


High jet multiplicity physics at the Large Hadron Collider

Mireia Crispín Ortuzar
of Balliol College, Oxford



Thesis submitted in partial fulfilment of the requirements
for the degree of Doctor of Philosophy
at the University of Oxford

Hilary Term, 2015

We are a way for the Cosmos to know itself

Declaration

This dissertation is the result of the author's own work, except in the few cases 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 page limit.

Mireia Crispín Ortuzar

High jet multiplicity physics at the Large Hadron Collider

Mireia Crispín Ortuzar

The Large Hadron Collider at CERN completed its first data-taking phase in 2013, after three years of remarkable performance. The high-energy proton-proton collisions recorded by the ATLAS experiment provide a gateway to the world of subatomic particles. This thesis presents two analyses of the full 8 TeV dataset taken by ATLAS, inspired by two of the major physics goals of the experiment. The first analysis is a search for new phenomena that could explain the nature of Dark Matter and solve the hierarchy problem. In particular, the search is optimised to look for heavy supersymmetric particles decaying to large numbers (7 to ≥ 10) of jets. The events are further classified according to the number of jets identified as originating from a b quark. No evidence is found for physics beyond the Standard Model, so the results are interpreted in terms of exclusion limits on various simplified supersymmetry-inspired models where gluinos are pair produced, as well as a mSUGRA/CMSSM model. The main background to the search is due to multi-jet production via the strong force. This motivates the second analysis presented in this thesis, which is a measurement of the cross section of four-jet events. The measurement is performed differentially in a series of variables which describe the kinematics and spatial configuration of the events. The results are compared to existing theoretical predictions.

Acknowledgements

First and foremost, I would like to thank Alan Barr, my supervisor, for guiding me through my DPhil. He is one of the most inspiring and brilliant scientists I know, an extremely supportive and attentive mentor, and the most efficient ATLAS member I have ever met.

My DPhil was funded by a combination of generous scholarships from the Caja Madrid Foundation, Balliol College (through a Foley-Béjar Scholarship), Telefonica and the British Spanish Society. I was also supported by the Particle Physics Sub-department of the University of Oxford, especially during my time in Geneva. I owe them a huge debt of gratitude.

I have had the pleasure of working with some great teams of physicists in the ATLAS Collaboration. I learnt a great deal from the wisdom and experience of Chris Young, Anna Sfyrlla, Zach Marshall and David Miller. Anna spent many hours with me in my first year, showing me the intricacies of the trigger system, and has remained an invaluable collaborator throughout my DPhil. I had the immense luck to work closely with Hernán Reisin on the SUSY multi-jet analysis, who made those long days of cutflow-checking actually joyful. Sabrina Sacerdoti followed on Hernán's steps in the Standard Model multi-jet measurement, responding to my hundreds of emails asking for more tests even in the middle of the holidays. I have also learnt a lot from Caterina Pizio and Paolo Francavilla, mostly about the details of the missing transverse momentum (and some very basic Italian). Daniel Maitre and Tuomas Hapola kindly provided additional theoretical predictions for the cross section measurement.

The Oxford SUSY group provided the perfect working environment. In particular, Andrée Robichaud-Véronneau and Alex Dafinca were my consecutive officemates, and both answered all my annoying questions and always made my time in the office an enjoyable one. With Will Kalderon I discovered the first evidence for direct spanda production, and James Scoville patiently explained some of the subtleties of particle theory.

The IT officers at Oxford, in particular Sean Brisbane, have been incredibly helpful: they have maintained and upgraded the systems, responded to every single request at any random time of the day or week, rescued lost files and helped run jobs on the cluster. This thesis could very well not exist without them.

Beyond my closest collaborators, I have been surrounded by a wonderful group of scientists both during my time in Oxford and in Geneva. My college tutor, Armin Reichold, has been a reference for me and a constant source of support. Kate and Nazim were my housemates for two years (and in two countries) and managed to turn our house into a home. The rest of the lovely Oxford group were always ready to discuss physics and have fun. My time in Geneva (and beyond) would not have been the same without Dani Cámpora and Tom Gillam. Both Dani and Tom filled my days with music, and there is not a piece of code I have which Tom has not saved from disaster at some point in time. At Balliol College I have shared countless musical projects and general fun with Hilary, Ben, Andrew, and many others.

I would not be here without the selfless, unconditional and loving support that my parents, Visi and Vicente, have given me and my brother Marcos throughout my life. I embarked on this DPhil journey with Tomeu, and we have shared all the challenges and small victories. He has encouraged me beyond words, listened to every talk I have given and read every paper I have written with unshakable patience and love. This thesis is dedicated to the four of them.

Preface

As a member of the ATLAS Collaboration I was involved in various projects, some of which are included in this thesis. The full list is included here in reverse chronological order, for completeness.

- **Measurements of four-jet differential cross sections from $\sqrt{s} = 8$ TeV proton-proton collisions using the ATLAS experiment.** I led most of the aspects of the analysis. See Chapter 4 of this thesis. The paper, of which I am contact editor, is currently going through the ATLAS approval process.
- **Limits on metastable gluinos from ATLAS SUSY searches at 8 TeV.** I produced all the results corresponding to the multi-jet analysis (one of the two searches included in the note). Published in 2014, Ref. [1].
- **Performance of E_T^{miss} at high luminosity.** I performed a new parametrisation of the E_T^{miss} in the high luminosity scenario, which has been used in the upgrade physics analyses since 2013. Part of the results were published in Ref. [2].
- **Search for new phenomena in final states with large jet multiplicities and missing transverse momentum at $\sqrt{s} = 8$ TeV proton-proton collisions using the ATLAS experiment.** I was responsible for the b -jet stream of the analysis, described in Chapter 3 of this thesis. I performed the optimisation of the analysis strategy, calculated the signal and background contributions and uncertainties, and processed the data. See Chapter 3 of this thesis. Published in 2013, Ref. [3].
- **Searches for supersymmetry at the high luminosity LHC with the ATLAS detector.** I produced all the results for the strong production section. Published in 2012, Ref. [4].

Contents

Introduction	1
1. Theoretical overview	3
1.1. The particle content of the Standard Model	4
1.2. Dynamics of the Standard Model	6
1.2.1. Gauge theories in the Standard Model	7
1.2.2. Renormalisation	9
1.2.3. The description of mass	10
1.3. Quantum chromodynamics	10
1.3.1. Structure of a hadronic event	11
1.3.2. The perturbative regime	12
1.3.3. Non-perturbative effects	16
1.4. Going beyond the Standard Model	18
1.4.1. Foundations of SUSY	19
1.4.2. Implications of SUSY	22
1.4.3. Experimental constraints	25
1.5. Conclusion and motivation	27
2. The Large Hadron Collider and the ATLAS experiment	29
2.1. The Large Hadron Collider	29
2.2. ATLAS detector overview	32
2.2.1. Coordinate system	33
2.2.2. Inner detector	34
2.2.3. Calorimetry	36
2.2.4. Muon spectrometer	38
2.3. Trigger	39
2.3.1. Jet triggers	40
2.4. Detector simulation	41

2.5. Object reconstruction	42
2.5.1. Tracks	42
2.5.2. Electrons	42
2.5.3. Muons	43
2.5.4. Jets	43
2.5.5. b -tagging	46
2.5.6. Missing transverse momentum	47
3. Search for new phenomena in events with large jet multiplicities	49
3.1. Introduction	49
3.1.1. Overview of the analysis	50
3.2. Analysis strategy	52
3.2.1. Signal and backgrounds	52
3.2.2. Dataset and trigger	57
3.2.3. Physics object definitions	59
3.2.4. Kinematic variable definitions	61
3.2.5. Event cleaning	62
3.2.6. Optimisation of the signal region definitions	63
3.3. Estimation of the Standard Model background	66
3.4. Estimation of the multi-jet background	69
3.4.1. Soft energy correction	74
3.4.2. The template method in events with b -jets	75
3.4.3. Sensitivity to pile-up effects	77
3.5. Systematic uncertainties of the multi-jet background prediction	79
3.5.1. Closure systematic uncertainty	79
3.5.2. Heavy-flavour systematic uncertainty	80
3.5.3. Soft energy and pile-up systematic uncertainties	82
3.5.4. Trigger systematic uncertainty	82
3.5.5. <i>Leptonic</i> background subtraction	87
3.6. Estimation of the leptonic backgrounds	87
3.6.1. Estimation of the leptonic $t\bar{t}$ and W +jets backgrounds	88
3.6.2. Estimation of other minor backgrounds	95
3.7. Systematic uncertainties on the leptonic backgrounds	95
3.7.1. Experimental systematic uncertainties	98
3.7.2. Theoretical systematic uncertainties	100

3.8. Results and interpretation	102
3.8.1. Fitting procedure	102
3.8.2. Fit results and interpretation	108
3.9. Summary and outlook	114
4. Measurement of the cross section of four-jet events	119
4.1. Introduction	119
4.1.1. Overview of the analysis	120
4.1.2. Useful concepts	121
4.2. Object definitions and event cleaning	122
4.2.1. Jets	122
4.2.2. Event cleaning and general cuts	123
4.3. Trigger	124
4.3.1. Single-jet triggers	125
4.3.2. Four-jet triggers	129
4.3.3. Combination	130
4.4. Variables of interest	132
4.4.1. Definitions and truth-level distributions	133
4.4.2. Truth- to reconstructed-level correlations	143
4.5. Binning	147
4.5.1. Statistical uncertainties and rounding	150
4.6. Unfolding of detector effects	152
4.6.1. Method description	152
4.6.2. Choosing the nominal Monte Carlo	155
4.6.3. Shape systematic uncertainty	157
4.6.4. Selection efficiency systematic uncertainty	161
4.6.5. Statistical and systematic uncertainties	163
4.6.6. Optimisation of the unfolding algorithm	167
4.7. Theoretical calculations and uncertainties	168
4.7.1. Leading-order predictions	169
4.7.2. Next-to-leading-order predictions	170
4.7.3. Other predictions	170
4.8. Results	171
4.9. Summary and outlook	190
5. Conclusions	191

A. Trigger efficiencies	193
B. Deriving variable-width binnings	195
Bibliography	203

Introduction

*‘Explain all that,’ said the Mock Turtle.
‘No, no! The adventures first,’ said the Gryphon
in an impatient tone: ‘explanations take such a dreadful time.’*

The goal of particle physics is to understand the nature and interactions of the most elemental constituents of the Universe, the fundamental particles. It is common for particle physics experiments to find out what things (particles) are made of by smashing them together and breaking them apart. As technology advances, particle accelerators and colliders reach higher and higher energies, which allows them to potentially produce increasingly heavier particles as part of the collision products. It is the particle physicists’ job to reconstruct the full history of the interactions and decays happening in a collision, starting from a single electronic ‘photograph’ of the final state. Particle detectors have allowed us to discover a plethora of particles, all of which fit nicely into the theoretical framework of the Standard Model. But there are a number of reasons why it is believed that this model is not complete, and that new particles must appear at the energy threshold of $\mathcal{O}(\text{TeV})$. The Large Hadron Collider (LHC) sits exactly on that energy frontier.

This thesis presents two analyses of the collision data measured by the ATLAS detector at the LHC. The two analyses may be seen as two sides of the same coin: both of them focus on the same type of final state particle configuration, but each has a different – though complementary – goal.

The events of interest for both analyses are those in which a large number of *jets* are produced. Jets are roughly conical sprays of particles, and they are omnipresent at the LHC. However, they present a number of challenges, both from the experimental perspective – as they are hard to reconstruct and calibrate – and from the theory point of view, because the complexity of the calculations grows rapidly with the number of jets in the final state. Why should we be interested in such a complicated scenario? The

main reason is that new, very heavy and strongly-interacting particles, of the sort that one would hope to be able to discover at the LHC, decay in cascades producing large numbers of jets. Moreover, these types of events also provide an excellent ground in which to test the subtleties of the theory of strong interactions.

The first study (Chapter 3) is a direct search for new physics phenomena beyond the Standard Model. The estimation of the multi-jet backgrounds will rely precisely on the mismeasurement of the energy of the jets, providing a way to describe some of the components of very high jet multiplicity environments without the need to resort to theoretical calculations. Although, alas, no new signals are found, strong limits are set on a variety of new physics models.

The second study (Chapter 4) is a measurement of the cross section of four-jet events, studied differentially in several variables that describe the dynamics of the event. The measurement provides discrimination between different state-of-the-art theoretical calculations.

Chapters 1 and 2 will set the basic theoretical and experimental concepts needed for the rest of the thesis.

1 | Theoretical overview

*‘If I had a world of my own, everything would be nonsense.
Nothing would be what it is, because everything would be what it isn’t.
And contrary wise, what is, it wouldn’t be.
And what it wouldn’t be, it would. You see?’*

Particle physics is a relatively young branch of science, although its founding question – ‘what is the Universe ultimately made of, and can we identify one or several fundamental building blocks?’ – has been around for at least several millenia. With the advent of more sophisticated experimental techniques, and in particular the invention of particle accelerators in the early twentieth century, the field grew dramatically. Today we are familiar with a whole zoo of particles, a handful of which are thought to be fundamental (at least at the energy ranges accessible to experiment). Although the experimental progress has been crucial to the development of particle physics, it has always also relied on the parallel development of theories that could explain the discoveries – and ideally predict what should come next.

This thesis will focus mostly on experimental aspects of particle physics, for which a basic understanding of the underlying theory is essential. This chapter will cover some of the most fundamental concepts that will be used later on in the experimental chapters. It will also try to highlight where the theory is in fact incomplete, as some of these points will be precisely the ones the rest of the thesis will try to shed light on. The chapter is divided into four sections, three of which will cover the Standard Model of particle physics, and one which will explore one of the possible extensions of the model, supersymmetry (SUSY), which could address some of its shortcomings. This division attends to the two main subjects of this thesis: a search for new phenomena using collision data from the LHC –interpreted in the context of SUSY–, and a measurement of the cross section of one of the main Standard Model (SM) background processes in that search.

An overview of the SM will be given in Sections 1.1 and 1.2, to set the particle content and the most important ideas underlying particle dynamics. Section 1.3 will then focus on the theory of strong interactions and how it is used to make predictions for the LHC – predictions which the measurement in Chapter 4 will probe. Strong interactions are also pivotal to the search for new phenomena that will be presented in Chapter 3. The theoretical background to the search will be complemented by the discussion on SUSY given in Section 1.4.

This Chapter borrows largely from the excellent textbooks corresponding to Refs. [5–8].

1.1. The particle content of the Standard Model

The SM is the theory that describes three of the four fundamental interactions of Nature –strong, electromagnetic and weak, the fourth one being gravity– and the properties of all the known fundamental particles. Before going into the details of the interactions, this section will describe the elements that are most widely used from an experimental perspective: the different particles, their charges and main interactions.

Within the SM, fundamental particles are assumed to be point-like and have an internal angular momentum quantum number called spin. According to whether the spin is an integer or a half-integer (in units of $\hbar$), particles are classified as bosons or fermions, respectively. Bosons are the mediators in the interactions between fermions, and they can also interact with other bosons.

Leptons				Quarks		
Particle		Mass	Charge	Particle	Mass	Charge
electron	e	0.511 MeV	−1	up	u	2.3 MeV
e neutrino	ν_e	$m_{\nu_e}^{(\text{eff})} < 2 \text{ eV}$	0	down	d	4.8 MeV
muon	μ	105.658 MeV	−1	charm	c	1.275 GeV
μ neutrino	ν_μ	$m_{\nu_\mu}^{(\text{eff})} < 0.19 \text{ MeV}$	0	strange	s	95 MeV
tau	τ	1776.82 MeV	−1	top	t	173.07 MeV
τ neutrino	ν_τ	$m_{\nu_\tau}^{(\text{eff})} < 18.2 \text{ MeV}$	0	bottom	b	4.18 GeV

Table 1.1.: The fermions in the SM, with masses taken from the Review of Particle Physics [9]. They all have spin $\frac{1}{2}$.

Particle	Mass	Charge	Spin
photon γ	$< 10^{-18}$ eV	0	1
$W^\pm$	80.385 GeV	± 1	1
Z	91.1876 GeV	0	1
gluon g	0 ^(*)	0	1
Higgs h	125.7 GeV	0	0

Table 1.2.: The bosons in the SM, with masses taken from the Review of Particle Physics [9].
(*) = Theoretical value.

The known particle content of the SM is given in Tables 1.1 (fermions) and 1.2 (bosons). Each fermion has an antimatter counterpart with identical mass but opposite charge. The antiparticle of the electron is usually known as ‘positron’. Fermions are further separated into two categories, according to whether they interact via the strong force or not: if they do, they are called ‘quarks’, if not, they are classified as ‘leptons’. Neutrinos are neutral leptons and are massless according to the SM, although experiments show that they have a small, non-zero mass [10–14]. Within the SM, they only interact via the weak force. Charged leptons, on the other hand, can also interact via the electromagnetic force, and quarks interact via all three forces of the SM. Three generations of quarks and leptons are known to exist, and they appear to be identical in all properties other than mass. In addition, there are three ‘colour’ versions of each quark.

In the family of bosons, photons are the carriers of the electromagnetic interaction, and they have no mass and null electric charge. The weak interaction is carried by the neutral Z bosons and the charged $W^\pm$ bosons. The strong force is mediated by gluons, of which there are eight types with different colour combinations. The Higgs boson is responsible for giving mass to the other particles, as will be discussed in Section 1.2.3.

There are 61 different particles in the SM: 6 leptons and 6 anti-leptons; 6 flavours of quark, each with 3 different colour versions, and their 18 corresponding anti-quarks; 8 gluons; 2 $W^\pm$, 1 Z and 1 Higgs h ¹. They have all been seen experimentally (alone or in bound states), the most recent one being the Higgs boson, discovered by the ATLAS and CMS collaborations in 2012 [15–17].

Some of the main interaction vertices of the SM are depicted in Figures 1.1 and 1.2. Figure 1.1 contains the fermion-boson interaction vertices, and Figure 1.2 the

¹Gluons and quarks are sometimes referred to collectively as *partons*.




Figure 1.1.: Three diagrams representing the SM interactions between fermions and bosons.




Figure 1.2.: Four diagrams representing the SM self-interactions between gauge bosons. V and V' in the second diagram correspond to any two vector bosons such that the total electric charge is conserved.

self-interaction vertices of gauge bosons. Higgs vertices and neutrino oscillations are not represented.

1.2. Dynamics of the Standard Model

The theories that make up the SM belong to the class of relativistic quantum field theories (QFT). In QFT, particles are seen as excitations of fields that permeate all space, and crucially, these fields are quantised, which naturally gives place to Fermi-Dirac and Bose-Einstein statistics, as appropriate. The Lagrangian formalism is used to derive the dynamics of the different particle fields.

Predictions in particle physics are typically made in the form of cross sections, which represent the probability for a particle reaction to occur. Transition probabilities are usually computed using perturbation theory, with the perturbation parameter being the coupling strength for the corresponding interaction. The transition amplitude is also referred to as the ‘matrix element’ (ME). Each term in the perturbative expansion can be interpreted graphically following the ‘Feynman rules’. The resulting pictures are also known as ‘Feynman diagrams’, and provide a simple way to compute the contributions to a process of interest. For example, Figure 1.3 shows some of the contributions to the lepton scattering process, including leading- and higher-order graphs. The loops in the two right diagrams of Figure 1.3 are characteristic of high-order terms, and will




Figure 1.3.: Three diagrams representing different contributions to the scattering process of two leptons. The first diagram on the left represents one of the contributions to the lowest order in perturbation theory. The two diagrams on the right represent higher order terms.

be mentioned again in the discussion of the divergences of the theory in Section 1.2.2. Diagrams without loops are also called ‘tree-level’ graphs.

The three types of interactions encompassed by the SM are treated on the same footing. The unifying feature is subtle: they are all ruled by the same category of symmetry principles. Symmetries in a theory² are connected with conserved physical quantities, as proved by Emmy Noether [18]. Some of the most widely known properties of particle interactions (e.g. the conservation of electric charge in electromagnetism) are indeed a direct consequence of these symmetry principles.

The SM arises from the requirement that all interactions be invariant under *local gauge transformations*, as well as Poincaré group transformations (translations, rotations and boosts). Electromagnetic and weak interactions are thus found to be manifestations of a single ‘electroweak’ gauge theory (EW), and the strong interactions are described by the gauge theory of quantum chromodynamics (QCD).

Experimentally, electromagnetic, weak and strong forces are vastly different. One of the reasons is that although formally they are based on equivalent principles, each of them satisfies a particular type of local gauge symmetry, which may also be ‘spontaneously’ broken. In addition, the sizes of the couplings are different. All of this has very important phenomenological effects that will be discussed in the next section.

1.2.1. Gauge theories in the Standard Model

A theory is said to have a gauge symmetry if its Lagrangian is invariant under a continuous group of local transformations. An example of an abelian gauge transformation is

$$\psi(x) \rightarrow \psi'(x) = e^{i\alpha(x)}\psi(x), \quad (1.1)$$

²Laws of physics are said to be symmetric when they remain invariant under a certain transformation.

where the phase $\alpha(x)$ depends on space and time in an arbitrary way. One could try to apply this transformation to the Dirac Lagrangian, which describes fermions:

$$\mathcal{L} = i\bar{\psi}\gamma^\mu\partial_\mu\psi - m\bar{\psi}\psi, \quad (1.2)$$

where ψ is the fermion field, m is the fermion mass, and γ is a Dirac matrix, which contains spin information.

However, it turns out that the Dirac Lagrangian is only invariant under local gauge transformations if the partial derivatives are replaced by covariant derivatives with an extra term with ad-hoc transformation properties,

$$\begin{aligned} D_\mu &= \partial_\mu - ieA_\mu, \\ A_\mu &\rightarrow A_\mu + \frac{1}{e}(\partial\alpha), \end{aligned} \quad (1.3)$$

where D_μ is the covariant derivative, and A_μ is a new field in the theory, which couples to the Dirac field $\psi(x)$ with strength equal to the electric charge $-e$. This new field, accompanied by an appropriate gauge invariant kinetic term, satisfies all the requirements to represent the force carrier of the electromagnetic interaction: the photon. It also follows that the photon must be massless, because a mass term for the photon field would break the local gauge invariance. These are the key elements for the interacting field theory of QED³.

This same idea can be generalised by noting that imposing gauge invariance on a free Lagrangian leads to the appearance of interacting gauge bosons. The fundamental difference between this example and the other more complicated gauge theories is that the symmetry transformations may involve more than one state or field at a time, instead of just one as is the case here. Consequently, the phase factors $\alpha(x)$ become $n \times n$ matrices, where n is the dimension of the group representation.

The complexity of the transformation is ruled by the underlying symmetry. QCD is ruled by a local SU(3) symmetry which describes the strong interactions between quarks. This symmetry is generated by the colour degree of freedom, so it is usually denoted SU(3)_c. The theory of electroweak interactions includes U(1) and SU(2) symmetries that together govern QED and the weak interactions of quarks and leptons. The SU(2) symmetry is generated by the ‘weak isospin’ degree of freedom. Overall, the symmetry of the EW theory is SU(2)_L × U(1)_Y, where U(1) is generated by the ‘weak hypercharge’

³Note that in the full electroweak theory, the photon is actually a mix of SU(2)_L and U(1)_Y states.

Y^4 ; and L refers to the fact that only the left-handed *chiral*⁵ parts of the quark and lepton fields enter into weak interactions. This property is directly connected to the breaking of the parity symmetry by weak interactions.

1.2.2. Renormalisation

It was said previously that transition probabilities were usually computed using perturbation theory. However, beyond leading order (LO) divergent integrals enter the calculation. There are two types of divergences: those which arise from the high-energy limit of the theory (ultraviolet, or UV), and those caused by the low-energy limit (infrared, or IR). Infrared divergences will be discussed in the context of strong interactions in Section 1.3.2, and this section will focus on UV divergences. The procedure of removing the UV-divergences of a theory to obtain a UV-finite result is called ‘renormalisation’. All theories in the SM have UV divergences; but they can be circumvented, as renormalisability is a property of local gauge theories.

Technically, the essence of renormalisation consists on absorbing the divergent terms in the definition of physical quantities like the coupling strength or the mass of the interacting particles. But the meaning of this process goes beyond a pure technicality, as the higher order corrections added to the *bare* coupling strength or mass via the renormalisation procedure are always present in nature. It is therefore only the renormalised values of these physical quantities that are accessible by experiment.

The second important consequence of renormalisation is that the (renormalised) coupling strength –as well as other physical observables– is not constant, but instead evolves with the momentum scale μ , introduced as an arbitrary parameter of the renormalisation procedure. This phenomenon is known as the ‘running’ of the coupling strengths. The observable predictions of the theory should not depend on the renormalisation scale μ if they are calculated at all orders; as this is often not possible, there might be a residual dependence on μ of the final result. For practical purposes, μ can be understood as a measure of the momentum transfer in the interaction, although the choice of the exact value is in fact arbitrary.

⁴The electric charge Q is related to the weak hypercharge Y and the weak isospin T_3 by the relation $Q = T_3 + \frac{Y}{2}$.

⁵Two chiral states are connected by a parity transformation. For massless particles, the chirality is equivalent to the physical helicity, which is defined as the projection of the spin onto the direction of the linear momentum. The left- and right-handed chiral states of a particle can be obtained by applying the projection operators $P_L^R = \frac{1 \pm \gamma_5}{2}$.

The QED coupling strength grows as μ increases (or the distance decreases). This can be understood as the charge screening of electrodynamics, in which as the distance increases (or the momentum transfer decreases) the effective charge seen by a electromagnetically-interacting probe gets smaller. However, the strong coupling does exactly the opposite. The striking consequences of this fact will be discussed in Section 1.3.

1.2.3. The description of mass

Section 1.2.1 showed that the gauge boson generated by the U(1) (electromagnetic) symmetry has to be massless. This is a generalisable statement; gauge invariance forbids mass terms also for the QCD and EW gauge bosons. Gluons are thought to be massless, but W and Z bosons have very large masses that have been measured with high precision (see Table 1.2).

Gauge bosons are given mass via the Higgs mechanism [19–22]. The idea is to introduce a new field in the theory whose potential has a continuum of degenerate minima. The field has to ‘choose’ one minimum out of all of these: this choice is said to break the symmetry ‘spontaneously’ (which is why the process is also known as spontaneous electroweak symmetry breaking (EWSB)). After EWSB, the Lagrangian of electroweak theory contains mass terms for the gauge bosons, as well as an additional massive scalar field called the Higgs boson.

The same Higgs field that explains the origin of the mass of gauge bosons can be used to explain the masses of fermions via the postulation of some additional gauge-invariant interaction terms. In the SM, all neutrinos are massless. However, several experiments have confirmed the existence of neutrino oscillations, which can only be explained if neutrinos have non-zero masses [10–14]. Recent studies found an upper bound on the total, combined mass of neutrinos of 0.32 eV [23].

1.3. Quantum chromodynamics

QCD is the theory that governs strong interactions between quarks and gluons, as well as their properties. Quarks and gluons carry the strong charge or *colour*, controlled by the $SU(3)_c$ symmetry, as was already mentioned in Section 1.2.1. One of the consequences

of this is that there are three types of colour charges, denoted as red, green and blue. Quarks carry colour, and antiquarks carry anti-colour; gluons (the gauge bosons of QCD) are not colour-neutral, but instead carry one of eight possible combinations of colour and anti-colour. This means that gluons, unlike photons, can self-interact.

The most crucial difference between QCD and the other gauge theories is the behaviour of its running coupling strength α_s . At high momentum scales, α_s tends to zero – the opposite of what happens with the electromagnetic coupling strength. Conversely, at low momentum scales α_s blows up. This is the origin of one of the most peculiar properties of QCD, namely *asymptotic freedom*, which refers to the fact that at short distances quarks and gluons are essentially free due to α_s being very small. The other important feature of QCD is known as *confinement*, and it corresponds to the low-momentum extreme. As a quark and an antiquark separate, the potential energy *increases*, which has the empirical consequence that quarks and gluons are never found in isolation. The observed particle spectrum of QCD consists solely of colourless bound states of quarks called hadrons. In high energy physics experiments, quarks and gluons are seen as approximately conical sprays of hadrons usually known as *jets*.

The theory of QCD was motivated by a series of symmetry observations of different hadronic states, and it has been verified in many fixed-target and collider experiments [24–26]. However, solving the QCD equations is challenging, and sometimes approximations are needed in order to make useful predictions to compare with experiments.

The next sections explore some of the difficulties that arise in the calculation of observables from QCD processes at high-energy colliders. They are relevant both for the new physics search described in Chapter 3, as multi-jet QCD production is one of the major backgrounds, and for the measurement of four-jet events described in Chapter 4, whose main aim is precisely to test theoretical QCD calculations.

For a more detailed discussion of the topics covered here, the reader should consult the review papers [27] and [28], upon which this Section is based.

1.3.1. Structure of a hadronic event

The structure of a proton–proton collision is complex. In order to simplify the simulations, the collisions are normally divided into several regimes characterised by the typical size of the momentum transfer. The process of separating such regimes is called ‘factorisation’. This section describes the basic structure of a hadronic event. The different parts are

illustrated in Figure 1.4, and they will be described in more detail in the following sections.

The large momentum transfer involved in the collision of the two protons results in the interaction of their constituents, the partons. This part of the event is usually called the *hard process* or *hard scatter*. The transition amplitude corresponding to the hard process can usually be calculated perturbatively. Protons are included in the calculation by means of a convolution of the partonic process with the *parton distribution functions*, which represent the probabilities of finding a parton within the proton, thus englobing the corresponding non-perturbative processes.

Partons and gluons present in both the initial and the final state of the hard process can in turn radiate more gluons, quarks and antiquarks, generating showers of outgoing partons. This process can be generated using a *parton shower algorithm* in which one parton is added to the final state at a time, using probabilistic methods. Parton showers involve an evolution in the momentum scale, therefore linking the high momentum transfer of the hard scatter with the low momentum scales related to the incoming protons and the outgoing hadrons.

Hadrons are formed when the momentum scale of the parton shower reaches a cut-off value (of the order of 1 GeV) where QCD becomes non-perturbative. The process of confining partons into colourless hadrons is called *hadronisation*, and it can be described via phenomenological models.

As the incoming protons are bound states of partons, it is possible that more than one pair of partons interacts in the collision. These additional, multiple interactions produce additional partons throughout the event and are usually said to constitute the *underlying event*.

1.3.2. The perturbative regime

QCD predictions for the LHC usually contain a combination of perturbative calculations and other approximations to deal with non-perturbative effects or difficult regions of phase space. At least part of the parton-level hard scattering calculation is usually performed perturbatively, and it may or may not be accompanied by a model of the subsequent parton shower. The following subsections will discuss the calculations performed at fixed order of perturbation theory, the parton shower simulations, and how the two approaches can be combined.




Figure 1.4.: Main stages of a proton–proton collision, illustrated step by step. The example hard process corresponds to $t\bar{t}$ production. Figures adapted from Ref. [29].

As in QED, higher-order loop corrections contain UV divergences that have to be removed by means of the renormalisation of the theory. The renormalisation results in the running of the coupling strength, as discussed in the previous section.

One of the most important features of QCD calculations is the appearance of logarithmic divergences when a parton is radiated with a small angle (‘collinear’) or small momentum (‘soft’). These are called ‘infrared’ divergences. The soft and collinear behaviour is ubiquitous in QCD, as will be seen next.

Fixed-order predictions

Fixed-order predictions contain the first few terms in the perturbative expansion in the strong coupling α_s of the cross section. An observable can only be calculated in fixed-order perturbation theory if it is insensitive to the emission of soft or collinear gluons. This is the motivation behind jets, whose definition is such that if an event is modified through soft and collinear emission their form will not change. Jets are discussed further in Section 2.5.4.

Fixed-order predictions can be obtained at leading order (LO) using only tree-level diagrams (that is, not including loops) if infrared-safe observables are used. There is a large number of tools which provide LO QCD calculations; some of them evaluate Feynman diagrams directly (like MADGRAPH [30]), and some use other techniques which allow them to be efficient at high parton multiplicities where the evaluation of all diagrams would be too costly (e.g. Alpgen [31] or SHERPA [32]). The phase-space integrals involved in the calculations are performed using Monte Carlo (MC) techniques, which is why these programs are usually called fixed-order Monte Carlo generators.

Infrared divergences must be treated with care when calculating total cross sections in higher-order perturbation theory. Due to unitarity constraints, the divergences from real emission terms have to cancel out with those from virtual loop correction terms. However, loop and tree-level diagrams are handled differently due to the need to regularise the divergences in the loop terms. In order to get a finite result, next-to-leading-order (NLO) generators have to use methods to reshuffle the divergences, called ‘subtraction’ methods. One of the historical reasons why NLO predictions were difficult to obtain was that the loop amplitude calculations had to be obtained semi-manually until recently. New techniques are now available that make possible the automatic calculation of loop amplitudes, for example the BLACKHAT library [33, 34] or NJET [35, 36], which are currently able to calculate the NLO cross section of QCD events with up to four or even five jets in the final state.

An example of the types of diagrams that intervene at various orders in α_s is shown in Figure 1.5. The example given is the production of Z +jets from two arbitrary incoming partons i, j , and it is taken from Ref. [27]. The columns represent the number of partons in the final state. The rows represent the number of loops in the diagrams. From the figure it is possible to read off the contributions needed for an $N^p\text{LO}$ calculation of $ij \rightarrow Z+n$ partons, by taking all entries with at least n partons up to order α_s^{n+p} . Then,




Figure 1.5.: Illustration of the contributions that are known for $ij \rightarrow Z + n$ partons, where i and j are arbitrary incoming partons, according to the number of outgoing partons, the number of loops and the number of powers of the coupling. An ‘ $\times$ ’ represents a squared tree-level diagram, an ‘ o ’ represents the interference of a 1-loop diagram with a tree-level diagram, and a ‘ $\emptyset$ ’ represents the interference of a two-loop diagram with a tree-level diagram or the square of a 1-loop diagram. Entries in black are known and used; entries in grey are known but have not been used. The entries in the shaded ellipses are those that are relevant for the NLO calculation of the cross section for the production of a Z-boson with a jet. Caption and figure taken from Ref. [27].

for example, for the calculation of $Z+1$ jet at NLO one would need the contributions circled in green, which include the interference of 1-loop and tree-level diagrams.

Parton shower programs

Fixed-order predictions can describe the momenta of the outgoing hard jets, but they may be insufficient to describe the substructure of jets or the distributions of other particles. The parton radiation process can be simulated to all orders through a *parton shower* algorithm.

The simulation of parton showers is based on the calculation of the probability for a parton to radiate a second parton, and it proceeds as a random process. This probability is computed in the soft and collinear limit, which reduces the calculation to the cases where the radiated partons are close to the original one or have low momentum. The approximation is justified by the fact that these configurations result in divergent cross sections, and therefore dominate the emission. The divergences themselves are cut off by introducing a resolution criterion for the emitted partons. The contributions from the different radiated partons are calculated successively in decreasing order of a particular

ordering variable q^2 , for example the virtuality $q^2 = p^2 - m^2$ of the parton, where p is its four-momentum. The procedure stops when a certain non-perturbative cutoff is reached.

Matching procedures

Parton-shower generators can describe many of the features of hadronic events, but as they rely on the soft and collinear approximation, they may sometimes fail at reproducing the large-angle radiation correctly or the multi-jet structure. Fixed-order programs, on the other hand, are able to calculate such configurations, but they are limited by the number of partons in the final state.

A compromise solution involves combining multi-leg matrix element calculations with parton showers, by letting a matrix element generator compute the hard process, and asking a parton shower program to evolve the final states. Care must be taken when doing this to avoid double counting of emissions; some of the existing ‘matching’ procedures are CKKW [37] or MLM [38] matching, both of which involve the choice of a matching scale between the hard and soft regimes controlled by the matrix element and the parton shower, respectively. The choice of matching scale must be such that the transition is smooth, but physics processes happening at scales significantly larger than the matching scale should not be affected by it in any case.

The MLM procedure consists on checking for each event that the jets formed after parton showering can be matched to one of the hard partons, and that there are no additional jets above a certain momentum threshold. If any of these requisites is not satisfied, the event is rejected. In MADGRAPH, a variant of MLM matching called ‘shower k_T ’ is available, which considers immediately after the first shower branching if the event should be rejected.

1.3.3. Non-perturbative effects

Despite the large center-of-mass energy of LHC collisions, there are some parts of an event in which the energies are in fact much lower. In QCD, this means that the coupling strength is much larger, such that the physics becomes non-perturbative. The initial state of the collisions, for example, involves protons, which are non-perturbative objects. The momentum transfers in the collisions are high, but as the final state quarks and gluons evolve and lose energy, hadrons form, entering again the realm of low-momentum, non-

perturbative physics. Non-perturbative effects are included in simulations via different mechanisms, including hadronisation algorithms, parton distribution functions and a description of the underlying event:

- **Hadronisation** must be understood in order to transform the final state partons obtained after shower evolution into observable hadrons. Monte Carlo event generators use phenomenological hadronisation models, as the transition between the two regimes is not well understood. The Lund string model is one of the widely used hadronisation algorithms [39]. It involves the presence of a colour ‘string’ across quarks and gluons, which breaks up into hadrons as the energy decreases or the distance increases. The ‘cluster’ model [40, 41] breaks the gluons up into $q\bar{q}$ pairs which are grouped into colour-neutral clusters. These clusters then decay, giving the hadrons. Hadronisation models include parameters that require data to be tuned. This includes the non-perturbative cutoff used in parton showers to mark the end of the perturbative regime.
- **Parton distribution functions** (PDFs) correspond to the number densities of the partons carrying a certain fraction of the momentum of a hadron, and they are particularly relevant to translate the initial-state partons used in perturbative calculations into the actual hadrons used in collisions. The choice of a ‘factorisation’ scale between the non-perturbative PDF and a hard part, as well as the different techniques used to fit the theoretical predictions to data, can have an important effect on the final result of the cross section calculations. Multiple PDFs exist to be used at hadron colliders [42–44].
- The **underlying event** refers to all the additional activity that does not come from the hard partonic scattering. It includes, for example, the additional interactions from the remnants of the incoming protons. The description depends on the momentum of the partons involved; at soft p_T , the perturbative QCD (pQCD) description must be complemented with soft models and tuned to data. The physics of the underlying event is complex and will not be discussed here. Ref. [45] may be consulted for a thorough discussion of the subject.

1.4. Going beyond the Standard Model

Although the SM explains successfully a wide range of phenomena, there are some experimental observations that remain unanswered and theoretical questions yet to be addressed. For example, the SM contains several free parameters whose values have no theoretical basis; it does not provide an explanation for the experimental evidence for neutrino masses or Dark Matter; and it is not known why there should be three generations of fermions or why the mass spectrum of particles is so vast. In addition, many physicists believe that all fundamental interactions of Nature should unify when described under the right grand unified theory (GUT).

No hints have been found so far that point in the direction of a particular extension of the SM. Many theories have been devised, but none so far has been experimentally verified – though a number have been severely constrained [46]. The next step in the search for a beyond the Standard Model (BSM) theory will come in the form of an experimental discovery, or indeed a refutation of many of the existing theoretical paradigms if no significant find is made.

One of the goals of this thesis is to search for new phenomena, and discover or set limits on a series of BSM models. In particular, it focuses on the phenomenology of the supersymmetry theory [47]. SUSY is based upon an extension of the symmetries of the SM. This alone is already a strong point in its favour, since (as was shown in Section 1.1) most of the dynamics of the SM arises precisely from symmetry principles. One of the original attempts to extend the SM was to find a big symmetry group that would englobe the symmetries of special relativity (or Poincaré symmetries, including translations, rotations and boosts) and the internal symmetries such as the gauge symmetries of the SM. These attempts were stopped by Coleman and Mandula’s 1967 theorem [48], which implied that such a thing could not be done. However, in 1975 an extension of the theorem was published [49] which considered the possibility of having generators relating particles of different spins. The new result showed that the most general symmetry of scattering amplitudes was the direct product of *super*-Poincaré and internal symmetries. The *super*-Poincaré symmetries now included, as well as the Poincaré symmetries, transformations of bosons into fermions and vice-versa – also known as *supersymmetric* transformations.

There are also a number of phenomenological reasons which make SUSY an attractive theory, including solving the hierarchy problem – or in other words, the crucial *cancellation of the divergent radiative corrections to the Higgs mass*. This will be discussed in Section

1.4.2, only after the foundations of the theory have been set in Section 1.4.1. No experimental signs of SUSY have been found yet, though it has been constrained from numerous direct and indirect searches. This will be the topic of Section 1.4.3.

This discussion on SUSY is based on the much more detailed reviews [50] and [51].

