Long-Lived Particles

A Search for Displaced τ-Leptons with the ATLAS Experiment

Jack Gargan — The University of Edinburgh

Particle Physics Ph.D Thesis

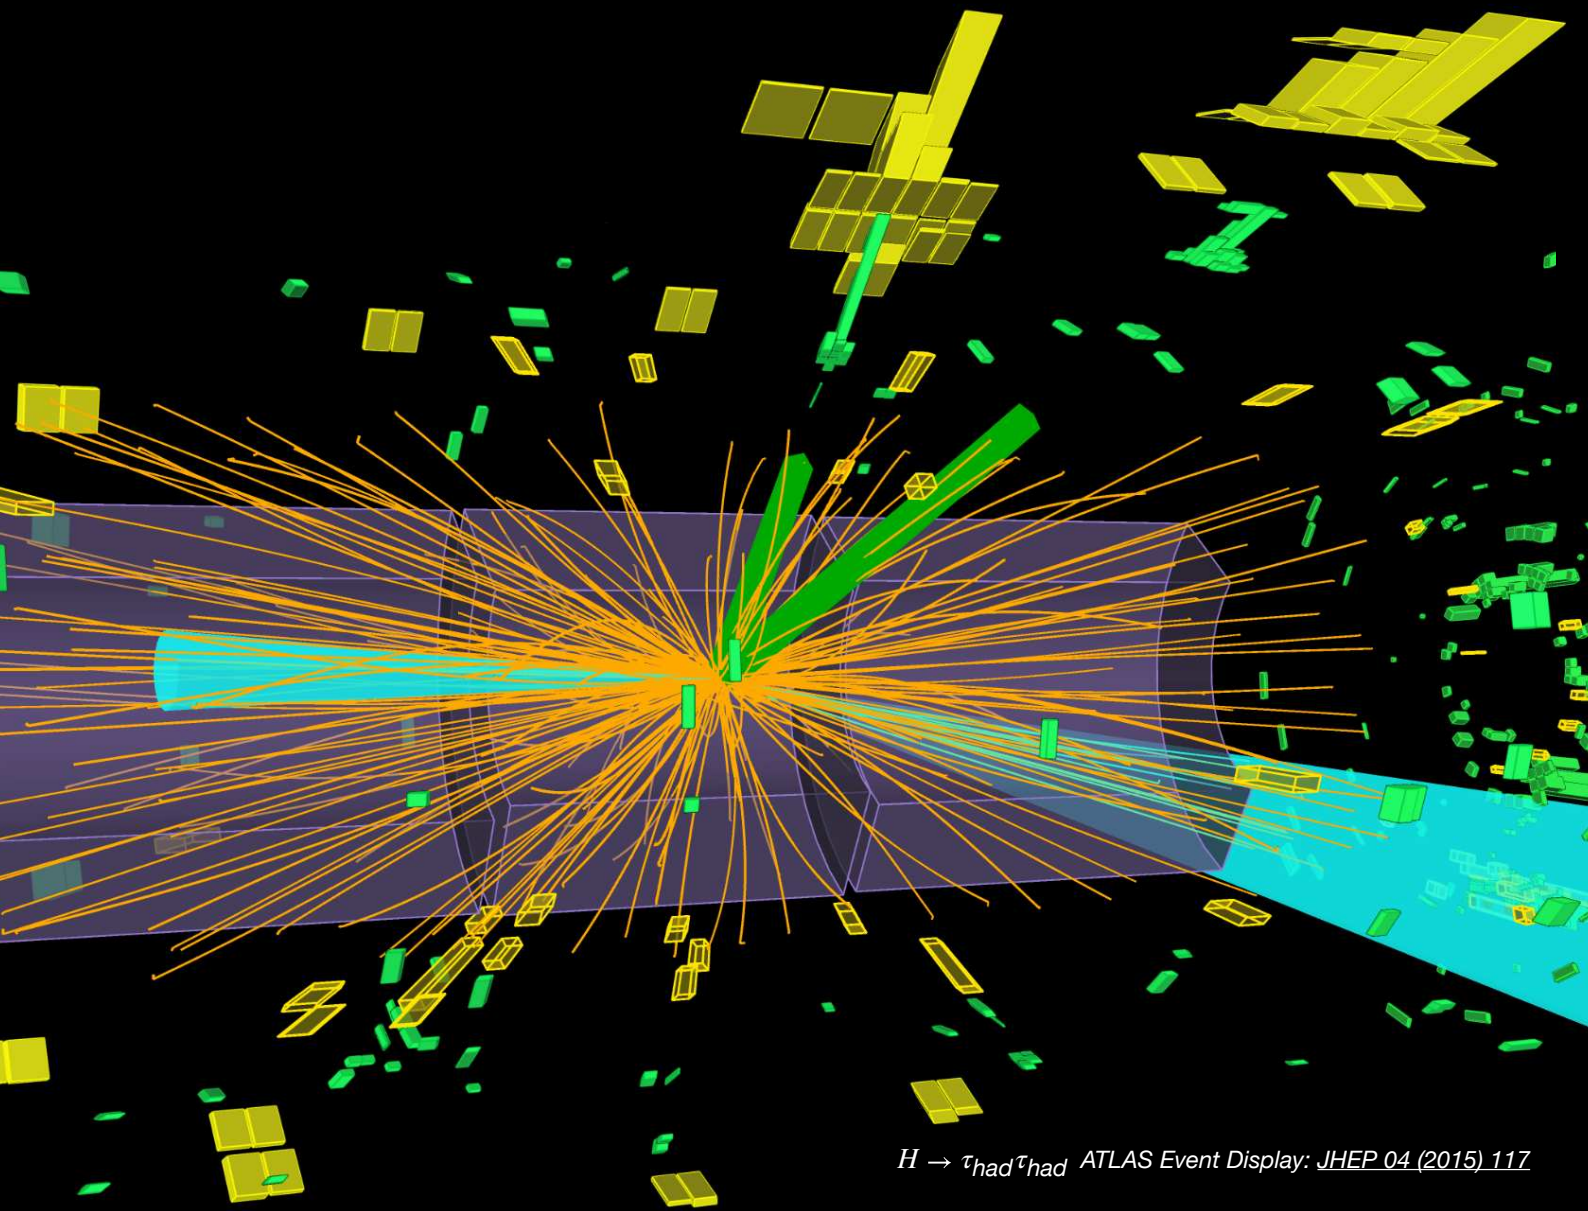

$H \rightarrow \tau_{had}\tau_{had}$ *ATLAS Event Display: JHEP 04 (2015) 117*

Long-Lived Particles Beyond the Standard Model

A Model-Independent Search for Hadronically-Decaying Displaced τ-Leptons in Run-2 LHC Data with the ATLAS Experiment

Jack Gargan

Doctor of Philosophy
The University of Edinburgh
October 2024

To my parents,
for their unsparing love and support,

and to Charly,
for always being by my side.

Thoroughly conscious ignorance is the prelude to every real advance in science.

— James Clerk Maxwell

Scotland's Greatest Scientist

The Path to Discovery

The earliest days in the arc of our species are characterised by moments of discovery: the discovery of fire, of the utility of primitive tools, of the art of agriculture. Such serendipitous events in our history planted the seeds for the development of our civilisation: of every advance in our technological capabilities and our shared wealth of knowledge. A drive to discover and to understand was born — a drive which would become innate to us all, as the resultant enhancement of our survival prospects compels evolution to embed it deep into our internal wiring diagram. Curiosity is in our nature; to be curious is to be human. Discovery is the only means by which we progress.

Fast-forward to the modern era, and our primary modes of discovery look somewhat more sophisticated. Our unquenchable thirst for an ever-deeper understanding of our Universe has culminated in the scientific method, irrefutably recognised as the central pilar of human ingenuity and the inextinguishable engine of human advancement. Moreover, science is the embodiment of curiosity for the sake of understanding and the satisfaction that comes therewith — technological advancements which arise from our inquisition are but a fortuitous byproduct.

It is not an aggrandisement to assert that that the Large Hadron Collider is the pinnacle of our modern-day quest for yet further knowledge. A cathedral to curiosity, the Large Hadron Collider proudly constitutes the largest machine built by mankind. From medical imaging to the world-wide-web, myriad unforeseen technological offspring owe their existence to the project. This, however, is not its true purpose. The Large Hadron Collider is a tool with which to extend the boundaries of our understanding; it exists to answer our questions regarding the Universe for no reason beyond scratching the scientific and human itch of curiosity. It represents the best hope in our age of deciphering and disentangling the most fundamental nature of reality.

In 2012, ATLAS and CMS — two experimental collaborations tasked with studying collisions at the Large Hadron Collider — announced the discovery of the Higgs boson. This achievement of historical proportions confirmed the final prediction of the theoretical framework through which we view the world, crowning its supremacy. From virtual particles to the origin of mass, the predictions admitted by our theoretical description of the world withstand experimental assault to a mind-boggling degree of numerical precision. Humanity would be forgiven for feeling a little smug.

