

SCHOOL OF PHYSICS, CHEMISTRY AND EARTH SCIENCES
FACULTY OF SCIENCES, ENGINEERING AND TECHNOLOGY

Aspects of Hadronic Final States and a Charming Search for New Physics with the ATLAS Detector



THE UNIVERSITY
of ADELAIDE

Edmund TING

Supervisors:

Prof. Martin WHITE

Prof. Paul JACKSON

Submitted towards the degree of:

Doctor of Philosophy

7th June, 2025

In memory of my grandfather

25/02/2025

TABLE OF CONTENTS

I	Prelude	1
1	The Subatomic Universe	3
1.1	Mathematical foundation of quantum theory	4
1.2	Towards a quantum field theory	5
1.3	Gauge symmetries	6
1.4	Standard Model	8
1.4.1	Quantum chromodynamics	9
1.4.2	Electroweak theory	10
1.4.3	Higgs mechanism	11
2	Beyond the Standard Model	13
2.1	Supersymmetry	14
2.1.1	An introduction to everything super	15
2.1.2	Minimal Supersymmetric Standard Model	16
2.1.3	Simplified models	18
2.1.4	Non-minimal flavour violation	18
2.2	Leptoquarks	19
2.2.1	Scalar leptoquark model	20
2.2.2	Vector U_1 models	20
3	From Theory to Experiment	21
3.1	Perturbation theory	21
3.2	Feynman diagrams	22
3.3	Cross-sections and decay rates	23
3.4	Event generation	24
3.5	Large Hadron Collider	25
3.5.1	Luminosity and event rate	26
3.5.2	Underlying event and pile-up	27
3.5.3	History and timeline	28
4	The ATLAS Experiment	31
4.1	ATLAS detector	31
4.1.1	Inner detector	32
4.1.2	Calorimeters	34
4.1.3	Muon spectrometer	36

4.2	Trigger and data acquisition	37
4.3	Signal reconstruction	38
4.3.1	Inner detector tracks	39
4.3.2	Calorimeter clusters	40
4.3.3	Muon chamber tracks	42
4.4	Physics objects	43
4.4.1	Hadronic final states	43
4.4.2	Electrons and photons	46
4.4.3	Muons	47
4.4.4	Taus	48
4.4.5	Missing transverse momentum and significance	49
4.4.6	Heavy flavour decays	50
4.5	Changes for Run 3	52
4.6	Towards the High-Luminosity era	53
4.7	Rise of the Machines	53
5	Particle Flow	55
5.1	Particle flow in ATLAS	55
5.1.1	Expected energy deposited by particles	56
5.1.2	Input objects	57
5.1.3	Track-cluster matching	57
5.1.4	Split shower recovery	58
5.1.5	Cell-based subtraction	58
5.1.6	Remnant removal	59
5.1.7	Post-subtraction refinements	59
5.2	Tracks for jet substructure	60
5.3	Unified Flow Objects	61
5.4	Global particle flow	62
5.5	Advancements and developments	62
II	On Aspects of Hadronic Final States	65
6	Jet-Object Overlap Removal	67
6.1	Simulated event samples	68
6.2	Object definitions and event selection	69
6.3	Preliminary analysis	69
6.4	Jet-photon overlaps in γ + jets	71
6.5	Jet-electron overlaps in $Z(\rightarrow ee)$ + jets	72
6.6	Jet-muon overlaps in $Z(\rightarrow \mu\mu)$ + jets	74
6.6.1	Subleading muon tracks in Particle Flow	74
6.6.2	Overlap removal performance comparison	75
6.7	Jet-lepton overlaps in $t\bar{t}$ events	77
6.8	Summary	78
6.A	Additional plots	79

7	Link-based Overlap Removal	83
7.1	First look at GPF links	83
7.2	Identifying true hadronic jets	86
7.3	Prospective GPF-based procedure	87
7.4	Jet–muon overlaps	88
7.5	Outlook	89
8	E/p Inside Jets	91
8.1	Dependence of E/p on $\Delta R(\text{jet}, \text{track})$	92
8.2	Comparisons using a simple Particle Flow	93
8.2.1	Deriving $\langle E/p \rangle$ reference values	95
8.2.2	Comparisons of $\langle E/p \rangle$ configurations	96
8.2.3	Comparisons to an ideal Particle Flow	98
8.3	Comparisons in Athena	101
8.3.1	Alternative ΔR -dependent $\langle E/p \rangle$ parametrisations	102
8.3.2	Modifications to eflowRec	103
8.3.3	Results	104
8.4	Conclusion	105
8.A	Additional plots	106
8.B	Simple vs ATLAS Particle Flow	110
8.C	Track–cluster resolution comparison	112
9	Overlapping Jets	115
9.1	Event samples and object definitions	115
9.2	Reconstructed jet categories	116
9.3	Effect of hard-scatter overlaps with pile-up	119
9.3.1	Overlaps with in-time pile-up	119
9.3.2	Overlaps with out-of-time pile-up	121
9.4	Outlook	122
III	A Search for Charms and Invisibles	125
10	Motivation and Setup	127
10.1	Data and simulated event samples	129
10.2	Event reconstruction	130
11	Analysis Regions	133
11.1	Recursive Jigsaw Reconstruction	134
11.2	Signal regions	135
11.2.1	High-Mass	137
11.2.2	Compressed	138
11.3	Background estimation and the control regions	140
11.3.1	Z + jets control regions	140
11.3.2	W + jets control regions	142
11.4	Validation of background estimates	144

12 Results	149
12.1 Systematic uncertainties	150
12.2 Results and interpretation	151
12.3 Conclusion	158
 Auxiliary Material	 161
A Signal acceptances and efficiencies	163
B Cutflow tables	167
C Auxiliary material for leptoquark models	171
D Charm-tagging efficiency	184
E Exclusion contours for individual signal regions	185
F Best expected signal region combination	189
G CL_s values	189
H Upper cross-section limits	190

ABSTRACT

This thesis presents a collection of studies related to the reconstruction of hadronic final states and a search for new physics with the ATLAS detector. This includes an investigation on ambiguities between jet and object reconstruction using traditional distance-based approaches and in the paradigm of Global Particle Flow; an alternative parametrisation of the expected energy deposited by charged particles in the calorimeters in the context of particle flow reconstruction; and the potential consequences of increased pile-up conditions at the High-Luminosity Large Hadron Collider. The search presented pertains to the pair production of charm or top squarks, each decaying to a charm quark and the lightest supersymmetric particle ($\tilde{\chi}_1^0$) to produce a final state containing missing transverse momentum and jets originating from charm quarks. This was conducted using the full Run 2 dataset of $\sqrt{s} = 13$ TeV proton–proton collisions collected by the ATLAS detector, corresponding to an integrated luminosity of 139 fb^{-1} . No significant excess of events over the expected Standard Model backgrounds is observed, and the pair production of the supersymmetric particles with masses up to 900 GeV is excluded in scenarios where the mass of the $\tilde{\chi}_1^0$ is below 50 GeV. The production of leptoquarks with masses up to 900 GeV is also excluded for the scenario where up-type leptoquarks decay into a charm quark and a neutrino.

STATEMENT OF ORIGINALITY

I certify that this work contains no material which has been accepted for the award of any other degree or diploma in my name in any university or other tertiary institution and, to the best of my knowledge and belief, contains no material previously published or written by another person, except where due reference has been made in the text. In addition, I certify that no part of this work will, in the future, be used in a submission in my name for any other degree or diploma in any university or other tertiary institution without the prior approval of the University of Adelaide and where applicable, any partner institution responsible for the joint award of this degree.

I give permission for the digital version of my thesis to be made available on the web, via the University's digital research repository, the Library Search and also through web search engines, unless permission has been granted by the University to restrict access for a period of time.

I acknowledge the support I have received for my research through the provision of an Australian Government Research Training Program Scholarship.

Edmund Ting

7th June, 2025

ACKNOWLEDGEMENTS

Firstly, a thank you to is clearly owed to my supervisors—Martin, ever the role model in pragmatism, and Jack, who has been invaluable for his expertise in RJR and SUSY searches at ATLAS. They have each supported me in various ways during my candidature, from providing physics insights into quirky distributions and teaching me the in-and-outs of ATLAS, to dealing with numerous administrative forms and paperwork on my behalf.

My technical supervisor, Chris, introduced me to the world of jets, and she was immensely helpful also with her expertise of the ATLAS software (despite her occasional claims of not being an expert). I am also grateful to Max and Mark for always being available on Mattermost to answer my questions about particle flow and its code implementation. A thank you also to Ariel, who I met during the Hadronic Final State Forum in Louisiana and who subsequently took on a supervisory role on two separate projects (which eventually turned into two chapters in this thesis). He imparted to me not only the finer philosophies of jet reconstruction, but also taught me much about the intricacies of the ATLAS tracker and calorimeters.

I also worked on the SUSY cc+MET analysis, and am grateful to all of the team members, especially Vangelis for his leadership and for teaching me about many aspects of analysis design. I really appreciate our chats on Slack and the impromptu calls to discuss finer technical details on anything from debugging code to statistics and uncertainty estimation.

The Adelaide ATLAS cohort, including those who have since graduated: this has now grown to be quite the large group so I will not list all the names, but it's been a fun place to work. Particular thanks to “sous-chef” Kong for cooking up the spiciest distractions in the office.

Last, but definitely not least: Silvana, Emily, and Sharon, for taking much of the stress out of travelling internationally and getting reimbursed.

Preface

During my candidature, I worked on several studies within the Jet/Etmiss and SUSY working groups in ATLAS. To facilitate the presentation of these distinct topics, this thesis is split into three parts. For the same reason, I will forgo a common “Conclusion” chapter in this thesis, with concluding remarks instead being presented below, and at the end of each chapter or part as apt.

Part I serves as a prelude to my own work. My vision for this part was to provide a bottom-up introduction to the theory of particle physics, covering a breadth wide enough to be relatively self-contained and accessible for future reference, while remaining unfettered by too many technical details or proofs that might obfuscate the discussion. Chapter 1 begins with a presentation of basic quantum theory and how it leads to the construction of the Standard Model. Chapter 2 continues with a discussion of the shortcomings of the Standard Model, and introduces a few supersymmetric extensions and leptoquark models that will be relevant for the search presented later. The transition from theory to experiment begins in Chapter 3, where I will describe the way in which abstract theories can be used to extract experimentally measurable quantities. This will include an introduction to the Large Hadron Collider, along with several essential concepts related to collider experiments. Chapter 4 presents a general overview of the ATLAS experiment, and finally, Chapter 5 describes particle flow and its implementation in ATLAS.

Then, we come to Part II of the thesis, which contains a collection of studies I conducted within the Jet/Etmiss working group to advance our understanding of hadronic final states and their reconstruction. To be clear, all of the results presented in this part are of my own direct work, achieved under the supervision of several key members of the group to whom I am very grateful. This includes defining the data pipelines and project workflow, developing and running analysis code, and generating all of the plots presented. The chapters each cover a different project, with motivations and outcomes as follows.

- Chapter 6 summarises the work I did as part of my authorship task. The goal was to validate the existing overlap removal recommendations when using the particle flow jet definition, and determine whether any re-optimisation was needed. In the end, similar overlap removal performances were found with particle flow jets (barring an unavoidable issue related to muons in the Run 2 reconstruction software), and no changes to the overlap removal recommendations were implemented as a result of these studies.
- Chapter 7 contains a closely related study in which I investigated the potential of a Global Particle Flow approach to overlap removal. This study was motivated by the introduction of an explicit correspondence between detector signals and physics objects in the reconstruction software for Run 3, which offers an attractive solution for overlap removal. I found that a simple algorithm based on the cut of a single variable could already perform well (as measured by an efficiency for retaining true jets), in addition to automatically tak-

ing into consideration factors such as whether or not a lepton was produced in the decay of a heavy-flavour hadron. However, the performance I achieved with these first studies was at a level comparable to the current procedures, and so no update to the core software was implemented. Nevertheless, this work paves the way for future studies (some already in progress within the collaboration) into revisiting event reconstruction in ATLAS, and how a Global Particle Flow paradigm might be of benefit in Run 4 and beyond.

- Chapter 8 presents a study on alternative parametrisations of the energy expected to be deposited by charged particles in the calorimeter. This parametrisation is core to the ATLAS particle flow reconstruction algorithm, and I investigated the impact of this on the reconstructed jet response and resolution. In particular, this study was motivated by the calorimeter clustering algorithm used in ATLAS, which allows for low-energy deposits to be measured provided that they are close in proximity to a cell above the noise threshold. Indeed, I found a strong dependence of the deposited energy on the surrounding local energy density, which resulted in a small (but positive) improvement to the jet resolution. However, comparisons to an idealised particle flow algorithm showed that there is still ample room for improvement. An interesting avenue of future study would be to explore whether machine learning algorithms might be able to achieve this.
- Chapter 9 explores the pile-up dense environment of the High-Luminosity LHC, and in particular, the consequences of multiple hard-scatter interactions occurring within a single bunch crossing. This can lead to the production of pile-up that is at a comparable scale as the hard-scatter, and I conducted a first study to characterise their impact on the reconstructed jet scale and resolution. A significant shift in the jet scale was found in cases where the jet is reconstructed from both hard-scatter and hard pile-up particles, which induces a corresponding increase in the jet energy resolution if not correctly accounted for. Thus, it is clear that a dedicated strategy for identifying and subsequently calibrating such jets will be essential for ensuring a robust reconstruction of jets at the HL-LHC, and a natural progression of this work would be to study potential methods for doing so (further remarks on this topic will be included in the concluding section of this chapter).

Finally, Part III will present a search for supersymmetry in final states containing two charm-tagged jets and large missing transverse momentum. The material in this part is based on the corresponding ATLAS publication of ref. [1]. The differences are as follows: I have restructured the discussion into chapters, and added exposition in a few places (along with some stylistic edits) to better fit the format of a thesis. Of course, such an analysis is conducted as a group effort; below, I provide a list of my major contributions.

- Validation of charm tagger working points: identifying charm jets is tricky, and we expect two of them in the final state of this analysis. I conducted studies looking into the impact of requiring two c -tagged jets in our regions (particularly with statistical limitations in mind), and how different working points affect our sensitivity over the signal grid. These studies informed the decision to adopt the 20% efficiency working point in the final analysis.
- Preparation of signal model MC requests: this involved migrating a job options file to a newer version of MADGRAPH, validating event kinematics between the versions, and requesting for MC events to be generated via standard ATLAS channels.
- Contribution to the design of all Compressed signal, control, and validation regions: in particular, my studies into correlations between jet flavour and the ISR system led to the

development of the $N_{c\text{-jets}}^S$ variable, which turned out to be vital for extending the sensitivity reach in the compressed signal regime.

- Calculation of signal theory uncertainties: these included the factorisation and renormalisation scale uncertainties, obtained by evaluating expected yields over six pairs of scale variation factors, along with initial-state radiation uncertainties, which involved generating private MC samples (with var1, var2, and var3a/b/c varied up/down) of representative points in each of the signal grids and estimating a conservative uncertainty envelope for different mass splittings.
- Performing the scan over branching fractions for SUSY signals: this corresponds to the two results shown in Figure 12.4.
- Carrying out the leptoquark interpretation: from preparing signal MC requests to calculating theoretical uncertainties and developing the fitting framework. This effort culminated in the results shown in Figures 12.5 and 12.6.
- Preparation of auxiliary material: developing the truth analysis code in SimpleAnalysis, calculating acceptances and efficiencies in the signal regions, generating the cutflow tables, and preparing the HEPData record. This encompasses all of the material included in appendices A, B and C.

Overall, the search found no significant excess of events over the expected Standard Model background processes, and limits were placed on the production cross-section of the BSM particles, including a first set of limits from any LHC experiment on leptoquark models in this final state. Still, statistical constraints form one of the leading limitations in this analysis, partly owing to the current difficulty in efficiently identifying jets that contain charm hadrons. It would be interesting to revisit this region of parameter space in the future to augment it with the integrated luminosity of the Run 3 dataset, along with any improvements in charm tagging that might have developed by then—no doubt imminent, considering the constant advancements offered by the rapidly-evolving landscape of machine learning.

PRELUDE

The Subatomic Universe

 HE development of quantum theory began at the turn of the 20th century, with pioneers such as Planck, Einstein, de Broglie and Bohr proposing several models to explain phenomena (in particular, black-body radiation and the photoelectric effect) that had been observed in the late 1800s. Building upon the work of de Broglie, Schrödinger published his famous equation in 1926, which models the time-evolution of quantum systems and has since become the foundation of modern quantum mechanics. Dirac and von Neumann quickly followed over the next few years: the former in formulating a quantum theory that was compatible with relativity, which led to the development of quantum field theory (QFT); and the latter in constructing a mathematically rigorous description of quantum mechanics.

This all culminated in the 1940s when a theory of quantum electrodynamics (QED) was devised through the work of Feynman, Dyson, Schwinger and Tomonaga. Then, in the 1950s, the experimental observation of parity violation in weak interactions led to attempts to unify the theory of the electromagnetic (EM) and weak forces; this was finally achieved by Weinberg in the late 1960s. Around the same period of time, efforts were underway to develop a theory that could describe the strong nuclear force. Significant contributors on this front were Yang and Mills, who in the 1950s published their eponymous theory wherein isospin conservation was invoked in a manner analogous to charge conservation in QED. The idea was initially shelved, as it predicted the existence of new massless particles that saw no experimental corroboration. However, everything came together nicely in 1964, when Higgs; Englert and Brout; and Guralnik, Hagen and Kibble independently published three papers within the span of three months showing how the spontaneous breaking of a global symmetry in a gauge theory could produce massive gauge bosons from the mixing of massless Nambu–Goldstone and gauge bosons. With this final ingredient, the theory of quantum chromodynamics (QCD) was born in 1973, by Fritzsch, Leutwyler and Gell-Mann.

What results from this remarkable half-century or so of effort is the most complete theory to date for describing the nature of the subatomic world: the Standard Model (SM) of particle physics. The theory has seen many successes on the experimental front, with the gradual discovery of its predicted particle content over the decades ensuing its formulation. At the same time, there is reason to believe that the SM does not provide a full description of our universe, with several experiments (particularly in more recent years) having produced results that are contradictory to it. Further discussion of this topic is deferred to Chapter 2.

Mathematics is the language of physical theory; thus, I will begin in section 1.1 by introducing the backbone of quantum theory, with some motivation as to *why* it is necessary in the first

1.1. Mathematical foundation of quantum theory

place. This will transition into a discussion of free quantum fields in section 1.2 (the method for handling interactions will be elaborated upon in Chapter 3). Then, section 1.3 will introduce the concept of gauge symmetries and their role in the definition of any theory that seeks to describe the fundamental forces. Now, as a disclaimer, the material presented in these sections is not intended to be a complete or particularly rigorous treatment of the subjects, but rather a broad overview to serve as a quick reference. As such, I will omit lengthy proofs and overly-detailed descriptions of mathematical particulars, instead focusing on conveying the physical consequences of the theory. All of this will culminate in section 1.4 with a presentation of the SM in all its glory.

1.1 Mathematical foundation of quantum theory

In classical mechanics, the initial position and force vectors that act upon an object can be used to calculate its trajectory through space (to some level of uncertainty). However, subatomic particles appear to behave non-deterministically: a chief example is the double slit experiment, wherein a series of individual electrons is projected onto a screen through a pair of slits. Classically, any electron that is observed on the screen (let us call this event C) will be expected to have passed through one of the two slits (denoted as event A or B , depending on which slit it passed through). In other words, the distributive law holds in the classical scenario: if A or B is true and C is true, then one can infer that the system is in a state of $(A \text{ and } C)$ or $(B \text{ and } C)$. Experimentally, however, one instead observes an interference pattern, and thus, true-false statements fail the distributive law: $(A \text{ or } B) \text{ and } C \not\Rightarrow (A \text{ and } C) \text{ or } (B \text{ and } C)$. Any theory that seeks to describe quantum particles needs to be able to account for such behaviour.

Consider a physical system with some properties—*observables*—for which a set of real values constitute a particular *state*. The set of all possible states of the system forms a *state space*, which in classical Hamiltonian theory is a symplectic manifold $\mathcal{M}$, with observables being real-valued functions on $\mathcal{M}$. To stay consistent with the double slit experiment example above, let us consider observables that correspond to true-false statements in what follows. This can always be done in general: instead of an observable being e.g. momentum p and its possible values in $\mathbb{R}$, i.e. $p : \mathcal{M} \rightarrow \mathbb{R}$, one considers the set corresponding to whether $p = a$ is true for $a \in \mathbb{R}$.

Now, in order to apply this construction to describe quantum phenomena, it turns out that the manifold has to be a *projective Hilbert space* $\mathbb{P}\mathcal{H}$, where the “projective” signifies that any pair of vectors in the Hilbert space that differ only up to a scalar factor can be thought of as being equivalent (i.e. representing the same state). The convention is to work with unit vectors in this space, such that the states are normalised. Then, true-false statements about the system correspond to closed linear subspaces of $\mathcal{H}$, and for a state represented by a unit vector $u = u_0 + u_1$, where u_1 is the component that lies in the range of its projection $\mathcal{R}(P)$ and u_0 the component in the nullspace $\mathcal{N}(P)$, one has that $\|u_0\|^2 + \|u_1\|^2 = 1$, and $\|u_1\|^2$ ($\|u_0\|^2$) can be interpreted as the probability that statement is true (false).

Given that state vectors are not necessarily orthogonal, what this means is that when two states u and v contain a non-zero component in the other’s projection, there is a *transition probability* $|\langle u|v \rangle|^2$ that the system can be found in one state or the other (in this case, the true-false observable is the statement that the system is in state u or v). This non-determinism is one of the unique features of quantum theory!

Under this construction, each true-false observable then corresponds to a projection-valued measure on $\mathbb{R}$; namely, a map from the values of the observable to the orthogonal projections

on $\mathcal{H}$. One important result of this, which drops out from the spectral theorem of functional analysis, is that observables can thus be associated with *self-adjoint operators* on $\mathcal{H}$. For such an operator A , the diagonal matrix elements $\langle u|A|u \rangle$ correspond to the expectation values of the observable to be in the state u .

Symmetries of a system are automorphisms of $\mathbb{P}\mathcal{H}$, i.e. bijective (one-to-one) transformations that leave the inner product invariant. By Wigner’s theorem, these will always correspond to *(anti-)unitary operators* on $\mathcal{H}$. Now, Stone’s theorem states that every strongly continuous one-parameter group of unitary operators can be written as

$$U(t) = e^{tS}, \quad (1.1)$$

where S is a skew-adjoint operator. However, note that self-adjoint ($A^\dagger = A$) and skew-adjoint ($S^\dagger = -S$) operators differ simply by a factor of an imaginary scalar ib , for $b \in \mathbb{R}$ (one conventionally invokes dimensionality arguments to set $b = -1/\hbar$, where $\hbar \approx 1.054 \times 10^{-34} \text{ kg m}^2 \text{ s}^{-1}$ is Planck’s constant). Thus, the above can also be written as

$$U(t) = e^{-itA/\hbar}, \quad (1.2)$$

and so one finds that observables (self-adjoint operators) are infinitesimal generators of the one-parameter groups of symmetries (unitary operators) of a system. The interpretation that it is the states that transform under these operators, $u \rightarrow u' = e^{-itA/\hbar}u$, with the observables held constant, is known as Schrödinger’s picture; while the interpretation in which states are fixed, with observables transforming as $B \rightarrow B' = U^\dagger B U = e^{itA/\hbar} B e^{-itA/\hbar}$, is Heisenberg’s picture.

1.2 Towards a quantum field theory

The introduction of the previous section forms the backbone of non-relativistic quantum mechanics. Early efforts of incorporating relativity considered systems of a fixed ensemble of particles, with attempts to render the Schrödinger equation Lorentz covariant leading to the Klein–Gordon and Dirac relativistic wave equations for describing free particles. However, these models were unsuccessful largely due to the experimental fact that particles can be created and annihilated, and partly also because they treated space and time on fundamentally different footing (position is an operator, while time is a parameter).

Thus, one might deduce that there are two ways forward: “promote” time into an observable (which eventually leads to string theory); or “demote” position into a parameter. It is the latter approach, which results in quantum field theory (QFT), that has been adopted in mainstream physics. The principle of QFT is to quantise the classical fields (sometimes referred to as “second quantisation”), rather than trying to quantise the wave function of a quantum particle (“canonical quantisation”). In this way, position returns to being a label on an ensemble of operators (i.e. a *field*) that are dependent on spacetime.

As an example, in QFT one would interpret the solutions of the (real) Klein–Gordon equation,

$$(\partial^2 + m^2)\varphi_{\text{clas}} = 0, \quad (1.3)$$

as a classical field instead of as the wave function of a quantum particle. Solutions can be found by presupposing the form $\varphi_{\text{clas}}(t, \mathbf{x}) = \sum q_j(t)f_j(\mathbf{x})$, where the coefficients $q_j(t)$ correspond to a classical harmonic oscillator. The act of quantising the theory then amounts to replacing

1.3. Gauge symmetries

these classical oscillators with quantum oscillators; in doing so, one obtains a quantum field $\varphi(t, \mathbf{x})$ that acts on a bosonic Fock space $F(\mathcal{H})$ (which is essentially a multi-particle state space constructed from the Hilbert space $\mathcal{H}$ of single particles).

The quantum field can be expressed as a Fourier integral of operators $a^\dagger(\mathbf{p})$ and $a(\mathbf{p})$ that respectively create and annihilate particles of momentum $\mathbf{p}$ at some position in spacetime, and obey commutation relations:

$$[a(\mathbf{p}), a^\dagger(\mathbf{p}')] = (2\pi)^3 \delta(\mathbf{p} - \mathbf{p}') I, \quad (1.4)$$

$$[a(\mathbf{p}), a(\mathbf{p}')] = [a^\dagger(\mathbf{p}), a^\dagger(\mathbf{p}')] = 0. \quad (1.5)$$

Thus, the quanta of the field are bosonic (spin-0) particles of mass m . One can do a similar exercise with the Dirac equation to find that it describes fermionic fields, with the creation and annihilation operators following anti-commutation relations instead (trying to enforce a commutation relation will produce an unphysical result of unbounded energy values).

The general procedure of quantising a field theory, then, is to start with a (classical) Lagrangian density, convert it into a Hamiltonian, then quantise the Hamiltonian into an operator. One feature of the QFT approach is that so long as one starts with a Lagrangian density that is Lorentz invariant, the theory will be relativistic from the outset. Any mathematically well-defined QFT satisfies a list of six conditions known as the Wightman axioms; for brevity, these will not be stated here, but one interesting result of this is the CPT theorem, which states that the theory will be invariant under the combined application of charge (C) conjugation, parity (P) inversion, and time (T) inversion operators (in any order). Initially, this was thought to be a trivial statement, as the belief was that C, P and T invariance each held separately; however, this viewpoint has changed with the experimental observation of both P and CP violation.

1.3 Gauge symmetries

“Gauge symmetry” is a term that pre-dates QFT—it was first introduced by Weyl to allow for a comparison of tangent vectors at different points on a spacetime manifold. This is the usage that relates to parallel transport from general relativity. Weyl suggested that 1-forms on the spacetime manifold could be interpreted as the EM potential, with their exterior derivative corresponding to the EM field. The choice of metric (i.e. “gauge”) at each point in spacetime then influences the EM potential, with a gauge transformation $g \rightarrow e^\chi g$ being accompanied by a “gauge” transformation of the potential, $\varphi \rightarrow \varphi - d\chi$. This thus produces a theory that unifies gravity and EM, and one often sees the common description that gauge degrees of freedom refer to those that can be arbitrarily chosen with no effect on the physical system described by the theory, stemming from the freedom of choice for the particular form of the spacetime metric.

While there were issues with the physical implications of Weyl’s formulation (it suggests that the EM field at some point in spacetime depends on the path taken to reach that point), there is nothing mathematically wrong with the idea. Indeed, it turned out that this principle of gauge invariance was physically relevant to quantum theory, which began to develop around a decade after Weyl’s original work. In particular, gauge theory was found to be necessary in order to define a renormalisable theory for vector fields [2].

In this application, quantum fields are sections of a vector bundle V associated with a (trivial) principal bundle $P = G \times \mathbb{R}^4$ over spacetime $\mathbb{R}^4$. Here, G is a compact subgroup—the “gauge” group, in the sense that it measures a “length” of spacetime—of the general linear group $GL(n, \mathbb{C})$,

and acts on the fields in a way that leaves the Lagrangian invariant, i.e. $\mathcal{L}(g \cdot \Phi, \partial(g \cdot \Phi)) = \mathcal{L}(\Phi, \partial\Phi)$ for $g \in G$. However, there is no unique choice for the trivialisation of P , and this is the reason for gauge invariance: a gauge field, in a particular gauge, simply corresponds to a particular trivialisation of P .

This becomes manifest in the Lagrangian of a theory when the symmetries are *local* in space-time (thus, in particle physics, gauge symmetry is often used to refer specifically to a local symmetry). Terms in the Lagrangian that involve a derivative will expand as

$$\partial_\mu(g(x) \cdot \Phi(x)) = g(x) \cdot \partial_\mu \Phi(x) + \partial_\mu g(x) \cdot \Phi(x), \quad (1.6)$$

yielding what appears to be an extra term that would break the invariance of the Lagrangian. However, note that the notion of a derivative doesn't make sense when the elements $g(x) \in G$ are spacetime dependent; one has to instead define a *covariant derivative*,

$$D_\mu = \partial_\mu + iA_\mu, \quad (1.7)$$

where iA_μ is known as the *gauge field*: a function that maps each global section σ of the principal bundle to a 1-form on $\mathbb{R}^4$, with values belonging to the Lie algebra $\mathfrak{g}$ of G . The gauge field transforms as $A_\mu \rightarrow gA_\mu g^{-1} + i(\partial_\mu g)g^{-1}$, such that the Lagrangian remains invariant under $\Phi(x) \rightarrow g(x) \cdot \Phi(x)$. The simultaneous transformations of $\Phi(x)$ and A_μ constitute a “gauge transformation”.

With these modifications, one thus obtains a Lagrangian where matter fields are coupled to gauge fields that describe vector particles. The final step is to introduce free-field terms for the gauge fields themselves, which should also respect Lorentz and gauge invariance. This can be done by considering the commutator of two covariant derivatives:

$$[D_\mu, D_\nu]\Phi = i(\partial_\mu A_\nu - \partial_\nu A_\mu + i[A_\mu, A_\nu])\Phi = iF_{\mu\nu}\Phi. \quad (1.8)$$

The $F_{\mu\nu}$ are $\mathfrak{g}$ -valued functions, and the requirement that $D_\mu D_\nu \Phi$ remains gauge invariant suggests that these transform as $F_{\mu\nu} \rightarrow gF_{\mu\nu}g^{-1}$. Thus, the Ad-invariant inner product $\langle F_{\mu\nu} | F^{\mu\nu} \rangle$ on $\mathfrak{g}$ is both Lorentz and gauge invariant, and a term that is proportional to this quantity can be used as the free-gauge field term in the Lagrangian. (The choice of coefficient is a matter of convention; different values amount to a rescaling of the gauge fields by a constant.) However, note that terms like $\langle A_\mu | A^\mu \rangle$ break gauge invariance; the field quanta are thus necessarily *massless*.

If the gauge group is simple or equal to the unitary group $U(1)$, the Ad-invariant product corresponds to the trace (up to some normalisation),

$$\langle F_{\mu\nu} | F^{\mu\nu} \rangle = -\text{tr}(F_{\mu\nu} F^{\mu\nu}). \quad (1.9)$$

Otherwise, if $\mathfrak{g} = \bigoplus_1^k \mathfrak{g}^j$, the inner product will include one normalisation factor for each of the simple or $U(1)$ factors:

$$\langle X | Y \rangle = -\sum \alpha_j^{-2} \text{tr}(X^j Y^j), \quad (1.10)$$

where $\alpha_j > 0$ and $X^j, Y^j \in \mathfrak{g}^j$. In these scenarios, the convention is typically to scale the gauge fields as $A_\mu^j \rightarrow \alpha_j A_\mu^j$ (where A_μ^j are the $\mathfrak{g}^j$ -valued functions, with $A_\mu = \sum A_\mu^j \rightarrow \sum \alpha_j A_\mu^j$), such that the normalisations instead appear in the form of the covariant derivative:

$$D_\mu = \partial_\mu + i \sum \alpha_j A_\mu^j. \quad (1.11)$$

1.4. Standard Model




Figure 1.1. The particles of the Standard Model. (*MissMf, Wikimedia Commons, cropped.*)

1.4 Standard Model

Let us now draw upon the concepts introduced in the previous sections to construct the SM, which provides a unified description of the EM, weak, and strong forces. Much like any physical theory, the development of the SM was driven and constrained by experiment—in particular, observations of the allowed interactions between the fundamental particles of our universe. However, the goal of the section is not so much to discuss the history of the SM, so let me present the known spectrum of fundamental particles from the outset.

A graphical depiction is shown in Figure 1.1, and can be summarised as follows: the particles can be categorised either as *fermions* or *bosons*, possessing half-integer or integer values of spin, respectively. The former includes the ordinary matter of our universe, comprising of three generations of *quarks* and *leptons*. The interactions between these particles operate under the exchange of the vector (spin-1) gauge bosons: the photon (γ) for the EM force, exchanged between electrically-charged particles; the $W^\pm$ and Z bosons for the weak force, between particles with weak charge (*weak isospin* and *hypercharge*); and the gluon (g) for the strong force, between particles with *colour* charge. Finally, there is the scalar (spin-0) Higgs boson, whose role will be discussed further in section 1.4.3.

So, how do these particles interact? Firstly, all of the matter particles are capable of interacting weakly; bar neutrinos, electromagnetically; and only quarks experience the strong force. The SM Lagrangian reflects this by introducing couplings between the fields as necessary. Massless vector bosons arise from imposing gauge invariance under $SU(3)_C \times SU(2)_L \times U(1)_{Y_W}$. The $SU(3)_C$ subgroup governs the theory of the strong force, with the subscript, C , denoting colour charge. The EM and weak forces are described in a unified electroweak (EW) theory under $SU(2)_L \times U(1)_{Y_W}$, where L indicates a coupling of the weak force to left-handed fields only, and Y_W is the weak hypercharge. Spontaneous symmetry breaking in the electroweak sector

yields the massive gauge bosons, as we will later see. For now, let us write the SM Lagrangian as a simple sum of QCD and EW contributions:

$$\mathcal{L}_{\text{SM}} = \mathcal{L}_{\text{QCD}} + \mathcal{L}_{\text{EW}}. \quad (1.12)$$

What follows will elucidate upon each of these two terms.

1.4.1 Quantum chromodynamics

QCD is non-Abelian gauge theory, invariant under transformations belonging to the gauge group $G = SU(3)$. The theory contains 18 quark fields (six flavours with three colours each), described by the components of six Dirac spinor fields $Q_f = (q_1^f, q_2^f, q_3^f)$, where f denotes the flavours, and the labels 1, 2 and 3 denote the colours, which are conventionally referred to as red, green, or blue (with antiquarks being antired, antigreen, or antiblue). The gauge fields of the theory produce the $\dim(\mathfrak{su}(3)) = 8$ massless gluons. These carry a colour-anticolour charge, but of the nine possible combinations, only eight are unique due to the constraint that $1\bar{1} + 2\bar{2} + 3\bar{3} = 0$.

The gauge invariant Lagrangian¹ can be written as discussed in section 1.3:

$$\mathcal{L}_{\text{QCD}} = -\frac{1}{4} G_{\mu\nu}^a G_{\mu\nu}^a + \sum_f \bar{Q}_f (i\gamma^\mu D_\mu - m_f) Q_f, \quad (1.13)$$

where the indices a span colour space, and $G_{\mu\nu}^a$ are the field strength tensors of the gluon fields:

$$G_{\mu\nu}^a = \partial_\mu G_\nu^a - \partial_\nu G_\mu^a + g_s f^{abc} G_\mu^b G_\nu^c. \quad (1.14)$$

Here, g_s corresponds to the strong coupling constant; f^{abc} are the structure constants of the symmetry group; and G_μ^a are the gluon gauge fields. The covariant derivative is given by

$$D_\mu = \partial_\mu + ig_s \frac{\lambda_a}{2} G_\mu^a, \quad (1.15)$$

with λ_a being the eight Gell-Mann matrices that span $\mathfrak{su}(3)$.

The basic interactions of QCD consist of vertices that connect two same-flavour quarks to a gluon, via the $\bar{Q}_f \gamma^\mu G_\mu^a Q_f$ terms in the Lagrangian; and three- or four-gluon vertices arising from the free gauge term, due to the non-Abelian nature of the gauge group. Strong interactions preserve colour, in that there needs to be an equal amount of each colour and its corresponding anticolour amongst the particles participating in the interaction.

There are two important aspects of QCD that warrant a brief discussion here: colour confinement and asymptotic freedom. Confinement refers to the experimental observation that all real quarks and gluons can only appear in colour-neutral bound states. These can be bosons of $q\bar{q}$ pairs known as *mesons*, or fermionic qqq states are known as *baryons* (together which are collectively referred to as *hadrons*). Colourless bound states of gluons—*glueballs*—are also theoretically possible, although only recently has there been any hint of such an object in an experiment [3].

Asymptotic freedom is the property wherein the effective coupling constant decreases with increasing energy (or equivalently, at shorter distances). This result is vital to the *modus operandi* of hadron collider experiments, which will be discussed in Chapter 3.

¹A mass term is explicitly included here; however, note that when QCD is considered in the context of the SM, quark masses actually arise via a coupling to the Higgs field, as we will see in section 1.4.3.

1.4. Standard Model

1.4.2 Electroweak theory

A theory that describes EW interactions has to account for the fact that weak interactions are *chiral*, i.e. only occurring between left-handed particles; and that neutrinos are uncharged under the EM force. There are two key operators involved in the theory, corresponding to the two forces. In the case of weak interactions, one takes inspiration from hadron theory, where due to their similarities, protons (p) and neutrons (n) can be considered as two states of a single particle (the “nucleon”). Then, the action of $SU(2)$ on the pair of fields (p, n) produce the so-called *isospin* $\pm 1/2$ eigenstates of the Pauli matrix. In analogy, EW theory thus defines the *weak isospin* operator, I_3 , which generates an action of $SU(2)$ that acts on the left-handed fermion fields. Then, for EM interactions, one defines the *weak hypercharge*² operator, $Y_W = 2(Q - I_3)$. This commutes with I_3 , and thus, one arrives at a combined action belonging to $SU(2)_L \times U(1)_{Y_W}$.

The Lagrangian of the theory can be split into several terms:

$$\mathcal{L}_{\text{EW}} = \mathcal{L}_{\text{gauge}} + \mathcal{L}_{\text{fermion}} + \mathcal{L}_{\text{Higgs}} + \mathcal{L}_{\text{Yukawa}}. \quad (1.16)$$

The first of these describes the pure gauge fields,

$$\mathcal{L}_{\text{gauge}} = -\frac{1}{4}W_{\mu\nu}^a W_a^{\mu\nu} - \frac{1}{4}B_{\mu\nu}B^{\mu\nu}, \quad (1.17)$$

with the field strength tensors now involving $\mathfrak{su}(2) \times \mathfrak{u}(1)$ -valued gauge fields:

$$W_{\mu\nu}^a = \partial_\mu W_\nu^a - \partial_\nu W_\mu^a - g_W \epsilon^{abc} W_\mu^b W_\nu^c, \quad (1.18)$$

$$B_{\mu\nu} = \partial_\mu B_\nu - \partial_\nu B_\mu, \quad (1.19)$$

where g_W is the weak coupling constant and ϵ^{abc} is the Levi-Civita symbol in three dimensions. The second term describes the interactions of the SM fermions (including quarks):

$$\begin{aligned} \mathcal{L}_{\text{fermion}} = & (\bar{u}_{iL}, \bar{d}_{iL}) i\gamma^\mu D_\mu \begin{pmatrix} u_{iL} \\ d_{iL} \end{pmatrix} + (\bar{\nu}_{iL}, \bar{l}_{iL}) i\gamma^\mu D_\mu \begin{pmatrix} \nu_{iL} \\ l_{iL} \end{pmatrix} \\ & + \bar{u}_{iR} i\gamma^\mu D_\mu u_{iR} + \bar{d}_{iR} i\gamma^\mu D_\mu d_{iR} + \bar{l}_{iR} i\gamma^\mu D_\mu l_{iR} + (\text{h.c.}), \end{aligned} \quad (1.20)$$

where $(u_{iL}, d_{iL})^\top$ and $(\nu_{iL}, l_{iL})^\top$ are $SU(2)_L$ doublets of left-handed quark and lepton fields, respectively, and u_{iR} , d_{iR} , and l_{iR} are $SU(2)_L$ singlets of right-handed fields. Note that explicit mass terms would involve a mixing of left- and right-handed fields, which would break gauge invariance; thus, the fermions are, at this moment, massless.

When acting on left-handed fields, the covariant derivative includes two coupling constants, one for each gauge field:

$$D_\mu = \partial_\mu + ig_W \frac{\sigma_a}{2} W_\mu^a + ig_Y \frac{Y_W}{2} B_\mu, \quad (1.21)$$

where σ_a are the three Pauli matrices that form a basis for $\mathfrak{su}(2)$. In the case of right-handed fields, the action of $\mathfrak{su}(2)$ vanishes, such that the covariant derivative effectively becomes:

$$D_\mu^{(R)} = \partial_\mu + ig_Y \frac{Y_W}{2} B_\mu. \quad (1.22)$$

²The reason why the electric charge, Q , is not used at this point (i.e. prior to spontaneous symmetry breaking) is due to the experimental observation that the $W^\pm$ bosons are electrically charged.

Now, the two terms that have been discussed thus far corresponds to a theory that contains massless fermions and gauge bosons (it is also clear, then, that the gauge fields in eq. (1.21) cannot correspond to the physical $W^\pm$ and Z gauge bosons). The role of the remaining terms in eq. (1.16) is precisely to admit a mechanism through which masses for the fermions and massive gauge bosons can be generated. Let us temporarily put on hold the present discussion to first describe how this works.

1.4.3 Higgs mechanism

The Higgs mechanism is the method through which masses are obtained for the SM vector gauge bosons (which are massless in a gauge invariant theory) and fermions (which are massless in a chiral theory). The procedure is essentially an application of Goldstone's theorem to theories in which gauge invariance has been imposed.

To aid the discussion, let me thus start with a brief sketch of this theorem using an example. Consider the Lagrangian:

$$\mathcal{L}(\Phi) = \frac{1}{2}(\partial_\mu \Phi) \cdot (\partial^\mu \Phi) - V(\Phi), \quad (1.23)$$

which remains invariant under some *global* symmetry group $G \subset O(n)$, where $\Phi = (\varphi_1, \dots, \varphi_n)$ is an n -tuple of scalar fields. The potential $V(\Phi) = V(\varphi_1, \dots, \varphi_n)$ is a polynomial on $\mathbb{R}^n$ that itself is invariant under G , with some minimum value $V_{\min}$. The idea then is to study what happens near a particular field configuration that minimises the potential—the *vacuum expectation value* (vev), which I will denote as ξ_0 in what follows.

If ξ_0 is invariant only under a *subset* $H = \{g \in G : g\xi_0 = \xi_0\} \subset G$ of the original symmetry group G , one says that the symmetry has been *spontaneously broken* (via the choice of ξ_0), and the symmetry group has been reduced from G to H . In such a case, setting $\tilde{\Phi} = \Phi - \xi_0$ and performing a Taylor expansion about the minimum configuration yields:

$$V(\Phi) = V(\tilde{\Phi} + \xi_0) = V_{\min} + \frac{1}{2} \sum_j m_j^2 \tilde{\varphi}_j^2 + [\text{higher order terms}], \quad (1.24)$$

where $\tilde{\varphi}_j$ are the components of $\tilde{\Phi}$ in an orthonormal eigenbasis for the Hessian $(\partial_i \partial_j V(\xi_0))$, with m_j^2 being the eigenvalues. However, the essence of Goldstone's theorem is that it is possible to choose a basis in which all the m_j vanish, with $\max(j) \geq \dim(G) - \dim(H)$. Then, what results in the post-symmetry breaking Lagrangian are terms like $(\partial_\mu \tilde{\varphi}_j)^2$, which look like a collection of at least $\dim(G) - \dim(H)$ massless free scalar fields—the *Goldstone bosons*.

Now, what if the Lagrangian were also invariant under a gauge transformation? As before, one can spontaneously break the global symmetry by choosing a particular minimum field configuration, before expanding the Lagrangian about this minimum to obtain some number of Goldstone fields. However, the key is that *it is possible to find a particular gauge in which all of these modes identically vanish*. In addition, during the exercise of expanding the Lagrangian, the $(D_\mu(\tilde{\Phi} + \xi_0))^2$ term will yield *mass terms* for the gauge fields, of the form $(A_\mu \xi_0) \cdot (A^\mu \xi_0)$. This is the Higgs mechanism—the Goldstone modes are no longer physical bosons, but instead act as the longitudinal modes of the initially massless gauge bosons of the theory, thus imparting mass to them.

With this preamble out of the way, let us return to the topic of the EW Lagrangian. The Higgs sector is given by

$$\mathcal{L}_{\text{Higgs}} = (D^\mu \Phi)^\dagger (D_\mu \Phi) - V(\Phi), \quad (1.25)$$

1.4. Standard Model

where Φ is an $SU(2)_L$ doublet of complex scalar fields,

$$\Phi = \begin{pmatrix} \varphi^+ \\ \varphi^0 \end{pmatrix} = \frac{1}{\sqrt{2}} \begin{pmatrix} \varphi_1 + i\varphi_2 \\ \varphi_3 + i\varphi_4 \end{pmatrix}, \quad (1.26)$$

and where the most general gauge invariant and renormalisable potential is of the form

$$V(\Phi) = \mu^2 \Phi^\dagger \Phi + \lambda (\Phi^\dagger \Phi)^2. \quad (1.27)$$

This corresponds to the famous Mexican hat potential of the Higgs field, which, subject to the constraints that $\mu^2 < 0$ and $\lambda > 0$, has a ring of degenerate minima at a radius of $v = \sqrt{-\mu^2/\lambda}$ (corresponding to the vev) around the origin of the complex plane.

Breaking the symmetry, we can take Φ to be of the form $(0, v + \varphi)^\top$ (which can be achieved through an appropriate gauge transformation), where φ is a real-valued field. The action of the covariant derivative (eq. (1.21)) on the field is then

$$\begin{aligned} D_\mu \Phi &= \begin{pmatrix} 0 \\ \partial_\mu \varphi \end{pmatrix} - \left[\frac{ig_W}{2} \begin{pmatrix} W_\mu^3 & W_\mu^1 - iW_\mu^2 \\ W_\mu^1 + iW_\mu^2 & -W_\mu^3 \end{pmatrix} + ig_Y \frac{Y_W}{2} B_\mu \right] \begin{pmatrix} 0 \\ v + \varphi \end{pmatrix} \\ &= -\frac{i}{2} \begin{pmatrix} v g_W (W_\mu^1 - iW_\mu^2) + g_W (W_\mu^1 - iW_\mu^2) \varphi \\ 2i\partial_\mu \varphi + v(-g_W W_\mu^3 + g_Y Y_W B_\mu) + (-g_W W_\mu^3 + g_Y Y_W B_\mu) \varphi \end{pmatrix}. \end{aligned} \quad (1.28)$$

Thus, the kinetic term in eq. (1.25) will yield new terms that are quadratic, cubic, and quartic in the gauge fields, the foremost of which constitute the mass terms of the physical gauge bosons, which we find to be given by the following linear combinations of the original gauge fields:

$$W_\mu^\pm = \frac{1}{\sqrt{2}} (W_\mu^1 \mp iW_\mu^2), \quad (1.29)$$

$$Z_\mu = W_\mu^3 \cos \theta_W - B_\mu \sin \theta_W, \quad (1.30)$$

$$A_\mu = W_\mu^3 \sin \theta_W + B_\mu \cos \theta_W, \quad (1.31)$$

with $W^\pm$ boson mass $m_W = v g_W / 2$, Z boson mass $m_Z = v \sqrt{g_W^2 + g_Y^2} / 2$, and Higgs boson mass $m_\varphi = -4\mu^2$ (the photon field A_μ remains massless, as expected). The weak mixing angle θ_W is related to the coupling constants via

$$\sin \theta_W = \frac{g_Y}{\sqrt{g_W^2 + g_Y^2}}. \quad (1.32)$$

The final term in eq. (1.16) deals with the Yukawa interaction of the SM fermions and the Higgs field:

$$\begin{aligned} \mathcal{L}_{\text{Yukawa}} &= -\Gamma_{ij}^u (\bar{u}_{iL}, \bar{d}_{iL}) \Phi^\dagger u_{jR} + \Gamma_{ij}^d (\bar{u}_{iL}, \bar{d}_{iL}) \Phi d_{jR} \\ &\quad + \Gamma_{ij}^l (\bar{\nu}_{iL}, \bar{l}_{iL}) \Phi l_{jR} + \Gamma_{ij}^\nu (\bar{\nu}_{iL}, \bar{l}_{iL}) \Phi^\dagger \nu_{jR} + (\text{h.c.}), \end{aligned} \quad (1.33)$$

where Γ_{ij}^f are the 3×3 matrix Yukawa couplings. Upon spontaneous symmetry breaking, this produces mass terms for the SM fermions (excluding neutrinos).

Beyond the Standard Model



WHILE the Standard Model (SM) has proven to be remarkably successful, it is generally accepted that the theory in its current state is not a complete one. The issues identified with the SM stem from a handful of experimental observations that are in contention with the phenomenology it predicts, as well as concerns regarding the technical niceties of the theory. These include:

- **the existence of dark matter:** only approximately 5% of universe is made up of the known matter of the SM, with the remainder comprising a mix of dark energy ($\sim 70\%$) and dark matter ($\sim 25\%$). The canonical motivation of the dark sector comes from astrophysical observations of galactic rotation curves: the rotational velocity of objects in spiral galaxies are found to be approximately independent of their radial position [4, 5], which is inconsistent with the observable matter in those galaxies. Thus, assuming that our understanding of gravity is not fundamentally flawed, this suggests the presence of invisible dark matter in the galaxies.

Another piece of evidence comes in the form of the cosmic microwave background (CMB), which is a background of microwave radiation that permeates all of our observable universe. Satellite maps of anisotropies in the CMB, which are sensitive to the presence of matter in the early universe, have facilitated calculations of the fraction of regular versus dark matter in the universe as quoted above [6].

The exact properties of dark matter remain a mystery: all we know is that it does not couple to the electromagnetic or strong forces (otherwise, we would be able to observe it easily). The potential DM particle could exist over a very wide mass range: from axions at the sub-eV scale, weakly-interacting massive particles (WIMPs) at the TeV scale, to primordial black holes at much larger mass scales.

- **non-zero neutrino masses:** the current formulation of the SM does not produce any mass terms for the neutrinos. However, neutrino flavour oscillations have now been observed in a multitude of experiments [7–12]. This is only possible if the neutrino flavour eigenstates are some linear combination of the mass eigenstates, and, crucially, if the latter *differs* for neutrinos of different flavours. The experimental results thus suggest that at least two of the three neutrinos have some small but non-zero mass.

One possible remedy for this is to introduce right-handed neutrinos into the SM such that they can gain mass in the Yukawa sector like all of the other fermions; however, the right-

2.1. Supersymmetry

handed neutrinos would be *sterile*—not participating in any interactions beyond those involving the Higgs field or gravity.

- **gravity and the hierarchy problem:** the SM currently does not incorporate gravity, which is generally fine as its effects are largely negligible at the microscopic level. Thus, one expects the theory to hold up to roughly the Planck scale, beyond which is when gravitational effects start to become significant. The core of the hierarchy problem is that the weak scale is many orders of magnitude ($\sim 10^{24}$) below the Planck scale.

One might posit that there are *surely* new particle(s) lying somewhere in this large mass range. However, this has implications that lead to the *Higgs fine-tuning problem*: given that the Higgs field couples to all massive particles within the current framework of the SM, it is uniquely sensitive (as the single SM scalar) to mass corrections, compared to the other SM fermions and bosons. The current measurements of the Higgs mass of around 125 GeV, then, is problematic: if there were new particles near the Planck scale, one would expect a Higgs mass at the order of the Planck mass ($\sim 10^{19}$ GeV). The fact that it appears to lie at the electroweak scale suggests very precise cancellations between any potential radiative corrections, which is not a particularly elegant justification (despite technically not being a fundamental flaw of the theory).

In this chapter, I will present a selection of attractive extensions to the SM (generally referred to as beyond-Standard Model (BSM) theories) that were formulated to address the shortcomings of the SM. Keeping to the theme of Part III, I will begin by discussing supersymmetry (SUSY¹) in section 2.1, before introducing leptoquark (LQ) models in section 2.2.

2.1 Supersymmetry

Let us consider the hierarchy problem in more detail: the 1-loop correction to the Higgs mass under the coupling to a new Dirac fermion, described by Lagrangian term $-\lambda_f H \bar{f} f$, is:

$$\Delta m_H^2 = -\frac{|\lambda_f|^2}{8\pi^2} \Lambda_{UV}^2 + \dots, \quad (2.1)$$

where Λ_{UV} is some ultraviolet (UV) momentum cutoff scale. Contrast this to the 1-loop correction when a new complex scalar couples to the Higgs field like $-\lambda_S |H|^2 |S|^2$. One finds that the quadratic correction in this case contains a minus sign relative to eq. (2.1):

$$\Delta m_H^2 = \frac{\lambda_S}{16\pi^2} \left[\Lambda_{UV}^2 - 2m_S^2 \log(\Lambda_{UV}/m_S) + \dots \right]. \quad (2.2)$$

Thus, the cancellations required to render the Higgs mass to be near the EW scale could be achieved by introducing a symmetry between fermions and bosons—this is the basis of SUSY.

Now, the SM contains external (Poincaré, or spacetime) and internal (gauge) symmetries. The idea behind SUSY is to extend regular spacetime into a “superspace”, and the Poincaré into a set of “super” Poincaré symmetries that contain additional supersymmetric transformations between fermionic and bosonic degrees of freedom. When elevating the SM into a SUSY model, all of the known particles (before electroweak symmetry breaking) gain a *superpartner*—each

¹I will use this shorthand when referring to the *theory* of supersymmetry. However, when talking of it as a *symmetry* of the theory (as we will soon see), it seems more appropriate to (and so I will) retain the full spelling of the word.

fermion a bosonic counterpart, with a naming convention that prefixes an *s* to the SM particle name (e.g. quarks $\rightarrow$ *squarks*); and for bosons, a fermionic partner suffixed with *-ino* (e.g. gluon $\rightarrow$ *gluino*).

In addition to the hierarchy problem, certain SUSY models are also able to provide natural solutions to the dark matter problem or a unification of gauge couplings at high energy. We shall see this later in section 2.1.2; but, before we get to that, let us take an introductory foray into the world of SUSY. What follows will be a summary of some key points from ref. [13].

2.1.1 An introduction to everything super

While supersymmetry can be formulated using the language of traditional quantum field theories (i.e. akin to that used in Chapter 1), there exists a more succinct and elegant set of notation for presenting the theory. This section will present a brief introduction of this specialised language, in which one often adds a preceding *super-* to the common words of QFT.

As already alluded to, the starting point is to extend the four spacetime coordinates x^μ into a *superspace* that additionally includes two complex anti-commuting two-component spinors θ^α and $\theta_{\dot{\alpha}}^\dagger$, where the indices α and $\dot{\alpha}$ are reserved specifically to denote the first and last two components of a Dirac spinor, respectively. (If an object possesses more than one such index, additional Greek letters will be introduced in alphabetical order.) Infinitesimal translations in this superspace correspond to a global supersymmetry transformation that transforms bosons into fermions and vice versa. This is generated by the following operators:

$$\hat{Q}_\alpha = i \frac{\partial}{\partial \theta^\alpha} - (\sigma^\mu \theta^\dagger)_\alpha \partial_\mu, \quad \hat{Q}^\alpha = -i \frac{\partial}{\partial \theta_\alpha} + (\theta^\dagger \bar{\sigma}^\mu)^\alpha \partial_\mu, \quad (2.3)$$

$$\hat{Q}^{\dagger\dot{\alpha}} = i \frac{\partial}{\partial \theta_{\dot{\alpha}}^\dagger} - (\bar{\sigma}^\mu \theta)^{\dot{\alpha}} \partial_\mu, \quad \hat{Q}_{\dot{\alpha}}^\dagger = -i \frac{\partial}{\partial \theta^{\dagger\dot{\alpha}}} + (\theta \sigma^\mu)_{\dot{\alpha}} \partial_\mu. \quad (2.4)$$

These operators act on *superfields* in superspace, and one can work out from the definitions above that they satisfy the following anti-commutator relations:

$$\{\hat{Q}_\alpha, \hat{Q}_{\dot{\beta}}^\dagger\} = -2\sigma_{\alpha\dot{\beta}}^\mu \hat{P}_\mu, \quad (2.5)$$

$$\{\hat{Q}_\alpha, \hat{Q}_\beta\} = \{\hat{Q}_{\dot{\alpha}}^\dagger, \hat{Q}_{\dot{\beta}}^\dagger\} = 0, \quad (2.6)$$

where $\hat{P}_\mu = -i\partial_\mu$ is the generator of translations in regular spacetime.