1.4.1. Foundations of SUSY

It has already been said that SUSY refers to the hypothetical symmetry between fermions and bosons. In particular, a SUSY transformation turns a bosonic state into a fermionic state, and vice-versa; and a supersymmetric extension of the SM would require doubling the number of particles in the spectrum. The SUSY operators Q and $Q^\dagger$ are constrained by the Haag-Lopuszansky-Sohnius extension [49] of the Coleman-Mandula theorem to satisfy the following commutation and anticommutation relations:

$$\begin{aligned}\{Q_\alpha, Q_{\dot{\alpha}}^\dagger\} &= -2\sigma_{\alpha\dot{\alpha}}^\mu P_\mu, \\ \{Q_\alpha, Q_\beta\} &= \{Q_{\dot{\alpha}}^\dagger, Q_{\dot{\beta}}^\dagger\} = 0, \\ [P^\mu, Q_\alpha] &= [P^\mu, Q_{\dot{\alpha}}^\dagger] = 0,\end{aligned}\tag{1.4}$$

where P^μ is defined as the four-momentum generator of spacetime translations, μ is a Lorentz index, and $\alpha, \dot{\alpha}$ are spinor indices. The algebra of SUSY is such that the single-particle states of a supersymmetric theory are organised in collections of fermion and boson states called *supermultiplets*. Within each supermultiplet, fermion and boson states are said to be *superpartners* of each other, and can be obtained by some combination of the Q and $Q^\dagger$ operators. Because Q and $Q^\dagger$ commute with P^μ (and consequently with the squared-mass operator $-P^2$), all the particles included in a supermultiplet have the same mass.

Each supermultiplet must contain the same number of fermionic and bosonic degrees of freedom. All the combinations of particles reduce to two simple cases: the combination of a two-component fermion and a complex scalar field, known as *chiral* or scalar supermultiplet, or the combination of a massless spin-1 vector boson and a massless spin-1/2 two-component fermion, known as *gauge* or vector supermultiplet. In a supersymmetric extension to the SM, each of the known fundamental particles is in one of these two types of multiplets, and must have a superpartner with spin differing by half a unit of spin.

The SM quarks and leptons form one component of a chiral supermultiplet, and their superpartners are spin-0 scalars. These scalars have an ‘s’ prepended to their name (*squarks* and *sleptons*) and they are denoted with a tilde, $\tilde{q}, \tilde{l}$. Gauge bosons, on the other hand, combine with spin-1/2 superpartners to form a gauge supermultiplet. In this case the names of the superpartners are formed by adding the suffix *-ino*, and they are globally referred to as *gauginos*. Finally, the Higgs boson is part of a chiral multiplet. There must be (at least) two such multiplets, one of which has the Yukawa couplings that give mass to up-type quarks, and the other one to down-type quarks.

All the concepts outlined above set the foundation for a class of supersymmetric models of varied complexity. The *minimal supersymmetric standard model* (MSSM) [52] is the simplest version of the theory. For every quark q there is a squark $\tilde{q}$, and for each lepton a slepton $\tilde{l}$. Gauge bosons (gluons, $W^\pm, W^0, B^0$)⁶ translate respectively into the gluino $\tilde{g}$, winos $\tilde{W}^\pm/\tilde{W}^0$ and bino $\tilde{B}^0$. The MSSM requires two Higgs doublets, H_u and H_d , whose superpartners are called higgsinos H . The neutral bino, wino and higgsino mix to form four neutral particles called neutralinos, and denoted $\tilde{\chi}_i^0$, with $i = 1 \dots 4$. Charged winos and higgsinos mix into two chargino states, $\tilde{\chi}_i^\pm$, with $i = 1, 2$.

R-parity

Some processes in the MSSM allow for proton decay, although experimental measurements set a lower bound on its lifetime of the order of 10^{33} years [53]. To avoid such processes, some supersymmetric models are required to conserve an ad-hoc symmetry called *R*-parity [54], defined for each particle as

$$R = (-1)^{3(B-L)+2s}, \quad (1.5)$$

where s is spin, B baryon number and L lepton number. Supersymmetric particles (or *sparticles*) are *R*-parity odd, whereas SM particles are *R*-parity even. The conservation of *R*-parity has important phenomenological consequences. Firstly, there cannot be any mixing between particles and sparticles. Furthermore, every vertex must contain an even number of sparticles. The consequences of these two facts are that the lightest particle with odd *R*-parity, called the ‘lightest supersymmetric particle’ (LSP), must be stable. The fact that no stable SUSY particle has been observed means that it should be neutral

⁶These are the electroweak gauge fields prior to electroweak symmetry breaking.

and weakly interacting. This is one of the crucial features of SUSY, as will be discussed in Section 1.4.2

In collider experiments, where the colliding particles are even states of R -parity, sparticles have to be produced in pairs, if R -parity is conserved. Each sparticle will in turn have to decay into a final state with an odd number of LSPs.

It is important to note that, though phenomenologically motivated, R -parity conservation is not required from the theoretical perspective. Supersymmetric models with R -parity violation (RPV) exist, and the results of the supersymmetric search presented in this thesis will also be interpreted within the context of one of them [55].

SUSY breaking

To date, no supersymmetric particle has ever been seen experimentally. If SUSY were a perfect symmetry of Nature, supersymmetric particles would have exactly the same mass as their SM partners and they would have already been seen experimentally. The fact that this has not happened indicates that SUSY is a broken symmetry.

The nature of SUSY breaking is unknown. However, one would want it to be such that the attractive properties of unbroken SUSY are retained – for example, ideally it should not re-introduce quadratic divergences to the corrections to the Higgs mass. Some restrictions on the SUSY-breaking terms must then be applied. It is common to consider the case of ‘soft’ SUSY breaking, in which divergences in Δm_H^2 are avoided by considering a SUSY-breaking Lagrangian containing only mass terms and coupling parameters with a positive mass dimension.

The MSSM lagrangian with soft SUSY breaking includes a total of 105 unknown masses, phases and mixing angles. However, the number of parameters can be reduced by applying constraints on the presence of flavour mixing and CP violating processes. These assumptions can bring the number of parameters down to 5, as in the model known as ‘minimal supergravity’ (mSUGRA) or ‘constrained MSSM’ (cMSSM). The five parameters mentioned above determine all the phenomenology of the model.

1.4.2. Implications of SUSY

So far, SUSY has been presented as a possible extension to the SM. There are a series of reasons why it has historically ranked amongst physicists' favourites. Three and a half important consequences of having a supersymmetric SM will be discussed here – some of them have been hinted at already. The first one is of a practical nature, the second one has cosmological implications, the third one is ruled by aesthetic principles, and the last half is an important motivation for BSM searchers.

The hierarchy problem

The hierarchy problem is the name given to the observation that the mass of the Higgs boson receives divergent quantum corrections from all the particles that couple to it through virtual loops. The corrections depend quadratically on an ultraviolet momentum cut-off, $\delta m_H^2 \sim \Lambda_{UV}^2$. If it is assumed that the SM is valid up to some value close to some GUT scale (for example, $\Lambda_{UV} \sim 10^{16}$ GeV), such corrections become gigantic. However, there is a strong reason why this cannot be the case: the Higgs mass has now been measured, and it has been found to be ~ 126 GeV [16].

There are different ways to address this problem from the theoretical perspective. It is technically possible to carefully cancel out these divergences, but having a cancellation precise to 32 orders of magnitude is widely regarded as unlikely. This solution is usually known as *fine tuning*. Another solution would be to assume that the momentum cutoff is actually smaller [56, 57]. A third possibility involves the addition of a new symmetry that naturally cancels out the divergent terms. Interestingly, SUSY is able to do precisely that.

The quadratically divergent terms in δm_H^2 change sign according to whether the particle the Higgs boson couples to is a boson or a fermion. This is the reason why unbroken SUSY, with its equal-mass supermultiplets, can naturally cancel the divergences out in all orders of perturbation theory. Even though SUSY has to be a broken symmetry, it is possible to concoct breaking mechanisms which maintain the cancellation of the quadratically divergent terms to all orders. These are the 'soft' SUSY breaking mechanisms mentioned in Section 1.4.1.

Dark Matter

One of the big unsolved problems of contemporary physics is the nature of dark matter. Dark matter is the name given to a substance which has been indirectly observed via its gravitational interaction only – hence the epithet ‘dark’, which refers to it being ‘non-luminous and non-absorbing’ [9]. Indirect astronomical measurements estimate dark matter to account for $\sim 24\%$ of the matter-energy content of the Universe [58], compared to $\sim 5\%$ of normal matter.

SUSY provides one of the most popular candidates for a stable, weakly-interacting massive particle (WIMP) which could explain dark matter. As discussed before, if R -parity is conserved then a neutral LSP would be stable. Supersymmetric models can be designed in such a way that the LSP has the right mass and interaction cross section to fit the cosmological constraints on dark matter. Models without R -parity conservation can also accommodate dark matter constraints; for example, a dark matter candidate could be a gravitino which, though unstable, would have a long lifetime due to the fact that it can only decay gravitationally [59, 60].

Unification of forces

If it is assumed that the Lagrangian of softly-broken SUSY has some underlying symmetry at some very high energy scale Q_0 , then the calculations of masses and cross sections at ordinary energy scales get very large contributions of the order $\log(Q_0/m_Z)$ from loop diagrams. These terms can be summed to all orders by treating the couplings and masses as running parameters and applying the appropriate renormalisation procedure, as discussed before for other gauge theories, for example in Section 1.2.2.

In the MSSM, the running gauge couplings can unify at a high scale $M_U \sim 2 \times 10^{16}$ GeV. It is not clear whether this is a happy coincidence or a sign of some underlying common principle, such as the ones proposed in grand unified theories. Figure 1.6 shows the running gauge couplings in the MSSM and the SM. The unification of the MSSM couplings is not perfect, but the small differences can be attributed to new particles that may appear near M_U .




Figure 1.6.: Evolution of the inverse running gauge couplings $\alpha_a^{-1}(Q)$, including two-loop effects, for the SM (dashed lines) and the MSSM (solid lines). The MSSM band (defined by the blue and red lines) is obtained by treating the particle masses as a common threshold varied between 500 GeV and 1.5 TeV. $\alpha_3(m_Z)$ is varied between 0.117 and 0.121. Taken from Ref. [50].

A great model-building machine

SUSY is not a model: it is a symmetry, a principle, a framework. Softly broken SUSY has $\mathcal{O}(100)$ parameters, which makes it computationally expensive to scan. Various SUSY models fix certain relations between parameters, based on experimental considerations, thereby reducing the number of degrees of freedom. The *phenomenological* MSSM (pMSSM), for example, has only 19 free parameters [61], and the *constrained* MSSM (CMSSM) has only 5 [62]. It is common to optimise and limit the interpretation of the searches to these more constrained (although well motivated) scenarios, which means that a large fraction of the supersymmetric phase space is yet to be explored [63]. Indeed, the phase space of SUSY is so vast that it can be seen as an almost endless source of BSM-like processes and topologies, many of which are common to other perhaps more limited models. For example, multi-jet cascade decays are one of the foremost features of the decays of very heavy supersymmetric particles – and also the main target of the search presented in this thesis. However, multi-jet final states are also characteristic of other BSM phenomena, for example, microscopic black holes [64].

In recent years, many searches for new physics at hadron colliders have used the so-called ‘simplified models’. These models are used to optimise the analysis and enhance

its sensitivity to the production of a certain particle or final state. In a simplified model, the masses of most SUSY particles are set at a scale of a few TeV, and only the LSP and a few particles of interest are left in the kinematically accessible mass range. In addition, branching fractions for the decays of interest are usually set to 100%. This framework, as well as simplifying the optimisation procedure, is useful to extrapolate exclusion limits to more complex models. For example, one could take a simplified model that only includes the production of sparticle A decaying 100% of the time to B , and compare it against a UV-complete SUSY model which also includes that process. The simplified model would be providing a best case scenario in terms of the exclusion limits on A ; this means that the exclusion limits of the second, more complex model B should be in principle weaker than those for the simplified model A . More sophisticated ways to apply simplified model results to general models exist [65].

The search presented in this thesis uses several simplified models to interpret the results obtained. More information about them will be given in Section 3.2.1.

1.4.3. Experimental constraints

SUSY can be constrained from many different angles. Here three types of constraints will be discussed: theoretical bounds, direct experimental constraints and indirect experimental constraints.

The main theoretical constraint on SUSY is related to the level to which the theory is fine tuned, or its degree of *naturalness*. As the lower bounds on the masses of sparticles are raised experimentally, it becomes increasingly harder to solve the hierarchy problem in a *natural* way. However, there is no consensus on the definition of naturalness, and the hierarchy problem itself depends on the value of Λ_{UV} , whose exact value is unknown. Moreover, the resulting upper bounds on SUSY masses vary significantly between the different sparticles. The problem of naturalness has been widely discussed [51], concluding in general that the current experimental limits are not yet strong enough to discard natural SUSY.

The experimental constraints come from both direct and indirect measurements. In the category of direct measurements, the limits in the majority of SUSY scenarios are ruled by the LHC (null) search results. They are always given with respect to a particular model assumption. The tightest bounds are set by the searches for strongly-interacting particles like squarks and gluinos, which have the largest cross sections for a fixed sparticle

mass. However, the bounds depend on the decay mechanisms. For example, gluinos with masses below 1.33 TeV are excluded in a simplified model where they decay to light squarks and an LSP [66]. However, naturalness constraints are more stringent for the heavy-flavour squarks – and the mass limits on direct production of gluinos decaying to stop quarks via $\tilde{g} \rightarrow t\bar{t}\tilde{\chi}_1^0$ are not so high. The SUSY search presented in this thesis is particularly sensitive to this channel. A previous search for new phenomena in events with high jet multiplicity – upon which the analysis presented in this thesis is based – was able to exclude gluinos decaying via $\tilde{g} \rightarrow t\bar{t}\tilde{\chi}_1^0$ up to $m_{\tilde{g}} = 1$ TeV [67]. These limits are significantly increased in this thesis, as will be shown in Chapter 3 on page 49.

Although most of the models studied in the context of this thesis conserve R -parity, one RPV model is also considered and a competitive limit is obtained, as will be seen in Section 3.8.

SUSY is also probed by indirect measurements. For example, some SUSY models are challenged by the discovery of the Higgs boson with $m_H \sim 125$ GeV. In the MSSM, for example, the Higgs mass is generically light. This value can be raised significantly by loop corrections, which can be tuned by modifying some of the MSSM parameters, increasing the amount of mixing between left-handed and right-handed stops, and/or increasing the stop mass, in tension with naturalness constraints [51]. The mSUGRA sample used in this thesis – which, as will be seen in Chapter 3, is severely constrained by the SUSY search presented here – satisfies the Higgs constraints. It is possible to avoid these problems in supersymmetric models beyond the MSSM [68, 69].

Some of the parameters introduced with SUSY breaking mechanisms imply flavour mixing or CP violating processes. In general, flavour mixing terms violate low energy constraints on flavour-changing neutral currents such as $\ell_i \rightarrow \ell_j \gamma$, where ℓ_i, ℓ_j are any two different leptons [70, 71]. Further to this, the CP-violating phases are in confrontation with the electric dipole moments of the electron and muon. There are several options to circumvent these problems, including making symmetry-based assumptions on the form of the guilty terms and phases; they are reviewed in Ref. [50]. These constraints motivate the pMSSM and mSUGRA/CMSSM scenarios mentioned above.

1.5. Conclusion and motivation

In this section the key ideas of the SM have been described, focusing on the theory of strong interactions and the limitations that exist in the calculation of strong processes. Then, the shortcomings of the SM as a whole have been outlined, describing one of the possible theoretical extensions, SUSY.

The LHC is the ideal facility to test these two crucial aspects of the understanding of fundamental physics. Firstly, as it is a hadronic collider, hadronic processes are dominant. The high energy and high luminosity of the collisions have produced billions of high-multiplicity hadronic events in which to test QCD. This thesis presents the first measurement of the cross section of four-jet events produced in 8 TeV collisions, in Chapter 4. The measurement is differential in a range of observables that describe the dynamics of the events. The goal is to compare multiple generators to assess where the different methods introduced in Section 1.3, including the latest developments in the field, perform better or worse.

A search for new physics in high jet multiplicity events is also presented in Chapter 3. The search is highly sensitive to gluino production, and sets limits on the masses of gluinos, neutralinos and stops, amongst others. Moreover, high jet multiplicity events are common to a wide variety of BSM models.

2 | The Large Hadron Collider and the ATLAS experiment

*‘Curiouser and curiouser!’ cried Alice (she was so much surprised, that for the moment she quite forgot how to speak good English);
‘now I’m opening out like the largest telescope that ever was!
Good-bye, feet!’*

This thesis uses data taken by the ATLAS (“A Toroidal LHC ApparatuS”) Experiment [72]. It is one of the experiments at the Large Hadron Collider (LHC) at the *Centre Européen pour la Recherche Nucléaire*, CERN.

The LHC is the highest-energy particle accelerator in the world, and delivers of the order of hundreds of millions of collisions per second. These two ingredients are essential for the physics goals of ATLAS. The first part of this chapter introduces the main features that give the LHC its unprecedented power.

ATLAS is a general-purpose detector composed of a series of sub-detectors, each of which is optimised to perform a different type of measurement. The second part of this chapter will focus on the different sub-components of ATLAS, and on the algorithms used to reconstruct and identify the physics objects used in the analyses that will be presented in this thesis.

2.1. The Large Hadron Collider

The LHC [73] is a particle collider situated at CERN in Geneva, Switzerland. Between 2010 and 2012, it collided proton beams at centre-of-mass energies $\sqrt{s} = 7$ TeV (in 2010

and 2011) and 8 TeV (in 2012) at four interaction points. It entered a shutdown phase in 2013 and will restart in 2015 at 13 TeV of centre-of-mass energy.

The LHC is a synchrotron 27 km in circumference, housed in the tunnel originally built for the Large Electron-Positron Collider (LEP) between 1984 and 1989. The tunnel crosses the French-Swiss border and lies at an average depth of 100 m underground. Along the circumference are 1232 superconducting dipole magnets kept at 1.9 K which bend the trajectory of the protons.

The LHC has been designed to collide protons and lead nuclei¹. All the data used in this thesis come from proton-proton collisions. Protons are produced by ionising hydrogen inside a *duoplasmatron*, which generates a plasma of electrons, protons and molecular ions. Protons are then extracted and injected into the first of a series of accelerators in which their energy increases progressively. A schematic of the CERN accelerator complex is shown in Figure 2.1. Protons reach an energy of 450 GeV before being injected in the LHC in two counter-circulating beams. In 2010 and 2011, each beam was accelerated to up to 3.5 TeV, and in 2012 they reached 4 TeV. The protons are arranged in up to 2808 bunches per beam with a nominal bunch spacing of 25 ns, although for most of 2011 and 2012 the bunch spacing was 50 ns.

The proton beams collide at four interaction points distributed around the ring. The experiments are situated at each of these points to detect and measure the particles produced in the collisions. The ATLAS experiment is one of the two general-purpose detectors, the other one being the Compact Muon Solenoid (CMS) [74]. The main physics goals of the ATLAS collaboration are the study of new physics phenomena at the TeV scale, the discovery of the Higgs boson, and the performance of precision measurements of the Standard Model of particle physics. It is a collaboration of roughly 3000 physicists from all over the world. The other two large experiments installed around the collision points are LHCb, designed to perform heavy-flavour physics measurements, and ALICE, designed for lead-lead collisions.

After the energy, the second most important figure of merit of the LHC is the luminosity. The luminosity $\mathcal{L}$ connects the rate of events $R = \frac{dN}{dt}$ to the cross section σ of a given process:

$$R = \mathcal{L}\sigma. \quad (2.1)$$

¹The centre-of-mass energy in this case is 2.76 TeV per nucleon.




Figure 2.1.: The accelerator complex at CERN (not to scale), including the LHC, its pre-accelerators, and other active facilities.

The luminosity can be understood as the quantity that measures the ability of the collider to produce the required number of interactions of a particular process. The physics goals of the LHC include searches for rare processes, and therefore require high luminosities. The peak instantaneous luminosity reached in 2012 was $\sim 7 \times 10^{33} \text{ cm}^{-2} \text{ s}^{-1}$. The design luminosity is $\sim 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$, which translates into ~ 22 simultaneous proton-proton interactions and $\mathcal{O}(1000)$ particles being produced in the central detector region in every bunch crossing [75, 76]. The presence of additional proton-proton interactions in a single measured event is called *pile-up*, and requires careful studies from physics analyses. Pile-up can be *in-time*, if the additional interactions came from the same bunch crossing, or *out-of-time*, if the electronic signals from different collisions overlapped. These conditions are challenging from the perspective of the design of the detectors, which must withstand very high levels of radiation. The properties of the ATLAS detectors will be described in Section 2.2.

Integrating the ‘instantaneous’ luminosity defined above over time gives the ‘integrated luminosity’, which is directly related to the total number of observed events, and is therefore commonly used to express the size of a dataset. A total integrated luminosity of 5.61 fb^{-1} (7 TeV data) was delivered in 2011, and 23.3 fb^{-1} (of 8 TeV data) in 2012 [77]. A fast selection system is needed to select the small fraction of the data that can be physically recorded and analysed. This system is known as the *trigger*, and will be presented in Section 2.3.

2.2. ATLAS detector overview

ATLAS has three main subdetectors designed to identify and measure the properties of different types of particles. Its general layout is shown in Figure 2.2. It follows the conventional layered design, with three main types of sub-detectors (particle tracking detectors, calorimeters and muon detectors), and is nominally forward-backward symmetric with respect to the interaction point. The presence of magnetic fields permits the measurement of the momenta of charged particles via the curvature of their trajectories; this is the purpose of the solenoid magnet surrounding the tracker and of the toroid magnets interleaved with the muon detectors. All of them are superconducting and cooled by liquid helium to a temperature of 4.5 K.

The **inner detector** (ID) system is located around the interaction point. It consists of a high-resolution pixel detector nearest the beamline, followed by the strip-based




Figure 2.2.: Layout and sub-detectors of the ATLAS detector. The detector is 44 m long and 25 m high, and weighs approximately 7 tonnes. From [78].

semiconductor tracker (SCT) and the transition radiation tracker (TRT). Its main goals are locating the interaction vertex, measuring displaced vertices from long-lived particles and measuring the track momentum with high precision. More details on the ID are given in Section 2.2.2. The **calorimeter system** measures the energy and position of electrons, taus, photons and hadrons. There is an electromagnetic calorimeter and a hadronic calorimeter, both of which are described in Section 2.2.3. The **muon system** is responsible for triggering on and measuring the muon momentum. High precision chambers are used for precise measurements of muons, and coarser chambers are used for triggering muon events. The muon system is described in Section 2.2.4.

The layout and sub-detectors of ATLAS are depicted in Figure 2.2.

2.2.1. Coordinate system

The coordinate axes used by ATLAS form a right-handed system in which the x -axis points to the centre of the LHC ring, the z -axis follows the beam direction and the y -axis points upwards. The azimuthal angle ϕ is measured around the beampipe, on the plane perpendicular to it. $\phi = 0$ corresponds to the positive x -axis and ϕ increases clockwise




Figure 2.3.: Layout and sub-components of the ATLAS Inner Detector. From [78].

looking into the positive z direction. The polar angle θ is the angle from the z beam axis. It is more common to use instead the pseudorapidity η , defined as $\eta = -\ln \tan(\theta/2)$. The pseudorapidity is an approximation of the physical rapidity Y in the limit when the particle is travelling close to the speed of light. The rapidity is defined as

$$Y = \frac{1}{2} \frac{E + p_z c}{E - p_z c}. \quad (2.2)$$

The pseudorapidity is a common choice of coordinate because the particle flux from the interaction point is approximately constant as a function of η . The pseudorapidity is 0 anywhere in the $z = 0$ plane and tends towards $\pm\infty$ in the forward and backward directions, respectively. The ID provides coverage up to $\eta = \pm 2.5$ (approximately 0.05π rad away from the beampipe) and the calorimeters up to around $\eta = \pm 5$ (approximately 0.004π rad away from the beampipe).

The detector is generally divided in three regions in η : barrel, endcap and forward. The exact η boundaries vary according to the context, but they are typically around $\eta \sim 1.5$ and $\eta \sim 3.2$.

2.2.2. Inner detector

The inner detector (ID) reconstructs the trajectories, or *tracks*, of electrically charged particles. Tracks can be used for particle identification, to measure particle momentum

(in combination with the solenoid magnet), to reconstruct the vertex from which a set of particles originate, and following from this, to measure the distance of a secondary set of tracks from the primary interaction vertex. This last measurement makes it possible to identify jets coming from the decay of a b quark, as most b -hadrons have lifetimes of ~ 1 ps, which implies decay lengths of $\mathcal{O}(\text{mm})$ [9]. Other particles that can be identified via this mechanism include c quarks or τ leptons.

An overview of the ID is shown in Figure 2.3. Two precision tracking detectors, pixel and SCT, cover the region $|\eta| < 2.5$. The principle of operation relies on the production of free charge carriers in silicon by the incoming particles, with carriers travelling to the electrodes under the influence of an electric field. The energy of the ionising particle can be inferred from the number of electron-hole pairs detected. Semiconductor trackers have very high time and energy resolution, and are small in size compared to gaseous detectors.

The pixel detector consists of a barrel made of concentric cylinders around the beam axis, and two end-cap regions with three disks each. The barrel provides the highest granularity of the ID, with a minimum pixel size of $50 \times 400 \mu\text{m}^2$. The pixel detector is capable of providing very accurate measurements of the impact parameter of tracks, thus helping detect particles with non-negligible lifetimes such as b -hadrons or τ leptons. The detecting units of the SCT are strips, instead of pixels. They are organised in a similar manner to the pixel detector, with four cylindrical layers in the barrel and nine disks in the two end-cap regions. It uses two sets of strips with a relative rotation of 40 mrad to measure all coordinates. The mean pitch of the strips is approximately $80 \mu\text{m}$. The SCT plays an important role in the momentum measurement, as well as the impact parameter and vertex position.

The TRT was designed to extend the tracking measurements to larger radii, as well as to provide electron identification via the detection of transition radiation. The tracking measurement is provided in the $r - \phi$ plane by straw drift tubes filled by a Xe-based gas mixture, which function as ionisation chambers in the proportional regime. The material surrounding the straws consists of multiple layers with different dielectric constants, which triggers the emission of transition radiation by charged particles traversing the medium. The intensity of the radiation is proportional to the Lorentz factor of the incoming particle, which permits the discrimination of ultra-relativistic electrons from other heavier hadronic particles. The straw tubes are situated parallel to the beam axis in the barrel region, and radially in wheels in the end-caps, reaching up to $|\eta| = 2.0$. It




Figure 2.4.: Layout of the ATLAS calorimeter system. From [78].

has a lower resolution per point compared to the silicon detectors, but provides larger numbers of measurements, and the measured track lengths are longer.

2.2.3. Calorimetry

The ATLAS calorimeters provide a measurement of the energy and momentum of particles, both electrically charged and neutral ones. A layout of the system is shown in Figure 2.4. The calorimeters cover up to $|\eta| = 4.9$ and combine different techniques to cover the characteristics of the different physics processes of interest. The distribution of the different calorimeters in ATLAS is as follows. The first layer of calorimetry is formed, in order of increasing rapidity, by the electromagnetic barrel calorimeter ($|\eta| < 1.475$), the electromagnetic end-cap calorimeters (EMEC, covering $1.375 < |\eta| < 3.2$), the hadronic end-cap calorimeters (HEC, $1.5 < |\eta| < 3.2$), and the forward calorimeters (FCal, covering the region $3.1 < |\eta| < 4.9$). The outer layer is formed by the hadronic tile calorimeters, with one central ($|\eta| < 1.0$) and two extended ($0.8 < |\eta| < 1.7$) barrels.

The energy measurement in calorimeters is based on the interaction of the incoming particle with the detector, producing a shower of daughter particles whose energy is deposited in the material, collected and measured. The segmentation of the calorimeter allows one also to obtain information about the direction of the particles and the shape

of the shower, which helps with particle identification. In addition, the response from the calorimeters is fast, so they are widely used for triggering.

The ATLAS calorimeters are sampling detectors, which means that the active material providing the signal is different from the medium responsible for absorbing the particle energy.

The electromagnetic calorimeters use liquid argon (LAr) as the active material. LAr is known to be intrinsically radiation hard, as well as have a linear and stable response over time. In addition, they use lead as an absorbing material, with readout electrodes situated in the gaps between the absorbers. They are distributed in an accordion shape, which provides full coverage in ϕ and a fast extraction of the signal. In the central region, where precision measurements are required for electrons and photons, the first layer is finely segmented in η to improve the position determination. The electrodes are etched to obtain projective² segments in η and in depth, whereas the segmentation in ϕ is achieved by grouping the signal from the appropriate electrodes. The smallest strips on the first layer are $\Delta\eta \times \Delta\phi \sim 0.003 \times 0.1$ in size, while the two following layers, organised in ‘towers’ pointing to the interaction point, have a size of $\Delta\eta \times \Delta\phi \sim 0.025 \times 0.025$ or larger.

The hadronic tile calorimeters use scintillator as active material, and steel as the absorber medium. The barrel hadronic calorimeter is divided into three sections (central and extended barrels), composed by wedges of size $\Delta\phi \sim 0.1$ and made of alternating, radially-oriented scintillator tiles and steel plates. The readout fibres are grouped into photomultiplier tubes, providing a projective segmentation in η . To minimise the loss of energy in the gaps between the central and extended barrels, those regions have special steel-scintillator modules. The hadronic end-cap calorimeters suffer harsher radiation conditions, and are therefore based on copper/LAr technology. They consist of two wheels, with $\Delta\eta \times \Delta\phi = 0.1 \times 0.1$ readout cells in the region $|\eta| < 2.5$ and $\Delta\eta \times \Delta\phi = 0.2 \times 0.2$ for the rest. The FCal is split into three modules. The first one is an electromagnetic module with copper as absorber material, and the other two are hadronic modules using tungsten as the absorbing material.




Figure 2.5.: Layout of the ATLAS muon system. From [78].

2.2.4. Muon spectrometer

The muon spectrometer is designed to measure the momentum of particles that escape the calorimeters in the region $|\eta| < 2.7$, and also to trigger on these particles in the region $|\eta| < 2.4$. This is achieved by means of different types of muon chambers optimised for trigger and high-precision measurements. A layout of the muon system is given in Figure 2.5.

The momentum measurement relies on the bending of the particle trajectory under the effect of the superconducting toroid magnets. The magnetic field is provided by different magnets across $|\eta|$, including the large barrel toroid in the central region, the end-cap magnets inserted into the ends of the barrel toroid, and a combination of both in the intermediate $1.4 < |\eta| < 1.6$ region. The resulting field is mostly orthogonal to the muon trajectories.

The muon chambers form three cylindrical shells around the beam axis in the barrel region, with the third layer being 10 m away from the interaction point. In the transition and end-cap regions, the muon chambers form three wheels perpendicular to the beam. There is a gap at $\eta \sim 0$ to service all the other inner sub-detectors. The momentum of

²‘Projective’ means in this context that the size of the elements grows proportionally to the distance from the interaction point.

muons is measured in the range $|\eta| < 2.7$ with high precision by Monitored Drift Tube chambers (MDTs). Drift chambers are a type of proportional counters which also include the precise timing measurements of the generated pulses, which provides an accurate determination of the position of the particle. The overall layout is projective. In the forward region, due to the higher muon flux, the inner layer is made of Cathode-Strip Chambers (CSCs). CSCs are multiwire proportional chambers in which the information on the position of the particle is obtained from the orthogonally segmented cathodes. CSCs provide a more robust, high-rate-safe measurement.

Complementing the precision-tracking chambers are the trigger chambers, which deliver tracking information within a few tens of nanoseconds. Resistive Plate Chambers (RPCs) cover the barrel region ($|\eta| < 1.05$), structured in three concentric cylindrical layers, and Thin Gap Chambers (TGCs) cover the end-caps ($1.05 < |\eta| < 2.4$). They measure both the η and the ϕ coordinates of the muon tracks. The trigger chambers also provide bunch-crossing identification and complement the measurement of MDTs and CSCs, which can only measure in the bending direction.

2.3. Trigger

The trigger system selects events passing some set of kinematic cuts, typically the transverse momentum of an object, the missing transverse momentum, or some slightly more complex variable. The challenge of the trigger is to reject background without biasing the signal selection, in a short time and at a high rate.

The system is subdivided into three levels. Level 1 (L1) is a hardware-based trigger, which uses coarse detector information from the calorimeter and the muon subsystems (RPCs and TGCs), and produces an answer in less than $2.5 \mu\text{s}$. At L1, the objects triggered on include high- p_T muons, electrons/photons (which cannot be distinguished at this stage), jets and hadronic τ leptons, and missing transverse momentum. The L1 calorimeter trigger works with approximately 7000 calorimeter trigger towers of $\Delta\eta \times \Delta\phi \sim 0.1 \times 0.1$ in most cases, or larger at high pseudorapidities. Geometric information about the objects is kept by the muon and calorimeter trigger processors, and sent to the next level in the form of ‘regions of interest’ (RoIs) if the event passes the L1 requirements. The input event rate at this stage is 20 MHz (for 50 ns bunch spacing), and the L1 trigger reduces it to a nominal value of 70 kHz.

After being selected by the L1 trigger, events are passed on to the software-based high level trigger (HLT), which has two stages: the Level 2 (L2) trigger and the Event Filter (EF), both of which can already use the full granularity of the data. The L2 trigger increases the level of detail available for the decision by looking at the RoIs defined by the L1. It uses information on coordinates, energy and type of signatures only, in order to minimise the data transfer. It has an average latency of 40 ms, and reduces the event rate to 6 kHz. Finally, successful events pass to the EF, which runs reconstruction algorithms very similar to the offline software. It takes about 4 seconds per event and reduces the event rate down to 400 Hz. The selected events can be recorded for offline analysis.

In order to keep low thresholds, some triggers are prescaled. In prescaled triggers, only a randomly selected fraction of the events passing the trigger cuts are passed onto the next level. The prescale value indicates that only a fraction $1/\text{prescale}$ of the events are accepted. ‘Streams’ of data are recorded according to whether the events passed one or more of the triggers included in a certain list. Data streams include the ‘JetTauEtmiss’, ‘Egamma’ and ‘Muon’ streams.

A special stream is the ‘delayed’ stream of data [79]. It is motivated by the realisation that EF triggers are constrained mainly by the processing capacity, and not the storing capabilities. The ‘delayed’ stream of 8 TeV data recorded additional events at 200 Hz, selecting them with the L1 and L2 triggers, and storing them to be reconstructed after data taking. Softer EF triggers with low or no prescales were run on that fraction of the data.

2.3.1. Jet triggers

The nominal triggers used in the two analyses presented in this thesis are based solely on jets. The effective trigger efficiency of interest for the offline analysis is a combination of the three stages of the trigger and, in particular, how they reconstruct the jets.

At the L1 level, jet elements are the sum of 2×2 trigger towers in the EM calorimeters and 2×2 towers in the hadronic calorimeters. A threshold is set on the sum of the transverse energy of the calorimeter cells within sliding windows covering 4×4 towers. The sums are compared with the pre-set threshold to decide if the event passes on to the next trigger level.

In the original design of the L2 jet trigger, L1 RoIs were used as seeds to define the region in which more refined calorimeter objects would be studied, using the full calorimeter granularity and running a simple cone-like jet algorithm. However, this approach is insufficient for multi-jet events, since the efficiency of the L1 jet trigger to identify close-by jets is very low. Full calorimeter reconstruction at the L2 level — with L1-level granularity, to keep within time constraints — was introduced in 2012. The improvement with respect to the L1 full scan is that sophisticated jet algorithms like anti- k_t [80] were run, providing a measurement closer to the offline result. The use of trigger towers at L2 to do a full detector scan is usually abbreviated *L2FS* [81].

At the EF level, the objects used to build the jets are topological clusters of cells. This improves the resolution with respect to the L2 measurement significantly. Jets are formed using the anti- k_t algorithm using topological clusters from across the whole detector. All the different types of jet algorithm and jet calibration used offline are also available at the EF level. Jet algorithms will be discussed in more detail in Section 2.5.4.

2.4. Detector simulation

In order to interpret the measurements taken by ATLAS, a full simulation of the particle interactions including the effects of the detector is needed. The particles are generated as explained in Section 1.3 and then propagated through a simulation of the ATLAS detector performed with the GEANT4 software toolkit [82, 83]. The energy deposits made by outgoing particles in the detector material are recorded with the same format as the standard ATLAS detector readout, and reconstructed using the same software as for real data.

Producing these datasets is computationally expensive, so a lightweight version of the simulation exists that reduces the simulation time by one order of magnitude. The performance improvement is achieved by parametrising the energy profiles of particle showers [84]. This type of detector simulation is called *AtlFastII* or simply *fast* simulation, as opposed to *full* simulation.

2.5. Object reconstruction

The electronic signals recorded in the different ATLAS sub-detectors are translated into particles and other physical objects after a series of complex reconstruction, identification and calibration processes. Each type of particle or derived physical object is studied individually. This is done centrally in ATLAS by means of the different performance groups. This section contains a summary of the reconstruction, identification and calibration mechanisms of the different objects relevant for the analyses presented in this thesis.

2.5.1. Tracks

The tracks of charged particles with $p_T > 0.5$ GeV and $|\eta| < 2.5$ can be reconstructed in the ID. The reconstruction process involves a succession of algorithms which define the basic tracker objects; perform the track seeding, fitting, and cleaning; remove fake tracks; extrapolate between the different sub-detectors; and perform a final, global re-fitting to provide the ultimate result.

Tracks are also reconstructed in the muon spectrometer. The procedure has elements in common with the ID track reconstruction process. In this case, track candidates are built from segments, which are straight lines in a single MDT or CSC station, and are joined together by a global track-fitting procedure.

2.5.2. Electrons

Electrons are reconstructed differently depending on which region of the detector they are produced in [85, 86]. In the central region ($|\eta| < 2.5$), a three-step process is performed. First, clusters are built in the EM calorimeter using a sliding-window algorithm of a fixed $\Delta\eta \times \Delta\phi$ size. Then, tracks with $p_T > 0.5$ GeV are extrapolated to the central layer of the EM calorimeter, and matched in η to the EM clusters. An electron candidate is formed when at least one track is matched to the cluster³. After reconstruction, the energy of the electron candidate is adjusted by optimising the cluster size in the different calorimeter regions. Other corrections include the energy deposit in front of, outside or beyond the EM calorimeter. The absolute energy scale is determined by exploiting

³If a cluster is not matched with a track, it is classified as an unconverted photon candidate.

benchmark processes such as $Z \rightarrow ee$, $J/\psi \rightarrow ee$ or $W \rightarrow e\nu$. In the forward region, as there is no tracking information, no distinction is possible between electrons and photons.

Electrons are identified using a set of sequential criteria on different calorimeter, tracking and combined tracking/calorimeter variables, binned in η and E_T . Variables used in the identification include shower shapes, hits on particular sections of the detector, or relative energies. Three benchmark selection criteria are set to provide increasing power of background rejection, at the same time that some identification efficiency is lost. They are referred to as the *loose*, *medium* and *tight* quality criteria.

2.5.3. Muons

Muons are reconstructed using a combination of the momentum measurements produced by the ID and the muon spectrometer [87]. Four different techniques are used to reconstruct muons with energies between 3 GeV and 3 TeV: ‘stand-alone’ muons, which only use data from the muon spectrometer; ‘combined’ muons, which combine tracks in the ID and muon spectrometer; ‘segment-tagged’ muons, which combine a track in the ID with a local track segment in the muon spectrometer (not reconstructed as a track); and ‘calorimeter-tagged’ muons, formed from a track in the ID and a minimum-ionising-particle-like energy deposit in the calorimeter. The standard class of muons is the combined type. The combined use of these different reconstruction mechanisms ensures that muons are not misreconstructed even if they have low energy or are produced near a transition region.

The momentum scale and resolution of muons are calibrated using large Monte Carlo samples of dimuon resonances, including $J/\psi \rightarrow \mu\mu$, $\Upsilon \rightarrow \mu\mu$ and $Z \rightarrow \mu\mu$.

2.5.4. Jets

Jets are narrow and approximately conical bundles of the particles produced by the hadronisation of quarks and gluons in a high-energy physics experiment. This may sound vague; in fact there is no unique jet definition. The different options are classified according to two considerations: what is the set of rules used to group together the particles into a common jet, or *jet algorithm*; and how the momenta of the particles inside the jet are combined, or *recombination scheme*.

