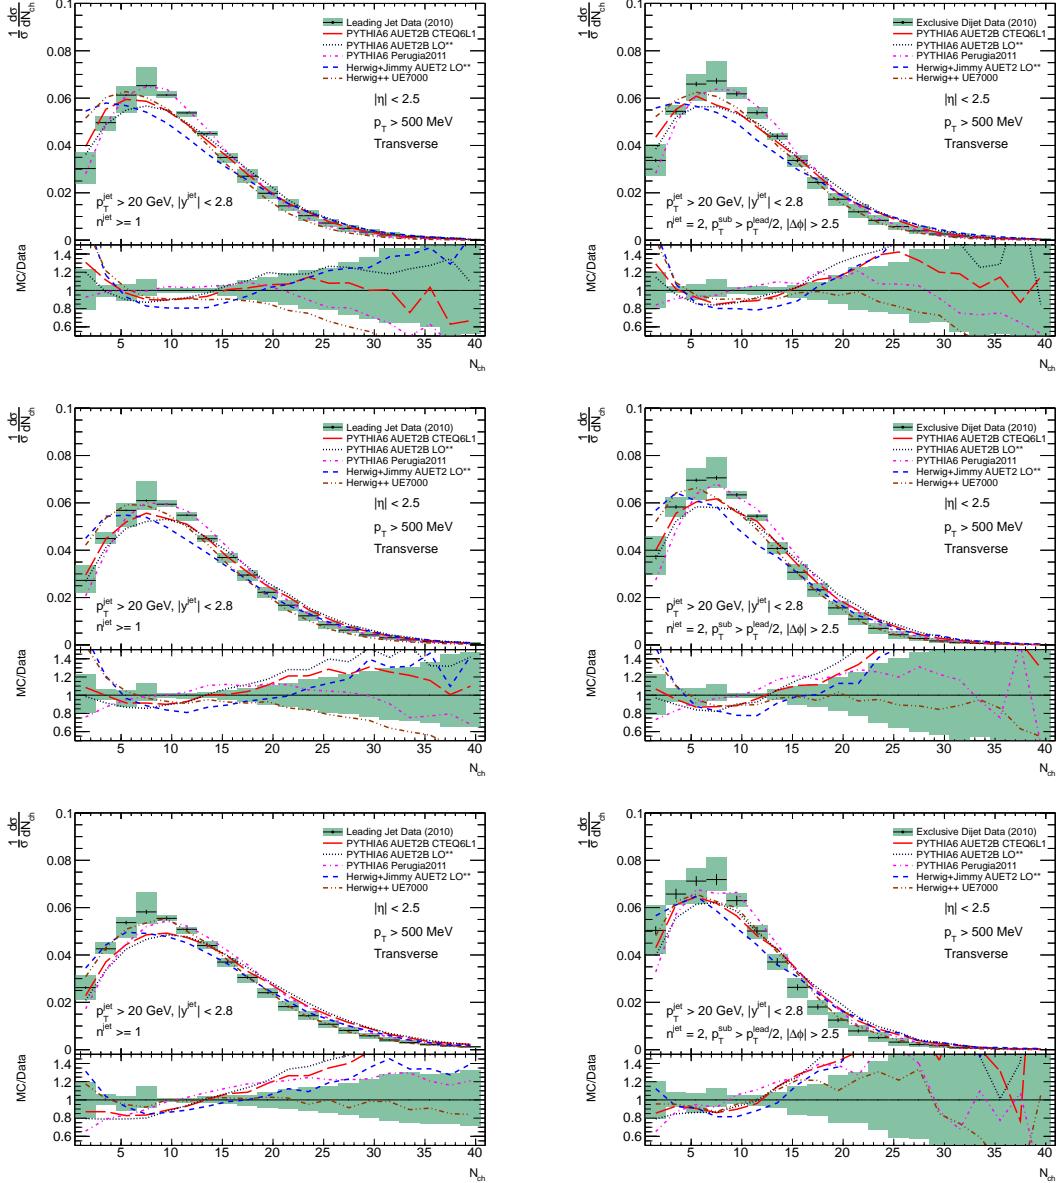


Figure D.7.: Charged particle N_{ch} distributions for the transverse region shown from top to bottom for p_T^{lead} from 20–60 GeV, 60–210 GeV, and > 210 GeV. Inclusive leading jet and exclusive di-jet selections in the left and right columns respectively.