Due to the anti-commuting properties of θ and $\theta^\dagger$, the most general form for a superfield is

$$S(x, \theta, \theta^\dagger) = a + \theta\xi + \theta^\dagger\chi^\dagger + \theta\theta b + \theta^\dagger\theta^\dagger c + \theta^\dagger\bar{\sigma}^\mu\theta v_\mu + \theta^\dagger\theta^\dagger\theta\eta + \theta\theta\theta^\dagger\zeta^\dagger + \theta\theta\theta^\dagger\theta^\dagger d, \quad (2.7)$$

where (assuming that the superfield is Grassmann-even and carries no additional Lorentz or spinor indices) a, b, c, d and v_μ are boson fields, and $\xi, \chi^\dagger, \eta$ and $\zeta^\dagger$ are two-component fermion fields. These fields correspond to the components of single particle states, which in the irreducible representation of the supersymmetry algebra are known as *supermultiplets* that consist of a pair of bosonic and fermionic states considered as superpartners of one another.

A superfield Φ that satisfies $\bar{D}_{\dot{\alpha}}\Phi = 0$ is said to be (*left*-)chiral (and its complex conjugate *anti*- or *right*-chiral), where the chiral covariant derivatives are:

$$D_\alpha = \frac{\partial}{\partial \theta^\alpha} - i(\sigma^\mu \theta^\dagger)_\alpha \partial_\mu, \quad D^\alpha = -\frac{\partial}{\partial \theta_\alpha} + i(\theta^\dagger \bar{\sigma}^\mu)^\alpha \partial_\mu, \quad (2.8)$$

2.1. Supersymmetry

with anti-chiral covariant derivatives given by $\bar{D}_{\dot{\alpha}} S^* \equiv (D_{\alpha} S)^*$. In general, a chiral superfield Φ can be obtained from a general superfield S by computing $\Phi = \bar{D}_{\dot{\alpha}} \bar{D}^{\dot{\alpha}} S$ (and an anti-chiral superfield via $\Phi^* = D^{\alpha} D_{\alpha} S^*$). Alternatively, superfields that satisfy $V = V^*$ are known as *vector* or *real* superfields; a superfield cannot simultaneously be chiral and real, unless it is completely independent of the superspace coordinates.

Real superfields are an analogue of the gauge fields in regular QFT; thus, any Lagrangian defined over superspace should be *supergauge invariant*. The key to constructing a Lagrangian is to note that the result of integrating a superfield over all superspace coordinates will remain invariant under supersymmetry transformations. The traditional form of a Lagrangian density (which yields an action when integrated over the regular spacetime dimensions) is thus given by an integral over the fermionic coordinates:

$$\mathcal{L}_{\text{SUSY}} = \int d^2\theta d^2\theta^{\dagger} S_i(x, \theta, \theta^{\dagger}). \quad (2.9)$$

In particular, the Lagrangian can be constructed from a combination of so-called D - and F -terms, with the former governing the dynamics of the theory:

$$[V]_D \equiv \int d^2\theta d^2\theta^{\dagger} V(x, \theta, \theta^{\dagger}), \quad (2.10)$$

for some real superfield V , while F -terms correspond to integrals of chiral superfields:

$$[\Phi]_F \equiv \int d^2\theta d^2\theta^{\dagger} \Phi \Big|_{\theta^{\dagger}=0} = \int d^2\theta d^2\theta^{\dagger} \delta^{(2)}(\theta^{\dagger}) \Phi. \quad (2.11)$$

However, note that D -terms can always be rewritten as F -terms; thus, in general, one can fully define the Lagrangian of a supersymmetric theory by specifying a holomorphic function $W(\Phi)$ of chiral superfields, as this itself would correspond to a chiral superfield. Such functions are referred to as *superpotentials*.

If a supersymmetric Lagrangian is invariant under a global $U(1)_R$ transformation $\theta \rightarrow e^{i\alpha}\theta$ and $\theta^{\dagger} \rightarrow e^{-i\alpha}\theta^{\dagger}$, then each superfield can be associated with a so-called R -charge, denoted by r_S and defined by the transformation required of the superfield for the Lagrangian to remain invariant:

$$S(x, \theta, \theta^{\dagger}) \rightarrow e^{ir_S\alpha} S(x, e^{-i\alpha}\theta, e^{i\alpha}\theta^{\dagger}). \quad (2.12)$$

2.1.2 Minimal Supersymmetric Standard Model

The simplest supersymmetric extension to the SM is known as the Minimal Supersymmetric Standard Model (MSSM), which prescribes one superpartner for each SM field. All SM fermions are chiral under electroweak interactions; hence, the left- and right-handed components of each are placed in separate chiral supermultiplets with their scalar superpartners, which are conventionally denoted with the addition of a tilde; e.g. e_L and $\tilde{e}_L$ for left-handed Dirac electrons and the corresponding “left-handed” selectron. The Higgs field also exists in a chiral supermultiplet; however, it turns out that *two* supermultiplets are required, where the fermionic Higgsinos need different values of weak hypercharge in order to avoid gauge anomalies, and to generate the correct Yukawa couplings for up-type quarks ($Y_W = 1/2$) and down-type quarks and leptons ($Y_W = -1/2$). SM gauge bosons (prior to electroweak symmetry breaking) are in gauge supermultiplets with their fermionic superpartner gauginos.

Now, if supersymmetry remained unbroken, the superpartner particles would have the same mass as their SM counterpart; however, in this case it would have been exceedingly trivial to observe them in experiment. Thus, supersymmetry must be a broken symmetry; the Lagrangian of the MSSM can then be summarised as:

$$\mathcal{L}_{\text{MSSM}} = [G]_D + [W_{\text{MSSM}}]_F + [K]_F + \mathcal{L}_{\text{soft}}. \quad (2.13)$$

The first of these terms contains the gauge fields that describe the interactions between the gauge bosons and gauginos,

$$G = \frac{1}{4} F^\alpha F_\alpha + \text{h.c.}, \quad (2.14)$$

where F is used to collectively denote the gauge supermultiplets of $SU(3)_C \times SU(2)_L \times U(1)_{Y_W}$. Then, one has the superpotential term that is responsible for generating the masses of the matter and Higgs fields:

$$W_{\text{MSSM}} = \bar{u} \mathbf{y}_u Q H_u - \bar{d} \mathbf{y}_d Q H_d - \bar{e} \mathbf{y}_e L H_d + \mu H_u H_d, \quad (2.15)$$

where $\{Q, L\}$ and $\{\bar{u}, \bar{d}, \bar{e}\}$ are the chiral superfields containing the left-handed and right-handed squarks and sleptons, respectively; H_u and H_d are the superfields for the two Higgs bosons; and $\mathbf{y}_u$, $\mathbf{y}_d$ and $\mathbf{y}_e$ are dimensionless Yukawa coupling parameters. The Kähler potential term describes the kinetic terms for the matter and Higgs fields:

$$K = \sum_{\text{chiral}} \Phi^\dagger e^{2gV} \Phi, \quad (2.16)$$

where V represents the scalar potentials associated with the gauge symmetries. Finally, $\mathcal{L}_{\text{soft}}$ is the crucial term that enables a soft breaking of supersymmetry in the MSSM. For brevity, I will not write this in full, but it essentially consists of mass terms for all of the superpartner particles, allowing for them to attain a different mass from their SM counterparts. In total, the MSSM contains 105 free parameters, most of which (sparticle masses, phases and mixing angles) arise within this supersymmetry-breaking term.

After electroweak symmetry breaking, the chiral and gauge supermultiplets of the MSSM mix to form various mass eigenstates. The two Higgs doublets give rise to two CP-even neutral bosons h^0 and H^0 (the lighter of which is assumed to correspond to the SM Higgs boson), one CP-odd neutral boson A^0 , and two charged Higgs bosons $H^\pm$. Then, excluding the gluinos, mixing between pure gaugino states form the zino ($\tilde{Z}$) and photino ($\tilde{\gamma}$), which are the superpartners of the Z -boson and photon, respectively. However, given that the neutral higgsinos and gauginos possess the same quantum numbers, one can also interpret their mixing as four neutral mass eigenstates called neutralinos (denoted $\tilde{\chi}_i^0$); similarly, two charginos ($\tilde{\chi}_i^\pm$) can be constructed by mixing the charged Higgsinos with the charged winos.

The form of the superpotential in eq. (2.15) has been chosen such that a discrete $\mathbb{Z}_2$ symmetry known as R -parity is conserved:

$$R_P = (-1)^{3(B-L)+2s}, \quad (2.17)$$

where B (L) is the baryon (lepton) number and s is spin. Within the framework of the MSSM, this was originally motivated to disallow B - and L -violating processes (e.g. proton decay), which are not observed experimentally. All SM particles have $R_P = +1$, while sparticles have $R_P = -1$. A consequence of R -parity conservation, then, is that particles and sparticles do not mix, and so any interaction vertex will involve an even number of sparticles. Thus, the final product in the decay chain of any sparticle has to include (an odd number of) the lightest supersymmetric particle

2.1. Supersymmetry

(LSP), which will be stable (there is nothing else it can decay into). The case of an electrically-neutral LSP constitutes an excellent candidate for dark matter.

2.1.3 Simplified models

The phenomenology of the MSSM (or indeed, any new physics model in general) is rich, and covers a broad spectrum of possible observables in an experimental environment. Within the context of collider experiments such as the Large Hadron Collider (LHC), this translates to many potential final states and signatures that could be observed. One can imagine, then, that when optimising a search to a specific signature, including all of the irrelevant degrees of freedom in the full theory (added to the fact that there are *many* different theories readily available) makes things unnecessarily complicated.

For this reason, the analysis of LHC data is typically conducted using *simplified models*, which are designed to produce the same topology as the full theory but with a greatly reduced number of free parameters, corresponding to only a few particle masses, production cross-sections, and decay modes of interest to the specific search. Upon any affirmative hint of new physics, the results derived in terms of simplified models can be transformed into the parameter space of a more general model; thus, while simplified models are (by definition) not exactly model-independent, they are still able to serve as a useful benchmark for reinterpretation within the wider theory community. A summary of the various simplified frameworks used by the LHC experiments can be found in refs. [14, 15].

2.1.4 Non-minimal flavour violation

In the MSSM, the mass splitting between SM particles and their corresponding sparticles is determined by the parameters that enter the soft supersymmetry-breaking term in the Lagrangian. While it might seem that we have arbitrary control over the values of these parameters, it turns out that viable models of SUSY (i.e. those that are able to reconcile the hierarchy problem) require that the particle–sparticle mass difference is around an order of magnitude, at most. Given that the W and Z -bosons have masses of $\mathcal{O}(100 \text{ GeV})$, one then expects the mass of any potential sparticle to lie somewhere in the TeV range.

Fortunately, the production of TeV-scale particles is within the kinematic reach of the LHC. Unfortunately, after over a decade of data-taking, there has yet to be any definite sign of supersymmetric particle production. An easy way to sidestep this issue is to say that the experimental limits aren't *really* reaching the TeV range yet. Indeed, the various limits currently placed on the sparticle masses are valid only within the context of the simplified models adopted, which in particular often assume a 100% branching fraction to the final state of interest. If the possible decays are instead split between multiple channels, the overall expected event yield in the data will decrease, and so too the mass limits.

The MSSM includes only minimal flavour violating (MFV) effects, which arise when diagonalising the fermion sector of the theory. However, what if one introduces non-minimal flavour violation (NMFV) to the squark sector? Mixing involving the first generation squarks is highly constrained by kaon experiments; however, that between the 2nd and 3rd generation squarks is still largely allowed [16]. This is the scenario that is relevant for the analysis that will be presented in Part III.

To implement this feature into a simplified model, one can introduce a mixing between right-handed top and charm squarks such that they form physical eigenstates $\tilde{u}_1$ and $\tilde{u}_2$. The strength




Figure 2.1. Visualisation of how the inverse strength of the three SM couplings run with energy. They *almost* (but unfortunately don't quite) meet at an energy of approximately 10^{14} GeV.

of the mixing is controlled by a mixing angle:

$$\begin{pmatrix} \tilde{u}_1 \\ \tilde{u}_2 \end{pmatrix} = \begin{pmatrix} \cos \theta_{tc} & \sin \theta_{tc} \\ -\sin \theta_{tc} & \cos \theta_{tc} \end{pmatrix} \begin{pmatrix} \tilde{c}_R \\ \tilde{t}_R \end{pmatrix}. \quad (2.18)$$

In addition to a bino-like neutralino $\tilde{\chi}_1^0$, the model then contains four free parameters of interest: the masses of the two squarks, $m(\tilde{u}_1)$ and $m(\tilde{u}_2)$, the mixing angle θ_{tc} , and the mass of the neutralino, $m(\tilde{\chi}_1^0)$.

2.2 Leptoquarks

In its current form, the SM describes three disjoint forces: the strong, the weak, and the electromagnetic. The individual couplings of these forces each vary with energy, and at an energy scale of $\sim 10^{14}$ GeV, their values almost coincide (this is loosely demonstrated in Figure 2.1). Now, if one could formulate a theory in which the coupling strengths *did* meet (in fact, the MSSM is one such theory), this opens up the appealing suggestion that the three forces appearing in the SM might actually be emergent from a single, more fundamental force above some energy scale. This is the basis of a *grand unified theory* (GUT).

Many different theories have been proposed in the quest for grand unification. However, a feature that is common to all of them is the existence of so-called *leptoquarks*—particles that carry both lepton and baryon numbers, which arise as a result of attempting to couple together the lepton and quark sectors of the SM. These have recently gained renewed interest due to hints of lepton flavour universality violation in semi-leptonic decays of B -hadrons [17–22].

The particular properties of LQs are model-dependent, but they are commonly classified based on their interaction mechanism and their electric charge. *Vector* LQs refer to those mediated by the gauge boson of the unified gauge group, while mediation by the Higgs sector give rise to *scalar* LQs. Orthogonal to this classification, LQs with positive electric charge are referred to as *up-type* LQs, with $Q = +5/3$ or $Q = +2/3$; or *down-type* if $Q = -1/3$ or $Q = -4/3$. The charge of the LQ determines what combination of quark and anti-lepton it can decay into; for example, LQs with $Q = +2/3$ can decay either into an up-type quark (u, c, t) and an anti-neutrino, or into a down-type quark (d, s, b) and a charged anti-lepton (e^+, μ^+, τ^+), based on the conservation of

2.2. Leptoquarks

electric charge. The following sections will describe two particular models that are relevant to the analysis presented in Part III.

2.2.1 Scalar leptoquark model

The simplified model for scalar LQs is described in refs. [23, 24]. In particular, the model for up-type LQs with $Q = +2/3$ has a Lagrangian that contains the following interaction term:

$$\mathcal{L}_{\text{int}} = \lambda_\ell (\sqrt{\eta_L} \bar{d}_R \ell_L^- + \sqrt{\eta_R} \bar{d}_L \ell_R^-) \ell_q + \lambda_\nu \bar{u}_R \nu_L^\ell \ell_q + \text{h.c.}, \quad (2.19)$$

which allows for the decays $\text{LQ} \rightarrow d + \ell^+$ or $\text{LQ} \rightarrow u + \bar{\nu}$, where d and u represent any down- or up-type quark, respectively. The parameter η_L controls the ratio of left- to right-handed charged leptons produced in the decay, with η_L and η_R summing to unity; however, given that collider experiments are generally agnostic to the lepton polarisation, one can choose $\eta_L = 1$ without any loss of generality. The couplings λ_ℓ and λ_ν are related via

$$\lambda_\ell^2 = \beta \lambda^2, \quad \lambda_\nu^2 = (1 - \beta) \lambda^2, \quad (2.20)$$

such that $0 < \beta < 1$ is a parameter whose value effectively determines the branching ratio of the possible decay channels.

2.2.2 Vector U_1 models

The simplified vector LQ model presented here is based on the one described in ref. [25], but where the degrees of freedom related to the UV completion of the model have been neglected. The relevant terms in the Lagrangian are:

$$\begin{aligned} \mathcal{L}_{\text{NP}} = & -\frac{1}{2} U_{1\mu\nu}^{a\dagger} U_1^{\mu\nu,a} + m_{\text{LQ}}^2 U_{1\mu}^{a\dagger} U_1^{\mu,a} - i g_s \frac{1}{2} (1 - \kappa) U_{1\mu}^{a\dagger} \lambda^b U_{1\nu}^a G^{\mu\nu,b} \\ & - i g_Y \frac{2}{3} (1 - \tilde{\kappa}) U_{1\mu}^{a\dagger} U_{1\nu}^a B^{\mu\nu} + \frac{g_U}{\sqrt{2}} \left[U_1^{\mu,a} \left(\beta_L^{ij} \bar{Q}_L^{i,a} \gamma_\mu L_L^i + \beta_R^{ij} \bar{d}_R^{i,a} \gamma_\mu \ell_R^j \right) + \text{h.c.} \right]. \end{aligned} \quad (2.21)$$

This includes the possible non-minimal interactions with the SM gauge fields, controlled by the κ and $\tilde{\kappa}$ couplings; however, given that contributions to the production of LQs from electroweak (the $\tilde{\kappa}$ term) and t -channel lepton exchange (the g_U terms) are small, they are thus neglected in favour of the more dominant strong production mechanisms. Two models are then defined based on the remaining terms: the Yang–Mills and minimal coupling scenarios, corresponding to choosing $\kappa = 0$ or $\kappa = 1$, respectively.

The possible decay channels of the LQs are controlled by the β_λ^{ij} parameters, where λ dictates the polarisation of the decay products while the indices i and j respectively indicate the generation of quark or lepton that is produced.

From Theory to Experiment



As discussed in Chapter 1, gauge-invariant quantum field theories provide a rigorous and consistent mathematical framework for modelling the fundamental phenomena of our universe. However, the presentation thus far has remained quite abstract: a theory can be specified by writing down a Lagrangian density that satisfies a set of mathematical requirements, but how exactly can one use this to infer practical consequences or predictions of the real world?

This is the question that motivates this chapter, the goals of which are twofold. Firstly, I will provide a bridge between theory and experiment by introducing an important tool of the trade: perturbation theory, and how this can be applied to derive experimentally measurable quantities. The second part of this chapter will discuss particle collider experiments, focusing in particular on CERN's Large Hadron Collider (LHC), which is the most ambitious scientific experiment realised to date.

3.1 Perturbation theory

The quantisation of free quantum fields can be performed in a mathematically rigorous manner. Unfortunately, once we begin to consider the more interesting systems that include dynamical interactions between the fields, things become somewhat less precise (much to the dismay of mathematicians). Still, this is unavoidable if we want to realise any theory that is to have any physical relevance; indeed, we have already seen how the construction of the Standard Model is motivated by the interactions between elementary particles.

The primary tool that facilitates calculations involving interacting systems is known as *perturbation theory*. The idea is that the Hamiltonian describing the behaviour of free fields, H_0 , can be separated from that describing interactions, H_I , such that the total Hamiltonian of the system can be written as:

$$H = H_0 + H_I . \quad (3.1)$$

In the case of a theory whose Lagrangian is known instead, one can find the Hamiltonian via $L = T - V \implies H = T + V$, where T represents the kinetic (i.e. free-field) terms and V are the potential (i.e. interaction) terms.

One can thus interpret H_I as a *perturbation* of the free field Hamiltonian, and the challenge comes down to calculating the time evolution operators for the system, $U(t) = e^{-itH/\hbar}$. This is typically done in the *interaction picture*, which lies between the Schrödinger and Heisenberg pictures introduced in section 1.1: in this picture, one allows the observables to evolve in time

3.2. Feynman diagrams

following the unperturbed H_0 (the time evolution of which is assumed to be known),

$$A(t) = e^{itH_0/\hbar} A e^{-itH_0/\hbar} \equiv U_0(-t) A U_0(t), \quad (3.2)$$

while states evolve according to the perturbation,

$$\psi(t) = U_0(-t) U(t) \psi \equiv V(t) \psi. \quad (3.3)$$

In this way, the description of the system at some time τ is then as one expects:

$$\langle \psi_2(\tau) | A(\tau) | \psi_1(\tau) \rangle = \langle \psi_2(t=0) | U(-\tau) A(t=0) U(\tau) | \psi_1(t=0) \rangle. \quad (3.4)$$

Since $U_0(t)$ has been assumed to be known, the exercise of calculating $U(t)$ boils down to finding the family of unitary operators $V(t)$. The way this can be done is by expanding it as a *Dyson series*, $V(t) \sim I + \sum_1^\infty V_n(t)$, where

$$V_n(t) = \frac{1}{i^n} \int_0^t \int_0^{\tau_n} \cdots \int_0^{\tau_2} H_I(\tau_n) H_I(\tau_{n-1}) \cdots H_I(\tau_1) d\tau_1 \cdots d\tau_{n-1} d\tau_n. \quad (3.5)$$

One can use the notion of a *time-ordered product* to more succinctly write this as

$$V(t) = \mathcal{T} \exp \left[-i \int_0^t H_I(\tau) d\tau \right], \quad (3.6)$$

where $\mathcal{T}[A_1(t_1) \cdots A_n(t_n)] = A_{i_1}(t_{i_1}) \cdots A_{i_n}(t_{i_n})$ for $t_{i_1} > t_{i_2} > \cdots > t_{i_n}$. This notation lends itself to a more heuristic interpretation of the Dyson series: any one of its terms then describes a freely-propagating particle that experiences an interaction at some point in time; continues propagating; interacts again at some later time; and so on, for as many times as H_I appears in the integral. The full interacting propagator—corresponding to the full Dyson series—accounts for all possible interactions at all possible times.

Now, all of this might still seem a little abstract, but the key lies in the following ansatz, which forms the basis of perturbation theory: one assumes that some physical quantity C can be written as $\sum_0^\infty c_n$, derived from $\sum V_n$; then, if the first few terms c_n decrease rapidly as n increases¹, a *partial sum of the series provides an approximation for C* . This result—which turns out to just *work* in practice—allows us to calculate approximate values for physical observables that can then be tested experimentally.

3.2 Feynman diagrams

The information contained within a Dyson series can be condensed into the form of a useful visual representation known as a *Feynman diagram*. As one might recall from the brief discussion in section 1.2, quantisation of a theory involves rewriting its fields in terms of creation and annihilation operators. The Dyson series, then, will consist of integrals involving these, along with field operators, according to the fields that interact within the theory. Each non-zero contribution (there are some constraints that render certain combinations of terms zero) will correspond to a particular visual element in the corresponding Feynman diagram.

¹Note that this requirement is not quite the same as simply needing the series to converge—it is possible to define a series where successive terms vary substantially, yet the sum of which still converges overall.

appearance	particle
	fermion
	scalar
	photon, $W^\pm$, Z
	gluon

Figure 3.1. Common line style convention for representing different particle types in Feynman diagrams.

Every instance of a creation or annihilation operator, along with each unique spacetime point at which a field operator acts, corresponds to a *vertex* in the diagram. *External* vertices are those that correspond to creation or annihilation operators, while *internal* vertices represent the spacetime points. Vertices are connected by *lines* that represent the particles corresponding to the operators involved; the convention is to use different line styles for different particle species (summarised in Figure 3.1), with directionality indicated via arrows to differentiate between particles and antiparticles if necessary.

Due to the sharp correspondence between the features of a Feynman diagram and the term in the Dyson series that it represents, it is possible to derive the mathematical form of the latter given the former, using a set of *Feynman rules* for the theory in question. Indeed, in practice, one often foregoes a calculation from first principles, instead leveraging the ansatz of perturbation theory to choose a set of valid diagrams for the process of interest (usually starting from the simplest possible, which correspond to those containing the fewest number of vertices) and using these to obtain an approximation of the full matrix element.

3.3 Cross-sections and decay rates

Now that we have seen the basic tools that allow us to model the time evolution of an interacting quantum system, let us see how they can be applied to obtain experimentally accessible quantities. In particular, consider particle collider experiments, where one generates a beam of some particle species and collides it with a target (be it another beam or a fixed target). In this scenario, there is an incoming $|\text{in}\rangle$ and outgoing $|\text{out}\rangle$ state of free particles, which interact in some way such that $|\text{in}\rangle \rightarrow |\text{out}\rangle$. This transformation is described by the *scattering operator*,

$$S = \lim_{t_0 \rightarrow -\infty, t_1 \rightarrow +\infty} U_0(-t_1) U(t_1 - t_0) U_0(t_0), \quad (3.7)$$

where the limits indicate that the particles are free to $t \rightarrow \pm\infty$ before and after the interaction. The matrix element $\langle \text{out} | S | \text{in} \rangle$ is known as the *scattering (S) matrix*, which can be expanded as a Dyson series, $\sum_0^N \langle \text{out} | S_n | \text{in} \rangle$, following the perturbative approach; for some finite choice of N , this yields an approximation for the full matrix element.

In general, for a collection of incoming and outgoing particles with four-momenta $p_1, \dots, p_m$ and $q_1, \dots, q_n$ respectively, the S-matrix can be written as

$$\langle q_1, \dots, q_n | S - I | p_1, \dots, p_m \rangle = i(2\pi)^4 \delta(\Sigma q_k - \Sigma p_j) M(q_1, \dots, q_n; p_1, \dots, p_m), \quad (3.8)$$

3.4. Event generation

where the δ -function serves to enforce four-momentum conservation through the interaction. The matrix element can be interpreted as a *transition amplitude*, and much like in non-relativistic quantum mechanics, the absolute square of it then corresponds to a *transition probability* for an interaction to occur. There is some mathematical freedom in how to square eq. (3.8) (it is not as trivial an exercise as one might initially think, because the square of a δ -function is ill-defined); fortunately, it turns out that the end result remains agnostic to the particular method of calculation employed.

The end result of this is that one can compute the rate at which interactions are expected to occur in an experimental setting. This is called a *cross-section*, denoted σ , and corresponds to the number of scattering events per unit time $[T^{-1}]$, per unit volume $[L^{-3}]$, per unit density of particles $[L^3]$, and per unit flux of the particle beam $[TL^2]$. For the case of two particles colliding, the differential cross-section can generally be written as:

$$d\sigma = (2\pi)^4 \delta(\Sigma q_k - \Sigma p_j) |\mathbf{v}|^{-1} |M(q_1, \dots, q_n; p_1, \dots, p_m)|^2 \frac{d^3 \mathbf{q}_1 \cdots d^3 \mathbf{q}_n}{(2\pi)^{3n}}, \quad (3.9)$$

where $\mathbf{v}$ represents the beam velocity. Typically, one also takes the sum over all spin states of the particles, since experiments are usually not conducted for any particular spin; and, in the case of hadron collisions, a sum over colour degrees of freedom and the introduction of *parton distribution functions* (PDFs) to account for the fractional momentum carried by each parton in the parent hadron. Cross-sections have dimensions of $[L^2]$, and are commonly expressed in units of *barns* (b), with $1 \text{ b} = 10^{-28} \text{ m}^2$.

An interesting case of “scattering” is when only *one* particle is involved in the initial ensemble. Imposing this assumption upon the methodology outlined above yields a transition rate for the *decay* of that particle.

3.4 Event generation

How can we use the concepts introduced thus far to obtain predictions for what we might expect to see in an experiment? Of course, an accurate simulation of the particular experiment’s setup (e.g. its detector geometry and electronics) forms an important part of this, but in this section let us focus on the theoretical aspects. Consider, for example, two hadrons colliding in a $2 \rightarrow 2$ process. Perturbative calculations can provide an approximation for the matrix elements for the interaction, but there are complications:

- hadrons are composite objects, so the differential cross-section will be given by the product of the hard-scattering matrix elements with the structure functions for the hadron
- incoming and outgoing partons can radiate, producing additional particles in the final state
- isolated partons cannot be observed; they fragment into hadrons, which, if unstable, will subsequently decay until stable products are reached

What results from all this is that the original $2 \rightarrow 2$ process that we were aiming to study is now a $2 \rightarrow n$ interaction in general, and obtaining a prediction for it requires the calculation also of all the other matrix elements—an exercise that grows very complicated, very quickly.

Fortunately (and crucially), it turns out that all of these different facets of a collision event can largely be factorised according to the momentum transfer scale at which they happen. At the highest scale lies the hard-scatter interaction, for which the matrix elements are perturbatively

calculable. As the scale evolves downwards, heavy resonances decay, initial- and final-state parton showers occur, and secondary low-energy interactions between other partons in the colliding hadrons are generated. Perturbation theory can be applied to predict some of these phenomena, but this will eventually break down at low scales, where one has to rely on tunable models instead. Finally, at the lowest scales, non-perturbative interactions involving the final state partons produce the stable hadrons in the event, which are simulated using hadronisation models inspired by QCD. The idea, then, is that one need only calculate the lowest-order² matrix element for the process of interest, and augment this basic scattering with the other separately-simulated effects to obtain a realistic prediction for the event.

All of the steps described above are achieved in practice using Monte Carlo (MC) integration. The essence of the technique lies in approximating an integral using the average of its integrand:

$$I = \int_{x_1}^{x_2} dx f(x) = (x_2 - x_1) \langle f(x) \rangle \approx (x_2 - x_1) \frac{1}{N} \sum_{i=1}^N f(x_i), \quad (3.10)$$

where x_i are values randomly drawn (N times) from the interval (x_1, x_2) . The uncertainty in this approximation scales as $1/\sqrt{N}$, irrespective of the number of dimensions being integrated over; this makes MC sampling the ideal (and perhaps only viable) choice for cross-section calculations that involve integrals over a very large phase space of initial and final state particles. Parton showers (which are modelled as Markov processes) and hadronisation models are both also well-suited to being simulated using MC techniques.

Of course, several assumptions need to be made in this formulation of event generation, and these will give rise to systematic uncertainties in the method. Firstly, the PDFs that enter the cross-section calculation cannot be derived completely from first principles; they rely on experimental measurements (e.g. deep inelastic scattering), and are thus subject to experimental uncertainties. Furthermore, their theoretical dependence on the momentum transfer scale depends on a factorisation scale, μ_F , which loosely represents the threshold that separates the perturbative and non-perturbative regimes. Parton-level cross-sections additionally depend on the choice of a renormalisation scale, μ_R , which is introduced when truncating the calculation at some order in perturbation theory. The factorisation and renormalisation scales are unphysical parameters with no “correct” value; the convention is to set $\mu_F = \mu_R = Q^2$ (where Q is an energy scale representative of the process under study), but different choices for these parameters can affect the outcome of event generation, since they are related to the starting scale for parton showers. These considerations give rise to the theoretical uncertainties that are later referenced in Part III.

Further reading on the intricacies of event generation can be found in refs. [26, 27].

3.5 Large Hadron Collider

Located at the Franco–Swiss border near Geneva, CERN’s Large Hadron Collider (LHC) [28] is an experiment that accelerates, collimates, and smashes together beams of protons (and occasionally heavier ions such as lead and oxygen) at extremely high energies. Due to the feature of asymptotic freedom in QCD, the bonds that hold together the constituent quarks and gluons in a proton (collectively known as *partons*) weaken under these collision conditions, facilitating direct parton–parton interactions against which our theories of nature can be tested.

²Of course, explicit calculations of higher-order contributions will yield more accurate results; state-of-the-art matrix element calculations now often go to next-to-leading order (NLO) or even next-to-next-to-leading order (NNLO).

3.5. Large Hadron Collider




Figure 3.2. Overview of the LHC and its experiments. (Image: CERN)

A schematic of the LHC is shown in Figure 3.2. A series of four accelerators are first used to gradually ramp up the energy of the proton beams in preparation for collisions. The chain begins with the 86-m-long Linac4 linear accelerator³ that takes negative hydrogen ions and boosts them to an energy of 160 MeV. These are stripped of their two electrons before being passed into the circular Proton Synchrotron Booster, which further increases their energy to 2 GeV. The Proton Synchrotron (PS) then boosts the beams to 26 GeV, followed by the Super Proton Synchrotron (SPS) to 450 GeV. Finally, the beams are injected into the main LHC accelerator, which is housed in a 27-km-long circular tunnel located approximately 100 m underground. Here, the proton beams travel in opposite directions in two separate beam pipes, guided by a collection of superconducting dipole magnets and focused using quadrupole magnets. Each beam reaches a final energy of 6.5 TeV (Run 2) or 6.8 TeV (Run 3) per beam, corresponding to a collision centre-of-mass energy of $\sqrt{s} = 13$ TeV or 13.6 TeV, respectively. Collisions occur at four main interaction points (IPs) where the detectors are situated: the ATLAS [29] and CMS [30] general purpose detectors, the ALICE [31] heavy ion experiment, and the LHCb [32] flavour physics experiment.

3.5.1 Luminosity and event rate

The rate at which interactions occur is given by the product of the cross-section and the *instantaneous luminosity*, which can loosely be thought of as a density of protons available to participate in interactions:

$$\frac{dN}{dt} = \sigma \mathcal{L}(t). \quad (3.11)$$

³The Linac4 began operating in 2020, and plays an important role in the plan to increase the luminosity of the LHC. It supersedes the Linac2, which previously accelerated protons to 50 MeV during the first two runs of the LHC.

Since the event count is dimensionless and σ has units of $[L^2]$, it follows that the luminosity has dimensions of $[L^{-2}T^{-1}]$.

Cross-sections of interesting hard-scattering processes are typically very low, so to achieve a reasonable event rate, the challenge becomes that of effectively increasing the instantaneous luminosity of the beams. In the nominal operating conditions of the LHC, this is achieved by separating each proton beam into 2 808 discrete *bunches*, with each bunch containing approximately 10^{11} protons. Bunches are spaced 25 ns (approximately 7.5 m) apart and revolve around the LHC ring at a frequency of 11.245 kHz. Upon approaching an interaction point, they are *squeezed* to a width of approximately $20 \mu\text{m}$ to further increase the probability for an interaction to occur. The value of the instantaneous luminosity is derived from these and a few other beam parameters:

$$\mathcal{L} = \frac{N^2 n_b f_{\text{rev}} \gamma}{4\pi \epsilon_n \beta^*} F, \quad (3.12)$$

where N denotes the number of particles in each bunch; n_b is the number of bunches; f_{rev} is the revolution frequency; γ is the relativistic Lorentz factor; ϵ_n is the normalised beam emittance; and β^* is the beta function at the interaction point, related to the transverse width of the beam. The geometric factor F accounts for the possibility of an incomplete overlap between the beams at the interaction point, although in practice this often has a value very close to unity:

$$F = \left(1 + \frac{\theta_c \sigma_z}{2\sigma^*}\right)^{-\frac{1}{2}}, \quad (3.13)$$

where θ_c is the crossing angle at the interaction point, σ_z is the rms bunch length, and σ^* is the transverse rms beam size.

At the LHC, each instance of a collision between two bunches of protons is referred to as a *bunch crossing*, and contains one hard-scatter interaction on average. This corresponds to a rate of up to 40 million interesting physics events occurring every second (although due to a few technical reasons, the bunch spacing is not always at 25 ns in practice, leading to an average rate of around 30 MHz instead). Another widely used metric is the *integrated luminosity*, which provides a measure of the total amount of data generated by the experiment over some operational duration:

$$\mathcal{L}_{\text{int}} = \int_{t_{\text{start}}}^{t_{\text{end}}} dt \mathcal{L}(t). \quad (3.14)$$

This is often expressed in units of inverse femtobarns (fb^{-1}).

3.5.2 Underlying event and pile-up

In addition to the hard-scatter interaction, other processes can occur in a bunch crossing. These make up the *underlying event* (UE), and arise from a variety of effects, including: multi-parton interactions (MPI) between other partons in the colliding protons; initial or final state radiation (ISR, FSR); colour reconnection (non-perturbative long-range gluon exchange); soft emissions of particles during hadronisation; beam remnants; and more. The examples listed here constitute real physical effects associated with collisions between protons, and in order to account for them in simulation, one typically incorporates an *underlying event tune* that is derived using minimum bias events in data (i.e. events collected independently of the hard-scatter).

Another prominent class of contributions to events at the LHC stems from the fact that collisions occur between bunches of many protons. The result of interactions between the other

3.5. Large Hadron Collider

protons in a bunch crossing—other than the pair involved in the hard-scatter—is known as *pile-up* (PU). The rate of such interactions typically dwarfs that of the hard-scatter; for example, in Run 3 of the LHC, every hard-scatter process is accompanied by an average of approximately 60 additional soft collisions. This results in many more soft particles being generated, contributing to the overall event noise (particularly in the forward regions). These can be electrically charged or neutral, so in the context of general purpose detectors such as ATLAS, this leads to many low-energy deposits within the calorimeter, and in the case of charged PU, also hits within the tracking volume. During the reconstruction of an event, PU particles can be reconstructed as independent objects (*QCD pile-up*); or, the detector signals from multiple PU particles can overlap and be collected into a single object (*stochastic pile-up*).

Contributions from PU interactions between protons in the same bunch crossing are referred to as *in-time* PU. However, the tight 25 ns spacing between bunches can also give rise to what is known as *out-of-time* PU. This latter category depends on the sensitivity of the detector apparatus: namely, whether the signals generated within the detector material persist for a duration that is longer than the time until the subsequent bunch crossing. The number of in-time PU contributions in an event can be estimated using the number of primary vertices in the event, N_{PV} , while the average number of interactions per bunch crossing, denoted $\langle\mu\rangle$, provides an indicator for the rate of out-of-time PU interactions.

Since PU interactions occur independently of the hard-scatter, it is often desirable to identify and consequently exclude PU from the event description. In simulation, their contributions are included by overlaying independently-generated events onto the simulated hard-scatter event. Chapter 4 will include a discussion of PU suppression techniques used in the ATLAS experiment.

3.5.3 History and timeline

The LHC was first proposed in the 1980s as an experiment that could make use of the tunnel that had already been excavated for the Large Electron–Positron (LEP) collider experiment. Construction of the LHC and its four main experiments was approved and carried out over the next two decades, and 10:28 am on 10 September 2008 saw the first successful circulation of a proton beam around the LHC. Unfortunately, an electrical failure just nine days after this resulted in serious damage to the magnets and a large helium leak in the tunnel, inducing a delay to the experiment’s schedule of around a year while repairs and safety measures were implemented for the affected components.

The year 2010 marked the beginning of data-taking for physics—the so-called Run 1 of the LHC, which operated at a collision centre-of-mass energy of $\sqrt{s} = 7$ TeV for the first two years before ramping up to $\sqrt{s} = 8$ TeV in 2012. Collisions reached a peak instantaneous luminosity of $4.67 \times 10^{32} \text{ cm}^{-2} \text{ s}^{-1}$ (beating the record previously held by the Tevatron), with a pile-up rate of $\langle\mu\rangle \approx 20$. Over the three-year period, approximately 5 fb^{-1} of 7 TeV and 21 fb^{-1} of 8 TeV collision data was collected by the experiments. The hallmark of Run 1 was the discovery of a particle consistent with the Higgs boson by both the ATLAS [33] and CMS [34] experiments.

After this first successful run, the LHC entered a two-year period (2013–2015) known as Long Shutdown 1 (LS1). During this time, various maintenance tasks were carried out, including checks of more than 10 000 interconnections between magnets and the replacement of over 100 km of cabling in the SPS. The plan was to increase the collision energy for the next physics run, and as such, a stronger magnetic field was required to continue guiding the proton beams. To facilitate this, the cryogenics of the superconducting magnets were improved with the addition of liquid helium allowing for an operating temperature of 1.9 K to be maintained. Due to




Figure 3.3. (a) Cumulative luminosity versus day delivered to ATLAS during stable beams for high energy pp collisions. (b) Integrated luminosity as a function of the mean number of interactions per crossing, $\langle\mu\rangle$, for the 2015–2018 pp collision data at $\sqrt{s} = 13$ TeV. All data recorded by ATLAS during stable beams are shown, including the integrated luminosity and the mean μ value for each year. The mean number of interactions per bunch crossing corresponds to the calculated mean of the Poisson distribution of the number of inelastic interactions per bunch crossing. Images from ref. [35].

time constraints for training these magnets (related to the phenomenon of quenching, in which superconductivity can be lost), the decision was made to begin operations at a collision energy of $\sqrt{s} = 13$ TeV so that physics data-taking could commence as scheduled in 2015.

Run 2 surpassed the $10^{34} \text{ cm}^{-2}\text{s}^{-1}$ design instantaneous luminosity of the LHC, managing to achieve a peak value of $2.1 \times 10^{34} \text{ cm}^{-2}\text{s}^{-1}$ in 2016. This was made possible not only by the higher 6.5 TeV beam energies, but also by a decreased bunch spacing, improved beam collimation tolerances (using the Batch Compression Merging and Splitting (BCMS) production scheme in the PS), the implementation of a crossing angle anti-levelling scheme (where the crossing angle at the IP is reduced as the bunch intensity decreased), and a β^* anti-levelling strategy (corresponding to an additional squeeze towards the end of each fill). In particular, the anti-levelling schemes allowed for an increased overall fraction of time that the LHC could operate near the peak luminosity. In total, the LHC delivered approximately 140 fb^{-1} of data during Run 2 (Figure 3.3a) at an average of $\langle\mu\rangle \approx 30$ interactions per bunch crossing (Figure 3.3b). This dataset is the one that will be relevant for the analysis presented in Part III of this thesis.

The cycle then repeated with the second Long Shutdown (LS2) over 2019–2022, which involved a consolidation campaign for the bypass diodes of the superconducting dipole magnets, the addition of the Linac4 to the injector chain (replacing the Linac2), and upgrades to the cryogenic systems. All of these served to facilitate yet another increase in the beam energy towards the LHC design value of $\sqrt{s} = 14$ TeV, with Run 3 of the LHC (which commenced in 2022 and is ongoing at the time of writing) achieving just shy of this target at an energy per beam of 6.8 TeV ($\sqrt{s} = 13.6$ TeV). The peak instantaneous luminosity remained as it was in Run 2 at around $2 \times 10^{34} \text{ cm}^{-2}\text{s}^{-1}$, limited by heat generation in the inner triplet magnet system that focuses the beams just before they collide. Instead, the aim was to run near the peak luminosity for even longer durations of up to 10 hours per LHC fill. Under these conditions, each bunch crossing features an average of around $\langle\mu\rangle \approx 60$ inelastic pp interactions.

So far, Run 3 has yielded approximately 196 fb^{-1} of data after a very successful 2024, and the plan is to continue the run until mid-2026, with a total integrated luminosity of over 250 fb^{-1}

3.5. Large Hadron Collider




Figure 3.4. Long term schedule for the LHC as of November 2024. (Image: CERN)

being projected for the full run. After this, the experiment will transition into the era of the High-Luminosity LHC (HL-LHC). The planned timeline for this, extending into the 2040s, is presented in Figure 3.4. The goal of the HL-LHC is to achieve the design collision energy of $\sqrt{s} = 14$ TeV and a greatly increased instantaneous luminosity, levelled at $5 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$ [36]. To accommodate these drastically more demanding conditions, new focusing magnets, along with major hardware upgrades for all of the detector experiments, will have to be installed during Long Shutdown 3 (LS3) beginning in July 2026. The final hardware commissioning will take place in January 2030, with the first beams to commence in June of the same year.

The HL-LHC will produce an annual dataset of around 250 fb^{-1} , accumulating to a total of 3000 fb^{-1} over its 10-year duration. One of the major challenges in the HL-LHC environment is a greatly increased average of $\langle \mu \rangle = 200$ interactions in each bunch crossing; an interesting consequence of this forms the motivation for the studies that will be presented in Chapter 9.

The ATLAS Experiment



ATLAS (“A Toroidal LHC ApparatuS”) is one of four major experiments at the LHC [29]. While its inception was motivated largely by the prospect of discovering the (at the time) theoretically-predicted Higgs boson, the design of its detector has been kept generic, allowing the experiment to maintain a diverse physics programme ranging from precision measurements of the Standard Model to searches for new and exotic phenomena.

As of 2023, the ATLAS collaboration has grown to over 5000 members affiliated with institutions located in more than 100 countries, with approximately 20% of its qualified authors being comprised of women [37]. The physics effort of the collaboration is broadly split into the two branches of Combined Performance (CP) and Physics Analysis (PA), with additional ancillary groups responsible for event simulation and providing recommendations for statistical techniques. The six CP groups (Jet/ E_T^{miss} , E/Gamma, Muon, Tau, Flavour Tagging, and Tracking) handle the reconstruction and identification of various objects in the detector, supporting common definitions and working points that are used by the wider collaboration. The slew of PA groups produce the bulk of the physics output of the experiment: each group targets a specific final state (e.g. the production of top-quarks, Higgs bosons, or B mesons and other light states), physics model assumption (e.g. the Standard Model, supersymmetry, or other exotic and unconventional signatures), or other focussed goal (e.g. modelling of various physics processes in Monte Carlo simulations, or upgrades in preparation for the high-luminosity LHC).

This chapter will present an overview of various facets of the experiment. I will begin by summarising the detector hardware (as relevant during Run 2 of the LHC) in section 4.1. The techniques used for reconstructing and interpreting collision events will be described in sections 4.3 and 4.4. A review of the upgrades introduced for Run 3, or planned for the high-luminosity era, is presented in sections 4.5 and 4.6, respectively. Finally, given its remarkable surge in ubiquity in recent years, section 4.7 will be dedicated to a discussion of machine learning techniques and their planned and potential applications within the experiment.

4.1 ATLAS detector

The ATLAS detector is designed to measure (or infer) the entire particle spectrum of the Standard Model (SM). Sitting at 46 m in length and 25 m in diameter, it manages to be the largest detector of any LHC experiment by a sizeable margin, whilst weighing in at a slim 7000 t (less than that of ALICE, and only half that of CMS). The detector consists of a precise vertexing and tracking volume, a high-resolution calorimeter system for energy measurements, and a dedicated muon

4.1. ATLAS detector




Figure 4.1. Cut-away visualisation of the ATLAS detector and its components (with little people to scale).

system. These components are arranged in cylindrical fashion around the beampipe, covering a near- 4π solid angle to maximise the acceptance of particles produced in collision events (Figure 4.1). The tracking system sits innermost, followed by the calorimeter layers (necessarily in this order, given the destructive *modus operandi* of the latter). Finally, the outermost layer corresponds to a detector subsystem dedicated to the measurement of muons, which have a low rate of interaction and typically traverse the extent of the calorimeters unimpeded. Each of these subsystems will be described further in the following sections.

ATLAS adopts a Cartesian right-handed coordinate system centred around the collision point. The positive x -axis is defined as pointing from the collision point towards the centre of the LHC ring; the positive y -axis points upwards; and the positive z -axis points anti-clockwise along the beam pipe. The azimuthal angle (ϕ) is measured around the beam axis, while the polar angle (θ) is measured with respect to the z -axis, although it is usually more common to work with the pseudorapidity, $\eta = -\ln[\tan(\theta/2)]$, or rapidity, $y = 0.5 \ln[(E + p_z)/(E - p_z)]$ where E is the energy and p_z is the z -component of momentum, due to their useful property of differences being invariant under Lorentz boosts in the z -direction (this is important since the z -component of the centre-of-mass frame is unknown in a pp collision).

4.1.1 Inner detector

The Inner Detector (ID) [38, 39] utilises a 2 T axial magnetic field, applied throughout the entirety of its volume, to provide measurements of charged particle momenta via the curvature of their trajectories. The system consists of two silicon-based and one straw-based detector, with barrel segments arranged as a series of cylinders concentric with the beam axis, and endcap detectors as circular discs centred on and perpendicular to the beampipe. The detector has an overall




Figure 4.2. Layout of the ATLAS Inner Detector.

diameter of 2.5 m and length of 5.3 m, and covers the pseudorapidity range up to $|\eta| = 2.5$. A summary of the ID geometry and layout is depicted in Figure 4.2.

A high-granularity silicon pixel detector is used in the vertex region where the particle density is highest in a collision event. It consists of three cylindrical layers in the barrel, with a fourth (innermost) insertable barrel layer (IBL) installed during Long Shutdown 1 (LS1) to improve tracking performance in higher luminosity conditions. The three outer layers contain 80 million readout channels with pixels $50 \mu\text{m} \times 400 \mu\text{m}$ in size, while the IBL provides an additional 12 million readout channels, with sensors of $50 \mu\text{m} \times 250 \mu\text{m}$ pixels. Three discs form each of the endcaps, using the same pixel modules as those in the barrel. A charged particle typically produces a hit in all four layers of the pixel detector and can be measured to an uncertainty of $10 \mu\text{m}$ in $r\phi$ and $115 \mu\text{m}$ in z (r) in the barrel (discs).

Beyond a radial distance of approximately 14 cm from the beam line, the precision pixel layers are replaced by a silicon strip detector called the SemiConductor Tracker (SCT), which covers the intermediate region up to a radius of ~ 60 cm. At these distances, the level of precision offered by a pixel geometry becomes less important, and the lesser material requirements and simpler design of a strip-based system lends itself to being the more economically favourable option when considering the costs associated with mass-producing silicon sensor elements [40]. Each SCT module consists of a heat transport plate with two wire-bonded microstrip detectors glued onto each face. In the four barrel layers, one set of strips is aligned parallel to the beam direction to provide measurements in $r\phi$, while the other is offset by a small 40 mrad angle to also allow for measurements in z . A similar configuration is used in the nine discs forming each endcap, with

4.1. ATLAS detector

strips aligned either radially or at a 40 mrad offset. The SCT provides approximately 6.3 million readout channels in total, with a resolution of $17\ \mu\text{m}$ in $r\phi$ and $580\ \mu\text{m}$ in z (r) in the barrel (discs).

Finally, the Transition Radiation Tracker (TRT) is a straw-based detector forming the outermost layer of the ID. The low resolution per point (compared to the silicon-based systems) is balanced by high-rate readouts and a larger number of measurements (~ 30) offered per track. The barrel is composed of approximately 50 000 straws, each 4 mm in diameter and containing a central wire immersed in a xenon gas mixture. The straws are oriented axially and separated in the middle to facilitate readouts at both ends of the barrel (improving its performance at high occupancy), while each endcap consists of 18 wheels constructed from approximately 320 000 radially-pointing straws in total. The TRT provides measurements up to $|\eta| = 2$, with a resolution of approximately $50\ \mu\text{m}$ when averaging over all straw hits.

4.1.2 Calorimeters

The tracking detectors are encapsulated by a system of non-compensating sampling calorimeters¹, which function by absorbing incident particles to produce electronic signals related to the energy of the particle. The detector contains two calorimeter systems: a liquid argon (LAr) calorimeter [41–43], which uses LAr as an active medium to obtain precise measurements of electromagnetic and hadronic particle energies in the central and forward regions; and a scintillator tile calorimeter [44–46] at larger radii from the interaction point, providing coarser granularity but being more cost-effective and sufficient for jet and missing energy measurements. The main components of the LAr and Tile calorimeter systems are depicted in Figure 4.3.

The LAr calorimeter consists of an electromagnetic barrel (EMB) and two endcaps, each placed in a vacuum-insulated aluminium cryostat to maintain a -184°C temperature necessary for argon to remain liquid. Lead sheet absorbers are folded into an accordion-like structure to increase the chance of stopping incident particles, segmented in the radial direction via copper-polyimide readout boards to form three ($|\eta| < 1.35$) or two ($1.35 < |\eta| < 1.475$) layers. There is also a thin presampler layer covering up to $|\eta| = 1.52$, mounted to the inside of the cryostat in front² of the EM layers, which is necessary for correcting for energy losses in materials upstream of the calorimeter [47]. Cells in (η, ϕ) are defined by the grouping and etching of the readout boards; in the first layer (EMB1), a combination of strip and square towers are used, with $\Delta\eta \times \Delta\phi$ cell granularities of $0.025/8 \times 0.1$ in $|\eta| < 1.40$, and 0.025×0.025 in $1.40 < |\eta| < 1.475$. Square towers are used in the second layer (EMB2), with a granularity of 0.025×0.025 ($|\eta| < 1.40$) or 0.075×0.025 ($1.40 < |\eta| < 1.475$). Finally, the third layer (EMB3) features 0.050×0.025 cells up to $|\eta| = 1.35$. In total, the EM barrel calorimeter provides 101 760 readout channels, with an additional 7 808 from the presampler sectors.

The endcaps of the LAr calorimeter each consist of an electromagnetic wheel (EMEC), two hadronic wheels (HEC), and the forward calorimeter (FCal) system. The EM endcap wheels are constructed from radially-pointing lead absorber plates, with a similar accordion structure as in the barrel. Arrangement of the absorber plates in this configuration imposes a limitation on the ratio of inner to outer radii, and as such, the component is constructed as two separate

¹Non-compensating calorimeters generate signals that are directly proportional to the energy of electromagnetic particles, but do not compensate for the invisible components in hadronic showers, resulting in a non-linear response to the latter. Sampling calorimeters generate signals using an active medium separate from the material used to absorb incoming particles (this is in contrast to homogeneous systems, in which both actions are achieved within the same material).

²In this case, “front” is used from the radial perspective, i.e. the presampler lies between the tracker and the first EM calorimeter layer.




Figure 4.3. Visualisation of the ATLAS calorimeter systems.

concentric wheels. The larger wheel, which covers the pseudorapidity range $1.375 < |\eta| < 2.5$, contains three sampling layers with cell granularities similar to those in the barrel layers. Beyond $|\eta| = 2.5$, precision physics can no longer be feasibly achieved, and so the smaller inner wheel was designed at a reduced cost target, covering $2.5 < |\eta| < 3.2$ with a coarser 0.1×0.1 cell granularity and split into only two sampling layers to reduce the number of readout channels to approximately 2 000. In addition to a presampler layer inserted prior to the first EMEC layer and covering $1.5 < |\eta| < 1.8$, the EMEC wheels contain $\sim 64\,000$ readout channels in total.

The final components of the LAr system are the HEC and FCal detectors, which provide sensitivity also to hadrons in the forward regions. The two independent HEC wheels utilise copper absorbers, each consisting of 32 identical modules placed around a central ring of readout boards that provide 5 632 readout channels in total, covering $1.5 < |\eta| < 3.2$ with 0.1×0.1 cells in $|\eta| < 2.5$ and 0.2×0.2 for $|\eta| > 2.5$. These wheels surround the FCal calorimeters, which are supported by external structural tubes integrated into each of the endcap cryostats to allow for placement as close to the beampipe as possible. However, as a consequence, the decision was made to displace the first layer of the FCal by approximately 1.2 m from the front of the endcap cryostat, in order to reduce the contamination of neutrons (produced in secondary interactions with the detector material) in the central cavity.

In total, the forward detectors contain three layers that cover the pseudorapidity range of $3.1 < |\eta| < 4.9$. The first layer utilises copper absorbers, while a denser sintered tungsten alloy is selected in the final two layers to increase the active interaction length of the system to compensate for its reduced physical length. A denser choice of material provides the additional benefit of reducing leakage and pile-up, particularly to the back wheel of the HEC. Cell granularities are on the order of 1×1 to 5×5 in $\Delta\eta \times \Delta\phi$, depending on pseudorapidity, with 3 524 readout channels (in total) providing a balance between response speed and noise performance.

4.1. ATLAS detector

Surrounding the LAr system is the scintillator tile calorimeter designed for the capture and measurement of hadronic particle showers, which tend to extend beyond the radial range of the LAr volume. The tile calorimeter utilises steel absorbers in conjunction with plastic scintillating tiles as its active mechanism. The detector is divided into a central barrel ($|\eta| < 1.0$) and two extended barrels ($0.8 < |\eta| < 1.7$), each consisting of 64 wedge-shaped modules oriented and arranged akin to the carpels of an orange. Modules are constructed from layers of alternating steel plates and scintillator tiles, with wavelength-shifting fibres placed along the tile edges to carry scintillation light to the readout photomultiplier tubes. In total, the tile system offers approximately 10 000 readout channels over three layers, with cell sizes of $\Delta\eta \times \Delta\varphi = 0.1 \times 0.1$ in the first two layers and 0.2×0.1 in the third.

4.1.3 Muon spectrometer

The Muon Spectrometer (MS) [48] consists of a magnet system, which in conjunction with readout and dedicated trigger chambers allow for the measurement of muon momenta via their deflection within the magnetic field. The detector is designed to achieve a 1% resolution on mass and momentum measurements over a wide $|\eta|$ range in order to reliably reconstruct the channels targeted in the LHC physics programme.