Not-withstanding the remarkable success of our existing theoretical machinery, many questions remain unanswered. Dark matter, dark energy, matter/anti-matter asymmetry — numerous observations stand at variance with our theoretical framework, and cannot be reconciled with our current view of the Universe using the conceptual toolkit it provides.

The Large Hadron Collider stands ready to assist us in our efforts to resolve this scientific impasse.

If we are to be thoroughly conscious of our ignorance, as Maxwell implores us to be, we must first recognise the full extent to which our understanding of the Universe may be incomplete. "There are known knowns, ...and there are known unknowns", as exclaimed by Rumsfeld. But there are also "unknown unknowns". It is in this spirit of unassuming exploration that the research which follows seeks to contribute to the scientific endeavour.

The approach adopted by this thesis is one of *model-independence* — that is to say, minimal presumptions have been made concerning the form which new physics may assume. The analysis undertaken in this thesis does not seek to test the predictions of a particular theoretical model of new physics. A plethora of analyses have hitherto performed extensive and sweeping searches for the experimental signatures predicted by our most promising theories, and have failed to yield the fruit of discovery. Instead, this research seeks to expand the space of possibilities to which we are experimentally sensitive — reducing the number of places where nature may choose to hide.

The frontier which this thesis serves to extend is that of particle *lifetime*. Fundamental particles are characterised by a number of key attributes, one of which is their lifetime: the average duration of time which will elapse before they decay into other species. In a practical sense, this corresponds to the distance a given particle will traverse within the large experimental structures tasked with

detecting their presence at the Large Hadron Collider. Current experimental procedures are largely optimised for the detection of promptly-decaying particles, such as the Higgs boson, and may not be sensitive to species of particle with extended lifetimes — particles which may act as harbingers of the new physics we seek.

There exist many reasons to believe that such long-lived particles may exist, which will be examined in the pages overleaf. While the search conducted in this thesis is independent of any given theoretical prediction and aims to generalise to a wide scope of possibilities, some assumptions must, of course, be implemented. It is assumed in this analysis that long-lived particles, should they exist, would decay preferentially to a particle known as the τ (read 'tau') lepton. The motivation for imposing this assumption is discussed in detail hereinafter. The statistically-significant observation of τ-leptons at some appreciable displacement beyond their typical length of decay would therefore represent the equivalent of a smoking gun in our detective-style search for evidence.

Before an astronomer can peer farther into the cosmos, they must first increase the magnification of their telescope. Analogously, the tools with which a particle physicist studies the quantum realm must be adequately optimised for the analysis of interest. The pages which follow describe the work undertaken to develop the tools required in order to peer deeper into the fundamental constituents of our Universe, and the analytical techniques employed to search for displaced τ-leptons — messengers of the new physics required to answer our outstanding questions about the nature of our reality.

Abstract

The standard model of particle physics (SM) represents the most sophisticated theoretical account of fundamental physics to-date, admitting predictions which stand in remarkable agreement with experimental measurement. Despite this, a number of questions remain unanswered by the SM. Long-lived particles (LLPs) beyond the SM offer a possible remedy to the theoretical and predictive shortcomings of the SM.

The conjectured existence of LLPs is herein contextualised against the landscape of contemporary theoretical thinking, before the experimental and statistical methods with which constraints are placed upon their production are introduced. A model-independent search performed with 'run-2' LHC data for the decay signatures expected of LLPs is subsequently presented, recognising the privileged role which may be assumed by third-generation leptons in the decay kinematics of such BSM states. While model-independent in nature, simulated Monte Carlo samples serve as a benchmark to the sensitivity of the analysis to well-motivated models of LLP states.

A high-level trigger algorithm was developed in the course of this thesis, which is shown to improve the efficiency with which events featuring displaced τ-lepton content may be captured with the ATLAS detector by $\sim$ 30% with respect to the acceptance of simulated data. The background event population owing to quantum-chromodynamic (QCD) dijet production is estimated through appeal to the data-driven 'hadronic fake factor' method. This thesis reports an estimated background population in the signal region which corresponds to the fully-hadronic decay of displaced τ-leptons, where systematic uncertainties are determined as a function of the non-closure error associated with the validation of the estimate as performed on an orthogonal data-sample. Further, expected sensitivity limits are presented at the 95% confidence-level, representing the exclusion of states in hitherto-uncharted phase-space.

Declaration

I hereby declare that this thesis was composed by myself, that the work contained herein is my own except where explicitly stated otherwise, and that this work has not been submitted for any other degree or professional qualification except as specified.

Parts of this work have been published in [1].

— Jack Gargan, October 2024

Acknowledgements

There are a great many people to whom I owe my thanks as I reflect on the completion of my doctoral studies. To endeavour to thank them all would surely exhaust my page limit — to those who I have not mentioned: I ask for your forbearance. You nonetheless have my sincere gratitude for your help and encouragement along the way.

Undoubtably, my greatest thanks go to my parents. As is true for so many things in my life, the love and support you have given me, and the sacrifices you have made, are what have brought me here today. For this I will be forever indebted. The support you have offered me over the course of my PhD has enabled me to divert my focus and energy to my studies without cause for concern in other domains of life. I don't express my gratitude nearly as often as I should for all you have done for me — I hope that, by some unspoken means of understanding, it has always been known. The greatest support you have provided, however, is through being my best friends. Mum and Dad: thank you.

I wish to express my heartfelt and unreserved thanks to my supervisor, Professor Sinead Farrington FRSE. Sinead, I could not have asked for a better induction into the world of High Energy Physics. Throughout my doctoral studies, you have without exception placed my academic and scientific development ahead of other considerations (including the budget!). Between summer school engagements and pursuing auxiliary research of interest, you have always afforded me the opportunity to follow my interests while simultaneously helping me to navigate the successful completion of my PhD. Moreover, thank you for having faith in my ability to undertake a PhD in the first instance — the last four years of horizon-expanding travel and a potential career in high-energy physics would not otherwise have been possible. Thank you for your guidance, and thank you for trusting me.

A compilation of acknowledgements would not be complete without issuing an immense thank you to Dr. Sara Alderweireldt. Sara, despite seemingly having one of the largest work-loads in ATLAS, you have unfailingly allocated generous proportions of your time to provide me with advice and tutorship. You have equipped me with the practical toolset with which to undertake data-intensive research with ATLAS, and given me valuable insight into the many factors pertinent to undertaking physics research at the coal-face of the LHC. The

example you set as a role model for young physicists is unparalleled. Thank you.

On the subject of post-doctoral guidance, I would like to add my profuse thanks to Dr. Santiago Saenz and Dr. Julia Silva. Since joining our research team, the ability to seek your assistance within the office has allowed my studies to change gear. Discussing the practical details of our analysis together has been among the most helpful and elucidating aspects of my studies. I am fortunate to have benefited from the excellent example you set for early career researchers. Thank you for your approachable and friendly support, and for always leaving the door to your office open.

I would like to add further thanks to Professor Frank Krauss. Frank, thank you for the opportunity to learn about MC theory, and for offering your valuable time to tutor me despite being under no obligation to do so. I am supremely grateful for the chance to learn from one of the intellectual giants of the field. My ignorance of this arena must have been a source of frustration at times. The thought of undertaking a project in theoretical physics was initially a source of some intimidation, yet your disarming sense of humour and relaxed manner put me at ease. Your passion for teaching is clear. Thank you for your enjoyable induction into the world of particle physics phenomenology.

Additional thanks are due to Dr. John Baines for his stewardship through the ATLAS authorship qualification process, and for offering his knowledge and expertise of the ATLAS trigger system. John, your advice is very much appreciated.

Of the numerous life-changing elements to come from my time as a PhD student, the most important and lasting is unquestionably the friendships I have acquired along the way. In particular, the experience of spending a year at CERN on long-term attachment has provided me with friendships and memories which I look forward to taking with me through my life. Thank you all for the smiles and the laughter.

Last but certainly not least, I wish to thank CERN for the successful operation of the LHC, without which my research would not have been possible, and for providing a friendly and accommodating environment to call home for the duration of 2022 despite the challenges still posed by the global pandemic. My year at CERN will forever remain a life-affirming memory.

Contents

List of Figures

List of Tables

Part I

Theoretical Framework

“Those who are not shocked when they first come across quantum theory cannot possibly have understood it.”

— Niels Bohr

1

Quantum Field Theory

Quantum field theory may reasonably be regarded as a momentous milestone in the intellectual efforts of humanity to comprehend the natural world. A comprehensive introduction to quantum field theory is beyond the scope of this thesis. Instead, this chapter serves to introduce the key conceptual innovations which underpin the success of quantum field theory, such that the construction of the standard model of particle physics may be presented on a solid foundation. Moreover, the practical calculations of interest to modern particle physics experiments are introduced in advance of their invocation in subsequent chapters.