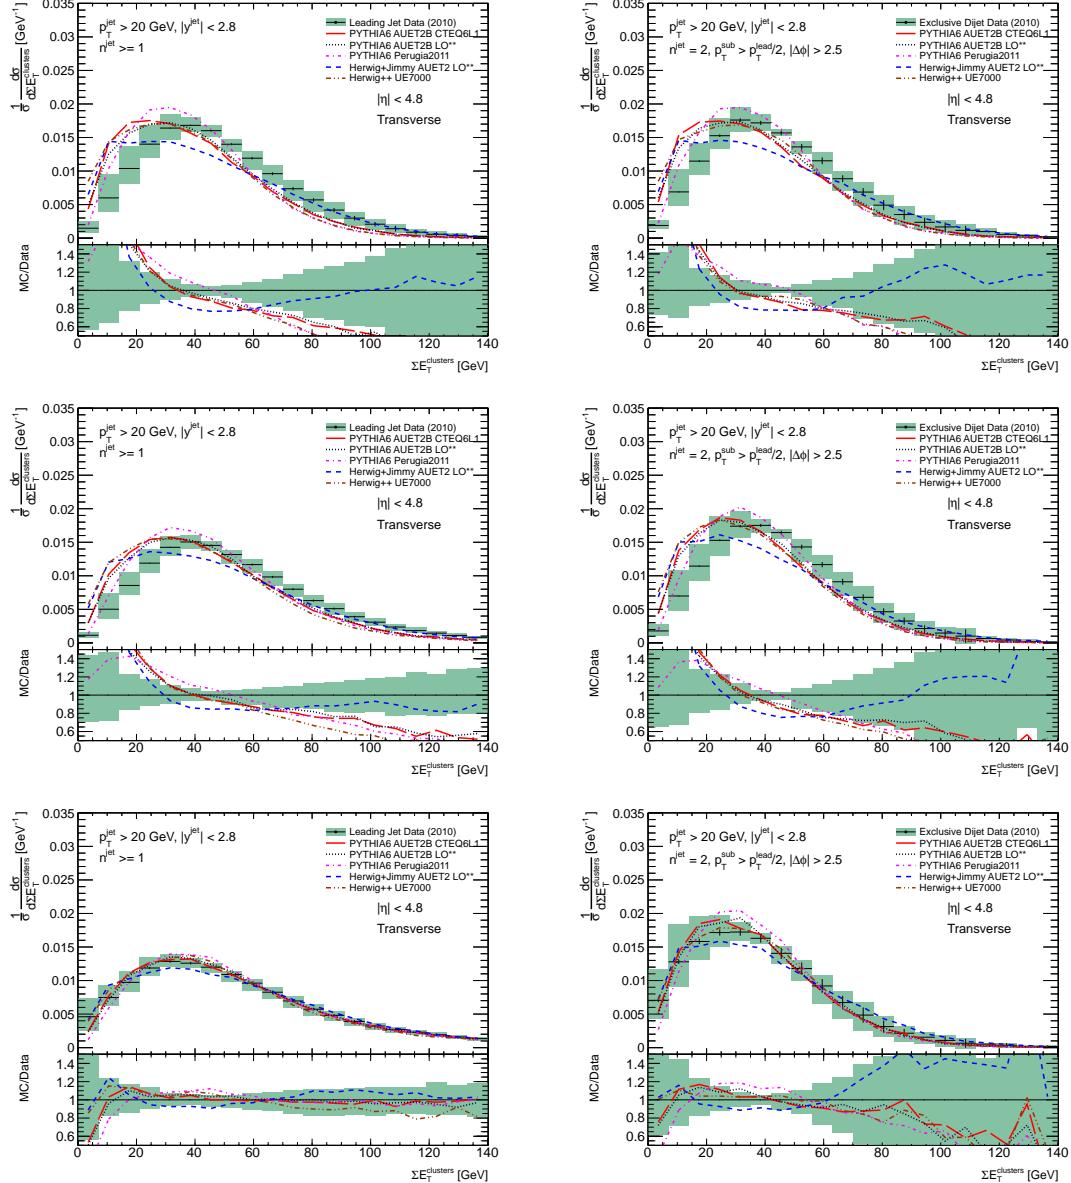


Figure D.8.: Topocluster $\sum E_T$ distributions ($|\eta| < 4.8$) for the transverse region shown from top to bottom for p_T^{lead} from $20 - 60 \text{ GeV}$, $60 - 210 \text{ GeV}$, and $> 210 \text{ GeV}$. Inclusive leading jet and exclusive di-jet selections in the left and right columns respectively.

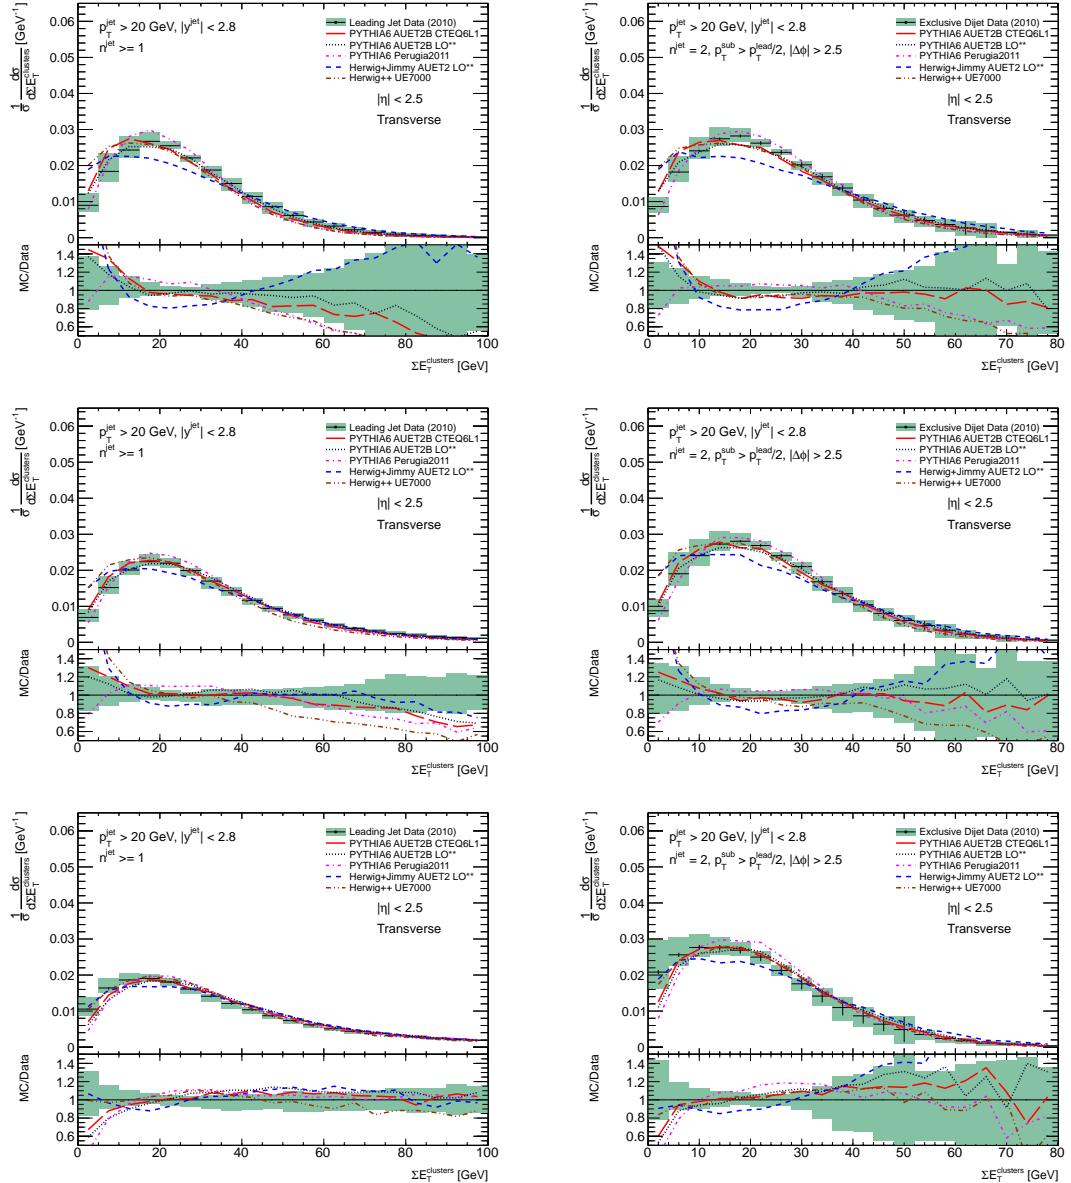


Figure D.9.: Topocluster $\sum E_T$ distributions ($|\eta| < 2.5$) for the transverse region shown from top to bottom for p_T^{lead} from $20 - 60 \text{ GeV}$, $60 - 210 \text{ GeV}$, and $> 210 \text{ GeV}$. Inclusive leading jet and exclusive di-jet selections in the left and right columns respectively.