Three superconducting air-core toroid magnets are employed to produce the magnetic field within the barrel and in each endcap. The barrel magnet is 25 m long, with an outer diameter of 20 m, and generates a magnetic field with an average bending power of 3 Tm in the central pseudorapidity range $|\eta| \leq 1.0$. Smaller endcap magnets, approximately 5 m long and 11 m in diameter, cover $1.4 \leq |\eta| \leq 2.7$ with typical bending powers of 6 Tm. Each toroid is made of eight flat coils symmetrically arranged around the beam axis. In the barrel, these coils are each contained within a cryostat and aligned with the beam direction, and act also as a support frame for the detection chambers. The smaller endcap coils are contained within a single cryostat and azimuthally offset by 22.5° in order to optimise the bending power in the transition region $1.0 \leq |\eta| \leq 1.4$. With this setup, the magnetic field is mostly orthogonal to muon trajectories across the full pseudorapidity range.

The rate and occupancy of expected backgrounds (originating either from the primary collision, or in the form of low energy radiation produced in secondary interactions with the detector material) drive the granularity requirements of the tracking chambers. For the most part, detection is provided by monitored drift tubes (MDTs), which are thin-walled aluminium tubes 30 mm in diameter with lengths varying between 70–630 cm. Each tube contains a tungsten–rhenium (W–Rh) wire immersed in a gas mixture of argon, methane and nitrogen ($\text{Ar}/\text{CH}_4/\text{N}_2$), and provides a measurement resolution of approximately $80 \mu\text{m}$. The barrel (endcap) consists of three layers (“stations”) of MDT readout chambers that cover $|\eta| < 1.0$ ($1.0 < |\eta| < 2.7$). Each chamber improves on the single-wire resolution by utilising two sheets of MDT elements attached to either side of a rigid support frame, with each sheet being four (for the inner station) or three (middle and outer stations) tubes thick. In the innermost ring of the inner endcap station ($2.0 < |\eta| < 2.7$), higher-granularity cathode strip chambers (CSCs) are used in favour of the MDTs in order to meet the higher expected rates in this region. Typical electron drift times in the CSC readouts are less than 30 ns (roughly an order of magnitude lower than that of the MDTs), leading to improved time resolutions (7 ns) over the MDT readouts.

The MS also implements a dedicated trigger system covering the pseudorapidity region up to $|\eta| \leq 2.4$. To be suitable for nominal physics sensitivity, the system has to be capable of selecting muons above p_T thresholds on the order of 10–20 GeV, down to approximately 5 GeV




Figure 4.4. The ATLAS Trigger and Data Acquisition system during Run 2 of the LHC. The expected peak event and data rates through each component is shown.

for probing CP-violating effects and *B*-physics phenomena, at the expected luminosity of the LHC. Two technologies are selected based on the expected rates in the barrel and endcap regions, physics requirements, and overall cost. In the barrel, this is fulfilled by resistive plate chambers (RPCs) located to either side of the middle and outer MDT stations, while three thin gap chamber (TGC) stations, placed near the middle MDT station, are used in each of the endcaps. RPC units are simple and cheap to manufacture, consisting of two resistive plates separated by a gas gap, and provide spatial and timing resolutions of $1\text{ cm} \times 1\text{ ns}$. TGCs feature longer electron drift times, but nevertheless remain capable of resolving arrival times to less than 25 ns at 99% efficiency, allowing for good bunch crossing identification.

4.2 Trigger and data acquisition

The LHC collides protons in bunches spaced 25 ns apart, corresponding to an event rate of 40 MHz. However, it is both unfeasible and undesirable to process and store a record of every collision event: the former due to limitations of the readout electronics, computational power available and data storage requirements; and the latter because most events are uninteresting from the physics perspective, anyway.

The filtering of events into manageable rates is handled by the trigger and data acquisition (TDAQ) system [49–53]. There are two main components to the trigger system, summarised in the left part of Figure 4.4. The first is a hardware-based “Level-1” (L1) trigger, which utilises

4.3. Signal reconstruction

reduced-granularity readouts from the LAr and tile calorimeters (L1Calo) and the dedicated RPC and TGC trigger chambers of the MS (L1Muon). These two systems respectively identify events with significant activity in the calorimeter (indicating the presence of photons, electrons, tau-leptons, hadronic objects or missing energy) or containing muon candidates. A third system added during Run 2, L1Topo, combines the signals from L1Calo and L1Muon to construct topological variables that target more complex final states potentially containing both calorimeter objects and muons. This enhances the signal-to-background ratio in selected events, and also allows for the triggering of lower p_T objects, which would otherwise be impossible due to the vastly increased rates at lower thresholds. L1 trigger signals from each component are forwarded to the Central Trigger Processor (CTP) for fast decision-making as to whether or not an event should be recorded. Overall, the event rate is reduced to a maximum of approximately 100 kHz by the L1 trigger system, split between various signatures defined by the physics goals of ATLAS, with a maximum processing time of $2.5\ \mu\text{s}$ per event.

Events that satisfy the L1 trigger decision are passed through to the software-based High-Level Trigger (HLT) for further processing. At this step, the full-granularity event information from all detector components is made available, along with regions of interest (RoI) in η - ϕ space identified by the L1 system. Object reconstruction at the HLT follows offline techniques where possible, in order to improve the harmonisation of objects defined between on- and offline processing [54]. ID track reconstruction follows a two-step procedure, with an initial fast-tracking algorithm [55] used as a seed for precision track reconstruction using the offline tracking software. In the calorimeter, energy clusters are reconstructed either in the RoIs (*RoI-based* mode) or over the full detector volume (*full-scan* mode) [56]. The former is used when defining localised objects such as electrons, photons and taus, while the latter is used for jet and E_T^{miss} reconstruction. Muons are reconstructed using a combination of MDT hits and ID tracks. Trigger selections based on b -jets or B -meson signatures are also available, the latter being particularly improved in Run 2 due to the introduction of the L1Topo system. In all cases, computational requirements are minimised by only considering narrow RoIs wherever the full global event information is unnecessary. In total, each event is processed by the HLT within 200–400 ms using a commercial CPU farm containing approximately 40 000 CPU cores. Events that pass the HLT selection are saved for offline analysis, occurring at a rate of 1.5 kHz.

The transport of information between the detector electronics and the trigger systems, and ultimately through to permanent storage for offline analysis, is facilitated by the data acquisition (DAQ) system (Figure 4.4, right). Readout drivers (RODs) send events that receive a positive L1 signal to the readout system (ROS) through optical links that support a peak data rate of approximately 160 GB/s. These events are stored in a memory buffer until computational resources become available, at which point the HLT will request for the data from the ROS. HLT output events are collected by the data logger (abbreviated SFO according to its historical name, “sub-farm output”) and delivered to the ATLAS Tier-0 facility [57] at a nominal throughput of 1.5 GB/s, with a buffer capacity of 340 TB to allow for approximately 24 hours of data to be stored locally while pending the transfer.

4.3 Signal reconstruction

Particles traversing the detector volume generate electronic readout signals: hits in the silicon elements or straws in the ID layers; energy deposits in the LAr or tile calorimeter cells; or measurements in the MDT chambers of the muon system. However, these signals, considered indi-

vidually, do not provide particular insight into the physics that occurred in the collision event. Instead, a more useful interpretation is obtained by considering an aggregation of the signals into several basic objects, with properties that can be associated with the physical particle that generated the signals. This has already been alluded to in the previous section, being a crucial ingredient for the ATLAS trigger system. In this section, I will describe in more detail the techniques used to reconstruct charged particle tracks in both the ID and MS subdetectors, and dense energy clusters within the combined calorimeter system.

4.3.1 Inner detector tracks

Signal reconstruction in the ID is crucial for identifying vertices, facilitating particle classification, and measuring the momenta and positions of charged particles to high precision [38]. Standard tracks are associated with five parameters: the transverse (d_0) and longitudinal (z_0) impact parameters corresponding to the distance between the vertex and the beamspot position in those respective directions; the azimuthal (φ) and polar (θ) angles; and the ratio of charge to momentum (q/p). Tracks are reconstructed using both an “inside-out” and an “outside-in” approach, with the latter improving the acceptance for tracks of charged particles produced away from the beamline (e.g. photon conversions into e^+e^- pairs).

The inside-out reconstruction begins with the formation of 3-dimensional “space points” consisting of clusters of adjacent hits in the pixel or SCT detector layers. Track seeds are defined using triplets of space points, with several requirements (e.g. $d_0 < 5$ mm and $z_0 < 150$ mm cuts, and the presence of a fourth “confirmation space point” in a different silicon layer) to avoid the reconstruction of low-quality tracks. For each seed, an expected trajectory through the detector layers is then calculated; this is used to identify a subset of modules in which the particle could have produced a hit. This defines a “search road” for the seed, with the track fitting algorithm being performed only along the road to vastly reduce the extent of possible combinatorics and improving the compute time overall.

Hits along the search road are used to progressively extend the seed via a Kalman filtering algorithm [58]. This operates on a principle similar to that of Bayesian inference, in which prior knowledge (i.e. that charged particles in a magnetic field follow helical trajectories) is combined with observations (i.e. the measurements within the precision layers). Deviations from a helical path are allowed if the track seed is compatible with a region of interest (ROI) in the calorimeter, improving the efficiency of reconstructing tracks from electrons that exhibited bremsstrahlung radiation. However, while the Kalman filter provides a computationally fast algorithm for track finding, it does not account for overlaps in the hits used to construct each track candidate. This is handled by a dedicated ambiguity solver, which assigns a score to each track based on quality criteria such as the number of hits shared with another candidate. Tracks evaluated to be of lower quality often correspond to fakes, and are consequently rejected.

After this, an attempt is made to include TRT hits in the track: each candidate is extrapolated into the TRT volume, with a Kalman filter again used for the track finding. A global χ^2 track fit is redone after all TRT hits have been added. If the track precision is found to improve with the addition of these hits, they are retained; otherwise, the track candidate is defined using only the silicon layer hits. This avoids the use of TRT measurements when significant outliers are found, or if there are too many missing hits along the track’s trajectory.

The next reconstruction pass is performed outside-in, and begins by identifying track segments in the TRT that are matched to a calorimeter ROI. Possible track trajectories in the TRT barrel (endcap) are constructed as projections in the 2-dimensional r - φ (z - φ) plane, in accor-

4.3. Signal reconstruction

dance with the alignment of the straw detection elements in the respective segments. The 2D representation of flight trajectories allows for a simple histogramming approach for track finding in the TRT, and is used to define a helix that forms the initial seed for global track finding. Hits from the silicon detector layers are then used to progressively extend the initial track road all the way to the vertex using the same Kalman filter approach as in the inside-out method. The track fit is updated with every set of new measurements from each successive ID layer considered, and after all of the precision layers have been included, an ambiguity processor is again used to reduce the fake rate.

An extrapolation back into the TRT and a final χ^2 fit reduces the uncertainties of the track. The algorithm takes into account multiple scattering and energy losses from ionisation (“ μ -fit”) by introducing corrections to the expected track curvature and additional terms in the covariance matrix of the fit. Corrections for bremsstrahlung losses (“ e -fit”) are also considered if the track is identified to have originated from an electron, although this configuration is only chosen over the μ -fit case if it leads to a greater number of associated precision layer hits.

Final track candidates are assigned a quality score based on the number of precision hits found, and any failing to meet a baseline set of criteria are discarded. At least seven hits in the precision layers are required (of which at least six should be unique, i.e. not shared with another track); and, unless nine or more precision hits were found, a minimum of nine hits in the TRT is also required, with hits registered in at least 70% of the straws that are crossed. A $p_T > 0.5$ GeV threshold rejects low momentum charged particles, which tend to spiral in the magnetic field and form “loopers” that persist over the duration of several bunch crossings.

4.3.2 Calorimeter clusters

Energy deposits in calorimeter cells are used to reconstruct all of the standard physics signatures in ATLAS, except muons. One important point of consideration when reconstructing objects in the calorimeter is that spurious signals can be generated by noise in the cells, which arise predominantly from the detector electronics and additional soft interactions between protons in each bunch crossing (i.e. pile-up):

$$\sigma_{\text{noise}} = \sqrt{\left(\sigma_{\text{noise}}^{\text{electronic}}\right)^2 + \left(\sigma_{\text{noise}}^{\text{pile-up}}\right)^2}. \quad (4.1)$$

On average, there are approximately $\langle\mu\rangle = 30$ interactions occurring per bunch crossing in the Run 2 operational conditions, leading to significant levels of pile-up noise particularly in the forward regions of the calorimeter.

In addition to this, the LAr calorimeter has a characteristic response time of approximately 500 ns, which is very slow when compared to the nominal 25 ns time interval between bunch collisions. In principle, this means that a measurement in the detector at any given time could be biased, or even entirely seeded, by interactions that occurred in previous bunch crossings. To mitigate these out-of-time effects, LAr calorimeter signals are digitally sampled with a shaping amplifier to have a bipolar pulse (Figure 4.5). The shape of this response is tuned to the nominal 40 MHz event rate and $\langle\mu\rangle$ to feature an overall net zero integral when summing up pile-up measurements over multiple bunch crossings. Of course, this will not always be fully efficient (given the dynamic nature of μ from event to event); but the approach nevertheless allows for us to leverage the good spatial resolution of LAr detectors, and forms an important first stage of pile-up suppression during data-taking.




Figure 4.5. The nominal triangular response of the LAr calorimeter, in contrast to the shaped bipolar response for mitigating out-of-time pile-up effects (adapted from [59]).

To infer physically significant objects in the calorimeter, cells containing energy depositions are used as input to a clustering algorithm [60] with considerations in place for the aforementioned contributions from pile-up and electronic noise. To facilitate the clustering process, each cell is assigned a significance,

$$\zeta_{\text{cell}}^{\text{EM}} = \frac{E_{\text{cell}}^{\text{EM}}}{\sigma_{\text{noise,cell}}^{\text{EM}}}, \quad (4.2)$$

which is the ratio of its energy (at the EM scale) to the expected contribution from noise. Formation of topological clusters proceeds by considering the cell significances, along with three parameters S , N , and P respectively corresponding to the thresholds for seeding, growth control, and cell filtering of the clusters (with $S > N \geq P$), as follows: firstly, cells with a significance of $|\zeta_{\text{cell}}^{\text{EM}}| > S$ are used to seed an initial set of proto-clusters. Then, an expansion in all three dimensions occurs by checking the cells that neighbour each proto-cluster (i.e. cells that are adjacent in the same sampling layer, or overlapping in (η, φ) in a different layer). Any neighbouring cell satisfying $|\zeta_{\text{cell}}^{\text{EM}}| > N$ is iteratively collected into the proto-cluster until the algorithm encounters a neighbouring cell below this threshold. If the final cell has a significance $|\zeta_{\text{cell}}^{\text{EM}}| > P$, it is also added; then, the growth is terminated. (If $P = N$, the boundary of the proto-cluster is simply defined by the cells with $|\zeta_{\text{cell}}^{\text{EM}}| > N$.)

Of course, given the potentially high particle multiplicity in hard-scattering collisions, it can be (and often is) the case that several high-significance seeds are situated close-by to one another. During the expansion phase, if any cell with $|\zeta_{\text{cell}}^{\text{EM}}| > N$ is found to neighbour two proto-clusters, then they are merged into one proto-cluster. However, this can lead to large clusters that encompass multiple particles, which is preferable to avoid if possible, given that the original motivation behind cell clustering was to produce a set of calorimeter signals associable with the individual particles that enter the system. Thus, after the set of proto-clusters are fully defined, a splitting step is performed: if a proto-cluster contains multiple local energy maxima (defined as a cell with $E_{\text{cell}}^{\text{EM}} > 500$ MeV and having at least four lower-energy neighbours) then it is split in the three spatial dimensions between each maximum. Any cell that lies on the boundary of two

4.3. Signal reconstruction

split clusters is geometrically weighted to share an appropriate fraction of its energy with each, as follows:

$$w_{\text{cell},1}^{\text{geo}} = \frac{E_{\text{clus},1}^{\text{EM}}}{E_{\text{clus},1}^{\text{EM}} + r E_{\text{clus},2}^{\text{EM}}}, \quad (4.3)$$

$$w_{\text{cell},2}^{\text{geo}} = 1 - w_{\text{cell},1}^{\text{geo}}, \quad (4.4)$$

where $E_{\text{clus},1}^{\text{EM}}$ and $E_{\text{clus},2}^{\text{EM}}$ denote the energy of the two clusters, and $r = \exp(d_1 - d_2)$, with d_1 and d_2 being the distance from the cell to the centre of gravity of the corresponding cluster.

ATLAS uses an $S = 4$, $N = 2$ and $P = 0$ clustering configuration optimised for the reconstruction of hadronic signatures. The choice of $P = 0$ allows for cells near the noise level to be clustered as long as they neighbour a significant cell. Another important feature of the algorithm is that the thresholds considered take into account the absolute value of the cell significance, rendering it agnostic to cells that contain negative signals (which is possible due to the bipolar shaping described above). The inclusion of both positive and negative energy cells facilitates a local cancellation of random noise fluctuations, mitigating their overall contribution to a cluster. For negative energy cells significant enough to seed a cluster, it is usually the case that its neighbouring cells received a residual signal from the same out-of-time particle, thus also possessing negative energy and resulting in a cluster with an overall energy $E_{\text{clus}}^{\text{EM}} < 0$. Such clusters can provide an average cancellation of contributions from in-time pile-up in global event observables such as $E_{\text{T}}^{\text{miss}}$, but are otherwise not considered when reconstructing physics objects.

Kinematic properties of a cluster are defined based on the cells it contains. Its position in the η - ϕ plane is calculated as the barycentre of the cell positions, weighted to favour the more energetic cells:

$$\eta_{\text{clus}} = \frac{\sum_{\text{cells}} w_{\text{cell}}^{\text{geo}} \cdot |E_{\text{cell}}^{\text{EM}}| \cdot \eta_{\text{cell}}}{\sum_{\text{cells}} w_{\text{cell}}^{\text{geo}} \cdot |E_{\text{cell}}^{\text{EM}}|}, \quad (4.5)$$

$$\phi_{\text{clus}} = \frac{\sum_{\text{cells}} w_{\text{cell}}^{\text{geo}} \cdot |E_{\text{cell}}^{\text{EM}}| \cdot \phi_{\text{cell}}}{\sum_{\text{cells}} w_{\text{cell}}^{\text{geo}} \cdot |E_{\text{cell}}^{\text{EM}}|}, \quad (4.6)$$

with the absolute value of cell energies being taken to avoid negative signal contributions from being interpreted as positive signals in the opposite hemisphere of the detector. The geometric weights (eqs. (4.3) and (4.4)) are relevant only for the cells that border another cluster. The total energy is a simple sum, correctly taking into account any negative signal contributions:

$$E_{\text{clus}}^{\text{EM}} = \sum_{\text{cells}} w_{\text{cell}}^{\text{geo}} \cdot E_{\text{cell}}^{\text{EM}}. \quad (4.7)$$

A plethora of additional moments can be defined by looking at the distribution of cells within each cluster, which can aid in the characterisation and calibration of the cluster (see e.g. ref. [60] for further reading).

4.3.3 Muon chamber tracks

Hits in the detection chambers of the MS are used to reconstruct tracks [61, 62]. This proceeds in two steps: firstly, local segments of a track are identified in each MDT or CSC chamber, with each typically providing four to eight measurements of the muon's position in η at a precision of




Figure 4.6. Display of a candidate di-Higgs to 3 lepton event in the transverse plane of the detector. Tracks are shown as cyan lines, and hits in the muon spectrometer as light blue blocks. Energy deposits in the LAr and tile calorimeters are shown as yellow boxes. From these signals, two muons (green lines), an electron (red line), a jet (grey cone), and missing transverse momentum (dashed magenta line) can be reconstructed. Image taken from the auxiliary material of ref. [63].

$40\ \mu\text{m}$. These local track segments are then combined in a global track-finding step, beginning from the outermost layer and performing the fit as inner-layer segments are iteratively included until the innermost layer is reached. The fitting of the track follows the same method used to reconstruct tracks in the ID.

4.4 Physics objects

The final step of reconstruction involves taking a combination of the various detector signals and using them to produce a representation of the event in terms of physical particles. In ATLAS, dedicated algorithms are employed to define the following final state objects: hadronic jets, taus, muons, electrons, photons, and missing transverse momentum. Several techniques are also used to further classify jets based on the flavour of hadrons that they contain (referred to as *tagging*). An example of a candidate three-lepton event from di-Higgs production is presented in Figure 4.6, illustrating the reconstruction of two muons, an electron, a jet, and missing transverse momentum from the more fundamental detector signals.

4.4.1 Hadronic final states

Collisions at the LHC are dictated largely by the theory of quantum chromodynamics (QCD), and the features of the theory essentially guarantee the ubiquity of partons in every event: any parton produced in a hard scatter collision will be unstable and initiate a *parton shower*, greatly

4.4. Physics objects

increasing the particle multiplicity of the final state via perturbative (e.g. hadronic decays, final state radiation) or non-perturbative QCD interactions (e.g. soft gluon emissions), or the QED radiation of photons or leptons. Each of the partons generated will then undergo *hadronisation* to produce the stable and observable final state particles.

In particle physics, the individual signature of these partonic final state particles are referred to as *jets*, although there can be several interpretations of what a “jet” constitutes depending on the context. From the theory perspective, *parton jets* refer to the final state partons themselves (prior to showering and hadronisation), arising from fixed-order calculations in perturbation theory. Phenomenologically, any partons produced in a collision will shower and hadronise, forming *particle jets* corresponding to the stable particles produced in those processes. Finally, there is the experimental viewpoint wherein particles are characterised simply by the signature they impart within a detector, resulting in *detector jets* that consist of many localised tracks and clusters of energy deposits in the calorimeters, with topological and kinematical properties that in principle reflect the original final state parton.

Reconstruction

There are several challenges associated with jet clustering. Firstly, clustering depends crucially on the geometry and granularity of the detector, and the algorithms adopted by experiments have historically been driven primarily by these considerations. However, given the expected kinematic reach achievable at the LHC, it also becomes desirable for jet finding to remain stable when performing calculations to higher order. This is important from the theoretical perspective to allow for particle jets to be simulated at any order whilst remaining an accurate prediction for experimental observations (i.e. detector jets). In particular, the algorithm should feature both infrared and collinear (IRC) safety: low-energy “contamination” (e.g. soft gluon emissions or the presence of pile-up particles) should not result in the merging or splitting of clustered jets (infrared safety); and the splitting of an energetic particle into lower-energy particles should not bias the reconstruction (collinear safety).

In ATLAS, the FastJet [64] implementation of the anti- k_t clustering algorithm [65] is used, with a radius parameter of $R = 0.4$ or $R = 1.0$ for defining small or large radius jets, respectively. Small- R jets are defined using Particle Flow Objects (PFOs) as input. (An alternative calorimeter-only definition, “EMTopo”, was the default for Run 1 and the early stages of Run 2; these jets use only topo-clusters (described in section 4.3.2) as input to the clustering.) Further discussion on particle flow and its particular implementation in ATLAS is presented in Chapter 5. The $R = 1.0$ configuration is used for reconstructing the decays of boosted (i.e. high momentum) heavy objects, e.g. $W \rightarrow q\bar{q}$ or $t \rightarrow bW \rightarrow bq\bar{q}$, which result in large showers of collimated particles that cannot be resolved individually. The substructure of constituents within these jets provides information for identifying the origins of the jet [66–69].

Jets defined by the anti- k_t algorithm are IRC-safe and resilient against additional soft radiation, having a regular (i.e. conical) shape in the majority of cases. The algorithm begins by defining a distance d_{iB} between each entity i and the beamline, and d_{ij} between every pair of entities i and j , as:

$$d_{iB} = k_{t,i}^2, \quad (4.8)$$

$$d_{ij} = \min(d_{iB}, d_{jB}) \frac{\Delta R_{ij}^2}{R^2}. \quad (4.9)$$




Figure 4.7. The ATLAS jet calibration chain. Image from ref. [73].

The quantity $\Delta R_{ij}^2 = \Delta y_{ij}^2 + \Delta \varphi_{ij}^2$ is a measure of the geometric separation between the particles (where y and φ denote the rapidity and azimuthal angle, respectively), and k_t is the transverse momentum of the particle. R and p are parameters of the algorithm, with the former representing the jet radius and the latter controlling the weight given to the energy of the particles. The anti- k_t algorithm sets a value of $p = -1$. The case of $p = 1$ corresponds to the inclusive k_t [70] algorithm, and $p = 0$ to the Cambridge–Aachen (C–A) [71, 72] algorithm.

In the first pass of the algorithm, entities correspond simply to the list of input objects (i.e. PFOs). Out of the list of all d_{iB} and d_{ij} values, if a d_{ij} has the smallest value then the corresponding particles i and j are combined and the distances are recalculated. If the smallest value instead corresponds to a d_{iB} then the particle i becomes a jet candidate and is removed from the list. In this way, soft objects will be clustered to nearby hard particles before they are grouped with other soft objects. This procedure repeats until no entities remain in the list.

Calibration

Reconstructed jets can have properties that differ from the particle level due to contributions from pile-up, the non-compensating nature of the ATLAS calorimeters, or the geometry and granularity of the detector. This necessitates a calibration step to correct the energy scale of reconstructed jets to match that of jets predicted by simulation. In addition, corrections derived based on the known detector geometry are applied to improve the overall jet resolution. The discussion here will focus on small radius jets, as these are the objects that are relevant to the studies presented in this thesis.

The jet calibration chain used in Run 2 is described in ref. [73] and summarised in Figure 4.7. In the first step, the jet p_T is corrected to remove the expected contributions from pile-up:

$$p_T^{\text{corr}} = p_T^{\text{reco}} - \rho \times A - \alpha \times (N_{PV} - 1) - \beta \times \mu. \quad (4.10)$$

The first term takes into account the average p_T density of pile-up contributions (ρ) in the y - φ plane, and the jet area (A). However, because ρ is estimated using simulations of pile-up jets in the central region of the detector ($|\eta| < 2$) but applied to jets in all regions of the detector, the p_T of forward jets can remain sensitive to the amount of pile-up in an event. The residual correction terms serve to remove these dependencies, using the number of primary vertices (N_{PV}) and the mean number of interactions per bunch crossing (μ) as estimators for the level of in-time and out-of-time pile-up contamination, respectively.

4.4. Physics objects

After correcting for pile-up contributions, an MC-based calibration step corrects the jet four-vector back to the particle-level energy scale. Truth jets in simulation, constructed by passing all stable truth particles (excluding muons, neutrinos and pile-up particles) as input to the anti- k_t algorithm, are used as the reference objects for this correction. The energy response of matched reconstructed and truth jet pairs, calculated as $\mathcal{R} = E_{\text{reco}}/E_{\text{truth}}$, is derived in bins of E_{truth} and η_{det} , where the latter represents the pseudorapidity of the jet as calculated with respect to the geometrical centre³ of the detector. A numerical inversion procedure [74] is then used to parametrise the mean response in E_{reco} , and the calibration factor is calculated as the inverse of this quantity.

The calibration steps described to this point constitute the MC-based jet energy scale (MCJES) corrections, and are applied ubiquitously to all reconstructed jets. However, jets originate from a variety of sources; in particular, gluon-initiated jets will typically exhibit a wider shower of softer particles when compared to quark-initiated jets, which tend to contain lower multiplicity but harder particles. These differences result in an increase of the jet *resolution*, which is a measure of the variance in the jet response.

The Global Sequential Correction (GSC) aims to account for these effects by using knowledge of the detector material and geometry. Six global jet observables⁴ are identified upon which the jet p_T response depends:

- f_{charged} , the fraction of jet p_T that its ghost-associated tracks⁵ carry
- f_{Tile0} , the fraction of jet's energy deposited in the first layer of the tile calorimeter
- f_{LAr3} , the fraction of jet's energy deposited in the third layer of the LAr calorimeter
- n_{trk} , the number of tracks ghost-associated with the jet
- w_{trk} , the average distance (“width”) between ghost-associated tracks and the jet axis
- n_{segments} , the number of ghost-associated muon track segments (“punch-through”)

Independent correction factors are derived for each of these observables and applied sequentially to the jet four-momentum. This correction improves the jet resolution without affecting the mean response.

Finally, an *in situ* calibration is applied to correct any remaining differences between jets in data and MC, which can arise from a variety of inaccuracies in the simulation. For this, the jet response is measured in data using topologies in which a well-calibrated object or system can be used as reference for the reconstructed jet; for example, in $\gamma + \text{jet}$ processes where the prompt photon can be balanced against the hadronic recoil in the event. This is then compared to the response expected from simulation to derive a calibration factor that is applied to jets in data:

$$c = \frac{\mathcal{R}_{\text{in situ}}^{\text{data}}}{\mathcal{R}_{\text{in situ}}^{\text{MC}}} . \quad (4.11)$$

4.4.2 Electrons and photons

The reconstruction of electrons (e) and photons (γ) are closely intertwined due to the possibility of bremsstrahlung radiation of a photon from an electron, or the conversion of a photon into an

³As opposed to being calculated with respect to the primary vertex.

⁴The observables are “global” in the sense that they account for multiple detector systems.

⁵This refers to tracks that fall within the jet area; in practice, this is implemented by passing tracks as input to the anti- k_t algorithm, but with momenta set to zero so that they do not impact the clustering.

electron–positron pair. In general, the signature of these particles will consist of a collection of clusters in the EM calorimeter, along with the presence of a track(s) in the case of electrons or converted photons.

Standard 4-2-0 topo-clusters are used as the starting point for reconstruction [75, 76]. However, since e/γ showers are not expected to extend into the tile calorimeter, a preselection and modification of the clusters is first performed by selecting only those with an EM fraction above a threshold of $f_{\text{EM}} > 0.5$, optimised to balance electron efficiency against pile-up rejection. EM clusters are then defined by removing contributions from tile calorimeter cells from the selected topo-clusters, with cluster moments being recalculated accordingly after this modification.

Bremsstrahlung and photon conversion can result in the formation of more than one cluster in the calorimeter. To reconstruct only one particle in these scenarios, an attempt is made to combine EM clusters into *superclusters*. In the order of descending p_T , clusters with $p_T > 1$ GeV that are matched to a track, or with $p_T > 1.5$ GeV with no matched tracks, are designated as “seed” clusters for electrons and photons, respectively. A scan for the presence of “satellite” clusters, originating from secondary EM showers from the primary electron or photon, is then performed within a window of $\Delta\eta \times \Delta\phi = 0.075 \times 0.125$ around each seed cluster. For photon seeds, satellite clusters are required to be compatible with a reconstructed conversion vertex. For electron seeds, a second scan is also performed over a larger $\Delta\eta \times \Delta\phi = 0.125 \times 0.3$ window, with a requirement that satellite clusters in this window are matched to the same track as the seed cluster. Each EM cluster can only be designated as a seed or satellite once; the algorithm continues until all clusters have been considered.

Reconstructed electrons and photons are calibrated using both simulation-based and data-driven techniques [76, 77]. Simulated single particle events are used to train a multivariate (MVA) regression algorithm, using deposits within the calorimeter as inputs for learning the energy of the electron or photon. The MVA calibration is applied to reconstructed objects in both data and simulation. After this, a correction to the energy scale and resolution is applied, the former to objects in data and the latter to simulation, to reduce any residual discrepancies between the two:

$$E^{\text{data,corr}} = \frac{E^{\text{data}}}{(1 + \alpha_i)}, \quad \left(\frac{\sigma_E}{E}\right)^{\text{MC,corr}} = \left(\frac{\sigma_E}{E}\right)^{\text{MC}} \oplus c_i. \quad (4.12)$$

The calibration parameters α_i and c_i are derived from template fits of simulated $Z \rightarrow ee$ events to the data.

To increase the purity of the reconstruction, several identification working points are defined. A likelihood discriminant is constructed for the identification of electrons, taking into account the features of associated clusters, the quality of matched tracks, and various track–cluster variables. The default “Tight” working point achieves an 80% efficiency for selecting prompt electrons at $E_T = 40$ GeV [78]. For photons, a cut-based approach based on shower shape parameters provides discrimination between prompt photons and fakes (predominantly arising from $\pi^0 \rightarrow \gamma\gamma$) [79]. An isolation condition is often also imposed on prompt photons; this requires that the amount of energy within a cone of $\Delta R = 0.2$ or $\Delta R = 0.4$ (depending on the working point) falls below a specified threshold.

4.4.3 Muons

Muons are minimally-ionising particles (MIPs) that can result in the presence of tracks in the MS and ID detectors, and occasionally also depositions of energy in the calorimeter. To account for the variety of possible signatures a muon can impart within the overall detector volume,

4.4. Physics objects

five different algorithms are employed for their reconstruction [80, 81]. The product of these approaches are referred to as combined (CB), inside-out combined (IO), MS extrapolated (ME), segment-tagged (ST), and calorimeter-tagged (CT) muons.

CB muons are reconstructed by extrapolating MS tracks into the ID and performing a combined fit using the hits measured in both systems. This can extend beyond $|\eta| = 2.5$ if short track segments in the pixel and SCT detectors are present; such cases define a subset of CB muons known as silicon-associated forward (SiF) muons. The IO approach is similar, but begins instead with tracks in the ID and attempts an extrapolation into the MS. A stand-alone track within the MS is not required in this case (only the presence of hits compatible with the ID track), increasing the overall muon reconstruction efficiency. ME muons are defined when a track reconstructed in the MS cannot be matched to a track in the ID. This can occur particularly when muons fall beyond the acceptance of the ID ($|\eta| > 2.5$) but within that of the MS (up to $|\eta| = 2.7$). For these cases, the muon is defined by extrapolating the MS track back to the beamline. ST muons correspond to ID tracks that can be matched to a reconstructed segment(s) in the MS; the parameters of the muon are taken to be those of the ID track, with no combined fit being performed. Finally, CT muons correspond to ID tracks that are matched to calorimeter energy deposits consistent with those expected of a MIP.

A calibration of the momentum scale is applied to reconstructed muons in data to account for shifts in the alignment of the ID or MS [81]. After this, a momentum calibration is applied to reduce any remaining discrepancy between simulation and data. This accounts for deficits in the simulation such as inaccuracies in the description of the magnetic field, the modelling of the detector geometry, or the simulation of interactions within the detector material. Corrections are derived by fitting candidate $J/\psi \rightarrow \mu\mu$ and $Z \rightarrow \mu\mu$ events in simulation to the data, and applied to the momentum scale and resolution of muons in simulation.

A series of identification criteria [80] serve to distinguish prompt muons from muons produced in the semileptonic decay of light hadrons. The “Medium” working point is the default recommendation for most analyses, providing a balance between efficiency and purity for prompt muons. Within $|\eta| < 2.5$, this accepts only CB and IO muons that contain at least three hits in two or more MS stations. A loose selection is also placed on the compatibility between the MS and ID tracks, based the charge-to-momentum ratio of the tracks in each system:

$$q/p \text{ compatibility} = \frac{|q/p_{\text{ID}} - q/p_{\text{MS}}|}{\sqrt{\sigma^2(q/p_{\text{ID}}) + \sigma^2(q/p_{\text{MS}})}} < 7. \quad (4.13)$$

ME and SiF muons extend the acceptance to $|\eta| < 2.7$, with the requirement that they contain at least three hits in three MS stations.

4.4.4 Taus

Tau leptons can decay leptonically via $\tau \rightarrow \ell \nu_\ell \nu_\tau$, where $\ell = \{e, \mu\}$, or hadronically into a τ -neutrino, some number of neutral pions, and one (“1-prong”) or three (“3-prong”) charged pions. The leptonic decay channel is covered by the standard reconstruction algorithms for electrons and muons, given that neutrinos impart no signal within the detector. The hadronic decays will result in jets, with characteristic features that allow for them to be distinguished from quark or gluon-initiated jets.

Reconstruction of the visible component of hadronically-decaying taus ($\tau_{\text{had-vis}}$) employs anti- k_t jets with $R = 0.4$, constructed from topo-clusters that are calibrated to the hadronic

scale [82]. Because tau leptons travel a measurable distance in the detector before decaying, it is possible for the nominal primary vertex in an event (normally chosen as the one with the highest $\sum p_T^2$ of associated tracks) to not coincide with the tau decay vertex. Thus, for each “seed” jet, the primary vertex to which the highest p_T sum of tracks within $\Delta R < 0.2$ point is instead designated as the tau vertex (TV). Tracks with $p_T > 1$ GeV within longitudinal and transverse distances of the TV of $|d_0| < 1.0$ mm and $|\Delta z_0 \sin \theta| < 1.5$ mm are associated with the $\tau_{\text{had-vis}}$ candidate. This approach correctly reconstructs 1-prong and 3-prong decays as 1-track and 3-track candidates approximately 70% of the time.

To identify $\tau_{\text{had-vis}}$ jets from the large background of hadronic jets, boosted decision trees (BDTs) are trained separately on the 1-prong and 3-prong scenarios, with discriminating variables taking into account the topo-cluster structure, track distributions, and various track-cluster compatibility metrics [83]. Constant-efficiency working points are defined by making selections on the BDT score, with the “Loose” working point corresponding to a 60% (50%) efficiency for identifying 1-prong (3-prong) decays. Electrons can also form a background for the 1-prong signature, and in the scenario that a 1-prong $\tau_{\text{had-vis}}$ candidate overlaps geometrically with a reconstructed electron, the likelihood score (described in section 4.4.2) is used to make the decision on which particle to veto, with thresholds chosen such that a 95% efficiency for retaining Loose 1-prong $\tau_{\text{had-vis}}$ objects is achieved.

After this, a calibration is applied to correct the energy scale of reconstructed $\tau_{\text{had-vis}}$ candidates:

$$E_{\text{calib}} = \frac{E_{\text{reco}} - E_{\text{pile-up}}}{\mathcal{R}}. \quad (4.14)$$

This incorporates two corrections, the first serving to remove the expected contributions from pile-up ($E_{\text{pile-up}}$), and the second a response correction to account for mis-reconstruction biases (e.g. energy being deposited in dead material, or out-of-cluster effects). The response calibration functions, $\mathcal{R}$, are derived separately for 1-prong and multi-prong decays, and in bins of pile-up-corrected energy, $|\eta|$, and the number of primary vertices in the event.

4.4.5 Missing transverse momentum and significance

Not all of the particles produced in a collision event are detectable. Within the SM, the neutrinos constitute a prime example—their weakly-interacting nature ensures that they pass through the detector volume without leaving any hits in the detector material. However, their presence can still be inferred from the principle of momentum conservation: since the initial pp system contains no momentum in the transverse plane prior to a collision, it follows that the total momentum of particles produced in the interaction will also sum vectorially to zero in this plane. Thus, a deviation of this quantity from zero provides an indicator for the presence of invisible particle(s) in the event.

The calculation of the missing transverse momentum, $\mathbf{p}_T^{\text{miss}}$, involves a “hard” and a “soft” term [84]. The hard term takes into account a selection of calibrated objects (jets, taus, muons, electrons, and photons), while the soft term includes residual detector signals that were not used in the reconstruction of hard objects. In the standard approach, the latter consists of unused ID tracks⁶ that are compatible with the hard-scatter vertex of the event, and the $\mathbf{p}_T^{\text{miss}}$ is calculated

⁶There is an alternative construction that instead utilises residual calorimeter clusters for the calculation of the soft term; however, this approach, while being more inclusive by its inclusion of neutral particle deposits, tends to be inferior due to the difficulty in disentangling signal contributions from pile-up.

4.4. Physics objects

as follows:

$$\mathbf{p}_T^{\text{miss}} = - \left(\underbrace{\sum_e \mathbf{p}_T^e + \sum_\gamma \mathbf{p}_T^\gamma + \sum_\tau \mathbf{p}_T^\tau + \sum_\mu \mathbf{p}_T^\mu + \sum_{\text{jets}} \mathbf{p}_T^{\text{jet}}}_{\text{hard term}} + \underbrace{\sum_{\text{tracks}} \mathbf{p}_T^{\text{track}}}_{\text{soft term}} \right). \quad (4.15)$$

Because the calibration procedure differs for each type of object, the $\mathbf{p}_T^{\text{miss}}$ calculation has to be tailored specifically to each analysis, in order for the final result (and the propagation of systematic uncertainties through the calculation) to be consistent with the choice of object definitions of the analysis.

An important ambiguity resolution step is implemented to avoid any potential double counting of detector signals between the various terms. This is done independently of the standard analysis overlap removal procedure: electrons and muons selected by the analysis are used as-is, but the remaining terms are modified to only consider “accepted” objects, with photons taking priority, followed by taus, jets, then the soft term, in that order. Muons typically do not overlap with the other objects as their signals are not calorimeter-based; their inclusion in the overlap removal serves mostly to influence the acceptance of jets.

The magnitude of the $\mathbf{p}_T^{\text{miss}}$ vector, conventionally denoted E_T^{miss} or “MET”, is the canonical observable through which the production of weakly-interacting particles can be inferred. Another variable that offers similar discriminatory power is the object-based E_T^{miss} significance [85], which can be expressed as follows:

$$E_T^{\text{miss}} \text{Sig.}^2 = 2 \ln \left(\frac{\max_{\mathbf{p}_T^{\text{miss, true}} \neq 0} \mathcal{L} \left(\mathbf{p}_T^{\text{miss}} \mid \mathbf{p}_T^{\text{miss, true}} \right)}{\max_{\mathbf{p}_T^{\text{miss, true}} = 0} \mathcal{L} \left(\mathbf{p}_T^{\text{miss}} \mid \mathbf{p}_T^{\text{miss, true}} \right)} \right) = \frac{|E_T^{\text{miss}}|^2}{\sigma_L^2 (1 - \rho_{LT}^2)}. \quad (4.16)$$

While the standard E_T^{miss} variable provides an estimation of the amount of missing momentum in an event, the E_T^{miss} significance attempts to quantify the likelihood for the observed E_T^{miss} to have originated from real invisible particles, as opposed to being emergent from mis-reconstruction or resolution effects. The right-most expression is how the variable is computed in practice, with σ_L representing the longitudinal E_T^{miss} resolution, and ρ_{LT} being the correlation between the transverse and longitudinal resolutions.

4.4.6 Heavy flavour decays

The identification of jets produced from the decay of heavy flavour (b or c) hadrons is important for many areas of the ATLAS physics programme. The phenomenology of these jets differs from “light” flavour jets (defined simply as those not containing a b - or c -hadron) due to the several characteristics of heavy-flavour hadron decays, foremost of which is that the mean lifetime of a b -hadron lies at the order of a picosecond (and for c -hadrons, at approximately 0.1 ps). This results in a mean flight length of a few millimetres in the detector before these hadrons decay, which is resolvable as a displaced secondary vertex given the $O(10 \mu\text{m})$ precision of vertex reconstruction. Also, the fragmentation of heavy quarks tends to produce higher particle multiplicity showers, which can be exploited via jet substructure observables.

ATLAS employs five low-level b -taggers for identifying various features of b -hadron decays [86]. The IP2D, IP3D and RNNIP algorithms are designed to identify b -jets based on the impact parameter (IP) of their associated tracks. IP2D [87] constructs log-likelihood ratio dis-

criminants for each pair of flavour hypotheses from the signed transverse IP significance (defined as the ratio of the IP to its uncertainty):

$$D_{\text{IP2D}} = \sum_{i=1}^{N_{\text{track}}} \ln \left(\frac{p_A(S_i)}{p_B(S_i)} \right), \quad (4.17)$$

where $A, B \in \{b, c, \text{light}\}$ denote the possible jet flavours, and N_{track} is the number of tracks associated with the jet. The probability density functions are derived from reference histograms of the IP significances associated with tracks from each jet flavour in simulation. IP3D defines a similar discriminant, but additionally takes into account the signed longitudinal IP significance.

The RNNIP approach [88] utilises a recurrent neural network (RNN) to exploit correlations between the IP of tracks in b -hadron decays, which are difficult to incorporate in the more traditional IP2D and IP3D algorithms due to the sheer number of raw events necessary to adequately parametrise the probability templates with this information. The RNN learns the dependencies between tracks associated to each jet, taking into account both the signed longitudinal and transverse IP significances, and additionally also the ratio of track-to-jet p_T , the geometric distance from each track to the jet axis, and the number of hits each track imparts in the various ID layers. The network produces three outputs that correspond to the probability for the jet to be a b -jet, c -jet, or light flavour jet.

The fourth of the low-level taggers, SV1, attempts to reconstruct a secondary vertex (SV) within each jet [89]. Starting with 2-track vertices in the jet, up to 25 of the highest- p_T associated tracks are iteratively fitted to each candidate SV to define multi-track vertices, with track cleaning and quality criteria imposed to reduce the rate of reconstructing fake vertices. This approach offers the advantage of being more computationally efficient than attempting to directly construct multi-track vertices, while also allowing for 2-track candidates that are compatible with photon conversions, or with K_S^0 or Λ^0 decays, to immediately be excluded. In the end, six discriminatory variables are defined: the number of SVs found; the number of tracks matched to each SV; the invariant mass of the SVs; the fraction of energy between tracks associated to an SV to that of all tracks in the jet; and the transverse decay length of each SV and its corresponding significance.

Lastly, the JetFitter algorithm [90] adopts a topological approach to b -tagging by attempting to find a common line that represents the flight path of the b -hadron, and assuming that the tracks from any decay products will intersect with this path. The algorithm starts by searching for tracks that are incompatible with the primary hard-scatter vertex of the event, and inconsistent with hadronic material interactions, photon conversions, and K_S^0 or Λ^0 decays. The initial position of candidate SVs are then obtained by approximating the flight path of the b -hadron with the direction of the jet axis. Subsequent iterations of the vertex fit refine these positions by merging or removing candidate SVs depending on their compatibility with the expected topology of the decay chain. The algorithm yields a suite of topological variables that can be used to discriminate between b -, c -, and light flavour jets, such as the number of SVs reconstructed, the track multiplicity of each of the vertices, and the displacement length of the SVs.

To leverage the discriminatory power afforded separately by each of the low-level taggers, all of their outputs are combined and used as inputs to a fully-connected multi-layer feed-forward neural network. Basic jet kinematic properties are also included to account for their potential correlations with the other input variables. The result of this is the default high-level tagger used for identifying b -jets in ATLAS, known as the DL1r algorithm [86]. Its three outputs, p_b ,

4.5. Changes for Run 3

p_c and p_{light} , can be interpreted as the probability for the jet to be a b -jet, c -jet or light flavour jet, respectively, and a discriminant for b -tagging is constructed from their combination:

$$D_{\text{DL1r}} = \ln \left(\frac{p_b}{f_c p_c + (1 - f_c) p_{\text{light}}} \right). \quad (4.18)$$

The f_c parameter nominally represents the expected fraction of c -jets in the jet sample, although in practice its value is tuned in conjunction with a selection on D_{DL1r} to achieve a target b -tagging performance.

Due to the shorter lifetimes and lower masses of c -hadrons, the experimental signature of c -jets tends to be very similar to that of light flavour jets, making their identification challenging. However, the driving principles behind each of the b -taggers discussed above are still relevant for c -hadron decays, given that the basic topology of b - and c -hadron decay chains are similar (in that one expects the presence of displaced secondary vertices, just to a lesser degree for the latter case). Thus, while not optimised for this purpose, the multi-dimensional nature of the DL1r algorithm can be exploited to construct an alternative discriminant suitable for c -tagging:

$$D_{\text{DL1r}}^c = \ln \left(\frac{p_c}{f_b p_b + (1 - f_b) p_{\text{light}}} \right). \quad (4.19)$$

The parameter f_b now represents the fraction of b -jets, but is again treated as a tunable parameter for defining c -tagging working points.

While charm taggers tend to be much less efficient than b -taggers due to the greatly increased rate of misidentifying light flavour jets, these advancements have begun to allow for direct probes of final states containing charm quarks, a chief example being the work of Part III.

4.5 Changes for Run 3

As prefaced at the beginning of this chapter, the description of the detector hardware and various reconstruction techniques presented thus far pertain to Run 2 of the LHC. With the advent of Run 3, the collision centre-of-mass energy has been increased to $\sqrt{s} = 13.6$ TeV with an average of $\langle \mu \rangle = 60$ interactions per bunch crossing. These changes are accompanied by several upgrades to the detector hardware, installed as part of the “Phase-I” upgrades during Long Shutdown 2 (LS2) between 2019–2022.

The main focus of the upgrade was the installation of the New Small Wheels (NSWs), which replaced the inner wheels of the muon endcap system. These new detectors utilise Micromegas (Micro-Mesh Gaseous Structure) and small-strip thin gap chamber (sTGC) technologies to achieve a manageable L1 muon trigger rate in higher-background environments without necessitating increased muon p_T thresholds or trigger pre-scales. This constitutes an important upgrade that is beneficial not only for Run 3 operations, but especially also for the high-luminosity phase beyond. Another crucial upgrade was the improvement of the DAQ system to facilitate a physics rate to disk of approximately 3 kHz. This event rate is double that of Run 2, and given the increased size per event (at approximately 2.1 MB/event at $\langle \mu \rangle = 60$), the system has been upgraded to support a maximum throughput of 8 GB/s with at least 1.4 PB of storage available at the Tier-0 facility as a buffer. New electronic boards were also installed into the LAr calorimeter as part of an improved digital trigger system. A more comprehensive overview of the Run 3 detector configuration can be found in ref. [35].

4.6 Towards the High-Luminosity era

Once Run 3 concludes in 2026, the LHC will enter a period known as Long Shutdown 3 (LS3), during which work will commence to install the various upgrades for the High-Luminosity LHC (HL-LHC). In parallel, ATLAS will begin installation of its Phase-II upgrades. At the forefront of this is the complete replacement of the ID with an all-new fully silicon-based Inner Tracker (ITk) [91, 92], driven by the higher pile-up conditions at the HL-LHC, and necessary due to the radiation damage that the current ID will have been subjected to over the course of the first three data-taking runs. Notably, the ITk will provide tracking up to $|\eta| = 4$, which will improve the forward physics programme. The TDAQ system will undergo an overhaul to contend with the more demanding HL-LHC environment [93], where an average of $\langle\mu\rangle = 200$ interactions are expected to occur every bunch collision. The target is for events to be saved at a rate of 10 kHz with the improved DAQ system. In order to keep up with the new trigger specifications and for the components to remain robust against increased radiation levels, the readout electronics of the calorimeters [94, 95] and MS [96] will also be replaced. A new novel High-Granularity Timing Detector (HGTD) [97] will supplement the ITk by offering the ability to measure the time of charged particles with a resolution of approximately 30 ps.

4.7 Rise of the Machines

The field of machine learning (ML) has experienced a remarkable growth over the past decade, and ML techniques are beginning to outperform traditional approaches in multiple areas of high energy physics. ATLAS is no exception, with many facets of the data processing and event reconstruction chains having the potential to benefit from these techniques. Indeed, as presented in section 4.4.6, deep neural networks are already employed as the standard approach for identifying b -jets, with an alternative tagger based on graph neural networks (GNNs) having also been studied [98]. To highlight a few more cutting-edge applications: the classification of image-based representations of calorimeter energy deposits using convolutional neural networks (CNNs) [99], or more recently with graph and point clouds [100], has proven capable of distinguishing between the deposits from charged and neutral pions, with positive implications for jet energy calibration. Bayesian neural networks (BNNs) can be used to calibrate topo-clusters to the hadronic energy scale [101], with an improved resolution over the current calibration procedure. A neural network can be used for defining optimal E_T^{miss} working points [102]. ML techniques also offer lightweight alternatives to full simulations of the detector, which is typically done using GEANT4 [103] and can be a computationally intensive task, particularly in the simulation of particle showers within the calorimeter volume. Variational auto-encoders (VAEs) and generative adversarial networks (GANs) greatly reduce computational requirements whilst producing predictions that are similar to a full simulation [104]. The use of GANs for calorimeter simulation has already been implemented in the newest ATLAS fast-simulation tool, AtlFast3 (AF3) [105].

Particle Flow



ARTICLE FLOW (PFlow) refers to the holistic approach to reconstruction wherein observations from multiple detector systems are optimally combined to improve the measurement of a physics object. The technique is particularly relevant for hadronic jets, which have traditionally been reconstructed primarily using calorimetric methods. However, a large fraction of the typical jet energy (approximately half, on average) originates from charged hadrons, and the design of general-purpose detectors such as ATLAS [29] or CMS [30], which contain a tracking volume in addition to calorimeters, lends itself to the more sophisticated alternative in which the calorimeters are used to measure the neutral component of the jet, while the tracking system measures the momenta of its charged constituents.

PFlow is by no means a new concept, with its potential already having been recognised in the early days of collider physics at the Positron–Electron Tandem Ring Accelerator (PETRA) [106] and during the era of the Large Electron–Positron Collider (LEP) [107, 108]. The CMS collaboration has been using a PFlow reconstruction since the beginning of the LHC [109], necessitated by the limited energy resolution of their calorimeter; and while ATLAS used a robust calorimeter-only jet definition in Run 1 of the LHC, a PFlow approach was found to be superior on several metrics and consequently adopted as the default during Run 2 [110]. The particular implementation of PFlow used in the ATLAS experiment will be described in section 5.1, a few details of which will be relevant to some of the studies later shown in Part II of this thesis. The Track CaloCluster (TCC) and Unified Flow Object (UFO) approaches, which are motivated by some of the shortcomings of the standard PFlow, will be summarised in sections 5.2 and 5.3. Finally, a discussion of ongoing developments towards a global particle flow, along with prospects for future advancements, will be presented in sections 5.4 and 5.5, respectively.

5.1 Particle flow in ATLAS

The reconstruction of PFlow signals, known as Particle Flow Objects (PFOs), supplements the calorimeter-based topo-clusters with charged particle tracks from the Inner Detector (ID) [110]. Despite the increase in computational requirements of such an approach, the inclusion of charged particle tracks provides several benefits:

- The ID operates on the principle of the Lorentz force, and as such, provides a superior momentum resolution for low momentum charged particles that exhibit greater curvature within the tracker’s magnetic field.

5.1. Particle flow in ATLAS




Figure 5.1. A flowchart depicting the logic of the ATLAS PFlow algorithm. Solid (dotted) arrows represent the flow of tracks (clusters) through the algorithm. Image from ref. [110].

- Tracks with a p_T as low as 500 MeV can still be reconstructed; such soft particles often fall below the noise thresholds required for seeding a cluster within the calorimeters.
- The ID provides a better angular resolution for charged particles, as each reconstructed track can be traced back to the point of interaction.
- Low momentum charged particles, which have a greater curvature within the magnetic field in the ID, can be swept out of the jet cone upon reaching the calorimeter. By including tracks in the jet definition, these particles can be correctly accounted for.
- Any track not found to be associated with the primary hard-scatter vertex can be excluded from the jet, mitigating the contributions from charged pile-up particles.

At the same time, calorimeter energy measurements offer the advantage of being more performant in the higher energy regime; having an acceptance up to $|\eta| < 4.9$ (as opposed to $|\eta| < 2.5$ for the ID); and, of course, also being able to capture the signals of neutral particles. A demonstration of the complementary nature of the two detector systems can be found in section 8.C.

The optimal strategy for improving the jet energy resolution, then, utilises charged particle tracks at low energies whilst switching to a calorimeter-only definition at high energies. This forms the basis of the PFlow jet definition in ATLAS. The main challenge in this approach lies in avoiding the double-counting of signals that are imparted by a charged particle, since they can both form a track and also deposit energy within the calorimeter. The goal of the ATLAS PFlow algorithm is to avoid this issue by applying a correction to the energy of calorimeter clusters, subtracting off the energy that each track¹ is expected to have deposited. An overview of the main algorithm logic is shown in Figure 5.1. In essence, tracks and topo-clusters from standard reconstruction are used to define output PFOs consisting of tracks, unmodified clusters, and cell-subtracted clusters. The following sections will describe the algorithm in further detail.

5.1.1 Expected energy deposited by particles

The bulk of the charged particle content in a typical jet is known to come from charged pions. Thus, simulations of single charged pion events are used to study their isolated interactions within the detector volume. The reference calorimeter energy in each of these events, $E_{\text{ref}}^{\text{clus}}$, is calculated as the summed energy of all topo-clusters found within a cone of $\Delta R = 0.4$ around the calorimeter-extrapolated position of the charged pion track. The ratio of this reference energy with the momentum of the reconstructed track in each event, $E_{\text{ref}}^{\text{clus}} / p_{\text{ref}}^{\text{track}}$, provides a measure of the calorimeter response to charged pions.

¹For brevity, I will often speak of energy that “tracks” deposit, although it should of course be understood that it is the charged particle (from which the track was reconstructed) that actually deposits the energy within the calorimeter.

The mean values of the $E_{\text{ref}}^{\text{clus}}/p_{\text{ref}}^{\text{track}}$ distributions, $\langle E_{\text{ref}}^{\text{clus}}/p_{\text{ref}}^{\text{track}} \rangle$, as well as their standard deviations are derived in bins of $p_{\text{ref}}^{\text{track}}$, $|\eta_{\text{ref}}^{\text{track}}|$, and the layer of first interaction (LFI) of the shower within the calorimeter. These values form a vital ingredient of PFlow: given a track of momentum p^{track} , the amount of energy that it is expected to deposit into the calorimeter, and thus that should be subtracted from any matched clusters (section 5.1.5), is calculated as:

$$\langle E_{\text{dep}} \rangle = p^{\text{track}} \left\langle \frac{E_{\text{ref}}^{\text{clus}}}{p_{\text{ref}}^{\text{track}}} \right\rangle. \quad (5.1)$$

Values of the standard deviation, denoted $\sigma(E_{\text{dep}})$, are used in the selection of tracks (section 5.1.2), split shower recovery (section 5.1.4), and remnant removal (section 5.1.6).