1.1 Overview

Quantum field theory (QFT) is widely regarded as the greatest intellectual advance of the twentieth century in the understanding of physics. Collating the most successful theories to precede it in a rigorous and self-consistent fashion, QFT is a testament to the success of humanity's scientific endeavour. This chapter seeks to introduce the central ideas which underpin QFT such that a modern view of fundamental particle physics may be presented in subsequent chapters.

QFT can be understood as the amalgamation of three major advances in theoretical physics: classical field theory, special relativity, and quantum mechanics. The result is a self-consistent framework with which to understand the Universe which is more intellectually complete than that provided by any of the constituent components. As an example of its predictive power, the quantum field theoretic description of light and matter interactions predicts that the magnetic moment of the electron, g_e, should take the value [2] of:

$$\left(\frac{g_e}{2}\right)_{\text{theory}} = 1.001\,159\,652\,180\,73(28), \tag{1.1}$$

which stands in remarkable agreement with the experimentally measured value [2] of:

$$\left(\frac{g_e}{2}\right)_{\text{exp.}} = 1.001\,159\,652\,181\,643(764). \tag{1.2}$$

This is the most accurately-verified prediction in the history of physics.

The standard model of particle physics is a quantum field theory which leverages the conceptual and mathematical advances charted by QFT to provide a more complete description of the natural world. Before the standard model is discussed, a number of core concepts from QFT must first be introduced.

The introduction which follows is intentionally brief in nature and limited in scope — the focus of this chapter is directed towards the concepts upon which the standard model of particle physics is constructed, and the formulae which prove to be of practical utility in the study of high energy physics at hadron colliders.

1.2 Relativistic Quantum Mechanics

The Dirac equation [3] represents the first fully successful attempt to wed special relativity to quantum mechanics, having addressed some of the conceptual hurdles presented by the Klein-Gordon equation [3].

While the Klein-Gordon equation secured Lorentz invariance as a function of second-order derivatives in both space and time, it suffered from the prediction of unphysical negative probability densities. This issue is resolved by Dirac's treatment [3]. Written in covariant form, the Dirac equation may be expressed succinctly as:

$$(i\gamma^{\mu}\partial_{\mu} - m)\,\psi = 0, \tag{1.3}$$

where $\partial_{\mu} \equiv \partial/\partial x^{\mu}$ is the covariant derivative, m is the particle mass, ψ is the quantum mechanical wavefunction, and the Dirac gamma matrices γ^{μ} take the form of:

$$\begin{aligned} \gamma^0 &= \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & -1 & 0 \\ 0 & 0 & 0 & -1 \end{pmatrix}, \quad \gamma^1 = \begin{pmatrix} 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \\ 0 & -1 & 0 & 0 \\ -1 & 0 & 0 & 0 \end{pmatrix}, \\ \gamma^2 &= \begin{pmatrix} 0 & 0 & 0 & -i \\ 0 & 0 & i & 0 \\ 0 & i & 0 & 0 \\ -i & 0 & 0 & 0 \end{pmatrix}, \quad \gamma^3 = \begin{pmatrix} 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & -1 \\ -1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \end{pmatrix}. \end{aligned} \tag{1.4}$$

The Dirac gamma matrices are known to adhere to the anti-commutation relation:

$$\{\gamma^{\mu},\, \gamma^{\nu}\} \equiv \gamma^{\mu}\gamma^{\nu} + \gamma^{\nu}\gamma^{\mu} = 2g^{\mu\nu} \tag{1.5}$$

for the metric tensor $g^{\mu\nu}$ and satisfy the additional relations $\left(\gamma^0\right)^{\dagger} = \gamma^0$ and $\left(\gamma^k\right)^{\dagger} = -\gamma^k$, where the Hermitian conjugate $(\gamma^{\mu})^{\dagger}$ is defined as the complex conjugate of the transpose of γ^{μ}.

The solutions to the Dirac equation may be packaged in the form of *Dirac spinors* [3]. A spinor, $\psi = (\psi^1, \psi^2, \psi^3, \psi^4)^\mathsf{T}$, harbours four components given by:

$$\begin{aligned} \psi^1 &= u^\uparrow \, e^{-ip^\mu x_\mu}, \quad \psi^2 = u^\downarrow \, e^{-ip^\mu x_\mu}, \\ \psi^3 &= \nu^\uparrow \, e^{ip^\mu x_\mu}, \quad \psi^4 = \nu^\downarrow \, e^{ip^\mu x_\mu}, \end{aligned} \tag{1.6}$$

where $p^\mu = (E/c, \vec{p}) = (E/c, p_x, p_y, p_z)$ and $x_\mu = (ct, -\vec{x}) = (ct, -x, -y, -z)$ are respectively the contra-variant momentum and covariant positional four-vectors of the particle solution, and the terms u and ν are defined as:

$$\begin{aligned} u^\uparrow(E, \vec{p}) &= \sqrt{E+m} \begin{pmatrix} 1 \\ 0 \\ \frac{p_z}{E+m} \\ \frac{p_x+p_y}{E+m} \end{pmatrix}, \quad u^\downarrow(E, \vec{p}) = \sqrt{E+m} \begin{pmatrix} 0 \\ 1 \\ \frac{p_x-ip_y}{E+m} \\ \frac{-p_z}{E+m} \end{pmatrix}, \\ \nu^\uparrow(E, \vec{p}) &= \sqrt{E+m} \begin{pmatrix} \frac{p_x-ip_y}{E+m} \\ \frac{-p_z}{E+m} \\ 0 \\ 1 \end{pmatrix}, \quad \nu^\downarrow(E, \vec{p}) = \sqrt{E+m} \begin{pmatrix} \frac{p_z}{E+m} \\ \frac{p_x+p_y}{E+m} \\ 1 \\ 0 \end{pmatrix}. \end{aligned} \tag{1.7}$$

In the above notation, E represents the energy of the particle solution provided by the respective spinor components, while p_x, p_y, and p_z represent the three spatial components of the particle's momentum vector, $\vec{p}$.

In addition to the regular Dirac spinor, the adjoint spinor may be defined as $\overline{\psi} = \psi^\dagger \gamma^0 = (\psi_1^*, \psi_2^*, -\psi_3^*, -\psi_4^*)^\mathsf{T}$, where the Hermitian conjugate $\psi^\dagger$ is defined as the complex conjugate of the transpose of ψ.

Beyond providing an excellent relativistic description for the behaviour of fermions, the Dirac equation foreshadows the existence of anti-matter. The negative energy solutions of the equation must have a physical interpretation if the quantum-mechanical condition of providing a complete set of basis states is to be satisfied [4].

The negative energy solutions to the Dirac equation are interpreted as *anti-particles* [3]. In practise, the anti-particles of charged states typically correspond to particles with the opposite sign of charge to that of their regular matter

counterparts, whereas neutral elementary particles typically serve as their own anti-particle. Such anti-particle states would otherwise manifest themselves without distinction if the picture were not complicated by the exception of their CP-symmetrical properties [5]. The components $u^{\uparrow}$ and $u^{\downarrow}$ of the spinor solution respectively represent the spin-up and spin-down particle states described by the solution, whereas $v^{\uparrow}$ and $v^{\downarrow}$ respectively denote the spin-up and spin-down states which correspond to anti-particles. The prediction of anti-matter has been extensively verified by experiment, whose existence presents a number of challenging concepts. For instance, an anti-matter particle moving forward in time is conceptually equivalent to a regular particle travelling backwards in time in the relativistic Minkowski view of space-time [4, 6, 7].

1.3 Towards a Quantum Theory of Fields

Classical field theory is observed to have enjoyed tremendous scientific success long before relativistic quantum mechanics was successfully devised. The most prominent testament to the power of field theory is found in Maxwell's theory of electromagnetism, published in 1861. With the newfound success of relativistic quantum mechanics, the efforts of physicists naturally converged upon the desire to devise a quantum mechanical field theory.

The theoretical framework which resulted from the combination of special relativity, quantum mechanics, and field theory enabled the development of the most successful theory in scientific history, providing the most fundamental account of reality to-date and withstanding experimental assault to a hitherto unthinkable degree of accuracy. The manner in which fields may be quantised is now considered.

1.3.1 Second Quantisation

In this field-theoretic framework, particles are promoted to fields. Just as the photon is classically considered to be a manifestation of the electromagnetic field, the electron is equivalently viewed as a manifestation of the electron field. In this view, particles are considered to be excitations of their underlying field. As Frank Wilczek remarks in his essay on quantum field theory [8], this readily explains the indistinguishability of fundamental particles — or, equivalently, why an electron

from the Andromeda galaxy cannot be distinguished from an electron within ourselves. This is easily understood if all electrons are viewed as excitations of a single underlying field.