Jet reconstruction

Two main types of jet algorithm exist: cone algorithms and sequential-recombination algorithms. The definition of a jet should be such that if an event is modified through soft and collinear radiation, the set of reconstructed jets stays the same. Most variants of cone algorithms are either collinear or infrared unsafe, so they will not be discussed here. In ATLAS sequential-recombination algorithms are used instead, which construct jets via a bottom-up approach in which the sequence of parton splittings is reversed.

The default algorithm used by ATLAS is the anti- k_t algorithm [80]. It is collinear and infrared safe, produces jets of a conical shape, and can be implemented efficiently [88]. The algorithm proceeds as follows. For every pair of particles, the distance measure d_{ij} is defined,

$$d_{ij} = \min \left(\frac{1}{p_{T,i}^2}, \frac{1}{p_{T,j}^2} \right) \frac{\Delta R_{ij}^2}{R^2}, \quad (2.3)$$

where $\Delta R_{ij}^2 = (y_i - y_j)^2 + (\phi_i - \phi_j)^2$, and y_i , ϕ_i and $p_{T,i}$, are respectively the rapidity, azimuth and transverse momentum of particle i . The radius parameter R is set to 0.4 by default. In addition, for every particle a ‘beam distance’ d_{iB} is also defined,

$$d_{iB} = p_{T,i}^2. \quad (2.4)$$

For every particle, its beam distance is compared with all the possible d_{ij} . If the smallest distance is one in the d_{ij} set, then particles i and j are recombined into a new particle. If the smallest distance is d_{iB} , then particle i is defined as a jet and removed from the list of input particles. The distances are recalculated and the procedure repeated until no input particles are left. The anti- k_t algorithm was born from a generalisation of other sequential-recombination algorithms called k_T [89, 90] and Cambridge/Aachen [91, 92], both of which produce jets with irregular shapes due to soft radiation. This is illustrated in Figure 2.6.

The default recombination scheme is to add the four-vectors of the constituent particles, which naturally gives mass to the resulting jet even if the input particles are massless.




Figure 2.6.: Example of a simulated parton-level event clustered with the k_T (left) and the anti- k_t (right) algorithms. The shaded regions correspond to the active areas of the jets in y - ϕ space. Figures taken from Ref. [80], where more details are provided.

Jets in ATLAS

In ATLAS, the inputs to the anti- k_t jet algorithm can be either calorimeter ‘topoclusters’, tracks, or stable simulated particles in the case of ‘truth’-level Monte Carlo simulations. Track jets are not used in this thesis, and truth-level jets will only be used in Chapter 4. Jets can be built with different radii; in this thesis, only $R = 0.4$ jets will be used.

Topoclusters are formed from topologically connected cells in the calorimeter containing an energy deposit above the noise threshold. Topoclusters are calibrated assuming that the energy was deposited by electromagnetic showers by default, and then corrected for the hadronic shower contributions in what is called the ‘local cell signal weighting’ (LCW) procedure⁴. The procedure includes a classification of the clusters as electromagnetic or hadronic, based mostly on the energy density and longitudinal shower depth.

Jets are then calibrated to restore their energy scale to that observed in simulated truth particle jets [93]. The procedure consists of several steps. First, pile-up contributions are suppressed using the jet areas pile-up subtraction mechanism [94,95]. Then, a residual pile-up correction is applied, parametrised as a function of the average number of interactions per bunch crossing, $\langle\mu\rangle$, and the number of primary vertices, N_{PV} . The second stage consists on redirecting the jet in such a way that the origin points to the

⁴The ATLAS calorimeter response to electrons is typically 1.3 times higher than the hadron response. This difference is due to energy losses in the hadron case, such as nuclear break-up, spallation and excitation, soft neutrons, neutrinos produced in hadron decays, etc.

hard-scatter vertex, instead of the detector vertex. Finally, a correction dependent on p_T and η is applied to take the reconstructed energy of the jet to that of the corresponding truth particle jet as seen in Monte Carlo simulations. A final, residual correction derived from benchmark physics processes measured in data is applied. This correction uses ‘in situ’ techniques that exploit the balance between jets and well-measured reference objects.

An extra calibration stage was developed at the end of Run 1 to improve the jet energy resolution and reduce the sensitivity of the calibration to the flavour of the jets, without affecting the mean jet energy. The corrections, based on several jet properties — including the number of muon segments behind a jet or the fraction of the energy deposited in different layers of the calorimeter —, are applied sequentially. The procedure is known as *global sequential calibration* (GSC) [96], and is applied after all the steps described above.

2.5.5. *b*-tagging

Hadronic jets formed from heavy flavour quarks can be identified and tagged with high efficiency using the so-called *b*-tagging algorithms. The most basic algorithms used in ATLAS are likelihood ratio tests based on the impact parameter (IP3D) or secondary vertices (SV1) [97, 98]. There also exist more sophisticated methods that exploit the topology of the weak decay of *b*-hadrons, like *JetFitter*. In all cases, the output of the algorithm is a weight computed for every jet in the following way:

$$w_{\text{jet}} = \sum_i w_i = \sum_i \log \frac{b(x_i)}{u(x_i)}, \quad (2.5)$$

where x_i is the discriminating variable, the index i represents either the tracks or the vertices associated with the jet, and $b(x)$ and $u(x)$ are probabilities — obtained by comparing the measured value of x_i with the Monte Carlo expectations — for the *b* and light jet hypotheses, respectively. Jets are tagged as coming from a *b*-hadron when their weight is above a certain value, which defines an efficiency working point.

It is still possible to optimise the result (that is, minimise the mistag rate without losing efficiency) by combining the jet weights obtained from different individual algorithms. In particular, the IP3D, SV1 and JetFitter weights are used as inputs to an artificial neural

network which produces new weight probabilities for b , c and light-flavour jets, referred to as the MV1 weights.

2.5.6. Missing transverse momentum

When the protons collide in ATLAS, their momentum is contained in the longitudinal z direction. By conservation of momentum, it is then expected that the total momentum in the plane transverse to the beam will also be null after the collision. A non-zero global momentum imbalance in the transverse plane can therefore be interpreted as a sign that one or more particles escaped detection. Neutrinos and other hypothetical weakly interacting particles are among the possible sources of missing transverse momentum.

The missing transverse momentum $\mathbf{p}_T^{\text{miss}}$ is computed as the negative vector sum of all the energy deposits in the detector. Its magnitude is usually referred to as ‘missing transverse energy’, or E_T^{miss} , and it is the most widely used form of the observable.

ATLAS has developed an optimised reconstruction of the E_T^{miss} which relies partly on all the other calibrated and pile-up suppressed physics objects [99, 100]. The objects are added in a particular order which defines the priorities at the time of removing the overlap between them: electrons with $E_T > 10$ GeV, photons with $E_T > 10$ GeV, jets with $E_T > 20$ GeV and muons with $E_T > 10$ GeV. All the locally calibrated and unmatched clusters within $|\eta| < 4.9$, and any jet with $p_T < 20$ GeV are included in a separate term known as the *soft* term, which is particularly sensitive to pile-up contributions.

E_T^{miss} can thus be written as

$$\begin{aligned} (E_T^{\text{miss}})^{\text{RefFinal}} = & (E_T^{\text{miss}})^{\text{Electron}} + (E_T^{\text{miss}})^{\text{Gamma}} + \\ & (E_T^{\text{miss}})^{\text{Jet}} + (E_T^{\text{miss}})^{\text{Muon}} + (E_T^{\text{miss}})^{\text{Soft}}, \end{aligned} \quad (2.6)$$

where RefFinal is the name of the algorithm described above [100]. A consequence of this reconstruction approach is that the mismeasurement of the physics objects may be a source of missing transverse momentum. In these cases it is said that the E_T^{miss} is *fake*.

The dependence of E_T^{miss} on calorimeter observables results in its resolution following a stochastic behaviour, which can be approximately parametrised as $\sigma \propto \sqrt{\Sigma E_T}$, where ΣE_T is defined as the scalar sum of the transverse energies of reconstructed and calibrated

calorimeter objects and of the soft term⁵. The dependence of the E_T^{miss} resolution on ΣE_T has been shown to hold in many situations, including events with different numbers of jets. This fact is the basis of the data-driven multi-jet background determination used in Chapter 3.

⁵The noise and constant terms in the resolution contribute at the very low and very high ΣE_T regions, respectively.

3

Search for new phenomena in events with large jet multiplicities

*‘Do you know, I always thought unicorns were fabulous monsters, too? I never saw one alive before!’
‘Well, now that we have seen each other,’ said the unicorn,
‘if you’ll believe in me, I’ll believe in you.’*

3.1. Introduction

One of the design goals of the LHC is to search for signs of new physics phenomena that could either explain some of the questions left open by the Standard Model of particle physics or widen the horizons of the field by unveiling new mysteries [101]. Historically, high hopes were set on SUSY due to its attractive theoretical properties (as discussed in Section 1.4), but there is currently no experimental evidence that favours this or any other BSM theory.

This chapter will present a search for new physics in events with large jet multiplicities (or simply ‘multi-jet’ events). Such scenarios are relevant, for example, in the decay of strongly-produced pairs of supersymmetric particles. When the decay chains are long, or include heavy intermediate particles, large numbers of jets are produced – with only small amounts of the total energy being contained in the E_T^{miss} . These scenarios are well motivated theoretically [102–104].

Searching for SUSY in multi-jet events has several advantages. Firstly, the SM background is small — although determining it accurately is challenging, for the reasons explained in the QCD Section 1.3. The analysis will rely heavily on data-driven techniques, which will be introduced in the next section and discussed in detail in Sections 3.3 and 3.4. Secondly, the targeted final state is a generic feature of many different BSM scenarios,

including RPC and RPV supersymmetric models. Both cases will be studied here through various representative models.

Performed for the first time in 2011 with 7 TeV data [105], the multi-jet search originated as a spin-off of the search for new physics in events with no leptons, E_T^{miss} and jets [106]; the main difference was that the new analysis had a lower (and indirect) E_T^{miss} cut and a high jet multiplicity requirement. Two subsequent iterations of the analysis [67, 107] increased the sensitivity by selecting events with a more refined background determination technique and applying tighter jet multiplicity cuts, a development which was made possible by the higher energy and luminosity of the LHC.

In this thesis the latest version of the search is presented. It is the most complex of the four, and it was the first SUSY search using the full Run-1 8 TeV dataset to be published [3]. The analysis combines 13 signal regions, all of them built on top of the basic multi-jet selection, characterised by cuts in the jet and b -jet multiplicities. Most of the signal regions are orthogonal to each other, and they are later combined into a global maximum-likelihood fit which constrains the backgrounds and reduces the uncertainties.

The more sophisticated event selection and statistical treatment helped to extend the systematically-limited reach of the previous version of the analysis. As no excess over the SM prediction was found, limits were set on a wide variety of physics models. These results, and the methods and materials used to obtain them, will be presented in the next section and discussed in detail in the rest of the chapter.

3.1.1. Overview of the analysis

The main challenge of the multi-jet analysis is the determination of the SM backgrounds. The basic selection cuts include a lepton veto and a jet multiplicity cut, using jets with $p_T > 50$ GeV or 80 GeV. Despite having a lepton veto, leptonic processes can still contribute to the backgrounds if the lepton is misidentified or out of acceptance. Therefore, the background processes can be divided into two categories: fully hadronic and non-fully-hadronic (or simply ‘leptonic’).

Fully hadronic multi-jet events will be described using a data-driven technique, introduced for the first time in Ref. [105], which exploits some properties of the E_T^{miss}




Figure 3.1: Schematic of the ABCD method used to determine the multi-jet background. The orange blob represents the signal, and the green area represents the multi-jet background.

significance variable $\mathcal{S}_T^{\text{miss}}$, defined as

$$\mathcal{S}_T^{\text{miss}} = E_T^{\text{miss}} / \sqrt{H_T}, \quad (3.1)$$

where H_T is the scalar sum of the jets in the event¹.

The method relies on two observed facts: (i) the shape of the significance-like variable $\mathcal{S}_T^{\text{miss}}$ is invariant with the jet multiplicity; and (ii) QCD events tend to dominate in regions of low $\mathcal{S}_T^{\text{miss}}$ and low jet multiplicity, while signal events tend to populate the region of high $\mathcal{S}_T^{\text{miss}}$ and high jet multiplicity. This same idea is illustrated in Figure 3.1. One can then build ‘templates’ of the $\mathcal{S}_T^{\text{miss}}$ at low jet multiplicities (after subtracting the leptonic contributions expected from MC) and extrapolate them to high jet multiplicities to be used as QCD background predictions. In addition, $\mathcal{S}_T^{\text{miss}}$ can be used as a strong discriminating variable at high jet multiplicities: indeed, after the lepton veto and jet multiplicity cuts, the $\mathcal{S}_T^{\text{miss}}$ cut will be the one that completes the definition of the signal regions.

In this version of the analysis, the jet multiplicity cut includes two steps: a first, flavour-blind cut on the total number of jets in the event, and a second cut on the number of b -jets. For each jet multiplicity cut, three different b -jet regions will be considered: those with exactly none, exactly one, or two or more b -jets, such that all the possibilities are covered. It was found that the assumption of the invariance of $\mathcal{S}_T^{\text{miss}}$ with respect to the jet multiplicity was still valid if both the low jet multiplicity template and the signal region of interest had the same, fixed number of b -jets. This is the basis of the fully-hadronic multi-jet background prediction in the signal regions of this analysis.

¹More details on variable definitions will be given in Section 3.2.4.

Leptonic backgrounds are obtained from Monte Carlo simulations. In order to improve the description of the data, ‘control’ regions are defined in which leptonic backgrounds dominate. These control regions are very similar to the signal regions: they have the same jet multiplicity, b -jet multiplicity and $\mathcal{S}_T^{\text{miss}}$ cuts, but they have a one-lepton requirement. This way, contributions from W and $t\bar{t}$ (the main leptonic backgrounds) are enhanced; the different b -jet requirements naturally distinguish between the two cases. All the control regions are included in a likelihood fit which simultaneously adjusts the backgrounds and their uncertainties. Fits are also used to set limits on models of interest.

The remainder of this chapter is organised as follows. Section 3.2 defines the analysis strategy. Sections 3.2.1 and 3.2.2 describe the signal, the data and the triggers used. Sections 3.2.3 and 3.2.4 define the physics objects used and the kinematic variables, and 3.2.5 outlines the main cleaning cuts. Section 3.3 describes the main backgrounds to the search and the methods used to estimate them. The multi-jet background is discussed in Section 3.4, and the leptonic backgrounds in Section 3.6. Systematic uncertainties are discussed in Section 3.7. Finally, Section 3.8 includes an overview of the statistical procedure used to interpret the result, and the corresponding model-independent and model-dependent exclusion limits on new physics.

3.2. Analysis strategy

3.2.1. Signal and backgrounds

Most BSM theories that predict final states with no leptons, large numbers of strongly interacting particles and missing momentum can be probed by the multi-jet analysis. The signal models discussed here are representative examples of different signal processes. Although the analysis focuses on SUSY, the results can also be interpreted in terms of generic, strongly-interacting, heavy particles.

The majority of the models that were used for the signal region optimisation and exclusion limits fall into the category of simplified models. These models focus on a particular production and decay process. Sparticles not participating in the process of interest are generally decoupled by taking them to very high masses, and the decay and branching ratios and other kinematic features of the chosen decay process are fixed to trivial values – e.g. 100% branching ratios for particular channels. This way, the focus is not so much on the details of the model, but rather on the final states the analysis

is sensitive to. The limits set on these models can be extrapolated to any other theory predicting the same final state.

The multi-jet analysis is interpreted in terms of the following simplified models:

1. ***Gtt* model.** The only process studied in this model is the pair production of gluinos, which then decay with 100% branching ratio to t and $\bar{t}$ and the lightest neutralino (which is also the LSP),

$$\tilde{g} \rightarrow t + \bar{t} + \tilde{\chi}_1^0. \quad (3.2)$$

The decay is mediated by an off-shell stop quark. The final state therefore comprises four top quarks and missing momentum. From the decay of the tops one would expect between four and twelve quarks, which corresponds to a high jet multiplicity, including four heavy-flavour jets. The mass splitting between the LSP and the gluino must be at least twice the top quark mass, and the rest of the sparticles are much heavier than the gluino. The *Gtt* model is used for the optimisation of the jet definition. For a Feynman diagram of the process, see Fig. 3.2a.

2. ***Gtt*-on-shell model.** It is similar to *Gtt*, but with $m_{\tilde{t}} < m_{\tilde{g}}$ so that the intermediate virtual stop can be on its mass shell, as shown in Fig. 3.2b. The full gluino decay chain is then

$$\tilde{g} \rightarrow \bar{t} + \tilde{t} \rightarrow \bar{t} + t + \tilde{\chi}_1^0. \quad (3.3)$$

The exact kinematic constraints are $m_{\tilde{g}} > m_{\tilde{t}} + m_t$, and also $m_{\tilde{t}} > m_t + m_{\tilde{\chi}_1^0}$. The branching ratios are again set to 100%. Although the optimisation of the analysis cuts has been derived for *Gtt*, it has been checked that the selection applied was appropriate for *Gtt*-on-shell too.

3. **One-step model.** This model contains again direct gluino pair production, but with an extra step in the decay of the gluino. First it decays to two light quarks and a chargino via an intermediate squark,

$$\tilde{g} \rightarrow \bar{q} + \tilde{q} \rightarrow \bar{q} + q' + \tilde{\chi}_1^\pm. \quad (3.4)$$

The chargino subsequently decays into a W and the LSP,

$$\tilde{\chi}_1^\pm \rightarrow W^\pm + \tilde{\chi}_1^0, \quad (3.5)$$

as shown in Fig. 3.2c. There are three free parameters in total: the masses of the gluino, the chargino and the LSP. In order to be able to draw the exclusion limits in two-dimensional planes, one degree of freedom must be fixed. Two choices are considered in the exclusion diagrams:

$$\text{a) } m_{\tilde{\chi}_1^\pm} = \frac{m_{\tilde{g}} + m_{\tilde{\chi}_1^0}}{2}$$

$$\text{b) } m_{\tilde{\chi}_1^0} = 60 \text{ GeV}$$

This model is characterised by the lack of heavy-flavour jets in its final state.

4. **Two-step model.** It is similar to the one-step model, but with a second intermediate step in the gluino decay,

$$\begin{aligned} \tilde{g} &\rightarrow q + \tilde{q}, \\ \tilde{q} &\rightarrow q + \tilde{\chi}_1^\pm \rightarrow q' + W + \tilde{\chi}_2^0 \rightarrow q' + W + Z + \tilde{\chi}_1^0 \end{aligned} \quad (3.6)$$

The intermediate particle masses, $m_{\tilde{\chi}_1^\pm}$ and $m_{\tilde{\chi}_2^0}$, are set to $(m_{\tilde{g}} + m_{\tilde{\chi}_1^0})/2$ and $(m_{\tilde{\chi}_1^\pm} + m_{\tilde{\chi}_1^0})/2$, respectively. Fig. 3.2d shows the full Feynman diagram of the process. Like in the one-step model, there are no heavy-flavour jets in the final state, and the number of partons in the final state can be as high as 12 if both W 's and Z 's decay hadronically.

5. **mSUGRA/CMSSM.** An mSUGRA/CMSSM model (see Section 1.4.2 on page 22) is also used for interpretation for historical reasons, as it was originally one of the preferred scenarios. The ratio of the vacuum expectation values of the two Higgs fields ($\tan \beta$) is set to 30, the universal trilinear scalar coupling is $A_0 = -2m_0$, and the sign of the higgsino mass parameter is $\mu < 0$. This choice of parameters accommodates a lightest Higgs boson compatible with the Higgs boson observed at the LHC [15–17]. Only strong production (squark and gluino production) and associated electroweak production (gluino-gaugino and squark-gaugino) processes are simulated.
6. **R-parity violating model.** This is a simplified model where only gluino pair production is allowed. The gluino decays, as in Gtt -on-shell, via

$$\tilde{g} \rightarrow \tilde{t} + t, \quad (3.7)$$

and the stop decays via the R-parity violating channel [108]

$$\tilde{t} \rightarrow b + s. \quad (3.8)$$

The process is depicted in Fig. 3.2e. The only missing momentum in the final states comes from the neutrinos produced in the decays of the b quarks. The multi-jet analysis can nevertheless set very stringent limits on this grid due to its low, indirect missing momentum requirement.




Figure 3.2.: Benchmark signal processes studied in the multi-jet analysis. No attempt is made to distinguish anti-particles from the corresponding particles. The diagrams are produced by the ATLAS Collaboration.

To test the different models, Monte Carlo simulations produced centrally by the ATLAS Collaboration were used. For Gtt , Gtt -on-shell, mSUGRA and the RPV sample the generator used was HERWIG++ [109]; the one-step and two-step samples were generated using MADGRAPH [30] for the matrix element calculation and PYTHIA [110] for the showering.

The cross sections are calculated centrally at next-to-leading order in the strong coupling constant, and the soft gluon emission is resummed at the next-to-leading-logarithmic accuracy (NLO+NLL), which results in smaller uncertainties related to




Figure 3.3.: Cross section of direct $\tilde{g}$ production as a function of the gluino mass. The coloured band represents the theoretical uncertainty on the cross section.

unknown higher perturbative orders and less dependence on physical scales [111]. The theoretical uncertainties include variations on the renormalisation and factorisation scale, PDF uncertainties from different sets [112, 113], and the value of the strong coupling constant. The final, symmetric uncertainty is obtained as an envelope of the cross sections obtained in all the previous variations, following the prescription proposed in [111]. Figure 3.3 shows the gluino pair production cross section and the corresponding uncertainty.

The main background to the search for any of these signal models is dominated by SM multi-jet production. This includes but is not limited to purely hadronic interactions and the subsequent QCD showers, or top quark pair production with the tops decaying via their purely hadronic modes. These backgrounds are determined in the multi-jet analysis by a data-driven technique, as will be discussed in Section 3.4.

The rest of the background comes from non-fully-hadronic electroweak processes with additional jets from initial or final state radiation, in which the lepton was not identified. This could be due to it being out of acceptance or due to an erroneous object reconstruction and/or identification. This last case is particularly relevant for tau leptons, which may be misreconstructed as jets up to 5% of the time [114]. Backgrounds that fall into this category include top quark pair production, W and Z bosons decaying to leptons and produced together with jets, single top production, or the associated production of a top quark pair and a vector boson. These are all determined from Monte Carlo simulations, and those with the highest yields $-t\bar{t}$ and $W+\text{jets}-$ are normalised in control

regions. The background from Z to neutrinos is irreducible, but its contribution for the large jet multiplicities of the signal regions is found to be negligible. The techniques used to predict the SM backgrounds are described in Section 3.3.

Several sets of selection cuts are used with different purposes. A set of selection cuts is said to delimit a *region* in phase space. Regions expected to have large contributions from signal are called ‘signal regions’, whereas regions used to probe the SM backgrounds receive different names according to their exact use – namely, confirmation, control and validation regions. The exact definition of each type of region will be given in Section 3.3.

Signal Trigger	L1 Item	L2 Chain	Int. Luminosity
EF_5j55_a4tchad_L2FS	L1_4J15	L2_5j15_a4TTem	20.3 fb ⁻¹
EF_6j45_a4tclcw_L2FS_5L2j15	L1_4J15	L2_5j15_a4TTem	20.3 fb ⁻¹
Control Region Triggers	L1 Item	L2 Chain	Int. Luminosity
EF_4j65_a4tchad_L2FS	L1_4J15	L2_4j20_a4TTem	1.3 fb ⁻¹
EF_5j45_a4tchad_L2FS	L1_4J15	L2_5j15_a4TTem	1.4 fb ⁻¹
EF_mu24i_tight	L1_MU15	L2_mu24_tight	20.3 fb ⁻¹
EF_e24vhi_medium1	L1_EM18VH	L2_e24vhi_medium1	20.3 fb ⁻¹

Table 3.1.: Triggers used in the signal and control regions of the multijet analysis.

3.2.2. Dataset and trigger

All the data taken by the ATLAS detector between the months of March and December of 2012 passing the quality criteria were analysed. After selecting the luminosity blocks included in the 2012 *Good Run List* (GRL) created centrally by the ATLAS Data Quality group, the total integrated luminosity corresponds to 20.3 fb⁻¹.

The first analysis-specific cut is the trigger selection. One of the strengths of the multi-jet analysis lies on the possibility to use multi-jet triggers, which have low enough rates to not need an additional E_T^{miss} cut at trigger level. Event Filter (EF) and offline jets are reconstructed using the anti- k_t algorithm with $R = 0.4$.

The triggers used to select the events in the signal regions required 5 and 6 jets with p_T cuts at the EF level of 55 and 45 GeV, respectively. The corresponding offline jet p_T thresholds used were 80 GeV and 50 GeV respectively, so that the efficiency for signal

region jet multiplicities is $> 99\%^2$. Figures 3.4 and 3.5 show the corresponding efficiency curves versus the fifth, sixth and seventh jet p_T , respectively. Details of the calculation of the efficiency curves are given in appendix A. Table 3.1 shows the complete list of triggers used in the analysis and the corresponding integrated luminosity.

Different triggers were used for the control, validation and confirmation regions. QCD control regions with lower jet multiplicity requirements were selected with 4-jet ($p_T > 65$ GeV) and 5-jet ($p_T > 45$ GeV) triggers, depending on whether the p_T cut of the offline jets had $p_T > 80$ GeV or $p_T > 50$ GeV. Non-QCD control regions, where one signal lepton is required and the jet multiplicity cut is kept the same as the corresponding signal regions, were selected with the lowest unprescaled single-electron and single-muon triggers available in the menu. Both triggers reach full efficiency at a lepton p_T of 25 GeV, which is used as the minimum p_T cut for control region leptons.




Figure 3.4.: Efficiencies of the 5j55 trigger versus the 5th jet offline p_T (left) and the 6j45 trigger versus the 6th jet offline p_T (right) calculated using the full dataset, and using the EF_e24vhi_medium1 trigger as a reference. The blue curve corresponds to the trigger efficiency calculated in the Sherpa semileptonic $t\bar{t}$ sample, and the black curve corresponds to the trigger efficiency calculated in data. The vertical dashed lines on the main plot indicate offline p_T thresholds.

²The small inefficiency seen in Figure 3.4 for events with the sixth jet p_T between 50 and 55 GeV could only affect the signal regions indirectly via the multi-jet background prediction. A dedicated systematic uncertainty is derived, and found to be negligible. The uncertainty will be discussed in Section 3.5.4.




Figure 3.5.: Efficiencies of the 5j55 trigger versus the 7th jet offline p_T (left) and the 6j45 trigger versus the 7th jet offline p_T (right) calculated using the full dataset, and using the EF_e24vhi_medium1 trigger as a reference. The blue curve corresponds to the trigger efficiency calculated in the Sherpa semileptonic $t\bar{t}$ sample, and the black curve corresponds to the trigger efficiency calculated in data. The vertical dashed lines on the main plot indicate offline p_T thresholds.

3.2.3. Physics object definitions

The definitions of the objects used by the analysis followed the recommendations of the corresponding ATLAS performance groups. The cuts described here ensure that the objects are well calibrated, the triggers are efficient and the purity of the selection is sufficiently high. The E_T^{miss} is built according to the description in Section 2.5.

Jets are reconstructed using the anti- k_t algorithm with $R = 0.4$ and calibrated following the procedure explained in Section 2.5.4. The baseline jet selection requires $p_T > 20$ GeV and $|\eta| < 2.8$. This selection ensures good closure³ of the jet calibration [93] and compatibility with the jets used by the trigger, which have $|\eta| < 3.2$. Each signal region requires a minimum number of jets above a certain p_T threshold and within $|\eta| < 2.0$ (see Section on optimisation, 3.2.6), but does not veto any additional jets. The jet collection may initially contain electrons, since they too deposit their energy in

³Closure here refers to the self-consistency of the calibration procedure, i.e. the recovery of the truth-level jet energy scale when the calibration procedure is applied on the simulated events used to derive it.

the electromagnetic calorimeter. The overlap is removed by discarding any jet within $\Delta R = \sqrt{\Delta\phi^2 + \Delta\eta^2} < 0.2$ of an electron. In addition, any electron or muon within $\Delta R < 0.4$ of any remaining jet is in turn discarded, as they may come from the semi-leptonic decay of a heavy-flavour hadron. Final jet quality cuts are applied to jets surviving the overlap removal to check that their origin is not from hardware problems, LHC beam related backgrounds or cosmic ray showers. All events containing one or more jets failing these cuts are rejected.

Electrons are classified according to two sets of criteria. The looser set defines the baseline electrons, which are required to have $p_T > 10$ GeV and $|\eta| < 2.47$. They must satisfy a series of requirements on the shower shape, track hits and alignment of the tracks and calorimeter clusters, as outlined in Ref. [85], but optimised for the conditions of the 2012 data; they are collectively known as the *medium* criteria. They must be separated by at least $\Delta R > 0.4$ from any jet, for the reasons explained above. Baseline electrons are used for the signal region lepton veto. The electrons used to define the leptonic control regions have to pass additional cuts: they must have $p_T > 25$ GeV, satisfy the *tight* quality criteria, have a transverse (longitudinal) impact parameter within 5σ (0.4 mm) of the primary vertex, and be isolated. An electron is said to be isolated when the momenta of the nearby tracks and calorimeter clusters are below a certain threshold. In this case, the scalar sum of the transverse momenta of tracks⁴ in a cone of $\Delta R = 0.3$ around the electron is required to be smaller than 16% of the electron p_T . The scalar sum of the transverse energy of the calorimeter clusters contained in that same cone is in turn required to be smaller than 18% of the electron p_T .

Muons are used in the same context as electrons. Baseline muons must have $p_T > 10$ GeV and $|\eta| < 2.5$, satisfy track quality criteria and be separated by at least $\Delta R > 0.4$ from any jet. Control region muons are required to have $p_T > 25$ GeV and $|\eta| < 2.4$, a transverse (longitudinal) impact parameter within 5σ (0.4 mm) of the primary vertex, and they have to be isolated. Similarly to electrons, the isolation requirements are defined in terms of the scalar sum of the transverse momenta of tracks and calorimeter clusters within a $\Delta R = 0.3$ cone, which here have to be less than 12% in both cases.

b-jets are defined as jets with $p_T > 40$ GeV within the region covered by the inner detector, $|\eta| < 2.5$, which have been identified by the MV1 algorithm as originating

⁴Other than its own.

from a b quark. The MV1 b-tagging algorithm is based on a neural network using the output weights of the JetFitter+IP3D, IP3D and SV1 algorithms as input; the details are described in Ref. [98] and summarised in Section 2.5.5 on page 46. MV1 assigns a value between 0 and 1 to each jet, where 0 corresponds to light-flavour and 1 to heavy flavour. The chosen working point is 0.7892, which corresponds to a b-tagging efficiency of 70%, purity of 92.28%, and rejection factors of 4.97, 13.24 and 136.6 for c-, tau- and light-jets respectively, as measured in $t\bar{t}$ events. Rejection factors are defined as the inverse of the fake rate.

3.2.4. Kinematic variable definitions

H_T H_T is defined as the scalar sum of the p_T of all the jets in the event which have $p_T > 40$ GeV and $|\eta| < 2.8$,

$$H_T = \sum_{\substack{p_T > 40 \text{ GeV} \\ |\eta| < 2.8}} p_T^{\text{jet}}. \quad (3.9)$$

The p_T and η cuts were determined in a previous version of the analysis as part of the studies to improve the data-driven prediction of the multijet background [67]. The jet p_T and η cuts are more inclusive than those used to define signal regions to improve the correlation with the E_T^{miss} , which is essential for the $\mathcal{S}_T^{\text{miss}}$ -based multi-jet background data-driven prediction.

$\mathcal{S}_T^{\text{miss}}$ Usually referred to as “ E_T^{miss} significance”, $\mathcal{S}_T^{\text{miss}}$ is a ubiquitous variable in the analysis, used to define both signal and control regions. It is defined as

$$\mathcal{S}_T^{\text{miss}} = E_T^{\text{miss}} / \sqrt{H_T}. \quad (3.10)$$

It is a significance-like variable, since the denominator is highly correlated with the resolution of E_T^{miss} [100]. Significance variables are known to have high discriminatory power [115]. In this analysis $\mathcal{S}_T^{\text{miss}}$ is also used to extrapolate the multi-jet background template from the QCD-dominated region to the signal-dominated region, as explained in Section 3.4.

3.2.5. Event cleaning

The cleaning cuts applied followed the recommendations from the relevant ATLAS groups. They are designed to maximise the purity of the analysed events – that is, to make sure that there were no problems with the detector and that the particles originated in the proton-proton collision. This is especially important in searches for new physics, as they are made to be sensitive to exotic features that could be easily faked by detector effects. The most relevant cleaning cuts applied are described in this section.

Firstly, only runs and luminosity blocks included in the GRL provided by the ATLAS Data Quality group are included in the analysis. The events contained therein are expected to be mostly free from hardware problems, but there may still be some issues for which additional cleaning cuts are required. Events with reported problems in the LAr or tile calorimeters, like noise bursts, data integrity errors or high voltage trips are flagged and removed from the event set.

Some of the areas of the detector had malfunctions or even became inoperative at different moments of data taking. Jets pointing in the direction of these regions could fake E_T^{miss} , so if a significant fraction of the E_T^{miss} is found to be associated to problematic detector areas the event is rejected. Events containing any single jet identified as being caused by detector effects –such as a calorimeter noise spike– or originating from non-collision backgrounds or cosmics are also rejected.

Some further cleaning cuts are applied to reduce the number of events with bad or non-collision data. They are as follows:

- Cosmic muon veto. Events are rejected if they contain a muon with $|z_0| > 1$ mm or $|d_0| > 0.2$ mm after the overlap removal. z_0 is defined as distance in the $r - z$ plane between the intersection of the track direction with the beam line and the primary vertex, and d_0 is the transverse impact parameter.
- Bad muon veto. Events are rejected if they contain a muon with a q/p (charge over momentum) fractional error larger than 0.2.
- Vertex cut. The vertex with the highest value of $\sum (p_T^{\text{track}})^2$, where the sum runs over all the associated tracks, must have at least 5 such tracks.
- Charge fraction cut. If the leading two jets after overlap removal, provided they have $p_T > 100$ GeV and $|\eta| < 2.0$, satisfy either $f_{ch} < 0.02$, or $f_{ch} < 0.05$ and

$f_{EM} > 0.9$, the event is rejected. Here f_{ch} is defined for each jet as

$$f_{ch}^{\text{jet}} = \frac{\sum_{\text{all associated tracks}} p_{\text{T}}^{\text{track}}}{p_{\text{T}}^{\text{jet}}} \quad (3.11)$$

and f_{EM} is the fraction of the energy measured in the electromagnetic calorimeter. This selection cut rejects events where jets were not generated by the hard interaction, be they cosmics, beam background or detector noise.

3.2.6. Optimisation of the signal region definitions

Previous versions of the multi-jet analysis [67, 105, 107] set the grounds of the search for new phenomena in multi-jet events using data from the ATLAS detector. In all of them, the search was structured around two fundamental variables: the jet multiplicity and $\mathcal{S}_{\text{T}}^{\text{miss}}$, introduced in Section 3.2.4. The 8 TeV of centre-of-mass energy achieved by the LHC in 2012, as well as the high luminosity delivered, made it possible this time to use the largest jet multiplicity cuts to date in an ATLAS analysis, and additional selection cuts according to the flavour of the jets in the event. The exact selections had to be optimised; the procedures used will be explained next.

It should be noted that the optimisation process, as well as the validation of the background prediction techniques, were performed with restricted access to data. The optimisation was performed using only a small fraction of the data, corresponding to 5.8 fb⁻¹. This fraction corresponds to the dataset used by a previous version of the analysis [67]. Potential signal regions (i.e., the high $\mathcal{S}_{\text{T}}^{\text{miss}}$ and high jet multiplicity regions) were not looked at in data during the development stage. This procedure is normally referred to as ‘blinding’. The analysis was unblinded once the signal region definitions were frozen and the description of the backgrounds was under control. This is done to avoid biases in the design of the analysis due to the presence (or absence) of a signal.

Optimisation of the jet definition. The jets in the Gtt and Gtt -on-shell grids are in general more central than those in the SM background processes. This can be observed in Figure 3.6 on page 70, which shows the η distributions for some Gtt and Gtt -on-shell models with different sparticle mass configurations, compared to the shape of the $t\bar{t}$ semileptonic background. There are several reasons why one would expect to find more central jets in the Gtt signal. Firstly, SUSY particles are in general more massive than the

SM ones and are therefore created with smaller boost in the z -direction, which leads to more central topologies. Secondly, SUSY production is dominated by s -channel processes, while $t\bar{t}$ +ISR is predominantly produced through a t -channel process (between $t\bar{t}$ and the ISR jet). Typically s -channel processes have lower $\Delta\eta$ between the jets.

The fact that signal jets typically have higher p_T than those coming from background processes can be seen in Figure 3.6e. However, the peak in the p_T distribution of the signal jets decreases with jet multiplicity, so that in regions with a sufficiently constraining jet multiplicity, a tight p_T cut would be too aggressive and could reduce the significance of the signal, as can be seen in Figure 3.6f.

Following the approach taken in previous versions of the analysis, two sets of signal regions with different jet p_T cuts are considered. The jet p_T cuts used in the past were 55 GeV and 80 GeV.

The optimal cuts in η , p_T and jet multiplicity were obtained by maximising the sensitivity to the gluino mass in the Gtt model for the low- p_T set of signal regions. In particular, the cuts tested were 2.8, 2.0 and 1.5 for η , 50 and 55 GeV for the jet p_T , and at least 7, 8, 9 and 10 jets respectively in jet multiplicity. The ranges were constrained by the multi-jet trigger used (see Section 3.2.2). The higher- p_T set of signal regions (using 80 GeV jets) was not considered in the optimisation. The overall conclusions are illustrated in Figure 3.7 (page 71), and can be summarised as follows:

- Regions with low jet multiplicity (7, 8) increase their sensitivity in the Gtt plane with tighter cuts in η and p_T , yielding the best results for $|\eta| < 1.5$ and $p_T > 55$ GeV.
- Regions with higher jet multiplicities (9, 10) are more sensitive to a reduction in the number of signal events. The 9 jet signal region is most sensitive for $|\eta| < 1.5$ and $p_T > 55$ GeV, but the 10 jet signal region reaches its maximum coverage at $|\eta| < 2.0$ and $p_T > 50$ GeV, due to the limited number of events.

The exclusion limits were obtained with the HistFitter package [116]. HistFitter uses the HistFactory tool [117], part of the RooStats package [118]. A more detailed discussion of the fitting procedure will be given in Section 3.8.1. The author calculated all the inputs to the exclusion limits, and the HistFitter code was run by a collaborator. The following assumptions on the size and nature of the systematic uncertainties were made:

- Signal uncertainty: 0.1 correlated, 0.2 uncorrelated.
- Background uncertainty:
 - ◊ ≤ 10 background events: $0.5/\sqrt{2}$ correlated, $0.5/\sqrt{2}$ uncorrelated.
 - ◊ 10 to 30 background events: $0.35/\sqrt{2}$ correlated, $0.35/\sqrt{2}$ uncorrelated.
 - ◊ > 30 background events: $0.1/\sqrt{2}$ correlated, $0.1/\sqrt{2} + 0.4/\sqrt{0.2 \times N}$ uncorrelated

N is the number of events expected in the signal region, the uncertainties are given normalised to 1, and the correlations refer to whether different regions have uncertainties originating from the same source. The total uncertainty is obtained by adding the correlated and uncorrelated components quadratically. These numbers capture approximately the sizes of the uncertainties seen in previous versions of the analysis. In Figure 3.7, the ‘observed’ limits are obtained by rounding the expected limits to the closest integer.

The most powerful signal region of the ones considered is that with a requirement of at least 10 jets with $p_T > 50$ GeV and $|\eta| < 2.0$. In order to define the final set of signal regions, the b -jet requirements must also be taken into account. This will be done in the next section.

Optimisation of the b -jet signal regions. A selection based on $\mathcal{S}_T^{\text{miss}}$ and jet multiplicity grants sensitivity to a wide variety of models predicting strongly-interacting cascades of particles in their final state. Often these final states have a fixed number of b -jets, including none. Subdividing (or *binning*) each of the regions with a fixed value of the jet multiplicity into sub-regions exclusive in the number of b -jets spanning all the different possibilities, retains the scope of previous versions of the analysis at the same time as it enhances the sensitivity to models with a well-defined number of b -jets in their final state.