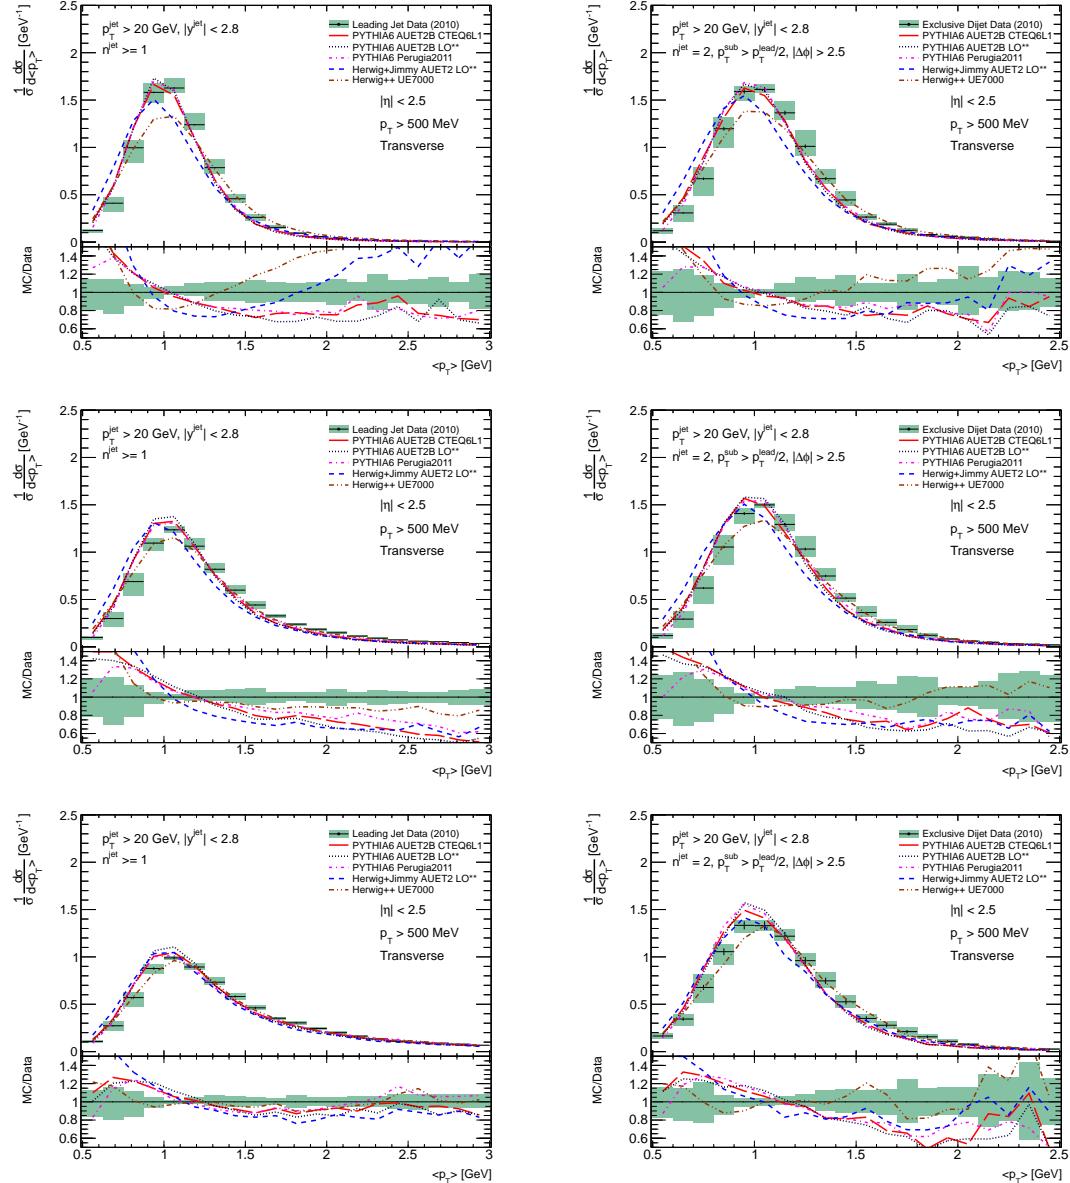


Figure D.10.: Charged particle $\langle p_T \rangle$ distributions for the transverse region shown from top to bottom for p_T^{lead} from $20 - 60 \text{ GeV}$, $60 - 210 \text{ GeV}$, and $> 210 \text{ GeV}$. Inclusive leading jet and exclusive di-jet selections in the left and right columns respectively.

Bibliography

- [1] *LEP design report: vol. 2. The LEP main ring*, Tech. Rep. CERN-LEP-84-01, Geneva, 1984. <https://cdsweb.cern.ch/record/102083>.
- [2] L. Evans and P. Bryant, *LHC Machine*, **JINST** **3** (2008) S08001.
- [3] ATLAS Collaboration, G. Aad et al., *The ATLAS experiment at the CERN Large Hadron Collider*, **JINST** **3** (2008) S08003.
- [4] CMS Collaboration, S. Chatrchyan et al., *The CMS experiment at the CERN Large Hadron Collider*, **JINST** **3** (2008) S08004.
- [5] ALICE Collaboration, K. Aamodt et al., *The ALICE experiment at the CERN Large Hadron Collider*, **JINST** **3** (2008) S08002.
- [6] LHCb Collaboration, J. Alves, A. Augusto et al., *The LHCb detector at the Large Hadron Collider*, **JINST** **3** (2008) S08005.
- [7] TOTEM Collaboration, G. Anelli et al., *The TOTEM experiment at the CERN Large Hadron Collider*, **JINST** **3** (2008) S08007.
- [8] LHCf Collaboration, O. Adriani et al., *The LHCf detector at the CERN Large Hadron Collider*, **JINST** **3** (2008) S08006.
- [9] J. Pinfold et al., *Technical Design Report of the MoEDAL Experiment*, Tech. Rep. CERN-LHCC-2009-006. MoEDAL-TDR-001, CERN, Geneva, Jun, 2009.
- [10] TeVI Collaboration, *Design Report Tevatron 1 project*, Tech. Rep. FERMILAB-DESIGN-1984-01, 1984.
- [11] C. Lefevre, *LHC: the guide (English version). Guide du LHC (version anglaise)*, Feb, 2009.
- [12] P. Collier, B. Goddard, R. Jung, K. Kissler, T. Linnecar, et al., *The SPS as injector for LHC: Conceptual design*, Tech. Rep. CERN-SL-97-007-DI, 1997.