5.1.2 Input objects

Standard EM-scale 4-2-0 topo-clusters and reconstructed ID tracks form the inputs to the PFlow algorithm. Tracks are required to have $0.5 \text{ GeV} < p_{\text{T}}^{\text{track}} < 100 \text{ GeV}$, fall within the ID acceptance region of $|\eta| < 2.5$, and satisfy the “Tight” quality criteria, corresponding to having at least nine hits in the silicon detector layers and including the presence of all pixel layer hits where expected from the track fit. Note that in this initial selection step, no requirements are demanded of the vertex of each track. The removal of pile-up-associated tracks is done only after the track has been passed through the subtraction algorithm; this allows for a more complete removal of charged pile-up contributions, with their expected calorimeter deposits also being accounted for. (This is described further in section 5.1.7.) Tracks with $p_{\text{T}}^{\text{track}} > 100 \text{ GeV}$ are not used, since the calorimeter provides a better energy resolution at these scales.

While all tracks satisfying the selections described above are used when building PFlow jets, the implementation of the cell-based subtraction strategy becomes ambiguous in scenarios where the charged particle is produced in close proximity to other particles. Thus, for the tracks that are deemed to be in a dense environment, the subtraction is not performed, and a scale factor is instead applied to the track energy to circumvent the double-counting (section 5.1.7 will elucidate this procedure). The criterion for being in a dense environment is based on the significance of the difference between the expected energy deposited by the track, $\langle E_{\text{dep}} \rangle$, and the actual energy, E^{clus} , that is measured in the calorimeter within a cone of $\Delta R = 0.15$ around the track [73]:

$$S(E^{\text{clus}}) = \frac{E^{\text{clus}} - \langle E_{\text{dep}} \rangle}{\sigma(E_{\text{dep}})} > 33.2 \times \log_{10}(40 \text{ GeV}/p_{\text{T}}^{\text{track}}). \quad (5.2)$$

The form of eq. (5.2) applies a gradual truncation of the subtraction algorithm as tracks approach the $p_{\text{T}}^{\text{track}} = 100 \text{ GeV}$ threshold.

5.1.3 Track–cluster matching

For selected tracks not within a dense environment, the energy that they are expected to deposit within the calorimeter is subtracted from the clusters that are reconstructed. The track–cluster matching step is important for ensuring that energy is removed from the cluster(s) that the track actually contributed to.

In the first pass, the algorithm attempts to match only a single cluster to each track. Only topo-clusters satisfying $E^{\text{clus}}/p^{\text{track}} > 0.1$, where E^{clus} is the cluster energy and p^{track} is the track momentum, are considered (below this threshold, the cluster is more likely to have orig-

5.1. Particle flow in ATLAS

inated from electronic or pile-up noise in the calorimeter). The closest cluster satisfying this requirement within a cone of $\Delta R' = 1.64$ around each track is taken to be the cluster that it deposited energy into; this threshold was derived using truth information in simulated events, and in principle depends on the typical shower width of a charged pion and the granularity of the calorimeter. The $\Delta R'$ metric is calculated from the track position (extrapolated to the second layer of the EM calorimeter) and the barycentre of each topo-cluster:

$$\Delta R' = \sqrt{\left(\frac{\Delta\varphi}{\sigma_\varphi}\right)^2 + \left(\frac{\Delta\eta}{\sigma_\eta}\right)^2}, \quad (5.3)$$

with $\Delta\varphi$ and $\Delta\eta$ being the separations in azimuthal angle and pseudorapidity, respectively, and σ_φ and σ_η the angular widths of the topo-cluster, included in this definition to account for the variation in granularity across different regions of the calorimeter. If no topo-clusters are found inside the cone of $\Delta R' = 1.64$, the track is assumed to have not formed a topo-cluster within the calorimeter, and no subtraction is performed for it.

5.1.4 Split shower recovery

It is possible for the energy shower of a particle in the calorimeter to be split and reconstructed as more than one cluster. To account for this effect, it may be necessary to associate additional clusters with each selected track. The significance of the difference between the matched track-cluster pair is used to control whether the split shower recovery procedure is performed:

$$S(E_{\text{matched}}^{\text{clus}}) = \frac{E_{\text{matched}}^{\text{clus}} - \langle E_{\text{dep}} \rangle}{\sigma(E_{\text{dep}})}. \quad (5.4)$$

Note that this differs slightly from the significance defined in eq. (5.2); here, $E_{\text{matched}}^{\text{clus}}$ represents the energy of only the cluster that was matched to the track. If $S(E^{\text{clus}}) < -1$, the list of clusters matched to the track is modified to include all of the topo-clusters within a cone of $\Delta R = 0.2$ around it, and the subtraction is performed using this list of matched topo-clusters.

5.1.5 Cell-based subtraction

Given the list of selected tracks and the set of topo-clusters matched to each, an energy subtraction step is then performed. In the simple case, if the topo-clusters matched to a track contain a total energy less than the expected energy deposited by the track, $\langle E_{\text{dep}} \rangle$, those clusters are simply removed.

If the matched clusters contain more energy than expected from the track, it is likely that they receive contributions from additional particles, which should be retained for the unified track-cluster description. In this scenario, a partial subtraction scheme is instead implemented: firstly, each calorimeter layer is divided into rings in η - φ space concentric about the position of the track in the layer of highest energy density (LHED). The energy density in the j^{th} cell of the i^{th} calorimeter layer is defined as:

$$\rho_{ij} = \frac{E_{ij}}{V_{ij}}, \quad (5.5)$$

where E_{ij} is the energy in the cell (in units of GeV) and V_{ij} is the cell volume (in units of cubic radiation lengths, X_0^3). The index i spans three layers in the EM calorimeter and one layer in the hadronic calorimeter (the Tile and HEC layers are treated as one). An average weighted energy

density in each layer is then calculated as:

$$\langle \rho' \rangle_i = \sum_j w_{ij} \rho_{ij}, \quad (5.6)$$

where w_{ij} is a weight corresponding to the integral over the cell area of a Gaussian distribution in η - ϕ space, with its mean at the extrapolated track position and a width equivalent to $\Delta R = 0.035$. The LHED is identified as the calorimeter layer for which

$$\Delta \rho'_i = \frac{\langle \rho' \rangle_i - \langle \rho' \rangle_{i-1}}{d_i - d_{i-1}} \quad (5.7)$$

is maximised, where d_i is the depth of the i^{th} layer (in units of X_0). The first layer is assigned the index $i = 1$, and $\langle \rho' \rangle_0$ and d_0 are both defined to be zero.

Rings are equally spaced within the same layer, with a radius chosen so that at least one calorimeter cell is contained within each. Ranking the rings in order of decreasing energy density, the subtraction proceeds as follows. Starting with the leading ring and moving successively down the ordered list, the calorimeter cells contained within each ring are removed, and the energy in those cells subtracted from $\langle E_{\text{dep}} \rangle$, until the algorithm finds a ring with an energy greater than the energy remaining in $\langle E_{\text{dep}} \rangle$. Then, the remaining $\langle E_{\text{dep}} \rangle$ is subtracted from this ring's energy, and the procedure halts.

5.1.6 Remnant removal

If the remnant energy in the set of post-subtraction topo-clusters is greater than $1.5\sigma(E_{\text{dep}})$, it is expected that additional particles in close vicinity to the track also contributed some energy to the clusters, and the topo-clusters are retained in the list of PFOs. Otherwise, the topo-cluster(s) is assumed to only contain deposits from the matched track, and that the remnant energy simply arises from a shower fluctuation. In this scenario, the residual topo-clusters are removed as well.

5.1.7 Post-subtraction refinements

A few additional corrections and selections are implemented before the collection of PFOs defined by the subtraction algorithm are used for jet building.

The energy of a charged particle inferred from the momentum of a reconstructed track lies close to the true energy of that particle, while calorimeter measurements are taken at the EM scale, which for hadrons will deviate from the true particle energy. To account for this difference, two scaling factors are applied to the collection of charged PFOs (i.e. selected tracks). The first applies only to tracks that were not used in the subtraction, which can happen if the track was found to be in a dense environment or not matched to any cluster. In this case, the assumption is that the track still deposited energy in the calorimeter; thus, to avoid double-counting, the track momentum is scaled by a factor of $(1 - \langle E_{\text{dep}} \rangle / p_{\text{track}})$. The second scaling is applied to all charged PFOs and serves to smoothen the transition between the particle-scale track and EM scale cluster energies: for $p_{\text{T}}^{\text{track}} > 60$ GeV a factor of $(1 - \langle E_{\text{dep}} \rangle / p_{\text{track}})$ is applied, and for $30 \text{ GeV} < p_{\text{T}}^{\text{track}} < 60$ GeV, with a factor that increases linearly to a value of unity at $p_{\text{T}}^{\text{track}} = 30$ GeV.

The second correction is motivated by the particle hypothesis of charged particles in a jet: because these are dominated largely by charged pions, all of the tracks used in PFlow are assumed to have been created from a charged pion. Thus, leveraging the independent approach to particle reconstruction, charged PFOs built from tracks that were also used for reconstructing candidate

5.2. Tracks for jet substructure

“Medium” electrons or muons are vetoed from being used as inputs to jet reconstruction. If these tracks were used in the subtraction algorithm, the subtracted energy is restored to the affected neutral PFOs; in cases where the neutral objects had been completely removed, new ones are created in their place.

Finally, the charged hadron subtraction (CHS) procedure removes charged PFOs that are inconsistent with the hard-scatter vertex of the event. This is done by requiring that the track lies within a distance of $|z_0 \sin \theta| < 2 \text{ mm}$ from the hard-scatter vertex, where z_0 is the longitudinal impact parameter and θ is the polar angle. This selection suppresses contributions from charged pile-up particles.

5.2 Tracks for jet substructure

The main focus of the standard application of PFlow in ATLAS is to improve the energy measurement of charged particles at low energies. However, high energy particle tracks can offer a different benefit: they maintain the good angular resolution of the tracker, which can be leveraged to obtain better measurements of the jet substructure. This has an advantage over the more classical calorimeter-based substructure techniques [60] particularly in the regime of boosted hadronic decays where, for example, an energetic vector boson can decay into a pair of highly collimated quarks that are reconstructed as a single large-radius jet.

The Track CaloCluster (TCC) jet definition [111] seeks to combine the energy resolution of the calorimeter with the angular resolution of the tracker. The first step of the algorithm involves a matching of topo-clusters to tracks, which is achieved under two conditions: firstly, a track will only be considered for the matching to a cluster if the uncertainty in its extrapolation into the calorimeter, σ_{track} , has a value less than the width of the cluster, which is calculated approximately as [60]

$$\sigma_{\text{clus}} \approx \arctan \left(\frac{\sqrt{\langle r^2 \rangle}}{|\vec{c}|} \right) \times \cosh(\eta), \quad (5.8)$$

where $\langle r^2 \rangle$ represents the mean of the squared radial distance from each cell in the cluster to the shower axis, and $\vec{c}$ is the centre of gravity of the cluster. Assuming that this condition is satisfied, the actual matching criterion is then based on the angular separation between the track and cluster:

$$\Delta R < \sqrt{\sigma_{\text{clus}}^2 + \sigma_{\text{track}}^2}. \quad (5.9)$$

For very high p_T tracks, the extrapolation uncertainty becomes very small, $\sigma_{\text{track}}^2 \ll \sigma_{\text{clus}}^2$, and the requirement approximately becomes $\Delta R \lesssim |\sigma_{\text{clus}}|$, i.e. that the extrapolated track position (loosely) falls within the extent of the cluster.

Three classes of TCC objects can be defined after the matching exercise: combined TCC objects, corresponding to matched tracks and clusters; and charged and neutral TCC objects, corresponding to unmatched tracks and clusters, respectively. Charged and neutral TCC objects simply inherit the 4-vector of the original object they were constructed from. For combined TCC objects, the simplest scenario is the one in which a single track is matched to a single cluster; in this case, the TCC 4-vector is constructed as the following combination of the cluster and track kinematics:

$$\text{TCC}_{1c,1t} = \left(p_T^{\text{clus}}, \eta^{\text{track}}, \varphi^{\text{track}}, m^{\text{clus}} = 0 \right). \quad (5.10)$$




Figure 5.2. A flowchart of the UFO reconstruction algorithm. Image from ref. [113].

The procedure becomes a little more involved when there are multiple track(s) matched to cluster(s). One TCC object is still created per track, with angular coordinates inherited from that of the track; however, the scale coordinates have to take into account all of the clusters that are matched. In general, each matched cluster will contribute a fraction of their p_T proportional to the total p_T of all the clusters that are matched, and further weighted based on the p_T fraction of the track, if multiple tracks have been matched.

The TCC objects constructed in this manner can be used as inputs to a clustering algorithm to define jets. There is some freedom available in choosing which objects to include for this, with each choice offering different benefits. Jets clustered using all TCC objects provide the best representation of the hadronic shower. However, clustering only the combined TCC objects, while inevitably missing part of the shower energy from the stand-alone charged or neutral TCC objects, suppresses contributions from pile-up and can result in jets that exhibit a well-behaved mass response even before the implementation of dedicated pile-up trimming procedures [112].

5.3 Unified Flow Objects

While TCC jets offer the advantage of improved classification of boosted hadronic decays, their performance in the lower p_T regime is generally inferior to standard PFlow reconstruction. This is mainly due to the less stringent track–cluster matching requirements of the TCC approach, which simply focuses on splitting the energy of a cluster between the tracks it was matched to. The complementarity of the kinematic range over which the PFlow and TCC approaches are each performant motivates the definition of Unified Flow Objects (UFOs) [113], which aim to combine the best features of the two algorithms to provide a description that performs well over the full kinematic range.

Figure 5.2 presents an overview of the UFO algorithm. In essence, the approach supplements the PFOs from standard PFlow reconstruction with all other reconstructed tracks in the event that are compatible with the hard-scatter vertex and did not satisfy the PFlow track selection criteria. The TCC algorithm is performed using these tracks, along with any charged PFOs that were not used in the cell-by-cell energy subtraction procedure; and, instead of the raw topo-clusters, the

5.4. Global particle flow

splitting is performed using the neutral PFOs that consist of both unmodified and cell-subtracted clusters. The UFO algorithm produces objects that benefit from the PFlow track–cluster matching and energy subtraction procedures when particles are in clean low-energy environments, while receiving the improved angular resolution of tracks at higher energies.

The multifaceted capabilities of UFOs make them suitable for the reconstruction of both small and large radius jets. UFOs have already been adopted for the default $R = 1.0$ jet definition, and are expected to also eventually supersede PFOs as the default object for defining $R = 0.4$ jets.

5.4 Global particle flow

Currently, the PFlow approach is only used in ATLAS for the reconstruction of jet constituents. However, the technique can of course be extended for the description of *all* physics objects. This paradigm of an event-wide PFlow treatment is referred to as *Global Particle Flow* (GPF).

The core idea of GPF is that every reconstructed signal in the detector—namely, clusters in the calorimeters and tracks in the ID and MS—originated from the interactions of a specific particle with the detector material, and should thus be identified accordingly during the reconstruction of the event. In contrast to the current strategy in ATLAS wherein each species of physics object is reconstructed independently of one another, with dedicated overlap removal procedures applied post-reconstruction to avoid the double-counting of signals, the GPF formulation seeks the simultaneous reconstruction of all physics objects given the full set of detector observables in each event.

The progression towards a GPF approach to event reconstruction remains one of the long term goals of the ATLAS experiment. The basic prerequisites for this have already been satisfied in Run 2, with the harmonisation of the electron, photon and jet reconstruction algorithms to use the same 4-2-0 topo-clusters as a basic calorimeter input (in Run 1, e/γ reconstruction had used a custom sliding window algorithm [78] for constructing calorimeter signals). A proposal for a staged implementation of GPF was formulated towards the end of Run 2 by a task force, formed with the goal of studying the feasibility of modifying the existing PFlow algorithm into a GPF algorithm. The first stage of this proposal identified the necessity for the reconstruction software to be able to determine whether two reconstructed physics objects were built from any common inputs. Over the period of LS2, this led to the development of so-called GPF links, which utilise the history of the fundamental inputs of each reconstructed object to draw a correspondence between the constituents of each object. This concept is illustrated in Figure 5.3.

GPF links have since been integrated into the ATLAS reconstruction software for Run 3 [114], and can be propagated through reconstruction to be exploited in the analysis step. One of the main benefits of these links lies in the potential simplification of overlap removal procedures, which have traditionally been based on the geometric ΔR separation between reconstructed objects and can become computationally cumbersome if the calculation has to be performed many times. The prospects of a GPF link-based overlap removal procedure will be the focus of the studies presented in Chapter 7.

5.5 Advancements and developments

The modular nature of the ATLAS PFlow algorithm offers the potential for its overall performance to be improved via the modification of any one of its components. In particular, the development of deep learning algorithms over the past decade, made possible by innovations in




Figure 5.3. Illustration of the initial implementation of Global Particle Flow in the ATLAS reconstruction software in the form of links between reconstructed physics objects ($e/\gamma/\mu/\tau$) and PFOs, represented by the red arrows.

modern computer hardware, allows for the possibility of an "advanced" PFlow wherein aspects of the existing algorithm are replaced by machine learning (ML) solutions. This could include, for example: ML-based alternatives for determining the amount of energy a track is expected to deposit into the calorimeter, which is currently calculated using a set of parametrised $\langle E/p \rangle$ values in bins of the track kinematics; refined track–cluster matching criteria, which currently rely on a simple selection based on the $\Delta R'$ metric; or modifications to the cell subtraction logic, which is currently driven by parametrised shower shape variables. Several of these possibilities are currently being studied internally within ATLAS.

In the case of $\langle E/p \rangle$ estimation, potential improvements might be possible even without the use of ML techniques, since the current workflow is derived from the behaviour of charged pions in an isolated environment (section 5.1.1). This becomes a poor assumption in higher-luminosity collisions (e.g. at the HL-LHC), and will be the topic of investigation in Chapter 8.

The applicability of ML towards classification and pattern recognition problems also makes it well-suited for exploring the GPF paradigm, especially as particles begin to overlap and merge into a single reconstructed signal in the detector. In these scenarios, simple distance metrics are no longer adequate for resolving the particles. However, the exercise can still be possible by exploiting differences in the expected shower of different particle species, which will manifest as differences in the overall distribution of energy across the cells of a reconstructed cluster. A first study to this effect was recently conducted in ref. [115], where a computer vision approach was used to analyse calorimeter images of overlapping charged and neutral pions (which constitute the most common charged and neutral objects produced in hadronic showers). The method was found to outperform traditional parametric methods in predicting the energy and direction of the initial particles. A follow-up study explored the prospects of an ML-based GPF using a recurrent hypergraph, with nodes corresponding to basic detector signals (tracks and clusters), and the hyperedges between nodes corresponding to the particles that contributed to that track or cluster [116]. The construction of the hypergraph in this way allowed for a direct relationship between particles and detector signals to be inferred from its incidence matrix, and despite not being optimised to do so, the approach was able to achieve superior jet momentum resolution over a parametrised PFlow.

ON ASPECTS OF HADRONIC FINAL STATES

Jet–Object Overlap Removal

THE reconstruction of hadronic jets and other physics objects—electrons, photons, muons, and taus—are optimised based on the signatures that they are respectively expected to impart within the detector material. This independent approach to particle reconstruction simplifies the reconstruction of the event as a whole; however, it comes with a drawback in that it is possible for any given signal in the detector to be used in the reconstruction of more than one particle, which would lead to an overestimation of the overall event activity.

To avoid this issue, the analysis of any event will typically feature a dedicated overlap removal step after all of the relevant objects have been reconstructed. There are two broad classes of overlap that can occur: the first concerns the *duplication* of objects, where the detector signals left by a particle are also mistakenly identified and reconstructed as another type of particle. An example of this is the electron, which forms a track in the Inner Detector (ID) and topo-clusters in the calorimeter—a signature that is often also well-clustered as a jet. The second type of overlap occurs when a collection of multiple real particles are produced close-by to each other in the detector, to the extent that their reconstruction cannot be disentangled. In principle, this is a problem that falls under the remit of particle *isolation*, but both of these scenarios are addressed by the overlap removal procedure, whose goal is ultimately to identify overlapping objects and consequently veto the extraneous particles (or, in some cases, the event entirely) in order to avoid the double-counting of detector signals.

Leading up to Run 2 of the LHC, a study group was formed within the ATLAS collaboration seeking to harmonise the overlap removal procedures used between different analysis teams. This effort led to a set of overlap recommendations to be formed, which were based primarily on the geometric distance between reconstructed objects:

$$\Delta R = \sqrt{\Delta\varphi^2 + \Delta\eta^2}, \quad (6.1)$$

where $\Delta\varphi$ and $\Delta\eta$ are the difference in azimuthal angle and pseudorapidity, respectively, of the pair of particles being considered. These recommendations were optimised using “EMTopo” jets, which are constructed from calorimeter-only inputs to the anti- k_t ($R = 0.4$) algorithm, with additional standard pile-up cleaning and quality cuts applied.

In general, however, the performance of any overlap removal procedure will depend on the definitions of the objects involved. With the transition to Particle Flow (PFlow) jet reconstruction in Run 2, it is thus important to revisit the recommendations and determine whether a re-optimisation is needed. This is precisely the goal of this chapter; the collection of validation studies presented here will cover jet overlaps with electrons, photons, and muons, and were con-

6.1. Simulated event samples

ducted as part of my authorship task for the ATLAS collaboration. Sections 6.1 and 6.2 will begin with a description of the event samples, the object definitions used, and the event and object selection criteria. A preliminary study of the objects and the extent of the overlap in the various event topologies will be presented in section 6.3. More detailed studies of jet overlaps with the different types of particles will then be presented in sections 6.4 to 6.7. In each case, the performance of PFlow jets will be compared to that of EMTopo jets as a test of whether any significant differences arise between the two jet definitions.

6.1 Simulated event samples

Monte Carlo (MC) samples of event topologies that are enriched with jets and electrons, photons, and/or muons are employed to facilitate the study of overlaps between the different types of particles. For jet–electron and jet–muon studies, these were events containing a Z boson decaying either to a pair of electrons or muons, along with jets from associated production. The same overlaps are studied also using semi- and di-leptonic $t\bar{t}$ events in which one or both of the top quarks decays as $t \rightarrow bW \rightarrow b\ell\nu$. These events contain an abundance of b -jets, allowing for the specific study of jet–lepton overlaps in the context of heavy flavour hadron decays. Lastly, the production of a prompt photon in association with jets provides a clean event for studying overlaps between jets and photons. Additional details on the simulation of these event samples are provided below.

The POWHEG BOX v1 MC generator [117–120] was used for the simulation at next-to-leading order (NLO) accuracy of the hard-scattering processes of Z boson production and decay in the electron and muon channels. It was interfaced to PYTHIA 8.186 [121] for the modelling of the parton shower, hadronisation, and underlying event, with parameters set according to the AZNLO tune [122]. The CT10NLO parton distribution function (PDF) set [123] was used for the hard-scattering processes, whereas the CTEQ6L1 PDF set [124] was used for the parton shower. The effect of quantum electrodynamical (QED) final-state radiation was simulated with PHOTOS++ 3.52 [125, 126]. The EVTGEN 1.2.0 program [127] was used to decay bottom and charm hadrons.

The production of $t\bar{t}$ events was modelled using the POWHEG BOX v2 [117–119, 128] generator, which provided matrix elements at NLO in the strong coupling constant α_s , and the NNPDF3.0NLO [129] PDF set. The h_{damp} parameter, which controls the matching in POWHEG and effectively regulates the high- p_T radiation against which the $t\bar{t}$ system recoils, was set to $1.5 m_{\text{top}}$ [130]. The functional form of the renormalisation and factorisation scales was set to the default scale $\sqrt{m_{\text{top}}^2 + p_T^2}$. The events were interfaced with PYTHIA 8.230 [131] for the parton shower and hadronisation, using the A14 set of tuned parameters [132] and the NNPDF2.3LO set of PDFs [133]. The decays of bottom and charm hadrons were simulated using the EVTGEN 1.6.0 program [127].

Prompt single-photon production was simulated with the SHERPA 2.2 [134] generator. In this setup, NLO-accurate matrix elements for up to two partons, and LO-accurate matrix elements for up to four partons were calculated with the Comix [135] and OPENLOOPS [136–138] libraries. They were matched with the SHERPA parton shower [139] using the MEPS@NLO prescription [140–143] with a dynamic merging cut [144] of 20 GeV. Photons were required to be isolated according to a smooth-cone isolation criterion [145]. Samples were generated using the NNPDF3.0NNLO PDF set [129], along with the dedicated set of tuned parton-shower parameters developed by the SHERPA authors.

Table 6.1. Truth object matching requirements for an event to be selected. These thresholds are imposed on the prompt leptons in Z + jets and $t\bar{t}$ events, and the prompt photon in γ + jets events.

	ΔR	$\Delta p_T^{\text{ratio}}$
$Z \rightarrow ee$	< 0.1	< 0.1
$Z \rightarrow \mu\mu$	< 0.05	< 0.1
$t\bar{t}$	< 0.1	< 0.1
γ + jets	< 0.1	< 0.2

6.2 Object definitions and event selection

Reconstructed jets correspond to anti- k_t $R = 0.4$ PFlow and EMTopo jets, with each collection calibrated according to their respective set of recommendations derived using the Run 2 dataset [73]. They are required to satisfy the Tight jet vertex tagger (JVT) selection [146] to reject contributions from pile-up, and the DL1 algorithm [147] was used to identify jets that contain a b -hadron with a 77% efficiency. Truth jets are defined using the same clustering configuration, with inputs consisting of all stable truth particles in the event except muons, neutrinos, leptons from W , Z , Higgs or tau decay (and any photons within $\Delta R < 0.1$ of these), and prompt photons from Higgs decays. Reconstructed electrons and photons are required to satisfy their respective Tight identification criteria [78, 79], while muons are required to be of Medium quality [80].

Selected Z + jets events are required to contain at least two reconstructed leptons of the corresponding flavour (i.e. electrons or muons), the p_T -leading pair of which should be matched to the truth leptons produced in the Z boson decay. This matching is performed by comparing the ΔR (eq. (6.1)) and $\Delta p_T^{\text{ratio}}$ between the truth and reconstructed particles, where the latter is defined as

$$\Delta p_T^{\text{ratio}} = \frac{|p_T^{\text{truth}} - p_T^{\text{reco}}|}{p_T^{\text{truth}} + p_T^{\text{reco}}}. \quad (6.2)$$

Similarly, in $t\bar{t}$ events the truth lepton(s) produced in the W decay is required to be well-matched to a reconstructed lepton of the corresponding flavour, and the leading reconstructed photon in γ + jets events is required to be matched to the truth-level prompt photon. The particular matching thresholds chosen for these variables in each of the event topologies is detailed in Table 6.1. These were chosen based on the distributions of the ΔR and $\Delta p_T^{\text{ratio}}$ variables, which are included in section 6.A together with comparisons of the truth- and reconstructed-level object kinematics after imposing the matching criteria.

6.3 Preliminary analysis

As a starting point, let us take a look at PFlow jets in Z + jets and γ + jets events. Figure 6.1a shows ΔR distributions (normalised to unity) when considering all possible jet-object combinations in each event. (In Z + jets events, objects correspond to the two p_T -leading leptons, while in γ + jets only the p_T -leading photon is considered.) The distributions share similar features, with a large excess at $\Delta R \approx 0$ that falls off to a minimum at $\Delta R \approx 0.35$. Given that the jet cone has a radius of $R = 0.4$, this indicates that the jet axis tends to be biased towards nearby objects, rather than

6.3. Preliminary analysis




Figure 6.1. The ΔR distribution between PFlow jets and electrons, muons or photons in $Z + \text{jets}$ and $\gamma + \text{jets}$ events, for (a) all jet–object combinations, and (b) only the closest jet to each object. The final bin of each histogram includes overflow entries.

simply encapsulating them near the boundary of the cone. The extent of this will, of course, be dependent on the relative hardness between the true jet and the object (it is a feature of the anti- k_t clustering algorithm), and hence on the particular event topology. In $Z + \text{jets}$ and $\gamma + \text{jets}$ events, the jets from associated production are typically on a softer scale than the prompt leptons/photon, so the bias is more pronounced.

The distributions rise again and quickly stabilise beyond the $\Delta R = 0.4$ threshold, although the entries in this region are largely an artifact of combinatorics, given the large jet multiplicity in each event. The alternative viewpoint in Figure 6.1b is perhaps more informative in showing the real extent of overlap: here, only the closest jet in ΔR to each object of interest is included, and one can see a clearer interpretation of how often the objects seed a fake jet.

The similarity between the jet–electron and jet–photon ΔR distributions suggests that these objects fake a jet equally well, likely due to their similar shower properties in the calorimeter. For the case of muons, one observes a flatter distribution, since they usually do not interact within the calorimeters and thus are less likely to fake a jet. However, there is still an excess around $\Delta R \approx 0$, which is a consequence of a known issue in the Run 2 reconstruction framework wherein subleading muon tracks are not vetoed from the list of inputs to the PFlow algorithm. This leads to an increased rate of muon-faked jets than would normally be expected (given that electron and muon tracks are supposed to be excluded from the inputs to jet reconstruction). This issue will be discussed further in section 6.6.1.

We can study the impact of the jet scale on these overlaps by restricting our considerations to only the n^{th} p_T -leading jet in each event as a proxy. The jet–object ΔR distributions for the $n = 1, 2, 3$ and 4 cases are shown in Figure 6.2. For $n \geq 3$ the distributions begin to look the same, since there are at most two prompt objects in the topologies considered, which are typically of a harder scale than the associated jets in the event. The majority of overlaps thus involve the p_T -leading and subleading jets, and there is an interesting observation that one can make by comparing, in particular, the jet–electron and jet–photon distributions in Figures 6.2a and 6.2b to the $n = 3$ and $n = 4$ cases in Figures 6.2c and 6.2d. In the first two distributions, there is a point of inflection in the lineshapes at $\Delta R \approx 0.15$ as one moves from the peak at $\Delta R = 0$ towards the




Figure 6.2. The ΔR distribution between the n^{th} -leading PFlow jet ($n = 1, 2, 3$ or 4) and electrons, muons or photons in $Z + \text{jets}$ and $\gamma + \text{jets}$ events. The final bin of each histogram includes overflow entries.

$\Delta R = 0.4$ threshold. This behaviour is related to the possible bias in reconstruction as discussed above, or more accurately, the absence of it: because these leading jets are (by definition) more energetic than the average jet in the event, they are less likely to be biased by the presence of additional objects in the periphery of the jet core. The inflection in these distributions can thus be attributed to the two different overlap scenarios: duplication at $\Delta R \approx 0$, and non-isolated particles between $0.15 \lesssim \Delta R \lesssim 0.35$.

6.4 Jet–photon overlaps in $\gamma + \text{jets}$

Analyses that select photons typically only consider those that are already isolated within a cone of $\Delta R < 0.4$. No such requirement has been made of the photons considered for this overlap study, of course; but because of this, the baseline recommendation for jet–photon overlap removal is just a simple one of removing any remaining jet—if any—within $\Delta R(j, \gamma) < 0.4$ of cone-isolated photons. Figure 6.3 shows the minimum ΔR between the prompt photon and PFlow jets in $\gamma + \text{jets}$ events against the ratio of their energies. The vast majority of photons produce a fake jet, as indicated by the peak at $\Delta R \approx 0$ and $E(j)/E(\gamma) \approx 1$. There is a small tail that extends to

6.5. Jet–electron overlaps in $Z(\rightarrow ee) + \text{jets}$




Figure 6.3. Minimum ΔR between prompt photons and PFlow jets in $\gamma + \text{jets}$ events against the ratio of their energy.

Table 6.2. Jet–photon overlap removal performance for PFlow and EMTopo jets in $\gamma + \text{jets}$ events.

	$\gamma + \text{jets}$	
	PFlow	EMTopo
N_{jets}	308 845 237	483 897 214
$N_{\text{jets rejected}}$	33 326 113	33 056 805

higher values of the energy ratio, but these entries remain very close to $\Delta R \approx 0$, suggesting that the jet reconstruction has probably been biased by the presence of the photon. Thus, the removal of all the jets within $\Delta R(j, \gamma) < 0.4$ makes sense. A comparison of the overlap removal performance for PFlow and EMTopo jets is presented in Table 6.2. While the two jet definitions admit a significantly different total number of jets (for reasons I shall not discuss here), a similar number is found to be removed in both cases. This allows for two conclusions to be drawn: firstly, that the PFlow and EMTopo definitions are equally susceptible to being faked by a photon; and secondly, that overlap removal performs similarly for both collections.

6.5 Jet–electron overlaps in $Z(\rightarrow ee) + \text{jets}$

The current jet–electron overlap removal recommendation is a two-step procedure. Firstly, any jet within $\Delta R < 0.2$ of an electron is rejected, as these are likely to have been seeded by the electron. The second step concerns jets and electrons separated by $0.2 < \Delta R < 0.4$, which typically corresponds to a real hadronic jet and electron in close proximity. For these cases, the reconstruction of the electron tends to have been biased by the nearby hadronic activity, so the electron is usually rejected in order to avoid complications related to the additional systematic uncertainties that this introduces. However, this choice can be analysis-dependent, with some choosing to remove the jet instead; in particular, for precision cross-section measurements of




Figure 6.4. Minimum ΔR between prompt electrons and PFlow jets in $Z(\rightarrow ee) + \text{jets}$ events against the ratio of their energy. Subfig. (a) considers all jets in the event, while (b) only includes truth-matched jets.

Table 6.3. Jet–electron overlap removal performance for PFlow and EMTopo jets in $Z(\rightarrow ee) + \text{jets}$ events.

	$Z(\rightarrow ee) + \text{jets}$	
	PFlow	EMTopo
N_{jets}	1 003 323 857	1 667 904 382
$N_{\text{jets}}^{\text{rejected}}$	281 736 128	283 590 199
$f_j^{\text{rej.}} (\text{truth-matched})$	4.16%	3.93%
$f_j^{\text{truth-matched}} (\text{rejected})$	6.29%	6.05%
$f_e^{\text{rej.}}$	0.10%	0.13%

high jet multiplicity regions, where the overall contribution of lepton identification systematics is small and the removal of a jet simply reclassifies the event as a lower multiplicity one. That said, what we will study in this section is the nominal choice of removing the electron in this intermediate region of overlap.

Distributions of the minimum ΔR between electrons and PFlow jets in $Z(\rightarrow ee) + \text{jets}$ events against the ratio of their energy are shown in Figure 6.4a. As with the jet–photon case, one finds a clear region of electron duplication peaking at $\Delta R \approx 0$ and $E(j)/E(e) \approx 1$, along with a heavily-depleted trough at $\Delta R \approx 0.3$ that separates these cases from the non-overlapping objects in $\Delta R > 0.4$. However, the duplication region also contains jets with real hadronic origin; this is illustrated in Figure 6.4b, where the minimum ΔR is calculated only considering the jets that can be matched to a truth jet via a $\Delta R(\text{truth}, \text{reco}) < 0.3$ requirement. While the core of the peak has vanished (as expected), the tail that extends to higher $E(j)/E(e)$ values largely remains. Thus, the removal of jets in the region $\Delta R(j, e) < 0.2$ will result in a small inefficiency of rejecting these hadronic jets.

The outcome of applying this overlap removal procedure to both PFlow and EMTopo jets in $Z(\rightarrow ee) + \text{jets}$ events is summarised in Table 6.3. Here, $f_j^{\text{rej.}} (\text{truth-matched})$ denotes the fraction

6.6. Jet–muon overlaps in $Z(\rightarrow \mu\mu) + \text{jets}$

of the rejected jets (i.e. those in $\Delta R(j, e) < 0.2$) that are truth-matched; $f_j^{\text{truth-matched}}$ (rejected) denotes the fraction of truth-matched jets that are rejected, which can be interpreted as an inefficiency of the overlap removal procedure for real hadronic jets; and f_e^{rej} denotes the fraction of electrons that would be rejected due to being in the $0.2 < \Delta R(j, e) < 0.4$ region. A slightly higher rate of rejecting hadronic jets (6.29%) is found when using the PFlow definition (compared to 6.05% for EMTopo), although this is balanced by fewer electrons being removed (0.10% vs 0.13%). However, the differences between the two are very small overall, and a reoptimisation of the overlap removal thresholds when using the PFlow jet definition was deemed unnecessary.

6.6 Jet–muon overlaps in $Z(\rightarrow \mu\mu) + \text{jets}$

Jet–muon overlap removal aims to distinguish prompt muons from muons that are produced in the decay of heavy flavour jets. Additionally, final state radiation (FSR) and the bremsstrahlung radiation of photons can result in a muon faking the signature of a jet (track accompanied by calorimeter deposits from the FSR or bremsstrahlung photon).

The general Run 2 recommendation is to apply overlap removal in the region of $\Delta R < 0.4$ between a jet and cone-isolated muon. The jet should be removed if its number of associated tracks, N_{trk} , is less than three (a typical characteristic of jets clustered as a result of FSR/bremsstrahlung), unless it is b -tagged (indicating that the muon was produced in the semi-leptonic decay of the b -hadron), in which case the muon is typically removed instead. In this section, we will study the performance of this procedure in $Z(\rightarrow \mu\mu) + \text{jets}$ events. The b -tagging condition will be neglected here, since these events do not contain many real b -jets; the impact of heavy flavour considerations will be examined in section 6.7 using a more appropriate event topology.

As with the previous sections, the ultimate goal will be to verify whether the recommendations perform similarly for both the PFlow and EMTopo jet definitions. However, before we do so, there is a caveat that has to be mentioned: analyses that select both muons and PFlow jets have to apply an additional overlap removal step, which is not necessary if using the EMTopo jet collection. Let us first discuss the reason for this, and the consequences it has on the jet–muon overlap removal performance for PFlow jets.

6.6.1 Subleading muon tracks in Particle Flow

In the Run 2 reconstruction software, there is an unintended feature in which subleading muon tracks are not masked from the PFlow subtraction algorithm. This leads to an increase in the rate of muons being mis-reconstructed as a jet. To mitigate this issue, any analysis that considers both muons and PFlow jets has to apply an additional algorithm at the start of the overlap removal sequence that removes jets based on their properties, e.g. the number of tracks they are associated to, the relative p_T of these tracks to the p_T of the muon, and the proximity of the jet to a muon.

The number of jets that overlap with a muon is significantly reduced after applying the preliminary overlap removal procedure (Figure 6.5a), and the peak in the leading jet p_T distribution due to these muon-faked jets is effectively removed (Figure 6.5b). Taking into account only the closest jet to each muon, Figures 6.5c and 6.5d shows the $\Delta R(j, \mu)$ against the ratio of their energy, $E(j)/E(\mu)$, before and after applying the algorithm. The core of the distribution at $\Delta R(j, \mu) \approx 0$ and $E(j)/E(\mu) \approx 1$, which corresponds to the muon track–faked jets, is effectively depleted.

The remaining “bad” jets after applying this overlap removal (i.e. those that are not matched to any truth jet) originate predominantly from pile-up contributions. While there will of course




Figure 6.5. Effect of applying the preliminary jet–muon overlap removal procedure in $Z(\rightarrow \mu\mu) + \text{jets}$ events. (a) The ΔR between jets and muons. (b) Leading jet p_T distribution. (c, d) The ΔR between jets and muons against the ratio of their energies, before (c) and after (d) the preliminary overlap removal.

be an inefficiency of this algorithm in retaining real hadronic jets, this had been determined, prior to the studies conducted in this chapter, to be minimal when compared to the efficiency loss of other cleaning algorithms (e.g. JVT), and thus generally at an acceptable level for most analyses especially when factoring in the improved efficiency for retaining prompt muons.

6.6.2 Overlap removal performance comparison

The $\Delta R(j, \mu)$ vs $E(j)/E(\mu)$ distributions in Figure 6.5 are markedly different from those we have previously seen for jet–photon (Figure 6.3) and jet–electron (Figure 6.4) overlaps for two reasons: firstly because muons are minimum ionising particles, and as such are not fully absorbed (if at all) in the calorimeter material; and secondly, because their tracks are nominally vetoed from PFlow reconstruction. Of course, muon tracks are sometimes unintentionally included, as we have just discussed in the previous section; thus, there are four possible scenarios:

- The objects do not overlap ($\Delta R(j, \mu) > 0.4$). In this region, the distribution falls smoothly along the $E(j)/E(\mu)$ axis.

6.6. Jet–muon overlaps in $Z(\rightarrow \mu\mu) + \text{jets}$




Figure 6.6. Minimum ΔR between prompt muons and truth-matched PFlow jets in $Z(\rightarrow \mu\mu) + \text{jets}$ events against the ratio of their energy.

- The objects overlap ($\Delta R(j, \mu) < 0.4$), but the muon did not interact within the calorimeter volume. In this case, the jet and muon are actually independent, and do not share any common constituents. Thus, one observes the same smoothly-falling distribution along $E(j)/E(\mu)$, and it follows that these jets are in fact clustered from real hadronic activity that just happens to be in close proximity to the muon.
- The objects overlap ($\Delta R(j, \mu) < 0.4$), and the muon deposited energy within the calorimeter, whether by ionising the material or by producing an FSR/bremsstrahlung photon. This will induce a positive bias in the energy scale, and possibly the axis, of any jet that overlaps with these clusters.
- The muon track is not vetoed from PFlow, and seeds a fake jet. Naturally, this will overlap nearly identically with the muon itself ($\Delta R(j, \mu) \approx 0$ and $E(j)/E(\mu) \approx 1$), although could also incorporate varying amounts of nearby hadronic energy, leading to a tail that extends to higher values of $E(j)/E(\mu)$.

We have already seen, in Figure 6.5d, that the preliminary overlap removal procedure effectively rejects the cases corresponding to this final point. However, in addition to the bulk of the fake jets around $\Delta R(j, \mu) \approx 0$ and $E(j)/E(\mu) \approx 1$, the band of jets near $E(j)/E(\mu) \approx 0$ is also affected. This region corresponds to real hadronic jets, as discussed in the second and third points above and illustrated in Figure 6.6, where the distribution is shown for truth-matched jets. Thus, the removal of this population of jets will contribute to the overall inefficiency of the overlap removal in retaining true hadronic jets.

A summary of the overall overlap removal performance for PFlow and EMTopo jets is presented in Table 6.4. As before, f_j^{rej} (truth-matched) denotes the fraction of the rejected jets that are truth-matched, $f_j^{\text{truth-matched}}$ (rejected) is an inefficiency of the overlap removal procedure for preserving real hadronic jets, and f_μ^{rej} is the fraction of muons that are rejected. Unlike the previous jet–photon and jet–electron results, we find that PFlow overlap removal rejects a higher percentage of true jets in the event (8.52%) compared to that of EMTopo jets (1.77%). The upside

Table 6.4. Jet–muon overlap removal performance for PFlow and EMTopo jets in $Z(\rightarrow \mu\mu) + \text{jets}$ events.

	$Z(\rightarrow \mu\mu) + \text{jets}$	
	PFlow	EMTopo
N_{jets}	843 615 464	1 469 710 515
$N_{\text{jets}} \text{ rej. (prelim.)}$	183 021 868	–
$N_{\text{jets}} \text{ rej. (total)}$	183 187 465	75 272 579
$f_j^{\text{rej.}} \text{ (truth-matched)}$	9.17%	4.58%
$f_j^{\text{truth-matched}} \text{ (rejected)}$	8.52%	1.77%
$f_\mu^{\text{rej.}}$	0.37%	6.46%

Table 6.5. Jet–lepton overlap removal performance for PFlow and EMTopo jets in $t\bar{t}$ events.

	$t\bar{t}_{[1e/2e]}$		$t\bar{t}_{[1\mu/2\mu]}$	
	PFlow	EMTopo	PFlow	EMTopo
N_{jets}	290 284 103	404 568 026	298 973 482	426 317 950
$N_{b\text{-jets}}$	41 016 460	39 277 752	47 408 656	45 503 385
$N_{\text{jets}} \text{ rej. (prelim.)}$	14 301 281	–	18 283 893	–
$N_{b\text{-jets}} \text{ rej. (prelim.)}$	1 973 534	–	2 088 982	–
$N_{\text{jets}} \text{ rej. (total)}$	38 602 186	28 933 533	18 316 991	7 396 288
$f_j^{\text{rej.}} \text{ (truth-matched)}$	25.03%	8.79%	48.08%	5.09%
$f_j^{\text{truth-matched}} \text{ (rejected)}$	5.73%	1.54%	4.52%	0.20%
$f_{e/\mu}^{\text{rej.}}$	0.85%	0.67%	4.26%	14.02%

is that a far greater muon efficiency is achieved (0.37% vs 6.46%). However, the overwhelming majority of the PFlow jets removed were a consequence of applying the preliminary overlap removal algorithm, whose performance has already been studied and deemed to be at an acceptable level. Thus, in the end, no further modification to the standard jet–muon overlap removal procedure was implemented for the PFlow case.

6.7 Jet–lepton overlaps in $t\bar{t}$ events

Jets originating from the shower of b -quarks form an important aspect of many analyses, and thus there is interest in studying how the overlap removal procedures affect or is affected by such jets. For this, semi- and di-leptonic $t\bar{t}$ events that contain a rich content of b -jets and leptons (electrons or muons) are utilised. A summary of the overlap removal performance is presented in Table 6.5, where $t\bar{t}_{[1e/2e]}$ and $t\bar{t}_{[1\mu/2\mu]}$ indicate that events have been selected to contain at least one electron or muon, respectively. The preliminary jet–muon overlap removal algorithm has been applied in both cases, and its impact on the b -jet efficiency is included.

6.8. Summary

In the case of jets overlapping with muons, the overall conclusions are consistent with those found in $Z(\rightarrow \mu\mu) + \text{jets}$ events. For PFlow jets, almost all of the jets that have been rejected are due to the preliminary overlap removal algorithm. Approximately a tenth of these comprise of b -tagged jets, although this is a relatively small fraction ($\sim 4\%$) of the total b -jet population. The algorithm does reject many true jets (around half of the jets failing its conditions can be truth matched), but in the end this amounts to an efficiency loss of only approximately 4% for both hadronic jets and muons. In contrast, the procedure for EMTopo jets results in a very high jet efficiency but a loss of approximately 14% of the muons, largely owing to b -tagged jets taking precedence over muons when they are found to overlap.

The addition of the preliminary jet–muon check also impacts the final performance of the jet–electron overlap removal procedure, with a much greater number of PFlow jets now being removed at an inefficiency of approximately 5% for retaining true jets compared to 1.5% in the EMTopo case. This comparison is obviously skewed by the fact that a different set of procedures are being applied to the two jet collections, but it further highlights the considerations that any analysis looking to select PFlow jets and muons might need to take. Beyond this, the electron efficiency is found to remain consistent between the two scenarios.

6.8 Summary

In this chapter, the Run 2 recommendations for performing overlap removal between jets and other physics objects (e, μ, γ) have been studied, and their performance compared when PFlow jets are used instead of the EMTopo definition for which they were optimised for.

When evaluating the jet–photon and jet–electron overlap removal procedures, a similar performance was found between both the PFlow and EMTopo jet definitions. However, the jet–muon case differs due to a known issue in PFlow reconstruction, where muon tracks are sometimes not masked from the subtraction algorithm and hence can seed a fake PFlow jet. This has since been amended in the Run 3 reconstruction framework, but to mitigate its effects in Run 2 analyses, an additional jet–muon overlap removal algorithm was developed with the intention for this to be implemented at the start of the overlap removal sequence if muons are selected in the analysis.

There are two consequences of this preliminary step on the studies that have been presented in this chapter. Firstly, the nominal jet–muon overlap removal procedure becomes largely unimportant when selecting PFlow jets, as most of the overlaps between these objects will already be removed in the preliminary check. Secondly, because the algorithm heavily favours the retention of the overlapping muon in order to more efficiently reject the muon track–faked PFlow jets, the overall hadronic jet efficiency is decreased, while the muon efficiency is increased, when compared to performing the nominal overlap removal with EMTopo jets. However, this performance was already validated in the past and deemed to be acceptable for most analyses, and given the similar performance seen for PFlow and EMTopo jets when overlapping with other types of objects, no re-optimisation of the existing overlap removal procedures is considered to be necessary.

6.A Additional plots




Figure 6.7. (a, b) ΔR and Δp_T^{ratio} between the truth and reconstructed electrons in $Z(\rightarrow ee) + jets$ events. (c, d) Comparison between the truth and reconstructed leading (c) and subleading (d) electron p_T distributions after requiring that they match within $\Delta R < 0.1$ and $\Delta p_T^{ratio} < 0.1$.

6.A. Additional plots




Figure 6.8. (a, b) ΔR and $\Delta p_T^{\text{ratio}}$ between the truth and reconstructed muons in $Z(\rightarrow \mu\mu) + \text{jets}$ events. (c, d) Comparison between the truth and reconstructed leading (c) and subleading (d) muon p_T distributions after requiring that they match within $\Delta R < 0.05$ and $\Delta p_T^{\text{ratio}} < 0.1$.




Figure 6.9. (a, b) ΔR and $\Delta p_T^{\text{ratio}}$ between the truth and reconstructed photon in γ +jets events. (c) Comparison between the truth and reconstructed photon p_T distributions after requiring that they match within $\Delta R < 0.1$ and $\Delta p_T^{\text{ratio}} < 0.2$.

6.A. Additional plots




Figure 6.10. (a, b) ΔR and $\Delta p_T^{\text{ratio}}$ between the truth and reconstructed electrons and muons in $t\bar{t}$ events. (c, d) Comparison between the truth and reconstructed electron (c) and muon (d) p_T distributions after requiring that they match within $\Delta R < 0.1$ and $\Delta p_T^{\text{ratio}} < 0.1$.

Link-based Overlap Removal

 HE introduction of Global Particle Flow (GPF) links (section 5.4) in the ATLAS reconstruction framework for Run 3, coupled with harmonisation of the basic inputs used for the reconstruction of different objects, means that it is now possible to explicitly check whether any constituents of reconstructed objects originate from the same detector signal. This lends itself to alternative overlap removal procedures that have the potential to simplify the traditional ΔR -based methods used in Run 1 and Run 2.

There are two broad approaches to utilising GPF links for overlap removal. The first approach is a bottom-up one, with the links being taken into account at the reconstruction step: for example, one might first construct the electrons, photons and muons, and apply calibrations and working point selections using the standard reconstruction and identification algorithms for these objects. Then, jet reconstruction would proceed with the caveat that any tracks or topo-clusters already used for reconstructing the other objects are excluded from the Particle Flow (PFlow) algorithm (section 5.1). In this way, the jets that are defined will be mutually exclusive of the other physics objects from the outset, and a dedicated overlap removal step becomes unnecessary. Unfortunately, while this is an elegant approach, an actual implementation is not so straightforward due to the jet definition being dependent on the definitions of the other physics objects. This makes it difficult to provide a standardised calibration or generic characterisations of common quality metrics (e.g. the jet energy resolution or hard-scatter efficiency), and also complicates the evaluation of jet-related systematic uncertainties.

The alternative approach is to preserve the existing workflow and simply replace the ΔR -based overlap removal logic with one based on GPF links. This is the approach that will be explored in this chapter. I will focus primarily on overlaps between jets and electrons, as this is the most common ambiguity in reconstruction, although a brief discussion of jet–muon overlaps (particularly in the context of muons from heavy flavour decay) will also be included in section 7.4. To begin, section 7.1 will present a showcase of GPF links and some of the useful information one can glean from them. Prospective link-based jet–electron overlap removal strategies will then be defined and compared to the existing ΔR -based method in section 7.3.

7.1 First look at GPF links

The infrastructure of GPF links centres around Particle Flow Objects (PFOs), which form the inputs to jet clustering. These can be charged (seeded from tracks) or neutral (from topo-clusters), and in the reconstruction software, are stored in two vectors referred to as the *PFO containers*.

7.1. First look at GPF links




Figure 7.1. (a) The ΔR separation between jets and electrons when charged PFOs are or are not shared. (b) The ΔR separation between jets and electrons when neutral PFOs are or are not shared.

Information on the jet constituents is retained such that one is able to fetch the list of PFOs that each jet was constructed from. With the advent of GPF links, each object (electrons, photons, muons or taus) is also decorated with vectors that point to elements of the PFO containers, based on whether the object’s constituents coincide with those used to define that PFO. Drawing a correspondence between jets and objects then amounts to a simple check of whether they link to any common PFOs.¹

As a starting point, let us consider jets and electrons in semi- or di-leptonic $t\bar{t}$ events that contain at least one prompt electron from W boson decay. Reconstructed electrons are required to satisfy the Tight identification criteria and also be matched to the truth-level prompt electrons in the event. We can check the extent of separation between the objects when they share any charged or neutral PFOs: this is shown in Figures 7.1a and 7.1b, respectively. In the former case, since charged PFOs linked to Medium leptons are vetoed from jet clustering, one finds that there are very few instances of overlap, as expected. The few overlaps that do occur comes about due to the different electron working point used here, although in these cases, the jet and electron almost always shares a neutral constituent as well.

The majority of jet–electron overlaps instead occur between the neutral constituents. The sharp excess of entries at $\Delta R \approx 0$ in the right-hand bin of Figure 7.1b corresponds to cases of electrons being duplicated as jets. The overlaps continue up to a separation of $\Delta R = 0.4$ (corresponding to the jet radius), becoming cases of non-isolated hadronic activity close-by to electrons. However, the gradient of this fall-off is relatively smooth, and the value of ΔR at which these two scenarios begin to transition is not particularly evident from this visualisation.

To better separate these two overlap regimes, we can take a closer look at the jets and electrons that share constituents. Let us start with the canonical plot of the ΔR between the objects against the ratio of their energies. This is shown for jet–electron pairs that share charged PFOs in Figure 7.2a, while Figure 7.2b shows the same for pairs that share neutral PFOs. Both distributions are constrained to low values of ΔR , with tails extending to larger values of the ratio of their energies. Note that when electron tracks are included as a jet constituent, the core of

¹Technically, there is a slight difference in the implementation: objects are linked to the PFOs defined by the particle flow algorithm, but a few additional corrections are applied to these before they are used in jet clustering (see section 5.1.7). Nevertheless, what all this amounts to is just an extra layer of links between the standard and modified PFO containers to obtain the correct correspondence between each object and jet.




Figure 7.2. (a) The ΔR separation between jets and electrons against the ratio of their energies when charged PFOs are shared. (b) The ΔR separation between jets and electrons against the ratio of their energies when neutral PFOs are shared. (c) The ΔR separation between jets and electrons against the fraction of neutral PFO energy shared by the electron. (d) The fraction of neutral PFO energy shared by jets against that shared by electrons.

the distribution is shifted to $E(j)/E(e) \approx 2$ because the electron’s energy is effectively counted twice: once from the track, and again from its calorimeter deposits. This occurs at a relatively low rate, as can be gauged from Figure 7.2b where these cases are also visible. Nevertheless, they complicate the interpretation of these variables: one might initially think that a greater ratio of jet to electron energy provides an indicator for the regime of overlap, but this is now obfuscated by the type of constituent that they share.

An advantage of the GPF approach is that since we know which constituents are common between objects, their properties can be used to open up a host of additional variables that can be leveraged in the overlap removal decision-making. For example, focusing on the more common case of neutral constituent overlaps, Figure 7.2c shows the fraction of energy that is shared by the electron, which is computed as the ratio of the energy of the shared neutral PFO(s) to the total energy of the reconstructed electron itself. Indeed, we can see that this seems to offer better separation than the jet–electron ΔR (plotted on the y-axis), although some ambiguity remains: an electron that shares all of its neutral constituents with a jet could correspond to an electron-duplicated jet, *but* can also be consistent with an electron that might have been

7.2. Identifying true hadronic jets




Figure 7.3. Fraction of jet energy arising from electrons (calculated using calibration hits) for jets that share constituent(s) with a reconstructed Tight electron (*signal*) and jets that do not (*control*).

produced within a hadronic shower (e.g. the semi-leptonic decay of a heavy flavour hadron). These cases can be disentangled by also considering the fraction of energy shared by the jet, which is computed similarly and shown along the y-axis of Figure 7.2d. Electrons that share most of their energy can overlap with jets that share varying fractions of their energy, with the fake jets reconstructed from electrons identifiable as those where nearly all of the neutral constituents are shared between the two objects. In practice, a selection on the fraction of energy shared by the jet is sufficient to distinguish these cases due to the correlation between the variables, and the decision to keep or reject a jet when it only shares part of its energy with an electron can be made according to the needs of the particular analysis.