The field associated with a given particle is given by the Fourier sum of each individual quantum wave-function associated with the particle. Coefficients are ascribed to each wave-function in the series which denote the creation and annihilation probability associated with each wavelength (or, via de Broglie, momentum). This procedure is known as *second quantisation* [4]. (The recognition of particle-wave duality and its applicability to matter is regarded as the first iteration of quantisation.)

Second quantisation may be implemented with a number of formalisms. The most primitive formalism is that of *canonical quantisation* [4], in which fields are promoted to quantum mechanical operators which obey a set of canonical commutation relations. A more elegant and convenient approach is offered by Feynman path integrals [4, 6], whose operation shall be elucidated in the following passages.

Mathematically, the picture which results from quantising fields may be viewed as equivalent to an infinite series of harmonic oscillators [4]. Each oscillator will vibrate at a particular frequency when excited — where each such vibration corresponds to a particular excitation of the quantum field in question, producing quanta (*i.e.* particles) at that given frequency, ν (or, *via* $E = h\nu$, energy).

1.3.2 Time-Dependent Perturbation Theory

A theory of *interacting* fields is required if the natural world as observed is to be described. The probability of a given interaction occurring between two quanta is related to the dynamics of the parent theory. In a manner analogous to classical field theories, the dynamics of a quantum field theory may be obtained through the calculus of variations [4]. The *action*, S, for a given theory of quantum field, ϕ, may be defined as:

$$S = \int \mathrm{d}t\, L = \int \mathrm{d}^4 x\, \mathcal{L}\left(\phi,\, \partial_\mu \phi\right), \tag{1.8}$$

where $\partial_\mu \equiv \partial / \partial x^\mu$ for the position four-vector x^μ. Here, the *Lagrangian density*,

$\mathcal{L}$, has been introduced as a more convenient and natural description of quantum fields than the Lagrangian, L. Quantities denoted by $\mathcal{L}$ which formally meet the definition of a Lagrangian density will simply be referred to as 'Lagrangians' hereinafter.

Invoking the *principle of least action* asserts that $\delta S = 0$, whereafter the *Euler-Lagrange* equation enables the equation of motion for the field, ϕ, to be determined:

$$\frac{\partial \mathcal{L}}{\partial \phi} - \partial_\mu \left(\frac{\partial \mathcal{L}}{\partial \left(\partial_\mu \phi \right)} \right) = 0. \tag{1.9}$$

In this manner, the dynamics of a given quantum theory of fields may be ascertained.

It is further advantageous to define the *Hamiltonian (density)*, $\mathcal{H} = \pi \dot{\phi} - \mathcal{L}$, where $\pi = \partial \mathcal{L} \, / \, \partial \dot{\phi}$ is the conjugate momentum density.

The Hamiltonian which governs the dynamics of a given quantum field theory will contain for each particle a *free term*, in addition to an *interaction term* which describes the permissible interactions between them. Decomposing a given Hamiltonian into its free terms, $\mathcal{H}_0$, and interaction terms, $\mathcal{H}_1$, it is possible to utilise time-dependent *perturbation theory* [4, 3] such that interactions may be approximated as perturbations between the free initial and final quantum states. The S-matrix may then be defined relating initial and final quantum states by such perturbative interactions, whose elements are instantiated to:

$$S_{fi} = \langle f; t = +\infty | \, \widehat{U}(t, t_0) \, | i; t_0 = -\infty \rangle \, , \tag{1.10}$$

for initial quantum state $|i; t_0 = -\infty\rangle$ and final quantum state $|f; t = +\infty\rangle$. In the Dirac picture [4], the unitary *time evolution operator*, $\widehat{U}(t, t_0)$, is defined as:

$$|\Psi(t)\rangle = \widehat{U}(t, t_0) \, |\Psi(t_0)\rangle \tag{1.11}$$

for a generic time-dependent quantum state, $|\Psi(t)\rangle$. The effect of the time evolution operator on an initial state $|i\rangle$ may be represented as an infinite series quantifying each possible 'path' between states $|i\rangle$ and $|f\rangle$. Each successive term in the expansion perturbs the solution to a lesser degree, rendering possible the

approximation of a solution with a finite number of terms.

The probability that a quantum system will transition from initial state $|i\rangle$ to final state $|f\rangle$ is then given by $|S_{fi}|^2$. This is known as the *amplitude* [3]. The element of the S-matrix which corresponds to a given interaction process of interest is conventionally referred to as the 'matrix element' for brevity.

Perturbation theory was originally leveraged as a means by which the amplitude for a process whose Hamiltonian has no exact solution may be approximated. This is achieved by expressing the Hamiltonian in question as a sum of two alternative Hamiltonians: one for which the exact solution is known, together with a small perturbation. The contribution of each component to the overall amplitude may then be determined individually.

Similar logic was applied to quantum field theory by Richard Feynman, who devised a method whereby the propagation of interacting field quanta is expressed as a sum of increasingly complex (and diminishingly weighted) sub-processes [4, 9]. Each sub-process describes the free field terms and the successively more complicated interactions between them. The initial term in the expansion is known as the 'leading order' (LO) contribution, with successive contributions termed next-to-leading order (NLO), NNLO, $\mathrm{N^3LO}$, and so forth. The notion of summing every possible interaction, or 'path', by which the amplitude may be arrived at is the logic which underpins the afore-mentioned Feynman path integral.

The rules associated with this procedure are encapsulated by the fully rigorous graphical calculus provided by *Feynman diagrams*. In addition to providing a convenient pictorial representation for each sub-process in the sum, Feynman diagrams enable the automatic determination of the LO matrix element by following the rules of the calculus, known as 'Feynman rules'. An introduction to Feynman diagrams and their associated rules is provided in Appendix 1.

This procedure is found in some circumstances to produce solutions which feature unphysical infinities. The manner in which this is resolved is known as *renormalisation* [4, 10], and will be discussed further in the subsequent chapter.

1.3.3 Virtual Processes and the Quantum Vacuum

The sub-processes which feature in a perturbative series of Feynman diagrams are, of course, constrained by energy and momentum conservation. Remarkably, however, the individual field quanta within a given sub-process — represented by internal propagation lines within the corresponding Feynman diagram — are *not* in isolation required to satisfy this conservation condition [10]. This freedom is afforded by the Heisenberg Uncertainty Principle [4], which states that energy is uncertain within ΔE for time Δt such that $\Delta E \Delta t \geqslant \hbar$. Additional particle states may therefore be excited from their underlying field provided that their existence is contained within the time window of Δt, thereby ensuring that energy and momentum conservation is preserved overall.

Particles of this nature are known as *virtual particles* and constitute the intermediate states of fundamental particle interaction processes. While virtual particles are not directly observable, their effects are tangible and must be accounted for. Virtual particles are said to be "off mass-shell" by extension of the fact that they do not satisfy $E^2 = p^2c^2 + m^2c^4$.

In contrast to the classical concept of empty space-time, quantum field theory predicts the existence of a much more complex and active vacuum. The 'quantum vacuum' is the quantum-mechanical state in which no field quanta are excited. Despite being in their ground state, quantum fields are subject to random quantum 'fluctuations' which arise due to the creation and annihilation of virtual states within the time window Δt, as permitted by the Heisenberg Uncertainty Principle. The quantum vacuum therefore contains fluctuating quantum fields which give rise to a *zero-point energy* [4, 10].

The physical consequences of the quantum vacuum may be deduced indirectly through a number of empirical observations. One such instance is the observation of the Casimir effect [11], predicted by Hendrik Casimir in 1948 before its experimental confirmation in 1958. Casimir postulated that two flat metallic plates ought to exert an attractive force on one-another with a magnitude inversely proportional to their separation to the fourth power. This arises due to the suppression of electromagnetic zero-point fluctuations whose wavelength exceeds the separation of the plates. The zero-point energy density between the plates is therefore less than that of the external vacuum, exerting a pressure on the plates which is perceived as an attractive force. A further example may be

found in the Hawking radiation [12] emitted from black holes.

1.4 Symmetries of Quantum Field Theory

With the central concepts interwoven throughout QFT now introduced, it is timely to present the mathematical machinery which facilitates the gauge-theoretic construction of the standard model of particle physics — itself a QFT. The mathematical concept of greatest import to the standard model is that of *symmetry*.

1.4.1 The Role of Symmetry

Throughout the history of its development, fundamental physics has long exhibited symmetrical properties. The laws of Newtonian mechanics satisfy Galilean invariance, while Maxwell's equations of electrodynamics were later found to exhibit both Lorentz and gauge invariance. The fundamental significance of symmetry to physics, however, was largely overlooked until the twentieth century.

The turning-point in the rise of symmetry is marked by the discovery of *Noether's theorem* [4]. Every continuous symmetry present in a physical theory was found by Noether in 1918 to correspond to a conservation law. This revelation played a critical role in forming the modern view of Einstein's General Relativity, and is central to the prediction of conserved charges in quantum field theories.