- [13] D. Boussard and T. Linnecar, *The LHC Superconducting RF System*, Adv. Cryog. Eng. **45A** (2000) 835–844.
- [14] O. S. Bruning, P. Collier, P. Lebrun, S. Myers, R. Ostojic, et al., *LHC Design Report. 1. The LHC Main Ring*, Tech. Rep. CERN-2004-003-V-1, 2004.
- [15] J. Pequenao, *Computer generated image of the whole ATLAS detector*, Mar, 2008.
- [16] ATLAS Collaboration, G. Aad et al., *Commissioning of the ATLAS Muon Spectrometer with Cosmic Rays*, Eur. Phys. J. **C70** (2010) 875–916, [arXiv:1006.4384 \[physics.ins-det\]](https://arxiv.org/abs/1006.4384).
- [17] O. Kortner, E. Moyse, and A. Salvucci, *Muon resolution plots*, Tech. Rep. ATL-COM-PHYS-2011-1504, CERN, Geneva, Nov, 2011.
- [18] J. Pequenao, *Computer Generated image of the ATLAS calorimeter*, Mar, 2008.
- [19] J. Pequenao, *Computer generated image of the ATLAS Inner Detector*, Mar, 2008.
- [20] *ATLAS Inner Detector: Technical Design Report, 1*. CERN, Geneva, 1997. <http://cdsweb.cern.ch/record/331063>.
- [21] ATLAS Collaboration, G. Aad et al., *ATLAS pixel detector electronics and sensors*, J. Instrum. **3** (2008) P07007.
- [22] J. Pequenao, *Computer generated image of the pixel detector, part of the ATLAS Inner Detector*, Mar, 2008.
- [23] P. Coe, D. Howell, and R. Nickerson, *Frequency scanning interferometry in ATLAS: Remote, multiple, simultaneous and precise distance measurements in a hostile environment*, Measur. Sci. Tech. **15** (2004) 2175–2187.
- [24] ATLAS TRT Collaboration, E. Abat et al., *The ATLAS Transition Radiation Tracker (TRT) proportional drift tube: Design and performance*, JINST **3** (2008) P02013.
- [25] ATLAS TRT Collaboration, E. Abat et al., *The ATLAS TRT barrel detector*, JINST **3** (2008) P02014.
- [26] S. Haywood, L. Rossi, R. Nickerson, and A. Romanouk, *ATLAS Inner Detector: Technical Design Report, 2*. CERN, Geneva, 1997. <https://cdsweb.cern.ch/record/331064>.

- [27] G. Lutz, *A Simplistic model for reverse annealing in irradiated silicon*, *Nucl. Instrum. Meth.* **B95** (1994) no. 1, 41 – 49.
- [28] M. Moll, *Radiation damage in silicon particle detectors: microscopic defects and macroscopic properties*. PhD thesis, Hamburg Univ., 1999.
<http://cdsweb.cern.ch/record/425274>.
- [29] D. Attree, B. Anderson, E. Anderssen, V. Akhnazarov, R. Apsimon, et al., *The evaporative cooling system for the ATLAS Inner Detector*, *JINST* **3** (2008) P07003.
- [30] ATLAS Collaboration, G. Aad et al., *The ATLAS Inner Detector commissioning and calibration*, *Eur. Phys. J.* **C70** (2010) 787–821, [arXiv:1004.5293](https://arxiv.org/abs/1004.5293) [physics.ins-det].
- [31] S. Glashow, *Partial Symmetries of Weak Interactions*, *Nucl. Phys.* **22** (1961) 579–588.
- [32] S. Weinberg, *A Model of Leptons*, *Phys. Rev. Lett.* **19** (1967) 1264–1266.
- [33] A. Salam and J. Ward, *Electromagnetic and weak interactions*, *Phys. Lett.* **13** (1964) no. 2, 168–171.
- [34] Wikipedia, *Standard Model of Elementary Particles*, 2012.
http://en.wikipedia.org/wiki/Standard_Model.
- [35] F. Englert and R. Brout, *Broken Symmetry and the Mass of Gauge Vector Mesons*, *Phys. Rev. Lett.* **13** (1964) 321–323.
- [36] P. W. Higgs, *Broken Symmetries and the Masses of Gauge Bosons*, *Phys. Rev. Lett.* **13** (1964) 508–509.
- [37] G. S. Guralnik, C. R. Hagen, and T. W. B. Kibble, *Global Conservation Laws and Massless Particles*, *Phys. Rev. Lett.* **13** (1964) 585–587.
- [38] P. W. Higgs, *Spontaneous Symmetry Breakdown without Massless Bosons*, *Phys. Rev.* **145** (1966) 1156–1163.
- [39] ATLAS Collaboration, G. Aad et al., *Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC*, *Phys. Lett.* **B716** (2012) no. 1, 1 – 29, [arXiv:1207.7214](https://arxiv.org/abs/1207.7214).
- [40] S. Chatrchyan et al., *Observation of a new boson at a mass of 125 GeV with the CMS experiment at the LHC*, *Phys. Lett.* **B716** (2012) no. 1, 30 – 61,

- [arXiv:1207.7235](https://arxiv.org/abs/1207.7235).
- [41] M. Roos, *Dark Matter: The evidence from astronomy, astrophysics and cosmology*, [arXiv:1001.0316 \[astro-ph.CO\]](https://arxiv.org/abs/1001.0316).
 - [42] H. Yukawa, *On the interaction of elementary particles*, *Proc. Phys. Math. Soc. Jap.* **17** (1935) 48–57.
 - [43] C. Lattes, H. Muirhead, G. Occhialini, and C. Powell, *Processes involving charged mesons*, *Nature* **159** (1947) 694–697.
 - [44] M. Gell-Mann and Y. Ne’eman, *The eightfold way*. Frontiers in Physics. Benjamin, New York, NY, 1964.
 - [45] M. Breidenbach, J. I. Friedman, H. W. Kendall, E. D. Bloom, D. Coward, et al., *Observed Behavior of Highly Inelastic electron-Proton Scattering*, *Phys. Rev. Lett.* **23** (1969) 935–939.
 - [46] M. Gell-Mann, *A Schematic Model of Baryons and Mesons*, *Phys. Lett.* **8** (1964) 214–215.
 - [47] Particle Data Group Collaboration, J. Beringer et al., *Review of Particle Physics*, *Phys. Rev.* **D86** (2012) 010001. <http://pdg.lbl.gov>.
 - [48] V. Barnes, P. Connolly, D. Crennell, B. Culwick, W. Delaney, et al., *Observation of a Hyperon with Strangeness -3*, *Phys. Rev. Lett.* **12** (1964) 204–206.
 - [49] F. Tkachov, *A contribution to the history of quarks: Boris Struminsky’s 1965 JINR publication*, [arXiv:0904.0343 \[physics.hist-ph\]](https://arxiv.org/abs/0904.0343).
 - [50] H. D. Politzer, *Reliable Perturbative Results for Strong Interactions?*, *Phys. Rev. Lett.* **30** (1973) 1346–1349.
 - [51] S. Weinberg, *Non-Abelian Gauge Theories of the Strong Interactions*, *Phys. Rev. Lett.* **31** (1973) 494–497.
 - [52] D. J. Gross and F. Wilczek, *Ultraviolet Behavior of Non-Abelian Gauge Theories*, *Phys. Rev. Lett.* **30** (1973) 1343–1346.
 - [53] J. Carrington, *QCD Feynman diagrams*, private communication, 2012.
 - [54] S. Bethke, *The 2009 World Average of alpha(s)*, *Eur. Phys. J.* **C64** (2009) 689–703, [arXiv:0908.1135 \[hep-ph\]](https://arxiv.org/abs/0908.1135).