7.2 Identifying true hadronic jets

A quantitative approach to determining the regime of overlap that a jet–electron pair belongs to is beneficial for evaluating the performance of any new overlap removal procedure. To this end, two methods will be employed to identify jets that arise from real hadronic activity. The first approach is the standard one of attempting to associate each reconstructed jet to a truth-level jet via a ΔR matching requirement. Truth jets are defined using the anti- k_t $R = 0.4$ clustering algorithm with inputs corresponding to all generator-stable truth particles in the event with the exception of: muons (which are minimum ionising particles); weakly-interacting particles (i.e. neutrinos); prompt leptons (and any radiated photons within a cone of $R = 0.1$ around these leptons) from W , Z , Higgs and tau decays; and prompt photons from Higgs decays. To ensure a clean ΔR match, only reconstructed jets that are isolated from other jets with $p_T > 7$ GeV within a cone of $\Delta R < 0.6$ are considered. A threshold of $\Delta R < 0.3$ between the truth and reconstructed jet is used as the matching requirement.

The second approach utilises truth information from the detector simulation. In particular, the interactions of each particle with the individual calorimeter cells—referred to as the *calibration hits*—can be recorded during the data processing workflow. This allows for the origin of each energy deposit in the calorimeter to be known; and, by extension, the origin of any calorimeter contributions to each reconstructed jet. In the context of fake jets from electrons, one can check




Figure 7.4. Performance of GPF link-based jet-electron overlap removal procedures compared to the traditional ΔR -based one. (a) Comparison of jets that pass or fail the overlap removal procedure against the fraction of their calorimeter energy attributed to electrons. (b) Correlation between the GPF 3 and ΔR procedures. In each cell, two values are printed: the number of isolated reconstructed jets (black), and the number of isolated and truth-matched jets (red).

how much of the jet’s energy is attributed to electron sources: a true hadronic jet is expected to contain very little contribution from electrons, while fake jets will clearly be dominated by them. This is illustrated in Figure 7.3, where the fraction of each jet’s calibration hit energy coming from electrons is plotted for two categories of jets: the “signal” jets are those that share at least one constituent (either charged or neutral) with an electron, while the “control” group consists of truth-matched jets that do not overlap with any electrons.

7.3 Prospective GPF-based procedure

In this section, I will present several prospective GPF link-based overlap removal procedures and evaluate their performance alongside the current ΔR -based method. Building upon the discussion of section 7.1, the procedures that will be compared are as follows:

- **ΔR :** the traditional approach wherein the jet is removed if $\Delta R(j, e) < 0.2$, while the electron is removed if $0.2 < \Delta R(j, e) < 0.4$.
- **GPF 1:** any jet that shares a constituent(s) with a reconstructed electron is removed.
- **GPF 2:** if shared constituents are found between a jet and electron: the jet is removed if it shares more than 40% of its neutral or all of its charged energy with the electron; otherwise, the electron is removed.
- **GPF 3:** if shared constituents are found between a jet and electron: the jet is removed if it shares more than 70% of its neutral or all of its charged energy with the electron; otherwise, the electron is removed.

The results are summarised in Figure 7.4a, where the jets that pass or fail each procedure are plotted against the fraction of their electron calibration hit energy. The GPF 1 approach (drawn in azure) has been included mostly for completeness: it constitutes the most extreme case in which the reconstructed electron always takes priority over any overlapping jets. This is often

7.4. Jet–muon overlaps

not the most optimal solution, since electron candidates can arise from the semi-leptonic decay of heavy flavour jets or be faked from the signature of light flavour jets. In addition, there can be a third conceptual scenario of non-isolation between a real jet and electron, in which case the reconstruction of the position or energy of one or both objects can be biased. Indeed, it can be seen that in addition to the fake jets that contain only or mostly calorimeter contributions from electrons, the GPF 1 method also removes many real hadronic jets with very little electron content.

The GPF 2 (green) and GPF 3 (magenta) procedures address these concerns by incorporating a selection on the fraction of neutral constituent energy shared by the jet. As discussed, this variable is a promising one for discriminating between the various regimes of overlap: the GPF 2 method uses an intermediate cut ($> 40\%$) on the shared energy fraction and is able to achieve a performance very similar to the ΔR -based approach, while the tighter cut ($> 70\%$) in the GPF 3 procedure allows for more jets of hadronic origin to be retained. This is illustrated in Figure 7.4b, which shows the number of jets in four quadrants corresponding to whether they pass or fail the GPF 3 and ΔR -based procedures. For this result, only isolated jets (as described in section 7.2) are considered, and the number of truth-matched jets in each quadrant is also printed in red. The bulk of the jets are placed into the same category (i.e. pass or fail) by both of the procedures, with the vast majority of passing jets being truth-matched and oppositely for the rejected jets. The off-diagonal contributions are subdominant, but the GPF 3 procedure is able to retain many truth-matched jets that the ΔR -based method would otherwise remove.

From these results, we can infer that a link-based approach to overlap removal can be viable, whilst also having the potential to slightly improve upon the hadronic jet efficiency. Of course, this latter benefit will obviously come with a trade-off against the efficiency for retaining prompt electrons; nevertheless, the use of GPF links is capable of, at the very least, performing overlap removal at a similar level as the traditional ΔR -based approach, with opportunities to tune the specific algorithm if one of the overlapping objects is to be favoured over the other.

7.4 Jet–muon overlaps

While this chapter focuses primarily on the study of jet–electron overlaps, let me also provide a brief discussion here on the potential of GPF approaches for handling overlaps between jets and muons. Unlike electrons, isolated muons do not commonly fake a jet since they deposit very little energy within the calorimeter. However, a muon can occasionally radiate a photon through bremsstrahlung or final state radiation, and the photon’s calorimeter deposits along with the muon’s track can be mistakenly interpreted as a hadronic jet. Of course, barring considerations of muon identification working points, muon tracks are vetoed from being used in jet reconstruction, so the jet that is clustered in this scenario will only consist of the neutral calorimeter deposits. These jets can thus be identified, and consequently rejected, by either checking whether all of their neutral constituents are common with the muon, or via a standard ΔR isolation check—both approaches yield a similar performance.

The second goal of a jet–muon overlap removal procedure is to resolve cases where a muon is produced in the decay of a heavy flavour (b or c) quark. Flavour considerations are core to many physics analyses, and the decision to keep the jet or the muon typically varies from analysis to analysis depending on whether they treat these muons as a signal. In the traditional approach, if a selected muon falls within a cone of $\Delta R = 0.4$ of a jet, the overlap removal decision is made based on the characteristics of the jet: the muon is considered as non-prompt if the jet is b -




Figure 7.5. Fraction of neutral PFO energy shared between overlapping jets and muons in $t\bar{t}$ events. Subfig. (a) considers all jets, while subfig. (b) considers only b -tagged jets.

tagged or is associated with 3+ tracks, otherwise the jet is treated as a fake. This relies both on the performance of the b -tagger in identifying b -jets, as well as the fact that the particle shower in a heavy flavour jet typically involves many charged particles in addition to the muon.

GPF links can be used to augment this procedure by allowing also for an explicit check for shared neutral constituents between the jet and muon, which will typically be the case if they lie within $\Delta R < 0.4$ of one another. The advantage of this is that the dependence on the particular tagging algorithm used is then reduced. The characteristics of jet–muon overlaps in $t\bar{t}$ events is shown in Figure 7.5, firstly for the inclusive case where all jets are considered (Figure 7.5a), then when only b -tagged jets using the 77% working point of the DL1r algorithm [86] are considered (Figure 7.5b). Due to the minimally-ionising nature of muons, their neutral constituents typically contain only a small fraction of their total energy; however, muon-faked jets can be seen to be well-separated from true heavy-flavour jets, with the latter tending to share a smaller fraction of their energy with the overlapping muon. Thus, as with the jet–electron case, the fraction of energy contained within the neutral constituents shared by the jet looks to be a performant variable for distinguishing between the different jet–muon overlap scenarios.

7.5 Outlook

The studies presented in this chapter constitute only a very preliminary foray into the potential uses of the GPF links available within the Run 3 ATLAS analysis software. In particular, the removal of reconstruction overlaps between different particle species is a promising area where GPF could offer improvements; even by leveraging only a few simple checks for shared constituents between objects, we have seen that a performance comparable to that of the traditional ΔR -based methods can be achieved. In addition, a GPF-based approach offers other benefits worth mentioning: it greatly simplifies the configuration of overlap removal code, and also avoids the potentially large combinatorics when calculating the ΔR between many objects. Indeed, for these reasons, the E_T^{miss} calculation in Run 3 now implements a GPF approach in its overlap removal decision tool, as this constitutes one of its major computational bottlenecks.

Still, on the analysis front, the prospect of a GPF-based overlap removal procedure that is also improved from the physics perspective is desirable, as the similarity in its current performance

7.5. Outlook

with respect to the established approaches does not warrant a paradigm shift in the overlap removal code. In the end, such improvements might be achievable only by utilising more sophisticated techniques at a more fundamental level; for example, by training machine learning algorithms to better identify and exploit correlations in the phase space of particles during event reconstruction. Studies in this direction are already underway within the collaboration, following from the work that has been presented here. In this regard, an implementation that leverages the utility that GPF links provide is no doubt imminent as ATLAS begins to transition into the era of the High-Luminosity LHC, where denser collision environments will necessitate an improved understanding and characterisation of reconstruction overlaps.

E/p Inside Jets



AN integral part of adopting a Particle Flow (PFlow) approach to object reconstruction is that one has to take care to avoid double-counting contributions from particles that interact with multiple detector systems. Pertinent to the Particle Flow of ATLAS is the case of electrically charged particles, which produce two reconstructible signals in the detector: a series of hits within the Inner Detector, forming a track from which the particle's momentum can be calculated; and depositions within the calorimeters, from which its total energy can be measured. The measurement of particle properties using both sets of detector observables, then, requires that calorimeter measurements are corrected to remove any contribution thought to originate from the same particle as a reconstructed and selected track.

The current approach to determining the correct amount of energy to subtract in such a scenario is by studying simulations of charged pions within the detector. Events are generated to contain a single $\pi^\pm$ so that the presence of any clusters in the calorimeter can immediately be attributed to the charged pion, providing a straightforward method for characterising the hadronic response of the calorimeter. Values of the charged pion $\langle E/p \rangle$ (“*e-over-p*”), where E is the energy deposited in the calorimeter and p the track momentum, are obtained in bins of the track momentum, pseudorapidity, and layer of first interaction (LFI) within the calorimeter, and used as reference during the subtraction steps in PFlow.

However, the isolated conditions currently used to derive the $\langle E/p \rangle$ values are not representative of the typical environment inside a jet, where one expects a collection of many localised particles produced in the hadronic shower. This provides the premise for the studies presented in this chapter: the amount of clustered energy deposited by low-momentum particles is expected to be greater in the realistic scenario because of the stricter requirements for topo-cluster seeding compared to that of growth in the standard 4-2-0 clustering approach. In the case of an isolated particle, an energy deposition that does not meet the threshold for seeding a topo-cluster results in zero calorimeter energy being measured, while the presence of additional particles in close proximity may allow for its inclusion if adjacent cells meet the cell significance requirements as a consequence.

As a first course of action, I will establish the relationship between the amount of energy deposited by charged particles in a jet, and the ΔR distance between those particles and the axis of the jet. The expectation is that larger E/p values will be observed for tracks that fall more closely to the core of the jet where the density of particles is greatest. The ΔR with respect to the jet axis acts as a proxy for the local energy density around the track of interest. The effect of accounting for this dependence on the reconstructed jet will then be studied by constructing two

8.1. Dependence of E/p on $\Delta R(\text{jet}, \text{track})$




Figure 8.1. The E/p of tracks as a function of the ΔR between the track and the jet it is matched to. The ΔR is calculated using track coordinates extrapolated to the calorimeter, and red lines denote the mean E/p value in each ΔR bin.

$\langle E/p \rangle$ maps: one parametrised in bins of track momentum and pseudorapidity, and the second additionally as a function of $\Delta R(\text{jet}, \text{track})$. These two maps will be used in a simple Particle Flow implementation based on the CMS jet-plus-tracks algorithm [148, 149] to study the impact on the reconstructed jet energy response and resolution. The same will then be studied within the standard ATLAS Particle Flow algorithm, comparing a 2-dimensional case (in which the $\langle E/p \rangle$ is parametrised only in bins of the track energy and pseudorapidity) to one in which a binning in ΔR is also included.

The studies are performed using samples of Monte Carlo (MC) simulated dijet events with matrix elements calculated to leading order (LO) in quantum chromodynamics (QCD), and corresponding to a centre-of-mass energy of $\sqrt{s} = 13.6$ TeV. The events were simulated with the leading truth jet required to have a p_T between 0 GeV and 400 GeV. No pile-up overlay was included, and calibration hit information was retained during the simulation step such that the amount of energy deposited by generator-stable truth particles into each calorimeter cell can be accessed; this will facilitate the parametrisation of the expected energy deposited by each charged particle as a function of its track's properties.

8.1 Dependence of E/p on $\Delta R(\text{jet}, \text{track})$

The relationship between the energy deposited by charged particles and their proximity to a jet is studied by considering all jets with $p_T > 20$ GeV in an event. Tracks are matched to the jet if $\Delta R(\text{jet}, \text{track}) < 0.4$, where the ΔR is defined as

$$\Delta R = \sqrt{\Delta\phi^2 + \Delta\eta^2}, \quad (8.1)$$

with the calorimeter-extrapolated track coordinate used in the calculation. For each matched track found to correspond to a charged pion (with the association being done using generator-level information), the E/p of the track is computed. E represents the amount of energy that the charged pion deposited within the calorimeter volume, and p is the momentum of the reconstructed track. A measure for the former quantity is obtained using the total amount of visible calibration hit energy deposited by the charged pion inside any clustered calorimeter cells. The relationship between this quantity and the ΔR of the track to the jet is visualised in Figure 8.1,

and in select bins of track momentum and $|\eta|$ in Figure 8.2. Distributions in additional kinematic bins are included in section 8.A (Figures 8.14 to 8.17).

Tracks are found to deposit, on average, a greater fraction of their energy the closer they are to the jet axis. The dependence is strongest for tracks of intermediate momentum (around 10 GeV), and also for higher momentum tracks (in the tens of GeV) in the endcap regions ($|\eta| \gtrsim 2$), with variations in the mean E/p of up to approximately 60% between tracks that were produced in the core and at the boundary of a jet. At very low momenta ($\lesssim 2$ GeV) the effect is less pronounced due to these particles being very soft and often depositing very little to no clustered energy within the calorimeter, especially in the more forward pseudorapidity regions where the cell granularity (and thus noise level) is greater. On the other end of the spectrum, depositions from high- p_T charged particles ($\gtrsim 25$ GeV) produced away from the jet core are often able to themselves seed a cluster, which suppresses the dependence on ΔR .

8.2 Comparisons using a simple Particle Flow

The impact of different $\langle E/p \rangle$ configurations on the final jet energy is studied using a simple implementation of particle flow inspired by the CMS jet-plus-tracks algorithm [148, 149]. The inputs to the algorithm consist of uncalibrated EMTopo jets, defined by using topological clusters in the calorimeter as input to the anti- k_t algorithm [65] with a radius parameter of $R = 0.4$; and all reconstructed tracks in the event¹ that satisfy the TightPrimary requirements [150].

The energy of each EMTopo jet is corrected depending on the set of tracks found within its proximity. Three correction steps are applied, and in what follows, $\text{track}_{\text{vertex}}$ denotes the track's position at the vertex while $\text{track}_{\text{extrap.}}$ refers to its coordinates after being extrapolated into the calorimeter volume. The dominant correction comes from considerations of in-cone tracks: for any track within $\Delta R(\text{jet}, \text{track}_{\text{extrap.}}) < 0.4$ of the jet, the expected energy deposited by the track (i.e. the $\langle E/p \rangle \times p_{\text{track}}$) is subtracted, and the momentum of the track is added to the jet energy. Then, edge effects are accounted for: if a track was produced within the jet cone ($\Delta R(\text{jet}, \text{track}_{\text{vertex}}) < 0.4$) but falls out of acceptance after reaching the calorimeter ($\Delta R(\text{jet}, \text{track}_{\text{extrap.}}) > 0.4$) the momentum of the track is added back to the jet energy; while for the opposite scenario where a particle produced independently of the jet happens to curve into the jet cone in the calorimeter (i.e. $\Delta R(\text{jet}, \text{track}_{\text{vertex}}) > 0.4$ and $\Delta R(\text{jet}, \text{track}_{\text{extrap.}}) < 0.4$) the expected energy deposited by the track is subtracted.

The algorithm can be summarised as follows:

$$\begin{aligned}
 E_{\text{jet}}^{\text{corrected}} = E_{\text{jet}}^{\text{baseline}} &+ \sum_{\text{in-in}} (p_{\text{track}} - \langle E/p \rangle \times p_{\text{track}}) \\
 &+ \sum_{\text{in-out}} p_{\text{track}} \\
 &- \sum_{\text{out-in}} \langle E/p \rangle \times p_{\text{track}} .
 \end{aligned} \tag{8.2}$$

The comparisons made in this chapter, then, amount essentially to applying these corrections using different sets of $\langle E/p \rangle$ reference values and seeing how the jet energy response and resolution is affected as a result. While this implementation of Particle Flow has been adopted for these studies due to its simplicity, it is worth noting that it represents the limiting case of the ATLAS

¹No additional z_0 cut was imposed since the studies are performed using MC samples containing no pile-up.

8.2. Comparisons using a simple Particle Flow




Figure 8.2. The E/p of tracks as a function of the ΔR between the track and the jet it is matched to, for tracks with momenta between 2–5 GeV (a, c, e, g) or 5–15 GeV (b, d, f, h), with the plots in each group corresponding to different track $|\eta|$ bins. The ΔR is calculated using track coordinates extrapolated to the calorimeter, and red lines denote the mean E/p value in each ΔR bin.




Figure 8.3. The E/p distribution of π^- tracks in the central barrel region of the detector.

PFlow in environments where particles are fully overlapping, i.e. when track–cluster matching becomes challenging and the nominal cell-based subtraction method cannot be applied. Thus, while not the focus of this chapter, it is interesting to compare this approach with the standard ATLAS PFlow reconstruction in different regions of phase space. These studies are presented in section 8.B.

8.2.1 Deriving $\langle E/p \rangle$ reference values

Four sets of $\langle E/p \rangle$ values were derived using charged pion tracks in simulated dijet events:

1. in bins of track momentum and $|\eta|$
2. in bins of track momentum, $|\eta|$, and $\Delta R(\text{jet}, \text{track})^2$
3. as in item 1, but including tracks with $E/p \sim 0$
4. as in item 2, but including tracks with $E/p \sim 0$

Item 1 is used to mimic the standard ATLAS PFlow configuration, with the difference of being agnostic to the LFI for simplicity. Items 3 and 4 are motivated by the current approach of deriving $\langle E/p \rangle$ reference values in ATLAS where tracks that are matched to no calorimeter deposits (i.e. with an E/p exactly equal to zero) are vetoed. In single charged pion simulations, this thus excludes soft pions that consequently form no clusters in the calorimeter. The fraction of such cases is not negligible, as shown in Figure 8.3 for simulated π^- in the central barrel of the calorimeter. In dijet events where the particle density is much greater, these $E/p = 0$ cases can attain a small non-zero value due to the presence of additional particles in close proximity, allowing for the minuscule energy deposits to be clustered. In practice, the difference between the value being exactly or approximately zero is inconsequential; however, it implies that the current methodology will tend to overestimate the average amount of energy deposited by charged pion tracks, and result in a greater amount of energy being subtracted from the calorimeter measurement in the PFlow algorithm. The inclusion of items 3 and 4 serves to quantify the impact of this choice.

Track momentum and $|\eta|$ bins are defined identically to what is currently used in standard PFlow. Explicitly, the binning used for each of the variables is as follows:

²The ΔR is calculated using track coordinates extrapolated into the calorimeter, unless indicated otherwise.

8.2. Comparisons using a simple Particle Flow

- track momentum bin edges: (0, 2, 5, 15, 25, 40, 100) [GeV]
- track $|\eta|$ bin edges: (0, 0.5, 1, 1.5, 2.5)
- $\Delta R(\text{jet}, \text{track})$ bin edges: (0, 0.1, 0.2, 0.3, 0.4)

Only tracks within $\Delta R < 0.4$ of a jet with $p_T > 20$ GeV are considered (even in the cases where no binning in $\Delta R(\text{jet}, \text{track})$ is performed). The mean of the 1-dimensional E/p histograms in each bin is taken as the $\langle E/p \rangle$ value to be used in the PFlow algorithm. A selection of these histograms, and how they shift as one moves to different $\Delta R(\text{jet}, \text{track})$ bins, is shown in Figure 8.4 with a veto on tracks below $E/p < 0.01$ as per the configuration described in item 2. In general, larger E/p values are seen in the more central $|\eta|$ regions and for tracks closer to the jet axis, with more Gaussian-like distributions for more energetic tracks. The $\langle E/p \rangle$ values themselves are summarised in Figure 8.5.

8.2.2 Comparisons of $\langle E/p \rangle$ configurations

After applying the algorithm of eq. (8.2) using the four different E/p configurations, the resulting jets are compared using the mean energy response and resolution as metrics. This is done in bins of jet p_T and $|\eta|$ (using the initial EMTopo jet kinematics to define the binning). Only jets with $p_T > 20$ GeV, satisfying an isolation criterion, and furthermore matched to an isolated truth jet are considered. Reconstructed jets are isolated if no other reconstructed jet with $p_T > 7$ GeV lies within a cone of $\Delta R < 0.6$. Candidate truth jets are required to have $p_T > 10$ GeV and to be isolated from other truth jets with $p_T > 7$ GeV within $\Delta R < 1.0$. The matching of reconstructed to truth jets is performed via a $\Delta R(\text{reco}, \text{truth}) < 0.3$ requirement.

The energy response of reconstructed jets is defined as

$$R_E = \frac{E_{\text{reco}}}{E_{\text{truth}}}, \quad (8.3)$$

where E_{reco} is the energy of the reconstructed jet (either unmodified, or post-PFlow corrections) and E_{truth} the energy of the truth jet it is matched to. The energy resolution, σ_E , measures the uncertainty in the response; it is calculated as the standard deviation of the response histogram. Because the absolute magnitude of the resolution is dependent on the scale of the response, comparisons between the different jet definitions are made using $\sigma_E / \langle R_E \rangle$ as the metric.

Response distributions for jets with $50 \text{ GeV} < p_T < 60 \text{ GeV}$, in bins of $|\eta|$, are shown in Figure 8.6. The PFlow corrections, described in items 1 to 4, result in a large shift of the average response towards unity when compared to the baseline EMTopo jet response. This shows an improved estimation of the jet energy when tracker information is taken into account, in line with the premise of the PFlow approach to reconstruction. However, the difference between the different PFlow configurations themselves is small, with minimal impact observed when switching from the 2-dimensional to the 3-dimensional $\langle E/p \rangle$ parametrisation. The $\langle E/p \rangle$ configuration that includes $E/p \sim 0$ tracks yields a slightly higher response compared to when they are vetoed; this is because less energy is subtracted from the calorimeter measurement in the former scenario, resulting in greater jet energy values after applying the correction.

These trends are summarised by the mean response and equal-scale resolutions ($\sigma_E / \langle R_E \rangle$), shown for each jet definition in Figures 8.7 and 8.8 as a function of the EMTopo jet p_T and in various $|\eta|$ bins. Similar distributions along $|\eta|$ and in bins of jet p_T are included in Figure 8.18 of section 8.A. In general, the PFlow corrections are found to improve the energy response and




Figure 8.4. E/p distributions for $\pi^\pm$ tracks in dijet events. Only tracks falling within the cone of a reconstructed jet ($\Delta R(\text{jet}, \text{track}) < 0.4$) and with $E/p > 0.01$ are included. Distributions are shown for tracks with momenta between 2–5 GeV (a, c, e, g) or 5–15 GeV (b, d, f, h), with the plots in each group corresponding to different track $|\eta|$ bins. The histograms in each subfigure correspond to different $\Delta R(\text{jet}, \text{track})$ bins.

8.2. Comparisons using a simple Particle Flow




Figure 8.5. $\langle E/p \rangle$ values for the configuration described in item 2. Each of these plots correspond to a different $\Delta R(\text{jet}, \text{track})$ bin, as labelled.

resolution over the EMTopo baseline. This is most significant in the central $|\eta|$ regions ($|\eta| \lesssim 2$), with improvements becoming smaller for more forward jets.

The 3-dimensional $\langle E/p \rangle$ parametrisation, which includes the binning in ΔR , provides a minimal improvement to the resolution over the 2D $\langle E/p \rangle$ in all kinematic bins. However, differences between the configurations that include or exclude the $E/p \sim 0$ tracks are systematically present; despite pushing the average response towards unity when including such tracks in the $\langle E/p \rangle$ calculation, Figure 8.8 shows that this impacts the absolute jet resolution negatively, with degradations of up to a percent in some bins. This suggests that it is more important to accurately estimate the energy deposited by “hard” tracks (those with significant E/p values) than to underestimate it in an attempt to account for tracks that deposit very little energy.

8.2.3 Comparisons to an ideal Particle Flow

In the previous section, we studied the impact of implementing a simple PFlow algorithm using different sets of $\langle E/p \rangle$ reference values. The results showed that while a marked improvement in the jet energy resolution was achieved over the baseline EMTopo definition, the use of different $\langle E/p \rangle$ setups was largely irrelevant, with very similar performances seen for the configurations tested. This raises the question of whether the various cases performed similarly because of a general insensitivity of the parametrisation to additional degrees of freedom, or whether it is because a 2-dimensional parametrisation already performs close to some “ideal” PFlow scenario.

To test this, the extent of improvement expected to be achievable with the simple PFlow algorithm is investigated. One can note that the corrections detailed in eq. (8.2) boil down to two basic operations: adding the momentum of tracks to the jet; and subtracting off the expected en-




Figure 8.6. Energy response of EMTopo and PFlow jets defined using different $\langle E/p \rangle$ configurations. These distributions correspond to jets with $50 \text{ GeV} < p_T < 60 \text{ GeV}$ in various $|\eta|$ bins.

ergy deposited by the corresponding charged particles. In principle, perfect reconstruction can thus be achieved by using the true value of these quantities when applying the correction, without relying on any reconstructed and/or parametrised estimations. Particularly in the context of the current algorithm under study, direct use of the calibration hits provides the truth value for the latter, while the energy of the truth particle associated with the reconstructed track can be used in substitute for the former.

Taking all combinations of possible implementations into account, three idealised approaches can be defined:

1. **E/p , ideal track:** use the calibration hit energy instead of $\langle E/p \rangle \times p_{\text{track}}$
2. **ideal PFlow:** use the truth particle energy instead of adding p_{track}
3. **ideal PFlow+track:** use calibration hits instead of $\langle E/p \rangle \times p_{\text{track}}$, and truth energy instead of p_{track}

To account for differences between the EM and visible energy scales, these corrections are applied on a baseline jet energy that is computed as the sum of calibration hits in all clusters matched to the jet within a cone of $\Delta R < 0.4$. Here, I have used “ideal PFlow” to denote the substitution of $\langle E/p \rangle$ for calibration hit energies, since this represents the tunable aspect of particle flow algorithms. The comparisons to ideal track scenarios are included for completeness, but

8.2. Comparisons using a simple Particle Flow




Figure 8.7. Mean energy response of EMTopo and PFlow jets defined using different $\langle E/p \rangle$ configurations, plotted as a function of the baseline EMTopo jet p_T in various $|\eta|$ bins.

it is important to recognise that tracks are simply an input to PFlow and that improving track reconstruction falls well outside of its remit.

The mean response and resolution of the idealised jet definitions, along with that for EMTopo jets and one parametrised $\langle E/p \rangle$ case (item 2) for comparison, are shown in Figures 8.9 and 8.10. In the barrel region, a significant improvement to the jet resolution (up to 2%) is seen when the ideal PFlow correction is applied, suggesting that there are still avenues through which the estimation of $\langle E/p \rangle$ could improve PFlow performance. The ideal track correction alone offers no improvement to the resolution (black dashed vs solid histograms).³ However, the use of ideal track corrections further improves the performance of the ideal PFlow case, especially in the central region ($|\eta| < 0.5$).

The situation changes in the endcap regions ($|\eta| > 2$), with the ideal PFlow not improving the jet resolution over the standard E/p approach. This behaviour is not fully understood and warrants potential further study. One possible explanation is that the estimated $\langle E/p \rangle$ and re-

³The slight degradation seen is caused by biases in the track momentum reconstruction. On average, the reconstructed momentum of a track tends to be greater than the momentum of its associated truth particle (section 8.C). This results in a slightly lower jet response, and consequently, larger values when dividing the resolution by the mean response.




Figure 8.8. Energy resolution of EMTopo and PFlow jets defined using different $\langle E/p \rangle$ configurations, plotted as a function of the baseline EMTopo jet p_T in various $|\eta|$ bins.

constructed track momentum each tend to be biased with respect to their true values in such a way that manifests in the final jet resolution when only one of the ideal corrections is applied. For example, if the $\langle E/p \rangle$ and reconstructed track momentum are both overestimated, a correction that involves subtracting off the former and adding the latter could result in partial cancellations of the individual biases. This could be related to the presence of tracks with $E/p \approx 0$, which are more ubiquitous in the forward region of the detector. Such tracks result in close to no energy being subtracted from the jet in the ideal PFlow scenario; the correction then amounts essentially to an addition of the track momentum to the jet. This induces a small tail in the response distribution of the ideal PFlow scenarios (Figure 8.19 in section 8.A). Regardless, an improvement to the jet resolution is seen when both ideal corrections are used in conjunction, suggesting that track reconstruction could be important for improving performance in the forward region.

8.3 Comparisons in Athena

In section 8.2, we studied a ΔR -dependent parametrisation of $\langle E/p \rangle$ using a simplified PFlow correction algorithm to draw comparisons of the resulting jet performance. However, ATLAS

8.3. Comparisons in Athena




Figure 8.9. Mean energy response of EMTopo and PFlow jets defined using idealised corrections, plotted as a function of the baseline EMTopo jet p_T in various $|\eta|$ bins.

employs a more sophisticated approach to PFlow, with a subtraction algorithm that removes cells from topo-clusters matched to selected tracks. In this section, the performance of the various $\langle E/p \rangle$ configurations described in section 8.2.1 will be compared using the standard ATLAS algorithms implemented within the Athena reconstruction framework [151].

8.3.1 Alternative ΔR -dependent $\langle E/p \rangle$ parametrisations

In the standard PFlow jet reconstruction chain, selected tracks are matched to clusters and the subtraction algorithm defines charged and neutral Particle Flow Objects (PFOs) based on these input objects. This is the step in which the binned $\langle E/p \rangle$ values are read as reference, and used to correct the energy of neutral objects matched to tracks. After this, PFlow jet finding proceeds by clustering these PFOs.

This poses the first hurdle in the current endeavour: the use of $\Delta R(\text{jet}, \text{track})$ in the $\langle E/p \rangle$ map parametrisation is not possible if the jet axis is yet to be defined. Obviously, this was not an issue with the simplified PFlow algorithm described in eq. (8.2) since it only applies corrections to existing jets; however, it necessitates the definition of an alternative variable, which exhibits the same (or similar) dependence of E/p as $\Delta R(\text{jet}, \text{track})$, for the purpose of this exercise. Of course,




Figure 8.10. Energy resolution of EMTopo and PFlow jets defined using idealised corrections, plotted as a function of the baseline EMTopo jet p_T in various $|\eta|$ bins.

the $\Delta R(\text{jet}, \text{track})$ variable itself was only used as a convenient proxy for the local energy density around a track of interest, so an alternative metric of this quantity that is calculable during the PFO construction step should provide an adequate parametrisation.

One such metric is the ΔR between each selected track and the closest other track to it. The other track is required to have $p_T > 1$ GeV and satisfy the same quality criteria as the candidate track. A low minimum ΔR value guarantees the presence of at least one other (relatively) hard particle in close proximity to the selected track, and in principle should allow for a greater chance of the selected track particle's deposit to be clustered. While this does not provide a measure of any neutral energy contributions (thus potentially reducing the E/p dependence of tracks that fall closer to a neutral particle than another charged particle), a dependence of the track E/p on this variable can still be seen (Figure 8.11).

8.3.2 Modifications to eflowRec

The construction of charged and neutral PFOs from standard reconstructed tracks and topo-clusters is handled by the eflowRec package in Athena. Modifications to the package (summarised in Figure 8.12) were implemented to allow for alternative $\langle E/p \rangle$ reference values to be

8.3. Comparisons in Athena




Figure 8.11. The E/p of tracks as a function of the ΔR between the track and the closest other track with $p_T > 1$ GeV. The ΔR is calculated using track coordinates extrapolated to the calorimeter, and red lines denote the mean E/p value in each ΔR bin.




Figure 8.12. Modifications made in efLowRec to implement ΔR -dependent $\langle E/p \rangle$ reference values.

read from a JSON input, and used in place of the default parametrisation used by ATLAS. These references were derived in the same momentum, $|\eta|$ and ΔR bins described in section 8.2.1. No parametrisation in the LFI was included, and shower shape parameters were left unchanged from the default configuration.

8.3.3 Results

Kinematic distributions of PFlow jets and their constituents are shown in Figure 8.13 for three cases: the default ATLAS reconstruction, using a 2D $\langle E/p \rangle$ parametrised in track momentum and $|\eta|$, and a 3D binning additionally in $\Delta R(\text{track})_{\min}$. Little difference is observed between the different parametrisations.




Figure 8.13. Comparison of PFlow jet and constituent kinematics when different E/p configurations are used. (a) The reconstructed jet energy, (b) jet p_T , (c) energy of neutral jet constituents, and (d) summed energy of all neutral jet constituents in each event.

8.4 Conclusion

In this chapter, an alternative approach to E/p for Particle Flow has been studied, with an additional parametrisation in the ΔR between the track and the jet axis. A dependence of the track E/p with this ΔR variable was seen, with lower momentum tracks typically depositing a greater fraction of their energy into clustered calorimeter cells. The effect of an additional parametrisation of the $\langle E/p \rangle$ in $\Delta R(\text{jet}, \text{track})$ was studied using a simple energy correction algorithm, which represents the limiting case of PFlow when particle showers fully overlap. A small improvement to the jet energy resolution is seen when the ΔR parametrisation is included.

8.A. Additional plots

8.A Additional plots




Figure 8.14. The E/p of tracks with momenta between 0–2 GeV as a function of the ΔR between the track and the jet it is matched to, in different track $|\eta|$ bins. The ΔR is calculated using track coordinates extrapolated to the calorimeter, and red lines denote the mean E/p value in each ΔR bin.




Figure 8.15. The E/p of tracks with momenta between 15–25 GeV as a function of the ΔR between the track and the jet it is matched to, in different track $|\eta|$ bins. The ΔR is calculated using track coordinates extrapolated to the calorimeter, and red lines denote the mean E/p value in each ΔR bin.




Figure 8.16. The E/p of tracks with momenta between 25–40 GeV as a function of the ΔR between the track and the jet it is matched to, in different track $|\eta|$ bins. The ΔR is calculated using track coordinates extrapolated to the calorimeter, and red lines denote the mean E/p value in each ΔR bin.




Figure 8.17. The E/p of tracks with momenta between 40–100 GeV as a function of the ΔR between the track and the jet it is matched to, in different track $|\eta|$ bins. The ΔR is calculated using track coordinates extrapolated to the calorimeter, and red lines denote the mean E/p value in each ΔR bin.

8.A. Additional plots




Figure 8.18. Mean energy response (a, c, e, g) and resolution (b, d, f, h) of EMTopo and PFlow jets defined using different $\langle E/p \rangle$ configurations, plotted as a function of the EMTopo jet $|\eta|$ in various p_T bins.




Figure 8.19. Energy response of EMTopo and PFlow jets defined using different $\langle E/p \rangle$ configurations. A small tail can be seen for the ideal PFlow cases (cyan), most prominent for jets in the barrel.

8.B. Simple vs ATLAS Particle Flow




Figure 8.20. Energy response distributions, scaled by $\langle R_E \rangle^{-1}$ along the x -axis, for (uncalibrated) EMTopo (magenta), standard ATLAS PFlow (azure), and simple PFlow (black) jets. These distributions correspond to jets with $50 \text{ GeV} < p_T < 60 \text{ GeV}$ in various $|\eta|$ bins.

8.B Simple vs ATLAS Particle Flow

The simple PFlow procedure introduced in section 8.2 is equivalent to the limit of the standard ATLAS PFlow when particles maximally overlap, such that their showers within the calorimeter cannot be disentangled. Thus, it is interesting to compare the performance of the two approaches.

Figure 8.20 shows response distributions of the two PFlow jet definitions, along with that for EMTopo jets for comparison, after scaling each to the same energy scale. This allows for a more direct comparison of the width of each distribution, and was achieved by first fitting the respective response histograms with a Gaussian functional form and scaling the result by the inverse of the mean response.

Figure 8.21 shows the mean energy response and resolution as a function of jet p_T in different $|\eta|$ bins. The ATLAS PFlow definition provides the largest improvement to the jet energy scale, consistent across the kinematic bins. The resolution of ATLAS PFlow jets are also an improvement over EMTopo jets for $|\eta| < 1.5$, although the performance degrades in the more forward regions. The simplified PFlow definition appears to provide the best resolution; it would be interesting to probe whether certain aspects of the approach result in this behaviour, but this will not be studied here.

In general, these comparisons suggest that there are still opportunities for improving the ATLAS PFlow jet energy scale in certain regions of phase space.




Figure 8.21. Mean energy response (a, c, e, g) and resolution (b, d, f, h) of EMTopo, ATLAS PFlow, and simple PFlow jets, plotted as a function of the jet p_T in various $|\eta|$ bins.

8.C Track–cluster resolution comparison

The improvement to the jet resolution in $|\eta| > 2$ when using the ideal PFlow+track correction (Figure 8.10) prompted a brief look into the tracker and calorimeter resolutions in different regions of the detector. In the current implementation of PFlow, only tracks with momentum up to 100 GeV are considered, with the argument being that the calorimeter provides a better measurement resolution above this threshold. In this section, we seek to reproduce this result in different bins of $|\eta|$ using a sample of dijet events where the leading truth jet p_T ranges from 160–400 GeV.

Tracks are required to satisfy the TightPrimary quality criteria and have a valid association to a truth particle. The track response is computed as the ratio of the reconstructed track momentum to the associated truth particle’s momentum. For clusters, the sum of visible calibration hit energy over all contained cells is used as the measure of its truth energy, and the ratio of its uncalibrated (i.e. EM-scale) energy to this sum is taken as the response. Since events are not simulated with pileup, the discrepancy between EM-scale and calibration hit energies can be attributed to the inherent detector response, differences between the EM and visible energy scales (chiefly due to calorimeter sampling fractions being different for EM particles and hadrons), and contributions from cell noise. The mean response and resolution are calculated as the mean and standard deviation of the response histograms defined in this way.

The left-most column in Figure 8.22 shows the mean response of tracks and clusters, plotted as a function of the reconstructed track momentum or cluster energy respectively, in five different bins of $|\eta|$. For clusters, a large positive deviation from unity is seen at lower energies due to contributions from cell noise having a greater impact. This quickly stabilises for clusters above approximately 50 GeV, although remains offset from unity to varying degrees depending on the $|\eta|$ bin. The greatest discrepancy is observed for clusters in $1 < |\eta| < 1.5$, likely due to this corresponding to a transition region where interactions with the detector material can produce many secondary particles. Above $|\eta| > 2$, the response converges to a value close to unity.

For tracks, the mean response exhibits similar behaviour in all $|\eta|$ bins: at the lowest momenta, it is centred at unity, trending upwards with increasing track momentum. This degradation occurs most quickly for tracks in the centre of the detector, and is more stable for particles at greater $|\eta|$; for example, the mean deviation for 300 GeV tracks in $|\eta| > 2$ is on the order of 1%, compared to approximately 30% in the central barrel region.

A comparison of the scaled resolution in the same $|\eta|$ bins is shown in the middle column of Figure 8.22. The resolution of the tracker is better than the calorimeter for low energy particles, as expected, with opposite trends for the two systems as the particle energy increases. In the centre of the detector, the crossing point where the two detector systems yield equal reconstruction resolution is situated at approximately 120 GeV, which agrees with the performance reported in their respective technical design specifications. For more forward particles, the intersection point shifts to higher energy values due to a combination of poorer calorimeter resolution and an improvement in the reconstruction of high energy particle tracks, with the resolution of the latter being approximately twice as good in the endcaps than tracks in the barrel of the same momentum. This behaviour seems to emerge partially due to how the resolution has been defined in these plots; response histograms for an illustrative momentum bin (240–260 GeV) show that tracks at small $|\eta|$ in fact exhibit a narrower distribution, but with a small positive tail that increases the standard deviation computed for the histogram (right-most column in Figure 8.22). The reason for this tail is not fully understood. Regardless, the conclusion remains that the tracker provides better measurements than the calorimeter to higher energies as $|\eta|$ increases.




Figure 8.22. Mean response (a, d, g, j, m) and resolution (b, e, h, k, n) of tracks and clusters in various $|\eta|$ bins. (c, f, i, l, o) Response histograms for 240–260 GeV tracks and clusters in various $|\eta|$ bins.

Overlapping Jets

THE increase in instantaneous beam luminosity at the High-Luminosity Large Hadron Collider (HL-LHC) ushers in an unprecedented era of data-taking, offering the potential of delivering up to 4000 fb^{-1} of data over its 12-year operational period. Facilitating such extreme conditions leads to a corresponding increase in the rate of pile-up (PU) interactions that contaminate each hard-scatter (HS) event of interest, with the mean number of interactions per pp bunch crossing growing from $\langle\mu\rangle \approx 60$ in Run 3 to an expected rate of approximately $\langle\mu\rangle \approx 200$ at the HL-LHC. In particular, each bunch crossing in the high-luminosity environment is expected to feature more than one hard collision on average, with the additional particles produced in these hard-QCD interactions also contributing to the PU in the event.

Currently, ATLAS jet reconstruction implements PU suppression techniques that account for the presence of soft stochastic objects with a correction to the jet energy based on the expected PU density in the event, and with a jet vertex tagger (JVT) for rejecting pure PU jets based on the vertex the jet is associated with [146]. However, the increase in the rate of hard PU activity at the HL-LHC presents a potential new challenge in the form of overlaps between HS and hard PU jets. Since hard PU activity is more likely to manifest as localised high-energy fluctuations within the detector, the current mitigation approach based on average p_T densities would not correct for such overlaps; in principle, a case-by-case correction applied to each reconstructed jet will instead be necessary.

The work presented in this chapter constitutes a first exploratory study into overlapping jets at the HL-LHC. I will begin by investigating how often hard PU jets are expected to occur and to overlap with HS jets, before aiming to quantify the impact on the final reconstructed jet in the scenario of an overlap. Studies on the prevalence of out-of-time PU jets¹ and the effect of their overlap with HS jets will also be presented. The chapter concludes with a discussion on the prospects for future studies following from this work.

9.1 Event samples and object definitions

The studies were conducted using Monte Carlo (MC) simulations of multijet production in pp collisions at a centre of mass energy of $\sqrt{s} = 14 \text{ TeV}$, using a GEANT4 [103] model of the HL-LHC detector configuration. Events were generated using PYTHIA 8.307 to leading order (LO) in α_s , with the NNPDF2.3LO set of parton distribution functions (PDFs) used for the generation of matrix-elements, the parton shower, and the simulation of multi-parton interactions [152].

¹An introduction to the different types of PU can be found in section 3.5.2.

9.2. Reconstructed jet categories

The underlying event was tuned to the A14 [132] set of parameters. The effect of multiple interactions in the same and neighbouring bunch crossings (i.e. PU) was modelled by overlaying the simulated hard-scattering event with inelastic pp events generated with PYTHIA 8.186 [121] using the NNPDF2.3LO PDF set [133] and the A3 set of tuned parameters [153].

Jets are defined using the anti- k_t algorithm [65] with a radius parameter of $R = 0.4$. For detector-level reconstructed jets, clusters of calorimeter cells are used as input to the algorithm. Corrections to the energy scale of these jets are applied according to the nominal calibration procedure adopted in ATLAS. In addition, three particle-level truth jet collections (hard-scatter (HS), in-time (i.t.) PU, and out-of-time (o.o.t.) PU) are also defined by using the particles in the MC truth record as input to the anti- k_t algorithm. These consider all stable particles in the event, excluding muons and weakly interacting particles (neutrinos), and prompt leptons and photons from decays of τ leptons or Higgs, W or Z bosons (as well as any final-state radiation photons around these prompt leptons). The o.o.t. jet collection considers particles produced in up to the five preceding and three succeeding bunch crossings, based on the extent to which o.o.t. effects are expected to impact the current event.

9.2 Reconstructed jet categories

While the production of hard QCD PU jets in higher luminosity conditions is heuristically understood, their impact on object reconstruction has yet to be studied. The goal of this section is to conduct a first quantitative study of the rate at which hard PU jets and their overlap with HS particles are expected to occur at the HL-LHC.

Reconstructed jets with $p_T > 30$ GeV are considered for the study, as the accurate calibration of jets below this threshold is expected to be challenging in PU-dense environments. A cone of $\Delta R < 0.4$ is used to perform the matching of each reconstructed jet to truth HS and truth i.t. and o.o.t. PU jets. To focus on contributions from relatively hard PU particles, only truth jets with $p_T > 15$ GeV are considered. Categories are then defined based on the truth jets that are matched:

- **hard-scatter:** exactly one truth HS jet and no truth PU jets
- **HS + i.t. PU:** exactly one truth HS jet and at least one truth i.t. PU jet
- **HS + o.o.t. PU:** exactly one truth HS jet, at least one truth o.o.t. PU jet, and no truth i.t. PU jet
- **hard PU (i.t.):** no truth HS jet and exactly one truth i.t. PU jet
- **hard PU (o.o.t.):** no truth HS jet and exactly one truth o.o.t. PU jet
- **stochastic:** no truth HS jet and not classified as hard PU

The categories are mutually exclusive and cover all possible matching combinations, with the exception of cases where multiple truth HS jets overlap, as this is process dependent and of tangential interest to the present study. Note that the i.t. overlap category (“HS + i.t. PU”) can include jets that overlap with both i.t. and o.o.t. PU particles. However, in these scenarios, the i.t. contributions are expected to have a greater impact on the resultant jet due to the suppressive nature of the LAr calorimeter towards o.o.t. signals²; the overall effect in these scenarios

²As discussed in section 4.3.2.




Figure 9.1. Display of the η - ϕ plane in a simulated multijet event, which includes examples of jets with contributions from hard-scatter, pile-up, or an overlap of both. Markers are placed at the position of reconstructed (truth) jets with $p_T > 30$ (15) GeV. The green circles represent the cone of $R = 0.4$ around each reconstructed jet.

will thus be similar to cases where the HS jet overlaps only with i.t. PU particles. Figure 9.1 displays an example of a multijet event that contains these various reconstruction scenarios, where different markers have been used to denote the positions of each type of jet in the η - ϕ plane, and transparent green circles represent the catchment area of each reconstructed jet.

The fraction of reconstructed jets falling into each category is shown in Figure 9.2. The pure HS and PU categories correspond to jets that the current PU suppression techniques account for; the former via a density-based subtraction, and the latter with the JVT algorithm. While these two broad categories remain the most abundant, a significant fraction of jets are also found to overlap with hard o.o.t. PU. The rate of overlaps with hard i.t. PU is lower, but still not negligible, with approximately 8% of the jets (out of those that are matched to a truth HS jet) falling into this category.

Of course, it is expected that the contribution of reconstructed jets to each category should vary with the properties of the jet. Figure 9.3 presents two results in which the fraction has been parametrised as a function of the jet p_T or $|\eta|$. In Figure 9.3a, the relative rate of jets reconstructed from HS particles is compared against those for which no truth HS was matched, with the inclusive HS category corresponding to the union of the HS and HS+PU categories defined above. PU jets dominate at low p_T , as expected; however, their contribution remains sizeable up to a jet p_T of approximately 100 GeV, unlike in the current LHC conditions.

Figure 9.3b considers only at the reconstructed jet categories that include a matched truth HS particle, plotted as a function of the jet $|\eta|$. Approximately 40% of HS jets are found to overlap with hard PU particles; this occurs more frequently in the barrel region, and remains the case across the jet p_T range. These results make it clear that overlaps between HS and PU jets (both i.t. and o.o.t.) are a common occurrence in high luminosity conditions, and it is thus important to understand the impact of this on the jet reconstruction.

9.2. Reconstructed jet categories




Figure 9.2. Fraction of reconstructed jets in the various categories defined based on the number and type of truth jets matched to them.




Figure 9.3. Fraction of reconstructed jets in various categories, defined based on the number and type of truth jets matched to them, (a) as a function of the jet p_T (with jets that contain truth HS combined into one inclusive category), and (b) as a function of $|\eta|$, considering only the categories that include a matched truth HS jet.




Figure 9.4. Number of reconstructed jets with $p_T > 20$ GeV matched to isolated truth HS jets with $p_T > 30$ GeV within a cone of $\Delta R < 0.4$.

9.3 Effect of hard-scatter overlaps with pile-up

When HS and PU particles are produced close-by in the detector, whether they are resolved separately or merged into a single reconstructed jet is determined by two factors: the separation between the particles (ΔR); and their relative hardness (e.g. the ratio of their transverse momenta). The goal of this section is to study the impact on the reconstructed jet in the cases where the particles are not resolved.

To begin, truth HS jets with $p_T > 30$ GeV and that are isolated from other truth HS jets (of any p_T) within a cone of $\Delta R < 0.8$ are selected. An attempt is then made to match these jets to reconstructed jets above 20 GeV (to allow for detector smearing effects) within a cone of $\Delta R < 0.4$. In the majority of cases, exactly one reconstructed jet is found (Figure 9.4), and the number of truth PU jets with $p_T > 15$ GeV in $\Delta R < 0.4$ of the reconstructed jet is used to classify the jet as non-overlapping (no truth PU jets are found); overlapping with only i.t. PU particles; or overlapping with only o.o.t. PU particles. Cases where more than one reconstructed jet is matched, or where both i.t. and o.o.t. PU jets contribute, are ignored.

Figure 9.5 shows the distribution of the number of truth i.t. and o.o.t. PU jets (above 15 GeV) matched to the reconstructed jets selected. Relatively hard i.t. PU jets are found within the catchment area of the reconstructed jet in approximately 10% of the cases. These merged jets arise typically when the HS and PU particles are separated by less than the specified jet radius (Figure 9.6a), with the possibility of the PU particles contributing a significant fraction of the jet scale. This can be seen in Figure 9.6b, which shows the ratio of PU to HS contributions to the overlapping jets. In the event that multiple PU jets are matched to the reconstructed jet, the scalar sum of p_T is taken as the numerator.

9.3.1 Overlaps with in-time pile-up

The performance of a jet reconstruction algorithm can be measured through the p_T response, defined as the ratio of the reconstructed jet p_T to that of the matched truth HS jet. After applying calibrations to recover the particle-level energy scale, a well-reconstructed jet will exhibit a response close to unity (by definition). The precision of the reconstruction is represented by the

9.3. Effect of hard-scatter overlaps with pile-up




Figure 9.5. Number of truth i.t. and o.o.t. PU jets with $p_T > 15$ GeV matched to reconstructed jets with $p_T > 20$ GeV in $\Delta R = 0.4$.




Figure 9.6. The ΔR between truth HS and matched truth PU jets (a) and the ratio of their p_T (b).

resolution, which is calculated in the following as the standard deviation of the response distribution scaled by the inverse of the mean response. The resolution of calorimeter jets improves with increasing particle energy.

These metrics are used to compare the performance of reconstructed jets that contain varying degrees of i.t. PU contributions. Distributions of the p_T response are shown in Figure 9.7, and the mean response and resolution in Figure 9.8 as a function of the p_T ratio between the PU and HS contributions. In Figure 9.8a, the scale of the reconstructed jet is seen to increase accordingly with the relative hardness of the PU contribution. While this result may be unsurprising at first glance, it highlights that the present jet calibration chain does *not* correct for the presence of hard PU particles.

The dependence of the jet resolution on the p_T ratio is shown in Figure 9.8b for jets with p_T between 60 GeV and 80 GeV. The restriction on the p_T range is essential to mitigate any effects from the correlation between the jet resolution and the jet p_T , since it is more likely for lower p_T jets to contain a larger PU fraction (due to the falling PU p_T spectrum). The jet resolution is found to be largely resilient against PU contributions; this remains the case in other p_T bins. Thus, despite the drastic effect on the reconstructed jet scale, this result indicates that the performance of reconstructing overlapping jets is recoverable provided that the jet can be accurately identified and corrected back to the particle-level scale.




Figure 9.7. Distributions of the p_T response for jets reconstructed from overlapping HS and i.t. PU particles. The histograms correspond to different bins of the p_T ratio between PU and HS particles.




Figure 9.8. Mean p_T response (a) and resolution (b) of jets reconstructed from overlapping HS and i.t. PU particles, as a function of the PU to HS jet p_T ratio. Results are shown for jets with $p_T \in (60, 80)$ GeV, but the trends remain similar in other p_T bins.

9.3.2 Overlaps with out-of-time pile-up

In general, the bimodal shape of the LAr calorimeter response suppresses signals produced in o.o.t. bunch crossings, such that their primary effect is to contribute to stochastic pile-up jets [154]. Energetic o.o.t. particles can still have deposited enough energy to seed a jet, particularly those from adjacent bunch crossings. Figure 9.9a shows the fraction of isolated truth o.o.t. PU jets that can be matched to a reconstructed jet using a $\Delta R < 0.3$ requirement. The isolation check is performed with respect to all other truth HS and truth PU (both i.t. and o.o.t.) jets in the event, and uses a matching radius of $\Delta R < 1.0$. Approximately 10% of the jets with $p_T > 70$ GeV produced in an adjacent bunch crossing are found to result in a reconstructed jet in the current bunch crossing. The fraction of reconstructed jets that this corresponds to in each event is shown in Figure 9.9b.

The comparisons of section 9.3.1 can be repeated to study the impact of o.o.t. PU particles overlapping with HS particles during jet reconstruction. The selection of isolated truth HS jets

9.4. Outlook




Figure 9.9. Efficiency for o.o.t. PU jets, produced in collisions up to five bunch crossings away, to seed a reconstructed jet in the current bunch crossing (a) and the fraction of reconstructed jets in an event this corresponds to (b).

and their matching to reconstructed calorimeter jets remains as previously described; however, now we only consider reconstructed jets matched to truth HS and o.o.t. PU jets above 15 GeV.

Figure 9.10 shows the p_T response distributions for different levels of relative o.o.t. PU contribution. Unlike cases of overlaps with i.t. PU particles, o.o.t. contributions do not affect the energy scale of the reconstructed jets, with the mean p_T response remaining centred at unity even when the overlapping o.o.t. particle is of a similar scale as the HS jet (Figure 9.11a). The apparent increase in the jet resolution stems from the correlation between the jet p_T and the PU–HS p_T ratio. When considering a restricted jet p_T range (Figure 9.11b), it can be seen that the jet resolution remains consistent irrespective of the magnitude of o.o.t. PU contributions.

9.4 Outlook

The collision environment at the HL-LHC presents a new challenge in the form of overlapping jets, which will become more prevalent due to the marked increase in the rate of hard-QCD pile-up interactions. This chapter has presented a first study of such scenarios using simulations of dijet production at the HL-LHC. Overlaps between HS and PU jets are found to occur at a rate of approximately 60% across the entire jet p_T range, and being slightly more prevalent in the barrel region of the detector. Pile-up jets (both stochastic and hard) dominate below $p_T < 60$ GeV. Jets that overlap with hard in-time PU experience an increase in the scale of their response, with the magnitude of the shift dependent on the relative p_T of the HS and PU particles. However, the performance of the reconstruction can be recovered by correcting for this shift in scale. For jets that overlap with significant out-of-time PU deposits, both the scale and resolution of the reconstructed jet remain largely unaffected due to the suppression of LAr calorimeter signals towards out-of-time energy deposits.

The significant impact of hard, in-time PU contributions highlights the importance of accounting for overlapping jets during event reconstruction at the HL-LHC. While techniques for




Figure 9.10. Distributions of the p_T response for jets reconstructed from overlapping HS and o.o.t. PU particles. The histograms correspond to scenarios of different relative hardness between the HS and PU contributions.