While important to almost all domains of physics, symmetries of various forms are critical to the ability of QFT to provide such a successful account of the natural world at the most fundamental scales. Accordingly, the standard model of particle physics — as introduced in the following chapter — makes routine appeal to the symmetrical concepts introduced herein.

The domain of mathematics known as *group theory* [13] lends itself naturally to capturing the symmetrical properties of transformations. Groups may be used to describe transformations of either a discrete or continuous nature — groups which exclusively describe continuous transformations are termed *Lie groups* [13].

1.4.2 Lie Groups

Lie groups prove to be a particularly powerful language with which to describe the symmetries of fundamental physics, forming a key pilar of the gauge-theoretic concepts which underpin the standard model of particle physics.

In general, a group is a set of mathematical objects, equipped with a well-defined multiplication rule, which satisfies the conditions of closure and associativity. Further, a group must contain an identity element, with each element in the group being in possession of an inverse element.

For the purposes of fundamental particle physics, groups of relevance are those whose elements are the matrices describing the rotations or transformations pertinent to the symmetrical arguments of the theory at hand. In the case where such matrices describe a continuous transformation, the group in question is a Lie group. The *generators* of the group are the elements within the group from which every other element of the group may be produced.

A number of specific Lie groups are of particular utility in the standard model which deserve attention.

The U(1) Group

The U(1) Lie group is the group of matrices which induce rotations about a single axis, by a single rotation angle. Here, 'U' refers to the fact that each matrix in this group is unitary.

The dimension of a group is the number of parameters required to describe all elements of the group. In the case of the U(1) Lie group,

$$\dim\left[\mathrm{U}(1)\right] = 1, \tag{1.12}$$

which conveys the fact that each matrix in the group is of dimension 1×1 — that is to say, a single number. The group U(1) will be shown in the subsequent chapter to be of use in describing the gauge symmetry of quantum electrodynamics. The elements of the group may take the form of $\exp\left(i\alpha\right)$ for constant rotation angle α. The U(1) Lie group is an example of an *Abelian* group [13], where any given pair of group elements will always mutually commute.

The SU(N) Groups

Groups of the form $\mathsf{SU}(N)$ are those whose elements are complex-valued matrices of dimension $N \times N$. Here, 'S' denotes that the matrices encapsulated by the group are 'special', meaning that they are of unit determinant: $\det(\mathsf{U}) = \mathbb{1}$.

Due to their possession of both a real and imaginary component for each matrix element, complex matrices of $N \times N$ dimensionality experience $2N^2$ degrees of freedom. For special unitary matrices, there exists a set of N^2 equations:

$$\left(U^{\dagger}\right)_{ij} U_{jk} = \delta_{ik} \qquad \text{for } \delta_{ij} = \begin{cases} 1 & : i = j \\ 0 & : \text{otherwise} \end{cases}, \tag{1.13}$$

where δ_{ij} is the Kronecker-delta symbol. Addressing the further restrictions imposed by the condition of $\det(\mathsf{U}) = \mathbb{1}$, the degrees of freedom available for $\mathsf{SU}(N)$ groups is found to be:

$$2N^2 - N^2 - 1 = N^2 - 1. \tag{1.14}$$

For instance, the dimensionality of the Lie group $\mathsf{SU}(3)$ is therefore $\dim[\mathsf{SU}(3)] = 8$. The introduction of the gauge-theoretic nature of quantum chromodynamics in Chapter 2, Section 2.3, will rely on this observation to explain the existence of 8, not 9, gluons with which nature mediates the strong force. The $\mathsf{SU}(3)$ group is an example of a *non-Abelian* group [13], where at least one pair of elements are non-commutative.

Lie Algebra

The number of generators with which a Lie group is equipped is equal to the dimensionality of the group. The set of generators of a group, G, is conventionally denoted by $\{T^a\}$ for $a = 1, ..., \dim(G)$. *Lie's theorem* [13, 9] dictates that the elements, U, of a group are related to the generators of the group by:

$$U = \exp\left[i\,\theta^a\,T^a\right], \tag{1.15}$$

where $\{\theta^a\}$ are the continuous parameters of the group. From this relation, it may

be shown by arguments of consistency and closure that:

$$\left[T^a,\ T^b\right] = i\, f^{abc}\ T^c. \tag{1.16}$$

This equation is known as the *Lie algebra* of the group, and applies to the generators of any Lie group. The terms $\left\{f^{abc}\right\}$ are known as the *structure constants* of the group, which, unlike generators, are uniquely fixed for a given group. A group may therefore be defined by its associated structure constants.

Conversely, as the generators of a Lie group are not uniquely-valued, any set of matrices which satisfies the Lie algebra may be chosen. The set of chosen matrices is known as the *representation* of the group generators. Where the chosen set contains $N \times N$ matrices of unit determinant, the representation is known as the *fundamental representation* — a concept which shall be invoked in subsequent discussions.

The mathematical concepts introduced in this section, while abstract in nature, provide the foundation upon which to examine the gauge-theoretic construction of the standard model. Beyond the rotational symmetries thus far discussed, consideration is now payed to a number of additional symmetries exhibited by QFT.

1.4.3 *CP* Symmetries

A number of other symmetries are manifest in quantum field theories. Should a quantum state be left invariant under inversion of the sign of its charge, it is said to exhibit charge symmetry. The charge conjugation operator, $\widehat{C}$, may be defined such that:

$$\psi' \equiv \widehat{C}\psi = i\gamma^2\psi^*, \tag{1.17}$$

for a quantum state ψ, where ψ^* is the complex conjugate of the state.

Similarly, a quantum state is said to exhibit parity symmetry if the reflection of its spacial components through the origin leaves the system unchanged. The

parity operator, $\widehat{P}$, is defined as:

$$\psi(x, y, z, t) \rightarrow \psi'(-x, -y, -z, t) = \widehat{P}\psi. \tag{1.18}$$

Accordingly, the double-application of the parity operator will return a quantum state ψ to its original condition if parity symmetry is observed. Physically-significant consequences emerge from the violation of parity symmetry, as will be shown to be true for electroweak unification in Section 2.4 of the following chapter. Further, the violation of the combined charge-parity symmetry — CP-violation — is an important topic in high energy physics with cosmological implications, as discussed in [5].

1.5 Quantum Field Theory for Collider Physics

A command of the conceptual foundations of QFT as afore-introduced enables the introduction of the key calculational tools of utility to the experimentalist. Such equations allow the expected properties of various observables to be determined. It is the dual-pursuit of the experimentalist to open the widest possible aperture through which to discover the unexpected, and to interrogate the predictions of the theorist as advanced by a given theory.

A key component of such experimental interrogations is the measurement (and simulation) of the rate at which one quantum state transitions to another. The probability per unit time that an initial quantum state, $|i\rangle$, will transition to a final quantum state, $|f\rangle$, is given by *Fermi's Golden Rule* [3, 9]:

$$\Gamma_{fi} = 2\pi |\mathcal{M}_{fi}|^2 \rho(E_i), \tag{1.19}$$

where $|\mathcal{M}_{fi}|^2$ is the amplitude (*i.e.* square of the matrix element) for the process, and $\rho(E_i)$ is the *density of states*. The density of states may be understood as the density of accessible final momentum states available — determined by the kinematics of the process and the Pauli Exclusion Principle [3]. The density of

states may therefore be expressed as:

$$\rho\left(E_i\right) = \left|\frac{\mathrm{d}n}{\mathrm{d}E}\right|_{E_i} = \int \frac{\mathrm{d}n}{\mathrm{d}E}\,\delta\left(E_i - E\right)\,\mathrm{d}E, \tag{1.20}$$

where $\mathrm{d}n$ is the number of accessible states in the energy range given by $E \to E + \mathrm{d}E$. The density of states is typically referred to as 'phase-space' in high energy physics, a nomenclature which is employed throughout this thesis. Appeal has been made to the Dirac delta function [14] such that energy conservation is enforced. Fermi's Golden Rule then becomes:

$$\Gamma_{fi} = 2\pi \int |\mathcal{M}_{ij}|^2\,\delta\left(E_i - E\right)\,\mathrm{d}n. \tag{1.21}$$

The number of accessible quantum states within the momentum-volume $\mathrm{d}^3\vec{p}$ is given by $\mathrm{d}n = \mathrm{d}^3\vec{p}/(2\pi)^3$. For a particle decay process $A \to 1 + 2 + ... + N$, this becomes:

$$\mathrm{d}n = (2\pi)^3 \prod_{i=1}^{N} \frac{\mathrm{d}^3\vec{p}_i}{(2\pi)^3}\,\delta^3\left(\vec{p}_A - \sum_{i=1}^{N}\vec{p}_i\right), \tag{1.22}$$

where momentum conservation is once again assured with the invocation of the Dirac delta term. This expression may then be formulated in a Lorentz invariant manner:

$$\begin{aligned}\Gamma_{fi} &= \frac{(2\pi)^4}{2E_A}\int |\mathcal{M}_{fi}|^2\,\delta\left(E_A - E_1 - E_2\right)\delta^3\left(\vec{p}_A - \vec{p}_1 - \vec{p}_2\right)\,\frac{\mathrm{d}^3\vec{p}_1}{(2\pi)^3 2E_2}\cdot\frac{\mathrm{d}^3\vec{p}_2}{(2\pi)^3 2E_2} \\ &= \frac{|\vec{p}^*|}{32\pi^2 m_A^2}\int |\mathcal{M}_{fi}|^2\,\mathrm{d}\Omega,\end{aligned} \tag{1.23}$$

such that it may generalise to relativistic two-body phase-space, where consideration has been restricted to processes of the form $A \to 1 + 2$ for brevity. Here, $\mathrm{d}\Omega$ denotes the solid angle, and $|\vec{p}^*|$ represents the momentum as measured in the rest frame of reference of particle A. Equation 1.23 is valid for all two-body decays.