- [55] E. Eichten, K. Gottfried, T. Kinoshita, J. Kogut, K. D. Lane, and T. M. Yan, *Spectrum of Charmed Quark-Antiquark Bound States*, *Phys. Rev. Lett.* **34** (1975) 369–372.
- [56] J. C. Collins and D. E. Soper, *The Theorems of Perturbative QCD*, *Annual Review of Nuclear and Particle Science* **37** (1987) no. 1, 383–409.
- [57] K. G. Wilson, *Confinement of quarks*, *Phys. Rev.* **D10** (1974) 2445–2459.
- [58] A. Martin, W. Stirling, R. Thorne, and G. Watt, *Parton distributions for the LHC*, *Eur. Phys. J.* **C63** (2009) 189–285, [arXiv:0901.0002 \[hep-ph\]](https://arxiv.org/abs/0901.0002).
- [59] NNPDF Collaboration, R. D. Ball et al., *A determination of parton distributions with faithful uncertainty estimation*, *Nucl. Phys.* **B809** (2009) 1–63, [arXiv:0808.1231 \[hep-ph\]](https://arxiv.org/abs/0808.1231).
- [60] M. Bahr, S. Gieseke, and M. H. Seymour, *Simulation of multiple partonic interactions in Herwig++*, *JHEP* **0807** (2008) 076, [arXiv:0803.3633 \[hep-ph\]](https://arxiv.org/abs/0803.3633).
- [61] M. Bahr, J. M. Butterworth, S. Gieseke, and M. H. Seymour, *Soft interactions in Herwig++*, [arXiv:0905.4671 \[hep-ph\]](https://arxiv.org/abs/0905.4671).
- [62] Axial Field Spectrometer Collaboration, T. Åkesson et al., *Double parton scattering in pp collisions at $\sqrt{s} = 63$ GeV*, *Z. Phys.* **C34** (1986) 163–174.
- [63] D0 Collaboration, V. Abazov et al., *Double parton interactions in photon+3 jet events in $p\bar{p}$ collisions $\sqrt{s} = 1.96$ TeV*, *Phys. Rev.* **D81** (2010) 052012, [arXiv:0912.5104 \[hep-ex\]](https://arxiv.org/abs/0912.5104).
- [64] Y. L. Dokshitzer, *Calculation of the Structure Functions for Deep Inelastic Scattering and $e^+ e^-$ Annihilation by Perturbation Theory in Quantum Chromodynamics.*, Sov. Phys. JETP **46** (1977) 641–653.
- [65] V. N. Gribov and L. N. Lipatov, *Deep inelastic ep scattering in perturbation theory*, Sov. J. Nucl. Phys. **15** (1972) no. 4, 438–450.
- [66] G. Altarelli and G. Parisi, *Asymptotic Freedom in Parton Language*, *Nucl. Phys.* **B126** (1977) 298.
- [67] T. Sjöstrand, *A Model for Initial State Parton Showers*, *Phys. Lett.* **B157** (1985) 321.
- [68] T. Lee and M. Nauenberg, *Degenerate Systems and Mass Singularities*, *Phys. Rev.*

- 133 (1964) B1549–B1562.
- [69] T. Kinoshita and A. Ukawa, *New Approach to the Singularities of Feynman Amplitudes in the Zero Mass Limit*, Phys. Rev. **D13** (1976) 1573.
- [70] M. Cacciari, G. P. Salam, and G. Soyez, *The Anti- $k(t)$ jet clustering algorithm*, JHEP **0804** (2008) 063, [arXiv:0802.1189 \[hep-ph\]](https://arxiv.org/abs/0802.1189).
- [71] S. Catani, Y. L. Dokshitzer, M. Seymour, and B. Webber, *Longitudinally invariant K_t clustering algorithms for hadron hadron collisions*, Nucl. Phys. **B406** (1993) 187–224.
- [72] Y. L. Dokshitzer, G. Leder, S. Moretti, and B. Webber, *Better jet clustering algorithms*, JHEP **9708** (1997) 001, [arXiv:hep-ph/9707323 \[hep-ph\]](https://arxiv.org/abs/hep-ph/9707323).
- [73] M. Cacciari, G. P. Salam, and G. Soyez, *FastJet User Manual*, Eur. Phys. J. **C72** (2012) 1896, [arXiv:1111.6097 \[hep-ph\]](https://arxiv.org/abs/1111.6097).
- [74] L. Asquith et al., *Performance of Jet Algorithms in the ATLAS Detector*, Tech. Rep. ATL-PHYS-INT-2010-129, CERN, Geneva, Dec, 2010.
- [75] G. P. Salam and G. Soyez, *A Practical Seedless Infrared-Safe Cone jet algorithm*, JHEP **0705** (2007) 086, [arXiv:0704.0292 \[hep-ph\]](https://arxiv.org/abs/0704.0292).
- [76] W. Stirling, *Standard Model cross sections as a function of collider energy, with 125 GeV Higgs*, private communication, 2012.
- [77] R. D. Harrington, *Measurement of the top quark mass in lepton+jets events with secondary vertex tagging*. PhD thesis, 2007. FERMILAB-THESIS-2007-25.
- [78] A. Buckley, J. Butterworth, S. Gieseke, D. Grellscheid, S. Hoche, et al., *General-purpose event generators for LHC physics*, Phys. Rept. **504** (2011) 145–233, [arXiv:1101.2599 \[hep-ph\]](https://arxiv.org/abs/1101.2599).
- [79] ATLAS tunes of PYTHIA 6 and Pythia 8 for MC11, tech. rep., CERN, Geneva, Jul, 2011. <https://cdsweb.cern.ch/record/1363300>.
- [80] T. Gleisberg, S. Höche, F. Krauss, M. Schönherr, S. Schumann, et al., *Event generation with SHERPA 1.1*, JHEP **0902** (2009) 007, [arXiv:0811.4622 \[hep-ph\]](https://arxiv.org/abs/0811.4622).
- [81] T. Sjöstrand, S. Mrenna, and P. Z. Skands, *PYTHIA 6.4 physics and manual*, JHEP **0605** (2006) 026, [arXiv:hep-ph/0603175 \[hep-ph\]](https://arxiv.org/abs/hep-ph/0603175).