Figure 9.11. Mean p_T response (a) and resolution (b) of jets reconstructed from overlapping HS and o.o.t. PU particles, as a function of the (sum of) PU to HS jet p_T ratio. Results are shown for jets with $p_T \in (60, 80)$ GeV, but the trends remain similar in other p_T bins.

tagging such jets have not been explored in this chapter, previous studies have been conducted under similar motivations, e.g. in ref. [155]. The idea is that by computing a JVT-like score for each jet, one could form a discriminant to identify jets that receive contributions from multiple primary vertices. Jet substructure moments (e.g. n -subjettiness [69]) could also prove powerful in conjunction, as the shape and evolution of the shower would be different for jets that are reconstructed from the signals of a single versus multiple partons. Corrections to the overlapping jet could then be applied based on the relative hardness of the contributions from each vertex. This could be extended such that independent jet collections are reconstructed for each HS vertex in the event, with a different correction being applied to the overlapping jet depending on which vertex is being analysed. If a good reconstruction of the “event” corresponding to each vertex can be achieved, this offers the possibility of independently analysing the physics processes behind each HS vertex, allowing for the potential to exploit the multiple HS interactions expected to occur in each bunch crossing at the HL-LHC.

A SEARCH FOR CHARMS AND INVISIBLES

Motivation and Setup



UPERSYMMETRIC (SUSY) extensions of the Standard Model (SM) of particle physics predict new particles that are partners of SM particles [156–161]—the partners of SM fermions being bosons and those of SM bosons being fermions—and provide solutions to the hierarchy problem [162–165]. R -parity-conserving SUSY models [166] are considered for this work; these predict that the supersymmetric partners of quarks (squarks) will be produced in pairs at the Large Hadron Collider (LHC) and that the lightest SUSY particle (LSP) is stable. The lightest neutralino, $\tilde{\chi}_1^0$, is considered to be the LSP and is a dark-matter candidate [167, 168].

This part of the thesis will present a search that targets SUSY models where top or charm squarks ($\tilde{t}_1/\tilde{c}_1$) are the lightest squarks, with masses less than about one TeV [169]. For the top squarks considered in this search, it is assumed that flavour-violating effects allow the top squark to decay into a charm quark and an LSP, as shown in Figure 10.1a. This results in a final state with two charm quarks and missing transverse momentum from the LSPs, which escape detection. The extent of flavour-violating effects is model dependent, but when there is significant mixing between the top squark and the charm squark [170] (which can be due to renormalisation group equations, even if minimal flavour violation is assumed at a certain scale), the branching fraction of $\tilde{t}_1 \rightarrow c\tilde{\chi}_1^0$ can exceed the branching fraction of $\tilde{t}_1 \rightarrow t^*\tilde{\chi}_1^0$, where t^* can be either an on-shell or off-shell top quark.

In supersymmetric models with minimal flavour-violating effects, charm squarks could be considerably lighter than other squarks. This motivates searches for pair production of charm squarks, with each charm squark decaying into a charm quark and an LSP as shown in Figure 10.1b. Notably, the model of charm squark pair production in Figure 10.1b shares the same experimental signature and production cross-section as for the top squark pairs in Figure 10.1a. The experimental signature includes jets that contain charm hadrons but no bottom hadrons and momentum imbalance caused by the two LSPs. The search discussed here was optimised for the case where the branching fraction of the decay of top or charm squark into a charm quark plus the LSP is unity. However, the results of the search are also interpreted in the scenario of top squark production with a varying branching fraction of the top squark decay into a charm quark plus the LSP ranging from 0.1 to 1, with the alternative decay being to a top quark plus the LSP.

Leptoquark (LQ) models [171–179] with up-type scalar LQs (LQ^u) or vector LQs (vLQ) predict their pair production and subsequent decay into $c\nu c\nu$, with the same experimental signature as in top and charm squark production models (Figure 10.1c). The following models are considered: U_1 vLQ models that couple to second-generation quarks and third-generation leptons (vLQ_{23}^u , where the subscript indicates the quark and lepton generations) in both the minimal and Yang–




Figure 10.1. Representative diagrams for pair production of (a) top squarks ($\tilde{t}_1$) and (b) charm squarks ($\tilde{c}_1$), with subsequent decay into charm quarks and neutralinos ($\tilde{\chi}_1^0$); and (c) pair production of leptoquarks decaying into charm quarks and neutrinos.

Mills coupling scenarios [25, 180], assuming a branching fraction ($\mathcal{B}$) of $\mathcal{B}(\nu LQ_{23}^u \rightarrow c\nu_\tau) = 0.5$ in both cases with the alternative decay being $\nu LQ_{23}^u \rightarrow s\tau$; and models where scalar LQs couple to second-generation quarks and first- or second-generation leptons [24, 181, 182], allowing a range of $\mathcal{B}(LQ^u \rightarrow c\nu_{e/\mu})$ values (here the alternative decay is $LQ^u \rightarrow se/\mu$).

This search for top squark, charm squark, and leptoquark pair production uses pp collisions from the LHC at a centre-of-mass energy of $\sqrt{s} = 13$ TeV recorded by the ATLAS experiment from 2015 to 2018. Of course, while I aim to provide a full overview of the analysis in this part of the thesis, it goes without saying that the realisation of these results was only possible through a collaborative effort. My personal contributions to this work are as follows: studying the charm tagger performance; preparing and validating new simulated signal event samples; designing the Compressed signal, control, and validation regions (notably including the development of the $N_{c\text{-jets}}^S$ variable); calculating signal theory (scale and radiation) uncertainties; conducting the branching fraction scan for particular SUSY signal assumptions; implementing the leptoquark interpretation; and preparing the auxiliary material (truth analysis implementation, acceptance and efficiency calculations, cutflow tables, and the HEPData record).

A previous search by ATLAS [183], based on the identification of jets containing charm hadrons but no bottom hadrons (c -tagging), used 36.1 fb^{-1} of collisions and was able to probe top and charm squark masses up to 850 GeV. In addition to a larger dataset, the search herein also benefits from improvements in c -tagging and from an advanced technique, Recursive Jigsaw Reconstruction [184], which makes use of jets originating from initial state radiation (ISR) to provide sensitivity to models with small mass splittings between the top squark and the LSP (further discussion of this technique will be included in Chapter 11).

Other ATLAS results also have sensitivity to the models considered in this analysis: a search using final states with an energetic jet (but no c -tagging) and large missing transverse momentum [185] was particularly sensitive to models with small mass differences between the top/charm squark and the LSP, while a search for single squark production that did not utilise c -tagging [186] had good sensitivity to models with large top/charm squark masses. A recent search for top squark pair production focused on squarks decaying into either a top quark or charm quark plus an LSP [187]. The most recent search performed by the CMS collaboration used 137 fb^{-1} of $\sqrt{s} = 13$ TeV pp collisions from the LHC, but only considered mass differences between the top squark and the LSP of less than 80 GeV, while the present ATLAS search explores mass differences ranging from 20 GeV to about 1 TeV. Finally, for the considered leptoquark models, this work is the first to report results for the described final state from an LHC experiment.

Table 10.1. Overview of the simulated background event samples. The simulated vector bosons, V , are massive W and Z bosons. The NNPDF3.0NLO [129] parton distribution functions (PDF) set is used.

Background process	Matrix element generator	PDF accuracy	Parton showering and hadronisation	Underlying event tune	Cross-section calculation accuracy
$pp \rightarrow V + \text{jets}$	SHERPA 2.2.11 [194]	NNLO	SHERPA	Default	NNLO [195]
$pp \rightarrow t\bar{t}V + \text{jets}$	AMC@NLO 2.3.3 [196]	NLO	PYTHIA 8.210 [131]	A14 [132]	NLO [196]
$pp \rightarrow t\bar{t} + \text{jets}$	POWHEG Box v2 [119]	NNLO	PYTHIA 8.230	A14	NNLO+NNLL [197–202]
$pp \rightarrow t/tW + \text{jets}$	POWHEG Box v2	NNLO	PYTHIA 8.230	A14	NNLO+NNLL [203–205]
$pp \rightarrow VV + \text{jets}$	SHERPA 2.2.1–2.2.2 [134]	NNLO	SHERPA	Default	NLO

10.1 Data and simulated event samples

The data used in this analysis were recorded with a fully operational detector from stable proton-beam collisions, and are clear of a significant amount of cosmic-ray or beam-induced background [188]. The total integrated luminosity of this dataset is $139.0 \pm 2.4 \text{ fb}^{-1}$ [189], determined using the LUCID-2 detector [190] for the primary luminosity measurements, complemented by measurements from the Inner Detector (ID) and calorimeters.

The LHC accelerates and collides protons in bunches, resulting in multiple pp collisions during each bunch crossing, alongside the much rarer hard-scattering pp collisions of interest. The bunch crossings happen every 25 ns. The average number of additional proton–proton interactions per bunch crossing (pile-up) grew from 13 in 2015 to a value of about 38 in 2017.

Missing transverse momentum triggers [191] were used to select the data for this analysis. The triggers had thresholds of 70 GeV in 2015 and early 2016, 100 GeV in late 2016 and 2017, and 110 GeV in 2018. These triggers reach an efficiency plateau for offline missing transverse momenta greater than 200 GeV.

The predictions for signal and background contributions are obtained with Monte Carlo (MC) simulated event samples. To match the LHC beam conditions, these simulated event samples also contain pile-up interactions. The simulated events were processed through the ATLAS detector simulation [192] and then subjected to the same reconstruction algorithms as the data. All samples were processed with the full GEANT4 [103] detector model with the exception of the samples with leptoquarks or direct top or charm squark production, which were processed with ATLFAST-II detector simulations [193]. Finally, the simulated event samples were corrected to improve the accuracy of the pile-up modelling, jet momenta, lepton efficiency and momenta, trigger efficiencies, missing transverse momentum, and c -tagging efficiency.

Standard Model background samples were produced with various event generators as shown in Table 10.1. All samples, except the ones that use the SHERPA event generator [194], were processed with EVTGEN 1.7.0 [127] for the b - and c -hadron decays. The cross-sections, which were used to normalise the samples, were calculated separately with higher-order corrections in the strong coupling (α_s) for all samples except for the diboson production ($pp \rightarrow VV + \text{jets}$) sample, where the cross-section calculated by the SHERPA event generator was retained. For the SHERPA sample of single vector bosons produced in association with jets ($V + \text{jets}$), virtual electroweak loop-terms were included at next-to-leading-order (NLO) accuracy.

The SUSY event samples were generated with MADGRAPH5_AMC@NLO 2.8.1 [196] at leading order (LO) in the strong coupling constant (α_s) with up to two additional partons using the NNPDF3.0NLO [129] parton distribution function (PDF) set and were showered and hadronised with PYTHIA 8.244 [131] using the A14 set of tuned parameters [132]. Samples with a small mass

10.2. Event reconstruction

difference between the top squark and $\tilde{\chi}_1^0$ (up to 140 GeV) were generated at a later point in time with MADGRAPH5_AMC@NLO 2.9.5 and then interfaced with PYTHIA 8.306. The matrix element (ME) calculation was performed at tree level including the emission of up to two additional partons. The matching between the ME and the parton shower was implemented using the CKKW-L prescription [206], with the matching scale set to a quarter of the top/charm squark mass. These samples were also processed with EVTGEN 1.7.0.

The top squark and charm squark signal cross-sections were calculated from pair production of scalar coloured particles and with the gluino assumed to be massive enough to not significantly contribute to the production process. Consequently, each production cross-section depends only on the squark mass. These cross-sections are computed at approximate next-to-next-to-leading-order (NNLO) accuracy in α_s with resummation of next-to-next-to-leading logarithmic (NNLL) soft gluon terms [207–210].

The samples with pair-produced scalar LQs were generated with matrix elements calculated at NLO accuracy in α_s with MADGRAPH5_AMC@NLO 2.6.0 [196], using the method described in ref. [24], in which NLO matrix elements [181, 182] are interfaced to PYTHIA 8.230. The total cross-sections of scalar LQ pair production were computed at approximate NNLO accuracy in α_s with resummation of NNLL soft gluon terms [207–210]. The cross-sections do not include lepton t -channel contributions, which are neglected in ref. [24] and may lead to corrections at the percent level [211]. For vector-LQ pair production, the matrix elements were calculated at LO in α_s following the model described in ref. [25], and events were generated using MADGRAPH5_AMC@NLO 2.9.5 in conjunction with PYTHIA 8.306. Parton luminosities were provided by the five-flavour scheme NNPDF3.0NLO PDF set and the underlying event was modelled with the A14 tune in PYTHIA. These samples were also processed with EVTGEN 1.7.0.

10.2 Event reconstruction

Events are required to have a primary vertex reconstructed from at least two tracks with transverse momentum $p_T > 0.5$ GeV. If more than one such vertex is found, the one with the largest sum of squares of transverse momenta of associated tracks is selected as the hard-scattering collision [212]. The other vertices are considered as pile-up.

Hadronic jet candidates are reconstructed using the anti- k_t jet algorithm [65, 64] with radius parameter $R = 0.4$, using particle flow objects (PFOs) [110] as inputs. PFOs consist of charged particle tracks matched to the hard-scatter vertex with the requirement $|z_0 \sin(\theta)| < 2.0$ mm (where z_0 is the longitudinal impact parameter), and calorimeter energy clusters surviving an energy subtraction algorithm that removes the energy contributions deposited by good quality tracks from any vertex. Jet energy scale corrections, derived from MC simulation and data, are used to calibrate the average energies of jet candidates to the scale of their constituent particles [73]. Jets with $p_T > 20$ GeV and $|\eta| < 4.5$ are used to calculate the missing transverse momentum [213], while jets with $|\eta| < 2.8$ are considered to be “signal” jets. Jets containing a large particle momentum contribution from pile-up vertices, as measured by the jet vertex tagger (JVT) discriminant [146], are rejected if they have $p_T < 60$ GeV, $|\eta| < 2.5$ and a discriminant value of JVT < 0.5 .

Jets are eligible to be c -tagged or b -tagged if they lie within the ID acceptance of $|\eta| < 2.5$. To identify jets containing c -hadrons, the charm-tagging algorithm DL1r_c, which was optimised for a similar search [187], is used. The DL1r_c algorithm is based on the ATLAS DL1r algorithm [86], which uses a selection of inputs associated with the jet (see section 4.4.6 for a more detailed

description). The algorithm provides three probabilities for a jet to either contain b -hadrons, c -hadrons, or light-flavour hadrons, which are then combined in a similar way to the b -tagging algorithm, but with fine-tuned parameters specifically optimised to identify jets containing c -hadrons. The DL1r algorithm has a 77% identification efficiency for b -jets and 20% and 0.9% misidentification probabilities for c -jets and light-flavour jets, respectively, evaluated in a sample of simulated SM $t\bar{t}$ events. The c -tagging algorithm only considers jets that are not b -tagged (referred to as “ c -tagging with b -veto”), and is tuned to have an efficiency of 20% by requiring that the final DL1r_c discriminant be greater than 1.315, with $f_b = 0.28$. This corresponds to rejection factors of 29 for b -jets and 57 for light-flavour jets, and a τ -lepton misidentification efficiency of 15%. These efficiencies and rejection factors are determined with a sample of simulated $t\bar{t}$ events.

Both the simulated c -jet identification efficiency and misidentification probabilities are corrected to match those measured in data, using an approach similar to that in b -tagging [147, 214, 215]. Between the p_T range of 20–250 GeV, the scale factors are compatible with unity for all jet-flavour corrections. For c -jets, the systematic uncertainties range from 17% for $p_T < 65$ GeV to a few percent for higher- p_T jets (up to 3 TeV), while for b -jets and light-flavour jets, the uncertainties are 5–7% and approximately 13%, respectively, for all p_T ranges below 3 TeV. For b - and c -jets with $p_T > 3$ TeV, the uncertainties on the scale factors are up to 30%.

Electron candidates are constructed from energy deposits in the EM calorimeter matched to an ID track. “Baseline” electron candidates, which are used to veto events with electrons, must pass the Loose likelihood-based identification selection [76] and have $p_T > 4.5$ GeV and $|\eta| < 2.47$. The longitudinal impact parameter of the ID track associated with a “baseline” electron is required to satisfy $|z_0 \sin(\theta)| < 0.5$ mm. Electrons used for regions rich in the dominant SM background processes with prompt electrons (i.e. W + jets and Z + jets) must meet further selection criteria: $p_T > 10$ GeV, Loose isolation [76], $|d_0/\sigma(d_0)| < 5$ (where d_0 is the transverse impact parameter and $\sigma(d_0)$ is its uncertainty), and the Tight likelihood-based identification selection [76]. Electrons satisfying these criteria are referred to as “signal” electrons.

Muon candidates are found by combining track segments from the ID and muon spectrometer (MS), or by extrapolating from the forward sections of the MS, in the region $2.5 < |\eta| < 2.7$. “Baseline” muon candidates, which are used to veto events with muons, are required to pass Medium identification [80] and have $p_T > 4$ GeV, $|\eta| < 2.7$, and $|z_0 \sin(\theta)| < 0.5$ mm. Muons passing the more stringent Loose_VarRad isolation [80], $|d_0/\sigma(d_0)| < 3$, and $p_T > 10$ GeV requirements are referred to as “signal” muons.

Hadronically-decaying τ -leptons are reconstructed from jets with $p_T > 10$ GeV and either $|\eta| < 1.37$ or $1.52 < |\eta| < 2.5$ [82]. In addition, τ -lepton candidates must have $p_T > 20$ GeV, and one or three charged particle tracks with a total electric charge of ± 1 times the electron charge. The τ -leptons are also required to meet the Loose criterion of the neural network-based identification algorithm [216].

The overlap removal sequence proceeds as follows. Jets within $\Delta R = 0.2$ of an electron are rejected unless the jet is b -tagged, in which case the electron is removed if it has $p_T < 100$ GeV (otherwise, the jet is removed regardless of it being b -tagged). Muons are not counted if they are within $\Delta R = (0.04 + 10 \text{ GeV}/p_T(\text{jet}))$ of a jet. This criterion is applied to remove muons from decays of b - and c -hadrons whilst retaining high- p_T muons resulting from the decays of high- p_T massive particles, which tend to be less isolated than muons from decays of low- p_T massive particles. Electrons that share an ID track with a muon are discarded. If a τ -lepton candidate is found to overlap ($\Delta R < 0.2$) with a muon or electron, the τ -lepton is removed. Finally, jets within

10.2. Event reconstruction

$\Delta R = 0.4$ of a τ -lepton are removed. This removal of overlapping objects is performed using the baseline object definitions.

The missing transverse momentum $\mathbf{p}_T^{\text{miss}}$, with magnitude E_T^{miss} , is calculated as the negative vector sum of the transverse momenta of jets, τ -leptons, electrons, and muons [85]. Photons are counted as jets since they are not considered in this search. To account for the underlying event, the missing transverse momentum computation also considers tracks with $p_T > 0.5$ GeV that are associated with the primary vertex but not with any reconstructed particle or jet.

The significance of the missing transverse momentum, denoted by E_T^{miss} Sig., is the ratio of E_T^{miss} to its variance (see section 4.4.5) [217]. The variance is calculated using the p_T resolution of each component in the E_T^{miss} , including any correlations between the components.

Analysis Regions



At the core of designing a search for new physics, one constructs so-called *signal regions* in phase space that are expected to be sensitive to the signature of the new physics process. Counting experiments can then be conducted in each of the signal regions, with possible outcomes being either the observation of an excess of events over the expectation of some null hypothesis, or the absence of such an excess.

It follows, then, that there are two crucial points that will impact the outcome of a search. The first is obviously the optimisation of signal regions to be sensitive to the target process; after all, it is just a *little* difficult to draw meaningful conclusions about the new physics model when no events are observed in a region where none are expected. In this regard, the present analysis utilises simulated event samples of the supersymmetric (SUSY) models of interest (as described in section 10.1) to ensure that the signal regions constructed are indeed sensitive to the topologies expected of these processes. In particular, Recursive Jigsaw Reconstruction (RJR) [184] is used to increase the sensitivity to signal models with compressed mass spectra, and I will begin by presenting an overview of this technique in section 11.1. The collection of signal regions defined for this search will then be described in section 11.2.

The second important aspect of analysis design hinges on our understanding of the null hypothesis: namely, the Standard Model (SM) processes with signatures similar to that of the new physics process. These SM processes will constitute a background to the new physics events in the signal regions. In this analysis, a semi-data-driven approach is adopted to estimate the contributions of major SM backgrounds in each signal region. This involves the construction of *control regions* containing a high purity of these backgrounds. A more detailed description of how control regions are utilised for constraining the background estimations will be provided in Chapter 12, but the gist is that the normalisation of each background is determined by the data before being extrapolated to the signal regions using simulation. As such, control regions are designed to not overlap with the signal regions, to have selection criteria similar to those of the signal regions (to reduce uncertainties from the extrapolation), and to have only a small contamination from signal processes. These regions will be described in section 11.3.

To ensure that the extrapolation of background normalisations into the signal regions is sensible, a suite of *validation regions* that sit in the phase space between the signal and control regions are also defined. These are used to validate the background modelling before any recorded data in the signal regions is examined (unblinding); as such, they are similarly required to be orthogonal to the signal regions, and to contain negligible contamination from any signal model targeted by the search. These regions will be described in section 11.4.




Figure 11.1. The ISR decay tree for compressed scenarios, consisting of a collection of ISR jets, and a sparticle (S) system that decays into visible (V) and invisible (I) products.

11.1 Recursive Jigsaw Reconstruction

RJR is an approach to event reconstruction that aims to resolve kinematic and combinatoric ambiguities in order to facilitate the reconstruction of intermediate resonances, which feature commonly in many beyond the Standard Model (BSM) theories. Pertinent to the present analysis is the pair production of top or charm squarks that each decay into a charm quark and the lightest supersymmetric particle (LSP). In this case, combinatoric ambiguities arise within the collection of jets reconstructed in each event, as these can originate either from the sparticle decay or from initial state radiation (ISR), while the system of LSPs that contribute to the overall missing transverse momentum can constitute a kinematic ambiguity¹. In what follows, I will discuss the essence of RJR as relevant to this work; a more thorough description of the technique, demonstrated for a broader selection of examples, can be found in ref. [184].

Firstly, the standard event and object reconstruction algorithms are carried out as usual. The four-vectors of the reconstructed objects obtained from them form the inputs to the RJR algorithm: one begins by defining the expected event topology, and the idea is to see how well the collection of reconstructed objects fit this assumption. For the compressed signal models targeted in the present search, we impose the ISR decay tree depicted in Figure 11.1, where one expects the proton–proton collision in the centre-of-mass (CM) frame to produce a system of ISR jets recoiling against a sparticle (S) system, the latter of which decays into visible (V) and invisible (I) components. No attempt is made to distinguish between the jets produced in the decay of different parent squarks, nor to disentangle the measured $\mathbf{p}_T^{\text{miss}}$ into two LSPs. As such, there are no kinematic ambiguities, and the problem to be solved reduces to one of classifying each reconstructed jet as either belonging to the ISR or visible systems.

¹This latter case of ambiguity would arise if we wanted to estimate the kinematic properties of each LSP by splitting the measured $\mathbf{p}_T^{\text{miss}}$ into two components. However, as we will promptly discuss, the present search does not attempt to resolve the LSPs individually, so we only need to handle the combinatoric ambiguities.

In general, RJR addresses each case of ambiguity within the event description with a configurable set of physically-motivated assumptions, referred to as the *Jigsaw Rules* (JR). These rules are applied recursively to boost reconstructed object four-vectors from the laboratory frame (being the one in which they were measured) into the rest frame of each subsequent decay expected, continuing until the entirety of the decay tree has been traversed. In the case of the ISR decay tree of Figure 11.1, the *combinatoric mass minimisation* JR was used: in the CM frame, one expects a balance of momentum between the ISR and sparticle systems, and thus the group of objects associated with each should minimise the overall transverse mass in this frame of reference:

$$M_{\text{CM}} = \sqrt{M_{\text{ISR}}^2 + (p_{\text{ISR}}^{\text{CM}})^2} + \sqrt{M_{\text{S}}^2 + (p_{\text{S}}^{\text{CM}})^2}. \quad (11.1)$$

In practice, this amounts to taking the collection of all jets and the missing transverse momentum in the event (which is treated as an additional massless jet) and partitioning them into two mutually exclusive sets J_a and J_b , chosen such that the sum of jet momenta in each group, $|\sum_{j \in J_a} \mathbf{p}_j| + |\sum_{j \in J_b} \mathbf{p}_j|$, is maximised. Given that the missing transverse momentum is known to arise from the sparticle system, its inclusion as a massless jet in this procedure allows for a unique assignment of objects to the ISR and visible frames of the decay tree—the latter group is identified as the one that aligns more closely with the direction of the $\mathbf{p}_T^{\text{miss}}$.

After an event has been analysed using the RJR framework, variables can be derived based on the properties of the objects assigned to each node of the decay tree, with distributions that allow for discrimination between SM background processes and potential signal-like events. In particular, the total transverse momentum of objects assigned to the ISR system, $\mathbf{p}_T^{\text{ISR}}$, can provide an indicator for the presence of massive LSPs. As the LSP mass tends towards that of its parent squark $\tilde{P}$, the decay would occur (in the rest frame of the squark) with no momentum being imparted to the LSP, and the missing transverse momentum generated in the event will be approximately given by:

$$\mathbf{p}_T^{\text{miss}} \sim -\mathbf{p}_T^{\text{ISR}} \times \frac{m(\tilde{\chi}_1^0)}{m(\tilde{P})}. \quad (11.2)$$

Thus, the quantity

$$R_{\text{ISR}} \equiv \frac{|\mathbf{p}_T^{\text{I}} \cdot \hat{\mathbf{p}}_T^{\text{ISR}}|}{|\mathbf{p}_T^{\text{ISR}}|}, \quad (11.3)$$

can be used as an estimator for $|\mathbf{p}_T^{\text{miss}}|/|\mathbf{p}_T^{\text{ISR}}|$. This is calculated using the total momenta (in the CM frame) of objects assigned to the invisible and ISR systems by RJR; for compressed signal-like events with small squark–LSP mass differences, R_{ISR} will exhibit values that tend strongly towards $m(\tilde{\chi}_1^0)/m(\tilde{P})$.

This application of RJR to compressed scenarios is further described in ref. [218], and the technique has been successful in achieving largely orthogonal sensitivity to signal models in this regime as compared to more traditional analysis methods [186]. The full suite of RJR variables found to be useful for the present analysis will be further discussed in section 11.2.2.

11.2 Signal regions

To maximise the discovery potential, the search uses two sets of signal regions to target two drastically different kinematic regions: *High-Mass* and *Compressed*. Each signal region was designed to maximise the expected sensitivity towards a specific signal topology, as characterised by the

11.2. Signal regions




Figure 11.2. Analysis signal regions and the mass splittings that they target relative to one another.

mass difference between the squark and LSP. This is illustrated in Figure 11.2, where any overlap between rectangles in each tier additionally signifies the extent of overlap in the definitions of the corresponding regions.

High-Mass signal regions are designed to be sensitive to leptoquark and SUSY models with large leptoquark or squark masses ($\gtrsim 600$ GeV) and a large difference between the squark and LSP masses ($\gtrsim 200$ GeV). Here, E_T^{miss} -based variables (discussed further in section 11.2.1) are relied upon to discriminate between signal and background events. The Compressed signal regions are sensitive to SUSY models with smaller squark masses and a more compressed spectra; in the context of the present analysis, this refers to models with mass splittings below the approximate threshold of the top quark mass ($\Delta m(\tilde{t}_1, \tilde{\chi}_1^0) \lesssim 175$ GeV). For these signal models, the RJR technique is used to identify the presence of high-momentum jets originating from ISR, yielding an additional suite of variables (to be discussed in section 11.2.2) for the suppression of expected SM background contributions. The High-Mass and Compressed regions are kept orthogonal via a requirement for the leading jet to be c -tagged (High-Mass) or not c -tagged (Compressed).

The signal models targeted in this search are expected² to produce events that have significant E_T^{miss} ; in accordance, a requirement for events to have $E_T^{\text{miss}} > 250$ GeV is imposed. The signal processes do not produce any leptons, so any event containing reconstructed electrons, muons, or τ -leptons is vetoed. All regions require at least one jet with $p_T > 250$ GeV, which is expected to correspond to the c -jet produced in the sparticle decay (High-Mass), or to an energetic ISR jet (Compressed). For the High-Mass (Compressed) signal regions, all jet-related variables are computed using jets with $p_T > 40$ GeV (20 GeV). At least two jets must be present, with at least one being identified as a c -tagged jet. To reject backgrounds originating from top-quark-pair ($t\bar{t}$) decays, events with b -tagged jets are discarded. Finally, the minimum azimuthal angle separation between the $\vec{p}_T^{\text{miss}}$ and up to four p_T -leading jets in each event, $\Delta\phi(\vec{p}_T^{\text{miss}}, \vec{p}_T(j_4))$, is required to be greater than 0.4. This serves to reject multijet backgrounds where significant E_T^{miss} can arise from mis-reconstruction effects; the selection reduces these backgrounds to negligible levels, such that they need not be considered further in this search.

²In the case of compressed signals, this is made possible by ISR activity recoiling against the sparticle system.

Table 11.1. Requirements of the SR-HM1–3 High-Mass regions, which are aimed at signal models with high squark mass ($\gtrsim 600$ GeV), and SR-HM-Disc which is aimed at a broad range of models.

Variable	SR-HM1	SR-HM2	SR-HM3	SR-HM-Disc
Number of leptons	= 0			
Jet p_T [GeV]	> 40			
Number of b -tagged jets	= 0			
$p_T(j_1)$ [GeV]	> 250			
$\Delta\varphi(\mathbf{p}_T^{\text{miss}}, \vec{p}_T(j_4))$	> 0.4			
Number of c -tagged jets	≥ 2			
Leading jet is c -tagged	yes			
$m_T^{\text{min}}(c)$ [GeV]	> 300			
$p_T(j_2)$ [GeV]	> 200	> 150		
m_{cc} [GeV]	> 200	200–400	400–600	> 200
E_T^{miss} Sig.	> 22	15–22		> 15

11.2.1 High-Mass

Three non-overlapping High-Mass signal regions, named SR-HM1, SR-HM2 and SR-HM3, are defined as shown in Table 11.1. The SUSY scenario with $\mathcal{B}(\tilde{t}_1/\tilde{c}_1 \rightarrow c + \tilde{\chi}_1^0) = 1$ and $m(\tilde{t}_1, \tilde{\chi}_1^0) = (1000, 1)$ GeV is used to optimise the selection for SR-HM1, while the scenario with $m(\tilde{t}_1, \tilde{\chi}_1^0) = (700, 400)$ GeV is used to optimise the selection for SR-HM2 and SR-HM3. All signal models are expected to produce two c -tagged jets, so all three regions require the presence of at least two c -tagged jets. Additionally, the signal models that are used to optimise the High-Mass signal regions are expected to have highly energetic c -tagged jets, due to the large difference between the mass of the squark and neutralino. This motivates the requirement of the leading jet to be c -tagged and that the subleading jet satisfies $p_T(j_2) > 150$ GeV.

Standard Model background contributions with two c -tagged jets originating from real charm quarks are expected to be produced by processes such as $g \rightarrow c\bar{c}$, where g is a gluon. As such, the invariant mass of the two highest- p_T c -tagged jets in these events, denoted m_{cc} , is expected to have a falling spectrum. In contrast, the c -tagged jets in signal events are produced in the decay of the squarks, and the m_{cc} will tend to be shifted depending on the squark mass. Thus, SR-HM1, SR-HM2, and SR-HM-3 require m_{cc} to be at least 200 GeV, with different m_{cc} requirements made in SR-HM2 and SR-HM3 to target signal models with $\Delta m(\tilde{t}_1, \tilde{\chi}_1^0) \approx 300$ GeV and ≈ 400 GeV, respectively. An upper bound on m_{cc} was added to SR-HM3 to reject V + jets backgrounds where m_{cc} falls less steeply than in the signal samples above $m_{cc} \approx 600$ GeV.

The transverse mass between a c -tagged jet and the $\mathbf{p}_T^{\text{miss}}$,

$$m_T(c) = \sqrt{2p_T(c) \times E_T^{\text{miss}} \times (1 - \cos \Delta\varphi(c, \mathbf{p}_T^{\text{miss}}))}, \quad (11.4)$$

is calculated for each c -tagged jet in the event. The minimum value of this, denoted as $m_T^{\text{min}}(c)$, is required to be at least 250 GeV. This selection rejects $W(\rightarrow \tau\nu) + \text{jets}$ backgrounds that contain hadronically-decaying τ -leptons misidentified as c -tagged jets. Finally, various E_T^{miss} Sig. requirements are imposed to target signal models that are expected to have different kinematics.

11.2. Signal regions

Table 11.2. Requirements of the Compressed signal regions that target signal models with squark–neutralino mass differences between 20 GeV and 175 GeV.

Variable	SR-Comp1	SR-Comp2	SR-Comp3	SR-Comp-1c
Number of leptons	= 0			
Jet p_T [GeV]	> 20			
Number of b -tagged jets	= 0			
$p_T(j_1)$ [GeV]	> 250			
$\Delta\varphi(\vec{p}_T^{\text{miss}}, \vec{p}_T(j_4))$	> 0.4			
Number of c -tagged jets	≥ 2			= 1
Leading jet is c -tagged	no			
$m_T^{\text{min}}(c)$ [GeV]	–	> 80	> 120	< 250
m_{cc} [GeV]	–		> 100	–
p_T^{CM} [GeV]	–			< 10
m_V [GeV]	–			< 80
$\Delta\varphi(\vec{p}_T^{\text{miss}}, j_1)$	–			> 2
$N_{c\text{-jets}}^S$	≥ 2			= 1
p_T^{ISR} [GeV]	> 550	> 500	> 400	> 700
R_{ISR}	> 0.85	0.75–0.85	> 0.65	0.9–1.0

In addition to SR-HM1, SR-HM2, and SR-HM3, a signal region SR-HM-Disc with less stringent requirements is defined. This is only used, along with SR-HM1, to set model-independent cross-section limits on any BSM contributions. SR-HM-Disc has requirements similar to those in the other SR-HM regions, but it has a less stringent requirement on m_{cc} ($m_{cc} > 200$ GeV) and lacks the upper bound on E_T^{miss} Sig. applied in SR-HM2 and SR-HM3 (see Table 11.1) in order to facilitate interpretations for more general BSM models.

11.2.2 Compressed

Four Compressed signal regions, named SR-Comp1, SR-Comp2, SR-Comp3 and SR-Comp-1c, are defined as shown in Table 11.2. These regions are constructed using the RJR method assuming an ISR event topology (Figure 11.1), which, as discussed in section 11.1, is effective for increasing the sensitivity to signals with compressed spectra [219]. The Compressed regions target signal models with mass splittings of $\Delta m(\tilde{t}_1, \tilde{\chi}_1^0) \lesssim 175$ GeV, and are kept orthogonal to the High-Mass signal regions by requiring that the leading jet is not c -tagged (which is usually a valid assumption, since it is expected to arise from ISR).

The most compressed SUSY scenarios, with $\Delta m(\tilde{t}_1, \tilde{\chi}_1^0) \approx 20\text{--}50$ GeV, are targeted with SR-Comp1 and SR-Comp-1c. SUSY scenarios with $\Delta m(\tilde{t}_1, \tilde{\chi}_1^0) \approx 80$ GeV are targeted with SR-Comp2, which is designed in conjunction with SR-Comp1 to be binned in the R_{ISR} variable (discussed further below). Finally, a more general signal region with looser selection criteria, SR-Comp3, is designed to target a wide range of larger mass splittings $50 \text{ GeV} \lesssim \Delta m(\tilde{t}_1, \tilde{\chi}_1^0) \lesssim 175$ GeV. There is some overlap between the Compressed signal regions: SR-Comp3 overlaps with SR-Comp1

and SR-Comp2, but SR-Comp1, SR-Comp2, and SR-Comp-1c are all orthogonal to one another, and SR-Comp-1c is orthogonal to SR-Comp3.

At least two c -tagged jets are required in SR-Comp1, complemented by a selection that requires exactly one c -tagged jet in SR-Comp-1c to increase signal contributions in the low- p_T regime where the c -tagging efficiency decreases. SR-Comp-1c also has a requirement of $\Delta\varphi(\mathbf{p}_T^{\text{miss}}, j_1) > 2$ on the azimuthal angle separation between the leading jet and the $\mathbf{p}_T^{\text{miss}}$. This has little impact on its expected sensitivity, instead serving to reduce the computational workload when processing events in this region (which is of a greater volume, owing to its looser c -tagging selection). In SR-Comp3, selections on $m_T^{\text{min}}(c)$ and m_{cc} improve background rejection, with the latter serving to reduce background contributions (e.g. $g \rightarrow c\bar{c}$). In SR-Comp1 and SR-Comp2, selections on these variables are either relaxed or removed in order to reduce the statistical uncertainty in their expected event yield.

In the presence of an ISR system, which consists of one or more jets produced by initial-state radiation, the $\tilde{\ell}_1\tilde{\ell}_1$ or $\tilde{c}_1\tilde{c}_1$ system is boosted in the transverse plane. Kinematic variables are then defined appropriately for the assignment of objects to either the ISR system or the sparticle (squarks plus LSPs) system. This method is equivalent to grouping the event objects according to the axis of maximum back-to-back p_T in the event's centre-of-mass (CM) frame, where the p_T of all objects sums vectorially to zero. The RJR variables that were found to increase the sensitivity to the targeted signal models, and consequently used in the design of the Compressed signal regions, include:

- R_{ISR} , the projection of the invisible system's p_T direction vector onto the ISR system's p_T direction vector. The ISR and sparticle systems tend to be approximately back-to-back in very compressed SUSY events, resulting in R_{ISR} values close to unity, while SM backgrounds tend to populate a wider range of R_{ISR} values. The selections chosen in each of the Compressed signal regions reflect this, with the regions targeting more compressed topologies requiring a tighter selection.
- p_T^{ISR} , the vector sum of the transverse momenta of jets in the ISR frame. For SM backgrounds, p_T^{ISR} has a smoothly falling distribution, while for compressed SUSY signal models, a peak at higher p_T^{ISR} values is seen, with the location of the peak typically being related to the mass of the parent SUSY particle. The cut values chosen in each of the signal regions were optimised based on considerations of both the expected sensitivity to signal benchmarks as well as the integrated luminosity of the Run 2 dataset (looser selections were imposed where statistical limitations were expected).
- $N_{c\text{-jets}}^{\text{S}}$, the number of c -tagged jets assigned to the sparticle (S) frame of the ISR decay tree. For SM processes, RJR is equally likely to assign c -tagged jets to either of the visible frames, while for compressed SUSY signal events whose topology is in agreement with that of the ISR decay tree imposed, a more accurate assignment of objects is achieved. As such, a requirement of at least two c -tagged jets in the S frame is imposed on the Compressed signal regions, except SR-Comp-1c which requires the single c -tagged jet to be assigned.

To suppress the increased SM contributions from loosening the selection on the c -jet multiplicity, more stringent selections are placed on R_{ISR} and p_T^{ISR} in SR-Comp-1c, and additional selections are also imposed on p_T^{CM} , the transverse momentum of the CM frame as evaluated in the laboratory frame, and m_V , the invariant mass of objects assigned to the visible (V) system.

Model-independent cross-section limits are set in the signal region with the least stringent requirement, i.e. SR-Comp3.

11.3. Background estimation and the control regions




Figure 11.3. Phase-space partitioning between the signal, control, and validation regions for the (a) High-Mass signal regions, (b) Compressed signal regions SR-Comp1, SR-Comp2 and SR-Comp3, and (c) SR-Comp-1c. Control regions differ from the signal and validation regions by the electron and muon multiplicities (exactly one electron or muon in the W +jets control regions and two electrons or muons in the Z +jets regions) while the validation regions veto events with leptons as in the signal regions. The asterisk in “SR-Comp*” is a wildcard character to denote all the Compressed signal regions that require two c -tagged jets.

11.3 Background estimation and the control regions

The dominant backgrounds in all the signal regions are from processes with W or Z bosons produced in association with jets, with the bosons decaying as $W \rightarrow \ell\nu$ or $Z \rightarrow \nu\nu$ and producing real missing transverse momentum. (The charged lepton produced in the W boson decay corresponds either to a hadronically-decaying τ -lepton, or to an electron or muon that is misreconstructed or falls outside of the detector acceptance.) The Z +jets background has similar contributions from $Z + c\bar{c}$, $Z + cj$, and $Z + jj$ (where j indicates a non-charm quark) in all signal regions except SR-Comp-1c, where the $Z + jj$ background dominates. The W +jets background has larger contributions from the $W + cj$ and $W + jj$ components than from $W + c\bar{c}$. To reduce the systematic uncertainties, these backgrounds are estimated by extrapolating production-rate correction factors from the control regions (CRs) to the signal regions using simulated event samples. The extrapolation is performed in lepton multiplicity, as illustrated in Figure 11.3, while other requirements in the CRs are made to be as close as possible to those in the signal regions.

The normalisation of the Z +jets and W +jets backgrounds is performed *in situ* with the two-lepton (CRZ) and one-lepton (CRW) control regions, respectively. Three groups of CRs are designed for each background contribution: one set of CRs for the High-Mass signal regions, one set for the Compressed signal regions with two c -tagged jets, and one set for the Compressed signal region with one c -tagged jet. Contamination from signal processes was checked and found to be negligible in the defined CRs. Background processes that contribute less significantly are not normalised using control regions, but using calculated cross-sections. These include processes involving top quarks and multiple vector bosons, and will collectively be referred to as “Other” processes hereafter.

11.3.1 Z + jets control regions

The data for the Z +jets control regions (CRZ) were collected using the lowest unprescaled single-electron and single-muon triggers. These triggers have a near-constant efficiency for leptons above a p_T ($p_T(\ell_1)$) threshold of 27 GeV. To ensure a high-purity sample of Z + jets events, the

Table 11.3. Requirements of the $Z + \text{jets}$ control regions for SR-HM1, SR-HM2, and SR-HM3 (CRZ-HM), SR-Comp1, SR-Comp2, and SR-Comp3 (CRZ-Comp), and SR-Comp-1c (CRZ-Comp-1c). Variables that are calculated using E_T^{miss} (including E_T^{miss} Sig.) use an E_T^{miss} variant that includes contributions from electrons/muons that are treated as invisible, and are indicated by a prime (').

Variable	CRZ-HM	CRZ-Comp	CRZ-Comp-1c
Used for	SR-HM	SR-Comp	SR-Comp-1c
$p_{\text{T}}(j_1)$	$> 250 \text{ GeV}$		$> 350 \text{ GeV}$
Number of c -tagged jets	≥ 2		$= 1$
Leading jet is c -tagged	yes	no	
Number of signal muons/electrons	$= 2$		
Passed single-electron/muon triggers	yes		
$p_{\text{T}}(\ell_1) \text{ [GeV]}$	> 30		
$m_{\ell\ell} \text{ [GeV]}$	76–106		
$E_{\text{T}}^{\text{miss}} \text{ [GeV]}$	< 200		
$E_{\text{T}}^{\text{miss}'} \text{ [GeV]}$	> 200		
$m_{cc} \text{ [GeV]}$	> 150	–	
$m_{\text{T}}^{\text{min}'}(c) \text{ [GeV]}$	> 150	–	< 250
$E_{\text{T}}^{\text{miss}'} \text{ Sig.}$	> 10	–	
$N_{c\text{-jets}}^{S'}$	–	≥ 1	$= 1$
R'_{ISR}	–	> 0.75	0.9–1.0
$p_{\text{T}}^{\text{ISR}'} \text{ [GeV]}$	–		> 550
$\Delta\varphi(\boldsymbol{p}_{\text{T}}^{\text{miss}'}, j_1)$	–		> 2

two signal leptons are required to be of the same flavour and oppositely-charged (i.e. e^+e^- or $\mu^+\mu^-$), and their invariant mass, denoted $m_{\ell\ell}$, is required to be consistent with the mass of the Z boson ($76 \text{ GeV} < m_{\ell\ell} < 106 \text{ GeV}$). Further details on the CRZ selections are shown in Table 11.3.

For all of the $Z + \text{jets}$ control regions, E_T^{miss} and all other variables derived from it³ treat the leptons as invisible by subtracting their contribution from the $\mathbf{p}_T^{\text{miss}}$ computation; these variables are indicated with a prime ('). The treatment of the two leptons as invisible particles emulates the dominant contribution of $Z \rightarrow \nu\nu$ decays in the signal regions. In the CRZ, the real E_T^{miss} is required to be less than 200 GeV to ensure that the bulk of the $E_T^{\text{miss}'}$ comes from the addition of the leptons; in conjunction with the selection on $m_{\ell\ell}$, this serves to further reduce the contributions from $t\bar{t}$ and diboson events that can contain real invisible particles.

To avoid significant extrapolation uncertainties between the High-Mass signal and control regions related to jet flavour composition, the requirements on the number of c -tagged jets ($= 1$ or ≥ 2) and on $m_T^{\text{min}'}(c)$ are kept as close as possible in CRZ-HM to those in the High-mass signal

³These include the conventional variables such as the E_T^{miss} Sig., and also all of the RJR variables, which were recomputed with the $\mathbf{p}_T^{\text{miss}'}$ assumption.

11.3. Background estimation and the control regions




Figure 11.4. Distributions showing the level of agreement between data (points) and the SM expectation (stacked histograms, after simultaneously fitting the $Z + \text{jets}$ and $W + \text{jets}$ backgrounds) in the $Z + \text{jets}$ control regions for some variables which have different selection criteria in the signal regions: (a) E_T^{miss} Sig. in CRZ-HM, (b) m_{cc} in CRZ-HM, (c) $p_T^{\text{ISR}'}$ in CRZ-Comp, and (d) $p_T^{\text{ISR}'}$ in CRZ-Comp-1c. The hatched uncertainty band around the total SM expectation includes theory-based and detector-related systematic uncertainties and MC statistical uncertainties. Processes with top quarks and multiple vector bosons are included in “Other”. The $W + \text{jets}$ background contribution is negligible and not visible in the plots. The right-most bin in each histogram does not include the overflow entry but the x-axis range is chosen to include all observed data.

regions while also ensuring that the statistical uncertainty in the $Z + \text{jets}$ production rate remains sufficiently small. In the Compressed CRZ, selections on the $p_T^{\text{ISR}'}$ or $N_{c\text{-jets}}^S$ are relaxed in order to reduce the statistical uncertainty on the events. However, care was taken to ensure that the extrapolation uncertainties remained minimal. Distributions of selected kinematic observables are shown in Figure 11.4 to illustrate the accuracy of the simulation in modelling the data.

11.3.2 $W + \text{jets}$ control regions

The $W + \text{jets}$ control regions (CRW) require exactly one signal electron or muon. To keep the kinematic selection in the CRWs and the corresponding signal regions similar, all the CRWs require events to pass the E_T^{miss} trigger and to have $E_T^{\text{miss}} > 250$ GeV.

Table 11.4. Requirements of the W +jets control regions for SR-HM1, SR-HM2, and SR-HM3 (CRW-HM), SR-Comp1, SR-Comp2, and SR-Comp3 (CRW-Comp), and SR-Comp-1c (CRW-Comp-1c).

Variable	CRW-HM	CRW-Comp	CRW-Comp-1c
Used for	SR-HM	SR-Comp	SR-Comp-1c
$p_T(j_1)$ [GeV]	> 250		> 350
Number of c -tagged jets	≥ 2		$= 1$
Number of signal muons/electrons	$= 1$		
Passed E_T^{miss} trigger	yes		
E_T^{miss} [GeV]	> 250		
m_{cc} [GeV]	> 200	–	
$m_T^{\text{min}}(c)$ [GeV]	> 200	> 120	< 250
E_T^{miss} Sig.	> 14	–	
Leading jet is c -tagged	yes	no	
$N_{c\text{-jets}}^S$	–	≥ 1	$= 1$
p_T^{ISR} [GeV]	–	> 400	> 550
R_{ISR}	–	> 0.65	$0.9\text{--}1.0$
τ veto	–	yes	–
$\Delta\varphi(\mathbf{p}_T^{\text{miss}}, j_1)$	–		> 2

Note that the selection on $m_T^{\text{min}}(c)$ serves a different purpose in the control regions than in the signal regions, where the cut is used to reject the $W(\rightarrow \tau\nu) + \text{jets}$ background containing a hadronically-decaying τ -lepton that has been misidentified as a c -jet. The control regions do not receive this $W + \text{jets}$ contribution due to the requirement for a reconstructed lepton to be present; however, they can contain $t\bar{t}$ events in which one of the top quarks has decayed as $t \rightarrow bW \rightarrow b\ell\nu$, while the other via $t \rightarrow bW \rightarrow b\tau\nu$ with the τ -lepton being c -tagged. The $m_T^{\text{min}}(c)$ selection in the control regions serves to reject this $t\bar{t}$ contribution and increase the purity for $W + \text{jets}$ events.

Still, to minimise the extent of extrapolation, the requirement on $m_T^{\text{min}}(c)$ in each control region has been chosen to match the selection made in their corresponding signal regions. Because of this, CRW-Comp has a higher contamination from $t\bar{t}$ than is seen in CRW-HM, due to the former's lower $m_T^{\text{min}}(c)$ requirement. Thus, to maintain a high purity of $W + \text{jets}$ events and to avoid an increase in the contamination from $t\bar{t}$ processes, CRW-Comp also includes a veto on events that include a τ -lepton. CRW-Comp-1c contains events that have different kinematics due to its one- c -tag and high p_T^{ISR} requirements, and does not require additional $t\bar{t}$ rejection from an explicit τ -lepton veto.

A detailed description of all the CRW selections is presented in Table 11.4, and distributions of the main kinematic variables used for the extrapolation from each CRW to their respective signal regions are shown in Figure 11.5.

11.4. Validation of background estimates




Figure 11.5. Distributions showing the level of agreement between the data (points) and the SM expectation (stacked histograms, after simultaneously fitting the Z+jets and W+jets backgrounds) in the W+jets control regions for some variables which have different selection criteria in the signal regions: (a) $E_T^{\text{miss}} \text{ Sig.}$ in CRW-HM, (b) p_T of leading lepton in CRW-HM, (c) p_T^{ISR} in CRW-Comp, and (d) p_T^{ISR} in CRW-Comp-1c. The hatched uncertainty band around the total SM expectation includes theory-based and detector-related systematic uncertainties and MC statistical uncertainties. Processes with top quarks and multiple vector bosons are included in “Other”. The Z+jets background contribution is negligible and not visible in the plots. The right-most bin in each histogram does not include the overflow entry but the x-axis range is chosen to include all observed data.

11.4 Validation of background estimates

Dedicated validation regions are used to test the accuracy of the procedure used to extrapolate the backgrounds from the control regions to the signal regions before unblinding the data in the signal regions. The variable that is used to extrapolate between the control and signal regions is the electron and muon multiplicity. Therefore, to remain close to the signal regions, the validation regions apply the same veto on events with leptons, and are designed to have a similar background composition as that of the signal regions (which consist mainly of Z + jets and W + jets). The validation regions are kept orthogonal to the signal and control regions, and are ensured to contain sufficient event yields in order to facilitate an adequate test of the

Table 11.5. Requirements of the validation regions.

Variable	VR-HM	VR-Comp	VR-Comp-1c
Validates backgrounds in	SR-HM	SR-Comp	SR-Comp-1c
Number of c -tagged jets	≥ 2		$= 1$
Passed E_T^{miss} trigger	yes		
E_T^{miss} [GeV]	> 250		
m_{cc} [GeV]	> 200	–	
$m_T^{\text{min}}(c)$ [GeV]	> 200	> 80	> 100
Leading jet is c -tagged	yes	no	
$N_{c\text{-jets}}^S$	–	≥ 1	$= 1$
p_T^{ISR} [GeV]	–	< 400	
R_{ISR}	–	> 0.65	$0.85\text{--}1.05$
τ veto	yes		
$\Delta\phi(\mathbf{p}_T^{\text{miss}}, j_1)$	–	> 2	
p_T^{CM} [GeV]	–	< 10	
m_V [GeV]	–	< 80	

extrapolation procedure with reasonable accuracy. They are also designed to avoid large contamination from any signal model considered in this search, with the requirement that less than 20% of the expected event yield is attributed to the contamination from any signal model not already excluded by previous ATLAS searches.

To validate the extrapolation of the backgrounds from the control regions CRW-HM and CRZ-HM to the High-Mass signal regions SR-HM1, SR-HM2 and SR-HM3, a single validation region VR-HM is used, with an expected Z +jets and W +jets contribution of 45% and 26%, respectively. The requirements for VR-HM are kept the same as in the signal region, with the exception of a reversal of the m_{cc} requirement, the removal of any selection on the E_T^{miss} Sig. and $p_T(j_2)$ variables, and a less stringent requirement of $m_T^{\text{min}}(c) > 200$ GeV (as compared to $m_T^{\text{min}}(c) > 300$ GeV in the signal regions). These choices were made to reduce the overall signal contamination in VR-HM and to maintain orthogonality to the signal regions.

The normalisation procedure for Compressed regions that require at least two charm-tagged jets, i.e. SR-Comp1, SR-Comp2, and SR-Comp3, is validated in one region, VR-Comp. This region is orthogonal to the signal regions via a reversed p_T^{ISR} requirement ($p_T^{\text{ISR}} < 400$ GeV), which also serves to reduce the contamination from signal events. The other requirements, including the veto on events with leading c -tagged jets that keeps the Compressed regions orthogonal to the High-Mass regions, are kept the same as in the signal regions, with the following four exceptions:

- The c -tagged jet multiplicity in the sparticle frame, $N_{c\text{-jets}}^S$, is loosened to be greater than or equal to one (rather than two in the signal regions). This increases the overall event yield expected in VR-Comp, with little impact on the jet flavour composition in the region.

11.4. Validation of background estimates

- The requirement on R_{ISR} in VR-Comp ($R_{\text{ISR}} > 0.65$) is chosen to be the same as in SR-Comp3. This is less stringent than the selection in SR-Comp1 ($R_{\text{ISR}} > 0.85$) and SR-Comp2 ($0.75 < R_{\text{ISR}} < 0.85$).
- The $m_{\text{T}}^{\text{min}}(c)$ variable is required to be greater than 80 GeV in VR-Comp, despite there being no requirement on $m_{\text{T}}^{\text{min}}(c)$ in SR-Comp1. This choice was made to keep VR-Comp aligned with the other two Compressed signal regions, with SR-Comp2 also requiring a minimum of 80 GeV, and SR-Comp3 requiring $m_{\text{T}}^{\text{min}}(c) > 120$ GeV.
- No requirement is made on m_{cc} in VR-Comp, as is the case in SR-Comp1 and SR-Comp2; this differs from the $m_{\text{cc}} > 100$ GeV requirement in SR-Comp3.

The one- c -tag validation region, VR-Comp-1c, is identical to SR-Comp-1c with the exception of an inverted $p_{\text{T}}^{\text{ISR}}$ requirement (which, as before, serves to orthogonalise the two regions), along with a slight widening of the R_{ISR} window from $0.9 < R_{\text{ISR}} < 1.0$ in SR-Comp-1c to $0.85 < R_{\text{ISR}} < 1.05$ in VR-Comp-1c. The looser R_{ISR} selection has minimal impact on the extrapolation, but comes with the benefit of roughly doubling the expected event yield in the validation region.

The full set of selections for all validation regions is shown in Table 11.5, while distributions of the main kinematic variables that differ between the validation regions and the signal regions are shown in Figure 11.6.




Figure 11.6. Distributions, in various validation regions, showing the level of agreement between the data (points) and the SM expectation (stacked histograms, after simultaneously fitting the Z+jets and W+jets backgrounds) for variables which have different selection criteria in the signal regions: (a) E_T^{miss} Sig. in VR-HM, (b) $m_T^{\text{min}}(c)$ in VR-HM, (c) R_{ISR} in VR-Comp, and (d) R_{ISR} in VR-Comp-1c. The hatched uncertainty band around the total SM expectation includes theory-based and detector-related systematic uncertainties and MC statistical uncertainties. Processes with top quarks and multiple vector bosons are included in “Other”. The right-most bin in each histogram does not include the overflow entry but the x -axis range is chosen to include all observed data.

Results

 HE fitting procedures used in the analysis are implemented using the pyHF [220, 221] framework, which is a purely-pythonic implementation of the HistFactory tool for building parametrised probability density functions (PDFs) [222]. This framework performs binned profile likelihood fits using the number of expected (i.e. Monte Carlo (MC) simulated) events for each process in each control and signal region whilst accounting for any statistical and systematic uncertainties associated with these, along with the number of observed (i.e. data) events in each region. The building and steering of the template PDF fits is configured using the cabinetry library [223].

The likelihood function used in the fits corresponds to the product of the Poisson distributions of the number of events (both MC and data) in each of the signal (P_{SR}) and control (P_{CR}) regions, with additional Gaussian factors (C_{syst}) for modelling systematic uncertainty variations:

$$\begin{aligned} L(n, \theta^0 | \mu_{\text{sig}}, b, \theta) &= P_{\text{SR}} \times P_{\text{CR}} \times C_{\text{syst}} \\ &= P(n_{\text{SR}} | \lambda_{\text{SR}}(\mu_{\text{sig}}, b, \theta)) \times \prod_{i \in \text{CR}} P(n_i | \lambda_i(\mu_{\text{sig}}, b, \theta)) \times C_{\text{syst}}(\theta^0, \theta), \end{aligned} \quad (12.1)$$

where the n_i (λ_i) represent the observed (expected) event counts in region i , with μ_{sig} (the “signal strength” parameter) and b determining the normalisation of the expected signal and background contributions, respectively. Additional nuisance parameters are denoted by θ .