The lifetime of a particle may be defined as the mean duration of time elapsed

before it decays into a lighter species. The total decay rate of a given particle is reciprocally related to its proper lifetime:

$$\tau = \frac{1}{\Gamma_{\text{total}}}, \tag{1.24}$$

where Γ_{total} encompasses all possible decay modes of the particle in question. From a relativistic perspective, the proper lifetime is that whose measurement is agreed upon in the inertial rest frame of the particle itself. The quantity Γ_{total} is often termed the 'width' of the particle. Particle lifetime is a central concept to which this thesis will make routine reference.

The fraction of time by which a particle state, A, decays via a particular decay mode to the final state ij is given by the *branching ratio* for that process, defined as:

$$\text{Br}(A \to ij) = \frac{\Gamma_i}{\Gamma_{\text{total}}}. \tag{1.25}$$

In addition to the decay rate of a given particle, it is of further utility to the experimentalist to measure the probability that a given interaction will occur between two particle states, known as the *cross-section* [3, 9]: $\sigma = \int (\mathrm{d}\sigma/\mathrm{d}\Omega) \cdot \mathrm{d}\Omega$. Fermi's Golden Rule may be amended to provide a result in cross-sectional dimensions, yielding:

$$\sigma = \frac{1}{64\pi^2 \widehat{s}} \frac{|\vec{p}_f^{\,*}|}{|\vec{p}_i^{\,*}|} \int |\mathcal{M}_{fi}|^2 \ \mathrm{d}\Omega^*, \tag{1.26}$$

which provides a Lorentz-invariant metric for the cross-section in terms of the Mandelstam variable, $\widehat{s} = (p_1 + p_2)^2$. Asterisks denote the valuation of quantities in the centre of mass/momentum frame of reference.

This chapter has surveyed the foundational concepts which underpin the elevation of quantum mechanics to a relativistic field theory, and has further introduced the tools with which the observables of interest to an experimentalist may be calculated. Such tools shall be leveraged in the chapters which follow.

An elementary appreciation of quantum field-theoretic concepts is required if the

standard model of particle physics is to be grasped. An introduction to the standard model is now presented overleaf.

amsmath

> *Why should things be easy to understand?*
>
> — Thomas Pynchon

2

The Standard Model of Particle Physics

The Standard Model of particle physics (SM) represents one of the greatest intellectual triumphs to arise from the scientific method, whose predictions have stood up to experimental scrutiny to a remarkable degree of accuracy. The predicted value of the electromagnetic fine-structure constant admitted by quantum electrodynamics (QED) —– the SM theory of light-matter interactions —– agrees with precision measurement to within ten parts in a billion [15]. In 2012, the ATLAS and CMS collaborations announced the discovery of the Higgs boson with a mass of ~ 125 GeV [16, 17], in so doing establishing empirical evidence for the final pillar of the SM. The chapter which follows presents the basic theoretical construction of the SM, before proceeding to connect such theoretical concepts to the phenomenological considerations pertinent to performing physics measurements at hadron colliders. Experimental coverage of SM predictions is subsequently presented.

2.1 Overview

The Standard Model of Particle Physics (SM) is widely regarded as the pinnacle of humanity's endeavour to understand the natural world. With the notable exception of gravity, the SM presents an elegant encapsulation of the most fundamental constituents of the Universe and the interactions which may occur between them. From the unification of electromagnetism and the weak interaction to the complex phenomenology produced by collisions at the LHC, the SM can account for observations across an astonishing range of scales.

The predictions which the SM admit have been subjected to rigorous and continuous experimental assault since the conception of the theory, rendering it the most heavily tested theory in science. The axiomatic concepts of the SM, taken algebraically to their logical limit, admit predictions for the numerical values of physical parameters — values which may be measured by experiment. The agreement of the predicted value of such parameters with experimental measurement is a remarkable testament to the intellectual stature of the SM. For instance, the most recent measurement [2] of the magnetic moment of the electron agrees with theoretical prediction to approximately one part in a trillion, or, equivalently, 0.1 billionths of a per-cent.

Elementary particles are the main actors on the SM stage. Fundamental particles have no discernible size, revealing no structure in experimental observation down to a resolution of 10^{-18} m — elementary particles are therefore considered to be point-like [4, 7, 6]. The particles of the SM can be delineated into two main categories: *fermions* (named after Enrico Fermi), which constitute the ordinary matter familiar in every-day life, and *bosons* (named after Satyendra Nath Bose), which serve to mediate the fundamental forces of nature. Fermions can be further decomposed into *leptons* (a historical term for 'light particle'), and *quarks*.

Both leptons and quarks are composed of three generations. Electrons, e, muons, μ, and τ-leptons (together with their associated neutrinos) are respectively of the first, second, and third generation. Each generation is also home to two types, or 'flavours', of quark: respectively, the up (u) and down (d) quark, the charm (c) and strange (s) quark, and the top (t) and bottom (b) quark. Those who study the properties of the bottom quark typically refer to it as the 'beauty quark', presumably in an effort to avoid the charge of being "bottom scientists". Since the study of bottoms is better left to the medical profession, this can be

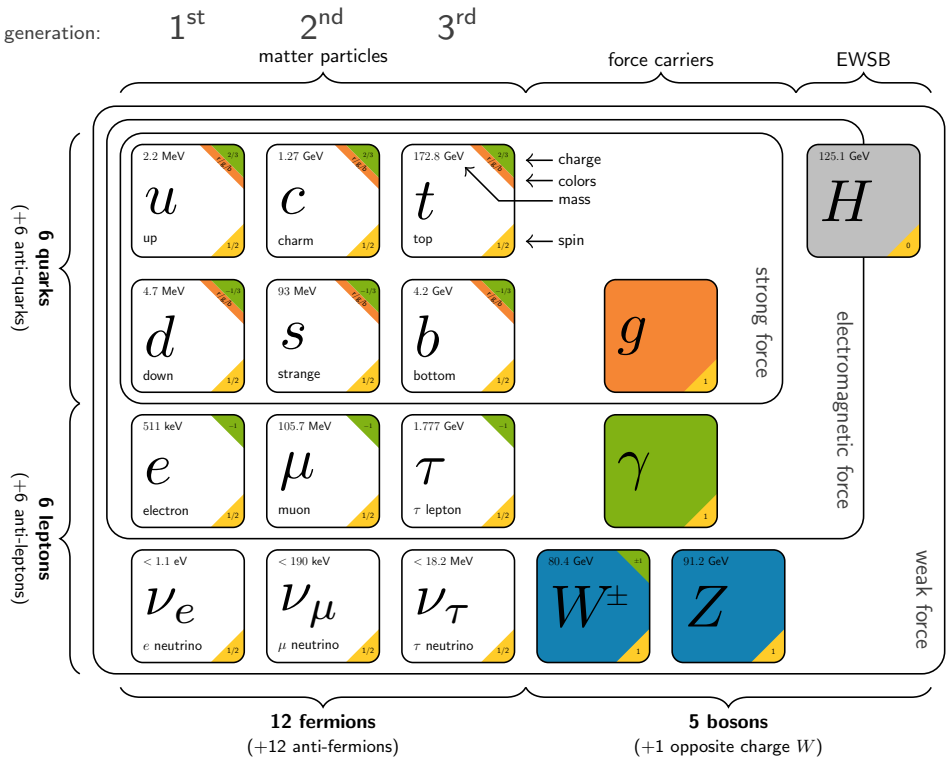

Figure 2.1 particulate content of the Standard Model of particle physics. Particles, divided into fermions and bosons, are annotated with the characteristic quantum numbers of electric charge, colour charge, mass, and spin. Image extracted from [18].

forgiven. Nonetheless, the more common-place parlance of 'bottom quark' — or simply b-quark — shall be adopted herein. The bosonic content of the SM is comprised of the neutral Z^0 and charged $W^{\pm}$ bosons, collectively responsible for the mediation of the weak force, the gluon (g), responsible for the mediation of the strong force, the photon (γ), carrier of the electromagnetic force, and the Higgs boson, a quantum mechanical excitation of the fundamental field responsible for the mechanism of spontaneous electroweak symmetry breaking (EWSB). Figure 2.1 provides an illustrative summary of the particle content of the SM. From these elementary building blocks, the visible Universe may be constructed.