Taking into account the conclusions from the p_T and η optimisation, several possibilities for the b -jet binning were studied. The idea underlying all the combinations is to exploit the statistical combination of exclusive signal regions to increase the significance of the signal. Due to the nature of the Gtt model used for the optimisation signal regions with high number of jets and/or b -jets are expected to dominate the exclusion power.

The main combinations considered are found on Table 3.2 and Figure 3.8 (page 72). They were so designed to assess the effect on the overall sensitivity of:

- Including a signal region with 10 jets (A vs B).
- Tightening p_T (B vs E).
- Tightening η for the same p_T cut (D vs E and B vs C).
- Lowering the η cut and including the 10j region (D or E vs F).
- Using 50 or 55 GeV jets (with tight η cuts) (C vs D).

Figure 3.8 shows that Set A provides the highest sensitivity to the Gtt grid. The 7-jet regions were found to not contribute to the sensitivity, so they were removed from the final selection.

The 80 GeV regions used in previous versions of the analysis were kept for completeness, but were not individually optimised; nonetheless, the jet multiplicity was increased with respect to Ref. [67] to 7 and 8 jets with $p_T > 80$ GeV, the η cut constrained to 2.0, and the regions were made exclusive, to be consistent with the 50 GeV regions.

The final set of signal regions is summarised in Table 3.3.

3.3. Estimation of the Standard Model background

One of the advantages of having very tight jet multiplicity requirements in the signal regions is that the contribution from SM processes is small. However, the sensitivity to new physics relies on the accurate description of the few expected background events, and typically small yields are associated to large uncertainties. It then becomes important to ensure that all the contributing processes are well understood, by measuring them in adjacent, less-constrained regions, where new physics processes are expected to be sub-dominant at most.

The SM processes relevant to the analysis can be divided into three different categories:

1. Fully-hadronic multi-jet background, comprising mostly QCD processes and fully hadronic $t\bar{t}$.
2. Non-fully-hadronic (or ‘leptonic’) multi-jet backgrounds, where leptons, if any, were misreconstructed or emitted out of acceptance. This type of background is dominated by non-fully-hadronic $t\bar{t}$ and W +jets.

Jet Multiplicity	η cut	B-jet cut
Set A		
7ej50	2.0	0e/1e/2i
8ej50	2.0	0e/1e/2i
9ej50	2.0	0e/1e/2i
10ij50	2.0	–
Set B		
Set A without 10j and and changing 9ej to 9ij		
Set C		
7ej50	1.5	0e/1e/2i
8ej50	1.5	0e/1e/2i
9ij50	1.5	0e/1e/2i
Set D		
7ej55	1.5	0e/1e/2i
8ej55	1.5	0e/1e/2i
9ij55	1.5	0e/1e/2i
Set E		
7ej55	2.0	0e/1e/2i
8ej55	2.0	0e/1e/2i
9ij55	2.0	0e/1e/2i
Set F		
7ej55	2.8	0e/1e/2i
8ej55	2.8	0e/1e/2i
9ej55	2.8	0e/1e/2i
10ij55	2.8	–

Table 3.2.: Combinations of $|\eta|$, p_T and jet and b -jet multiplicity tested in the Signal Region optimisation. The fourth column indicates the cut in b -jet multiplicity; xe should be read as “exactly x b -jets”, and yi as “at least y b -jets”. The signal sensitivity resulting from these combinations can be seen in Figure 3.8.

$ \eta $ selection for $R = 0.4$ jets	< 2.0						
$R = 0.4$ jet multiplicity & p_{T}	8j50			9j50			10+j50
b -jet requirement	0	1	2+	0	1	2+	Inclusive
$\mathcal{S}_{\text{T}}^{\text{miss}}$	$> 4 \text{ GeV}^{1/2}$						

$R = 0.4$ jet multiplicity & p_{T}	7j80			8+j80		
b -jet requirement	0	1	2+	0	1	2+
$\mathcal{S}_{\text{T}}^{\text{miss}}$	$> 4 \text{ GeV}^{1/2}$					

Table 3.3.: Signal region definitions.

3. Other minor backgrounds, including Z +jets, $t\bar{t} + V$ (where V can be either W or Z) and single top.

The fully-hadronic multi-jet background is one of the most challenging to describe, and will be discussed in Section 3.4. To determine it, a data-driven technique was used. For the leptonic multi-jet backgrounds, which can be dominant in regions with large numbers of b -jets, Monte Carlo simulations were used. However, the large cross sections of $t\bar{t}$ and W +jets made it possible to define specific background-enhanced regions to evaluate and normalise the MC description by comparing it with data. Minor backgrounds from the third category were taken straight from Monte Carlo.

The regions of phase space relevant for background determination were further subdivided into three types of regions, according to how they were used in the analysis: ‘*control region*’, ‘*validation region*’ and ‘*confirmation region*’, in addition to the usual ‘*signal regions*’. These are briefly summarized below, and used extensively in the following sections of this chapter.

- **Control region:** Region used either for normalising the SM background prediction (in the case of leptonic backgrounds), or to determine the uncertainty associated with the background prediction (in the case of the multi-jet data-driven method).
- **Validation region:** Region used to validate the performance of Monte Carlo simulations of the leptonic backgrounds.
- **Confirmation region:** Region used to confirm the accuracy of the full background estimation in events very similar to those in the signal regions but with slightly

smaller jet multiplicities. Confirmation regions have large background yields and small expected signal contamination (as opposed to negligible, in the validation and control regions). The ‘7j50’ and ‘6j80’ event selections are confirmation regions in this version of the analysis⁵.

3.4. Estimation of the multi-jet background

One of the major contributions to the background comes from multi-jet events where jet energy mismeasurement has led to momentum imbalance and hence ‘fake’ E_T^{miss} . The determination of this background is performed using a data-driven method proposed in Ref. [105] and adapted to the current conditions of the analysis.

The method corresponds to the wider class of data-driven background prediction techniques known as ‘ABCD’ methods. The name comes from the fact that the information to build the background prediction in a particular region of phase space, D , is taken from three other regions, A , B and C . The four regions are defined in terms of two uncorrelated variables. Figure 3.1 on page 51 shows schematically the division of the phase-space in terms of regions A , B , C and D .

The two variables chosen for the ABCD method used to predict the multi-jet background are the jet multiplicity and $\mathcal{S}_T^{\text{miss}}$. This ratio is closely related to the significance of E_T^{miss} . In events dominated by multi-jet production, the presence of genuine E_T^{miss} is small, and it is jet mismeasurement that leads to the magnitude of the missing transverse momentum not being exactly zero. Due to the sampling nature of the ATLAS calorimeter, the jet resolution is approximately proportional to $\sqrt{p_T}$ of the jet being measured for a large range of jet p_T . In particular,

$$\sigma(E_T^{\text{miss}}) = A\sqrt{\Sigma p_T^{\text{jets}}} \oplus B\sqrt{\Sigma E_T^{\text{soft}}}, \quad (3.12)$$

where ΣE_T^{soft} is the scalar sum of the energies of the objects included in the soft term of E_T^{miss} , and Σp_T^{jets} is denoted as H_T , as described earlier in Section 3.2.3. The approximate proportionality relation between the E_T^{miss} resolution and $\sqrt{H_T}$ independently of the jet multiplicity has been observed and exploited in E_T^{miss} performance studies [99, 100].

⁵The signal fraction in these regions could have been significant in the first iterations of the analysis. As SUSY phase space is getting excluded, the cross sections of the available phase space get smaller. This is why the signal fraction is now expected to be very small in these regions, and they are considered simply as confirmation regions.







(a) Maximum $|\eta|$ of leading 6 jets, Gtt -on-shell.(b) Maximum $|\eta|$ of leading 6 jets, Gtt .(c) $|\eta|$ of all jets, Gtt -on-shell.(d) $|\eta|$ of all jets, Gtt .(e) p_T of the 6th jet, Gtt .(f) p_T of the 9th jet, Gtt .

Figure 3.6.: p_T and η distributions for four Gtt and Gtt -on-shell benchmark points. All the jets have a $p_T > 20$ GeV cut and all events must pass the trigger EF_6j45_a4tc1cw_L2FS. The signal distributions are compared to a $t\bar{t}$ semileptonic sample. All distributions are normalised to unity.




Figure 3.7.: Expected and observed exclusion regions, used for the optimisation of the jet definition in the different signal regions. Representative figures are chosen to show the effects of tighter η and p_T cuts in the lower and higher multiplicity signal regions. Plots made by collaborators.




Figure 3.8.: Exclusion limits for the combination of the possible different b -jet SRs, explained in Table 3.2 and in the text. Options G and H correspond to the 9j55 region with and without the splitting in b -jet subregions. The limits are obtained with HistFitter using a semi-empirical approximation of the systematic uncertainties, as explained in the text. Plot made by a collaborator.




Figure 3.9.: Expected and observed exclusion limits obtained in the previous version of the analysis [67], which used 5.8fb^{-1} of 8 TeV data.

Therefore, ignoring the soft term for now, $\mathcal{S}_T^{\text{miss}}$ is a significance variable whose shape is expected to be independent of the number of jets in the event.

This sets the basis of the ABCD method used in the analysis to describe the multi-jet background. H_T is computed using jets with $p_T > 40$ GeV and $|\eta| < 2.8$, to maximise the correlation with the E_T^{miss} . The jets used for the multiplicity count have $|\eta| < 2.0$. Depending on the signal region, the p_T cut may be 50 or 80 GeV. To predict the multi-jet background in region D (see Figure 3.1), a *template* is formed using multi-jet events from region C , and normalised by the ratio of the regions A and B . This is why the method is sometimes referred to as the ‘template method’.

In practice, the full $\mathcal{S}_T^{\text{miss}}$ spectrum is computed in data for the two different jet multiplicities, and normalised to the higher multiplicity in the low $\mathcal{S}_T^{\text{miss}}$ region. The background template is purified from non-QCD contributions by subtracting bin-by-bin the expected yield from electroweak processes, as calculated from Monte Carlo. In particular, the processes included in the electroweak category are:

- $t\bar{t} \rightarrow l^- l^+ \nu \nu b\bar{b}$.
- $t\bar{t} \rightarrow l \nu j j b\bar{b}$.
- $W \rightarrow l \nu + \text{jets}$.
- $Z \rightarrow l l + \text{jets}$.
- $t\bar{t} + W/Z$, referred to as $t\bar{t} + V$ in the following.
- Single top.

Regions A and B have $\mathcal{S}_T^{\text{miss}} < 1.5$, and C and the signal region D have $\mathcal{S}_T^{\text{miss}} > 4.0$ $\text{GeV}^{1/2}$. These values are the same as in previous versions of the analysis.

This whole procedure can be summarised as follows:

$$\begin{aligned}
 [\text{Multijet background}]_{\mathcal{S}_T^{\text{miss}} > 4.0}^{\text{n jets}} &= \left(N_{\mathcal{S}_T^{\text{miss}} > 4.0}^{6(5) \text{ jets}} - N_{\mathcal{S}_T^{\text{miss}} > 4.0}^{\text{EWK}, 6(5) \text{ jets}} \right) \\
 &\times \frac{\left(N_{\mathcal{S}_T^{\text{miss}} < 1.5}^{\text{n jets}} - N_{\mathcal{S}_T^{\text{miss}} < 1.5}^{\text{EWK}, \text{n jets}} \right)}{\left(N_{\mathcal{S}_T^{\text{miss}} < 1.5}^{6(5) \text{ jets}} - N_{\mathcal{S}_T^{\text{miss}} < 1.5}^{\text{EWK}, 6(5) \text{ jets}} \right)}, \quad (3.13)
 \end{aligned}$$

where the low-jet-multiplicity template is taken from 5- or 6-jet events for the 80 GeV and 50 GeV regions respectively; N^{EWK} are always extracted from Monte Carlo, and

the rest from data. In previous versions of the analysis the cuts in jet multiplicity were inclusive (that is, they were simply lower bounds on the number of jets), whereas here they were made exclusive (that is, they require an exact number of jets), except in the ≥ 10 jets signal region. The template method was found to work also in this case, as will be seen in the next sections.

3.4.1. Soft energy correction

The previous discussion ignored the second term in equation 3.12 to simplify the basic ideas. This simplification was found to be accurate enough for 7 TeV data [105, 107], but the soft contributions to the hadronic activity became important with the increased levels of pile-up of the 8 TeV run (Figure 3.12). The expression for the resolution can be rewritten as

$$\sigma(E_T^{\text{miss}}) = \sqrt{H_T} \left(A \oplus B \sqrt{\frac{\Sigma E_T^{\text{soft}}}{H_T}} \right). \quad (3.14)$$

The invariance of $\mathcal{S}_T^{\text{miss}}$ under changes in jet multiplicity can be restored if $\Sigma E_T^{\text{soft}}/H_T$ can be made invariant too. Figure 3.10a shows that this is in fact not the case, so a reweighting procedure is needed in order to adjust the $\Sigma E_T^{\text{soft}}/H_T$ distribution of the low-jet-multiplicity template to that of the high-jet-multiplicity signal region. Figure 3.10b shows that the shape of ΣE_T^{soft} is not affected by changes in the number of hard jets of the event, which means that it is reasonable to assume that the hard and soft terms of E_T^{miss} are uncorrelated.

Variations in $\Sigma E_T^{\text{soft}}/H_T$ are accounted for by a weighting procedure as follows. First, histograms of $\Sigma E_T^{\text{soft}}/H_T$ are obtained for all the different jet multiplicity regions. The binning is chosen so as to reduce statistical fluctuations without dissolving the shape of the distribution, with the bin edges separated by 0.1 between $\Sigma E_T^{\text{soft}}/H_T \in [0, 1.6]$, and a final overflow bin. The weight for the multi-jet contribution in events with n jets is simply the ratio of $\Sigma E_T^{\text{soft}}/H_T$ for events with n jets divided by that for the reference jet multiplicity (5 or 6, depending on the jet p_T). The final multi-jet templates correspond to the $\mathcal{S}_T^{\text{miss}}$ distribution for events with exactly 5 or 6 jets (that is, the reference jet




Figure 3.10.: (a) $\Sigma E_T^{\text{soft}}/H_T$ distribution for different values of the jet multiplicity. (b) ΣE_T^{soft} distribution for different values of the jet multiplicity.

multiplicity), with each event weighted by

$$w = \frac{(\Sigma E_T^{\text{soft}}/H_T)^{\text{SR jet multiplicity}}}{(\Sigma E_T^{\text{soft}}/H_T)^{\text{Ref. jet multiplicity}}}, \quad (3.15)$$

where ‘SR’ means signal region and ‘Ref.’ means reference region. Figure 3.11b shows the $\mathcal{S}_T^{\text{miss}}$ distributions for different jet multiplicities after applying the reweighting, confirming that the method restores the invariance.

3.4.2. The template method in events with b -jets

B quarks decay to up or charm quarks via the weak interaction, emitting a W boson which may subsequently decay leptonically or hadronically. When the W boson decays leptonically, a neutrino is produced, which turns b quarks into sources of genuine (albeit small) missing transverse momentum. The presence of different numbers of b -jets in different events, and thus different amounts of real E_T^{miss} , breaks the invariance of the $\mathcal{S}_T^{\text{miss}}$ versus jet multiplicity, as seen in Figure 3.11a.

This problem is circumvented by deriving an individual template for each b -jet multiplicity studied in the analysis, namely exactly 0, exactly 1, or ≥ 2 b -jets, so that the multi-jet prediction for a signal region with a particular number of b -jets will come




Figure 3.11.: S_T^{miss} distribution in events with exactly 6 nominal jets and either 0, 1 or ≥ 2 b -jets, on the left; and in events with 6 or 8 nominal jets and exactly 1 b -jet. The 6-jet curve is shown before and after applying the soft term reweighting. Fixing the number of b -jets in turn fixes the amount of genuine E_T^{miss} in the event and permits the extrapolation of the S_T^{miss} shape from low to high jet multiplicities, combined with the soft term reweighting.




Figure 3.12: Recorded luminosity as a function of the mean number of interactions per crossing for the 2011 and 2012 data. Taken from [75].

from the control region with that same number of b -jets. The need to account for the differences in $\sum E_{\text{T}}^{\text{soft}}/H_{\text{T}}$ and b -jet multiplicity effectively means that a different multi-jet template is derived for each of the 13 signal regions.

Figure 3.11b shows $\mathcal{S}_{\text{T}}^{\text{miss}}$ distributions for different numbers of jets and exactly one b -jet. Distributions with any jet multiplicity cut, but with the same b -jet multiplicity requirement, are shown to coincide.

The degree to which all the approximations made in the method are correct is evaluated by looking at the agreement of the background prediction with the data in all the control and confirmation regions with low $\mathcal{S}_{\text{T}}^{\text{miss}}$ and/or low jet multiplicity. The discrepancy is taken as a systematic uncertainty, as will be explained Section 3.5.1.

The validity of the background prediction in the $\mathcal{S}_{\text{T}}^{\text{miss}} > 4.0 \text{ GeV}^{1/2}$ area is particularly important for templates with b -jets; a systematic uncertainty derived from Monte Carlo is used in that case. In addition, systematic uncertainties are derived to account for possible effects from trigger inefficiencies, pile-up dependence, and the subtraction of the leptonic backgrounds. The following sections explain each systematic uncertainty in detail.

3.4.3. Sensitivity to pile-up effects

The average number of interactions per bunch crossing grew from ~ 8 in 2011 in 7 TeV data to ~ 20 in 2012, in 8 TeV data. Although the calibration of most objects was re-optimised in 2012 taking into account these conditions, it is still important to check that the variables and techniques used in the analysis are not sensitive to pile-up effects which could appear in analysis-specific regions of phase space. Pile-up dependence studies

were performed for three aspects of the analysis: stability of the $\mathcal{S}_T^{\text{miss}}$ shape as a function of jet multiplicity, stability of the number of jets versus the number of interactions per bunch crossing, and dependence of the multi-jet background prediction on the number of vertices of the event.

The stability of the $\mathcal{S}_T^{\text{miss}}$ shape as a function of jet multiplicity has already been discussed in Section 3.4.1. The shape of the $\mathcal{S}_T^{\text{miss}}$ spectrum is found not to be affected by changes in the jet multiplicity when the appropriate reweighting scheme is applied.

The dependence of the multi-jet background prediction on the number of vertices of the event was estimated by splitting the analysis into three categories of events according to the number of primary vertices. The study is explained in more detail in Section 3.5.3.

Finally, the stability of the jet multiplicity as a function of the number of interactions per bunch crossing was also studied. Figure 3.13 shows that the average jet multiplicity in events with at least 7 jets with $p_T > 50$ GeV remains stable as a function of the average number of interactions per bunch crossing, for data as well as Monte Carlo $t\bar{t}$ and signal events.




Figure 3.13.: Stability of the average jet multiplicity versus the average number of pp collisions per bunch crossing in events with at least 7 jets with $p_T > 50$ GeV for a $Gt\bar{t}$ point (red), semileptonic $t\bar{t}$ (green, left) and data (black, right).

3.5. Systematic uncertainties of the multi-jet background prediction

3.5.1. Closure systematic uncertainty

The dominant systematic uncertainty on the template method is quantified from the degree to which the method is observed to work in regions where the signal contribution is expected to be subdominant. A systematic uncertainty is defined in terms of the different $\mathcal{S}_T^{\text{miss}}$ and jet multiplicity regions, as illustrated in Table 3.4. This uncertainty is referred to as the ‘closure’ uncertainty because it is based on the evaluation of the background prediction in regions where it is expected to agree well with the data. In general, the control regions from which the uncertainty is derived correspond to either low $\mathcal{S}_T^{\text{miss}}$ and/or low jet multiplicity.

The region below $1.5 \text{ GeV}^{1/2}$ is used to normalise the template, so it is not included in the control regions. The control regions stop at $3.5 \text{ GeV}^{1/2}$ to ensure good blinding of the signal region, since an excess above $4.0 \text{ GeV}^{1/2}$ could start showing as a positive bias at lower $\mathcal{S}_T^{\text{miss}}$ values. The confirmation regions are not actually used to build any template, but they contribute to the closure systematic uncertainty at low $\mathcal{S}_T^{\text{miss}}$ values. Since these confirmation regions were used as signal regions in the earlier 8 TeV analysis [67], the high $\mathcal{S}_T^{\text{miss}}$ values were blinded as a precaution. Higher jet multiplicities are treated in the same way.

A colour-coded illustration of the procedure can be found, together with the closure uncertainty, in Tables 3.4 and 3.5 on pages 84 and 85. Green indicates a confirmation region, red a signal region, and blue control regions.

The symmetrical systematic uncertainty on any signal region is given by the maximal deviation in any of the closure regions of the same jet multiplicity or lower, for the same b -tagging requirements. For example, the systematic uncertainty on the signal region requiring exactly 9 jets with $p_T > 50 \text{ GeV}$ and 0 b -jets corresponds to the highest value in any of the cells in the first sub-table in Table 3.4. Regions with tight cuts in jet multiplicity have large statistical uncertainties, which also result in larger closure deviations due to statistical fluctuations.

3.5.2. Heavy-flavour systematic uncertainty

The heavy-flavour systematic uncertainty is included to account for the possibility that the assumptions made to build the multi-jet template at high jet multiplicities fail due to the b -jet constraints applied to the signal regions. Three cases are identified and treated individually.

1- and 2- b -jet regions. For the 1- b -jet and ≥ 2 - b -jet regions the main issue with the template method is whether it still holds despite the presence of genuine E_T^{miss} coming from the leptonic decays of b quarks. The degree of closure of the method in the low $\mathcal{S}_T^{\text{miss}}$ region due to the presence of real E_T^{miss} is included in the standard closure systematic uncertainty described before, which is calculated for all jet multiplicities. However at high jet multiplicities no closure systematic uncertainty exists for high E_T^{miss} , as those regions are blinded. These regions are instead tested in Monte Carlo samples. A closure systematic uncertainty is derived by comparing the $\mathcal{S}_T^{\text{miss}}$ shapes at low and high jet multiplicity, using Sherpa fully hadronic $t\bar{t}$ events and MADGRAPH $t\bar{t} + V$ events, combined with their corresponding cross sections. The $t\bar{t} + V$ events are filtered at truth level to veto any events where the W bosons from the top quarks or the additional vector bosons decay to leptons or to a pair of b quarks. W bosons coming from other b quark decays are not filtered. The relative difference between the $\mathcal{S}_T^{\text{miss}} > 4 \text{ GeV}^{1/2}$ regions of 6-jet and 7-jet templates with $p_T > 50 \text{ GeV}$ and ≥ 1 b -jet is found to be 25%. This is taken as a flat systematic uncertainty for all the signal regions with b -jets.

For the 1- b -jet region the possibility that a light jet be mistagged as a b -jet is also considered. The probability of this happening is roughly 1%, so by creating a template using 10% of the 0- b -jet one the systematic uncertainty obtained is very conservative. This systematic uncertainty is combined with the one described in the previous paragraph by taking the largest of the two, and it is always found to be subdominant.

0- b -jet regions. For the 0- b -jet selections the main possible source of uncertainty is the fact that the overall number of b -jets is expected to increase as one extrapolates to higher jet multiplicities, causing an increase in the probability of a true jet being missed by the tagging algorithm. The b -tagging algorithm is used at its 70% operating point, so the probability of failing to tag a b -jet is 30%, and the probability of failing to tag 2 b -jets is $\sim 10\%$. A systematic uncertainty is then obtained by comparing the default data 0- b -jet template with one built by combining 10% of the 2- b -jet template and 90%

of the 0- b -jet one. This new template represents an estimate of the effect on the template of the increased chance of missing a b -jet pair. The exact values of the uncertainties per signal region are listed in Table 3.6 (page 86) and they range between 7% and 15%. They are obtained by first computing the integrals of the regions above $4.0 \text{ GeV}^{1/2}$ for the systematically varied sample and the nominal sample, and then taking the ratio of the two. The results of this procedure are illustrated in Figure 3.14.




Figure 3.14.: Examples of the data templates used for the calculation of the heavy flavour uncertainties in the 0 b -jet and b -jet blind regions. (a) and (b) Black line: Multi-jet template obtained from data to describe the 0- b -jet 8ej50 and 7ej80 regions respectively. Red line: Multi-jet template obtained by combining 10% of the 2- b -jet template and 90% of the 0- b -jet template for the 8ej50 and 7ej80 regions respectively. (c) Black line: Multi-jet template obtained from data to describe the inclusive 10-jet region. Red line: Multi-jet template obtained by combining 10-jet inclusive 0, 1 and ≥ 2 b -jet templates.

b -jet-agnostic regions. The 10 jet inclusive region is not split by b -jet multiplicity. In this case the heavy flavour systematic uncertainty is obtained by considering an alternative template formed by adding the 0, 1, and ≥ 2 b -jet templates separately, instead of not having any requirements on b -jets, as in the nominal case. The results are illustrated in Figure 3.14. The uncertainty is found to be 9%.

3.5.3. Soft energy and pile-up systematic uncertainties

The dependence of the template method on pile-up is tested in two different ways. To assess whether the template shape itself changes as a function of pile-up, the signal regions are binned in primary vertex multiplicity. The template prediction is found to be stable independently of the number of primary vertices in the event. Figure 3.15 shows the signal region with exactly 8 jets and no b -jets as an example. The closure systematic uncertainty already covers this effect, so these results are only used as a cross-check.

In addition, the sensitivity of the soft-energy reweighting to the p_T cut of the jets used to build the soft term is also checked. This is done to make sure that all the pile-up sensitive objects which could potentially degrade the shape of the $\mathcal{S}_T^{\text{miss}}$ template are included in the soft term used for the soft-energy reweighting. In particular, a systematic uncertainty is computed by calculating a new template including jets between 20 and 30 GeV in the soft term, instead of stopping at 20 GeV. The uncertainty is obtained by comparing the signal yields predicted by both templates above $\mathcal{S}_T^{\text{miss}} > 4.0 \text{ GeV}^{1/2}$. Table 3.6 shows that the uncertainty ranges between 3% and 21%.

3.5.4. Trigger systematic uncertainty

The 6j45 trigger is not fully efficient for the jet multiplicity values used to create the multi-jet template, as seen previously in Figure 3.4 on page 58. To evaluate the effect on the multi-jet background prediction, events where the 6th jet has a transverse momentum between 50 and 55 GeV are given a weight of $1/\epsilon = 1.2$, where ϵ is the trigger efficiency estimated from figure 3.4. The difference between the reweighted and nominal templates is taken as a systematic uncertainty. It is found to be 1% or smaller in all signal regions.




Figure 3.15.: S_T^{miss} distribution and prediction in the signal region with exactly 8 jets with $p_T > 50$ GeV, for different cuts in the number of reconstructed vertices.

0 b	$\mathcal{S}_T^{\text{miss}}$				
N. jets $p_T > 50$ GeV	[1.5,2]	[2,2.5]	[2.5,3.5]		[4.0, ∞]
9	5.5%	-14.8%	-1.2%		(blinded)
8	2.4%	1.7%	-1.5%		(blinded)
7	1.8%	0.8%	-1.9%		(blinded)
6	0.8%	-0.3%	-1.3%		1.3%

1 b	$\mathcal{S}_T^{\text{miss}}$				
N. jets $p_T > 50$ GeV	[1.5,2]	[2,2.5]	[2.5,3.5]		[4.0, ∞]
9	-3.0%	-6.7%	-15.0%		(blinded)
8	5.2%	0.0%	-10.8%		(blinded)
7	-0.0%	-0.3%	-4.8%		(blinded)
6	2.0%	-2.2%	-2.4%		-5.0%

≥ 2 b	$\mathcal{S}_T^{\text{miss}}$				
N. jets $p_T > 50$ GeV	[1.5,2]	[2,2.5]	[2.5,3.5]		[4.0, ∞]
9	-16.2%	-28.6%	13.7%		(blinded)
8	1.6%	-4.0%	-7.6%		(blinded)
7	-3.2%	0.1%	-2.8%		(blinded)
6	1.4%	-2.0%	-5.3%		-13.0%

	$\mathcal{S}_T^{\text{miss}}$				
N. jets $p_T > 50$ GeV	[1.5,2]	[2,2.5]	[2.5,3.5]		[4.0, ∞]
≥ 10	-16.3%	-12.4%	16.9%		(blinded)

Table 3.4.: Degree of closure observed in the various multi-jet control regions for events with 0, 1, ≥ 2 b -jets or no b -jet requirement (top to bottom), for the 50 GeV signal regions. The prediction is given by the sum of the prediction from the template method and the ‘leptonic’ background prediction taken straight from Monte-Carlo. The columns correspond to different $\mathcal{S}_T^{\text{miss}}$ regions in units of $\text{GeV}^{1/2}$, and the rows to different 50 GeV-jet multiplicity cuts. Green indicates a confirmation region, red a signal region, and blue a control region.

0 b	$\mathcal{S}_T^{\text{miss}}$				
N. jets $p_T > 80$ GeV	[1.5,2]	[2,2.5]	[2.5,3.5]		[4.0, ∞]
≥ 8	14.8%	15.4%	-72.6%		(blinded)
7	-0.3%	-10.2%	0.0%		(blinded)
6	0.8%	0.3%	0.4%		(blinded)
5	-1.8%	-1.8%	-0.2%		-3.1%

1 b	$\mathcal{S}_T^{\text{miss}}$				
N. jets $p_T > 80$ GeV	[1.5,2]	[2,2.5]	[2.5,3.5]		[4.0, ∞]
≥ 8	-26.0%	23.2%	-23.0%		(blinded)
7	-0.3%	2.0%	-0.4%		(blinded)
6	1.1%	1.1%	-2.5%		(blinded)
5	0.1%	1.2%	-6.2%		-25.2%

≥ 2 b	$\mathcal{S}_T^{\text{miss}}$				
N. jets $p_T > 80$ GeV	[1.5,2]	[2,2.5]	[2.5,3.5]		[4.0, ∞]
≥ 8	-10.5%	-54.2%	-54.3%		(blinded)
7	3.4%	1.8%	-19.4%		(blinded)
6	0.0%	-0.9%	-6.5%		(blinded)
5	-5.4%	1.0%	5.4%		-13.5%

Table 3.5.: As for Table 3.4 but for the 80 GeV jet regions.

Signal region			Soft term	Leptonic up	Leptonic down	HF
50 GeV	8 jets	0b	-0.03	-0.05	0.05	0.07
		1b	-0.07	-0.10	0.10	-0.25
		$\geq 2b$	-0.18	-0.20	0.20	-0.25
	9 jets	0b	-0.05	-0.05	0.05	0.11
		1b	-0.11	-0.09	0.09	-0.25
		$\geq 2b$	-0.21	-0.18	0.18	-0.25
	10 jets	–	-0.14	-0.09	0.09	0.09
80 GeV	7 jets	0b	-0.04	-0.05	0.05	0.09
		1b	-0.07	-0.07	0.07	-0.25
		$\geq 2b$	-0.14	-0.13	0.13	-0.25
	≥ 8 jets	0b	-0.05	-0.04	0.04	0.15
		1b	-0.05	-0.06	0.06	-0.25
		$\geq 2b$	-0.17	-0.12	0.12	-0.25

Table 3.6.: Soft term, leptonic and heavy flavour (HF) uncertainties of the data-driven multi-jet background.

3.5.5. *Leptonic* background subtraction

The multi-jet background prediction relies on a precise determination of the leptonic backgrounds, which have to be subtracted from the data to form the template. To assess the extent to which the template is sensitive to uncertainties on the leptonic backgrounds, these are scaled up and down by 20% to form two new templates. These templates are then compared to the nominal one to form a systematic uncertainty. The size of the variation is motivated by the average uncertainty of the $t\bar{t}$ and W backgrounds in the control regions. The systematic uncertainty obtained ranges between 5% and 18%, as seen in Table 3.6. The larger values correspond typically to the regions with several b -jets, as they have a more significant contribution from leptonic backgrounds.

3.6. Estimation of the leptonic backgrounds

Some multi-jet SM processes with leptons in their final state (mostly $t\bar{t}$, W +jets or Z +jets) may enter the signal region selections under different circumstances. These backgrounds are here referred to as *leptonic* (or electroweak) backgrounds, although to contribute to the signal regions they must have all found a way to pass the lepton veto. There are several reasons why a non-zero contribution from leptonic backgrounds is expected:

1. Hadronically decaying τ leptons from top and W decays are reconstructed as jets, so these events evade the lepton veto, which only looks for electrons or muons. In this case $E_{\text{T}}^{\text{miss}}$ comes from the tau neutrinos, and/or from leptonic decays of additional W bosons, in the $t\bar{t}$ case.
2. Electrons and muons from $t\bar{t}$ and W +jets decays may be produced out of detector acceptance. The neutrino from the W decay and the missed lepton contribute to the $E_{\text{T}}^{\text{miss}}$.
3. Electrons and muons from $t\bar{t}$ and W +jets, especially with low momenta, may be misreconstructed. The neutrino from the W decay and the missed lepton contribute to the $E_{\text{T}}^{\text{miss}}$.
4. Z bosons can decay into a neutrino pair and jets, creating a jet + $E_{\text{T}}^{\text{miss}}$ signature.




Figure 3.16.: S_T^{miss} distributions for events with exactly 7 jets with $p_T > 50$ GeV, $S_T^{\text{miss}} > 4.0$ GeV $^{1/2}$ and any number of b -jets. The points correspond to Sherpa $t\bar{t}$ and W +jets either with or without taus in the final state.

Hadronic taus are generally among the most relevant contributions, as seen in Figure 3.16. Z +jets, when the Z decays to neutrinos, is an irreducible background – but at high jet multiplicities its contribution is almost negligible. In addition, there may be contributions from $t\bar{t} + V$ or single top.

The two dominant leptonic backgrounds, $t\bar{t}$ and W +jets, are studied separately and compared with data in Section 3.6.1. Minor leptonic backgrounds are discussed in Section 3.6.2. The experimental and theoretical systematic uncertainties are described in Section 3.7.

3.6.1. Estimation of the leptonic $t\bar{t}$ and W +jets backgrounds

The $t\bar{t}$ and W +jets background descriptions are obtained from Monte Carlo simulations. Top pair events are generated using SHERPA with the CT10 PDF set, and up to four additional partons in the matrix element. The fraction of events initiated by gluon fusion relative to other processes is modified according to the reweighting scheme introduced in Ref. [67], which improves the agreement with data in $t\bar{t}$ control and validation regions. The weights are expected to correct for higher-order terms not present in the simulation. W +jets events were also simulated using SHERPA, with up to five partons in the matrix

element. In all cases, the predictions were normalised to the theoretical cross sections, calculated at next-to-next-to-leading order (NNLO) [119–122].

The $t\bar{t}$ and W +jets predictions are improved by performing two types of checks:

- Definition of validation regions where there is a significant contribution from $t\bar{t}$ and W +jets, but where the selection cuts are loose enough to have enough statistics to test the simulations. This is a qualitative inspection based on the comparison with data of different kinematic distributions.
- Definition of control regions by tightening the validation region cuts, bringing them closer to the signal region definitions. The similarity of the signal and control regions factorises out partly the theoretical and systematic uncertainties, and it also allows one to adjust the nominal predictions to the data. This is done as part of the global likelihood fit described in Section 3.8.1.

Validation Regions. The validation regions for $t\bar{t}$ and W +jets are defined as follows. To ensure orthogonality with the signal regions and enhance background contributions, one single isolated lepton (electron or muon) is required. Events are triggered using the lowest unprescaled lepton triggers (see Section 3.2.2), and the lepton is required to have $p_T > 25$ GeV and pass the isolation and quality criteria described in Section 3.2.3. Events with additional baseline leptons are vetoed. To remove any possible contribution from QCD events, $E_T^{\text{miss}} > 30$ GeV is required. There is a cut on $\mathcal{S}_T^{\text{miss}} > 2.0 \text{ GeV}^{1/2}$, which brings the validation region closer to the phase space of interest at the same time it reduces QCD contamination. Finally, an upper bound on the transverse mass built from the lepton and the missing momentum, $m_T < 120$ GeV, reduces possible contamination from new physics, which would be expected to appear at high m_T values. The transverse mass, taking the limit of zero mass for the lepton, is calculated as

$$m_T = \sqrt{2 \left(|p_T^{\text{lep}}| E_T^{\text{miss}} - p_x^{\text{lep}} E_x^{\text{miss}} - p_y^{\text{lep}} E_y^{\text{miss}} \right)}, \quad (3.16)$$

where the index ‘lep’ refers to the lepton. The full set of cuts is summarised in Table 3.7.

The kinematic distributions are studied in bins of jet and b -jet multiplicity. The jet multiplicity bins studied are 5 and 6 jets with $p_T > 50$ GeV, and 4 and 5 jets with $p_T > 80$ GeV. b -jet multiplicity bins are the same as the signal regions, namely exactly 0, exactly 1, or 2 or more b -jets. Figure 3.17 shows the m_T distributions for the electron validation region selection (apart from the m_T cut), at least 5 jets with $p_T > 50$ GeV,




Figure 3.17.: m_T distributions for events with at least 5 jets with $p_T > 50$ GeV, 0, 1 or 2 or more b -jets, and passing all the validation region cuts for electrons (left) and muons (right) apart from the m_T cut. The jet count does not include the lepton.




Figure 3.18.: E_T^{miss} distributions for events with at least 6 jets with $p_T > 50$ GeV, 0, 1 or 2 or more b -jets, and passing all the validation region cuts for electrons (left) and muons (right). The jet count includes the lepton.

Cut	Value
Validation Region	
Trigger	EF_e24vhi_medium1, EF_mu24i_tight
Lepton p_T	$> 25 \text{ GeV}$
E_T^{miss}	$> 30 \text{ GeV}$
$\mathcal{S}_T^{\text{miss}}$ (not including lepton)	$> 2.0 \text{ GeV}^{1/2}$
Transverse mass	$< 120 \text{ GeV}$
Control Region (additional criteria)	
Jet Count	Include lepton in jet count if it passes cuts
$\mathcal{S}_T^{\text{miss}}$ (including lepton)	$> 4.0 \text{ GeV}^{1/2}$

Table 3.7.: The selection criteria for the validation and control regions for the $t\bar{t}$ and W +jets backgrounds.

and the usual b -jet multiplicity bins. Higher jet multiplicities are not included here due to the large associated statistical uncertainties. Good agreement is seen in all cases. The E_T^{miss} distributions are drawn in Figure 3.18 after applying the $m_T < 120 \text{ GeV}$ cut. The variable is well described in both the electron and the muon streams.

Control Regions. Control regions are built starting from the validation regions defined above. The additional selection criteria are based around the fact that the dominant source of background entering the signal regions are hadronically decaying taus. To mimic the signal region scenario, the lepton tagged in the control regions is treated as a jet – that is, it is included in the jet count if it passes the jet p_T and η cuts. In addition, $\mathcal{S}_T^{\text{miss}} > 4 \text{ GeV}^{1/2}$ is required, where in the definition of $\mathcal{S}_T^{\text{miss}}$ the lepton is included in the computation of H_T . Finally, the events are binned in jet multiplicity and b -jet multiplicity in exactly the same way as in the signal regions. The full selection criteria are summarised in Table 3.7.

To examine the effects of treating the lepton as a jet, the relevant kinematic variables are shown in Figures 3.19 and 3.20. They show, respectively, the lepton p_T and the $\mathcal{S}_T^{\text{miss}}$ of the event, for events with at least 6 jets with $p_T > 50 \text{ GeV}$ and all the validation region cuts. The jet multiplicity requirement in these figures is low to avoid being dominated by statistical uncertainties. The jet count includes the lepton when it passes the jet cuts in both Figures 3.19 and 3.20. Similarly, the lepton is added to H_T in the $\mathcal{S}_T^{\text{miss}}$ plotted in 3.20 when it passes the jet cuts required to form H_T . The lepton p_T distribution shows a




Figure 3.19.: Lepton p_T distributions for events with at least 6 jets with $p_T > 50$ GeV, 0, 1 or 2 or more b -jets, and passing all the validation region cuts for electrons (left) and muons (right). The jet count includes the lepton.




Figure 3.20.: S_T^{miss} distributions for events with at least 6 jets with $p_T > 50$ GeV, 0, 1 or 2 or more b -jets, and passing all the validation region cuts for electrons (left) and muons (right). The jet count and the H_T in the denominator of S_T^{miss} may include the lepton.

step at 50 GeV, where the lepton starts contributing to the jet count, thus increasing the acceptance. The effect is well described by the Monte Carlo simulations, as are the $\mathcal{S}_T^{\text{miss}}$ spectra shown in Figure 3.20.

All control regions (with the same jet multiplicity requirements as the signal regions) will enter the global maximum likelihood fit explained in Section 3.8.1. Figures 3.21 and 3.22 show the jet multiplicity distributions for events passing all the control region cuts, with the electron and muon streams combined, and for 50 GeV and 80 GeV jet counts respectively. The grey shaded band is the experimental uncertainty on the Monte Carlo predictions. The exact contributions to the total experimental systematic uncertainty will be discussed in Section 3.7.1.