- [82] T. Sjöstrand, S. Mrenna, and P. Z. Skands, *A brief introduction to PYTHIA 8.1*, *Comput. Phys. Commun.* **178** (2008) 852–867, [arXiv:0710.3820 \[hep-ph\]](https://arxiv.org/abs/0710.3820).
- [83] G. Corcella, I. Knowles, G. Marchesini, S. Moretti, K. Odagiri, et al., HERWIG6: *An event generator for hadron emission reactions with interfering gluons (including supersymmetric processes)*, *JHEP* **0101** (2001) 010, [arXiv:hep-ph/0011363 \[hep-ph\]](https://arxiv.org/abs/hep-ph/0011363).
- [84] M. Bahr, S. Gieseke, M. Gigg, D. Grellscheid, K. Hamilton, et al., HERWIG++ *physics and manual*, *Eur. Phys. J.* **C58** (2008) 639–707, [arXiv:0803.0883 \[hep-ph\]](https://arxiv.org/abs/0803.0883).
- [85] J. Butterworth, J. R. Forshaw, and M. Seymour, *Multiparton interactions in photoproduction at HERA*, *Z. Phys. C* **72** (1996) 637–646, [arXiv:hep-ph/9601371 \[hep-ph\]](https://arxiv.org/abs/hep-ph/9601371).
- [86] M. L. Mangano, M. Moretti, F. Piccinini, R. Pittau, and A. D. Polosa, *ALPGEN, a generator for hard multiparton processes in hadronic collisions*, *JHEP* **0307** (2003) 001, [arXiv:hep-ph/0206293 \[hep-ph\]](https://arxiv.org/abs/hep-ph/0206293).
- [87] R. Field, *Min-bias and the underlying event at the Tevatron and the LHC*, A talk presented at the Fermilab MC Tuning Workshop, Fermilab, Oct, 2002.
- [88] CTEQ Collaboration, H. Lai et al., *Global QCD analysis of parton structure of the nucleon: CTEQ5 parton distributions*, *Eur. Phys. J.* **C12** (2000) 375–392, [arXiv:hep-ph/9903282 \[hep-ph\]](https://arxiv.org/abs/hep-ph/9903282).
- [89] J. Pumplin, D. Stump, J. Huston, H. Lai, P. M. Nadolsky, et al., *New generation of parton distributions with uncertainties from global QCD analysis*, *JHEP* **0207** (2002) 012, [arXiv:hep-ph/0201195 \[hep-ph\]](https://arxiv.org/abs/hep-ph/0201195).
- [90] *Summary of ATLAS Pythia 8 tunes*, tech. rep., CERN, Geneva, Aug, 2012. <https://cdsweb.cern.ch/record/1474107>.
- [91] H.-L. Lai, M. Guzzi, J. Huston, Z. Li, P. M. Nadolsky, et al., *New parton distributions for collider physics*, *Phys. Rev. D* **82** (2010) 074024, [arXiv:1007.2241 \[hep-ph\]](https://arxiv.org/abs/1007.2241).
- [92] A. Sherstnev and R. S. Thorne, *Parton distributions for LO generators*, *Eur. Phys. J.* **C55** (2008) 553–575, [arXiv:0711.2473 \[hep-ph\]](https://arxiv.org/abs/0711.2473).
- [93] S. Gieseke, C. Rohr, and A. Siodmok, *Multiple Partonic Interaction Developments*

- in Herwig++*, [arXiv:1110.2675 \[hep-ph\]](https://arxiv.org/abs/1110.2675).
- [94] M. L. Mangano, M. Moretti, and R. Pittau, *Multijet matrix elements and shower evolution in hadronic collisions: $Wb\bar{b} + n$ jets as a case study*, *Nucl. Phys.* **B632** (2002) 343–362, [arXiv:hep-ph/0108069 \[hep-ph\]](https://arxiv.org/abs/hep-ph/0108069).
 - [95] ATLAS Collaboration, G. Brandt, *Charged particle multiplicities in inelastic pp interactions with ATLAS and the ATLAS Minimum Bias Tune 1*, .
 - [96] T. Sjöstrand and M. van Zijl, *A Multiple Interaction Model for the Event Structure in Hadron Collisions*, *Phys. Rev.* **D36** (1987) 2019.
 - [97] V. Sudakov, *Vertex parts at very high-energies in quantum electrodynamics*, *Sov. Phys. JETP* **3** (1956) 65–71.
 - [98] B. Andersson, *The Lund model*, Camb. Monogr. Part. Phys. Nucl. Phys. Cosmol. **7** (1997) 1–471.
 - [99] D. Amati and G. Veneziano, *Preconfinement as a Property of Perturbative QCD*, *Phys. Lett.* **B83** (1979) 87.
 - [100] B. Webber, *A QCD Model for Jet Fragmentation Including Soft Gluon Interference*, *Nucl. Phys.* **B238** (1984) 492.
 - [101] CDF Collaboration, D. Acosta et al., *The underlying event in hard interactions at the Tevatron $\bar{p}p$ collider*, *Phys. Rev.* **D70** (2004) 072002, [arXiv:hep-ex/0404004 \[hep-ex\]](https://arxiv.org/abs/hep-ex/0404004).
 - [102] CDF Collaboration, T. Aaltonen et al., *Studying the Underlying Event in Drell-Yan and High Transverse Momentum Jet Production at the Tevatron*, *Phys. Rev.* **D82** (2010) 034001, [arXiv:1003.3146 \[hep-ex\]](https://arxiv.org/abs/1003.3146).
 - [103] CDF Collaboration, T. Aaltonen et al., *Measurement of particle production and inclusive differential cross sections in $p\bar{p}$ collisions at $\sqrt{s} = 1.96$ TeV*, *Phys. Rev.* **D79** (2009) 112005, [arXiv:0904.1098 \[hep-ex\]](https://arxiv.org/abs/0904.1098).
 - [104] ATLAS Collaboration, G. Aad et al., *Charged-particle multiplicities in pp interactions measured with the ATLAS detector at the LHC*, *New J. Phys.* **13** (2011) 053033, [arXiv:1012.5104 \[hep-ex\]](https://arxiv.org/abs/1012.5104).
 - [105] J. Pumplin, *Hard underlying event correction to inclusive jet cross-sections*, *Phys. Rev.* **D57** (1998) 5787–5792, [arXiv:hep-ph/9708464 \[hep-ph\]](https://arxiv.org/abs/hep-ph/9708464).