The elementary particles of the SM are characterised by a set of quantum numbers which collectively determine their properties. These include the mass and electric charge of a given particle. An additional quantum number quantifies the degree of intrinsic angular momentum, known as 'spin', attributable to a given particle. This is represented by the quantum number S. All SM fermions have a fractional spin value of $S = \frac{1}{2}$, while bosons are characterised by an integer spin value of $S = 1$ — with the notable exception of the Higgs boson, which is distinguished

as the only spin-0 particle of the SM. A yet further quantum number in the form of 'colour charge' is exhibited by particles which experience the strong force, and will be introduced in the section which introduces *Quantum Chromodynamics*. In accordance with the relativistic formulation of quantum mechanics set out in the preceding chapter, an anti-matter partner exists for each charged SM fermion — in practise, this corresponds to a particle with a charge value equal in magnitude and opposite in sign to that of its regular matter counter-part, with all remaining parameters otherwise indistinguishable.

The mathematical framework of the SM takes the form of a relativistic quantum field theory built upon the local gauge symmetry group $\mathsf{SU(3)_C \times SU(2)_L \times U(1)_Y}$, where $\mathsf{SU(3)_C}$ captures the symmetries of Quantum Chromodynamics, giving rise to the strong force, and $\mathsf{SU(2)_L \times U(1)_Y}$ describes the electroweak interactions. The sub-scripts C, L, and Y respectively denote that the gauge bosons of quantum chromodynamics couple only to colour-charged particles (*i.e.* quarks), that the symmetries of the $\mathsf{SU(2)}$ group are experienced only by left-handed particles, and that the reach of the $\mathsf{U(1)}$ group extends only to those particles to which a non-zero value of weak hypercharge, Y, may be ascribed.

The SM is constructed as a combination of distinct components, each of which seeking to explain a different domain of fundamental physics. Such components of the SM are now considered in turn.

2.2 Quantum Electrodynamics

Quantum electrodynamics (QED) is the quantum field theoretic account of electromagnetic interactions, coupled to electrons and positrons. While QED will be shown to be usurped by a more widely encompassing theory of electroweak unification, an examination of the theoretical underpinnings of QED is first presented — this will lay the foundation for some of the key conceptual pillars of the SM to which reference will be made throughout this chapter.

Developed largely by Dyson, Feynman, Schwinger, and Tomanaga towards the middle of the twentieth century, QED stands as the first successful attempt to fully reconcile special relativity with quantum mechanics, and builds upon the foundations developed by Dirac in the 1920s to achieve a fully consistent description of all phenomena pertaining to electrically charged particles. The

electromagnetic force is explained in terms of the exchange of (virtual) photons.

Within the framework of QED, the full confluence of electromagnetic phenomena can be viewed as an emergent consequence of *gauge symmetry*. This can be illustrated with consideration of the Dirac Lagrangian, containing only free electrons and positrons:

$$\mathcal{L} = \overline{\psi}\,(i\,\gamma^{\mu}\,\partial_{\mu} - m)\,\psi. \tag{2.1}$$

(While a Lagrangian should formally be denoted in terms of set notation, $\in$, this thesis will adopt the more relaxed convention of expressing the Lagrangian as an equality.) Here, m represents the mass of the electron or positron in-keeping with convention, and ψ is a four-component spinor field whose components define the two degrees of freedom for the spin of the electron and the two spin states of the positron. The Dirac matrix γ^{μ} is a 4×4 matrix in spinor-space which satisfies the anti-commutator given by:

$$\{\gamma^{\mu}, \gamma^{\nu}\} = 2\,g^{\mu\nu}. \tag{2.2}$$

The *adjoint spinor*, $\overline{\psi}$, takes the form $\overline{\psi} \stackrel{\text{def}}{=} \psi^{\dagger}\,\gamma^{0}$, where $\psi^{\dagger}$ is the Hermitian conjugate (the complex conjugate of the transpose) of the spinor.

It may be observed that the Dirac Lagrangian is invariant under the transformation

$$\psi \to \mathrm{e}^{i\alpha}\,\psi, \tag{2.3}$$

for a constant $\alpha \in \mathbb{R}$. It follows that $\overline{\psi} \to \mathrm{e}^{-i\alpha}\,\overline{\psi}$ also leaves the Lagrangian invariant, such that:

$$\mathcal{L} \to \left(\mathrm{e}^{-i\alpha}\,\overline{\psi}\right)(i\gamma^{\mu}\,\partial_{\mu} - m)\left(\mathrm{e}^{i\alpha}\,\psi\right) = \mathcal{L}, \tag{2.4}$$

where some basic algebraic spade-work has leveraged the laws of indices in order to cancel the exponential with its complex conjugate.

An intuition for this result may be obtained from its geometric interpretation.

The field $\psi(x)$ must necessarily possess a phase at every point in space, owing to its nature as a complex-valued field. Given that this phase, θ, must sit within the range $\theta \in (0, 2\pi)$, it may be visualised as a vector constrained by a unit circle at every point in space-time. The field ψ can be considered as the set of all such vectors. The influence of the spinor transformation may now be viewed intuitively as a rotation by α rad. The field which arises from the transformation $\psi \to e^{i\alpha}\,\psi$ can therefore be considered as the field resultant from the rotation of all vectors by constant α.

The calibration of the zero-point on the unit circle containing each arrow is free to be set as desired, since this is merely a question of setting the frame of reference. Resetting the calibration of said reference frame, or the 'gauge' of the position of all arrows, is known as a *global gauge transformation*. Given the result obtained in Equation 2.4, the Dirac Lagrangian is therefore be said to be invariant under global gauge transformation — or, equivalently, to manifest the desirable property of global gauge invariance.

The phase rotation of such transformations can be promoted to include a spatial dependency, as a function of the position four-vector, x. This may be visualised as a rotation of the complex phase to a differing degree at each point in space-time. Such transformations may be implemented as:

$$\psi \to e^{i\alpha(x)}\,\psi, \tag{2.5}$$

$$\overline{\psi} \to e^{-i\alpha(x)}\,\overline{\psi}. \tag{2.6}$$

Lagrangians which are left unchanged by such transformations are said to exhibit the property of *local gauge invariance*.

The local gauge transformation of the Dirac Lagrangian unfolds as:

$$\begin{aligned}
\mathcal{L} &\to e^{-i\alpha(x)}\,\overline{\psi}\,\left[i\gamma^{\mu}\,\partial_{\mu}\left(e^{i\alpha(x)}\,\psi\right) - m\,e^{i\alpha(x)}\,\psi\right] \\
&= e^{-i\alpha(x)}\,\overline{\psi}\,\left[i\gamma^{\mu}\,e^{i\alpha(x)}\,\partial_{\mu}\psi - \gamma^{\mu}\left(\partial_{\mu}\alpha\left(x\right)\right)\,e^{i\alpha(x)}\,\psi - m\,e^{i\alpha(x)}\,\psi\right] \\
&= \overline{\psi}\,\left(i\,\gamma^{\mu}\,\partial_{\mu} - m\right)\,\psi - \left(\partial_{\mu}\alpha\left(x\right)\right)\overline{\psi}\,\gamma^{\mu}\,\psi \\
&= \mathcal{L} - \left(\partial_{\mu}\alpha\left(x\right)\right)\overline{\psi}\,\gamma^{\mu}\,\psi.
\end{aligned} \tag{2.7}$$

The unwanted term perturbing the Lagrangian from its un-transformed formulation may be eliminated by the introduction of the *covariant derivative*, D_μ, such that the derivative transformation

$$D_\mu \psi \rightarrow e^{i\alpha(x)}\, D_\mu \psi \tag{2.8}$$

is satisfied. By comparison to the behaviour of the unaltered partial derivative under transformation, the form required of the covariant derivative in order to satisfy the requirement of local gauge invariance is deduced to be:

$$D_\mu \psi = \left[\partial_\mu + ieA_\mu\left(x\right)\right]\psi. \tag{2.9}$$

The utility of the added term $i\,e$, where e is a real constant which should not be conflated with the exponential function, will reveal itself in subsequent algebraic manoeuvres. The field $A_\mu\left(x\right)$ is a function of space-time, for which a definition much be found such that Equation 2.8 is obeyed.