3.6.2. Estimation of other minor backgrounds

All other SM processes have smaller contributions to the signal regions, and are obtained directly from Monte Carlo.

Z +jets events are generated with SHERPA, with up to 5 additional partons in the matrix element calculation. $t\bar{t} + V$ events are simulated using MADGRAPH interfaced with PYTHIA. Single top events produced in association with a W boson (Wt channel) and through the s -channel are simulated using MC@NLO interfaced with HERWIG for the fragmentation and JIMMY [123] for the underlying event. t -channel single top production is simulated with ACERMC [124] interfaced with PYTHIA. The theoretical uncertainties on these processes are summarised in Section 3.7.2.

The contribution from rarer processes such as di-boson and $t\bar{t}$ +Higgs production was evaluated and found to be negligible.

3.7. Systematic uncertainties on the leptonic backgrounds

Section 3.4 discussed the systematic uncertainties related with the data-driven determination of the multi-jet background. This section will discuss the uncertainties arising from the physics objects calibration and reconstruction, and how well they are described by the simulations; and from the different choices made at the theory level when generating the

MC samples. The first class of systematics is referred to as ‘experimental’ systematics, and the second one as ‘theoretical’ systematics. Experimental and theoretical uncertainties are evaluated for both background and signal samples, but the exact components taken into account may vary for each sample.




Figure 3.21.: 50 GeV jet multiplicity distributions for events with 0, 1 or 2 or more b -jets, and passing all the control region cuts for electrons and muons combined. The grey shaded area shows the experimental uncertainties on the Monte Carlo prediction.




Figure 3.22.: 80 GeV jet multiplicity distributions for events with 0, 1 or 2 or more b -jets, and passing all the control region cuts for electrons and muons combined. The grey shaded area shows the experimental uncertainties on the Monte Carlo prediction.

3.7.1. Experimental systematic uncertainties

In most cases, experimental systematics are evaluated by re-computing the MC predictions after varying some parameters. This is the case for the jet energy scale and resolution, E_T^{miss} , b -tagging calibration and pile-up reweighting. The resulting histograms are compared with the nominal ones and relative uncertainties are derived.

Jet Energy Scale (JES). The jet calibration scheme of ATLAS is explained in detail in references [93] and [95]. The calibration procedure relies partly on Monte Carlo simulation, and partly on data. The calibration in data is performed by studying the momentum balance between a well-measured reference object and the recoiling jet(s). Consequently, the associated systematic uncertainties originate both from the MC (differences between MC generators, physics modelling, etc.) and from experimental aspects and assumptions made in the data analysis (jet flavour, trigger inefficiencies, pile-up, etc.) In addition, uncertainties from the calibration of the reference objects used in the data-based jet calibration are also propagated. The final set of JES uncertainties includes 54 individual components. In this analysis, a single asymmetric term resulting from the combination of all the individual components is considered. To ensure that the total systematic uncertainty was not artificially increased by this combination, the final exclusion limits of the analysis were re-calculated considering only the dominant components of the jet energy scale uncertainty for signal and background, and no significant difference was observed. The JES uncertainty is one of the dominant uncertainties, with a typical size of 20% in $t\bar{t}$ and W +jets, but it can be as large as 80% in statistically limited regions.

Jet Energy Resolution (JER). The JES calibration takes the average measured jet energy back to the MC value at hadron level. Since it is a correction on the average, there is still a spread of reconstructed jet energies around the true one; this spread is characterised by the jet energy resolution. The jet energy resolution is measured in data by exploiting the momentum balance in di-jet events [125]. The MC simulations are found to agree with the data within 10%. This difference is corrected for by offsetting the MC resolution; in practice, this means that the jet energy in MC events is smeared with a Gaussian with a standard deviation that accounts for the data-MC difference. The jet energy resolution uncertainties arise from the in-situ measurements. They result from propagating the JES uncertainties as well as from the additional sources of error introduced by the method, including the choice of selection cuts, the soft radiation

correction, pile-up and MC modelling. In this analysis the event yield is found to vary typically between 10% and 30% due to the JER systematic uncertainty, with larger values in regions with few events.

Soft $E_{\text{T}}^{\text{miss}}$ term. $E_{\text{T}}^{\text{miss}}$ is defined, as seen in Section 2.5, as the sum of several terms corresponding to different types of reconstructed objects. Two terms are important in this analysis: the jet term and the soft term. The experimental uncertainties related to the jet term are evaluated by propagating through the JES and JER uncertainties. The soft term is made of objects that are only relevant to $E_{\text{T}}^{\text{miss}}$, so the uncertainty arises from the evaluation of its scale and resolution as measured in 8 TeV samples using $Z \rightarrow \mu\mu$ events [99, 100]. It is propagated as an asymmetric uncertainty, which is generally found to be negligible.

b -tagging scale factors. The flavour tagging group within the ATLAS Collaboration provide scale factors and associated errors for the tagging efficiency of jets of different true flavour for several operating points of the tagging algorithms. These are determined by measuring the tagging rate of different types of events in data and then scaling the Monte Carlo tagging efficiencies to match the observations. The systematic uncertainties on the b -tagging scale factors are defined by the ATLAS b -tagging performance group. They are obtained by replacing the nominal MC sample used to calculate the scale factors with a systematically modified sample. The variations observed on the efficiencies from all the systematic uncertainties are summed in quadrature to obtain the total systematic uncertainty on the result. The main sources of systematic uncertainties are related to the modelling of the top and background processes, and they are estimated by studying different MC generators, and hadronisation and parton shower models [126–129]. Three asymmetric systematic uncertainties are derived from this, corresponding to the b -tagging efficiency, and the c - and light-mistagging rates. b -tagging uncertainties range between 5% and 20%; c - and light-mistagging uncertainties are typically smaller than 10%.

Pile-up reweighting MC events are weighted such that the distribution of $\langle\mu\rangle$ of the sample matches that of the measured data. A one-sided uncertainty is generated by scaling the weights up by 10% and measuring the difference with respect to the ‘nominal’ number. In backgrounds not statistically limited, this uncertainty is typically of the order of a few percent.

Luminosity. The uncertainty on the integrated luminosity is $\pm 2.8\%$. It is derived, following the same methodology as that detailed in Ref.[77], from a preliminary calibration of the luminosity scale derived from beam-separation scans performed in November 2012.

3.7.2. Theoretical systematic uncertainties

The theoretical uncertainties on the signal samples include variations on the renormalisation and factorisation scale, PDF uncertainties from different sets [112,113], and the value of the strong coupling constant. The final, symmetric⁶ uncertainty is obtained as an envelope of the cross sections obtained in all the previous variations, following the prescription proposed in [111].

The theoretical uncertainties on the SM background predictions are computed by comparing the baseline MC samples to other simulations. Some of the comparisons (mostly those for the most important uncertainties) are computed using samples with a full (or *fast*) simulation of the detector; however, in many cases the comparison is done ignoring reconstruction effects for simplicity, and instead the truth-level stable particles are used. The recommendations from the ATLAS SUSY group are followed to derive all the systematic uncertainties.

This section describes briefly the different contributions to the theoretical systematics for each of the backgrounds computed using MC simulations. The calculation of these uncertainties was only partially performed by the author, but it is included here for completeness, as they enter the final results presented.

$t\bar{t}$. The default generator used for $t\bar{t}$ is SHERPA, with up to 4 additional partons in the matrix element. Systematic variations considered include using an alternative generator (ALPGEN showered by HERWIG with underlying event provided by JIMMY), reweighting the sample to an alternative PDF set (CTEQ6L1 [113] instead of CT10 [42]), and varying the factorisation scale, the scale of the additional parton splittings, the CKKW matching scale [37], and the number of partons in the matrix element. Samples already generated with the variations needed were used in most cases to assess the systematic. Finally, samples with and without weighting of events initiated by gluon fusion relative to other processes are used to provide a systematic uncertainty on this procedure.

⁶The convention used by ATLAS and CMS is to take the central value defined by the positive and negative cross section uncertainties as the best estimate of the cross section. This approach also results in a symmetric uncertainty.

W +jets. The default generator for W +jets is SHERPA, with up to 5 additional partons in the matrix element. For W +jets the theoretical variations considered were the use of an alternative generator (in this case, ALPGEN +PYTHIA), different parton showers (ALPGEN +PYTHIA versus ALPGEN +HERWIG), and different number of partons in the matrix element.

$W + b/\bar{b}/b\bar{b}$ +jets. The default generator for this process is SHERPA, with massive b -quark treatment and up to 4 partons in the ME. At the time of developing the analysis, the $W + b$ cross section had been measured to a precision of $\sim 24\%$ [130]. Therefore for this background the same systematic uncertainty as for W +jets was taken, added in quadrature to an uncertainty of 24%.

$Z \rightarrow \nu\nu$ +jets. The default generator for this process is SHERPA, with up to 5 additional partons in the ME calculation. No large samples were available to obtain Z -specific systematic variations, but as this is a very similar process to W +jets, the same theoretical systematic uncertainties were used.

Single top. Single top production through the Wt and s -channels is simulated using MC@NLO +HERWIG, while the t -channel process is simulated using ACERMC interfaced with PYTHIA. The Wt sample dominates both the signal and control regions (almost 100%) such that systematics are computed only for this process. Systematic variations considered include an alternative generator (POWHEG +PYTHIA), different parton shower models (PYTHIA and HERWIG), and different methods to treat the interference with $t\bar{t}$.

$t\bar{t} + V$. This process is by default simulated with MADGRAPH. A generator systematic was computed by comparing it to ALPGEN.

For $t\bar{t}$, the gluon fusion reweighting and the parton splitting scale systematics are the dominant ones, with typical values of 25–30% each, leading to a total theoretical uncertainty on the $t\bar{t}$ background of $\approx 40\%$. The size of the theoretical systematics for the other smaller backgrounds is found to be similar to those for the $t\bar{t}$ background.

It should be noted that the analysis has been designed such that the systematic uncertainties tend to cancel out when comparing control and signal regions in the

statistical fit. Section 3.8 describes the procedure and includes the new values of the total uncertainties after performing the fit, as well as before.

3.8. Results and interpretation

Figure 3.23 shows the $\mathcal{S}_T^{\text{miss}}$ spectra for the 8 and 9 $p_T > 50$ GeV jet signal regions, and either 0, 1 or 2 or more b -jets; Figure 3.24 shows the $\mathcal{S}_T^{\text{miss}}$ spectrum in the 10-jet inclusive signal region; and finally Figure 3.25 includes the $\mathcal{S}_T^{\text{miss}}$ distributions for all the 80 GeV signal regions.


Unfortunately, no obvious excess of the data over the SM prediction can be seen in any of the signal regions. The background-only hypothesis must then be tested quantitatively and exclusion limits set on new physics phenomena. Limits are given independently of any model prediction for each signal region, and also globally for different, specific supersymmetric predictions. The author calculated most of the inputs to the fits, and the fitting algorithm was run by a collaborator.


3.8.1. Fitting procedure

In order to check the consistency of the data with the background-only and signal hypotheses, a simultaneous profile maximum likelihood fit [131] is performed in the control and signal regions.

The fit is used, first of all, to derive scale factors for each background that improve the agreement with the data. In addition, uncertainties are reduced when they contribute to both control and signal regions. The background normalisation that best describes the data is obtained by maximising a likelihood function. The likelihood contains a Poisson term which includes background normalisation factors and variable factors (‘nuisance parameters’) that account for the effect of systematic uncertainties; and a Gaussian term to make sure that the nuisance parameters do not deviate far from the nominal.


The framework used to perform the fits is HistFitter [116]. Three different types of fits are performed:

Figure 3.23.: S_T^{miss} distributions for the 8- and 9-jet 50 GeV signal region selections.



(a) ≥ 10 jets with $p_T > 50 \text{ GeV}$, $|\eta| < 2.0$, no b -jet cut

Figure 3.24.: $\mathcal{S}_T^{\text{miss}}$ distributions for the 10-jet signal region selection.

Figure 3.25.: S_T^{miss} distributions for the 7- and 8-jet 80 GeV signal region selections.

- **Background-only fit.** The background-only fit is used simply to adjust the normalisation of the dominant backgrounds. It only considers the control regions, where background contributions are enhanced; and the Poisson term in the likelihood considers exclusively the sum of events from SM backgrounds (multiplied by the appropriate normalisation factors and nuisance parameters). Control regions are only taken into account if there are at least two expected events associated with it. The uncertainties and normalisation factors obtained for each background as a result of maximising the likelihood function are then applied to the background predictions in the signal regions. Tables 3.8 and 3.9 on pages 116 and 117 are obtained by performing a background-only fit.
- **Exclusion fit.** If a significant excess is not observed after performing the background-only fit, the results can be used to derive upper limits on the number of new physics events for each signal region. In order to do this, an extra parameter must be added to the likelihood to account for the normalisation (or ‘strength’) of the signal. A test statistic is then defined to be able to ascertain how in agreement with the data a particular value of the signal strength is; in particular, it implicitly compares the likelihoods of the signal strength in question with the one that best fits the data. This is normally called a ‘log-likelihood ratio’. In general, the test statistic is calculated many times using pseudo-datasets, and its distribution is used to define a p -value. A p -value is the probability of obtaining a test statistic result at least as extreme as the one actually observed.

Two types of exclusion fits are performed: model-independent and model-dependent fits.

- *Model-independent limits* give an estimate of the minimum number of events that a new physics model should predict in a particular signal region in order to be excluded by the analysis. They can also be translated into cross section limits. The fit includes one signal region at a time, along with all the control regions; the possible signal contamination in the control regions is neglected; and it is assumed that the signal events have no systematic uncertainties. A scan is performed over possible values of the signal strength, until the p -value of the signal plus background hypothesis is smaller than 0.05 – a convention that is usually known as ‘95% confidence level exclusion’.
- *Model-dependent limits* consider the contribution from a particular new physics model to all the signal regions, and also the possible contamination of the

control regions. The procedure is similar to the one described before, but this time several signal regions are considered simultaneously (see section 3.8.1); and no scan in the signal strength is needed. In addition, signal contamination is included.

It is important to note that, within this framework, an underfluctuation in the signal regions could lead to the exclusion of the signal plus background hypothesis. The convention in ATLAS is to penalise the exclusion power for the signal plus background hypothesis if the data is inconsistent with the background, by normalising the signal plus background p -value by the p -value of the background-only hypothesis. This new number is known as the CL_s [132], and it is what is actually used to set the limits.

The model-independent and the *observed* model-dependent limits are calculated using the observed yields in the signal and control regions to calculate the likelihood. The *expected* model-dependent limits are calculated using all the regions in which no significant signal contribution is expected⁷.

Simultaneous fit for model-dependent limits

In the model-dependent fits, the seven $p_T > 50$ GeV signal regions (and similarly the six $p_T > 80$ GeV signal regions) are fitted to the background and signal predictions. The two fit results (for 50 GeV and 80 GeV regions) are combined using the best expected limit at each point in the parameter space. Correlations from sample to sample and region to region are taken into account, separately for the $p_T > 50$ GeV and $p_T > 80$ GeV signal regions. Systematic uncertainties arising from the same source are treated as fully correlated.

The fit considers several independent background components:

- $t\bar{t}$ and $W + \text{jets}$. One control region is defined for each signal region, as described in Table 3.7; the normalisation of each background component is allowed to vary freely in the fit.

⁷Although the ‘vanilla’ approach would be to only use the data in the control regions for this calculation, for the fitting algorithm control and signal regions are by construction equivalent. The main difference between regions is the relative contribution from signal and background, which is what determines whether they can be used to constrain the background; this way, *all* the data available is used in order to obtain the best possible background description. From this perspective the use of the terms ‘control’ and ‘signal’ partially loses its meaning (inherited from older statistical approaches), although the nomenclature is still useful at the time of developing the analysis.

- Less significant backgrounds ($Z + \text{jets}$, $t\bar{t} + W$, $t\bar{t} + Z$, and single top) are determined using Monte Carlo simulations. These are individually allowed to vary within their uncertainties.
- Multi-jet background. Being data-driven, it is not constrained in the fit by any control region. It is constrained in the signal regions by its uncertainties, which are described in section 3.4.

The systematic effects, described in Sections 3.4 and 3.6, are treated as nuisance parameters in the fit. For the signal, the dominant systematic effects are included in the fit; these are the jet energy scale and resolution uncertainties, the b -tagging efficiency uncertainties, and the theoretical uncertainties.

3.8.2. Fit results and interpretation

Background-only fit

Tables 3.8 and 3.9 (pages 116 and 117) summarise the background-only fit results; the number of events observed in each of the signal regions, as well as their SM background expectations, are reported before and after the background-only fit to the control regions. In each of the signal regions, agreement is found between the SM prediction and the data. The fit results are checked for stability and consistency with the background modelling based on the predictions described in Sections 3.4 and 3.6. There is no indication of a systematic mismodelling of any of the major backgrounds; the fitted values are in all cases consistent with the Monte Carlo simulation predictions.

Model-independent limits

In addition to the event yields, tables 3.8 and 3.9 also include the probability (p_0 -value) that a background-only pseudo-experiment is more signal-like than the observed data is given for each individual signal region. To obtain these p_0 -values, the fit in the signal region proceeds in the same way as the control-region-only fit, except that the number of events observed in the signal region is included as an input to the fit. Then, as explained in Section 3.8.1, an additional parameter for the non-Standard-Model signal strength, constrained to be non-negative, is fitted. The significance (σ) of the agreement between data and the SM prediction is given. No significant deviations from the SM

prediction are found. The 95% confidence level (CL) upper limit on the number of events ($N_{\text{BSM}}^{95\%}$) and the cross section times acceptance times efficiency ($\sigma_{\text{BSM,max}}^{95\%} \cdot A \cdot \epsilon$) from non-Standard-Model production are also provided, neglecting in the fit possible signal contamination in the control regions. The acceptance is defined as the fraction of the events that pass the kinematic and geometric cuts of the analysis. The efficiency is defined as the fraction of the events measured by the detector after experimental effects such as energy resolution or identification efficiency are taken into account.

Model-dependent limits

In the absence of significant discrepancies, exclusion limits at 95% CL are set in the context of several simplified supersymmetric models and an mSUGRA model, all described in Section 3.2.1. Theoretical uncertainties on the SUSY signals are estimated as described in Section 3.7.2. Combined experimental systematic uncertainties on the signal yield from the jet energy scale, resolution and b -tagging efficiency range from 15% to 25%.

All uncertainties on the SM expectation are used, including those which are correlated between signal and background (for instance jet energy scale uncertainties) and all but theoretical cross section uncertainties (PDF and scale) on the signal expectation. The resulting exclusion regions are obtained using the CL_s prescription, as explained in Section 3.8.1. A fit is performed in all the signal regions for each of the two values of $p_{\text{T}}^{\text{min}}$, and the two fit results are combined using the better expected limit per point in the parameter space, as described in Section 3.8.1.

Limits on sparticle masses quoted in the text are those from the lower edge of the 1σ signal cross section uncertainty band (from PDF and scale uncertainties) rather than the central value of the observed limit.

As shown in the rest of this section, the analysis substantially extends previously published exclusion limits on various models, from ATLAS [107, 133] and CMS [134, 135].

In the published version of the analysis [3], a second set of signal regions is considered, which add constraints on the mass of large radius jets and do not bin in number of b -jets. The exclusion limits in the paper are obtained by combining the two streams of the analysis. This is done by choosing the stream with the better expected limit at each point in parameter space. The b -jet stream of the analysis (the one presented in this thesis) typically has stronger expected exclusion limits than the large-radius jet stream,




Figure 3.26.: 95% CL exclusion curve for the simplified Gtt model. The dashed grey and solid red curves show the 95% CL expected and observed limits, respectively, including all uncertainties except the theoretical signal cross section uncertainty (PDF and scale). The shaded yellow band around the expected limit shows the $\pm 1\sigma$ result. The $\pm 1\sigma$ lines around the observed limit represent the result produced when moving the signal cross section by $\pm 1\sigma$ (as defined by the PDF and scale uncertainties). The diagonal dashed line is the kinematic limit for this decay channel.

so the exclusion limits in the paper and the ones included here are almost identical. Only in models with very large numbers of objects in the final state, and more so in boosted topologies, does the large-radius jet stream become competitive.

Gtt

The analysis result is interpreted in a simplified model that contains only a gluino octet and a neutralino $\tilde{\chi}_1^0$ within kinematic reach, and decaying with unit probability according to $\tilde{g} \rightarrow t + \bar{t} + \tilde{\chi}_1^0$, via an off-shell $\tilde{t}$. The results are presented in the $(m_{\tilde{g}}, m_{\tilde{\chi}_1^0})$ plane in Figure 3.26, which shows the combined exclusion. Within the context of this simplified model, the 95% CL exclusion bound on the gluino mass is 1.1 TeV for the lightest neutralino mass up to 350 GeV.




Figure 3.27.: 95% CL exclusion curve for the simplified Gtt -on-shell model. Other details are as $\tilde{\chi}_1^0$ in Figure 3.26.

Gtt -on-shell

In this simplified model, each pair-produced gluino decays as $\tilde{g} \rightarrow \tilde{t} + \bar{t}$; $\tilde{t} \rightarrow \tilde{\chi}_1^0 + t$. The mass of $\tilde{\chi}_1^0$ is fixed to 60 GeV. The results are presented in the $(m_{\tilde{g}}, m_{\tilde{t}})$ plane in figure 3.27 which shows the combined exclusion limits. Within the context of this simplified model, the 95% CL exclusion bound on the gluino mass is 1.15 TeV for $\tilde{t}$ masses up to 750 GeV.

One-step model

In this simplified model, each pair-produced gluino decays as $\tilde{g} \rightarrow q + \tilde{q}$ and the squark as $\tilde{q} \rightarrow q' + \tilde{\chi}_1^\pm \rightarrow q' + W + \tilde{\chi}_1^0$. Two versions of this model are evaluated, and the combined exclusion results are shown in figure 3.28. In Figure 3.28a, the fractional mass splitting, x , defined as $x = (m_{\tilde{\chi}_1^\pm} - m_{\tilde{\chi}_1^0}) / (m_{\tilde{g}} - m_{\tilde{\chi}_1^0})$, is set to 1/2, while the $\tilde{\chi}_1^0$ mass varies, and the results are shown in the $(m_{\tilde{g}}, m_{\tilde{\chi}_1^0})$ plane. In 3.28b, the $\tilde{\chi}_1^0$ mass is fixed to 60 GeV while x varies, and the results are presented in the $(m_{\tilde{g}}, x)$ plane. Gluino masses are excluded below 1 TeV at 95% CL, for $\tilde{\chi}_1^0$ masses below 200 GeV, in the case of $x = 1/2$.




Figure 3.28.: 95% CL exclusion curve for the simplified one-step model, for the two versions on the model; fixed $x = 1/2$, where $x = (m_{\tilde{\chi}_1^\pm} - m_{\tilde{\chi}_1^0}) / (m_{\tilde{g}} - m_{\tilde{\chi}_1^0})$, and varying $\tilde{\chi}_1^0$ mass on the left, and $\tilde{\chi}_1^0$ mass fixed to 60 GeV and varying x on the right. The region with gluino masses between 400 GeV and 550 GeV at small x has no signal Monte Carlo simulation. Other details are as in Figure 3.26.

Two-step model

In this simplified model, each pair-produced gluino decays as $\tilde{g} \rightarrow q + \tilde{q}$ and the squark as $\tilde{q} \rightarrow q + \tilde{\chi}_1^\pm \rightarrow q' + W + \tilde{\chi}_2^0 \rightarrow q' + W + Z + \tilde{\chi}_1^0$. The intermediate particle masses, $m_{\tilde{\chi}_1^\pm}$ and $m_{\tilde{\chi}_2^0}$, are set to $(m_{\tilde{g}} + m_{\tilde{\chi}_1^0})/2$ and $(m_{\tilde{\chi}_1^\pm} + m_{\tilde{\chi}_1^0})/2$, respectively. The results are presented in the $(m_{\tilde{g}}, m_{\tilde{\chi}_1^0})$ plane in Figure 3.29, which shows the combined exclusion limits for this model. Gluino masses are excluded below 1.1 TeV at 95% CL, for $\tilde{\chi}_1^0$ masses below 300 GeV.

mSUGRA

An mSUGRA model with parameters $\tan \beta = 30$, $A_0 = -2m_0$ and $\mu > 0$ is also used to interpret the analysis results. The exclusion limits are presented in the $(m_0, m_{1/2})$ plane in Figure 3.30. For large universal scalar mass m_0 , gluino masses smaller than 1.1 TeV are excluded at 95% CL.




Figure 3.29.: 95% CL exclusion curve for the two-step model. Other details are as in Figure 3.26.




Figure 3.30.: 95% CL exclusion curve for the mSUGRA model, generated with parameters $\tan \beta = 30$, $A_0 = -2m_0$ and $\mu > 0$. Other details are as in Figure 3.26.




Figure 3.31.: 95% CL exclusion curve for the simplified RPV model. Other details are as in Figure 3.26.

RPV model

In this simplified model, each pair-produced gluino decays as $\tilde{g} \rightarrow \tilde{t} + \bar{\tilde{t}}$; and the $\tilde{t}$ -squark decays via the R-parity- and baryon-number-violating decay $\tilde{t} \rightarrow s + b$. The results are presented in the $(m_{\tilde{g}}, m_{\tilde{t}})$ plane in Figure 3.31. Within the context of this simplified model, the 95% CL exclusion bound on the gluino mass is 900 GeV for $\tilde{t}$ -squark masses ranging from 400 GeV to 1 TeV.

3.9. Summary and outlook

This chapter presented the b -jet stream of the search for new phenomena with large jet multiplicities (from 7 to ≥ 10) and missing transverse momentum using 20.3 fb^{-1} of 8 TeV pp collision data collected by the ATLAS experiment at the LHC. The sensitivity to new physics is enhanced by considering the number of b -tagged jets, as shown in Section 3.2.6. The description of the SM backgrounds is studied in Section 3.3. A data-driven background is used for the fully-hadronic multi-jet background, as explained in Section 3.4, and Monte Carlo is used for the leptonic backgrounds, as discussed in Section 3.6.

The SM predictions are found to be consistent with the data, as shown in Section 3.8. In section 3.8.2, the results are interpreted in the context of an mSUGRA model and various simplified models resulting in final states with large jet multiplicity and E_T^{miss} . The exclusion limits substantially extend results published before the analysis was released. In a model where both of the pair-produced gluinos decay via $\tilde{g} \rightarrow t + \bar{t} + \tilde{\chi}_1^0$, gluino masses smaller than 1.1 TeV are excluded for neutralino masses below 350 GeV.

The LHC will start taking data again in 2015 at an increased centre-of-mass energy of 13 TeV. The cross section of BSM models will increase, in some cases more rapidly than that of the SM backgrounds. It might therefore be possible to extend the reach of the search only with a couple of months' worth of data. The robustness of the multi-jet analysis makes it an ideal candidate to be one of the first searches to be performed, with few core changes. More improvements could be added on a subsequent iteration; these could include a tau-jet veto, the use of a top-tagging algorithm or a more refined definition of the resolution used in the E_T^{miss} significance.

Signal region	8j50			9j50			10j50
b -jets	0	1	≥ 2	0	1	≥ 2	—
Observed events	40	44	44	5	8	7	3
Total events after fit	35 ± 4	40 ± 10	50 ± 10	3.3 ± 0.7	6.1 ± 1.7	8.0 ± 2.7	1.37 ± 0.35
Fitted $t\bar{t}$	2.7 ± 0.9	11.8 ± 3.0	23.0 ± 5.0	0.36 ± 0.18	1.5 ± 0.5	3.2 ± 1.1	$0.06^{+0.09}_{-0.06}$
Fitted W +jets	$2.0^{+2.6}_{-2.0}$	$0.62^{+0.81}_{-0.62}$	$0.20^{+0.28}_{-0.20}$	—	$0.24^{+0.65}_{-0.24}$	—	—
Fitted others	$2.9^{+1.8}_{-1.8}$	$1.7^{+1.5}_{-1.2}$	$2.8^{+2.3}_{-2.0}$	0.03 ± 0.03	0.38 ± 0.25	$0.40^{+0.60}_{-0.24}$	0.08 ± 0.08
Total events before fit	36	48	59	3.4	6.6	8.9	1.39
$t\bar{t}$ before fit	3.5	15	30	0.41	1.8	4	0.08
W +jets before fit	2.9	1.0	0.29	—	0.40	—	—
Others before fit	2.4	1.8	2.8	0.03	0.34	0.4	0.08
Multi-jets	27 ± 3	30 ± 10	26 ± 10	3.0 ± 0.6	4.0 ± 1.4	4.4 ± 2.2	1.23 ± 0.32
$N_{\text{BSM}}^{95\%}$ (exp)	16	23	26	5	7	8	4
$N_{\text{BSM}}^{95\%}$ (obs)	20	23	22	7	9	7	6
$\sigma_{\text{BSM,max}}^{95\%} \cdot A \cdot \epsilon$ (exp) [fb]	0.8	1.2	1.3	0.26	0.36	0.40	0.19
$\sigma_{\text{BSM,max}}^{95\%} \cdot A \cdot \epsilon$ (obs) [fb]	0.97	1.1	1.1	0.34	0.43	0.37	0.29
p_0	0.24	0.5	0.7	0.21	0.28	0.6	0.13
Significance (σ)	0.7	−0.02	−0.6	0.8	0.6	−0.28	1.14

Table 3.8.:

Number of observed and expected (fitted) events for the seven $p_{\text{T}}^{\text{min}} = 50$ GeV signal regions of the multi-jet + flavour stream. The category indicated by ‘others’ includes the contributions from Z + jets, $t\bar{t}$ + W , $t\bar{t}$ + Z , and single top. The table also contains for each signal region the probability, p_0 , that a background-only pseudo-experiment is more signal-like than the observed data; the significance, σ , of the agreement between data and the Standard Model prediction; the 95% CL upper limit on the number of events, $N_{\text{BSM}}^{95\%}$, originating from sources other than the Standard Model; and the corresponding cross section times acceptance times efficiency, $\sigma_{\text{BSM,max}}^{95\%} \cdot A \cdot \epsilon$.

Signal region	7j80			8j80		
b -jets	0	1	≥ 2	0	1	≥ 2
Observed events	12	17	13	2	1	3
Total fitted events	11.0 ± 2.2	17 ± 6	25 ± 10	0.9 ± 0.6	1.5 ± 0.9	3.3 ± 2.2
Fitted $t\bar{t}$	$0.00^{+0.26}_{-0.00}$	5.0 ± 4.0	12 ± 9	$0.10^{+0.14}_{-0.10}$	$0.32^{+0.67}_{-0.32}$	$1.5^{+1.9}_{-1.5}$
Fitted W +jets	$0.07^{+0.38}_{-0.07}$	$0.29^{+0.37}_{-0.29}$	—	—	—	—
Fitted others	$1.9^{+1.1}_{-0.9}$	$0.71^{+0.31}_{-0.25}$	$2.6^{+1.7}_{-1.1}$	0.02 ± 0.02	0.02 ± 0.02	$0.32^{+0.36}_{-0.21}$
Total events before fit	11.7	16	23	0.8	1.8	3.3
$t\bar{t}$ before fit	0.34	4	10	0.08	0.6	1.5
W +jets before fit	0.46	0.29	—	—	—	—
Others before fit	1.8	0.89	3.0	0.02	0.02	0.35
Multi-jets	9.1 ± 1.6	11 ± 4	10 ± 4	0.75 ± 0.56	1.2 ± 0.5	1.4 ± 1.0
$N_{\text{BSM}}^{95}(\text{exp})$	10	17	14	4	4	6
$N_{\text{BSM}}^{95}(\text{obs})$	10	16	12	5	3.5	6
$\sigma_{\text{BSM,max}}^{95\%} \cdot A \cdot \epsilon(\text{exp})$ [fb]	0.5	0.8	0.7	0.18	0.18	0.31
$\sigma_{\text{BSM,max}}^{95\%} \cdot A \cdot \epsilon(\text{obs})$ [fb]	0.5	0.8	0.6	0.24	0.17	0.31
p_0	0.5	0.6	0.8	0.19	0.6	0.5
Significance (σ)	0.05	-0.14	-1.0	0.9	-0.28	-0.06

Table 3.9.: As for table 3.8 but for the six signal regions for which $p_T^{\min} = 80$ GeV.

4

Measurement of the cross section of four-jet events

*‘How can I have done that?’ she thought.
‘I must be growing small again.’ She got up
and went to the table to measure herself by it.*

4.1. Introduction

Hadron colliders such as the LHC provide a fertile testing ground for the theory of QCD. The large cross sections for QCD processes allow experiments to probe into the corners of phase space. At high p_T , the QCD predictions take the form of the scattering of hard partons in a perturbative calculation followed by their subsequent hadronisation. In this regime experimental jet measurements are directly related to the scattering of these partons, and unfolded distributions of jet variables can be used to test the underlying theory.

The analysis presented here concerns events where four hard jets are produced. These events are of particular interest as several vertices are required in the leading-order (LO) diagrams such that even this calculation is complicated. Consequently, while the best theoretical predictions are at next-to-leading order perturbative QCD (NLO pQCD) [33] and have recently been interfaced to parton showers [32], experimental collaborations continue to use the less computationally expensive LO pQCD generators (e.g. SHERPA, MADGRAPH and ALPGEN). Sometimes, they even use generators which provide a matrix element for the hardest $2 \rightarrow 2$ process and leave the rest of the parton splittings to the parton shower (e.g. PYTHIA 8 and HERWIG++). This chapter will present an analysis designed to be able to determine which kinematic regimes are sensitive to higher order

effects, such that the simple summation of logarithms by the parton shower is insufficient. To accomplish this aim a variety of different kinematic and topological distributions are explored.

Four-jet pure QCD events form the background to many other processes at hadron colliders, and therefore the accuracy of the tools used to predict the shapes of the corresponding distributions is of general interest. While many analyses use data driven techniques to estimate the contribution from pure QCD events, these methods are tested in Monte-Carlo simulations, so the accuracy of the theoretical predictions remains important. Such is the case, for example, of the multi-jet search for new phenomena presented previously in Chapter 3.

Multi-jet cross sections have been measured previously at ATLAS [136], CMS [137], CDF [138, 139] and D0 [140, 141], though most of these studies focused on different observables and/or different energies.

4.1.1. Overview of the analysis

This chapter presents the measurement of the cross section of four-jet events using all the 8 TeV data taken by the ATLAS detector in 2012. The measurement is inclusive in the number of jets – that is, it does not veto events with additional jets, as long as the four leading ones pass the analysis cuts.

The opening sections contain all the studies performed to define the analysis strategy, including object definitions and selection cuts. The analysis uses a single type of physics object: jets. They will be defined again in Section 4.2, as there are some differences with respect to the jet definition used for the multi-jet search. The analysis targets events with at least four high- p_T jets, but the exact event selection is ruled by the trigger requirements. A combination of one- and four-jet triggers is used, and cuts are derived to ensure that the triggers are more than 99% efficient. This is explained in Section 4.3.

The main cuts resulting from this optimisation are that the leading jet p_T must be above 100 GeV, and the fourth leading jet p_T above 64 GeV. In addition, the four jets have to be well separated. No additional requirements are applied on the second and third jets, apart from the inherited $p_T > 64$ GeV cut.

The cross section is studied as a function of several variables which describe the dynamics of the event; they are introduced and described in Section 4.4. Several classes

of variables are explored, including individual jet p_{TS} , H_{T} , minimum and maximum angles between any two or three jets, invariant masses, and the total momentum of two central jets when the other two are separated by a certain rapidity interval.

The measured distributions include a detector component which must be removed in order to be able to compare them with particle-level theoretical predictions. The deconvolution or ‘unfolding’ of the detector effects relies on the robustness of the variables against experimental effects, which is tested in Section 4.4.2. The binning of the distributions must be optimised to improve the correlation between the measured distributions and those calculated at particle level. The binning for each variable is derived in Section 4.5. The unfolding process itself and the related uncertainties are described in Section 4.6. Section 4.7 describes the theoretical predictions that will be compared against the unfolded results, which will be shown in Section 4.8.

The focus of this chapter is primarily on the experimental aspects of the analysis. Some comparisons to various theoretical predictions will be made at the end, showing in which regions of phase space NLO effects become important. The conclusions will be summarised in Section 4.9.

4.1.2. Useful concepts

- **Reconstructed level.** The observables measured at the LHC are the result of the interaction of the particles produced with the material of the detector. This interaction may vary the energy (both the scale and the resolution) and the direction of particles. The MC simulations that account for all these effects are said to be at the ‘reconstructed’ or *reco* level. Real data is always, by definition, at reconstructed level, unless detector effects are subtracted using so-called **unfolding** techniques.
- **Truth level.** When experimental effects are not accounted for, simulations are said to be at *truth* level. There are different types of truth-level simulations, depending on how many effects are included: parton level, with or without underlying event, with or without hadronisation, etc. Most of the truth-level data used in this chapter includes parton showers, hadronisation and underlying event, and excludes pile-up and the interaction with the detector. The only exceptions are the HEJ [142] and BLACKHAT [33] predictions, which do not include parton showers, hadronisation or underlying event. They will be used in the Results section only (Section 4.8).

- *Truth- or reco-level variables* are built using truth- or reco-level particles, respectively.
- *Truth- or reco-level cuts* are built using truth- or reco-level observables, respectively. Reco-level cuts include acceptance cuts (for example, p_T and rapidity of the jets), trigger requirements and cleaning cuts. Truth-level cuts only restrict the truth-level acceptance.
- *Truth- or reco-level binnings* refer to the bins used to study truth- or reco-level variables, respectively.
- **Transfer matrix.** The transfer matrix is a two-dimensional histogram depicting the correspondence between the truth-level and the reco-level values of a particular observable, for a fixed set of cuts. It is built using MC events including detector reconstruction, therefore it can only be used to translate from the truth level to the reconstructed level, and not vice-versa.

4.2. Object definitions and event cleaning

All the observables used in the analysis are defined in terms of jets, as they are the natural probes of high-energy QCD. Jets are reconstructed, calibrated and selected according to the standard ATLAS prescriptions. In order to be analysed, events must pass a series of generic cleaning cuts to ensure that the detector was operating correctly and the data were measured properly. In addition, there are analysis-specific selection cuts which define the phase space of interest. The next sections will deal with these two important aspects of the analysis: Section 4.2.1 defines and characterises the main physics objects used, and Section 4.2.2 summarises the main cleaning and physics cuts. Further justification for the physics cuts will be given in Section 4.3.

4.2.1. Jets

Reconstructed-level jets are built using the anti- k_t jet algorithm [80] with radius parameter $R = 0.4$. Truth-level jets are built from stable simulated particles, as explained in Section 2.5.4 on page 43.

The full jet calibration procedure used in this analysis includes the final recommendations from the jet performance group for the data taken in 2012, which provide significant

improvements with respect to previous versions. In particular, the scheme known as ‘global sequential calibration’ (GSC) [96] reduces the flavour dependence of the JES and improves the resolution in pseudorapidity through the use of a correction for the spatial origin of jets¹.

Two minimal jet acceptance cuts are required: $p_T > 20$ GeV and $|Y| < 2.8$. All jets passing this loose selection are considered when applying the event cleaning cuts, which are described below.

4.2.2. Event cleaning and general cuts

The cuts applied in the analysis –to reject problematic data but also to select the events of interest– are the following:

Good Run List (GRL): The official GRL produced by the Data Quality team is used [143]. After selecting the luminosity blocks included in the GRL, the total integrated luminosity corresponds to 20.3 fb^{-1} .

Trigger: The triggers used in the analysis and the p_T ranges in which they are applied are listed in Table 4.1. In addition, all events must be separated by $dR_{\min}^{4j} > 0.65$; the exact definition of $dR_{\min}^{4j}$ will be given in Section 4.3.2 on page 129. The trigger is the constraining factor at the time of defining the analysis cuts; in fact, it is precisely the p_T and $dR_{\min}^{4j}$ cuts included in Table 4.1 that define the phase space of the analysis. The trigger bits are applied in both data and MC. Details on the motivations behind these cuts will be given in Section 4.3.

Vertex Requirement: The primary vertex, i.e. the vertex with the highest sum of the squared p_T values of the associated tracks in the event, must have ≥ 2 tracks.

LAr and Tile Error Requirements (data only): In accordance with the recommendations from the Data Preparation group, events with errors flagged in the LAr or Tile calorimeters are vetoed.