- [106] ATLAS Collaboration, G. Aad et al., *Measurement of underlying event characteristics using charged particles in pp collisions at $\sqrt{s} = 900$ GeV and 7 TeV with the ATLAS detector*, *Phys. Rev.* **D83** (2011) 112001, [arXiv:1012.0791 \[hep-ex\]](https://arxiv.org/abs/1012.0791).
- [107] ALICE Collaboration, B. Abelev et al., *Underlying Event measurements in pp collisions at $\sqrt{s} = 0.9$ and 7 TeV with the ALICE experiment at the LHC*, *JHEP* **1207** (2012) 116, [arXiv:1112.2082 \[hep-ex\]](https://arxiv.org/abs/1112.2082).
- [108] J. S. Virzi and M. D. Shapiro, *A Measurement of the Underlying Event Distributions in Proton-Proton Collisions at $\sqrt{s} = 7$ TeV in Charged-Particle Jet Events using the ATLAS Detector at the Large Hadron Collider*. PhD thesis, UC, Berkeley, Nov, 2012. <https://cdsweb.cern.ch/record/1454635>.
- [109] CMS Collaboration, S. Chatrchyan et al., *Measurement of the Underlying Event Activity at the LHC with $\sqrt{s} = 7$ TeV and Comparison with $\sqrt{s} = 0.9$ TeV*, *JHEP* **1109** (2011) 109, [arXiv:1107.0330 \[hep-ex\]](https://arxiv.org/abs/1107.0330).
- [110] S. D. Drell and T.-M. Yan, *Massive Lepton-Pair Production in Hadron-Hadron Collisions at High Energies*, *Phys. Rev. Lett.* **25** (1970) 316–320.
- [111] CMS Collaboration, S. Chatrchyan et al., *Measurement of the underlying event in the Drell-Yan process in proton-proton collisions at $\sqrt{s} = 7$ TeV*, *Eur. Phys. J.* **C72** (2012) 2080, [arXiv:1204.1411 \[hep-ex\]](https://arxiv.org/abs/1204.1411).
- [112] D. Kar, A. Miagkov, V. Nikolaenko, B. Wynne, and O. Zenin, *Measurement of observables sensitive to the underlying event in inclusive Z-boson events at $\sqrt{s} = 7$ TeV*, tech. rep., CERN, Geneva, Jul, 2012. <https://cdsweb.cern.ch/record/1463358>.
- [113] ATLAS Collaboration, G. Aad et al., *Measurement of inclusive jet and dijet production in pp collisions at $\sqrt{s} = 7$ TeV using the ATLAS detector*, *Phys. Rev.* **D86** (2011) 014022, [arXiv:1112.6297 \[hep-ex\]](https://arxiv.org/abs/1112.6297).
- [114] S. Baker et al., *Measurement of inclusive jet and dijet production in pp collisions at $\sqrt{s} = 7$ TeV using the ATLAS detector*, tech. rep., CERN, Geneva, Jun, 2011. <https://cdsweb.cern.ch/record/1360174>.
- [115] W. Lampl et al., *Calorimeter Clustering Algorithms: Description and Performance*, tech. rep., CERN, Geneva, Apr, 2008. <https://cdsweb.cern.ch/record/1099735>.

- [116] J. Monk, E. Nurse, R. Prabhu, and P. Wijeratne, *Measurements of the pseudorapidity dependence of the total transverse energy in proton-proton collisions at $\sqrt{s} = 7$ TeV with ATLAS*, tech. rep., CERN, Geneva, Oct, 2011. <https://cdsweb.cern.ch/record/1394237>.
- [117] ATLAS Collaboration, G. Aad et al., *The ATLAS Simulation Infrastructure*, *Eur. Phys. J.* **C70** (2010) 823–874, [arXiv:1005.4568](https://arxiv.org/abs/1005.4568) [physics.ins-det].
- [118] GEANT4 Collaboration, S. Agostinelli et al., *GEANT4: A simulation toolkit*, *Nucl. Instrum. Meth.* **A506** (2003) 250–303.
- [119] F. Colecchia, J. Monk, E. Nurse, R. Prabhu, and P. Wijeratne, *Validation of the calorimeter energy response with $\pi^0 \rightarrow \gamma\gamma$ candidates*, tech. rep., CERN, Geneva, Aug, 2011. <https://cdsweb.cern.ch/record/1377411>.
- [120] I. Nedelkov, *Improper problems in computational physics*, *Computer Physics Communications* **4** (1972) no. 2, 157 – 164.
- [121] V. Blobel, *An Unfolding method for high-energy physics experiments*, [arXiv:hep-ex/0208022](https://arxiv.org/abs/hep-ex/0208022) [hep-ex].
- [122] I. Hadamard, *Lectures on Cauchy's problem in linear partial differential equations*. Oxford University Press, 1923. <http://archive.org/details/lecturesoncauchy00hadauoft>.
- [123] J. W. Monk and C. Oropeza-Barrera, *The HBOM Method for Unfolding Detector Effects*, *Nucl. Instrum. Meth.* **A701** (2013) 17–24, [arXiv:1111.4896](https://arxiv.org/abs/1111.4896) [hep-ex].
- [124] A. Hocker and V. Kartvelishvili, *SVD approach to data unfolding*, *Nucl. Instrum. Meth.* **A372** (1996) 469–481, [arXiv:hep-ph/9509307](https://arxiv.org/abs/hep-ph/9509307) [hep-ph].
- [125] G. D'Agostini, *A multidimensional unfolding method based on Bayes' theorem*, *Nucl. Instrum. Meth.* **A362** (1995) 487–498.
- [126] B. Wynne, *Imagiro: an implementation of Bayesian iterative unfolding for high energy physics*, [arXiv:1203.4981](https://arxiv.org/abs/1203.4981) [physics.data-an].
- [127] R. Brun and F. Rademakers, *ROOT: An object oriented data analysis framework*, *Nucl. Instrum. Meth.* **A389** (1997) 81–86.
- [128] T. Adye, *Unfolding algorithms and tests using RooUnfold*, [arXiv:1105.1160](https://arxiv.org/abs/1105.1160) [physics.data-an].

- [129] S. Wahrmund, *ATLAS leading jet underlying event plots*, private communication, 2012.
- [130] *Data-Quality Requirements and Event Cleaning for Jets and Missing Transverse Energy Reconstruction with the ATLAS Detector in Proton-Proton Collisions at a Center-of-Mass Energy of $\sqrt{s} = 7$ TeV*, tech. rep., CERN, Geneva, Jul, 2010.
<https://cdsweb.cern.ch/record/1277678>.
- [131] ATLAS Collaboration, G. Aad et al., *Jet energy measurement with the ATLAS detector in proton-proton collisions at $\sqrt{s} = 7$ TeV*, *Eur. Phys. J.* **C73** (2011) 2304, [arXiv:1112.6426 \[hep-ex\]](https://arxiv.org/abs/1112.6426).
- [132] *Jet energy resolution and selection efficiency relative to track jets from in-situ techniques with the ATLAS Detector Using Proton-Proton Collisions at a Center of Mass Energy $\sqrt{s} = 7$ TeV*, tech. rep., CERN, Geneva, Jul, 2010.
<https://cdsweb.cern.ch/record/1281311>.
- [133] D. Kar, *Combined ATLAS underlying event plots*, private communication, 2013.
- [134] O. Zenin, *Combined ATLAS underlying event plots*, private communication, 2013.
- [135] M. Whalley and J. Bentham, *The Durham HepData Project*.
<http://durpdg.dur.ac.uk/>.
- [136] A. Buckley, *The heptesis L^AT_EX class*, 2010.