The background processes are divided into two categories for the fitting purposes. For the processes to be estimated in a semi-data-driven fashion (i.e. Z + jets and W + jets), the shape of the PDF is taken from MC simulations while the normalisation is extracted from the fit to the data. This normalisation is heavily influenced by the various control regions defined for these backgrounds. Processes that do not have an associated control region (namely, the “Other” backgrounds comprising mostly of $t\bar{t}$ and diboson processes) are not normalised in the fit; rather, their normalisation is determined from theoretical calculations. However, when performing the fit, the normalisation of these “Other” backgrounds are allowed to float within the range of their respective uncertainties.

Three different fitting procedures can be performed, depending on the set of regions that are considered in the likelihood function and on whether any parameters are fixed. Each procedure serves a different purpose, as follows:

- **Background-only fit:** This is used to extract normalisation parameters for each of the backgrounds that have dedicated control regions. The purpose of this fit is to produce an

12.1. Systematic uncertainties

estimate of the total background expectation in the signal (and validation) regions without making any assumptions on the signal model under consideration ($\mu_{\text{sig}} = 0$).

The background-only fit is performed by fitting the likelihood function (eq. (12.1)) after removing the parameters from the signal and validation regions. In this way, the normalisation parameters are calculated using only the expected and observed event yields in the dedicated control regions. These normalisation parameters are then introduced back into the likelihood function to estimate the background contributions in the signal and validation regions.

- **Model-dependent (or exclusion) fit:** This is a frequentist fit performed to study a specific signal model. The likelihood function includes all of the control and signal regions, along with normalisation factors for the relevant backgrounds and the signal model under investigation (μ_{sig}). If no significant excess of events is seen in the signal regions after the background expectations have been normalised (using the results of background-only fits), one can determine whether the signal model is excluded or not by performing a hypothesis test with the assumption that $\mu_{\text{sig}} = 1$. Alternatively, if an excess is observed, one can measure the signal strength by allowing μ_{sig} to float during the fit.

Usually, the exclusion fit is performed for each signal on a grid of signal models. In the case of this analysis, the grid corresponds to signals with different $m(\tilde{t}_1)$ and $m(\tilde{\chi}_1^0)$ assumptions. The result of these exclusion fits can then be used to construct an exclusion contour in the $m(\tilde{t}_1, \tilde{\chi}_1^0)$ mass plane, based on whether the signal model is considered to be excluded or not at the 95% confidence level (CL) (corresponding to a CL_S value below 0.05). In this context, one can interpret the CL_S value as a probability for rejecting the background-only hypothesis.

- **Model-independent (or discovery) fit:** This fit provides a model-independent limit on the additional number of events in the signal regions, beyond the expected number of background events, whilst remaining compatible with the observed number of events. This enables the results of the analysis to be interpreted in the context of any generic new physics model which could be expected to contribute to the phase space defined by the signal regions.

The procedure for this fit proceeds generally in the same manner as the model-dependent fit, with the difference being that the “signal model” is an artificial one that simply predicts one event in each SR (with no associated uncertainty). No signal contamination is taken into account in the control regions. The number of observed events in the signal regions is added to the fit to enable the signal strength (of the dummy signal) to be investigated. The value of the signal strength where the CL_S value falls below the 0.05 threshold can then be converted into an upper limit on the visible cross-section of a generic new physics production mechanism, σ_{vis} , which will correspond to the product of the cross-section for the process with the selection efficiency and the detector acceptance.

12.1 Systematic uncertainties

The effect of systematic uncertainties is taken into account for the limited accuracy of the background and signal predictions. To properly treat correlations of uncertainties across the signal and control regions, the uncertainties are separated by source. There are uncertainties for the

simulated detector response, integrated luminosity [189], number of simulated events, predicted cross-sections, parton distribution functions, underlying event and minimum-bias pp collisions, parton hadronisation models, matrix element and parton shower calculations, and experimental effects. A discussion of the most relevant sources of uncertainty affecting the search is presented below.

The uncertainty in the measured c -tagging efficiency is one of the dominant sources of uncertainty because all the signal and control regions require events with one or more c -tagged jets. The simulated event samples are reweighted so that the predicted and measured identification efficiencies of c -tagged jets are identical and so the uncertainties are propagated. The uncertainty in the c -tagging efficiency also depends on the jet p_T and whether the tagged jet is initiated by a b -quark, c -quark, or light quark [147, 214, 215].

The uncertainties in the predicted jet energies are driven by the accuracy of the measured jet energies and resolution [73]. There is an uncertainty in the contribution from jets that originate from the overlapping pp interactions (i.e. pile-up) but that are mistakenly counted as coming from the hard scatter [146]. This includes the efficiency of identifying jets originating from secondary pp interactions and the accuracy of simulated minimum-bias interactions.

The theoretical uncertainties for the W + jets and Z + jets backgrounds are estimated with SHERPA. Uncertainties due to inaccuracies in the matrix element calculation and parton showering are estimated by varying the generator's parameters, with the envelope of the associated predicted distributions taken as the uncertainty. The matrix element matching scale is varied between 15 GeV and 30 GeV from its nominal value of 20 GeV. The renormalisation and factorisation scales are independently varied by factors of 0.5 and 2. The impact of the PDF uncertainties is taken into account in addition to the generator uncertainties.

The “Other” backgrounds, such as top-quark pairs produced in association with jets [224] and vector-boson pairs produced in association with jets, have a minor effect on the total background estimate. A normalisation uncertainty of 25% is assigned and treated as fully correlated across regions. This uncertainty estimate is meant to cover cross-section normalisation and modelling uncertainties and does not contribute significantly to the total uncertainty. Treating the uncertainty to be uncorrelated between regions was found to increase the total expected background contribution by at most one standard deviation in SR-Comp2 and SR-Comp3 with a negligible effect on the final result.

For the SUSY and LQ signal predictions, theoretical uncertainties are taken as an envelope of the results obtained when varying the factorisation and renormalisation scales independently by factors of 0.5 and 2. The Var3c parameter (which corresponds to the strong coupling parameter, α_s , used for ISR) of the A14 underlying-event tune is also varied. The contribution from these uncertainties is largest in the compressed regions, where it reaches 20%.

12.2 Results and interpretation

A statistical analysis of the data is performed with a simultaneous likelihood fit [220, 223, 225] to the number of signal and background events in the control and signal regions; effects of systematic uncertainties are included through nuisance parameters that change the expected contributions across regions in a correlated manner. The likelihood is calculated by combining Poisson likelihood terms for the signal, validation, and control regions with a Gaussian probability distribution for each systematic uncertainty. The W + jets and Z + jets background normalisations

12.2. Results and interpretation

Table 12.1. Post-fit event yields for the High-Mass control, validation, and signal regions. The Z + jets and W + jets backgrounds are normalised by fits within the control regions. Processes with top quarks and multiple vector bosons are included in “Other”. Pre-fit event yields for a representative signal point for each of the signal regions are also shown.

	CRW-HM		CRZ-HM		VR-HM
Z +jets	0.66 $\pm$ 0.12		74 $\pm$ 9		31 $\pm$ 5
W +jets	90 $\pm$ 11		0.030 $\pm$ 0.009		22 $\pm$ 4
Other	15 $\pm$ 4		6.6 $\pm$ 1.8		17 $\pm$ 5
Total SM	106 $\pm$ 10		81 $\pm$ 9		70 $\pm$ 8
Observed	106		81		65

	SR-HM1	SR-HM2	SR-HM3	SR-HM-Disc
Z +jets	5.2 $\pm$ 1.3	12.3 $\pm$ 3.4	14.7 $\pm$ 2.7	58 $\pm$ 9
W +jets	3.2 $\pm$ 0.7	9.8 $\pm$ 1.7	12.5 $\pm$ 2.1	43 $\pm$ 6
Other	0.88 $\pm$ 0.33	1.6 $\pm$ 0.6	2.2 $\pm$ 0.7	9.0 $\pm$ 2.5
Total SM	9.3 $\pm$ 1.6	24 $\pm$ 4	29 $\pm$ 4	110 $\pm$ 12
Observed	13	31	27	133
$m(\tilde{t}_1, \tilde{\chi}_1^0) = (1000, 1)$ GeV	6.6 $\pm$ 1.0			12.7 $\pm$ 2.1
$m(\tilde{t}_1, \tilde{\chi}_1^0) = (750, 450)$ GeV		17.3 $\pm$ 2.6	12.6 $\pm$ 2.1	41 $\pm$ 7

are free-floating in the fit, with three normalisation factors for each (corresponding to the High-Mass, Compressed, and one- c -tag regions).

A background-only fit is performed for the High-Mass (CRZ-HM, CRW-HM), Compressed (CRZ-Comp, CRW-Comp), and one- c -tag (CRZ-Comp-1c, CRW-Comp-1c) regions using the likelihood in the control regions and extrapolating the normalisation of Z + jets and W + jets contributions into the validation or signal regions (thus not considering the data yield in the signal and validation regions). The fitting strategy reduces systematic uncertainties in the signal regions because the control regions are chosen to minimise the extrapolation uncertainties. Thus, the effect on the W + jets and Z + jets background estimates due to systematic uncertainties is similar in the control and signal regions and partially cancels out. This strategy is especially effective in the High-Mass regions, where the total uncertainty of the SM backgrounds is reduced from 30%–37% to 13%–17%. The systematic uncertainties that contribute most to the total post-fit uncertainty are related to c -tagging and the V + jets theoretical uncertainties. The latter is dominated by variations of the scales, with contributions to the total uncertainty ranging from 14% to 36%, while the largest contributions from c -tagging-related systematic uncertainties range from 8% to 14% in the signal regions.

No significant excesses over the expected SM backgrounds in the validation and signal regions are observed when considering all statistical and systematic (experimental and theoretical) uncertainties, as shown in Tables 12.1 and 12.2 and Figure 12.1. The extracted normalisation factors for the W + jets and Z + jets backgrounds are shown in Table 12.3 and are generally consistent with unity given their associated uncertainties, with the normalisation factors for W + jets background being larger across all control regions. Selected kinematic distributions in the signal regions are shown in Figure 12.2.

Table 12.2. Post-fit event yields for the Compressed control, validation, and signal regions. The Z + jets and W + jets backgrounds are normalised by fits within the control regions. Processes with top quarks and multiple vector bosons are included in “Other”. Pre-fit event yields for a representative signal point in each signal region are also shown.

	CRW-Comp	CRZ-Comp	CRW-Comp-1c	CRZ-Comp-1c	VR-Comp	VR-Comp-1c
Z +jets	0.33 ± 0.12	26 ± 6	0.65 ± 0.18	38 ± 7	48 ± 12	48 ± 13
W +jets	41 ± 10	–	89 ± 11	–	71 ± 25	32 ± 11
Other	22 ± 6	3.9 ± 1.2	15 ± 4	4.0 ± 1.1	21 ± 6	2.8 ± 0.8
Total SM	63 ± 8	30 ± 5	105 ± 10	42 ± 6	139 ± 27	83 ± 17
Observed	63	30	105	42	141	76

	SR-Comp1	SR-Comp2	SR-Comp3	SR-Comp-1c
Z +jets	0.9 ± 0.4	1.6 ± 0.9	3.1 ± 0.9	4.3 ± 1.2
W +jets	2.2 ± 1.0	1.3 ± 0.7	0.9 ± 0.6	2.8 ± 0.9
Other	0.56 ± 0.19	0.62 ± 0.20	1.04 ± 0.34	0.68 ± 0.29
Total SM	3.7 ± 1.2	3.5 ± 1.2	5.0 ± 1.2	7.9 ± 1.7
Observed	3	6	8	12
$m(\tilde{t}_1, \tilde{\chi}_1^0) = (600, 550)$ GeV	6.1 ± 1.3			
$m(\tilde{t}_1, \tilde{\chi}_1^0) = (550, 470)$ GeV		6.0 ± 0.8		
$m(\tilde{t}_1, \tilde{\chi}_1^0) = (550, 375)$ GeV			7.9 ± 2.1	
$m(\tilde{t}_1, \tilde{\chi}_1^0) = (450, 430)$ GeV				11.2 ± 1.7

Table 12.3. Summary of the extracted normalisation factors when considering only the likelihood in the control regions for the High-Mass (HM) regions, the Compressed regions that require two c -tagged jets (Comp), and the Compressed regions requiring exactly one c -tagged jet (Comp-1c). Uncertainties include both the statistical and systematic components but do not include the extrapolation into the signal region, which can reduce the total uncertainty. The W + jets normalisation factors have larger uncertainties because of the larger effect of theoretical uncertainties.

Region	Normalisation factor
HM Z + jets	1.0 ± 0.4
Comp Z + jets	1.0 ± 0.4
Comp-1c Z + jets	1.1 ± 0.4
HM W + jets	1.4 ± 0.6
Comp W + jets	1.3 ± 0.7
Comp-1c W + jets	1.6 ± 0.7

12.2. Results and interpretation




Figure 12.1. Comparison of event yields for the data (points) and SM expectation (stacked histograms) in all signal and validation regions after the background-only fits. The hatched uncertainty band around the SM prediction represents the total uncertainty, i.e. combining the detector-related and theory-based systematic uncertainties and the MC statistical uncertainties. The significance of the difference between the data and the SM background prediction, calculated with the profile-likelihood method described in ref. [226] and using the total uncertainty, is shown in the bottom panel. Processes with top quarks and multiple vector bosons are included in “Other”.

The exclusion fit is performed for each signal model individually, using the signal and control regions. For SUSY models, the two statistical combinations considered use the three High-Mass signal regions with non-overlapping Compressed signal regions: the first combines SR-HM1, SR-HM2, SR-HM3, SR-Comp1, SR-Comp2, and SR-Comp-1c, while the second combines SR-HM1, SR-HM2, SR-HM3, SR-Comp3, and SR-Comp-1c. For both combinations, all six control regions are used to estimate the background contributions. To maximise the exclusion power of the two configurations, the combination with the lowest expected CL_s [227] value is chosen for the final result. The expected sensitivity is calculated by generating pseudo-experiments that assume an absence of BSM signals. Exclusion limits are set at 95% CL and are shown for the considered SUSY models in Figure 12.3. The expected exclusion limits from this search exceed those from the previous ATLAS search that used c -tagging [183] (grey-filled contour and dashed line in Figure 12.3), while being complementary to an ATLAS search for single squark production that did not utilise c -tagging [186] (green-filled contour and dashed line in Figure 12.3a) for $m_{\tilde{t}_1/\tilde{c}_1} \lesssim 920$ GeV and $m(\tilde{\chi}_1^0) \gtrsim 400$ GeV, and is also complementary to an ATLAS single-jet search [185] (cyan-filled contour and dashed line in Figure 12.3b) for $\Delta m(\tilde{t}_1, \tilde{\chi}_1^0) \gtrsim 30$ GeV.

Exclusion limits are also presented in Figure 12.4 for an alternative interpretation in which top squarks are pair produced and each decays into a neutralino and either a top or a charm quark, with $\mathcal{B}(\tilde{t}_1 \rightarrow c + \tilde{\chi}_1^0)$ being varied between 0.1 and 1. Due to the b -jet veto imposed, none of the signal and control regions are sensitive to the decay pathway producing a top quark; nonetheless, these decay modes are included in the simulated signal event samples. Two cases are studied: the mass of the neutralino is fixed to either 1 GeV or 200 GeV. The strongest limits are set when the top squark decays only via the $c + \tilde{\chi}_1^0$ channel, as expected, with top squark masses




Figure 12.2. Distributions for the signal regions showing the data (points) and the expected backgrounds (stacked histograms), after simultaneously fitting all the control regions. The histograms are binned in the key variables for each signal region: (a) E_T^{miss} Sig. for SR-HM1, (b) $m_T^{\text{min}}(c)$ for SR-HM2, (c) R_{ISR} for SR-Comp1, (d) R_{ISR} for SR-Comp2, (e) R_{ISR} for SR-Comp3, and (f) R_{ISR} for SR-Comp-1c. The hatched band is the total systematic uncertainty for the expected backgrounds. Distributions of representative $\tilde{t}_1\bar{\tilde{t}}_1$ signals for different mass assumptions (or similar) are shown as dashed lines using pre-fit signal contributions. Processes with top quarks and multiple vector bosons are included in “Other”. The right-most bin in each histogram does not include the overflow entry but the x -axis range is chosen to include all observed data.

12.2. Results and interpretation




Figure 12.3. Exclusion limits at 95% CL for pair production of top squarks decaying into charm quarks and neutralinos. The expected limit is shown as a dashed line while the $\pm 1\sigma$ interval is shown as a yellow band. The observed limit is shown as a red line and the effect of varying the signal cross-section by $\pm 1\sigma$ of its predicted theoretical uncertainty is shown as red dashed lines. The contours are shown as two-dimensional projections for masses of squarks and neutralinos (a) and their differences (b). The exclusion contour from the previous ATLAS result that used charm tagging (ref. [183]) is shown as a grey-dashed line (expected) and grey-filled contour (observed). For (a), the single-squark-production exclusion contour from the ATLAS squarks and gluinos search that did not utilise charm tagging (ref. [186]) is shown as a green-dashed line (expected) and green-filled contour (observed), while for (b), the contour from the ATLAS single-jet search (ref. [185]) is shown as a cyan-dashed line (expected) and cyan-filled contour (observed).

up to approximately 900 GeV being excluded at 95% CL in both cases. The scenario in which $m(\tilde{\chi}_1^0) = 1$ GeV presents a boosted decay topology over the entire range of top squark masses considered, and as such, the major constraining power comes from the High-Mass signal regions. A monotonic decrease in sensitivity is observed with decreasing branching fraction to $c + \tilde{\chi}_1^0$. For $m(\tilde{\chi}_1^0) = 200$ GeV neutralinos, the Compressed signal regions gradually become more sensitive as the top squark mass decreases, causing the change in behaviour near $m(\tilde{t}_1) = 500$ GeV.

Additional limits are set in terms of LQ models after performing a combined fit of the LQ signals using only the High-Mass regions because of similarities in kinematics in the scenario where the LSP is massless. In the case of scalar-LQ models, events were reweighted to span the full range of branching fractions of the LQ decay into $c\nu$. Expected and observed exclusion limits on the $\mathcal{B}(\text{LQ}^u \rightarrow c\nu_e)$ and $\mathcal{B}(\text{LQ}^u \rightarrow c\nu_\mu)$ at 95% CL are presented in Figure 12.5 as a function of the leptoquark mass, with the observed upper limit on the production cross-section shown on the z-axis. Leptoquarks with masses up to approximately 900 GeV are excluded.

Vector LQ models are excluded at 95% CL up to masses of 950 GeV in the minimal coupling scenario, and 1150 GeV for the Yang–Mills case, for an assumed branching fraction of $\mathcal{B}(\text{vLQ}_{23}^u \rightarrow c\nu_\tau) = 0.5$. Upper limits on the production cross-section for both the minimal and Yang–Mills coupling scenarios are shown in Figure 12.6, with the theoretical cross-section at LO and the $\pm 1\sigma$ interval (encompassing the effects of PDF, α_s , renormalisation and factorisation scale variations) shown in blue.

A model-independent fit is performed for signal regions with the least stringent requirements: SR-HM1, SR-Comp3, and SR-HM-Disc. For the model-independent interpretation, the SM background estimates are fit to the observed yields in both the CRs and SRs, testing for the potential presence of any BSM events in the SRs. The observed cross-section limit, $\langle \epsilon\sigma \rangle_{\text{obs}}^{95}$, is




Figure 12.4. The expected and observed exclusion contours for pair production of top squarks decaying into top or charm quarks and neutralinos. The contours are shown as two-dimensional projections in the plane of the squark mass and assumed branching fraction, whose value is varied between 0 (all decays to $t + \tilde{\chi}_1^0$) and 1 (all decays to $c + \tilde{\chi}_1^0$). The mass of the neutralino is fixed to either (a) 1 GeV or (b) 200 GeV. The cross-section limits for the contours are calculated by combining the signal regions as discussed in the text. The branching fractions above and masses below the lines (high branching fraction and low mass, and thus high production cross-section) are excluded. The expected limit is shown as a dashed line, while the $\pm 1\sigma$ interval is shown as a yellow band. The observed limit is shown as a red line and the effect of varying the signal cross-section by $\pm 1\sigma$ of its predicted theoretical uncertainty is shown as red dashed lines.




Figure 12.5. Expected (white) and observed (red) exclusion limits at 95% CL for up-type scalar LQs coupled to (a) first- and (b) second-generation leptons. The observed upper limit on the production cross-section is plotted on the z -axis as a function of (a) $\mathcal{B}(\text{LQ}^u \rightarrow c\nu_e)$ with the alternative decay being $\text{LQ}^u \rightarrow s + e$ and (b) $\mathcal{B}(\text{LQ}^u \rightarrow c\nu_\mu)$ with the alternative decay being $\text{LQ}^u \rightarrow s + \mu$ (y -axis) and the leptoquark masses (x -axis). The branching fractions above and masses below the lines (high branching fraction and low mass, and thus high production cross-section) are excluded. The $\pm 1\sigma$ interval of the expected limit is shown as dashed white lines while the effect of varying the signal cross-section by $\pm 1\sigma$ of its predicted theoretical uncertainty is shown as red dashed lines.

12.3. Conclusion




Figure 12.6. Expected (dashed) and observed (solid) cross-section upper limits as a function of leptoquark mass for $U(1)$ vector-LQ models in the (a) minimal and (b) Yang–Mills coupling scenarios with $\beta_{23} = 1$, corresponding to a branching fraction $\mathcal{B}(\text{vLQ}_{23}^u \rightarrow c\nu_\tau) = 0.5$. The $\pm 1\sigma$ interval of the expected limit is shown as a yellow band. The theoretical prediction and its $\pm 1\sigma$ interval band are shown in blue.

Table 12.4. Model-independent upper limits on the BSM event yields and cross-sections. The $\langle \varepsilon \sigma \rangle_{\text{obs}}^{95}$ is the observed upper limit on the visible BSM cross-section at 95% CL, and S_{obs}^{95} is the observed upper limit on the number of BSM events. Similarly, S_{exp}^{95} is the expected upper limit on the number of BSM events but it is computed assuming that the observed number of events is identical to the number of background events. The discovery p -value ($p_{\text{obs}}(S = 0)$), which tests the compatibility of the observed data with the background-only (zero signal yield) hypothesis relative to fluctuations of the background, is shown in the last column.

Signal Region	$\langle \varepsilon \sigma \rangle_{\text{obs}}^{95}$ [fb]	S_{obs}^{95}	S_{exp}^{95}	$p_{\text{obs}}(S = 0)$
SR-HM1	0.09	12.8	$9.7^{+4.2}_{-2.9}$	0.19
SR-HM-Disc	0.36	50.7	34^{+13}_{-9}	0.09
SR-Comp3	0.07	9.2	$6.5^{+3.2}_{-2.0}$	0.15
SR-Comp-1c	0.08	11.6	$8.0^{+3.8}_{-2.4}$	0.13

calculated by dividing the observed limit on the signal strength, S_{obs}^{95} , by the integrated luminosity. These limits on the number of BSM events are shown in Table 12.4 and can be reinterpreted for a generic BSM model by calculating the identification efficiency and detector acceptance for the model of interest.

12.3 Conclusion

We have presented a search targeting BSM models contributing to final states with missing transverse momentum and charm-tagged jets, based on a 139 fb^{-1} dataset of proton–proton collisions at $\sqrt{s} = 13 \text{ TeV}$ recorded by the ATLAS experiment at the LHC. The observed experimental data are found to agree with the SM background estimate. The results are interpreted as upper limits at 95% CL on the production cross-section of selected SUSY and leptoquark models and as

model-independent upper limits on the production cross-section for BSM particles. Top/charm squark masses are excluded up to ~ 900 GeV, assuming each squark decays into a charm quark and an LSP. This is a ~ 100 GeV improvement on the previous ATLAS search that used c -tagging. The increase in sensitivity is driven by the larger experimental dataset, the sophisticated design of the signal regions (which includes different strategies for large and small mass splittings between the top squark and the LSP), and improvements in c -tagging. The search also excludes, for the first time at the LHC, scalar leptoquarks with a mass below 900 GeV that decay into a charm quark and either a first- or second-generation neutrino. Finally, vector leptoquark models in the minimal and Yang–Mills coupling scenarios are excluded at 95% CL up to masses of 950 GeV and 1150 GeV, respectively, assuming a branching fraction of $\mathcal{B}(\text{vLQ}_{23}^u \rightarrow c\nu_\tau) = 0.5$.

AUXILIARY MATERIAL

A Signal acceptances and efficiencies

The acceptance ($\mathcal{A}$) of signal events in each of the analysis signal regions is defined as the fraction of generator-level events that pass the selection criteria. This was computed using the SimpleAnalysis framework [228]. The reconstruction efficiency (ϵ), which is a measure of reconstruction and detector effects, can be calculated as the ratio of reconstructed to truth-level event yields in each signal region, where the latter can be inferred from the acceptance and an assumed integrated luminosity ($\mathcal{L}$):

$$\epsilon_{\text{SR}} = \frac{N_{\text{SR}}^{\text{reco}}}{N_{\text{SR}}^{\text{truth}}} = \frac{N_{\text{SR}}^{\text{reco}}}{\mathcal{A}_{\text{SR}} \sigma k \epsilon_{\text{filter}} \mathcal{L}}. \quad (12.2)$$

The acceptance and its product with the reconstruction efficiency of each point in the $m(\tilde{t}_1, \tilde{\chi}_1^0)$ signal grid are presented in this section for all of the analysis signal regions.




Figure A.1. Acceptance across the signal grid for the High-Mass SRs: (a) SR-HM1, (b) SR-HM2, (c) SR-HM3, and (d) SR-HM-Disc.

A. Signal acceptances and efficiencies







Figure A.2. Acceptance times efficiency across the signal grid for the High-Mass SRs: (a) SR-HM1, (b) SR-HM2, (c) SR-HM3, and (d) SR-HM-Disc.




(a)



(b)



(c)



(d)

Figure A.3. Acceptance across the signal grid for the compressed SRs: (a) SR-Comp-1c, (b) SR-Comp1, (c) SR-Comp2, and (d) SR-Comp3.

A. Signal acceptances and efficiencies




Figure A.4. Acceptance times efficiency across the signal grid for the compressed SRs: (a) SR-Comp-1c, (b) SR-Comp1, (c) SR-Comp2, and (d) SR-Comp3.

B Cutflow tables

Table 12.5. Cutflow table for SR-HM1 using the $m(\tilde{l}_1, \tilde{\chi}_1^0) = (1000, 1)$ GeV signal point as the benchmark. All jet-related variables use jets with $p_T > 40$ GeV, however, the E_T^{miss} is calculated with jets with $p_T > 20$ GeV. The “Preselection” requirement includes: $E_T^{\text{miss}} > 250$ GeV; events containing electrons or muons are vetoed; at least two jets must be present of which at least one must be identified as a c -tagged jet; events with jets that have been identified as b -tagged are discarded; the minimum azimuthal angular separation ($\Delta\phi$) between the (up to) four leading jets and the $\mathbf{p}_T^{\text{miss}}$ is required to be greater than 0.4.

SR-HM1			
Requirement	Weighted Yields	Reduction [%]	Raw Yields
$\sigma \times \mathcal{L}$	924.15	—	—
Preselection	162.96	−82.37	1354
Number of c -tagged jets (20% WP) ≥ 2	19.51	−88.03	165
τ veto	19.51	−0	165
Leading jet is c -tagged	15.52	−20.45	134
$p_T(j_2) > 200$ GeV	12.97	−16.43	113
$m_T^{\text{min}}(c) > 300$ GeV	12.28	−5.32	108
$m_{cc} > 200$ GeV	12.18	−0.81	107
E_T^{miss} Sig. > 22	6.68	−45.16	56

Table 12.6. Cutflow table for SR-HM2 using the $m(\tilde{l}_1, \tilde{\chi}_1^0) = (750, 450)$ GeV signal point as the benchmark. All jet-related variables use jets with $p_T > 40$ GeV, however, the E_T^{miss} is calculated with jets with $p_T > 20$ GeV. The “Preselection” requirement includes: $E_T^{\text{miss}} > 250$ GeV; events containing electrons or muons are vetoed; at least two jets must be present of which at least one must be identified as a c -tagged jet; events with jets that have been identified as b -tagged are discarded; the minimum azimuthal angular separation ($\Delta\phi$) between the (up to) four leading jets and the $\mathbf{p}_T^{\text{miss}}$ is required to be greater than 0.4.

SR-HM2			
Requirement	Weighted Yields	Reduction [%]	Raw Yields
$\sigma \times \mathcal{L}$	6222.92	—	—
Preselection	705.73	−88.66	1958
Number of c -tagged jets (20% WP) ≥ 2	127.44	−81.94	359
τ veto	127.03	−0.32	358
Leading jet is c -tagged	97.24	−23.45	284
$p_T(j_2) > 150$ GeV	58.57	−39.77	171
$m_T^{\text{min}}(c) > 300$ GeV	51.02	−12.89	147
E_T^{miss} Sig. > 15	43.92	−13.92	125
E_T^{miss} Sig. < 22	42.25	−3.80	119
$m_{cc} > 200$ GeV	39.94	−5.47	113
$m_{cc} < 400$ GeV	17.28	−56.74	48

B. Cutflow tables

Table 12.7. Cutflow table for SR-HM3 using the $m(\tilde{t}_1, \tilde{\chi}_1^0) = (750, 450)$ GeV signal point as the benchmark. All jet-related variables use jets with $p_T > 40$ GeV, however, the E_T^{miss} is calculated with jets with $p_T > 20$ GeV. The “Preselection” requirement includes: $E_T^{\text{miss}} > 250$ GeV; events containing electrons or muons are vetoed; at least two jets must be present of which at least one must be identified as a c -tagged jet; events with jets that have been identified as b -tagged are discarded; the minimum azimuthal angular separation ($\Delta\phi$) between the (up to) four leading jets and the $\mathbf{p}_T^{\text{miss}}$ is required to be greater than 0.4.

SR-HM3			
Requirement	Weighted Yields	Reduction [%]	Raw Yields
$\sigma \times \mathcal{L}$	6222.92	—	—
Preselection	705.73	−88.66	1958
Number of c -tagged jets (20% WP) ≥ 2	127.44	−81.94	359
τ veto	127.03	−0.32	358
Leading jet is c -tagged	97.24	−23.45	284
$p_T(j_2) > 150$ GeV	58.57	−39.77	171
$m_T^{\text{min}}(c) > 300$ GeV	51.02	−12.89	147
E_T^{miss} Sig. > 15	43.92	−13.92	125
E_T^{miss} Sig. < 22	42.25	−3.80	119
$m_{cc} > 400$ GeV	22.66	−46.37	65
$m_{cc} < 600$ GeV	12.54	−44.66	36

Table 12.8. Cutflow table for SR-HM-Disc using the $m(\tilde{t}_1, \tilde{\chi}_1^0) = (750, 450)$ GeV signal point as the benchmark. All jet-related variables use jets with $p_T > 40$ GeV, however, the E_T^{miss} is calculated with jets with $p_T > 20$ GeV. The “Preselection” requirement includes: $E_T^{\text{miss}} > 250$ GeV; events containing electrons or muons are vetoed; at least two jets must be present of which at least one must be identified as a c -tagged jet; events with jets that have been identified as b -tagged are discarded; the minimum azimuthal angular separation ($\Delta\phi$) between the (up to) four leading jets and the $\mathbf{p}_T^{\text{miss}}$ is required to be greater than 0.4.

SR-HM-Disc			
Requirement	Weighted Yields	Reduction [%]	Raw Yields
$\sigma \times \mathcal{L}$	6222.92	—	—
Preselection	705.73	−88.66	1958
Number of c -tagged jets (20% WP) ≥ 2	127.44	−81.94	359
τ veto	127.03	−0.32	358
Leading jet is c -tagged	97.24	−23.45	284
$p_T(j_2) > 150$ GeV	58.57	−39.77	171
$m_T^{\text{min}}(c) > 300$ GeV	51.02	−12.89	147
E_T^{miss} Sig. > 15	43.92	−13.92	125
$m_{cc} > 200$ GeV	41.26	−6.06	118

Table 12.9. Cutflow table for SR-Comp-1c using the $m(\tilde{t}_1, \tilde{\chi}_1^0) = (450, 430)$ GeV signal point as the benchmark. All jet-related variables use jets with $p_T > 20$ GeV. The “Preselection” requirement includes: $E_T^{\text{miss}} > 250$ GeV; events containing electrons or muons are vetoed; at least two jets must be present of which at least one must be identified as a c -tagged jet; events with jets that have been identified as b -tagged are discarded; the minimum azimuthal angular separation ($\Delta\phi$) between the (up to) four leading jets and the $\mathbf{p}_T^{\text{miss}}$ is required to be greater than 0.4.

SR-Comp-1c			
Requirement	Weighted Yields	Reduction [%]	Raw Yields
$\sigma \times \mathcal{L}$	40253.38	–	–
Preselection	1164.95	–97.11	4670
Number of c -tagged jets (20% WP) = 1	1071.96	–7.98	4315
$\Delta\phi(\mathbf{p}_T^{\text{miss}}, j_1) > 2$	1065.07	–0.64	4288
$p_T(j_1) > 350$ GeV	609.05	–42.82	2487
τ veto	591.73	–2.84	2419
Leading jet is not c -tagged	499.74	–15.55	2049
$m_T^{\text{min}}(c) > 100$ GeV	335.21	–32.92	1365
$N_{c\text{-jets}}^S = 1$	238.95	–28.72	980
$m_V < 80$ GeV	120.37	–49.63	479
$p_T^{\text{CM}} < 10$ GeV	52.27	–56.58	210
$p_T^{\text{ISR}} > 700$ GeV	12.24	–76.58	48
$R_{\text{ISR}} > 0.9$	11.72	–4.25	46
$R_{\text{ISR}} < 1.0$	11.19	–4.52	44

Table 12.10. Cutflow table for SR-Comp1 using the $m(\tilde{t}_1, \tilde{\chi}_1^0) = (600, 550)$ GeV signal point as the benchmark. All jet-related variables use jets with $p_T > 20$ GeV. The “Preselection” requirement includes: $E_T^{\text{miss}} > 250$ GeV; events containing electrons or muons are vetoed; at least two jets must be present of which at least one must be identified as a c -tagged jet; events with jets that have been identified as b -tagged are discarded; the minimum azimuthal angular separation ($\Delta\phi$) between the (up to) four leading jets and the $\mathbf{p}_T^{\text{miss}}$ is required to be greater than 0.4.

SR-Comp1			
Requirement	Weighted Yields	Reduction [%]	Raw Yields
$\sigma \times \mathcal{L}$	9121.23	–	–
Preselection	484.27	–94.69	2079
Number of c -tagged jets (20% WP) ≥ 2	71.12	–85.31	304
τ veto	69.68	–2.02	298
Leading jet is not c -tagged	53.77	–22.83	232
$N_{c\text{-jets}}^S \geq 2$	19.46	–63.81	82
$p_T^{\text{ISR}} > 550$ GeV	10.14	–47.89	41
$R_{\text{ISR}} > 0.85$	6.07	–40.14	23

B. Cutflow tables

Table 12.11. Cutflow table for SR-Comp2 using the $m(\tilde{t}_1, \tilde{\chi}_1^0) = (550, 470)$ GeV signal point as the benchmark. All jet-related variables use jets with $p_T > 20$ GeV. The “Preselection” requirement includes: $E_T^{\text{miss}} > 250$ GeV; events containing electrons or muons are vetoed; at least two jets must be present of which at least one must be identified as a c -tagged jet; events with jets that have been identified as b -tagged are discarded; the minimum azimuthal angular separation ($\Delta\phi$) between the (up to) four leading jets and the $\mathbf{p}_T^{\text{miss}}$ is required to be greater than 0.4.

SR-Comp2			
Requirement	Weighted Yields	Reduction [%]	Raw Yields
$\sigma \times \mathcal{L}$	17465.59	—	—
Preselection	883.77	−94.94	7372
Number of c -tagged jets (20% WP) ≥ 2	147.27	−83.34	1217
τ veto	144.41	−1.94	1194
Leading jet is not c -tagged	114.64	−20.61	944
$m_T^{\min}(c) > 80$ GeV	102.39	−10.69	845
$N_{c\text{-jets}}^S \geq 2$	26.18	−74.43	225
$p_T^{\text{ISR}} > 500$ GeV	16.80	−35.83	146
$R_{\text{ISR}} > 0.75$	10.68	−36.43	94
$R_{\text{ISR}} < 0.85$	6.03	−43.54	53

Table 12.12. Cutflow table for SR-Comp3 using the $m(\tilde{t}_1, \tilde{\chi}_1^0) = (550, 375)$ GeV signal point as the benchmark. All jet-related variables use jets with $p_T > 20$ GeV. The “Preselection” requirement includes: $E_T^{\text{miss}} > 250$ GeV; events containing electrons or muons are vetoed; at least two jets must be present of which at least one must be identified as a c -tagged jet; events with jets that have been identified as b -tagged are discarded; the minimum azimuthal angular separation ($\Delta\phi$) between the (up to) four leading jets and the $\mathbf{p}_T^{\text{miss}}$ is required to be greater than 0.4.

SR-Comp3			
Requirement	Weighted Yields	Reduction [%]	Raw Yields
$\sigma \times \mathcal{L}$	33371.72	—	—
Preselection	1578.44	−95.27	3392
Number of c -tagged jets (20% WP) ≥ 2	279.07	−82.32	596
τ veto	275.39	−1.32	589
Leading jet is not c -tagged	139.75	−49.25	281
$m_T^{\min}(c) > 120$ GeV	124.15	−11.16	247
$m_{cc} > 100$ GeV	115.94	−6.61	228
$N_{c\text{-jets}}^S \geq 2$	14.19	−87.76	29
$p_T^{\text{ISR}} > 400$ GeV	14.19	−0	29
$R_{\text{ISR}} > 0.65$	7.93	−44.12	16

C Auxiliary material for leptoquark models




Figure C.5. Acceptance across the LQ_{21}^u signal grid for the High-Mass SRs: (a) SR-HM1, (b) SR-HM2, (c) SR-HM3, and (d) SR-HM-Disc.

C. Auxiliary material for leptoquark models







Figure C.6. Acceptance across the LQ_{22}^u signal grid for the High-Mass SRs: (a) SR-HM1, (b) SR-HM2, (c) SR-HM3, and (d) SR-HM-Disc.




(a)



(b)



(c)



(d)

Figure C.7. Acceptance times efficiency across the LQ_{21}^u signal grid for the High-Mass SRs: (a) SR-HM1, (b) SR-HM2, (c) SR-HM3, and (d) SR-HM-Disc.

C. Auxiliary material for leptoquark models




Figure C.8. Acceptance times efficiency across the LQ_{22}^u signal grid for the High-Mass SRs: (a) SR-HM1, (b) SR-HM2, (c) SR-HM3, and (d) SR-HM-Disc.

Table 12.13. Acceptances and efficiencies for the $\nu\text{LQ}_{23}^{\text{u,min}}$ signal model in SR-HM1. Due to the limited statistical power of the signal samples, efficiencies can fluctuate and can drop to zero.

$m(\nu\text{LQ}_{23}^{\text{u,min}})$ [GeV]	$\mathcal{A}$	ϵ	$\mathcal{A} \times \epsilon$
300	6.11×10^{-5}	0.00	0.00
400	6.20×10^{-4}	2.31×10^{-2}	1.43×10^{-5}
500	2.56×10^{-3}	6.10×10^{-2}	1.56×10^{-4}
600	8.54×10^{-3}	3.50×10^{-2}	2.99×10^{-4}
700	1.82×10^{-2}	3.78×10^{-2}	6.90×10^{-4}
800	3.03×10^{-2}	3.26×10^{-2}	9.90×10^{-4}
850	3.65×10^{-2}	3.02×10^{-2}	1.10×10^{-3}
900	4.17×10^{-2}	3.06×10^{-2}	1.27×10^{-3}
950	4.79×10^{-2}	2.95×10^{-2}	1.41×10^{-3}
1000	5.36×10^{-2}	2.99×10^{-2}	1.60×10^{-3}
1050	5.87×10^{-2}	2.72×10^{-2}	1.59×10^{-3}
1100	6.26×10^{-2}	2.55×10^{-2}	1.60×10^{-3}
1150	6.75×10^{-2}	1.92×10^{-2}	1.29×10^{-3}
1200	7.16×10^{-2}	2.18×10^{-2}	1.56×10^{-3}
1250	7.54×10^{-2}	2.17×10^{-2}	1.64×10^{-3}
1300	7.85×10^{-2}	2.14×10^{-2}	1.68×10^{-3}
1350	8.27×10^{-2}	2.12×10^{-2}	1.75×10^{-3}
1400	8.58×10^{-2}	2.14×10^{-2}	1.84×10^{-3}
1450	8.96×10^{-2}	1.72×10^{-2}	1.54×10^{-3}
1500	9.11×10^{-2}	1.43×10^{-2}	1.30×10^{-3}
1550	9.34×10^{-2}	1.69×10^{-2}	1.58×10^{-3}
1600	9.64×10^{-2}	1.08×10^{-2}	1.04×10^{-3}
1700	1.00×10^{-1}	1.23×10^{-2}	1.23×10^{-3}
1800	1.04×10^{-1}	9.29×10^{-3}	9.67×10^{-4}
1900	1.08×10^{-1}	1.28×10^{-2}	1.38×10^{-3}
2000	1.12×10^{-1}	6.72×10^{-3}	7.56×10^{-4}

C. Auxiliary material for leptoquark models

Table 12.14. Acceptances and efficiencies for the $\nu\text{LQ}_{23}^{\text{u,min}}$ signal model in SR-HM2. Due to the limited statistical power of the signal samples, efficiencies can fluctuate and can drop to zero.

$m(\nu\text{LQ}_{23}^{\text{u,min}})$ [GeV]	$\mathcal{A}$	ϵ	$\mathcal{A} \times \epsilon$
300	4.97×10^{-4}	1.02×10^{-1}	5.08×10^{-5}
400	2.30×10^{-3}	7.36×10^{-2}	1.70×10^{-4}
500	5.29×10^{-3}	4.62×10^{-2}	2.44×10^{-4}
600	6.07×10^{-3}	6.01×10^{-2}	3.65×10^{-4}
700	5.60×10^{-3}	6.99×10^{-2}	3.92×10^{-4}
800	4.54×10^{-3}	7.67×10^{-2}	3.48×10^{-4}
850	4.14×10^{-3}	5.34×10^{-2}	2.21×10^{-4}
900	3.59×10^{-3}	7.26×10^{-2}	2.61×10^{-4}
950	3.39×10^{-3}	7.01×10^{-2}	2.37×10^{-4}
1000	3.39×10^{-3}	5.50×10^{-2}	1.87×10^{-4}
1050	2.79×10^{-3}	5.25×10^{-2}	1.46×10^{-4}
1100	2.39×10^{-3}	6.06×10^{-2}	1.45×10^{-4}
1150	2.44×10^{-3}	4.42×10^{-2}	1.08×10^{-4}
1200	2.00×10^{-3}	7.43×10^{-2}	1.48×10^{-4}
1250	1.86×10^{-3}	6.48×10^{-2}	1.20×10^{-4}
1300	1.95×10^{-3}	4.30×10^{-2}	8.40×10^{-5}
1350	1.86×10^{-3}	5.16×10^{-2}	9.58×10^{-5}
1400	1.58×10^{-3}	1.36×10^{-1}	2.15×10^{-4}
1450	1.50×10^{-3}	3.52×10^{-2}	5.28×10^{-5}
1500	1.52×10^{-3}	8.29×10^{-2}	1.26×10^{-4}
1550	1.62×10^{-3}	5.59×10^{-2}	9.06×10^{-5}
1600	1.73×10^{-3}	1.23×10^{-2}	2.14×10^{-5}
1700	1.17×10^{-3}	1.53×10^{-2}	1.79×10^{-5}
1800	1.30×10^{-3}	6.76×10^{-2}	8.79×10^{-5}
1900	1.17×10^{-3}	4.07×10^{-2}	4.75×10^{-5}
2000	1.53×10^{-3}	9.86×10^{-2}	1.51×10^{-4}

Table 12.15. Acceptances and efficiencies for the $\nu\text{LQ}_{23}^{\text{u,min}}$ signal model in SR-HM3. Due to the limited statistical power of the signal samples, efficiencies can fluctuate and can drop to zero.

$m(\nu\text{LQ}_{23}^{\text{u,min}})$ [GeV]	$\mathcal{A}$	ϵ	$\mathcal{A} \times \epsilon$
300	1.69×10^{-3}	6.62×10^{-2}	1.12×10^{-4}
400	6.15×10^{-3}	5.49×10^{-2}	3.38×10^{-4}
500	1.05×10^{-2}	5.98×10^{-2}	6.26×10^{-4}
600	1.23×10^{-2}	6.09×10^{-2}	7.46×10^{-4}
700	1.12×10^{-2}	6.45×10^{-2}	7.24×10^{-4}
800	9.00×10^{-3}	6.00×10^{-2}	5.40×10^{-4}
850	7.61×10^{-3}	6.58×10^{-2}	5.01×10^{-4}
900	6.97×10^{-3}	5.65×10^{-2}	3.94×10^{-4}
950	5.88×10^{-3}	5.83×10^{-2}	3.43×10^{-4}
1000	5.15×10^{-3}	7.98×10^{-2}	4.11×10^{-4}
1050	4.48×10^{-3}	6.68×10^{-2}	2.99×10^{-4}
1100	3.87×10^{-3}	5.69×10^{-2}	2.20×10^{-4}
1150	3.33×10^{-3}	6.34×10^{-2}	2.11×10^{-4}
1200	3.08×10^{-3}	3.68×10^{-2}	1.13×10^{-4}
1250	2.70×10^{-3}	4.63×10^{-2}	1.25×10^{-4}
1300	2.39×10^{-3}	4.63×10^{-2}	1.11×10^{-4}
1350	2.37×10^{-3}	6.01×10^{-2}	1.42×10^{-4}
1400	2.34×10^{-3}	3.42×10^{-2}	8.00×10^{-5}
1450	1.80×10^{-3}	4.54×10^{-2}	8.16×10^{-5}
1500	1.78×10^{-3}	0.00	0.00
1550	1.30×10^{-3}	3.80×10^{-2}	4.94×10^{-5}
1600	1.00×10^{-3}	1.82×10^{-1}	1.82×10^{-4}
1700	1.20×10^{-3}	2.98×10^{-2}	3.57×10^{-5}
1800	1.37×10^{-3}	2.14×10^{-2}	2.93×10^{-5}
1900	1.23×10^{-3}	0.00	0.00
2000	1.07×10^{-3}	0.00	0.00

C. Auxiliary material for leptoquark models

Table 12.16. Acceptances and efficiencies for the $\text{vLQ}_{23}^{\text{u,min}}$ signal model in SR-HM-Disc. Due to the limited statistical power of the signal samples, efficiencies can fluctuate and can drop to zero.

$m(\text{vLQ}_{23}^{\text{u,min}})$ [GeV]	$\mathcal{A}$	ϵ	$\mathcal{A} \times \epsilon$
300	4.72×10^{-3}	5.72×10^{-2}	2.70×10^{-4}
400	1.90×10^{-2}	5.17×10^{-2}	9.81×10^{-4}
500	4.03×10^{-2}	4.45×10^{-2}	1.79×10^{-3}
600	6.07×10^{-2}	4.03×10^{-2}	2.45×10^{-3}
700	7.90×10^{-2}	3.99×10^{-2}	3.16×10^{-3}
800	9.38×10^{-2}	3.49×10^{-2}	3.27×10^{-3}
850	1.00×10^{-1}	3.01×10^{-2}	3.02×10^{-3}
900	1.06×10^{-1}	3.09×10^{-2}	3.26×10^{-3}
950	1.11×10^{-1}	2.86×10^{-2}	3.17×10^{-3}
1000	1.15×10^{-1}	2.95×10^{-2}	3.41×10^{-3}
1050	1.19×10^{-1}	2.69×10^{-2}	3.21×10^{-3}
1100	1.21×10^{-1}	2.51×10^{-2}	3.05×10^{-3}
1150	1.25×10^{-1}	2.12×10^{-2}	2.65×10^{-3}
1200	1.29×10^{-1}	2.12×10^{-2}	2.72×10^{-3}
1250	1.31×10^{-1}	2.09×10^{-2}	2.75×10^{-3}
1300	1.34×10^{-1}	1.99×10^{-2}	2.66×10^{-3}
1350	1.37×10^{-1}	1.93×10^{-2}	2.65×10^{-3}
1400	1.40×10^{-1}	2.06×10^{-2}	2.87×10^{-3}
1450	1.41×10^{-1}	1.52×10^{-2}	2.15×10^{-3}
1500	1.42×10^{-1}	1.23×10^{-2}	1.74×10^{-3}
1550	1.43×10^{-1}	1.64×10^{-2}	2.35×10^{-3}
1600	1.44×10^{-1}	1.31×10^{-2}	1.89×10^{-3}
1700	1.47×10^{-1}	1.15×10^{-2}	1.68×10^{-3}
1800	1.51×10^{-1}	7.74×10^{-3}	1.17×10^{-3}
1900	1.53×10^{-1}	1.02×10^{-2}	1.57×10^{-3}
2000	1.56×10^{-1}	7.29×10^{-3}	1.13×10^{-3}

Table 12.17. Acceptances and efficiencies for the $\nu\text{LQ}_{23}^{\text{u,YM}}$ signal model in SR-HM1. Due to the limited statistical power of the signal samples, efficiencies can fluctuate and can drop to zero.

$m(\nu\text{LQ}_{23}^{\text{u,YM}})$ [GeV]	$\mathcal{A}$	ϵ	$\mathcal{A} \times \epsilon$
300	4.72×10^{-5}	1.97×10^{-1}	9.28×10^{-6}
400	5.63×10^{-4}	6.16×10^{-2}	3.47×10^{-5}
500	2.43×10^{-3}	5.87×10^{-2}	1.43×10^{-4}
600	8.41×10^{-3}	3.93×10^{-2}	3.31×10^{-4}
700	1.99×10^{-2}	3.79×10^{-2}	7.53×10^{-4}
800	3.20×10^{-2}	3.02×10^{-2}	9.67×10^{-4}
850	3.83×10^{-2}	3.05×10^{-2}	1.17×10^{-3}
900	4.52×10^{-2}	3.00×10^{-2}	1.36×10^{-3}
950	5.07×10^{-2}	2.97×10^{-2}	1.51×10^{-3}
1000	5.57×10^{-2}	2.23×10^{-2}	1.24×10^{-3}
1050	6.07×10^{-2}	2.69×10^{-2}	1.63×10^{-3}
1100	6.63×10^{-2}	2.81×10^{-2}	1.86×10^{-3}
1150	7.00×10^{-2}	2.37×10^{-2}	1.66×10^{-3}
1200	7.42×10^{-2}	2.16×10^{-2}	1.61×10^{-3}
1250	7.98×10^{-2}	2.23×10^{-2}	1.78×10^{-3}
1300	8.20×10^{-2}	2.30×10^{-2}	1.88×10^{-3}
1350	8.55×10^{-2}	1.88×10^{-2}	1.60×10^{-3}
1400	8.80×10^{-2}	2.00×10^{-2}	1.76×10^{-3}
1450	9.07×10^{-2}	2.25×10^{-2}	2.04×10^{-3}
1500	9.35×10^{-2}	1.50×10^{-2}	1.40×10^{-3}
1550	9.59×10^{-2}	1.62×10^{-2}	1.55×10^{-3}
1600	1.01×10^{-1}	1.38×10^{-2}	1.39×10^{-3}
1700	1.03×10^{-1}	1.23×10^{-2}	1.27×10^{-3}
1800	1.09×10^{-1}	9.03×10^{-3}	9.81×10^{-4}
1900	1.09×10^{-1}	1.37×10^{-2}	1.50×10^{-3}
2000	1.15×10^{-1}	9.25×10^{-3}	1.06×10^{-3}

C. Auxiliary material for leptoquark models

Table 12.18. Acceptances and efficiencies for the $\nu\text{LQ}_{23}^{\text{u,YM}}$ signal model in SR-HM2. Due to the limited statistical power of the signal samples, efficiencies can fluctuate and can drop to zero.

$m(\nu\text{LQ}_{23}^{\text{u,YM}})$ [GeV]	$\mathcal{A}$	ϵ	$\mathcal{A} \times \epsilon$
300	4.44×10^{-4}	8.38×10^{-2}	3.73×10^{-5}
400	2.71×10^{-3}	5.41×10^{-2}	1.46×10^{-4}
500	6.25×10^{-3}	7.06×10^{-2}	4.41×10^{-4}
600	8.12×10^{-3}	5.22×10^{-2}	4.24×10^{-4}
700	6.76×10^{-3}	6.62×10^{-2}	4.47×10^{-4}
800	5.50×10^{-3}	6.07×10^{-2}	3.34×10^{-4}
850	5.03×10^{-3}	6.14×10^{-2}	3.09×10^{-4}
900	4.38×10^{-3}	5.47×10^{-2}	2.40×10^{-4}
950	3.77×10^{-3}	6.70×10^{-2}	2.53×10^{-4}
1000	3.39×10^{-3}	6.55×10^{-2}	2.22×10^{-4}
1050	3.07×10^{-3}	6.56×10^{-2}	2.01×10^{-4}
1100	2.75×10^{-3}	4.82×10^{-2}	1.33×10^{-4}
1150	2.50×10^{-3}	3.03×10^{-2}	7.56×10^{-5}
1200	2.24×10^{-3}	5.13×10^{-2}	1.15×10^{-4}
1250	2.00×10^{-3}	2.74×10^{-2}	5.50×10^{-5}
1300	1.85×10^{-3}	6.04×10^{-2}	1.12×10^{-4}
1350	2.03×10^{-3}	4.40×10^{-2}	8.93×10^{-5}
1400	1.74×10^{-3}	9.80×10^{-2}	1.71×10^{-4}
1450	1.60×10^{-3}	6.23×10^{-2}	9.97×10^{-5}
1500	1.44×10^{-3}	2.62×10^{-2}	3.78×10^{-5}
1550	1.24×10^{-3}	7.04×10^{-2}	8.73×10^{-5}
1600	1.23×10^{-3}	0.00	0.00
1700	1.37×10^{-3}	1.02×10^{-2}	1.40×10^{-5}
1800	1.27×10^{-3}	1.84×10^{-2}	2.33×10^{-5}
1900	1.17×10^{-3}	0.00	0.00
2000	9.34×10^{-4}	6.68×10^{-2}	6.24×10^{-5}

Table 12.19. Acceptances and efficiencies for the $\nu\text{LQ}_{23}^{\text{u,YM}}$ signal model in SR-HM3. Due to the limited statistical power of the signal samples, efficiencies can fluctuate and can drop to zero.

$m(\nu\text{LQ}_{23}^{\text{u,YM}})$ [GeV]	$\mathcal{A}$	ϵ	$\mathcal{A} \times \epsilon$
300	1.17×10^{-3}	7.54×10^{-2}	8.84×10^{-5}
400	5.29×10^{-3}	6.76×10^{-2}	3.57×10^{-4}
500	1.04×10^{-2}	5.00×10^{-2}	5.20×10^{-4}
600	1.27×10^{-2}	5.35×10^{-2}	6.77×10^{-4}
700	1.19×10^{-2}	5.98×10^{-2}	7.12×10^{-4}
800	1.01×10^{-2}	6.38×10^{-2}	6.42×10^{-4}
850	8.82×10^{-3}	5.91×10^{-2}	5.21×10^{-4}
900	7.72×10^{-3}	5.79×10^{-2}	4.47×10^{-4}
950	6.60×10^{-3}	6.13×10^{-2}	4.05×10^{-4}
1000	5.54×10^{-3}	4.71×10^{-2}	2.61×10^{-4}
1050	4.93×10^{-3}	6.98×10^{-2}	3.44×10^{-4}
1100	4.45×10^{-3}	4.87×10^{-2}	2.17×10^{-4}
1150	3.59×10^{-3}	6.82×10^{-2}	2.45×10^{-4}
1200	3.43×10^{-3}	5.11×10^{-2}	1.75×10^{-4}
1250	2.79×10^{-3}	5.11×10^{-2}	1.42×10^{-4}
1300	2.63×10^{-3}	3.57×10^{-2}	9.39×10^{-5}
1350	2.49×10^{-3}	3.03×10^{-2}	7.54×10^{-5}
1400	2.08×10^{-3}	1.29×10^{-1}	2.69×10^{-4}
1450	2.02×10^{-3}	4.40×10^{-2}	8.90×10^{-5}
1500	1.54×10^{-3}	5.83×10^{-2}	8.98×10^{-5}
1550	1.24×10^{-3}	0.00	0.00
1600	1.43×10^{-3}	3.43×10^{-2}	4.91×10^{-5}
1700	1.33×10^{-3}	0.00	0.00
1800	1.20×10^{-3}	3.89×10^{-2}	4.67×10^{-5}
1900	1.27×10^{-3}	3.04×10^{-2}	3.85×10^{-5}
2000	9.34×10^{-4}	0.00	0.00

C. Auxiliary material for leptoquark models

Table 12.20. Acceptances and efficiencies for the $\text{vLQ}_{23}^{\text{u,YM}}$ signal model in SR-HM-Disc. Due to the limited statistical power of the signal samples, efficiencies can fluctuate and can drop to zero.