¹See Section 2.5.4 for a more detailed summary of how jets are reconstructed and calibrated in ATLAS.

Min. p_{T}^1 [GeV]	Max. p_{T}^1 [GeV]	Min. p_{T}^4 [GeV]	Max. p_{T}^4 [GeV]	dR_{min}^{4j}	Trigger
100	320	64	76	0.65	EF_4j45_a4tchad_L2FS
100	320	76	-	0.65	EF_4j65_a4tclcw_L2FS_delayed
320	410	64	-	0.65	EF_j280_a4tchad_delayed
410	-	64	-	0.65	EF_j360_a4tchad

Table 4.1.: Analysis trigger strategy. p_{T}^1 corresponds to the leading jet p_{T} and p_{T}^4 to the fourth leading jet p_{T} . The only jets considered are those with $|Y| < 2.8$.

Incomplete Data Events (data only): Again in accordance with the recommendations from Data Preparation, events are removed if they contain incomplete or corrupted data.

Jet Cleaning (both data and MC): Events with *bad* jets, according to the *loose* definition of bad jets, are removed. This definition includes requirements on the fraction of energy on the different layers of the calorimeter and the direction of the jets, and protects against calorimeter spikes, non-collision background or cosmics.

Tile Module Cleaning: Different corrections designed by the Jet and $E_{\text{T}}^{\text{miss}}$ performance group are applied to account for the effects from tile modules that were temporarily or permanently masked.

4.3. Trigger

Choosing the appropriate trigger for an analysis must be a combination of maximising the total luminosity and covering as much of the phase space as possible. The trigger collection used by ATLAS in 2012 included single- and multi-jet triggers, both of which are of interest for this analysis. Two main issues must be taken into account when deciding the trigger strategy: firstly, only triggers with a sufficiently high p_{T} or jet multiplicity requirement were able to record all the events; looser triggers had to be prescaled. Secondly, triggers asking for more than one jet may present counting mismatches between the trigger levels due to the coarser granularity of L1. This effectively results in selection inefficiencies at the analysis (offline) level.

At the event filter level, several four-jet triggers were available in 2012, with the lowest unscaled four-jet trigger requiring at least four jets with $p_T > 80$ GeV. Lower transverse momenta were reached by prescaled triggers and also by triggers active in the *delayed* stream of data, defined in Section 2.3 on page 39. Four-jet triggers are naturally indicated for this analysis, the only drawback being that they suffer from inefficiencies related to jet counting when jets are close to each other.

The trigger collection also included one-jet triggers, both prescaled (for low jet p_T) and unscaled (for high jet p_T). These triggers are of interest for the analysis for two reasons: firstly, because they have no jet counting problems, and therefore can be used to trigger efficiently events with high leading jet p_T ; and secondly, because they can be used as a reference to measure the efficiency of four-jet triggers. Table 4.2 shows the L1 and L2 elements of the trigger chains, and the total luminosity recorded by the single- and four-jet triggers of interest. These triggers are, respectively, EF_4j65_a4tchad_L2FS_delayed and EF_4j45_a4tchad_L2FS, henceforth referred to as ‘4j65’ and ‘4j45’; and EF_j360_a4tchad, EF_j280_a4tchad and EF_j80_a4tchad, which will be referred to as ‘j360’, ‘j280’ and ‘j80’, respectively.

Nominal four-jet	L1 Item	L2 Chain	Stream	Luminosity
4j45	L1_4J15	L2_4j15_a4TTem	JetTauEtmis	95.2 pb ⁻¹
4j65	L1_4J15	L2_4j20_a4TTem	HadDelayed	17.4 fb ⁻¹
Nominal single-jet	L1 Item	L2 Chain	Stream	Luminosity
j280	L1_J75	L2_j165_c4cchad	HadDelayed	17.4 fb ⁻¹
j360	L1_J75	L2_j165_c4cchad	JetTauEtmis	20.3 fb ⁻¹
Reference single-jet	L1 Item	L2 Chain	Stream	Luminosity
j80	L1_J30	L2_j75_c4cchad	HadDelayed	2.3 pb ⁻¹
j45	L1_J10	L2_j40_c4cchad	JetTauEtmis	0.04 pb ⁻¹

Table 4.2.: Triggers used for the analysis and reference triggers used to calculate the efficiencies of the former.

4.3.1. Single-jet triggers

Figure 4.1 shows the efficiencies of the j360 and j280 triggers as a function of the leading jet p_T . j360 is the lowest unscaled single-jet trigger. j280 is used in its ‘delayed’ version, which recorded 18 fb⁻¹ of data, as shown in Table 4.2.

Trigger	Sample	99% integral efficiency [GeV]	99% differential efficiency [GeV]
EF_j360_a4tchad	PYTHIA	380	400
	HERWIG++	380	400
	MADGRAPH	380	400
	Data	400	410
EF_j280_a4tchad	PYTHIA	300	310
	HERWIG++	300	310
	MADGRAPH	300	310
	Data	310	320
EF_j80_a4tchad	PYTHIA	90	100
	HERWIG++	90	90
	MADGRAPH	90	100
	Data	90	100

Table 4.3.: Values of the leading jet p_T at which the efficiency of the single-jet triggers reaches 99%. The values are obtained from histograms with bins of 10 GeV, and they correspond to the low edge of the bin where the efficiency is first reached. The histograms are drawn in Figures 4.1 and 4.2.

Single-jet triggers are only fully efficient for offline jets with a p_T slightly above the one required at the EF level. In particular, the minimum leading jet p_T required to use each of the triggers is chosen to correspond to the value at which 99% of differential efficiency is reached. This is a conservative requirement, since the differential efficiency of the trigger is calculated in exclusive bins of jet p_T . The integral efficiency curves reach 99% at lower values, providing a safety margin. Table 4.3 shows the values of the leading jet p_T at which 99% efficiency is reached both in the differential and integral curves, for j360 and j280. A detailed explanation of how trigger efficiency curves are computed is given in Appendix A.

Single-jet triggers other than j360 and j280 are only used in the analysis to select the sample events used to calculate trigger efficiencies. In particular, the reference triggers used are j80 and j45. j80 is used as a reference for all the nominal jet trigger efficiencies. Its own efficiency is calculated with respect to j45 in Figure 4.2, and its 99% efficiency points are included in Table 4.3. j80 is found to be fully efficient above 100 GeV of leading jet p_T –which is one of the basic analysis cuts–, which justifies its use as a reference trigger for the higher jet multiplicity triggers.




Figure 4.1.: Differential (left) and integral (right) efficiencies of the j280 and j360 triggers with respect to j80, as a function of the leading jet p_T . The bottom plots contain the ratios of the efficiency curves calculated with three different MC generators (PYTHIA, HERWIG++ and MADGRAPH) with respect to data.




Figure 4.2.: Differential (left) and integral (right) efficiencies of the j80 trigger with respect to j45, as a function of the leading jet p_T . The bottom plots contain the ratios of the efficiency curves calculated with three different MC generators (PYTHIA, HERWIG++ and MADGRAPH) with respect to data.

4.3.2. Four-jet triggers

Four-jet triggers suffer from inefficiencies from jet splitting and merging due to the coarse granularity of the L1-level trigger, as discussed earlier. For example, what appears to be one jet at L1 may correspond to two close-by offline jets; this may result in the L1 trigger not firing and the event being missed.

A useful variable to quantify the inefficiency is the minimum separation between the four leading jets in the event, defined as

$$dR_{\min}^{4j} = \min_{\substack{i,k=1,\dots,4 \\ i \neq k}} dR(\text{jet}_i, \text{jet}_k), \quad (4.1)$$

where $dR = \sqrt{\Delta\phi^2 + \Delta\eta^2}$. Jets in high energy events tend to be more collimated, and therefore the trigger inefficiency increases on average. This can be seen in Figure 4.3, which shows the integral turn-on curve of the 4j45 trigger for different bins of leading truth jet p_T .




Figure 4.3.: Efficiency of the 4j45 trigger calculated for PYTHIA events for different ranges of the truth leading jet p_T . The ranges correspond to the ones used to split the sample at the generation stage. In particular, ‘J2’ corresponds to truth anti- k_t jets with $R=0.6$ and $20 < p_T^1 < 80$ GeV, and ‘J3’ corresponds to truth anti- k_t jets with $R=0.6$ and $80 < p_T^1 < 200$ GeV.

To optimise the $dR_{\min}^{4j}$ cut needed to remove the inefficiency due to nearby jets, the differential efficiency of 4j45 and 4j65 is studied as a function of the 4th jet p_T . The optimal cut is defined as the minimum $dR_{\min}^{4j}$ value for which:

- (a) the trigger efficiency plateaus in p_T close to 100%, and
- (b) the efficiency plateau value remains stable despite increasing the $dR_{\min}^{4j}$ cut.

The optimisation is done using PYTHIA with $dR_{\min}^{4j}$ values between 0.5 and 0.8. The optimal value is found to be $dR_{\min}^{4j} = 0.65$, for both 4j45 and 4j65.




Figure 4.4.: Combination in the p_T^1 - p_T^4 plane of the four analysis triggers.

4.3.3. Combination

Single-jet triggers are efficient at high p_T^1 , whereas four-jet triggers can be efficient at lower p_T^1 . In addition, four-jet triggers require $dR_{\min}^{4j}$ and p_T^4 cuts. The combination of the triggers is done in an exclusive way, to avoid overlaps, and taking into account all these requirements.

The exact cuts in $dR_{\min}^{4j}$, p_T^1 and p_T^4 that define the four exclusive regions covered by the four triggers are shown in Table 4.1 (page 124) and depicted in Figure 4.4. The p_T^4 cuts are obtained from the 4j45 and 4j65 efficiencies, calculated only for events satisfying $dR_{\min}^{4j} > 0.65$ and $100 < p_T^1 < 320$ GeV. The efficiency curves are shown in Figure 4.5.

In the analysis, events selected with one of the prescaled triggers are reweighted by the trigger prescale.




Figure 4.5.: Differential and cumulative efficiencies of the 4j45 and 4j65 triggers versus the fourth jet p_T .

Name	Hard process	PDF	Parton shower	Underlying event	Tune
PYTHIA	PYTHIA 8	CT10	PYTHIA 8	PYTHIA 8	AU2-CT10
HERWIG++	HERWIG++	CTEQ6L1	HERWIG++	HERWIG++	UE-EE-3-CTEQ6L1
MADGRAPH	MADGRAPH	CTEQ6L1	PYTHIA 6	PYTHIA 6	AUET2B-CTEQ6L1

Table 4.4.: The different MC generators used for comparison against the data are listed, together with the parton distribution functions, parton shower algorithms, underlying event and parameter tunes.

4.4. Variables of interest

The main motivation of this analysis is to understand how well multi-jet events are described by the existing QCD calculations and approximations, and in which situations they disagree with one another and with the data. In particular, the measurement intends to explore:

- the differences between leading-order and next-to-leading order matrix elements,
- pure parton showers ($2 \rightarrow 2$ matrix element) versus matrix-element calculations matched to parton-showers ($2 \rightarrow n$ matrix element),
- how well the MC describes processes with different splitting scales,
- how well the MC describes processes with large ΔY between two jets,
- different topologies of four-jet events; for example, one jet recoiling against three or two against two.

With this end in mind, the cross section of four-jet events is measured differentially in several variables sensitive to the effects mentioned above. In addition, different regions of phase space are explored by binning some of the variables in either leading jet p_T or the invariant mass of the leading four jets.

The next sections justify the choice of the variables used in the measurement and the additional kinematic cuts. Section 4.4.1 defines the variables and includes the spectra at truth level. Section 4.4.2 discusses the correlation between the variables at the truth and reconstructed level. Three generators are compared in all cases: HERWIG++, PYTHIA and MADGRAPH. More information about the simulation is given in Table 4.4, and also in Section 4.7 on page 168, where additional theoretical predictions will be introduced.

4.4.1. Definitions and truth-level distributions

Jet kinematics. The p_T of the four leading jets (sorted in descending p_T) are studied individually. Individual transverse momenta are used as inputs to other more complex variables that will be studied later. In addition, Figure 4.6 shows that the description of the p_T distributions given by different MC generators at truth level varies significantly.




Figure 4.6.: Truth-level p_T of the 4 leading jets in events with a minimal selection of $p_T^4 > 64$ GeV, $p_T^1 > 100$ GeV and $dR_{\min}^{4j} > 0.65$ (all the cuts at truth level). The three distributions correspond to HERWIG++, MADGRAPH and PYTHIA samples. The ratios are all taken with respect to PYTHIA. The small ‘kinks’ are due to the subdivision of the sample into subsamples, which speeds up the generation process; they become negligible with the coarser binning used in the measurement.

H_T . H_T is defined as:

$$H_T = \sum_{i=1}^4 p_T^i. \quad (4.2)$$

H_T is not affected by the p_T ordering of the four leading jets. It is sensitive to the total momentum scale of the four leading jets and extends from a region where jets have a similar, low p_T scale to the high- p_T region. Figure 4.7 shows the truth-level spectrum of H_T and the discrepancies between the three generators studied.




Figure 4.7.: Truth-level H_T in events with a minimal selection of $p_T^4 > 64$ GeV, $p_T^1 > 100$ GeV and $dR_{\min}^{4j} > 0.65$ (all the cuts at truth level). The three distributions correspond to HERWIG++, MADGRAPH and PYTHIA samples. The ratios are all taken with respect to PYTHIA.

M_{jjjj} . M_{jjjj} is defined as:

$$M_{jjjj} = \sqrt{\left(\sum_{i=1}^4 E^i\right)^2 - \left(\sum_{i=1}^4 \mathbf{p}^i\right)^2}. \quad (4.3)$$

M_{jjjj} is an inclusive variable, therefore not affected by the p_T ordering of the four leading jets. It is sensitive to the angular structure between the four jets. Figure 4.8 shows that at truth level the differences are in the tail of the distribution.

$M_{jj}^{\min}/M_{jjjj}$. $M_{jj}^{\min}/M_{jjjj}$ is defined as the minimum invariant mass between any two jets in the event (out of the four with highest p_T) divided by M_{jjjj} . $M_{jj}^{\min}$ is sensitive to the smallest splitting scale in the event, whereas M_{jjjj} sets the hardest scale; therefore $M_{jj}^{\min}/M_{jjjj}$ probes the different splitting scale regimes. Different generators tend to disagree in the two extremes of the distribution, as seen in Figure 4.9. The difference gets larger in higher M_{jjjj} bins.

$\Delta\phi_{ij}^{\min}$ and $\Delta Y_{ij}^{\min}$. The two-jet angular variables are defined as:

$$\Delta\alpha_{ij}^{\min} = \min_{\substack{i,j \in [1,4] \\ i \neq j}} (|\alpha_i - \alpha_j|), \quad (4.4)$$

where $\alpha = \phi, Y$. These variables indicate the minimum angular separation in ϕ or rapidity Y between any two jets in the event. They are binned in leading jet p_T , as angles between jets tend to become smaller in high- p_T events and therefore a change in the spectrum is expected. The evolution of the shape of $\Delta\phi_{ij}^{\min}$ with p_T^1 can be seen in Figure 4.10, as well as the differences between generators. The kink in the distribution is caused by the $dR_{\min}^{4j} > 0.65$ cut. Figure 4.11 shows that $\Delta Y_{ij}^{\min}$ differs between generators in the high end of the spectrum.




Figure 4.8.: As for Figure 4.7 but for M_{jjjj} .

$\Delta\phi_{ijk}^{\min}$ and $\Delta Y_{ijk}^{\min}$. The three-jet angular variables are defined as:

$$\Delta\alpha_{ijk}^{\min} = \min_{\substack{i,j,k \in [1,4] \\ i \neq j \neq k}} (|\Delta\alpha_{ij}| + |\Delta\alpha_{jk}|) \quad (4.5)$$




Figure 4.9.: Truth-level $M_{jj}^{\min}/M_{jjj}$ in events with a minimal selection of $p_T^4 > 64$ GeV, $p_T^1 > 100$ GeV and $dR_{\min}^{4j} > 0.65$ (all the cuts at truth level), for different bins of M_{jjj} . The three distributions correspond to HERWIG++, MADGRAPH and PYTHIA samples. The ratios are all taken with respect to PYTHIA.

where $\alpha = \phi, Y$. These variables indicate the minimum angular separation in ϕ or rapidity Y between any three jets in the event. The ϕ case is particularly interesting, as the four jets are expected to approximately balance in p_T . Two extremes can be identified:

- Two jets recoiling against two jets, which gives large $\Delta\phi_{ijk}^{\min}$ values.




Figure 4.10.: Truth-level $\Delta\phi_{ij}^{\min}$ in events with a minimal selection of $p_T^4 > 64$ GeV, $p_T^1 > 100$ GeV and $dR_{\min}^{4j} > 0.65$ (all the cuts at truth level), for different bins of leading jet p_T . The three distributions correspond to HERWIG++, MADGRAPH and PYTHIA samples. The ratios are all taken with respect to PYTHIA.

- One jet recoiling against three jets, which gives small $\Delta\phi_{ijk}^{\min}$ values.

These variables are binned in leading jet p_T , which allows one to see the transition between the two populations in $\Delta\phi_{ijk}^{\min}$. Figure 4.12 shows the evolution of the shape of $\Delta\phi_{ijk}^{\min}$ as a function of p_T^1 . Figure 4.13 shows the discrepancies in $\Delta Y_{ijk}^{\min}$ between the MC generators in the different p_T^1 bins.




Figure 4.11.: As for Figure 4.10 but for $\Delta Y_{ij}^{\min}$.

$\Delta Y_{ij}^{\max}$. $\Delta Y_{ij}^{\max}$ complements the two-jet angular variables defined above. It is calculated as

$$\Delta Y_{ij}^{\max} = \max_{\substack{i,j \in [1,4] \\ i \neq j}} (|Y_i - Y_j|). \quad (4.6)$$

This variable is sensitive to large rapidity intervals in the event. It is used as an input to the $\Sigma p_T^{\text{central}}$ variables, discussed next. It is studied in bins of leading jet p_T .




Figure 4.12.: As for Figure 4.10 but for $\Delta \phi_{ijk}^{\min}$.

$\Sigma p_T^{\text{central}}$ for a given $\Delta Y_{ij}^{\text{max}}$. These variables are designed to study the kinematics of the jets when there is a rapidity interval in the event. $\Sigma p_T^{\text{central}}$ is built by first finding $\Delta Y_{ij}^{\text{max}}$ and identifying the two jets responsible for the largest rapidity interval. $\Sigma p_T^{\text{central}}$ is then simply the scalar sum of the other two jets, here denoted as “central”. By definition, the two central jets will be between the other two, although they might not necessarily be in central rapidities – for instance, they may be almost collinear with any of the forward jets. Four values of the rapidity interval are studied, each of them binned in p_T . The p_T




Figure 4.13.: As for Figure 4.10 but for $\Delta Y_{ijk}^{\text{min}}$.

bins allow one to see population changes, as the leading jet moves from being forward to being central. Figures 4.15 to 4.18 show the $\Sigma p_T^{\text{central}}$ spectrum for $\Delta Y_{ij}^{\text{max}} > 1, 2, 3, 4$. Figure 4.15 is a good example of the shape evolution as a function of leading jet p_T . In 4.15a, the lowest p_T combination is two 64 GeV jets, which corresponds to the third and fourth jets being ‘central’. This is very close to the combination 64 + 100 GeV, which




Figure 4.14.: Truth-level $\Delta Y_{ij}^{\text{max}}$ in events with a minimal selection of $p_T^4 > 64$ GeV, $p_T^1 > 100$ GeV and $dR_{\text{min}}^{4j} > 0.65$ (all the cuts at truth level), for different bins in leading jet p_T . The three distributions correspond to HERWIG++, MADGRAPH and PYTHIA samples. The ratios are all taken with respect to PYTHIA.

corresponds to the threshold for the first and fourth jets being ‘central’, so the spectrum is roughly single-peaked. In 4.15b, the leading jet p_T cut is 250 GeV, which means that the leading jet p_T can contribute to the variable only above 314 GeV. That is the point where the population transition occurs, and here it can be more easily distinguished due to the higher leading jet p_T cut. The same applies to 4.15c and 4.15d.




Figure 4.15.: Truth-level $\Sigma p_T^{\text{central}}$ in events with a minimal selection of $\Delta Y_{ij}^{\text{max}} > 1$, $p_T^4 > 64$ GeV, $p_T^1 > 100$ GeV and $dR_{\text{min}}^{4j} > 0.65$ (all the cuts at truth level), for different bins of leading jet p_T . The three distributions correspond to HERWIG++, MADGRAPH and PYTHIA samples. The ratios are all taken with respect to PYTHIA.

4.4.2. Truth- to reconstructed-level correlations

The process of reconstructing particles in the detector systematically smears the four-momenta of the original jets. In addition, the truth simulations used here do not include pile-up interactions. This means that the translation of measurements into truth-level results is not straightforward.




Figure 4.16.: As for Figure 4.15, but for $\Delta Y_{ij}^{\text{max}} > 2$.

The process of deconvolving the detector effects depends on the correlation between the reconstructed-level and the truth-level variables. The experimental effects are captured in transfer matrices for the variables of interest. The transfer matrix $\mathbf{T}$ is defined as

$$x_i^{\text{reco}} = T_{ij} x_j^{\text{truth}}, \quad (4.7)$$

where T_{ij} is an element of the transfer matrix $\mathbf{T}$, and x_i is the value of the i th bin of the distribution of the variable x .




Figure 4.17.: As for Figure 4.15, but for $\Delta Y_{ij}^{\text{max}} > 3$.

Some variables are more sensitive than others to reconstruction effects — i.e. they have larger off-diagonal elements. The robustness of the unfolding process relies partially on the variables not being over-sensitive to such reconstruction effects, or in other words, on their transfer matrices being sufficiently diagonal.

When the bin containing an event calculated using truth-level objects is different from the reconstructed-level bin, it is said that there has been a ‘migration’. Migrations in this analysis are mostly due to:




Figure 4.18.: As for Figure 4.15, but for $\Delta Y_{ij}^{\text{max}} > 4$.

- *Jet energy scale and resolution.* The jet energy scale correction rectifies the *average* jet energy, and therefore there is always a certain spread around this value that corresponds to the jet energy resolution.
- *Jet swapping.* Due to the fluctuations in the jet energy measurement, the p_T ordering of jets may change between truth jets and reco jets. For example, what at truth level was the third leading jet in p_T could be the fourth at reconstructed level if its p_T fluctuates.
- *Merging and splitting of jets.* Two truth jets have the potential to merge into one single reconstructed jet, or vice-versa, if their angular separation is small. This effect is completely removed in the analysis by the $dR_{\min}^{4j}$ cut.

Energy fluctuations in jets, and the resulting change in the p_T order, are the main causes of migrations between bins of the transfer matrix. Variables which treat the four leading jets on the same footing are the least sensitive, though there is still an effect from the fifth truth jet being counted as one of the leading four reconstructed jets about 6% of the time. Having robust variables is an advantage, but a good choice of binning and an unfolding algorithm capable of dealing with migrations are also required.

Figures 4.19 to 4.21 contain the transfer matrices for a representative set of variables built with HERWIG++, MADGRAPH and PYTHIA events, with a very fine binning that allows one to compare the shapes of the correlations. The MADGRAPH sample has fewer events, which is the major cause of the apparent differences between this and the other two generators. It can be seen that H_T and M_{jjjj} have an identical linear correlation coefficient for all three generators, with $r = 0.99$. The linear correlation coefficient is a useful way to characterise the strength of the linear dependence between two variables X and Y . It is defined as:

$$r = \frac{\sum_{i=1}^n (X_i - \bar{X})(Y_i - \bar{Y})}{\sqrt{\sum_{i=1}^n (X_i - \bar{X})^2} \sqrt{\sum_{i=1}^n (Y_i - \bar{Y})^2}}, \quad (4.8)$$

where n is the number of bins. It can take values between $+1$ and -1 inclusive, where 1 is the maximum positive correlation, 0 corresponds to the case of no correlation, and -1 represents the maximum anticorrelation.

Jet swapping is most relevant in angular variables, since interchanging two jets with very similar p_T but at different regions of the detector can result in very different angles. This is the origin of the structure seen for $\Delta\phi_{ijk}^{\min}$ in Figure 4.21. For example, a 2 vs. 2 configuration can easily turn into a 1 vs. 3 configuration if the fourth and fifth jets

are produced back-to-back and have very similar p_T ; if the fourth and fifth jets swap, $\Delta\phi_{ijk}^{\min} \sim \pi$ at truth level would turn into $\Delta\phi_{ijk}^{\min} \sim 0$ at reconstructed level. However, there is an order of magnitude of difference between the bins along the diagonal and the off-diagonal ones, so the truth–reco correlation remains strong at around $r = 0.95$. Figure 4.22 shows the transfer matrices used in the unfolding of $\Delta\phi_{ijk}^{\min}$ with the binning derived in Section 4.5, for different cuts in p_T^1 . It can be seen that the good correlation holds, and in fact improves, in the higher p_T^1 bins.

In conclusion, the large linear correlation coefficient and the agreement between the different MC generators tested supports the robustness of the variables proposed for the measurement.




Figure 4.19.: Transfer matrices of H_T in events with $p_T^4 > 64$ GeV, $p_T^1 > 100$ GeV and $dR_{\min}^{4j} > 0.65$ (at both truth and reco level). The matrices are normalised globally to a fixed, arbitrary value. The three histograms correspond (left to right) to HERWIG++, MADGRAPH and PYTHIA samples.

4.5. Binning

The reconstruction of physical observables in the detector may result in the migration of events between bins, which complicates the unfolding of the detector effects. Larger bins mitigate the problem, but also dilute the spectrum of the variable of interest, so a compromise between the two extremes has to be reached.




Figure 4.20.: Transfer matrices of M_{jjjj} in events with $p_T^4 > 64$ GeV, $p_T^1 > 100$ GeV and $dR_{\min}^{4j} > 0.65$ (at both truth and reco level). The matrices are normalised globally to a fixed, arbitrary value. The three histograms correspond (left to right) to HERWIG++, MADGRAPH and PYTHIA samples.




Figure 4.21.: Transfer matrices of $\Delta\phi_{ijk}^{\min}$ in events with $p_T^4 > 64$ GeV, $p_T^1 > 100$ GeV and $dR_{\min}^{4j} > 0.65$ (at both truth and reco level). The matrices are normalised globally to a fixed, arbitrary value. The three histograms correspond (left to right) to HERWIG++, MADGRAPH and PYTHIA samples.

The optimal binning is determined by considering the purity and stability of the bin contents. Purity is defined as:

$$P = \left[\frac{\text{N. events that stay in same bin}}{\text{N. events that pass reco cuts}} \right]_{\text{Passed truth}}. \quad (4.9)$$




Figure 4.22.: Truth-reco transfer matrices of $\Delta\phi_{ijk}^{\min}$ in events with $p_T^4 > 64$ GeV, $p_T^1 > 100$ GeV and $dR_{\min}^{4j} > 0.65$, and different additional p_T^1 requirements (at both truth and reco level). The matrices are normalised globally to a fixed, arbitrary value. The histograms use the final binning used for the unfolding.

Stability is defined as:

$$S = \left[\frac{\text{N. events that stay in same bin}}{\text{N. events that pass truth cuts}} \right]_{\text{Passed reco}}. \quad (4.10)$$

The purity is a property of the reconstructed-level bin contents, while the stability is a property of the truth-level bin contents.

For some variables, a constant binning is found to give a good, flat purity and stability: this is the case for all the event-level angular variables. In these cases the (constant) bin widths are adjusted by hand. For all the other variables, the binning is obtained in an automated manner, following the method described in Appendix B. The goal of the algorithm is to derive a binning that yields an approximately flat purity with an average of $P = 80$ %. In addition, the method imposes the conditions that the statistical uncertainty predicted in data be smaller than 10%, and that the purity in any of the bins be 70% or larger. The procedure is applied to derive the binning of all but the angular variables.

The bins of angular variables are chosen by hand so that the average purity is approximately 80%. As before, bins are merged if their purity is below 70 % or if their predicted yield for 20.3 fb^{-1} of data is smaller than 100 events.

Figure 4.23 shows the purities obtained for H_T using the automated method and for $\Delta Y_{ij}^{\max}$ using fixed-width bins. These binnings are preliminary, as they are subject to rounding and final statistical corrections, as will be explained in the next Section.




Figure 4.23.: Purity and stability of H_T and $\Delta Y_{ij}^{\max}$.

4.5.1. Statistical uncertainties and rounding

Some regions of phase space have fewer simulated events, such that the constraint on the 10% statistical uncertainty on the data prediction may not be accurate. Some bins must then be merged by hand. After studying the statistical uncertainties of real data, it is found that $\Delta Y_{ij}^{\min}$ and $\Delta Y_{ijk}^{\min}$ do indeed exceed the 10% statistical uncertainty limit in data for large rapidity separations, so they are rebinned accordingly.

Finally, the exact values of the bin edges are rounded for simplicity. For angular variables and $M_{jj}^{\min}/M_{jjjj}$, two decimal places are taken. For all the variables based on p_T , the values are rounded to the closest multiple of 5 GeV.

The complete set of bin edges can be found in Table 4.5. Note that the optimisation of the binning is done only for the distributions with the baseline analysis cuts; distributions with tighter p_T^1 and M_{jjjj} cuts inherit the binning from them.

Variable	Bin edges
H_T	290, 485, 705, 950, 1225, 1530, 1875, 2265, 2705, 3205, 7000
p_T^1	100, 155, 235, 325, 420, 530, 650, 790, 950, 1130, 1350, 1630, 4000
p_T^2	64, 145, 255, 385, 535, 715, 930, 1175, 1460, 3000
p_T^3	64, 120, 205, 305, 425, 570, 740, 2000
p_T^4	64, 85, 135, 190, 255, 330, 415, 515, 1500
M_{jjjj}	100, 545, 735, 935, 1150, 1375, 1620, 1880, 2160, 2460, 2780, 3115, 3460, 3810, 7000
$M_{jj}^{\min}/M_{jjjj}$	0.0, 0.04, 0.08, 0.13, 0.17, 0.23, 0.3, 0.4
$\Delta Y_{ij}^{\max} > 1, \Sigma p_T^{\text{central}}$	120, 170, 240, 315, 395, 480, 575, 680, 795, 930, 1085, 1260, 1465, 1705, 1980, 2300, 3075, 5000
$\Delta Y_{ij}^{\max} > 2, \Sigma p_T^{\text{central}}$	120, 185, 270, 365, 465, 575, 700, 845, 1005, 1195, 1410, 1665, 1960, 2305, 5000
$\Delta Y_{ij}^{\max} > 3, \Sigma p_T^{\text{central}}$	120, 190, 285, 385, 490, 605, 735, 880, 1040, 1225, 1430, 1655, 1905, 2790, 5000
$\Delta Y_{ij}^{\max} > 4, \Sigma p_T^{\text{central}}$	120, 190, 285, 385, 490, 605, 730, 865, 1010, 1170, 1340, 1525, 2165, 5000
$\Delta\phi_{ij}^{\min}$	0.0, 0.15, 0.3, 0.45, 0.6, 0.75, 0.9, 1.05, 1.2, 1.35, 2.5
$\Delta\phi_{ijk}^{\min}$	0.0, 0.25, 0.5, 0.75, 1.0, 1.25, 1.5, 1.75, 2.0, 2.25, 2.5, 2.75, 3.0, 3.5
$\Delta Y_{ij}^{\min}$	0.0, 0.15, 0.3, 0.45, 0.6, 0.75, 0.9, 1.05, 1.2, 1.35, 1.5, 2.5
$\Delta Y_{ijk}^{\min}$	0.0, 0.25, 0.5, 0.75, 1.0, 1.25, 1.5, 1.75, 2.0, 2.25, 2.5, 2.75, 3.0, 3.25, 3.5, 3.75, 4.0, 5.0
$\Delta Y_{ij}^{\max}$	0.0, 0.4, 0.8, 1.2, 1.6, 2.0, 2.4, 2.8, 3.2, 3.6, 4.0, 4.4, 4.8, 5.2, 5.6

Table 4.5.: Bin edges used for the different variables of interest. The procedure used to define them is described in Section 4.5 and Appendix B.

4.6. Unfolding of detector effects

The effects of the object reconstruction in the detector on the measurement are deconvoluted by applying an unfolding algorithm on the data. The data are unfolded to the truth jet level, using as a reference the truth-level MC simulations which include the underlying event, parton showering and hadronisation.

The next subsections will give a description of the method, the transfer matrix, and the basic tests performed to ensure that the method implementation is correct. Section 4.6.2 explains the studies performed to choose the MC sample used to construct the unfolding matrix. The sample chosen, generated with PYTHIA, is the one with the largest number of generated events, and it is shown to be compatible in terms of the jet response with other types of –in principle more accurate– predictions. Beyond this, to check that there is no bias in the unfolding process due to shape differences between data and MC, a data-driven test is performed in Section 4.6.3. The test shows that no significant bias exists for any of the variables of interest. In addition, a systematic uncertainty is derived to account for differences in the MC description of the efficiency with which the events that populate the transfer matrix are selected. This will be done in Section 4.6.4. Section 4.6.5 explains how the experimental uncertainties are propagated through the unfolding; this is also an important input to the optimisation of the parameters used in the setup of the unfolding algorithm, which is explained in Section 4.6.6.

4.6.1. Method description

The algorithm used for the unfolding is known as Bayesian Iterative [144,145], and the version used is the one implemented in the RooUnfold package [146]. The goal of the algorithm is to build an unfolding matrix. The unfolding matrix is the inverse of the transfer matrix defined in Section 4.4.2. The transfer matrix can be obtained from Monte Carlo, but the inversion process is mathematically unstable. The Bayesian Iterative method uses a probabilistic approach to build the unfolding matrix iteratively.

Truth- and reconstructed-level observables can be thought of in this context as *causes* and *effects*, respectively. This means that the unfolding matrix must contain the probabilities for each of the *causes given the observed effects*. These probabilities are calculated using Bayes' theorem: the likelihoods for the *effects given the causes* are

obtained from Monte Carlo simulations, and the prior probabilities are improved with each iteration of the algorithm.

The process can be expressed mathematically in a compact way. Let C_i ($i = 1, \dots, n_C$) be a set of independent causes that can cause a particular effect E_j . The goal of the unfolding procedure is to determine the number of events $\hat{n}(C_i)$ due to each of the causes (C_i), which can be calculated as

$$\hat{n}(C_i) = \frac{1}{\epsilon_i} \sum_{j=1}^{n_E} n(E_j) P(C_i|E_j), \epsilon_i \neq 0, \quad (4.11)$$

where ϵ_i is the efficiency for C_i to produce an effect, and $n(E_j)$ is the number of events in the effect bin j .

Assuming that we know the initial probability $P_0(C_i)$ and the conditional probability of C_i to produce E_j , $P(C_i|E_j)$, then the probability that an *observed* effect E_j was caused by C_i is

$$P(C_i|E_j) = \frac{P(E_j|C_i)P_0(C_i)}{\sum_{l=1}^{n_C} P(E_j|C_l)P_0(C_l)}. \quad (4.12)$$

The matrix $P(E_j|C_i)$ (usually referred to as transfer matrix) can be obtained from Monte Carlo, as well as the initial probabilities $P_0(C_i)$ and the efficiencies ϵ_i . Crucially, $P_0(C_i)$ can be improved iteratively by using the estimators $\hat{n}(C_i)$ obtained from applying equation 4.11,

$$\hat{P}(C_i) = \frac{\hat{n}(C_i)}{\sum_{i=1}^{n_C} \hat{n}(C_i)}. \quad (4.13)$$

Within the Bayesian framework, effects cannot exist without a prior cause, so background events –or *fakes*– must be added to the transfer matrix as coming from a single, extra, unknown cause. On the other hand, causes need not have a measured effect, and this is accounted for by the reweighting of each bin of the unfolded distribution by the inverse of the corresponding efficiency.

Transfer matrices obtained from Monte Carlo are essential ingredients to the unfolding process. They contain the information about jet smearing due to reconstruction effects. One transfer matrix is defined for each variable of interest and each p_T^1 and M_{jjjj} bin. Its two axes correspond to the truth and reconstructed values of the variable in question. The transfer matrix (and other histograms) used in the unfolding are provided by the

linear, cross section-weighted combination of the matrices obtained from the individual MC sub-samples.

The number of iterations of the algorithm has to be optimised, and it typically varies between 1 and 3. A lower number of iterations results in a higher dependence on the MC, whereas more iterations give larger statistical uncertainties. The optimisation of the number of iterations will be described in Section 4.6.6 on page 167.

No spatial matching is applied on the jets used to build the transfer matrix. However, the events must pass the analysis cuts at both reconstructed and truth level. This means that events must satisfy:

- Reconstructed-level cuts: as explained in Section 4.2.2 on page 123.
- Corresponding truth-level cuts: $p_T^4 > 64$ GeV, $p_T^1 > 100$ GeV, $dR_{\min}^{4j} > 0.65$.

Events that pass the reconstructed cuts but not the truth-level cuts, introduced earlier by the name of *fakes*, are included in the transfer matrix, forming an extra column for events with no well defined truth cause. Events that pass the truth cuts but not the reconstructed cuts are translated into bin-by-bin inefficiency corrections. The fake rate and inefficiency curves for two representative variables (H_T and $\Delta\phi_{ij}^{\min}$) are shown in Figures 4.24 and 4.25. The differences between Monte Carlo simulations are taken as a systematic uncertainty, as will be explained in Section 4.6.4.




Figure 4.24.: Left: Fraction of events passing analysis cuts at reconstructed level which also pass them at truth level, as a function of H_T . Right: Fraction of events passing analysis cuts at truth level which also pass them at reconstructed level, as a function of H_T . The selection cuts are $p_T^4 > 64$ GeV, $p_T^1 > 100$ GeV and $dR_{\min}^{4j} > 0.65$. The grey band represents the systematic uncertainty derived from the difference between the ratio curves, as will be explained in Section 4.6.4.




Figure 4.25.: Left: Fraction of events passing analysis cuts at reconstructed level which also pass them at truth level, as a function of $\Delta\phi_{ij}^{\min}$. Right: Fraction of events passing analysis cuts at reconstructed level which also pass them at truth level, as a function of $\Delta\phi_{ij}^{\min}$. The selection cuts are $p_T^4 > 64$ GeV, $p_T^1 > 700$ GeV and $dR_{\min}^{4j} > 0.65$. The differences between Monte Carlo simulations are taken as a systematic uncertainty, as will be explained in Section 4.6.4.

Technical closure In order to test the implementation of the unfolding algorithm, the reconstructed-level spectrum of the same MC sample used to build the unfolding matrix is itself unfolded. The resulting unfolded distributions are found to be identical to the truth-level spectra, as should be the case by construction if the implementation of the algorithm is correct.

4.6.2. Choosing the nominal Monte Carlo

Ideally, the transfer matrix should represent the jet smearing due to detector effects, and as such should be independent of the MC sample used for the training. The same can be said about the efficiency corrections. However, there are some reasons why this may in fact not be the case: for example, different numbers of events generated in different regions of phase space, or differences in the jet response due to the flavour composition of the sample, could lead to different transfer matrices and selection efficiencies between generators.

Three different MC samples are available with a full simulation of the ATLAS detector, as shown in Table 4.4 on page 132. Of those, the PYTHIA sample has 10 times more events generated than MADGRAPH and 4 times more than HERWIG++. However, MADGRAPH has a more sophisticated tree-level calculation with up to 4 jets in the matrix element.

Three factors determine the choice of unfolding MC sample:

- *Number of generated events.* Having a large number of generated events is an advantage, since it reduces the statistical uncertainty of the unfolded spectrum. Of the three samples mentioned above, PYTHIA is the largest one –in number of events– by an order of magnitude.
- *Detector response.* The transfer matrices for HERWIG++ and PYTHIA are compatible, as shown in Section 4.4.2. The comparison with the MADGRAPH sample is harder due to the smaller number of generated events. Instead, it is useful to compare the jet response of the three samples. The jet response is used to derive the jet energy scale, and it is defined as:

$$R = \frac{p_{\text{T}}^{\text{jet, reco}}}{p_{\text{T}}^{\text{jet, truth}}}. \quad (4.14)$$

Figure 4.27 shows the jet response as a function of the true jet p_{T} for all truth jets matched to reconstructed jets with $p_{\text{T}}^{\text{truth}} > 20$ GeV and $|Y^{\text{truth}}| < 2.8$. The events are preselected with the analysis truth-level cuts (i.e. $p_{\text{T}}^4 > 64$ GeV, $p_{\text{T}}^1 > 100$ GeV and $dR_{\text{min}}^{4j} > 0.65$), and in addition the four leading reconstructed jets must be matched to some truth jet, as follows. Truth jets are matched to reconstructed jets by looping over the latter and finding the closest truth jet within a cone of radius $R = 0.4$. The bias in the jet response seen at low truth jet p_{T} is due to the event preselection; the values are only meaningful above approximately 70 GeV.