List of Figures

2.1. Diagram of the LHC	5
2.2. Cross-section of an LHC dipole magnet	7
2.3. The CERN accelerator chain	8
2.4. The ATLAS detector	10
2.5. The ATLAS coordinate system	10
2.6. Illustration of the variation of pseudorapidity with angle	11
2.7. ATLAS muon tracker momentum resolution	14
2.8. The ATLAS calorimeters	15
2.9. Distribution of material within ATLAS	16
2.10. The ATLAS electromagnetic calorimeter	19
3.1. The ATLAS Inner Detector	24
3.2. Example geometry for measuring the momentum of a charged particle	25
3.3. The ATLAS Pixel detector	28
3.4. Pixel detector material distribution	29
3.5. Annealing of silicon detector defects	32
3.6. Diagram of an Inner Detector cooling loop	34
3.7. Cooling loop leak rates	36
3.8. Flowchart for beam conditions monitoring	38

3.9. Cross-section through the Inner Detector	39
3.10. Variation of track parameter resolution with p_T	41
4.1. Particle content of the Standard Model	43
4.2. Baryons with spin 1/2 and 3/2	45
4.3. Feynman diagrams of the vertices of QCD	45
4.4. The running coupling of QCD	46
4.5. LO and NLO diagrams for QCD quark scattering	48
4.6. Example behaviour from the MSTW PDF set	50
4.7. Evolution of parton cross-section with \sqrt{s}	51
4.8. Example performance of different jet algorithms	55
4.9. Evolution of cross-sections with \sqrt{s}	56
5.1. Example output from different MC generators and tunes	58
5.2. Breaking a colour string	65
5.3. Example of Lund string hadronisation	66
6.1. Topology of the underlying event	71
6.2. Efficiencies of the jet triggers	76
6.3. Track reconstruction efficiency	77
6.4. Dependence of track-vertex distance on $\sin\theta$	77
6.5. Correction to simulated topocluster energies	80
7.1. Example smearing matrix	83
7.2. Example of iterative unfolding convergence	89
7.3. Example of unfolding two correlated variables	95
7.4. Example of unfolding with different prior distributions	96

8.1. Charged particle $\sum p_T$ vs. p_T^{lead} profiles	99
8.2. Charged particle N_{ch} vs. p_T^{lead} profiles	100
8.3. Topocluster $\sum E_T$ vs. p_T^{lead} profiles ($ \eta < 4.8$)	102
8.4. Topocluster $\sum E_T$ vs. p_T^{lead} profiles ($ \eta < 2.5$)	103
8.5. Charged particle $\sum p_T$ and central topocluster $\sum E_T$ ratio vs. p_T^{lead} . . .	104
8.6. Charged particle $\langle p_T \rangle$ profiles	104
8.7. $\sum p_T$ vs p_T^{lead} profiles for inclusive and exclusive selections	106
8.8. $\sum p_T$ distribution for inclusive and exclusive selections	107
8.9. $\sum p_T$ distribution for three ranges of p_T^{lead}	108
8.10. Transition from inclusive to exclusive selection	109
8.11. Jet multiplicity and p_T distributions in the transverse region	110
8.12. $\sum p_T$ distributions in the trans-min region	111
8.13. The dependence of p_T^{lead} on data period	112
8.14. Charged particle $\sum p_T$ distribution for different data periods	114
8.15. The number of primary vertex tracks for a selected event	116
8.16. Merged pile-up vertex contribution	117
8.17. Comparing results from different UE analyses	121
8.18. Comparing uncorrected results from different UE analyses	123
A.1. MC closure for inclusive charged particle $\sum p_T$ vs. p_T^{lead}	128
A.2. MC closure for exclusive charged particle $\sum p_T$ vs. p_T^{lead}	129
A.3. MC closure for inclusive charged particle N_{ch} vs. p_T^{lead}	130
A.4. MC closure for exclusive charged particle N_{ch} vs. p_T^{lead}	131
A.5. MC closure for inclusive charged particle $\langle p_T \rangle$ vs. p_T^{lead}	132
A.6. MC closure for exclusive charged particle $\langle p_T \rangle$ vs. p_T^{lead}	133

A.7. MC closure for inclusive charged particle $\langle p_T \rangle$ vs. N_{ch}	134
A.8. MC closure for exclusive charged particle $\langle p_T \rangle$ vs. N_{ch}	135
A.9. MC closure for inclusive topocluster $\sum E_T$ for $ \eta < 4.8$ vs. p_T^{lead}	136
A.10. MC closure for exclusive topocluster $\sum E_T$ for $ \eta < 4.8$ vs. p_T^{lead}	137
A.11. MC closure for inclusive topocluster $\sum E_T$ for $ \eta < 2.5$ vs. p_T^{lead}	138
A.12. MC closure for exclusive topocluster $\sum E_T$ for $ \eta < 2.5$ vs. p_T^{lead}	139
B.1. Data stability for charged particle $\sum p_T$	141
B.2. Data stability for charged particle N_{ch}	142
B.3. Data stability for charged particle $\langle p_T \rangle$	143
B.4. Data stability for topocluster $\sum E_T$ distributions ($ \eta < 4.8$)	144
B.5. Data stability for topocluster $\sum E_T$ distributions ($ \eta < 2.5$)	145
C.1. Systematic error profiles for charged particle $\sum p_T$ vs. p_T^{lead}	147
C.2. Systematic error profiles for charged particle N_{ch} vs. p_T^{lead}	148
C.3. Systematic error profiles for topocluster $\sum E_T$ vs. p_T^{lead} ($ \eta < 4.8$)	149
C.4. Systematic error profiles for topocluster $\sum E_T$ vs. p_T^{lead} ($ \eta < 2.5$)	150
C.5. Systematic error profiles for charged particle $\langle p_T \rangle$ vs. p_T^{lead}	151
C.6. Systematic error profiles for charged particle $\langle p_T \rangle$ vs. N_{ch}	152
D.1. Charged particle $\sum p_T$ distributions	154
D.2. Charged particle N_{ch} distributions	155
D.3. Topocluster $\sum E_T$ distributions ($ \eta < 4.8$)	156
D.4. Topocluster $\sum E_T$ distributions ($ \eta < 2.5$)	157
D.5. Charged particle $\langle p_T \rangle$ distributions	158
D.6. Charged particle $\sum p_T$ distributions in restricted p_T^{lead} ranges	159

D.7. Charged particle N_{ch} distributions in restricted $p_{\text{T}}^{\text{lead}}$ ranges	160
D.8. Topocluster $\sum E_{\text{T}}$ distributions ($ \eta < 4.8$) in restricted $p_{\text{T}}^{\text{lead}}$ ranges	161
D.9. Topocluster $\sum E_{\text{T}}$ distributions ($ \eta < 2.5$) in restricted $p_{\text{T}}^{\text{lead}}$ ranges	162
D.10. Charged particle $\langle p_{\text{T}} \rangle$ distributions in restricted $p_{\text{T}}^{\text{lead}}$ ranges	163

List of Tables

2.1. ATLAS muon system technologies	12
2.2. ATLAS hadronic calorimeter resolution	17
2.3. ATLAS electromagnetic calorimeter resolution	19
3.1. Inner Detector combined resolution	40
6.1. Underlying event observables	70
6.2. The triggers used in the underlying event analysis	75
8.1. Table of systematic uncertainties for inclusive leading jet profiles vs. p_T^{lead}	119
8.2. Table of systematic uncertainties for exclusive di-jet profiles vs. p_T^{lead}	120
8.3. Table of systematic uncertainties for $\langle p_T \rangle$ vs N_{ch} profiles	120