$m(\text{vLQ}_{23}^{\text{u,YM}})$ [GeV]	$\mathcal{A}$	ϵ	$\mathcal{A} \times \epsilon$
300	3.51×10^{-3}	6.43×10^{-2}	2.26×10^{-4}
400	1.69×10^{-2}	5.12×10^{-2}	8.64×10^{-4}
500	3.84×10^{-2}	4.61×10^{-2}	1.77×10^{-3}
600	6.11×10^{-2}	3.98×10^{-2}	2.43×10^{-3}
700	8.02×10^{-2}	4.02×10^{-2}	3.22×10^{-3}
800	9.55×10^{-2}	3.44×10^{-2}	3.28×10^{-3}
850	1.02×10^{-1}	3.36×10^{-2}	3.42×10^{-3}
900	1.07×10^{-1}	3.24×10^{-2}	3.48×10^{-3}
950	1.13×10^{-1}	3.18×10^{-2}	3.59×10^{-3}
1000	1.16×10^{-1}	2.48×10^{-2}	2.87×10^{-3}
1050	1.20×10^{-1}	2.82×10^{-2}	3.37×10^{-3}
1100	1.25×10^{-1}	2.63×10^{-2}	3.27×10^{-3}
1150	1.27×10^{-1}	2.31×10^{-2}	2.93×10^{-3}
1200	1.31×10^{-1}	2.23×10^{-2}	2.91×10^{-3}
1250	1.34×10^{-1}	2.16×10^{-2}	2.91×10^{-3}
1300	1.35×10^{-1}	2.15×10^{-2}	2.90×10^{-3}
1350	1.38×10^{-1}	1.88×10^{-2}	2.59×10^{-3}
1400	1.39×10^{-1}	2.04×10^{-2}	2.84×10^{-3}
1450	1.41×10^{-1}	1.89×10^{-2}	2.67×10^{-3}
1500	1.41×10^{-1}	1.54×10^{-2}	2.18×10^{-3}
1550	1.44×10^{-1}	1.44×10^{-2}	2.08×10^{-3}
1600	1.49×10^{-1}	1.34×10^{-2}	2.01×10^{-3}
1700	1.50×10^{-1}	1.25×10^{-2}	1.87×10^{-3}
1800	1.52×10^{-1}	9.34×10^{-3}	1.42×10^{-3}
1900	1.53×10^{-1}	1.25×10^{-2}	1.92×10^{-3}
2000	1.56×10^{-1}	7.98×10^{-3}	1.25×10^{-3}

Table 12.21. Cutflow table for SR-HM1 using the $m(\text{LQ}_{21}^{\text{u}}) = 600$ GeV signal point, with a 50% branching fraction to $c\nu$, as the benchmark.

SR-HM1			
Requirement	Weighted Yields	Reduction [%]	Raw Yields
$\sigma \times \mathcal{L}$	28487.82	—	—
Preselection	998.23	−96.50	6656
Number of c -tagged jets (20% WP) ≥ 2	157.34	−84.24	1014
τ veto	155.87	−0.93	999
Leading jet is c -tagged	133.92	−14.08	824
$p_{\text{T}}(j_2) > 200$ GeV	70.40	−47.43	449
$m_{\text{T}}^{\text{min}}(c) > 300$ GeV	59.11	−16.04	381
$m_{\text{cc}} > 200$ GeV	55.21	−6.60	359
$E_{\text{T}}^{\text{miss}} \text{ Sig.} > 22$	8.54	−84.53	42

Table 12.22. Cutflow table for SR-HM2 using the $m(\text{LQ}_{21}^u) = 600$ GeV signal point, with a 50% branching fraction to $c\nu$, as the benchmark.

SR-HM2			
Requirement	Weighted Yields	Reduction [%]	Raw Yields
$\sigma \times \mathcal{L}$	28487.82	—	—
Preselection	998.23	−96.50	6656
Number of c -tagged jets (20% WP) ≥ 2	157.34	−84.24	1014
τ veto	155.87	−0.93	999
Leading jet is c -tagged	133.92	−14.08	824
$p_T(j_2) > 150$ GeV	104.07	−22.29	645
$m_T^{\min}(c) > 300$ GeV	88.31	−15.14	551
E_T^{miss} Sig. > 15	76.50	−13.37	454
E_T^{miss} Sig. < 22	63.30	−17.25	389
$m_{cc} > 200$ GeV	58.42	−7.71	363
$m_{cc} < 400$ GeV	22.96	−60.70	127

Table 12.23. Cutflow table for SR-HM3 using the $m(\text{LQ}_{21}^u) = 600$ GeV signal point, with a 50% branching fraction to $c\nu$, as the benchmark.

SR-HM3			
Requirement	Weighted Yields	Reduction [%]	Raw Yields
$\sigma \times \mathcal{L}$	28487.82	—	—
Preselection	998.23	−96.50	6656
Number of c -tagged jets (20% WP) ≥ 2	157.34	−84.24	1014
τ veto	155.87	−0.93	999
Leading jet is c -tagged	133.92	−14.08	824
$p_T(j_2) > 150$ GeV	104.07	−22.29	645
$m_T^{\min}(c) > 300$ GeV	88.31	−15.14	551
E_T^{miss} Sig. > 15	76.50	−13.37	454
E_T^{miss} Sig. < 22	63.30	−17.25	389
$m_{cc} > 400$ GeV	35.46	−43.98	236
$m_{cc} < 600$ GeV	24.28	−31.53	135

Table 12.24. Cutflow table for SR-HM-Disc using the $m(\text{LQ}_{21}^u) = 600$ GeV signal point, with a 50% branching fraction to $c\nu$, as the benchmark.

SR-HM-Disc			
Requirement	Weighted Yields	Reduction [%]	Raw Yields
$\sigma \times \mathcal{L}$	28487.82	—	—
Preselection	998.23	−96.50	6656
Number of c -tagged jets (20% WP) ≥ 2	157.34	−84.24	1014
τ veto	155.87	−0.93	999
Leading jet is c -tagged	133.92	−14.08	824
$p_T(j_2) > 150$ GeV	104.07	−22.29	645
$m_T^{\min}(c) > 300$ GeV	88.31	−15.14	551
E_T^{miss} Sig. > 15	76.50	−13.37	454
$m_{cc} > 200$ GeV	70.19	−8.25	421

D Charm-tagging efficiency




Figure D.9. Tagging efficiency for identifying jets that contain c -hadrons (upper panel) and c -jet tagging efficiency scale factors (SF_c) (bottom panel) for the 20% operating point (OP) of the DL1 r_c algorithm in bins of jet p_T as measured by the likelihood-based method and as evaluated from $t\bar{t}$ events generated using POWHEG BOX v2 plus PYTHIA 8. The 20% OP is defined with the c -tagging with b -veto technique as described in section 10.2. The b -tagging OP used for the veto is the 77% of the DL1 r algorithm. SFs are defined as the ratio between the measured and simulated efficiencies. The uncertainty on the SFs includes both statistical and systematic components but the statistical one is negligible thus not shown independently.




Figure D.10. Uncertainties on data vs MC simulation scale factors to correct the mistag rate of light jets (SF_{light}) to be identified as c -jets. These scale factors are derived using the full Run 2 dataset recorded by ATLAS. The dashed vertical lines indicate the bin boundaries while the coloured bars show the different components of the uncertainty.

E Exclusion contours for individual signal regions




Figure E.11. Exclusion limits at 95% CL for pair-production of top squarks decaying into charm quarks and neutralinos when only considering SR-HM1. The expected limit is shown as a dashed line while its $\pm 1\sigma$ uncertainty is shown as a yellow band. The observed limit is shown as a red line and the effect of varying the signal cross-sections by $\pm 1\sigma$ of its predicted theoretical uncertainty is shown as red dashed lines. The contours are shown as two-dimensional projections for masses of squarks and neutralinos.




Figure E.12. Exclusion limits at 95% CL for pair-production of top squarks decaying into charm quarks and neutralinos when only considering SR-HM2. The expected limit is shown as a dashed line while its $\pm 1\sigma$ uncertainty is shown as a yellow band. The observed limit is shown as a red line and the effect of varying the signal cross-sections by $\pm 1\sigma$ of its predicted theoretical uncertainty is shown as red dashed lines. The contours are shown as two-dimensional projections for masses of squarks and neutralinos.

E. Exclusion contours for individual signal regions




Figure E.13. Exclusion limits at 95% CL for pair-production of top squarks decaying into charm quarks and neutralinos when only considering SR-HM3. The expected limit is shown as a dashed line while its $\pm 1\sigma$ uncertainty is shown as a yellow band. The observed limit is shown as a red line and the effect of varying the signal cross-sections by $\pm 1\sigma$ of its predicted theoretical uncertainty is shown as red dashed lines. The contours are shown as two-dimensional projections for masses of squarks and neutralinos.




Figure E.14. Exclusion limits at 95% CL for pair-production of top squarks decaying into charm quarks and neutralinos when only considering SR-Comp1. The expected limit is shown as a dashed line while its $\pm 1\sigma$ uncertainty is shown as a yellow band. The observed limit is shown as a red line and the effect of varying the signal cross-sections by $\pm 1\sigma$ of its predicted theoretical uncertainty is shown as red dashed lines. The contours are shown as two-dimensional projections for masses of squarks and the differences between squark and neutralino masses.




Figure E.15. Exclusion limits at 95% CL for pair-production of top squarks decaying into charm quarks and neutralinos when only considering SR-Comp2. The expected limit is shown as a dashed line while its $\pm 1\sigma$ uncertainty is shown as a yellow band. The observed limit is shown as a red line and the effect of varying the signal cross-sections by $\pm 1\sigma$ of its predicted theoretical uncertainty is shown as red dashed lines. The contours are shown as two-dimensional projections for masses of squarks and the differences between squark and neutralino masses.




Figure E.16. Exclusion limits at 95% CL for pair-production of top squarks decaying into charm quarks and neutralinos when only considering SR-Comp3. The expected limit is shown as a dashed line while its $\pm 1\sigma$ uncertainty is shown as a yellow band. The observed limit is shown as a red line and the effect of varying the signal cross-sections by $\pm 1\sigma$ of its predicted theoretical uncertainty is shown as red dashed lines. The contours are shown as two-dimensional projections for masses of squarks and the differences between squark and neutralino masses.

E. Exclusion contours for individual signal regions




Figure E.17. Exclusion limits at 95% CL for pair-production of top squarks decaying into charm quarks and neutralinos when only considering SR-Comp-1c. The expected limit is shown as a dashed line while its $\pm 1\sigma$ uncertainty is shown as a yellow band. The observed limit is shown as a red line and the effect of varying the signal cross-sections by $\pm 1\sigma$ of its predicted theoretical uncertainty is shown as red dashed lines. The contours are shown as two-dimensional projections for masses of squarks and the differences between squark and neutralino masses.

F Best expected signal region combination




Figure F.18. Combination of signal regions with the best expected exclusion limits at 95% CL for pair-production of top squarks decaying into charm quarks and neutralinos. The two possible configurations either combine the High-Mass regions with SR-Comp1, SR-Comp2, and SR-Comp-1c (blue dots) or combine the High-Mass regions with SR-Comp3 and SR-Comp-1c (green dots).

G CL_S values




Figure G.19. Best (of the two fit configurations) expected CL_S values for pair-production of top squarks decaying into charm quarks and neutralinos.

H. Upper cross-section limits




Figure G.20. Best (of the two fit configurations using the expected CL_s) observed CL_s values for pair-production of top squarks decaying into charm quarks and neutralinos.

H Upper cross-section limits




Figure H.21. Upper limits on the observed cross-section at 95% CL for pair-production of top squarks decaying into charm quarks and neutralinos for the fit with the best expected CL_s .

Bibliography

- [1] ATLAS collaboration, *Search for supersymmetry in final states with missing transverse momentum and charm-tagged jets using 139 fb^{-1} of proton-proton collisions at $\sqrt{s} = 13\text{ TeV}$ with the ATLAS detector*, [2410.17824](#).
- [2] G. 't Hooft, *Why Do We Need Local Gauge Invariance in Theories With Vector Particles? An Introduction*, *NATO Sci. Ser. B* **59** (1980) 101.
- [3] BESIII collaboration, *Determination of Spin-Parity Quantum Numbers of $X(2370)$ as 0^{-+} from $J/\psi \rightarrow \gamma KS^0 KS^0 \eta'$* , *Phys. Rev. Lett.* **132** (2024) 181901 [[2312.05324](#)].
- [4] A. Bosma, *21-cm line studies of spiral galaxies. 2. The distribution and kinematics of neutral hydrogen in spiral galaxies of various morphological types.*, *Astron. J.* **86** (1981) 1825.
- [5] V. C. Rubin, D. Burstein, W. K. Ford, Jr. and N. Thonnard, *Rotation velocities of 16 SA galaxies and a comparison of Sa, Sb, and SC rotation properties*, *Astrophys. J.* **289** (1985) 81.
- [6] PLANCK collaboration, *Planck 2018 results. VI. Cosmological parameters*, *Astron. Astrophys.* **641** (2020) A6 [[1807.06209](#)].
- [7] SUPER-KAMIOKANDE collaboration, *Evidence for oscillation of atmospheric neutrinos*, *Phys. Rev. Lett.* **81** (1998) 1562 [[hep-ex/9807003](#)].
- [8] SNO collaboration, *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.* **87** (2001) 071301 [[nucl-ex/0106015](#)].
- [9] OPERA collaboration, *Observation of a first ν_τ candidate in the OPERA experiment in the CNGS beam*, *Phys. Lett. B* **691** (2010) 138 [[1006.1623](#)].
- [10] DOUBLE CHOOZ collaboration, *Indication of Reactor $\bar{\nu}_e$ Disappearance in the Double Chooz Experiment*, *Phys. Rev. Lett.* **108** (2012) 131801 [[1112.6353](#)].
- [11] DAYA BAY collaboration, *Observation of electron-antineutrino disappearance at Daya Bay*, *Phys. Rev. Lett.* **108** (2012) 171803 [[1203.1669](#)].
- [12] T2K collaboration, *Evidence of Electron Neutrino Appearance in a Muon Neutrino Beam*, *Phys. Rev. D* **88** (2013) 032002 [[1304.0841](#)].
- [13] S. P. Martin, *A Supersymmetry primer*, *Adv. Ser. Direct. High Energy Phys.* **18** (1998) 1 [[hep-ph/9709356](#)].
- [14] J. Alwall, P. C. Schuster and N. Toro, *Simplified models for a first characterization of new physics at the LHC*, *Phys. Rev. D* **79** (2009) 075020 [[0810.3921](#)].
- [15] LHC NEW PHYSICS WORKING GROUP collaboration, *Simplified Models for LHC New Physics Searches*, *J. Phys. G* **39** (2012) 105005 [[1105.2838](#)].
- [16] A. Chakraborty, M. Endo, B. Fuks, B. Herrmann, M. M. Nojiri, P. Pani et al., *Flavour-violating decays of mixed top-charm squarks at the LHC*, *Eur. Phys. J. C* **78** (2018) 844 [[1808.07488](#)].

Bibliography

- [17] BABAR collaboration, *Measurement of an Excess of $\bar{B} \rightarrow D^{(*)}\tau^-\bar{\nu}_\tau$ Decays and Implications for Charged Higgs Bosons*, *Phys. Rev. D* **88** (2013) 072012 [[1303.0571](#)].
- [18] LHCb collaboration, *Test of lepton universality using $B^+ \rightarrow K^+\ell^+\ell^-$ decays*, *Phys. Rev. Lett.* **113** (2014) 151601 [[1406.6482](#)].
- [19] LHCb collaboration, *Measurement of the ratio of branching fractions $\mathcal{B}(\bar{B}^0 \rightarrow D^{*+}\tau^-\bar{\nu}_\tau)/\mathcal{B}(\bar{B}^0 \rightarrow D^{*+}\mu^-\bar{\nu}_\mu)$* , *Phys. Rev. Lett.* **115** (2015) 111803 [[1506.08614](#)].
- [20] BELLE collaboration, *Measurement of the τ lepton polarization and $R(D^*)$ in the decay $\bar{B} \rightarrow D^*\tau^-\bar{\nu}_\tau$* , *Phys. Rev. Lett.* **118** (2017) 211801 [[1612.00529](#)].
- [21] LHCb collaboration, *Test of lepton universality with $B^0 \rightarrow K^{*0}\ell^+\ell^-$ decays*, *JHEP* **08** (2017) 055 [[1705.05802](#)].
- [22] LHCb collaboration, *Test of Lepton Flavor Universality by the measurement of the $B^0 \rightarrow D^{*-}\tau^+\nu_\tau$ branching fraction using three-prong τ decays*, *Phys. Rev. D* **97** (2018) 072013 [[1711.02505](#)].
- [23] T. Mandal, S. Mitra and S. Seth, *Single Productions of Colored Particles at the LHC: An Example with Scalar Leptoquarks*, *JHEP* **07** (2015) 028 [[1503.04689](#)].
- [24] T. Mandal, S. Mitra and S. Seth, *Pair Production of Scalar Leptoquarks at the LHC to NLO Parton Shower Accuracy*, *Phys. Rev. D* **93** (2016) 035018 [[1506.07369](#)].
- [25] M. J. Baker, J. Fuentes-Martín, G. Isidori and M. König, *High- p_T signatures in vector-leptoquark models*, *Eur. Phys. J. C* **79** (2019) 334 [[1901.10480](#)].
- [26] A. Papaefstathiou, *How-to: write a parton-level Monte Carlo particle physics event generator*, *Eur. Phys. J. Plus* **135** (2020) 497 [[1412.4677](#)].
- [27] A. Buckley et al., *General-purpose event generators for LHC physics*, *Phys. Rept.* **504** (2011) 145 [[1101.2599](#)].
- [28] LHC Machine, *JINST* **3** (2008) S08001.
- [29] ATLAS collaboration, *The ATLAS Experiment at the CERN Large Hadron Collider*, *JINST* **3** (2008) S08003.
- [30] CMS collaboration, *The CMS Experiment at the CERN LHC*, *JINST* **3** (2008) S08004.
- [31] ALICE collaboration, *The ALICE experiment at the CERN LHC*, *JINST* **3** (2008) S08002.
- [32] LHCb collaboration, *The LHCb Detector at the LHC*, *JINST* **3** (2008) S08005.
- [33] ATLAS collaboration, *Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC*, *Phys. Lett. B* **716** (2012) 1 [[1207.7214](#)].
- [34] CMS collaboration, *Observation of a New Boson at a Mass of 125 GeV with the CMS Experiment at the LHC*, *Phys. Lett. B* **716** (2012) 30 [[1207.7235](#)].
- [35] ATLAS collaboration, *The ATLAS experiment at the CERN Large Hadron Collider: a description of the detector configuration for Run 3*, *JINST* **19** (2024) P05063 [[2305.16623](#)].

- [36] *High-Luminosity Large Hadron Collider (HL-LHC) : Preliminary Design Report*, CERN-2015-005.
- [37] ATLAS collaboration, *Evolution of Regional, Age and Gender Demographics in the ATLAS collaboration*, *PoS LHCP2022* (2023) 279.
- [38] ATLAS collaboration, *ATLAS inner detector: Technical design report*. Vol. 1, CERN-LHCC-97-16, ATLAS-TDR-4.
- [39] ATLAS collaboration, *ATLAS inner detector: Technical design report*. Vol. 2, CERN-LHCC-97-17.
- [40] A. Ahmad et al., *The Silicon microstrip sensors of the ATLAS semiconductor tracker*, *Nucl. Instrum. Meth. A* **578** (2007) 98.
- [41] ATLAS collaboration, *ATLAS liquid-argon calorimeter: Technical Design Report*, Technical design report. ATLAS. CERN, Geneva, 1996, [10.17181/CERN.FWRW.FOOQ](#).
- [42] ATLAS collaboration, *The ATLAS liquid argon electromagnetic calorimeter*, *Nucl. Instrum. Meth. A* **384** (1996) 230.
- [43] Y. Tayalati, *ATLAS electromagnetic calorimeter, construction and test of a large scale system*, *Nucl. Instrum. Methods Phys. Res., A* **525** (2004) 178.
- [44] ATLAS collaboration, *ATLAS tile calorimeter: Technical Design Report*, Technical design report. ATLAS. CERN, Geneva, 1996, [10.17181/CERN.JRBJ.7O28](#).
- [45] ATLAS, RD34 collaboration, *The ATLAS scintillating tiles hadronic calorimeter*.
- [46] ATLAS collaboration, *The ATLAS Tile Calorimeter*, tech. rep., CERN, Geneva, 2015. [10.1109/ANIMMA.2015.7465554](#).
- [47] A. Ferrari, *A study of the ATLAS barrel presampler performance*, tech. rep., CERN, Geneva, 1999.
- [48] ATLAS collaboration, *ATLAS muon spectrometer: Technical Design Report*, Technical design report. ATLAS. CERN, Geneva, 1997.
- [49] ATLAS collaboration, *ATLAS level-1 trigger: Technical Design Report*, Technical design report. ATLAS. CERN, Geneva, 1998.
- [50] ATLAS collaboration, P. Jenni, M. Nessi, M. Nordberg and K. Smith, *ATLAS high-level trigger, data-acquisition and controls: Technical Design Report*, Technical design report. ATLAS. CERN, Geneva, 2003.
- [51] J. G. Panduro Vazquez, *The ATLAS Data Acquisition System: from Run 1 to Run 2*, tech. rep., CERN, Geneva, 2016. [10.1016/j.nuclphysbps.2015.09.146](#).
- [52] ATLAS collaboration, *The ATLAS Data Acquisition system in LHC Run 2*, Tech. Rep. 3, CERN, Geneva, 2017. [10.1088/1742-6596/898/3/032017](#).
- [53] ATLAS collaboration, *The ATLAS Level-1 Topological Processor: experience and upgrade plans*, tech. rep., CERN, Geneva, 2020.

Bibliography

- [54] C. Schiavi, *ATLAS High-Level Trigger algorithms for Run-2 data-taking*, tech. rep., CERN, Geneva, 2015.
- [55] I. T, *ATLAS FTK: The Fast Tracker*, tech. rep., CERN, Geneva, 2014.
- [56] ATLAS collaboration, *The ATLAS High-Level Calorimeter Trigger in Run-2*, tech. rep., CERN, Geneva, 2019. 10.1088/1742-6596/1162/1/012038.
- [57] M. Elsing, M. Goossens, A. Nairz and G. Negri, *The ATLAS Tier-0: Overview and operational experience*, *J. Phys.: Conf. Ser.* **219** (2010) 072011.
- [58] E. Lund, L. Bugge, T. G. Cornelissen, M. Elsing, I. Gavrilenko, W. Liebig et al., *Refitting of combined inner detector and muon spectrometer tracks from Monte Carlo samples by using the Kalman fitter and the STEP algorithm in the ATLAS experiment*, tech. rep., CERN, Geneva, 2008.
- [59] ATLAS collaboration, *Monitoring and data quality assessment of the ATLAS liquid argon calorimeter*, *JINST* **9** (2014) P07024 [[1405.3768](#)].
- [60] ATLAS collaboration, *Topological cell clustering in the ATLAS calorimeters and its performance in LHC Run 1*, *Eur. Phys. J. C* **77** (2017) 490 [[1603.02934](#)].
- [61] S. Hassani, L. Chevalier, E. Lancon, J. F. Laporte, R. Nicolaidou and A. Ouraou, *A muon identification and combined reconstruction procedure for the ATLAS detector at the LHC using the (MUONBOY, STACO, MuTag) reconstruction packages*, *Nucl. Instrum. Meth. A* **572** (2007) 77.
- [62] D. Adams, K. A. Assamagan, M. Biglietti, G. Carlino, G. Cataldi, F. Conventi et al., *Track reconstruction in the ATLAS Muon Spectrometer with MOORE 007*, tech. rep., CERN, Geneva, 2003.
- [63] ATLAS collaboration, *Search for non-resonant Higgs boson pair production in final states with leptons, taus, and photons in pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector*, *JHEP* **08** (2024) 164 [[2405.20040](#)].
- [64] M. Cacciari, G. P. Salam and G. Soyez, *FastJet User Manual*, *Eur. Phys. J. C* **72** (2012) 1896 [[1111.6097](#)].
- [65] M. Cacciari, G. P. Salam and G. Soyez, *The anti- k_t jet clustering algorithm*, *JHEP* **04** (2008) 063 [[0802.1189](#)].
- [66] A. J. Larkoski, G. P. Salam and J. Thaler, *Energy Correlation Functions for Jet Substructure*, *JHEP* **06** (2013) 108 [[1305.0007](#)].
- [67] A. J. Larkoski, I. Moult and D. Neill, *Power Counting to Better Jet Observables*, *JHEP* **12** (2014) 009 [[1409.6298](#)].
- [68] S. Marzani, G. Soyez and M. Spannowsky, *Looking inside jets: an introduction to jet substructure and boosted-object phenomenology*, vol. 958. Springer, 2019, 10.1007/978-3-030-15709-8, [[1901.10342](#)].
- [69] J. Thaler and K. Van Tilburg, *Identifying Boosted Objects with N-subjettiness*, *JHEP* **03** (2011) 015 [[1011.2268](#)].

- [70] S. D. Ellis and D. E. Soper, *Successive combination jet algorithm for hadron collisions*, *Phys. Rev. D* **48** (1993) 3160 [[hep-ph/9305266](#)].
- [71] Y. L. Dokshitzer, G. D. Leder, S. Moretti and B. R. Webber, *Better jet clustering algorithms*, *JHEP* **08** (1997) 001 [[hep-ph/9707323](#)].
- [72] M. Wobisch and T. Wengler, *Hadronization corrections to jet cross-sections in deep inelastic scattering*, in *Workshop on Monte Carlo Generators for HERA Physics (Plenary Starting Meeting)*, pp. 270–279, 4, 1998, [hep-ph/9907280](#).
- [73] ATLAS collaboration, *Jet energy scale and resolution measured in proton–proton collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector*, *Eur. Phys. J. C* **81** (2021) 689 [[2007.02645](#)].
- [74] ATLAS collaboration, *Jet energy measurement with the ATLAS detector in proton-proton collisions at $\sqrt{s} = 7$ TeV*, *Eur. Phys. J. C* **73** (2013) 2304 [[1112.6426](#)].
- [75] ATLAS collaboration, *Electron and photon reconstruction and performance in ATLAS using a dynamical, topological cell clustering-based approach*, ATL-PHYS-PUB-2017-022.
- [76] ATLAS collaboration, *Electron and photon performance measurements with the ATLAS detector using the 2015–2017 LHC proton-proton collision data*, *JINST* **14** (2019) P12006 [[1908.00005](#)].
- [77] ATLAS collaboration, *Electron and photon energy calibration with the ATLAS detector using data collected in 2015 at $\sqrt{s} = 13$ TeV*, ATL-PHYS-PUB-2016-015.
- [78] ATLAS collaboration, *Electron reconstruction and identification in the ATLAS experiment using the 2015 and 2016 LHC proton-proton collision data at $\sqrt{s} = 13$ TeV*, *Eur. Phys. J. C* **79** (2019) 639 [[1902.04655](#)].
- [79] ATLAS collaboration, *Measurement of the photon identification efficiencies with the ATLAS detector using LHC Run 2 data collected in 2015 and 2016*, *Eur. Phys. J. C* **79** (2019) 205 [[1810.05087](#)].
- [80] ATLAS collaboration, *Muon reconstruction and identification efficiency in ATLAS using the full Run 2 pp collision data set at $\sqrt{s} = 13$ TeV*, *Eur. Phys. J. C* **81** (2021) 578 [[2012.00578](#)].
- [81] ATLAS collaboration, *Studies of the muon momentum calibration and performance of the ATLAS detector with pp collisions at $\sqrt{s} = 13$ TeV*, *Eur. Phys. J. C* **83** (2023) 686 [[2212.07338](#)].
- [82] ATLAS collaboration, *Reconstruction, Energy Calibration, and Identification of Hadronically Decaying Tau Leptons in the ATLAS Experiment for Run-2 of the LHC*, ATL-PHYS-PUB-2015-045.
- [83] ATLAS collaboration, *Identification and energy calibration of hadronically decaying tau leptons with the ATLAS experiment in pp collisions at $\sqrt{s}=8$ TeV*, *Eur. Phys. J. C* **75** (2015) 303 [[1412.7086](#)].
- [84] ATLAS collaboration, *Performance of missing transverse momentum reconstruction with the ATLAS detector using proton-proton collisions at $\sqrt{s} = 13$ TeV*, *Eur. Phys. J. C* **78** (2018) 903 [[1802.08168](#)].

Bibliography

- [85] ATLAS collaboration, *The performance of missing transverse momentum reconstruction and its significance with the ATLAS detector using 140 fb^{-1} of $\sqrt{s} = 13\text{ TeV}$ pp collisions*, [2402.05858](#).
- [86] ATLAS collaboration, *ATLAS flavour-tagging algorithms for the LHC Run 2 pp collision dataset*, *Eur. Phys. J. C* **83** (2023) 681 [[2211.16345](#)].
- [87] ATLAS collaboration, *Optimisation and performance studies of the ATLAS b -tagging algorithms for the 2017-18 LHC run*, ATL-PHYS-PUB-2017-013.
- [88] ATLAS collaboration, *Identification of Jets Containing b -Hadrons with Recurrent Neural Networks at the ATLAS Experiment*, ATL-PHYS-PUB-2017-003.
- [89] ATLAS collaboration, *Secondary vertex finding for jet flavour identification with the ATLAS detector*, ATL-PHYS-PUB-2017-011.
- [90] ATLAS collaboration, *Topological b -hadron decay reconstruction and identification of b -jets with the JetFitter package in the ATLAS experiment at the LHC*, ATL-PHYS-PUB-2018-025.
- [91] ATLAS collaboration, *Technical Design Report for the ATLAS Inner Tracker Pixel Detector*, CERN-LHCC-2017-021, ATLAS-TDR-030.
- [92] ATLAS collaboration, *Technical Design Report for the ATLAS Inner Tracker Strip Detector*, CERN-LHCC-2017-005, ATLAS-TDR-025.
- [93] ATLAS collaboration, *Technical Design Report for the Phase-II Upgrade of the ATLAS TDAQ System*, CERN-LHCC-2017-020, ATLAS-TDR-029.
- [94] ATLAS collaboration, *ATLAS Liquid Argon Calorimeter Phase-II Upgrade : Technical Design Report*, CERN-LHCC-2017-018, ATLAS-TDR-027.
- [95] ATLAS collaboration, *Technical Design Report for the Phase-II Upgrade of the ATLAS Tile Calorimeter*, CERN-LHCC-2017-019, ATLAS-TDR-028.
- [96] ATLAS collaboration, *Technical Design Report for the Phase-II Upgrade of the ATLAS Muon Spectrometer*, CERN-LHCC-2017-017, ATLAS-TDR-026.
- [97] ATLAS collaboration, *A High-Granularity Timing Detector for the ATLAS Phase-II Upgrade: Technical Design Report*, CERN-LHCC-2020-007, ATLAS-TDR-031.
- [98] ATLAS collaboration, *Graph Neural Network Jet Flavour Tagging with the ATLAS Detector*, ATL-PHYS-PUB-2022-027.
- [99] ATLAS collaboration, *Deep Learning for Pion Identification and Energy Calibration with the ATLAS Detector*, ATL-PHYS-PUB-2020-018.
- [100] ATLAS collaboration, *Point Cloud Deep Learning Methods for Pion Reconstruction in the ATLAS Experiment*, ATL-PHYS-PUB-2022-040.
- [101] ATLAS collaboration, *The application of neural networks for the calibration of topological cell clusters in the ATLAS calorimeters*, ATL-PHYS-PUB-2023-019.
- [102] ATLAS collaboration, *METNet: A combined missing transverse momentum working point using a neural network with the ATLAS detector*, *PoS EPS-HEP2021* (2022) 625.

- [103] S. Agostinelli et al., *GEANT4 – a simulation toolkit*, *Nucl. Instrum. Meth. A* **506** (2003) 250.
- [104] M. Faucci Giannelli and R. Zhang, *CaloShowerGAN, a generative adversarial network model for fast calorimeter shower simulation*, *Eur. Phys. J. Plus* **139** (2024) 597 [2309.06515].
- [105] ATLAS collaboration, *AtlFast3: The Next Generation of Fast Simulation in ATLAS*, *Comput. Softw. Big Sci.* **6** (2022) 7 [2109.02551].
- [106] CELLO collaboration, *An Analysis of the Charged and Neutral Energy Flow in e^+e^- Hadronic Annihilation at 34-GeV, and a Determination of the QCD Effective Coupling Constant*, *Phys. Lett. B* **113** (1982) 427.
- [107] ALEPH collaboration, *Performance of the ALEPH detector at LEP*, *Nucl. Instrum. Meth. A* **360** (1995) 481.
- [108] J.-C. Brient, *Improving the Jet Reconstruction with the Particle Flow Method; an Introduction*, 12, 2004, [physics/0412149](#), DOI.
- [109] CMS collaboration, *Particle-flow reconstruction and global event description with the CMS detector*, *JINST* **12** (2017) P10003 [1706.04965].
- [110] ATLAS collaboration, *Jet reconstruction and performance using particle flow with the ATLAS Detector*, *Eur. Phys. J. C* **77** (2017) 466 [1703.10485].
- [111] ATLAS collaboration, *Improving jet substructure performance in ATLAS using Track-Caloclusters*, ATL-PHYS-PUB-2017-015.
- [112] D. Krohn, J. Thaler and L.-T. Wang, *Jet Trimming*, *JHEP* **02** (2010) 084 [0912.1342].
- [113] ATLAS collaboration, *Optimisation of large-radius jet reconstruction for the ATLAS detector in 13 TeV proton–proton collisions*, *Eur. Phys. J. C* **81** (2021) 334 [2009.04986].
- [114] ATLAS collaboration, *Software and computing for Run 3 of the ATLAS experiment at the LHC*, 2404.06335.
- [115] F. A. Di Bello, S. Ganguly, E. Gross, M. Kado, M. Pitt, L. Santi et al., *Towards a Computer Vision Particle Flow*, *Eur. Phys. J. C* **81** (2021) 107 [2003.08863].
- [116] F. A. Di Bello et al., *Reconstructing particles in jets using set transformer and hypergraph prediction networks*, *Eur. Phys. J. C* **83** (2023) 596 [2212.01328].
- [117] P. Nason, *A new method for combining NLO QCD with shower Monte Carlo algorithms*, *JHEP* **11** (2004) 040 [hep-ph/0409146].
- [118] S. Frixione, P. Nason and C. Oleari, *Matching NLO QCD computations with parton shower simulations: the POWHEG method*, *JHEP* **11** (2007) 070 [0709.2092].
- [119] S. Alioli, P. Nason, C. Oleari and E. Re, *A general framework for implementing NLO calculations in shower Monte Carlo programs: the POWHEG BOX*, *JHEP* **06** (2010) 043 [1002.2581].
- [120] S. Alioli, P. Nason, C. Oleari and E. Re, *NLO vector-boson production matched with shower in POWHEG*, *JHEP* **07** (2008) 060 [0805.4802].

Bibliography

- [121] T. Sjöstrand, S. Mrenna and P. Skands, *A brief introduction to PYTHIA 8.1*, *Comput. Phys. Commun.* **178** (2008) 852 [0710.3820].
- [122] ATLAS collaboration, *Measurement of the Z/γ^* boson transverse momentum distribution in pp collisions at $\sqrt{s} = 7$ TeV with the ATLAS detector*, *JHEP* **09** (2014) 145 [1406.3660].
- [123] H.-L. Lai, M. Guzzi, J. Huston, Z. Li, P. M. Nadolsky, J. Pumplin et al., *New parton distributions for collider physics*, *Phys. Rev. D* **82** (2010) 074024 [1007.2241].
- [124] J. Pumplin et al., *New Generation of Parton Distributions with Uncertainties from Global QCD Analysis*, *JHEP* **07** (2002) 012 [hep-ph/0201195].
- [125] P. Golonka and Z. Was, *PHOTOS Monte Carlo: a precision tool for QED corrections in Z and W decays*, *Eur. Phys. J. C* **45** (2006) 97 [hep-ph/0506026].
- [126] N. Davidson, T. Przedzinski and Z. Was, *PHOTOS interface in C++: Technical and physics documentation*, *Comput. Phys. Commun.* **199** (2016) 86 [1011.0937].
- [127] D. J. Lange, *The EvtGen particle decay simulation package*, *Nucl. Instrum. Meth. A* **462** (2001) 152.
- [128] S. Frixione, G. Ridolfi and P. Nason, *A positive-weight next-to-leading-order Monte Carlo for heavy flavour hadroproduction*, *JHEP* **09** (2007) 126 [0707.3088].
- [129] NNPDF Collaboration, R. D. Ball et al., *Parton distributions for the LHC run II*, *JHEP* **04** (2015) 040 [1410.8849].
- [130] ATLAS collaboration, “Studies on top-quark Monte Carlo modelling for Top2016.” ATL-PHYS-PUB-2016-020, 2016.
- [131] T. Sjöstrand, S. Ask, J. R. Christiansen, R. Corke, N. Desai, P. Ilten et al., *An Introduction to PYTHIA 8.2*, *Comput. Phys. Commun.* **191** (2015) 159 [1410.3012].
- [132] ATLAS collaboration, “ATLAS Pythia 8 tunes to 7 TeV data.” ATL-PHYS-PUB-2014-021, 2014.
- [133] NNPDF Collaboration, R. D. Ball et al., *Parton distributions with LHC data*, *Nucl. Phys. B* **867** (2013) 244 [1207.1303].
- [134] E. Bothmann et al., *Event generation with Sherpa 2.2*, *SciPost Phys.* **7** (2019) 034 [1905.09127].
- [135] T. Gleisberg and S. Höche, *Comix, a new matrix element generator*, *JHEP* **12** (2008) 039 [0808.3674].
- [136] F. Buccioni, J.-N. Lang, J. M. Lindert, P. Maierhöfer, S. Pozzorini, H. Zhang et al., *OpenLoops 2*, *Eur. Phys. J. C* **79** (2019) 866 [1907.13071].
- [137] F. Cascioli, P. Maierhöfer and S. Pozzorini, *Scattering Amplitudes with Open Loops*, *Phys. Rev. Lett.* **108** (2012) 111601 [1111.5206].
- [138] A. Denner, S. Dittmaier and L. Hofer, *COLLIER: A fortran-based complex one-loop library in extended regularizations*, *Comput. Phys. Commun.* **212** (2017) 220 [1604.06792].

- [139] S. Schumann and F. Krauss, *A parton shower algorithm based on Catani–Seymour dipole factorisation*, *JHEP* **03** (2008) 038 [[0709.1027](#)].
- [140] S. Höche, F. Krauss, M. Schönherr and F. Siegert, *A critical appraisal of NLO+PS matching methods*, *JHEP* **09** (2012) 049 [[1111.1220](#)].
- [141] S. Höche, F. Krauss, M. Schönherr and F. Siegert, *QCD matrix elements + parton showers. The NLO case*, *JHEP* **04** (2013) 027 [[1207.5030](#)].
- [142] S. Catani, F. Krauss, B. R. Webber and R. Kuhn, *QCD Matrix Elements + Parton Showers*, *JHEP* **11** (2001) 063 [[hep-ph/0109231](#)].
- [143] S. Höche, F. Krauss, S. Schumann and F. Siegert, *QCD matrix elements and truncated showers*, *JHEP* **05** (2009) 053 [[0903.1219](#)].
- [144] F. Siegert, *A practical guide to event generation for prompt photon production with Sherpa*, *J. Phys. G* **44** (2017) 044007 [[1611.07226](#)].
- [145] S. Frixione, *Isolated photons in perturbative QCD*, *Phys. Lett. B* **429** (1998) 369 [[hep-ph/9801442](#)].
- [146] ATLAS collaboration, *Performance of pile-up mitigation techniques for jets in pp collisions at $\sqrt{s} = 8$ TeV using the ATLAS detector*, *Eur. Phys. J. C* **76** (2016) 581 [[1510.03823](#)].
- [147] ATLAS collaboration, *ATLAS b-jet identification performance and efficiency measurement with $t\bar{t}$ events in pp collisions at $\sqrt{s} = 13$ TeV*, *Eur. Phys. J. C* **79** (2019) 970 [[1907.05120](#)].
- [148] CMS collaboration, *The Jet Plus Tracks Algorithm for Calorimeter Jet Energy Corrections in CMS*, tech. rep., 2009.
- [149] CMS collaboration, *Performance of Jet-Plus-Tracks algorithm in Run I*, tech. rep., CERN, Geneva, 2015.
- [150] ATLAS collaboration, “Track Reconstruction Performance of the ATLAS Inner Detector at $\sqrt{s} = 13$ TeV.” ATL-PHYS-PUB-2015-018, 2015.
- [151] ATLAS collaboration, *Athena*, Apr., 2019. [10.5281/zenodo.2641997](#).
- [152] ATLAS collaboration, “Multijet simulation for 13 TeV ATLAS Analyses.” ATL-PHYS-PUB-2019-017, 2019.
- [153] ATLAS collaboration, “The Pythia 8 A3 tune description of ATLAS minimum bias and inelastic measurements incorporating the Donnachie–Landshoff diffractive model.” ATL-PHYS-PUB-2016-017, 2016.
- [154] ATLAS collaboration, *Identification and rejection of pile-up jets at high pseudorapidity with the ATLAS detector*, *Eur. Phys. J. C* **77** (2017) 580 [[1705.02211](#)].
- [155] ATLAS collaboration, “Forward jet vertex tagging using the particle flow algorithm.” ATL-PHYS-PUB-2019-026, 2019.
- [156] Y. Golfand and E. Likhtman, *Extension of the Algebra of Poincare Group Generators and Violation of P Invariance*, *JETP Lett.* **13** (1971) 323.

Bibliography

- [157] D. Volkov and V. Akulov, *Is the neutrino a goldstone particle?*, *Phys. Lett. B* **46** (1973) 109.
- [158] J. Wess and B. Zumino, *Supergauge transformations in four dimensions*, *Nucl. Phys. B* **70** (1974) 39.
- [159] J. Wess and B. Zumino, *Supergauge invariant extension of quantum electrodynamics*, *Nucl. Phys. B* **78** (1974) 1.
- [160] S. Ferrara and B. Zumino, *Supergauge invariant Yang-Mills theories*, *Nucl. Phys. B* **79** (1974) 413.
- [161] A. Salam and J. Strathdee, *Super-symmetry and non-Abelian gauges*, *Phys. Lett. B* **51** (1974) 353.
- [162] N. Sakai, *Naturalness in supersymmetric GUTS*, *Z. Phys. C* **11** (1981) 153.
- [163] S. Dimopoulos, S. Raby and F. Wilczek, *Supersymmetry and the scale of unification*, *Phys. Rev. D* **24** (1981) 1681.
- [164] L. E. Ibáñez and G. G. Ross, *Low-energy predictions in supersymmetric grand unified theories*, *Phys. Lett. B* **105** (1981) 439.
- [165] S. Dimopoulos and H. Georgi, *Softly broken supersymmetry and SU(5)*, *Nucl. Phys. B* **193** (1981) 150.
- [166] G. R. Farrar and P. Fayet, *Phenomenology of the production, decay, and detection of new hadronic states associated with supersymmetry*, *Phys. Lett. B* **76** (1978) 575.
- [167] H. Goldberg, *Constraint on the Photino Mass from Cosmology*, *Phys. Rev. Lett.* **50** (1983) 1419.
- [168] J. Ellis, J. Hagelin, D. V. Nanopoulos, K. A. Olive and M. Srednicki, *Supersymmetric relics from the big bang*, *Nucl. Phys. B* **238** (1984) 453.
- [169] R. Mahbubani, M. Papucci, G. Perez, J. T. Ruderman and A. Weiler, *Light Nondegenerate Squarks at the LHC*, *Phys. Rev. Lett.* **110** (2013) 151804 [1212.3328].
- [170] R. Gröber, M. Mühlleitner, E. Popena and A. Wlotzka, *Light stop decays into $Wb\tilde{\chi}_1^0$ near the kinematic threshold*, *Phys. Lett. B* **747** (2015) 144 [1502.05935].
- [171] S. Dimopoulos and L. Susskind, *Mass without scalars*, *Nucl. Phys. B* **155** (1979) 237.
- [172] S. Dimopoulos, *Technicoloured signatures*, *Nucl. Phys. B* **168** (1980) 69.
- [173] E. Eichten and K. Lane, *Dynamical breaking of weak interaction symmetries*, *Phys. Lett. B* **90** (1980) 125.
- [174] V. Angelopoulos, J. Ellis, H. Kowalski, D. Nanopoulos, N. Tracas and F. Zwirner, *Search for new quarks suggested by the superstring*, *Nucl. Phys. B* **292** (1987) 59.
- [175] W. Buchmüller and D. Wyler, *Constraints on su(5)-type leptoquarks*, *Phys. Lett. B* **177** (1986) 377.
- [176] J. C. Pati and A. Salam, *Lepton number as the fourth "color"*, *Phys. Rev. D* **10** (1974) 275.

- [177] H. Georgi and S. L. Glashow, *Unity of all elementary-particle forces*, *Phys. Rev. Lett.* **32** (1974) 438.
- [178] B. Diaz, M. Schmaltz and Y.-M. Zhong, *The leptoquark Hunter’s guide: Pair production*, *JHEP* **10** (2017) 097 [1706.05033].
- [179] A. Belyaev, C. Leroy, R. Mehdiyev and A. Pukhov, *Leptoquark single and pair production at LHC with CalcHEP/CompHEP in the complete model*, *JHEP* **09** (2005) 005.
- [180] J. Blumlein, E. Boos and A. Kryukov, *Leptoquark pair production in hadronic interactions*, *Z. Phys. C* **76** (1997) 137 [hep-ph/9610408].
- [181] M. Krämer, T. Plehn, M. Spira and P. M. Zerwas, *Pair production of scalar leptoquarks at the CERN LHC*, *Phys. Rev. D* **71** (2005) 057503 [hep-ph/0411038].
- [182] M. Krämer, T. Plehn, M. Spira and P. M. Zerwas, *Pair production of scalar leptoquarks at the Fermilab Tevatron*, *Phys. Rev. Lett.* **79** (1997) 341 [hep-ph/9704322].
- [183] ATLAS collaboration, *Search for supersymmetry in final states with charm jets and missing transverse momentum in 13 TeV pp collisions with the ATLAS detector*, *JHEP* **09** (2018) 050 [1805.01649].
- [184] P. Jackson and C. Rogan, *Recursive Jigsaw Reconstruction: HEP event analysis in the presence of kinematic and combinatoric ambiguities*, *Phys. Rev. D* **96** (2017) 112007 [1705.10733].
- [185] ATLAS collaboration, *Search for new phenomena in events with an energetic jet and missing transverse momentum in pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector*, *Phys. Rev. D* **103** (2021) 112006 [2102.10874].
- [186] ATLAS collaboration, *Search for squarks and gluinos in final states with jets and missing transverse momentum using 139 fb⁻¹ of $\sqrt{s} = 13$ TeV pp collision data with the ATLAS detector*, *JHEP* **02** (2021) 143 [2010.14293].
- [187] ATLAS collaboration, *A search for top-squark pair production, in final states containing a top quark, a charm quark and missing transverse momentum, using the 139 fb⁻¹ of pp collision data collected by the atlas detector*, *JHEP* **07** (2024) 250 [2402.12137].
- [188] ATLAS collaboration, *ATLAS data quality operations and performance for 2015–2018 data-taking*, *JINST* **15** (2020) P04003 [1911.04632].
- [189] ATLAS collaboration, “Luminosity determination in pp collisions at $\sqrt{s} = 13$ TeV using the ATLAS detector at the LHC.” ATLAS-CONF-2019-021, 2019.
- [190] G. Avoni et al., *The new LUCID-2 detector for luminosity measurement and monitoring in ATLAS*, *JINST* **13** (2018) P07017.
- [191] ATLAS collaboration, *Performance of the missing transverse momentum triggers for the ATLAS detector during Run-2 data taking*, *JHEP* **08** (2020) 080 [2005.09554].
- [192] ATLAS collaboration, *The ATLAS Simulation Infrastructure*, *Eur. Phys. J. C* **70** (2010) 823 [1005.4568].

Bibliography

- [193] ATLAS collaboration, “The new Fast Calorimeter Simulation in ATLAS.” ATL-SOFT-PUB-2018-002, 2018.
- [194] T. Gleisberg, S. Höche, F. Krauss, M. Schönherr, S. Schumann, F. Siegert et al., *Event generation with SHERPA 1.1*, *JHEP* **02** (2009) 007 [[0811.4622](#)].
- [195] ATLAS collaboration, *Modelling and computational improvements to the simulation of single vector-boson plus jet processes for the ATLAS experiment*, *JHEP* **08** (2022) 089 [[2112.09588](#)].
- [196] J. Alwall, R. Frederix, S. Frixione, V. Hirschi, F. Maltoni, O. Mattelaer et al., *The automated computation of tree-level and next-to-leading order differential cross sections, and their matching to parton shower simulations*, *JHEP* **07** (2014) 079 [[1405.0301](#)].
- [197] M. Czakon, P. Fiedler and A. Mitov, *Total Top-Quark Pair-Production Cross Section at Hadron Colliders Through $O(\alpha_s^4)$* , *Phys. Rev. Lett.* **110** (2013) 252004 [[1303.6254](#)].
- [198] M. Czakon and A. Mitov, *NNLO corrections to top pair production at hadron colliders: the quark-gluon reaction*, *JHEP* **01** (2013) 080 [[1210.6832](#)].
- [199] M. Czakon and A. Mitov, *NNLO corrections to top-pair production at hadron colliders: the all-fermionic scattering channels*, *JHEP* **12** (2012) 054 [[1207.0236](#)].
- [200] P. Bärnreuther, M. Czakon and A. Mitov, *Percent-Level-Precision Physics at the Tevatron: Next-to-Next-to-Leading Order QCD Corrections to $q\bar{q} \rightarrow t\bar{t} + X$* , *Phys. Rev. Lett.* **109** (2012) 132001 [[1204.5201](#)].
- [201] M. Cacciari, M. Czakon, M. Mangano, A. Mitov and P. Nason, *Top-pair production at hadron colliders with next-to-next-to-leading logarithmic soft-gluon resummation*, *Phys. Lett. B* **710** (2012) 612 [[1111.5869](#)].
- [202] M. Czakon and A. Mitov, *Top++: A program for the calculation of the top-pair cross-section at hadron colliders*, *Comput. Phys. Commun.* **185** (2014) 2930 [[1112.5675](#)].
- [203] N. Kidonakis, *Next-to-next-to-leading-order collinear and soft gluon corrections for t-channel single top quark production*, *Phys. Rev. D* **83** (2011) 091503 [[1103.2792](#)].
- [204] N. Kidonakis, *Two-loop soft anomalous dimensions for single top quark associated production with a W^- or H^-* , *Phys. Rev. D* **82** (2010) 054018 [[1005.4451](#)].
- [205] N. Kidonakis, *Next-to-next-to-leading logarithm resummation for s-channel single top quark production*, *Phys. Rev. D* **81** (2010) 054028 [[1001.5034](#)].
- [206] L. Lönnblad and S. Prestel, *Matching tree-level matrix elements with interleaved showers*, *JHEP* **03** (2012) 019 [[1109.4829](#)].
- [207] W. Beenakker, C. Borschensky, M. Krämer, A. Kulesza and E. Laenen, *NNLL-fast: predictions for coloured supersymmetric particle production at the LHC with threshold and Coulomb resummation*, *JHEP* **12** (2016) 133 [[1607.07741](#)].
- [208] W. Beenakker, M. Krämer, T. Plehn, M. Spira and P. M. Zerwas, *Stop production at hadron colliders*, *Nucl. Phys. B* **515** (1998) 3 [[hep-ph/9710451](#)].

- [209] W. Beenakker, S. Brensing, M. Kramer, A. Kulesza, E. Laenen and I. Niessen, *Supersymmetric top and bottom squark production at hadron colliders*, *JHEP* **08** (2010) 098 [[1006.4771](#)].
- [210] W. Beenakker, C. Borschensky, R. Heger, M. Krämer, A. Kulesza and E. Laenen, *NNLL resummation for stop pair-production at the LHC*, *JHEP* **05** (2016) 153 [[1601.02954](#)].
- [211] C. Borschensky, B. Fuks, A. Kulesza and D. Schwartländer, *Scalar leptoquark pair production at hadron colliders*, *Phys. Rev. D* **101** (2020) 115017 [[2002.08971](#)].
- [212] ATLAS collaboration, “Vertex Reconstruction Performance of the ATLAS Detector at $\sqrt{s} = 13$ TeV.” ATL-PHYS-PUB-2015-026, 2015.
- [213] ATLAS collaboration, “ E_T^{miss} performance in the ATLAS detector using 2015–2016 LHC pp collisions.” ATLAS-CONF-2018-023, 2018.
- [214] ATLAS collaboration, *Measurement of the c -jet mistagging efficiency in $t\bar{t}$ events using pp collision data at $\sqrt{s} = 13$ TeV collected with the ATLAS detector*, *Eur. Phys. J. C* **82** (2022) 95 [[2109.10627](#)].
- [215] ATLAS collaboration, *Calibration of the light-flavour jet mistagging efficiency of the b -tagging algorithms with Z +jets events using 139 fb^{-1} of ATLAS proton–proton collision data at $\sqrt{s} = 13$ TeV*, *Eur. Phys. J. C* **83** (2023) 728 [[2301.06319](#)].
- [216] ATLAS collaboration, “Identification of hadronic tau lepton decays using neural networks in the ATLAS experiment.” ATL-PHYS-PUB-2019-033, 2019.
- [217] ATLAS collaboration, “Object-based missing transverse momentum significance in the ATLAS Detector.” ATLAS-CONF-2018-038, 2018.
- [218] P. Jackson, C. Rogan and M. Santoni, *Sparticles in motion: Analyzing compressed SUSY scenarios with a new method of event reconstruction*, *Phys. Rev. D* **95** (2017) 035031 [[1607.08307](#)].
- [219] ATLAS collaboration, *Searches for new phenomena in events with two leptons, jets, and missing transverse momentum in 139 fb^{-1} of $\sqrt{s} = 13$ TeV pp collisions with the ATLAS detector*, *Eur. Phys. J. C* **83** (2023) 515 [[2204.13072](#)].
- [220] L. Heinrich, M. Feickert, G. Stark and K. Cranmer, *pyhf: pure-python implementation of histfactory statistical models*, *Journal of Open Source Software* **6** (2021) 2823.
- [221] L. Heinrich, M. Feickert and G. Stark, “pyhf: v0.7.1.” [10.5281/zenodo.1169739](#).
- [222] ROOT COLLABORATION collaboration, *HistFactory: A tool for creating statistical models for use with RooFit and RooStats*, tech. rep., New York U., New York, Jan, 2012.
- [223] K. Cranmer and A. Held, *Building and steering binned template fits with cabinetry*, *EPJ Web Conf.* **251** (2021) 03067.
- [224] ATLAS collaboration, *Measurements of inclusive and differential fiducial cross-sections of $t\bar{t}$ production with additional heavy-flavour jets in proton–proton collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector*, *JHEP* **04** (2019) 046 [[1811.12113](#)].

Bibliography

- [225] G. Cowan, K. Cranmer, E. Gross and O. Vitells, *Asymptotic formulae for likelihood-based tests of new physics*, *Eur. Phys. J. C* **71** (2011) 1554 [[1007.1727](#)].
- [226] R. D. Cousins, J. T. Linnemann and J. Tucker, *Evaluation of three methods for calculating statistical significance when incorporating a systematic uncertainty into a test of the background-only hypothesis for a Poisson process*, *Nucl. Instrum. Meth. A* **595** (2008) 480 [[physics/0702156](#)].
- [227] A. L. Read, *Presentation of search results: The CL_s technique*, *J. Phys. G* **28** (2002) 2693.
- [228] ATLAS collaboration, *SimpleAnalysis: Truth-level Analysis Framework*, ATL-PHYS-PUB-2022-017.