After applying the jet calibration, the average jet response is very close to unity. However, the calorimeter response is dependent on the flavour of the partons produced in the sample –that is, whether they are quarks or gluons. This dependence is due to differences in the fragmentation and showering properties of quark and gluon jets. Effectively, this means that the jet response varies slightly between MC generators. The difference is the source of one of the components of the jet energy scale systematic uncertainty. Figure 4.27 shows that MADGRAPH agrees with PYTHIA within 1%, while HERWIG shows slightly larger discrepancies at low p_{T} .

- *Shape differences.* Figure 4.26 shows the H_{T} and $\Delta\phi_{ij}^{\text{min}}$ spectra for PYTHIA, HERWIG++ and MADGRAPH, compared with data before unfolding. All three simulations achieve a similar level of accuracy in their description of the data. The remaining shape differences could still bias the result of the unfolding, if the

algorithm was not robust or the binning was suboptimal. A data-driven test will be developed in Section 4.6.3. The results indicate that there is no bias due to the shape of the distributions in any of the variables of interest.




Figure 4.26.: Reconstructed-level distributions of p_T^1 and $\Delta\phi_{ij}^{\min}$ for three different MC generators compared against data. The selection cuts are $p_T^4 > 64$ GeV, $p_T^1 > 100$ GeV and $dR_{\min}^{4j} > 0.65$. The distributions are normalised to data in the region 500 GeV to 1000 GeV. The differences between generators in the low- p_T region cause the overall normalisation difference in $\Delta\phi_{ij}^{\min}$.

Consequently, the sample chosen to train the unfolding algorithm was PYTHIA, as it has a large number of available events, is compatible with the other generators in terms of the jet response, and the differences with respect to the data in terms of spectrum shapes –which are of similar size as the other generators tested– do not result in any unfolding bias. The next section describes the data-driven test, designed to derive a systematic uncertainty for shape biases in the unfolding procedure. The small differences in the jet response may also affect the efficiency with which the events used to build the transfer matrix are selected. This effect is not included in the data-driven shape systematic uncertainty, so a dedicated uncertainty will be derived in Section 4.6.4.

4.6.3. Shape systematic uncertainty

Since the differences between MC generators due to variations in the jet response are covered by the JES uncertainty, the main systematic uncertainty for the unfolding




Figure 4.27.: Average jet response versus truth jet p_T for PYTHIA, HERWIG and MADGRAPH events. The selection cuts are $p_T^4 > 64$ GeV, $p_T^1 > 100$ GeV and $dR_{\min}^{4j} > 0.65$. Truth and reconstructed jets are matched within a cone of $dR < 0.4$.

procedure itself is due to shape differences between the data and the reconstructed MC. A data-driven test is performed to assess this uncertainty. The test has been used before in other analyses [147], with a slightly different implementation.

The first step in the test is to reweight the MC to reproduce the shape of the data at the reconstructed level. The systematic uncertainty is defined from the ratio of the reweighted+unfolded MC distribution to the corresponding reweighted truth. In order to test the effects of the unfolding matrix only, some of the distributions have to be reweighted by the appropriate fake rates or efficiency curves taken from MC. The distributions obtained after applying these efficiencies, or equivalently, the distributions that only include events that pass both the truth and reconstructed cuts, are called here *partial truth* or *partial reco*, for simplicity. Figure 4.28 shows a schematic of the different elements used in the test, excluding the fake- and inefficiency-reweighting steps for simplicity. The MC sample used for the test is PYTHIA.

The full procedure is as follows:

- **Fake reweighting.** The fake rate is computed in MC, and applied to the data spectrum, yielding the ‘partial data’ spectrum.




Figure 4.28.: Schematic of the elements in the data-driven method used to determine the unfolding shape systematic. The fake- and inefficiency-reweighting steps have been omitted, such that all the spectra shown are already at the ‘partial’ level. The numbers are linked to the bullet points in the main body of Section 4.6.3.

- **Inefficiency reweighting.** The unfolded data distribution (obtained with the full, standard unfolding procedure with four iterations of the algorithm²) is weighted by the inefficiency due to truth events not passing reconstructed-level cuts. The resulting distribution is the ‘partial unfolded data’ spectrum.
- **Non-parametric regression reweighting.** [Step 1 in figure] The goal is to reweight each bin of the partial truth MC spectrum such that, after folding it, the reconstructed spectrum matches the partial data. The weights are taken from a smooth curve, to reduce the sensitivity to statistical fluctuations. To simplify the process of finding the right curve, it is hypothesised that such a curve will also take the partial truth spectrum to match the partial unfolded data spectrum. The reweighting curve is constructed as a Gaussian kernel regression³ of the ratio of the partial unfolded data and the partial truth MC. The standard deviation of the Gaussian kernel is set to $1/5$ of the range of the variable.

²A high number of iterations was chosen to maximise the shape accuracy in the unfolding. Since this unfolded spectrum is only used as a reference to derive the reweighting function, it is seen to provide good agreement at reconstructed level, and does not get propagated through the test in any way, the number of iterations does not need to be optimised.

³A kernel regression estimates the continuous distribution of a variable from a set of data points. The value of the distribution at each point is obtained from the weighted contributions of the data points; in this study, the weighting function – or kernel function – is a Gaussian.

- **Folding.** [Step 2 in figure] The folding matrix is built from the transfer matrix by fluctuating the contents of each bin with a Poisson weight centred at 1. In addition, the matrix is normalised to preserve unitarity.
- **Reweighting check.** [Step 3 in figure] The partial truth MC spectrum is reweighted, folded, and compared against the partial data. Good agreement is seen for all the variables of interest, apart from occasional fluctuations in the last bin.
- **Unfolding.** [Step 4 in figure] The folded-reweighted MC distribution is then *unfolded* using the Bayesian Iterative algorithm.
- **Derivation of the systematic uncertainty.** The unfolded MC is compared to the reweighted truth, and the ratio of the two distributions yields the systematic uncertainty. A statistical uncertainty on this systematic uncertainty is derived by using the Bootstrap Method (see Section 4.6.5) with 100 replicas [148–150]. The statistical uncertainty is used to determine whether the systematic is significant or not.

Figure 4.29 shows the multiple distributions obtained in the intermediate steps using H_T and $\Delta\phi_{ij}^{\min}$ as examples. The brown histogram shows that the folding procedure reproduces the reconstructed spectrum; and the black line is the reweighting function. The blue points correspond to the ratio between the folded-reweighted MC and the partial data spectrum, proving that the reweighting works correctly. The systematic uncertainty is derived from the ratios in Figure 4.30. To quantify the significance of the deviations from 1 seen in the ratio, it is useful to define a significance variable σ . If the ratio of two bins is r_i , and the statistical error on that value is e_i , then the significance σ_i of the deviation of r_i from 1 is defined as:

$$\sigma_i = \frac{r_i - 1}{e_i}. \quad (4.15)$$

A bin i is significantly biased if $\sigma_i > 2$. No 2σ deviations are found in any bins, therefore the shape systematic uncertainty is dropped from the calculation of the total systematic uncertainty.

The test is also used in the optimisation of the number of iterations of the unfolding algorithm in Section 4.6.6.

4.6.4. Selection efficiency systematic uncertainty

The data-driven systematic derived in Section 4.6.3 does not cover effects due to the different MC descriptions of the selection inefficiencies that occur at the time of building the transfer matrix, shown in Figures 4.24 and 4.25 on page 154. A new systematic uncertainty is derived to account for differences in these selection efficiency curves as obtained from two MC samples: HERWIG and PYTHIA. The differences are due in part




Figure 4.29.: Distributions used to derive the data-driven unfolding shape systematic for H_T and $\Delta\phi_{ij}^{\min}$ with $p_T^1 > 100$ GeV: the ratio of the folded and the reconstructed spectra (brown squares), the ratio between PYTHIA and data before the reweighting (green crosses), the ratio between PYTHIA and data after the reweighting (filled blue dots), and the reweighting function (black line, obtained as a regression of the open blue dots).




Figure 4.30.: Comparison of the unfolded, reweighted PYTHIA distribution and the original reweighted truth distribution. The ratio yields the shape systematic for H_T (left) and $\Delta\phi_{ij}^{\min}$ (right).

to the different flavour composition of the samples, which results in a residual difference in the jet response (as seen in Section 4.6.3). Other differences between the samples, listed in Table 4.4, account for the rest.

The systematic uncertainty is obtained as follows.

- The following efficiency curves are obtained, using the final binning, for each of the variables and selection criteria, and the three MC samples listed above:
 - Efficiency of selecting events passing truth- and reco-cuts, with respect to those passing only truth-cuts (*inefficiencies*). The efficiency is studied with respect to the truth-level variable.
 - Efficiency of selecting events passing truth- and reco-cuts, with respect to those passing only reco-cuts (*fakes*). The efficiency is studied with respect to the reco-level variable.
- The ratios of the HERWIG curves are calculated with respect to the corresponding PYTHIA curves, and the statistical uncertainty is propagated and symmetrised.
- The final systematic uncertainty corresponds to the envelope of the HERWIG/PYTHIA ratios corresponding to inefficiencies and fakes. The shape of the uncertainty is smoothed with a three-step procedure: first, the bins of the uncertainty distribution are combined if their statistical uncertainty is larger than 30% or the significance is smaller than 2σ . Second, the resulting uncertainty distribution is translated back into the original bins. Third, a gaussian kernel smoothing is applied, with a standard deviation of twice the bin width.

Figure 4.31 shows the resulting systematic uncertainty as a function of H_T and $\Delta\phi_{ij}^{\min}$. The corresponding efficiency curves can be found in Figures 4.24 and 4.25 on page 154. In general, the total systematic uncertainty is dominated by the JES and JER. Its typical value is around 5%, although it can be larger in some regions of phase space. In statistically dominated bins, it can reach values of the order of the JES and JER uncertainties or even larger. Overall, it stays of the same order as the total experimental systematic uncertainty.




Figure 4.31.: Systematic uncertainty accounting for differences in the selection efficiencies for events entering the transfer matrix, drawn as a function of H_T and $\Delta\phi_{ij}^{\min}$. The ‘envelope’ is drawn after applying the rebinning+smoothing procedure explained in the text.

4.6.5. Statistical and systematic uncertainties

The uncertainties on the final result have different origins. Statistical uncertainties are due to the finite size of both the data and the MC samples used to train the unfolding algorithm. Experimental uncertainties are evaluated at the reconstructed level and must be propagated through the unfolding procedure. Two systematic uncertainties derive from the unfolding process: the shape uncertainty and the selection efficiency uncertainty. Since the shape uncertainty was shown to be negligible (Section 4.6.3), the only contributions to the total systematic uncertainty are the jet energy scale, the jet energy resolution and the unfolding selection efficiency. Other sources of systematic uncertainty, like the trigger efficiency, are expected to be below the 1% level and are therefore not considered.

Statistical uncertainties. The statistical uncertainties are calculated using the Bootstrap Method [148, 149]. The method consists on generating new replicas of a given distribution, where each event count has been fluctuated with a Poisson weight. The statistical uncertainty of each bin is then simply the standard deviation of the corresponding bin content in all the replicas.

Replicas are generated for the data spectrum, the transfer matrix, and the full reconstructed- and truth-level spectra from MC (which are used to calculate fakes and inefficiencies). The nominal unfolded spectrum is obtained by using the nominal versions of all these histograms. 100 replicas of the unfolded spectrum are also obtained. The i th unfolded spectrum comes from using the i th replica of the data spectrum, the transfer matrix, and the full reconstructed- and truth-level spectra from MC.




Figure 4.32.: Optimisation of the number of unfolding iterations n as a function of H_T : change in the statistical uncertainty (top row), change in the systematic uncertainty (middle row), and shape uncertainty (bottom row). The dashed lines in the top plots correspond to the components of the total statistical uncertainty coming from data (red) and from Monte Carlo (black), for the case $n = 1$.




Figure 4.33.: Optimisation of the number of unfolding iterations as a function of $\Delta\phi_{ij}^{\min}$: change in the statistical uncertainty (top row), change in the systematic uncertainty (middle row), and shape uncertainty (bottom row). The dashed lines in the top plots correspond to the components of the total statistical uncertainty coming from data (red) and from Monte Carlo (black), for the case $n = 1$.

The Poisson weights are calculated with the RootCore Bootstrap package [150], and seeds are determined uniquely using the run and event numbers. This allows one to keep track of correlations, and would potentially permit the combination of these results with other measurements. Replicas are also computed during the unfolding of the distributions with systematic variations applied, such that statistical uncertainties on the systematic variations themselves can be calculated.

Jet Energy Scale (JES). The jet energy scale uncertainty was introduced in the context of the multi-jet search for new phenomena, in Section 3.7.1 on page 98. The main difference with respect to the approach followed in the multi-jet search is that here the uncertainty was split into 14 individual components, each of which was treated independently. In addition, the use of a more recent version of the jet energy calibration results in a reduction of some of the uncertainty components, as was explained in Section 4.2.1 on page 122.

The components of the JES uncertainty are propagated through the unfolding procedure one by one. The process is all done at the MC level: the PYTHIA spectrum is fluctuated up and down by one standard deviation for each uncertainty component, and the resulting distributions are unfolded and compared to the truth PYTHIA distribution. Any deviations correspond directly to the size of the systematic uncertainty. The unfolding algorithm itself is always trained using the nominal (i.e. not fluctuated by the JES uncertainty) PYTHIA events⁴.

To evaluate the significance of the unfolded JES systematic uncertainty, its statistical uncertainty is evaluated using the Bootstrap Method⁵. Then, following the same procedure discussed in Section 4.6.4, the distributions of each of the JES uncertainty components are rebinned and smoothed such that they are always more than 2σ significant and have a statistical uncertainty larger than 30%.

Jet Energy Resolution (JER). The jet energy resolution uncertainty was introduced in Section 3.7.1 as part of the multi-jet search (page 98). In this analysis, the JER systematic variation is applied to the events used to train the unfolding itself, while the spectrum to be unfolded is kept at the nominal value. This is the opposite to what is done for the JES. The energy of the jets is smeared using the ATLAS recommended tools. As

⁴Except for when the Jet Energy Resolution is evaluated. This is explained in the following subsection.

⁵Effectively this means that 100 replicas of the upward- and downward-fluctuated spectra are independently unfolded

for the JES, everything is done using PYTHIA only, and the systematic uncertainty comes from the deviation of the unfolded spectra from the expected truth PYTHIA distribution.

The combined size of JES and JER uncertainties is typically between 5% and 10%, although in regions limited in number of events it can be $> 50\%$.

Luminosity uncertainty. The uncertainty on the integrated luminosity is $\pm 2.8\%$. It is derived, following the same methodology as that detailed in Ref. [77], from a preliminary calibration of the luminosity scale derived from beam-separation scans performed in November 2012.

4.6.6. Optimisation of the unfolding algorithm

The Bayesian Iterative unfolding algorithm has one parameter to be optimised: the number of iterations. In general, it is expected that higher number of iterations will yield a result less sensitive to the MC simulation, but also affected by larger statistical uncertainties.

A scan over the number of iterations (1 to 4) is performed to choose the optimal value of the parameter. The following considerations apply:

- The statistical uncertainty of the unfolded spectrum.
- The total experimental systematic uncertainty of the unfolded spectrum.
- The shape uncertainty.

The same set of representative variables used in the previous section are shown as examples: H_T and $\Delta\phi_{ij}^{\min}$, with loose cuts $p_T^1 > 100$ GeV and $M_{jjjj} > 0$. Figures 4.32 and 4.33 on pages 164 and 165 show the statistical, experimental and shape uncertainties on the unfolded spectra. The statistical uncertainties are also shown divided by those obtained after only one iteration of the unfolding, in Figures 4.32e and 4.33e. They show that the statistical uncertainties of the unfolded spectrum grow with the number of iterations. The total systematic uncertainty, on the other hand, doesn't show such a clear trend: some bins (mostly at high p_T) have smaller systematic uncertainties with fewer iterations of the algorithm, but others show the opposite behaviour. In the angular variables the behaviour is mostly random, apart from bins where statistical uncertainties dominate, where smaller number of iterations are beneficial. Finally, the

shape uncertainty gets significantly smaller with increasing numbers of iterations. Even with one iteration, though, the shape uncertainty is typically smaller than 5%, which means that it is never dominant. With three or more iterations the uncertainty is consistent with zero within 1σ for all the variables studied⁶.

The optimal number of iterations should be chosen so as to minimise the uncertainties in the analysis. Systematic uncertainties are not very sensitive to the number of iterations, so they are left aside. Statistical uncertainties demand few numbers of iterations, and shape uncertainties the opposite. The latter are always subdominant, but they represent an important message - the *certainty* that the shape of the unfolded spectrum will be biased. They are therefore also taken into account. Using only one iteration will result in a significant bias in almost all variables. Two iterations, on the other hand, reduce the bias considerably while keeping the statistical uncertainties small. With two iterations the shape uncertainty is consistent with zero within 2σ for all variables and all bins. Consequently, two iterations is chosen as the optimal value.

4.7. Theoretical calculations and uncertainties

The measurements performed in this chapter have two main motivations: evaluating the robustness of LO calculations matched to parton showers, and testing perturbative NLO calculations. LO calculations are still used as nominal predictions in many searches for new physics, so it is also important to understand their performance.

In order to pin down the origin of the various differences between the different predictions and the data, one would need to isolate all of the intervening effects in different samples, including parton distribution functions, matrix element (ME) calculation, parton shower model or matching algorithm. Here the choice has been to consider only a small number of theoretical predictions and highlight the areas of phase space where differences appear, without attempting to discern the origin of the discrepancy.

The leading order samples considered were all produced within the ATLAS Collaboration. They include parton shower, hadronisation and underlying event. The uncertainties in the calculation are large due to the choice of renormalisation and factorisation scales, as well as the semiempirical tuning that enters the underlying event. For this reason, no theoretical systematic uncertainties are assigned to leading order predictions. In addition

⁶ σ was defined in the context of the shape uncertainty, in Section 4.6.3.

to the LO samples, two extra higher-order predictions produced by collaborators have also been included in the comparison.

The next subsections will describe the different LO and NLO theoretical predictions that will be compared against the measured data. For each of the predictions, a very short description of the generator will be given, together with some details of the implementation of the sample used in the analysis.

4.7.1. Leading-order predictions

One of the analysis objectives listed in Section 4.4 was studying the differences between a pure parton-shower (PS) calculation evolving from a $2 \rightarrow 2$ ME and a multi-leg calculation matched to a parton-shower (ME+PS), all at leading order. With this end in mind, the following samples are studied: one PYTHIA sample, a HERWIG++ sample, and a MADGRAPH sample. They are all listed in Table 4.4, page 132.

PYTHIA 8.160 [151] and HERWIG++ 2.5.2 [109] are both leading-logarithmic PS models; they use with different ordering variables in the PS, and different hadronisation models. The parton shower in the PYTHIA 8 sample is based on *dipole showering* and it is ordered in transverse momentum. The HERWIG++ PS is angular ordered. Hadronisation in PYTHIA is based solely on the Lund string fragmentation model, while HERWIG++ uses the cluster hadronisation model.

The models integrated in MC generators rely on a set of free parameters that must be tuned with data [152]. The tuning is performed iteratively, optimising a few parameters at a time, and ensuring that globally the parameters remain at their optimal values. Some of the models that require tuning include those for hadronisation, parton shower or multiple parton interactions. A “tune” refers to a specific set of parameter values. The tune employed by the PYTHIA sample is the ATLAS tune AU2 [153], which uses a range of 900 GeV and 7 TeV measurements from ATLAS. The HERWIG++ sample uses the UE-EE-3 tune [154], which was introduced together with some developments in the generator models, and describes the underlying event measured by ATLAS at 900 GeV and 7 TeV and by CDF at 1.8 TeV. The parton distribution functions are the NLO CT-10 [155] for PYTHIA and the LO distributions of CTEQ6L1 [156] for HERWIG, as shown in Table 4.4 on page 132.

The predictions from PYTHIA and HERWIG++ are compared to those from the multi-leg generator MADGRAPH 5 v1.5.12 [30], which have up to four partons in the ME.

MADGRAPH is interfaced with PYTHIA 6.427, which performs the parton shower. The ATLAS tune AUET2B-CTEQ6L1 [157] is employed. The ME has up to four partons, which are matched to the jets in the final state after the showering to avoid overlaps between the phase-space descriptions of the different pieces. The matching is performed using the shower k_T -jet MLM matching method [158].

4.7.2. Next-to-leading-order predictions

The cross section measurement is also compared to NLO QCD predictions for the production of four jets generated with the BLACKHAT library [33] in combination with SHERPA [159]. BLACKHAT performs the one-loop virtual corrections using methods that scale well as the number of external legs increases, which makes the four-jet calculation possible [34]. The remaining terms of the full NLO computation are obtained from AMEGIC++ [160, 161], part of SHERPA. The PDF used is NLO CT-10, and the simulation does not include parton shower, hadronisation or underlying event. Non-perturbative effects are expected to be small in the phase space of the analysis [162]. This is a preliminary result from collaborators outside ATLAS, which may be subject to modifications.

4.7.3. Other predictions

The results are also compared to a prediction obtained with the *High Energy Jets* (HEJ) perturbative framework [142, 163, 164]. HEJ provides an all-order description of processes with more than two hard jets. It calculates an approximation for both real and virtual corrections which also captures the hard, wide-angle emissions that parton showers may fail to reproduce. These hard, perturbative corrections are expected to be relevant at the LHC energies, where the increasing phase space allows for additional radiation produced between jets separated by a large rapidity interval [142]. The PDF used is NLO CT-10. The uncertainty of the predictions corresponds to variations in the choice of PDF and scale; it will not be included in the figures to avoid overcrowding them. It is typically of the order $^{+50\%}_{-30\%}$. Parton shower, hadronisation and underlying event are not included. It is a preliminary calculation provided by collaborators outside ATLAS, and should be taken as an estimate, as it may be subject to modifications.

4.8. Results

The differential inclusive four-jet cross sections are shown in Figures 4.34 to 4.49 for jets reconstructed with the anti- k_t algorithm with $R = 0.4$. The measurements extend from a jet transverse momentum of 64 GeV to several TeV, spanning two orders of magnitude in p_T and over seven orders of magnitude in cross section. The measured cross sections have been corrected for all detector effects using the unfolding procedure described in Section 4.6.1. The corrected results have been compared to the theoretical predictions described in Section 4.7.

In order to compare the data and the MC simulations, it is necessary to scale the MC to the dataset luminosity, 20.3 fb^{-1} . An additional scale factor is applied to facilitate the comparison between generators. This scale factor is obtained from the ratio of the data and MC integrals of the leading jet p_T distribution after applying all the analysis cuts described in Section 4.2.2. The integrals are computed in the range $500 < p_T^1 < 1000$ GeV, and rounded to 2 significant digits. A single scale factor is derived for each MC generator: 0.57 for PYTHIA, 0.97 for MADGRAPH, 1.30 for HERWIG++, 0.80 for HEJ and 0.75 for BLACKHAT. For HEJ and BLACKHAT, the normalisation is performed with respect to the unfolded data spectrum.

This choice of normalisation is to some extent arbitrary, and should not be taken as an accurate value of the correction to the MC cross section. For example, if the normalisation had been performed using the low- p_T^1 data as a reference (0–500 GeV), the corresponding BLACKHAT factor would have been 0.82. In addition, some studies in progress [165] show that the choice of renormalisation and factorisation scale in BLACKHAT may not be ideal, as the cross section drops to large negative values when the scales are reduced. This results in large negative cross section uncertainties, which would make the raw BLACKHAT cross section compatible with the data.

The results for the cross sections are presented in Figures 4.34 to 4.49. They contain the differential cross sections as a function of all the variables of interest, calculated from data and compared against different theoretical predictions in different bins of p_T^1 and M_{jjjj} . The total experimental systematic uncertainty is drawn as a grey band centred at the data points for the spectrum comparisons (top plots) and at 1 for the ratios (bottom plots). The different components of the uncertainty (JES and JER, unfolding selection efficiency, and luminosity), are drawn individually. The up and down components of the total experimental systematic uncertainty are obtained independently, by adding

up quadratically all the positive and negative systematic uncertainty components of the spectra, respectively. The total experimental uncertainty is typically 5 – 20%, and larger in statistically limited bins. Theoretical uncertainties due to scale and PDF variations are only available for the HEJ sample. They are typically $^{+50\%}_{-30\%}$, and will not be drawn to avoid the overcrowding of the figures. The uncertainties for LO predictions are not included either, as they are expected to be very large.





Figure 4.34.: Unfolded p_T^1 distribution, compared against different theoretical predictions. The HEJ and BLACKHAT predictions are at parton level. The grey shaded band corresponds to the total experimental systematic uncertainty.

Figure 4.35.: As in Figure 4.34 but for p_T^2 .




Figure 4.36.: As in Figure 4.34 but for p_T^3 .






Figure 4.37.: As in Figure 4.34 but for p_T^4 .

Figure 4.38.: As in Figure 4.34 but for H_T .

Figure 4.39.: As in Figure 4.34 but for M_{jjjj} .




Figure 4.40.: Unfolded $M_{jj}^{\min}/M_{jjjj}$ distributions for different cuts in M_{jjjj} , compared against the corresponding theoretical predictions. The HEJ and BLACKHAT predictions are at parton level. The grey shaded band corresponds to the total experimental systematic uncertainty.




Figure 4.41.: Unfolded $\Delta\phi_{ij}^{\min}$ distributions for different cuts in p_T^1 , compared against the corresponding theoretical predictions. The HEJ and BLACKHAT predictions are at parton level. The grey shaded band corresponds to the total experimental systematic uncertainty.





Figure 4.42.: As in Figure 4.41 but for $\Delta\phi_{ijk}^{\min}$.

Figure 4.43.: As in Figure 4.41 but for $\Delta Y_{ij}^{\min}$.





Figure 4.44.: As in Figure 4.41 but for $\Delta Y_{ijk}^{\min}$.

Figure 4.45.: As in Figure 4.41 but for $\Delta Y_{ij}^{\max}$.




Figure 4.46.: Unfolded $\Sigma p_T^{\text{central}}$ distributions for $\Delta Y_{ij}^{\text{max}} > 1$ and different cuts in p_T^1 , compared against the corresponding theoretical predictions. The HEJ and BLACK-HAT predictions are at parton level. The grey shaded band corresponds to the total experimental systematic uncertainty.




Figure 4.47.: As in Figure 4.46 but for $\Delta Y_{ij}^{\max} > 2$.




Figure 4.48.: As in Figure 4.46 but for $\Delta Y_{ij}^{\max} > 3$.




Figure 4.49.: As in Figure 4.46 but for $\Delta Y_{ij}^{\max} > 4$.

In general, the best description of the overwhelming majority of the variables is given by BLACKHAT. Next are MADGRAPH and HEJ, which describe the data well in many regions of phase space. HEJ is in most cases compatible with the data when theoretical uncertainties are taken into account. Parton shower programs – which are in principle not expected to be optimal for this type of events – are found to perform well in a variety of situations. The following discussion is based on the results obtained after applying the particular choice of normalisation of the theoretical histograms explained at the beginning of this section. BLACKHAT will only be discussed if any deviations from the data are present; in all other cases, it should be assumed that it is the one that provides the most accurate prediction.

Figures 4.34 to 4.37 show the p_T distributions of the leading four jets. All the LO generators show a slope with respect to the data in the leading jet p_T (Figure 4.34). HERWIG++ and HEJ are remarkably flat above ~ 300 GeV. MADGRAPH is within the experimental uncertainties from ~ 400 GeV and above, and it is the only one that (for the chosen normalisation) predicts fewer events with soft p_T^1 .

The subleading jet p_T (Figure 4.35) is well described by HEJ, and the LO generators show similar trends to those in p_T^1 . The distributions of p_T^3 and p_T^4 (Figures 4.36 and 4.37) are well described by PYTHIA. HEJ, and especially HERWIG++, overestimate the number of events with high p_T^3 and p_T^4 . H_T (Figure 4.38) shows features similar to those in p_T^1 .

The total invariant mass M_{jjjj} is studied in Figure 4.39. All LO generators describe the shape of the data between 1 TeV and 3–6 TeV, although the normalisation is off by 20–40%. HERWIG++ predicts significantly more events at high M_{jjjj} . HEJ is compatible with the data in all bins within its theoretical uncertainties.

The description of the splitting scales is tested in 4.40 through the variable $M_{jj}^{\min}/M_{jjjj}$. $M_{jj}^{\min}/M_{jjjj}$ is generally well described by PYTHIA for the first M_{jjjj} bins, but the shape deteriorates at higher M_{jjjj} . HERWIG++ has a similar but more discrepant shape, and it consistently overestimates the number of events with low values of $M_{jj}^{\min}/M_{jjjj}$. MADGRAPH provides a good description, with a flat ratio in all M_{jjjj} bins, ignoring overall normalisation issues. The HEJ prediction is mostly compatible with the data for low M_{jjjj} , but shows trends similar to those of PYTHIA at higher values of M_{jjjj} . BLACKHAT overestimates the number of events in the very first bin, and generally agrees well with the data in the rest.

Figure 4.41 compares the distributions of $\Delta\phi_{ij}^{\min}$ for different cuts in p_T^1 . PYTHIA has a downwards slope with respect to the data in all the p_T^1 bins. The other generators reproduce the data well.

The $\Delta\phi_{ijk}^{\min}$ spectrum is shown in Figure 4.42. In general, the description of the data improves as p_T^1 increases. For PYTHIA, at low p_T^1 (where events with three jets recoiling against one dominate) the number of events is significantly overestimated; as p_T^1 increases, the agreement improves such that the $p_T^1 > 1000$ GeV bin is very well described. MADGRAPH and HERWIG++ have some (less significant) features at low p_T^1 , and improve at higher p_T^1 too. HEJ is mostly in good agreement with data.

Figures 4.43 and 4.44 compare the distributions of $\Delta Y_{ij}^{\min}$ and $\Delta Y_{ijk}^{\min}$ with data. $\Delta Y_{ij}^{\min}$ is remarkably well described by PYTHIA, showing no significant trend. MADGRAPH mostly tends to underestimate high $\Delta Y_{ij}^{\min}$ values, while HERWIG++ does the opposite. HEJ overestimates the number of events with high $\Delta Y_{ij}^{\min}$ values at low p_T^1 , but describes the data very well at larger values of p_T^1 .

$\Delta Y_{ijk}^{\min}$ (in Figure 4.44) is not so well described by PYTHIA; the region with the best statistical significance ($p_T^1 > 400$ GeV, where the single-jet triggers enter but the momentum constraint is not too tight) shows a significant positive slope of PYTHIA with respect to the data. HERWIG++ behaves similarly even at larger p_T^1 . MADGRAPH reproduces the shape of the data well, as does HEJ for $p_T^1 > 400$ GeV; for smaller values of p_T^1 , it again overestimates the number of events at the end of the spectrum.

$\Delta Y_{ij}^{\max}$, shown in Figure 4.45 is very well described by HEJ above $p_T^1 > 400$ GeV. Both PYTHIA and HERWIG++ have upward slopes in all p_T^1 bins. MADGRAPH provides mostly a good description of the data.

Finally, the variables setting a minimum rapidity gap and measuring the total p_T of the central jets all show a similar behaviour. Figures 4.46 to 4.49 correspond to $\Delta Y_{ij}^{\max} > 1, 2, 3, 4$. In general, all generators have problems around the points where the contribution from different jet thresholds changes (for example, when the leading jet is allowed to be central).

MADGRAPH provides in general the most accurate description of the variables, especially at low p_T^1 . The agreement deteriorates with higher $\Delta Y_{ij}^{\max}$ and higher p_T^1 , but it is still the least affected by the changes between populations. PYTHIA is the most discrepant, with kinks at the transition points that reach differences of 70% at high p_T^1 . At low p_T^1 the shape is better described. In general the distributions are well

described by HEJ, especially the high $\Sigma p_T^{\text{central}}$ region; the low $\Sigma p_T^{\text{central}}$ region shows more shape differences. HERWIG++ describes the data well at low $\Delta Y_{ij}^{\text{max}}$, but as $\Delta Y_{ij}^{\text{max}}$ grows its normalisation worsens, as well as the shape – particularly at high p_T^1 . BLACKHAT significantly overestimates the number of events with low $\Sigma p_T^{\text{central}}$ for high p_T^1 , and provides a good description of the data otherwise.

4.9. Summary and outlook

This chapter has presented the methods used to perform first measurement of the cross section of four-jet events performed using 8 TeV data in ATLAS. The cross section has been measured differentially in several variables describing the kinematics and spatial configuration of the jets in the event, and compared against different theoretical predictions. The typical experimental systematic uncertainty on the result is $\mathcal{O}(10\%)$.

Several messages can be extracted from the results. Each of the samples studied was applied a global normalisation factor that ranged between 0.57 and 1.3 to facilitate the comparison with data. Within this normalisation scheme, PS generators are seen to describe the data well in a few cases, for example $\Delta\phi_{ijk}^{\text{min}}$ with HERWIG++ and $\Delta Y_{ij}^{\text{min}}$ with PYTHIA. Of the three LO samples considered, MADGRAPH (which has a four-jet ME calculation matched to a PS) provides a remarkably good description of the data. The preliminary HEJ prediction is mostly compatible with the data within scale and PDF uncertainties. Some regions are less accurately described, such as $\Delta Y_{ij}^{\text{min}}$ and $\Delta Y_{ijk}^{\text{min}}$ in the lowest p_T^1 bin. The preliminary BLACKHAT four-jet NLO prediction provides a very good description of the data throughout, only showing discrepancies at low $\Sigma p_T^{\text{central}}$ for the large p_T^1 bins.

With these experimental results now becoming publicly available, the baton will pass to the theoretical physicists whose job it will be to understand how better to model these high jet multiplicity events.

From the experimental perspective, further work could include a systematic study of additional theory predictions to accurately determine the origin of the data-MC discrepancies. The next step would be to analyse five-jet events. The ratio of the two jet multiplicities is expected to result in the cancellation of theoretical uncertainties, which will give better discriminating power between the different QCD calculations.

5 | Conclusions

*‘It seems very pretty,’ she said when she had finished it,
‘but it’s rather hard to understand! [...]’
Somehow it seems to fill my head with ideas —
only I don’t exactly know what they are!’*

The operation of the LHC since its switching on in 2010 has been remarkable. The Higgs boson was discovered in 2012, previous exclusion limits in searches for new physics have been significantly extended, and a great number of Standard Model processes have been measured to high precision.

This thesis has introduced two analyses of the complete 8 TeV dataset performed by the author. Both analyses focus on final states with large jet multiplicities, but they take different, complementary approaches. The first analysis looks for events with 8, 9, 10 or even more jets, since these could indicate the presence of new physics. Relying on a powerful data-driven technique to determine the dominant Standard Model backgrounds, the search is able to exclude supersymmetric gluinos with masses below 1.1 TeV in a simplified model where they decay via $\tilde{g} \rightarrow t + \bar{t} + \tilde{\chi}_1^0$. The second analysis focuses on events with jet multiplicities of 4 or more, performing the first measurement of their cross section at $\sqrt{s} = 8$ TeV. The measurement is performed differentially in several variables which provide discrimination between different theoretical predictions. The SUSY multi-jet search has already been published [3], and the SM measurement is undergoing ATLAS approval.

The LHC will turn on again in 2015 (only some months from the time of writing), and start taking data at 13 TeV. No one knows what kind of phenomena may occur at such high energies, and although searches for new physics such as the one presented in this thesis have been unsuccessful so far, this may very well change in the near future. Simplified versions of SUSY have been highly constrained by LHC searches, but as

discussed in the opening chapter of this thesis, more general SUSY models remain alive. Complementing the searches with measurements such as the one presented in this thesis allows the theoretical community to make progress in parallel with the calculation of the SM backgrounds. This is important both to deepen our understanding of the physics of the Standard Model, and to improve the sensitivity of new physics searches and the robustness of future discoveries.

This feedback process between searches, measurements and theory is at the heart of LHC physics, and guarantees that at the end of its lifetime, no matter what Nature is hiding at high energies, the experiment will have been a success.

A | Trigger efficiencies

Trigger efficiency curves allow one to study the fraction of events which fire a particular trigger as a function of a certain variable x . The inefficiencies are introduced by differences between the objects and decisions made at the different stages of the trigger and at the offline level. The sample of events used for the test must be carefully chosen so as not to introduce a bias. One option to measure the efficiency of a trigger A is to select a set of events which fired an orthogonal trigger B , in which case the efficiency of A , $\epsilon(A)$, is simply the bin-by-bin ratio

$$\epsilon(A, x) = \frac{\mathcal{H}(x)^{\text{Events triggered by A and B}}}{\mathcal{H}(x)^{\text{Events triggered by B}}}, \quad (\text{A.1})$$

where $\mathcal{H}(x)^\alpha$ is the histogram of the variable x for events selected according to α .

It is also possible to use a sample triggered by a looser trigger C , such that $\epsilon(C|A) = 1$, and then, according to Bayes' theorem,

$$\epsilon(A) = \frac{\epsilon(A|C) \epsilon(C)}{\epsilon(C|A)} = \epsilon(A|C) \epsilon(C), \quad (\text{A.2})$$

where $\epsilon(C)$ is implicit in the event selection, and $\epsilon(A|C)$ can also be obtained as

$$\epsilon(A|C, x) = \frac{\mathcal{H}(x)^{\text{Events triggered by A and C}}}{\mathcal{H}(x)^{\text{Events triggered by C}}}. \quad (\text{A.3})$$

A certain trigger C is said to be *looser* than a given trigger A if its cuts are less restrictive than the cuts implicit in A ; for example, C could have the same configuration as A but demand a lower p_T from a particular object.

The calculation of the efficiency is more delicate when the triggers are prescaled, since the correlation between the events that pass the reference and the test triggers is lost. Fortunately, the response from jet triggers can be easily emulated, as all the

information about the trigger objects needed (L1 cells, L2 clusters and Event Filter jets) is kept. The efficiency of a prescaled jet trigger can then be calculated as long as there exist an unbiased sample of events triggered by either an orthogonal trigger or a looser trigger, any of which could be prescaled as well. The prescale of the trigger of interest is effectively removed by reconstructing the trigger response for each of the events that passed the (prescaled or not) reference trigger. The prescale of the reference trigger is unavoidable, but the only effect it will have on the final result is to reduce the total luminosity.

In this thesis, efficiencies are typically computed bin-by-bin in the p_T of the N^{th} jet for an N -jet trigger, or in the minimum separation dR_{min}^{4j} between the jets, defined in Section 4.3.2 on page 129. Using a bin-by-bin notation, the efficiency $\epsilon^i(A)$ for a given trigger A in bin i can also be written as

$$\epsilon_i(A) = \frac{N_i^{\text{Events triggered by A and B}}}{N_i^{\text{Events triggered by B}}}, \quad (\text{A.4})$$

where the denominator is the number of events passing the reference trigger B in bin i , and the numerator is the number of events passing both the reference trigger B and the trigger of interest A in bin i . This defines a ‘differential’ efficiency curve, since the efficiency is calculated in each bin in an exclusive manner. To better illustrate certain trends in the turn-on curves, the ‘integral’ efficiency is also studied in most cases. The integral efficiency is defined as:

$$\epsilon_i^{\text{integ}}(A) = \frac{\sum_{j=i}^{\infty} N_j^{\text{Events triggered by A and B}}}{\sum_{k=i}^{\infty} N_k^{\text{Events triggered by B}}}. \quad (\text{A.5})$$

The integral efficiency informs of the *global*, inclusive efficiency of the trigger for a particular set of cuts, rather than for an exclusive bin.

In the figure legends, efficiencies will be indicated as ‘Nominal trigger wrt. reference trigger’. The data taken by the reference trigger are not corrected for prescales.

B | Deriving variable-width binnings

The goal of this method is to derive a binning that yields an approximately flat purity (as opposed to steeply falling, as would be the case for a p_T spectrum with constant binning) with an average of $P = 80$ %. The stability distributions are used as a cross-check. Purity and stability are defined in Section 4.5.

Let i be the index used to enumerate reco-level quantities. Let r^i represent the bin centres along the reco axis. The first step consists on projecting the contents of each reco bin such that they can be studied as individual one-dimensional histograms. The standard deviation of each of these distributions is then plotted versus r_i , as an estimate of the resolution of the variable.

The resolution at a particular value of r_i (σ_i^{Resol}) is in principle taken as the targeted width of the bin centred at r_i . The distribution of σ_i^{Resol} as a function of r_i therefore shows how the bin width changes as a function of r_i . However, what is needed is a discrete set of non-overlapping, successive bins, so the next step is to obtain a smooth distribution and sample from it.

The scatter distribution of σ_i^{Resol} as a function of r_i is smoothed by calculating a non-parametric Gaussian kernel regression. The regression curve is formed by a succession of points, each of which is computed as the weighted average of all the σ_i^{Resol} . For each point contributing to the average, the weight is determined from a Gaussian centred in r_i and evaluated at said point. The standard deviation (or ‘bandwidth’) of the Gaussian is fixed for each of the variables studied, and its value is chosen so as to be sensitive to the shape of the distribution of σ_i^{Resol} , while reducing the statistical fluctuations of the data. The left plots in Figure B.1 show the σ_i^{Resol} distributions (blue triangles) and the corresponding regressions (black line) for the leading and sub-leading jet p_T ¹.

¹The distributions are weighted by a factor k , for reasons that will be explained later.




Figure B.1.: σ_i^k distribution, purity and stability for p_T^1 and p_T^2 . In the σ_i^k distributions (left), the blue triangles correspond to the truth σ_i^k as calculated from PYTHIA for the central value of each of the reco bins, r_i ; the black curve is the regression of the σ_i^k versus the central value of each of the reco bins; the red curve is the regression of the σ_i^k versus $r_i - \sigma^k$ (this is the curve used to do the bin sampling); and the black points correspond to the centres and widths of the bins proposed. The grey histograms (right) correspond to the resulting variable distributions.

The binning is obtained by sampling the σ_i^{Resol} curve. In order to simplify the sampling algorithm, the distribution used is actually σ_i^{Resol} versus $r_i - \sigma_i^{\text{Resol}}$, defining the curve $C(r_i)$. In other words, rather than studying the bin widths as a function of the bin centres, we study them as a function of the lower bin edges. The algorithm then proceeds by iterating over the data and creating new bins in the following way:

1. The lower edge of the first bin is set manually.

2. The width of the first bin is $2 \times \sigma_1^{\text{Resol}}$ — i.e., instead of evaluating the fitted curve at r_i , the first raw value of σ_i^{Resol} is used. This is done because the regression curves tend to be higher than the data at the lower end of the spectrum, giving artificially large bin widths. This condition forces the first bin to be small. Figure B.1 shows that it is safe to do this for the momentum variables, as the statistical uncertainty is smallest in that region.
3. The upper edge of the first bin is then renamed as the lower edge of the second bin.
4. The bin width of the second bin is calculated as $2 \times C(b_{\text{low}})$, where b_{low} is the lower edge of the bin.
5. The new upper edge turns into the lower edge of the third bin, and so on.
6. The process finishes when the upper edge is bigger than a maximum r_i set manually.

Bins are then merged if:

- The width is smaller than a certain number set by hand.
- The number of Monte Carlo events equivalent to 20.3 fb^{-1} is smaller than 100.
- The purity is smaller than 70%.

Taking σ_i^{Resol} as the starting point to derive the binnings results in purities between 60% and 70%. In order to increase this value to roughly 80%, we multiply the values of σ_i^{Resol} by a constant. We define, for convenience, $\sigma_i^k = k \times \sigma_i^{\text{Resol}}$, where k is the multiplicative factor used to increase the purity.

Figure B.1 shows the σ_i^k curves, purity and stability for the leading and sub-leading jet p_T . The proposed bin centres, before applying the constraints on statistics or minimum purity, correspond to the black circles. They are obtained by sampling the red curve (σ_i^k vs $r_i - \sigma_i^k$) following the method described above, and are found to agree with the black curve (σ_i^k vs r_i), as expected. Table B.1 shows the values of x_{min} , x_{max} and bandwidth used to derive the binning.




Figure B.2.: σ_i^k distribution, purity and stability for H_T . In the σ_i^k distributions (left), the blue triangles correspond to the truth σ_i^k as calculated from PYTHIA for the central value of each of the reco bins, r_i ; the black curve is the regression of the σ_i^k versus the central value of each of the reco bins; the red curve is the regression of the σ_i^k versus $r_i - \sigma_i^k$ (this is the curve used to do the bin sampling); and the black points correspond to the centres of the bins proposed. The grey histograms (right) correspond to the resulting variable distributions.




Figure B.3.: σ_i^k distribution, purity and stability for $\Sigma p_T^{\text{central}}, \Delta Y_{ij}^{\text{max}} > 1$. In the σ_i^k distributions (left), the blue triangles correspond to the truth σ_i^k as calculated from PYTHIA for the central value of each of the reco bins, r_i ; the black curve is the regression of the σ_i^k versus the central value of each of the reco bins; the red curve is the regression of the σ_i^k versus $r_i - \sigma_i^k$ (this is the curve used to do the bin sampling); and the black points correspond to the centres of the bins proposed. The grey histograms (right) correspond to the resulting variable distributions.

Variable	$x_{\min}$	$x_{\max}$	Bandwidth	k factor
p_{T}^1	100 GeV	4000 GeV	300 GeV	2.0
p_{T}^2	64 GeV	3000 GeV	300 GeV	3.0
p_{T}^3	64 GeV	2000 GeV	200 GeV	3.0
p_{T}^4	64 GeV	1500 GeV	200 GeV	2.0
H_{T}	292 GeV	7000 GeV	700 GeV	3.0
M_{jjjj}	100 GeV	7000 GeV	2000 GeV	1.3
$M_{jj}^{\min} / M_{jjjj}$	0	0.4	0.05	1.5
$\Sigma p_{\text{T}}^{\text{central}}, \Delta Y_{ij}^{\max} > 1$	128 GeV	5000 GeV	700 GeV	1.0
$\Sigma p_{\text{T}}^{\text{central}}, \Delta Y_{ij}^{\max} > 2$	128 GeV	5000 GeV	700 GeV	1.3
$\Sigma p_{\text{T}}^{\text{central}}, \Delta Y_{ij}^{\max} > 3$	128 GeV	5000 GeV	700 GeV	1.5
$\Sigma p_{\text{T}}^{\text{central}}, \Delta Y_{ij}^{\max} > 4$	128 GeV	5000 GeV	700 GeV	1.7

Table B.1.: Parameters used to derive the variable-width bins for the momentum and mass variables.

Colophon

The quotes at the beginning of each chapter
are taken from ‘Alice’s Adventures in Wonderland’
and ‘Through the Looking Glass’,
by Lewis Carroll (Christ Church, 1851).
The opening quote is taken from the first episode
of ‘Cosmos’, by Carl Sagan.

This thesis was made in L^AT_EX
using the *hepthesis* class [166].

Bibliography

- [1] ATLAS Collaboration, “Limits on metastable gluinos from ATLAS SUSY searches at 8 TeV,” CERN, Geneva, Tech. Rep. ATLAS-CONF-2014-037, Jul 2014.
- [2] ATLAS Collaboration, “Performance assumptions based on full simulation for an upgraded ATLAS detector at a High-Luminosity LHC,” CERN, Geneva, Tech. Rep. ATL-PHYS-PUB-2013-009, Sep 2013.
- [3] ATLAS Collaboration, “Search for new phenomena in final states with large jet multiplicities and missing transverse momentum at $\sqrt{s}=8$ TeV proton-proton collisions using the ATLAS experiment,” *Journal of High Energy Physics*, vol. 2013, no. 10, pp. 1–50, 2013.
- [4] ATLAS Collaboration, “Searches for Supersymmetry at the high luminosity LHC with the ATLAS Detector,” CERN, Geneva, Tech. Rep. ATL-PHYS-PUB-2013-002, Feb 2013.
- [5] F. Halzen and A. D. Martin, *Quarks & Leptons: An introductory Course in Modern Particle Physics*. John Wiley & Sons, 2008.
- [6] I. J. Aitchison and A. J. Hey, *Gauge Theories in Particle Physics: A Practical Introduction, Volume 2: Non-Abelian Gauge Theories: QCD and The Electroweak Theory*. CRC Press, 2012.
- [7] I. J. Aitchison, A. J. Hey, and D. F. Brewer, *Gauge Theories in Particle Physics: A Practical Introduction, Volume 1: From Relativistic Quantum Mechanics to QED*. CRC Press, 2003.
- [8] G. L. Kane, *Modern elementary particle physics*. Addison-Wesley, 1993.
- [9] K. A. Olive *et al.*, “Review of Particle Physics,” *Chinese Physics C*, vol. 38, no. 9, 2014.
- [10] R. Davis, D. S. Harmer, and K. C. Hoffman, “Search for Neutrinos from the Sun,”

- Phys. Rev. Lett.*, vol. 20, pp. 1205–1209, 1968.
- [11] Y. Fukuda *et al.*, “Evidence for Oscillation of Atmospheric Neutrinos,” *Phys. Rev. Lett.*, vol. 81, pp. 1562–1567, 1998.
- [12] Q. R. Ahmad *et al.*, “Measurement of the Rate of $\nu_e + d \rightarrow p + p + e^-$ Interactions Produced by 8B Solar Neutrinos at the Sudbury Neutrino Observatory,” *Phys. Rev. Lett.*, vol. 87, p. 071301, 2001.
- [13] Q. R. Ahmad *et al.*, “Direct Evidence for Neutrino Flavor Transformation from Neutral Current Interactions in the Sudbury Neutrino Observatory,” *Phys. Rev. Lett.*, vol. 89, p. 011301, 2002.
- [14] K. Eguchi *et al.*, “First Results from KamLAND: Evidence for Reactor Antineutrino Disappearance,” *Phys. Rev. Lett.*, vol. 90, p. 021802, 2003.
- [15] ATLAS Collaboration, “Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC,” *Physics Letters B*, vol. 716, no. 1, pp. 1 – 29, 2012.
- [16] ATLAS Collaboration, “A Particle Consistent with the Higgs Boson Observed with the ATLAS Detector at the Large Hadron Collider,” *Science*, vol. 338, no. 6114, pp. 1576–1582, 2012.
- [17] ATLAS Collaboration, “Evidence for the spin-0 nature of the Higgs boson using ATLAS data,” *Physics Letters B*, vol. 726, pp. 120 – 144, 2013.
- [18] E. Noether, “Invariant variational problems,” *News from the Society of Sciences in Göttingen, Mathematical Physics Class*, pp. 235–257, 1918.
- [19] P. W. Higgs, “Broken symmetries, massless particles and gauge fields,” *Phys. Lett.*, vol. 12, p. 132, 1964.
- [20] P. W. Higgs, “Broken Symmetries and the Masses of Gauge Bosons,” *Phys. Rev. Lett.*, vol. 13, p. 508, 1964.
- [21] F. Englert and R. Brout, “Broken Symmetry and the Mass of Gauge Vector Mesons,” *Phys. Rev. Lett.*, vol. 13, p. 321, 1964.
- [22] G. Guralnik, C. Hagen, and T. Kibble, “Global Conservation Laws and Massless Particles,” *Phys. Rev. Lett.*, vol. 13, p. 585, 1964.
- [23] R. A. Battye and A. Moss, “Evidence for Massive Neutrinos from Cosmic Microwave

- Background and Lensing Observations,” *Phys. Rev. Lett.*, vol. 112, p. 051303, Feb 2014.
- [24] I. Park, S. Schnetzer *et al.*, “Experimental evidence for the non-Abelian nature of QCD from a study of multijet events produced in e^+e^- annihilation,” *Physical Review Letters*, vol. 62, no. 15, p. 1713, 1989.
- [25] G. Arnison, A. Astbury *et al.*, “Angular distributions and structure functions from two-jet events at the CERN SPS pp collider,” *Physics Letters B*, vol. 136, no. 4, pp. 294–300, 1984.
- [26] S. Chekanov *et al.*, “Dijet photoproduction at HERA and the structure of the photon,” *The European Physical Journal C*, vol. 23, no. 4, pp. 615–631, 2002.
- [27] G. P. Salam, “Elements of QCD for hadron colliders,” *arXiv:1011.5131*, 2010.
- [28] A. Buckley, J. Butterworth *et al.*, “General-purpose event generators for LHC physics,” *Physics Reports*, vol. 504, no. 5, pp. 145–233, 2011.
- [29] P. Richardson, “Introduction to Monte Carlo Event Generation. Lecture 1: Introduction to Monte Carlo Techniques,” 2013 MCNet School <https://indico.desy.de/contributionDisplay.py?sessionId=16&contribId=9&confId=7132>.
- [30] J. Alwall, M. Herquet, F. Maltoni, O. Mattelaer, and T. Stelzer, “MadGraph 5 : Going Beyond,” *JHEP*, vol. 06, p. 128, 2011.
- [31] M. L. Mangano, F. Piccinini *et al.*, “ALPGEN, a generator for hard multiparton processes in hadronic collisions,” *Journal of High Energy Physics*, vol. 2003, no. 07, p. 001, 2003.
- [32] Sherpa 2: <https://sherpa.hepforge.org/doc/SHERPA-MC-2.0.beta.html>.
- [33] C. Berger, Z. Bern *et al.*, “Automated implementation of on-shell methods for one-loop amplitudes,” *Physical Review D*, vol. 78, no. 3, p. 036003, 2008.
- [34] Z. Bern, G. Diana *et al.*, “Four-jet production at the Large Hadron Collider at next-to-leading order in QCD,” *Physical Review Letters*, vol. 109, no. 4, p. 042001, 2012.
- [35] S. Badger, B. Biedermann, P. Uwer, and V. Yundin, “Next-to-leading order QCD corrections to five jet production at the LHC,” *Phys. Rev. D*, vol. 89, p. 034019, Feb 2014.

- [36] S. Badger, B. Biedermann, P. Uwer, and V. Yundin, “NLO QCD corrections to multi-jet production at the LHC with a centre-of-mass energy of $\sqrt{s} = 8$ TeV,” *Phys.Lett.*, vol. B718, pp. 965–978, 2013.
- [37] S. Catani, F. Krauss, B. R. Webber, and R. Kuhn, “QCD matrix elements+ parton showers,” *Journal of High Energy Physics*, vol. 2001, no. 11, p. 063, 2001.
- [38] M. L. Mangano, M. Moretti, F. Piccinini, and M. Treccani, “Matching matrix elements and shower evolution for top-pair production in hadronic collisions,” *Journal of High Energy Physics*, vol. 2007, no. 01, p. 013, 2007.
- [39] B. Andersson, G. Gustafson, G. Ingelman, and T. Sjöstrand, “Parton fragmentation and string dynamics,” *Physics Reports*, vol. 97, no. 2, pp. 31–145, 1983.
- [40] J.-C. Winter, F. Krauss, and G. Soff, “A Modified cluster hadronization model,” *Eur.Phys.J.*, vol. C36, pp. 381–395, 2004.
- [41] B. R. Webber, “A QCD model for jet fragmentation including soft gluon interference,” *Nuclear Physics B*, vol. 238, no. 3, pp. 492–528, 1984.
- [42] H.-L. Lai, M. Guzzi *et al.*, “New parton distributions for collider physics,” *Phys. Rev.*, vol. D82, p. 074024, 2010.
- [43] J. Pumplin *et al.*, “New generation of parton distributions with uncertainties from global QCD analysis,” *JHEP*, vol. 07, p. 012, 2002.
- [44] P. M. Nadolsky *et al.*, “Implications of CTEQ global analysis for collider observables,” *Phys. Rev.*, vol. D78, p. 013004, 2008.
- [45] C. Buttar, J. Butterworth *et al.*, “The underlying event,” in *HERA and the LHC-A workshop on the implications of HERA for LHC physics: Proceedings Part A*, 2005, p. 192.
- [46] E. Halkiadakis, G. Redlinger, and D. Shih, “Status and Implications of BSM Searches at the LHC,” *arXiv:1411.1427*, 2014.
- [47] Yu.A. Golfand and E.P. Likhtman, “Extension of the algebra of Poincaré group generators and violation of P invariance,” *JETP Lett.* **13** (1971) 323-326.
A. Neveu and J.H. Schwartz, “Factorizable dual model of pions,” *Nucl. Phys.*, B31, 86-112, 1971.
A. Neveu and J.H. Schwartz, “Quark model of dual pions,” *Phys. Rev.*, D4, 1109-1111, 1971.

- P. Ramond, “Dual theory for free fermions,” *Phys. Rev.*, D3, 2415-2418, 1971.
- D.V. Volkov and V.P. Akulov, “Diffractive dissociation of composite particles,” *Phys. Lett.*, B46, 109-130, 1973.
- J. Wess and B. Zumino, “Light cone approach to positivity bounds on structure functions for deep inelastic lepton scattering in Weinberg’s theory,” *Phys. Lett.*, B49, 52-60, 1974.
- J. Wess and B. Zumino, “Supergauge transformations in four dimensions,” *Nucl. Phys.*, B70, 39-50, 1974.
- [48] S. Coleman and J. Mandula, “All possible symmetries of the S matrix,” *Physical Review*, vol. 159, no. 5, p. 1251, 1967.
- [49] R. Haag, J. T. Łopuszański, and M. Sohnius, “All possible generators of supersymmetries of the S -matrix,” *Nuclear Physics B*, vol. 88, no. 2, pp. 257–274, 1975.
- [50] S. P. Martin, “A supersymmetry primer,” *hep-ph/9709356*, 1997.
- [51] J. L. Feng, “Naturalness and the Status of Supersymmetry,” *arXiv:1302.6587*, 2013.
- [52] S. Dimopoulos and H. Georgi, “Softly broken supersymmetry and $SU(5)$,” *Nuclear Physics B*, vol. 193, no. 1, pp. 150–162, 1981.
- [53] P. Nath and P. Fileviez Pérez, “Proton stability in grand unified theories, in strings and in branes,” *Physics reports*, vol. 441, no. 5, pp. 191–317, 2007.
- [54] G. R. Farrar and P. Fayet, “Phenomenology of the production, decay, and detection of new hadronic states associated with supersymmetry,” *Physics Letters B*, vol. 76, no. 5, pp. 575–579, 1978.
- [55] B. Allanach and B. Gripaios, “Hide and seek with natural supersymmetry at the LHC,” *Journal of High Energy Physics*, vol. 2012, no. 5, pp. 1–25, 2012.
- [56] N. Arkani-Hamed, S. Dimopoulos, and G. Dvali, “The hierarchy problem and new dimensions at a millimeter,” *Physics Letters B*, vol. 429, no. 3, pp. 263–272, 1998.
- [57] O. Klein, “Zur Fünfdimensionalen darstellung der relativitätstheorie,” *Zeitschrift für Physik*, vol. 46, no. 3-4, pp. 188–208, 1928.
- [58] G. Hinshaw, D. Larson *et al.*, “Nine-year Wilkinson Microwave Anisotropy Probe (WMAP) observations: cosmological parameter results,” *The Astrophysical Journal Supplement Series*, vol. 208, no. 2, p. 19, 2013.

- [59] G. Cottin, M. A. Díaz, M. J. Guzmán, and B. Panes, “Gravitino dark matter in split supersymmetry with bilinear R-parity violation,” *The European Physical Journal C*, vol. 74, no. 11, pp. 1–17, 2014.
- [60] B. Batell, J. Pradler, and M. Spannowsky, “Dark matter from minimal flavor violation,” *Journal of High Energy Physics*, vol. 2011, no. 8, pp. 1–21, 2011.
- [61] A. Djouadi, J.-L. Kneur, and G. Moultaka, “SuSpect: A Fortran code for the supersymmetric and Higgs particle spectrum in the MSSM,” *Computer Physics Communications*, vol. 176, no. 6, pp. 426–455, 2007.
- [62] G. L. Kane, C. F. Kolda, L. Roszkowski, and J. D. Wells, “Study of constrained minimal supersymmetry,” *Phys. Rev.*, vol. D49, p. 6173, 1994.
- [63] M. W. Cahill-Rowley, J. L. Hewett, A. Ismail, and T. G. Rizzo, “More energy, more searches, but the phenomenological MSSM lives on,” *Physical Review D*, vol. 88, no. 3, p. 035002, 2013.
- [64] ATLAS Collaboration, “Search for Microscopic Black Holes in Multi-Jet Final States with the ATLAS Detector at $\sqrt{s} = 7$ TeV,” CERN, Geneva, Tech. Rep. ATLAS-CONF-2011-068, May 2011.
- [65] S. Kraml, S. Kulkarni *et al.*, “SModelS: a tool for interpreting simplified-model results from the LHC and its application to supersymmetry,” *The European Physical Journal C*, vol. 74, no. 5, pp. 1–23, 2014.
- [66] ATLAS Collaboration, “Search for squarks and gluinos with the ATLAS detector in final states with jets and missing transverse momentum using $\sqrt{s} = 8$ TeV proton–proton collision data,” *JHEP*, vol. 1409, p. 176, 2014.
- [67] ATLAS Collaboration, “Search for new phenomena using large jet multiplicities and missing transverse momentum with ATLAS in 5.8 fb^{-1} of $\sqrt{s} = 8$ TeV proton–proton collisions,” CERN, Geneva, Tech. Rep. ATLAS-COM-CONF-2012-142, Aug 2012.
- [68] L. J. Hall, D. Pinner, and J. T. Ruderman, “A natural SUSY Higgs near 125 GeV,” *Journal of High Energy Physics*, vol. 2012, no. 4, pp. 1–25, 2012.
- [69] S. Dimopoulos, K. Howe, and J. March-Russell, “Maximally Natural Supersymmetry,” *arXiv:1404.7554*, 2014.
- [70] J. Adam, X. Bai *et al.*, “New constraint on the existence of the $\mu \rightarrow e + \gamma$ decay,”

- Physical Review Letters*, vol. 110, no. 20, p. 201801, 2013.
- [71] F. Gabbiani, E. Gabrielli, A. Masiero, and L. Silvestrini, “A complete analysis of FCNC and CP constraints in general SUSY extensions of the standard model,” *Nuclear Physics B*, vol. 477, no. 2, pp. 321–352, 1996.
- [72] ATLAS Collaboration, “The ATLAS experiment at the CERN large hadron collider,” *Journal of Instrumentation*, vol. 3, no. 08, p. S08003, 2008.
- [73] L. Evans and P. Bryant, “LHC machine,” *Journal of Instrumentation*, vol. 3, no. 08, p. S08001, 2008.
- [74] S. Chatrchyan *et al.*, “The CMS experiment at the CERN LHC,” *JINST*, vol. 3, p. S08004, 2008.
- [75] ATLAS Luminosity Public Results, <https://twiki.cern.ch/twiki/bin/view/AtlasPublic/LuminosityPublicResults>.
- [76] C. Doglioni, *Measurement of the inclusive jet cross section with the ATLAS detector at the Large Hadron Collider*. Springer, 2012.
- [77] ATLAS Collaboration, “Improved luminosity determination in pp collisions at $\sqrt{s} = 7$ TeV using the ATLAS detector at the LHC,” *Eur.Phys.J.*, vol. C73, p. 2518, 2013.
- [78] ATLAS Photos, <http://www.atlas.ch/photos/index.html>.
- [79] A. Sfyrla, “ATLAS triggering on SUSY in 2012,” CERN, Geneva, Tech. Rep. ATL-DAQ-PROC-2013-001, Jan 2013.
- [80] M. Cacciari, G. P. Salam, and G. Soyez, “The Anti- $k(t)$ jet clustering algorithm,” *JHEP*, vol. 0804, p. 063, 2008.
- [81] F. Rubbo, “The Design and Performance of the ATLAS jet trigger,” *EPJ Web of Conferences*, vol. 60, p. 20053, 2013.
- [82] S. Agostinelli *et al.*, “GEANT4: A simulation toolkit,” *Nucl. Instrum. Meth.*, vol. A506, pp. 250–303, 2003.
- [83] ATLAS Collaboration, “The ATLAS Simulation Infrastructure,” *The European Physical Journal C*, vol. 70, no. 3, pp. 823–874, 2010. [Online]. Available: <http://dx.doi.org/10.1140/epjc/s10052-010-1429-9>

- [84] E. Richter-Was, L. Poggioli, and D. Froidevaux, “ATLFAST 2.0 a fast simulation package for ATLAS,” CERN-ATL-PHYS-98-131, Tech. Rep., 1998.
- [85] ATLAS Collaboration, “Electron reconstruction and identification efficiency measurements with the ATLAS detector using the 2011 LHC proton-proton collision data,” *Eur.Phys.J.*, vol. C74, no. 7, p. 2941, 2014.
- [86] ATLAS Collaboration, “Electron and photon energy calibration with the ATLAS detector using LHC Run 1 data,” *Eur.Phys.J.*, vol. C74, no. 10, p. 3071, 2014.
- [87] ATLAS Collaboration, “Measurement of the muon reconstruction performance of the ATLAS detector using 2011 and 2012 LHC proton-proton collision data,” *Eur.Phys.J.*, vol. C74, no. 11, p. 3130, 2014.
- [88] M. Cacciari, G. P. Salam, and G. Soyez, “FastJet user manual.”
- [89] S. Catani, Y. L. Dokshitzer, M. H. Seymour, and B. R. Webber, “Longitudinally-invariant $k_{\perp}$ -clustering algorithms for hadron-hadron collisions,” *Nuclear Physics B*, vol. 406, no. 1, pp. 187–224, 1993.
- [90] S. D. Ellis and D. E. Soper, “Successive combination jet algorithm for hadron collisions,” *Physical Review D*, vol. 48, no. 7, p. 3160, 1993.
- [91] Y. L. Dokshitzer, G. Leder, S. Moretti, and B. Webber, “Better jet clustering algorithms,” *Journal of High Energy Physics*, vol. 1997, no. 08, p. 001, 1997.
- [92] M. Wobisch and T. Wengler, “Hadronization corrections to jet cross sections in deep-inelastic scattering,” *arXiv preprint hep-ph/9907280*, 1999.
- [93] ATLAS Collaboration, “Jet energy measurement and its systematic uncertainty in proton-proton collisions at $\sqrt{s} = 7$ TeV with the ATLAS detector,” *The European Physical Journal C*, vol. 75, no. 1, 2015.
- [94] M. Cacciari and G. P. Salam, “Pileup subtraction using jet areas,” *Phys. Lett.*, vol. B659, p. 119, 2008.
- [95] ATLAS Collaboration, “Pile-up subtraction and suppression for jets in ATLAS,” CERN, Geneva, Tech. Rep. ATLAS-CONF-2013-083, Aug 2013.
- [96] S. Batista, M. Begel *et al.*, “Global Sequential Calibration with the ATLAS Detector in Proton-Proton Collisions at $\sqrt{s} = 8$ TeV with ATLAS 2012 data,” CERN, Geneva, Tech. Rep. ATL-COM-PHYS-2014-753, Jun 2014.

- [97] ATLAS Collaboration, “Expected performance of the ATLAS experiment - detector, trigger and physics,” CERN-OPEN-2008-020.
- [98] ATLAS Collaboration, “Commissioning of the ATLAS high-performance b-tagging algorithms in the 7 TeV collision data,” CERN, Geneva, Tech. Rep. ATLAS-CONF-2011-102, Jul 2011.
- [99] ATLAS Collaboration, “Performance of Missing Transverse Momentum Reconstruction in ATLAS studied in Proton-Proton Collisions recorded in 2012 at 8 TeV,” CERN, Geneva, Tech. Rep. ATLAS-CONF-2013-082, Aug 2013.
- [100] ATLAS Collaboration, “Performance of Missing Transverse Momentum Reconstruction in Proton-Proton Collisions at 7 TeV with ATLAS,” *Eur.Phys.J.*, vol. C72, p. 1844, 2012.
- [101] A. Collaboration *et al.*, “ATLAS detector and physics performance Technical Design Report,” *CERN/LHCC*, vol. 15, p. 1999, 1999.
- [102] P. Svantesson, A. Belyaev, J. Hall, and S. F. King, ““E6SSM vs MSSM gluino phenomenology”,” in *EPJ Web of Conferences*, vol. 28. EDP Sciences, 2012, p. 12014.
- [103] J. A. Maxin, T. Li, D. V. Nanopoulos, and J. W. Walker, “The Ultra-High Jet Multiplicity Signal of Stringy No-Scale F-SU (5) at the $\sqrt{s}=7$ TeV LHC,” *arXiv:1103.4160*, 2011.
- [104] D. Feldman, G. Kane, E. Kuflik, and R. Lu, “A new (string motivated) approach to the little hierarchy problem,” *Physics Letters B*, vol. 704, no. 1, pp. 56–61, 2011.
- [105] ATLAS Collaboration, “Search for new phenomena in final states with large jet multiplicities and missing transverse momentum using $\sqrt{s}=7$ TeV pp collisions with the ATLAS detector,” *Journal of High Energy Physics*, vol. 2011, no. 11, pp. 1–38, 2011.
- [106] ATLAS Collaboration, “Search for squarks and gluinos using final states with jets and missing transverse momentum with the ATLAS detector in $\sqrt{s} = 7$ TeV proton-proton collisions,” *Phys.Lett.*, vol. B701, pp. 186–203, 2011.
- [107] ATLAS Collaboration, “Hunt for new phenomena using large jet multiplicities and missing transverse momentum with ATLAS in 4.7 fb^{-1} of $\sqrt{s}=7$ TeV proton-proton collisions,” *Journal of High Energy Physics*, vol. 2012, no. 7, pp. 1–40, 2012.

-
- [108] B. Allanach and B. Gripaios, “Hide and Seek With Natural Supersymmetry at the LHC,” *JHEP*, vol. 1205, p. 062, 2012.
- [109] M. Bahr *et al.*, “Herwig++ physics and manual,” *Eur. Phys. J.*, vol. C58, pp. 639–707, 2008.
- [110] T. Sjöstrand, S. Mrenna, and P. Z. Skands, “PYTHIA 6.4 physics and manual,” *JHEP*, vol. 05, p. 026, 2006.
- [111] M. Kramer, A. Kulesza *et al.*, “Supersymmetry production cross sections in pp collisions at $\sqrt{s} = 7$ TeV,” CERN, Tech. Rep., 2012.
- [112] A. Martin, W. Stirling, R. Thorne, and G. Watt, “Parton distributions for the LHC,” *Eur.Phys.J.*, vol. C63, pp. 189–285, 2009.
- [113] P. M. Nadolsky, H.-L. Lai *et al.*, “Implications of CTEQ global analysis for collider observables,” *Phys.Rev.*, vol. D78, p. 013004, 2008.
- [114] ATLAS Collaboration, “Performance of the Reconstruction and Identification of Hadronic Tau Decays in ATLAS with 2011 Data,” CERN, Geneva, Tech. Rep. ATLAS-CONF-2012-142, Oct 2012.
- [115] B. Nachman and C. G. Lester, “Significance Variables,” *Phys.Rev.*, vol. D88, p. 075013, 2013.
- [116] M. Baak, G. Besjes *et al.*, “HistFitter software framework for statistical data analysis,” *arXiv:1410.1280*, 2014.
- [117] K. Cranmer, A. Shibata *et al.*, “HistFactory: A tool for creating statistical models for use with RooFit and RooStats,” CERN, Tech. Rep. CERN-OPEN-2012-016, Jan 2012.
- [118] L. Moneta, K. Belasco, K. Cranmer, S. Kreiss, A. Lazzaro, D. Piparo, G. Schott, W. Verkerke, and M. Wolf, “The roostats project,” *arXiv:1009.1003*, 2010.
- [119] S. Catani, L. Cieri *et al.*, “Vector boson production at hadron colliders: A Fully exclusive QCD calculation at NNLO,” *Phys. Rev. Lett.*, vol. 103, p. 082001, 2009.
- [120] M. Cacciari, M. Czakon *et al.*, “Top-pair production at hadron colliders with next-to-next-to-leading logarithmic soft-gluon resummation,” *Physics Letters B*, vol. 710, no. 4, pp. 612–622, 2012.
- [121] M. Aliev, H. Lacker *et al.*, “HATHOR: HAdronic Top and Heavy quarks crOss

- section calculatoR,” *Comput. Phys. Commun.*, vol. 182, p. 1034, 2011.
- [122] M. Czakon and A. Mitov, “NNLO corrections to top-pair production at hadron colliders: the all-fermionic scattering channels,” *Journal of High Energy Physics*, vol. 2012, no. 12, pp. 1–18, 2012.
- [123] J. M. Butterworth, J. R. Forshaw, and M. H. Seymour, “Multiparton interactions in photoproduction at HERA,” *Z. Phys.*, vol. C72, p. 637, 1996.
- [124] B. P. Kersevan and E. Richter-Was, “The Monte Carlo event generator AcerMC versions 2.0 to 3.8 with interfaces to PYTHIA 6.4, HERWIG 6.5 and ARIADNE 4.1,” *Computer Physics Communications*, vol. 184, no. 3, pp. 919–985, 2013.
- [125] ATLAS Collaboration, “Jet energy resolution in proton-proton collisions at $\sqrt{s}=7$ TeV recorded in 2010 with the ATLAS detector,” *The European Physical Journal C*, vol. 73, no. 3, pp. 1–27, 2013.
- [126] ATLAS Collaboration, “Measurement of the b-tag Efficiency in a Sample of Jets Containing Muons with 5 fb^{-1} of Data from the ATLAS Detector,” CERN, Geneva, Tech. Rep. ATLAS-CONF-2012-043, Mar 2012.
- [127] ATLAS Collaboration, “Measuring the b-tag efficiency in a top-pair sample with 4.7 fb^{-1} of data from the ATLAS detector,” CERN, Geneva, Tech. Rep. ATLAS-CONF-2012-097, Jul 2012.
- [128] ATLAS Collaboration, “ b -jet tagging calibration on c -jets containing D^{*+} mesons,” CERN, Geneva, Tech. Rep. ATLAS-CONF-2012-039, Mar 2012.
- [129] ATLAS Collaboration, “Measurement of the Mistag Rate with 5 fb^{-1} of Data Collected by the ATLAS Detector,” CERN, Geneva, Tech. Rep. ATLAS-CONF-2012-040, Mar 2012.
- [130] ATLAS Collaboration, “Measurement of the cross section for the production of a W boson in association with b-jets in pp collisions at with the ATLAS detector,” *Physics Letters B*, vol. 707, no. 5, pp. 418–437, 2012.
- [131] G. Cowan, K. Cranmer, E. Gross, and O. Vitells, “Asymptotic formulae for likelihood-based tests of new physics,” *Eur. Phys. J.*, vol. C71, p. 1554, 2011.
- [132] A. Read, “Presentation of search results: the CL_s technique,” *Journal of Physics G: Nucl. Part. Phys.*, vol. 28, pp. 2693–2704, 2002.

- [133] ATLAS Collaboration, “Search for top and bottom squarks from gluino pair production in final states with missing transverse energy and at least three b jets with the ATLAS detector,” *Eur.Phys.J.*, vol. C72, p. 2174, 2012.
- [134] CMS Collaboration, “Search for supersymmetry in hadronic final states with missing transverse energy using the variables α_T and b -quark multiplicity in pp collisions at $\sqrt{s} = 8$ TeV,” *The European Physical Journal C*, vol. 73, no. 9, pp. 1–31, 2013.
- [135] CMS Collaboration, “Interpretation of searches for supersymmetry with simplified models,” *Phys. Rev. D*, vol. 88, p. 052017, Sep 2013.
- [136] ATLAS Collaboration, “Measurement of multi-jet cross sections in proton-proton collisions at a 7 TeV center-of-mass energy,” *Eur. Phys. J.*, vol. C71, p. 1763, 2011.
- [137] CMS Collaboration, “Distributions of topological observables in inclusive three- and four-jet events in pp collisions at $\sqrt{s} = 7$ TeV,” 2015.
- [138] D. Acosta *et al.*, “Comparison of Three-Jet Events in pp-bar Collisions at $\sqrt{s} = 1.8$ TeV to Predictions from a Next-to-leading Order QCD Calculation,” *Phys. Rev.*, vol. D71, p. 032002, 2005.
- [139] F. Abe *et al.*, “Study of four-jet events and evidence for double parton interactions in p anti-p collisions at $\sqrt{s} = 1.8$ TeV,” *Phys. Rev. D*, vol. 47, p. 4857, 1993.
- [140] D0 Collaboration, “Multiple Jet Production at Low Transverse Energies in pp Collisions at $\sqrt{s} = 1.8$ TeV,” *Phys. Rev. D*, vol. 67, p. 052001, 2003.
- [141] D0 Collaboration, “Ratios of multijet cross sections in $p\bar{p}$ collisions at $\sqrt{s} = 1.8$ TeV,” *Phys.Rev.Lett.*, vol. 86, pp. 1955–1960, 2001.
- [142] J. R. Andersen and J. M. Smillie, “High energy description of processes with multiple hard jets,” *Nuclear Physics B-Proceedings Supplements*, vol. 205, pp. 205–210, 2010.
- [143] ATLAS Good Run Lists for Analyses, <https://twiki.cern.ch/twiki/bin/view/AtlasProtected/GoodRunListsForAnalysis>.
- [144] G. D’Agostini, “A multidimensional unfolding method based on Bayes’ theorem,” *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 362, no. 2, pp. 487–498, 1995.
- [145] G. D’Agostini, “Improved iterative Bayesian unfolding,” *arXiv:1010.0632*, 2010.

- [146] T. Adye, “Unfolding algorithms and tests using RooUnfold,” *arXiv:1105.1160*, 2011.
- [147] ATLAS Collaboration, “Measurement of inclusive jet and dijet production in pp collisions at $\sqrt{s}=7$ TeV using the ATLAS detector,” *Physical Review D*, vol. 86, no. 1, p. 014022, 2012.
- [148] B. Efron, “Bootstrap methods: another look at the jackknife,” *The Annals of Statistics*, vol. 7, pp. 1–26, 1979.
- [149] G. Jogesh Babu, P. Pathak, and C. Rao, “Second order correctness of the Poisson bootstrap,” *Annals of statistics*, vol. 27, no. 5, pp. 1666–1683, 1999.
- [150] Bootstrap Generator: <https://svnweb.cern.ch/trac/atlasphys/browser/Physics/StandardModel/Common/BootstrapGenerator>.
- [151] T. Sjostrand, S. Mrenna, and P. Z. Skands, “A Brief Introduction to PYTHIA 8.1,” *Comput.Phys.Commun.*, vol. 178, pp. 852–867, 2008.
- [152] J. M. Katzy, “QCD Monte-Carlo model tunes for the LHC,” *Progress in Particle and Nuclear Physics*, vol. 73, pp. 141–187, 2013.
- [153] ATLAS Collaboration, “Summary of ATLAS Pythia 8 tunes,” ATL-PHYS-PUB-2012-003, Geneva, 2012.
- [154] S. Gieseke, C. Rohr, and A. Siodmok, “Colour reconnections in Herwig++,” *Eur.Phys.J.*, vol. C72, p. 2225, 2012.
- [155] H.-L. Lai, M. Guzzi *et al.*, “New parton distributions for collider physics,” *Phys. Rev.*, vol. D82, p. 074024, 2010.
- [156] J. Pumplin *et al.*, “New generation of parton distributions with uncertainties from global QCD analysis,” *JHEP*, vol. 07, p. 012, 2002.
- [157] ATLAS Collaboration, “ATLAS tunes of PYTHIA 6 and Pythia 8 for MC11,” ATL-PHYS-PUB-2011-009, 2011.
- [158] S. Höche, F. Krauss *et al.*, “Matching parton showers and matrix elements,” *hep-ph/0602031*, 2006.
- [159] T. Gleisberg *et al.*, “Event generation with SHERPA 1.1,” *JHEP*, vol. 02, p. 007, 2009.
- [160] F. Krauss, R. Kuhn, and G. Soff, “AMEGIC++ 1.0, a matrix element generator in

- C++,” *Journal of High Energy Physics*, vol. 2002, no. 02, p. 044, 2002.
- [161] T. Gleisberg and F. Krauss, “Automating dipole subtraction for QCD NLO calculations,” *The European Physical Journal C*, vol. 53, no. 3, pp. 501–523, 2008.
- [162] H. Beauchemin, M. Bona *et al.*, “A measurement of the W production cross section and the ratio of W to Z production cross section in association with jets with the ATLAS detector,” CERN, Geneva, Tech. Rep. ATL-COM-PHYS-2013-590, May 2013.
- [163] J. R. Andersen and J. M. Smillie, “Constructing all-order corrections to multi-jet rates,” *Journal of High Energy Physics*, vol. 2010, no. 1, pp. 1–35, 2010.
- [164] J. R. Andersen and J. M. Smillie, “Factorization of the t-channel pole in quark-gluon scattering,” *Physical Review D*, vol. 81, no. 11, p. 114021, 2010.
- [165] S. Badger, private communication.
- [166] A. Buckley, “The hepthesis L^AT_EX class.”