The freedom for A_μ to vary under gauge transformation may be expressed as:

$$A_\mu\left(x\right) \rightarrow A'_\mu\left(x\right) \equiv A_\mu\left(x\right) + \delta A_\mu\left(x\right), \tag{2.10}$$

such that the covariant transformation proceeds as:

$$\begin{aligned} D_\mu \psi &\rightarrow \partial_\mu\left(e^{i\alpha}\,\psi\right) + ie\,A'_\mu\left(e^{i\alpha}\,\psi\right) \\ &= e^{i\alpha}\left[\partial_\mu\psi + i\left(\partial_\mu\alpha\right)\psi + ie\,\left(A_\mu + \delta A_\mu\right)\psi\right] \\ &= e^{i\alpha}\left[D_\mu\psi + i\left(\partial_\mu\alpha + e\delta A_\mu\right)\psi\right]. \end{aligned} \tag{2.11}$$

The demand imposed by the covariant transformation of Equation 2.8 therefore requires that:

$$\delta A_\mu = -\frac{1}{e}\partial_\mu\alpha, \tag{2.12}$$

provided that $\psi \rightarrow e^{i\alpha}\,\psi$.

Taking A_μ to be a field of physical significance rather than a merely mathematical tool, kinetic terms may be defined such that the equations of motion for this new field may be constructed. This aim is furthered by the introduction of the term:

$$F_{\mu\nu} = \partial_\mu A_\nu - \partial_\nu A_\mu. \tag{2.13}$$

This is known as the *electromagnetic field strength tensor*, which may be shown to exhibit gauge invariance as follows. Given that Equation 2.12 implies $\partial_\mu A_\nu \rightarrow \partial_\mu A_\nu - \frac{1}{e}\partial_\mu\partial_\nu\alpha$, the transformation for the field strength tensor

$$F_{\mu\nu} \rightarrow \partial_\mu A_\nu - \frac{1}{e}\partial_\mu\partial_\nu\alpha + \frac{1}{e}\partial_\nu\partial_\mu\alpha = F_{\mu\nu} \tag{2.14}$$

is readily obtained. The equations of motion for A_μ may now be constructed with the assistance of the kinetic term: $\mathcal{L}_{\text{kinetic}} = -\frac{1}{4}F_{\mu\nu}F^{\mu\nu}$, which enables the overall Lagrangian for QED to be expressed as:

$$\mathcal{L}_{\text{QED}} = -\frac{1}{4}F_{\mu\nu}F^{\mu\nu} + \overline{\psi}\left[i\gamma^\mu D_\mu - m\right]\psi. \tag{2.15}$$

In accordance with the calculus of variations set out in the previous chapter, the Euler-Lagrange equation may be employed in order to extract physical insight from $\mathcal{L}_{\text{QED}}$. A quantum gauge theory of electrodynamics is thereby attained.

The field A_μ, required to satisfy the necessary condition of local gauge invariance, is, in fact, the photon field. It is of profound remark that by following the conditions set by mathematical symmetry, a prediction for the physical photon field is naturally admitted by the mathematics. Maxwell's equations of classical electrodynamics are recovered from the quantum mechanical treatment of electromagnetism as a consequence of imposing the requirement of local gauge symmetry.

With reference to the mathematics introduced in Section 1.4.2, it should be noted that the local gauge transformations performed in this section may be described by the $\mathsf{U}(1)$ symmetry group. QED is therefore said to be a $\mathsf{U}(1)$ gauge theory. The remainder of the SM shall be shown in the proceeding sections to exhibit more complex symmetries.

2.3 Quantum Chromodynamics

The existence of the strong force was postulated as a possible explanation for the bound nature of nuclei, addressing the question of why nucleons do not repel one-another electrostatically. In the 1960s, a new fundamental force, much stronger than that of electromagnetism, was predicted, which was thought to overcome the force of electrostatic repulsion. The observations of many stable bound states emerging at this time could further be addressed by the existence of such a force.

It was later found that such stable hadrons behaved in a manner consistent with a collection of smaller constituent particles when examined at high energies. This measurement resulted in the acceptance of the *quark model* — earlier proposed independently by Murray Gell-Mann and George Zweig [19] — and its incorporation as an integral component of the SM.

The magnitude of the strong force has since been measured to be ~ 100 times that of electromagnetism, and $\sim 10^6$ times that of the weak force (at a distance of 10^{-15} m) [10]. The strong force acting between the constituent quarks of the proton accounts for the majority of its mass, by the principle of mass-energy equivalence.

The gauge-theoretic description of the strong interaction, developed in the 1970s, is known as quantum chromodynamics (QCD), and accounts for the behaviour of quarks and their interactions. Just as the electromagnetic force is mediated by gauge bosons in the form of photons, the strong force is carried by gauge bosons in the form of *gluons*. As the reigning gauge theory of hadronic interactions, a solid understanding of QCD in concert with accurate modelling of its phenomenological consequences is accordingly of acute import in the study of fundamental physics at hadron colliders.

The study of the quark sector requires the introduction of an additional conserved quantum number: *colour charge*. Each quark carries a colour charge, instantiated to either red (r), green (g), or blue (b). Each anti-quark, similarly, carries an anti-colour of either anti-red ($\bar{\text{r}}$), anti-green ($\bar{\text{g}}$), or anti-blue ($\bar{\text{b}}$). This quantum number may be thought of as the analogue of electric charge in QED.

2.3.1 Gauge Theory of the Strong Interaction

Quark fields possess the same phase at each point in space-time as was introduced in QED, which was shown to remain invariant under both global and local gauge transformations. As a consequence, quarks also experience the electromagnetic force. Colour charge may similarly be visualised by a space of vector arrows, where each arrow resides in an abstract *colour space* which is defined at every point in space-time. The magnitude and position of this vector-like object captures the overall colour charge — *i.e.* the degree of 'redness', 'greenness', and 'blueness' — present at that point. Such an abstract space is denoted an *internal* space (that is, internal to the quark) in order to emphasise its distinction from space-time.

As quarks are of spin $S = 1/2$, they may be described by Dirac spinors. A spinor, ψ_i, can therefore be defined for each quark colour: $\psi_i \equiv (\psi_r, \psi_g, \psi_b)$. Here, i is known as the *colour index*. This spinor resides in internal colour space, and can be considered as serving the purpose of the colour arrow in the geometrical description above.

In the construction of a Lagrangian with which to describe the strong force, global gauge invariance is assured by the demand that arbitrary recalibration of the internal space colour-axes does not exert influence on the physics which it describes. Local gauge invariance of the theory, as before, may be enforced through the introduction of a covariant derivative. It may be shown that the appropriate covariant derivative operator is defined as:

$$D_\mu = \partial_\mu + i\, g_s \frac{\lambda^a}{2} G^a_\mu, \tag{2.16}$$

which is matrix-valued in colour space, as is required for it to operate on a vector-like object. The term g_s denotes the strong coupling strength, while G^a_μ are the 8 gluon fields. Imposing the requirement that local gauge invariance be respected by the Lagrangian which describes the strong force has therefore induced the emergence of the gluon, just as the photon arose naturally from the local gauge invariance of QED.

By arguments of dimensionality, the operator D_μ must take the form of a 3×3 matrix such that it may operate on the quark fields. As such, it may be associated with the fundamental representation of the non-Abelian $\mathsf{SU}(3)$ symmetry group as introduced in Section 1.4.2. As discussed in Section 1.4.2, a group is uniquely

determined by its structure constants and not by the generators of the group — the covariant derivative is accordingly defined in terms of the *Gell-Mann matrices*, λ^a, as a convenient choice of generators which satisfy the Lie algebra of the group.

The dimension of the $\mathsf{SU}(3)$ group is found, via Equation 1.14, to be $\dim[\mathsf{SU}(3)] = 8$, with each of the eight group generators corresponding to a gauge mediator — the gluons. The existence of 8, rather than 9, gluons is thereby addressed by the symmetry of the gauge theory. A physical interpretation of this result is that the state corresponding to the combination of $(\mathsf{r\bar{r}} + \mathsf{g\bar{g}} + \mathsf{b\bar{b}})/\sqrt{3}$ is colourless overall, and therefore does not participate in the strong interaction.

The Gell-Mann matrices are given by:

$$\lambda_1 = \begin{pmatrix} 0 & 1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix}, \lambda_2 = \begin{pmatrix} 0 & -i & 0 \\ i & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix}, \lambda_3 = \begin{pmatrix} 1 & 0 & 0 \\ 0 & -1 & 0 \\ 0 & 0 & 0 \end{pmatrix}, \lambda_4 = \begin{pmatrix} 0 & 0 & 1 \\ 0 & 0 & 0 \\ 1 & 0 & 0 \end{pmatrix},$$
$$\lambda_5 = \begin{pmatrix} 0 & 0 & -i \\ 0 & 0 & 0 \\ i & 0 & 0 \end{pmatrix}, \lambda_6 = \begin{pmatrix} 0 & 0 & 0 \\ 0 & 0 & -1 \\ 0 & 1 & 0 \end{pmatrix}, \lambda_7 = \begin{pmatrix} 0 & 0 & 0 \\ 0 & 0 & -i \\ 0 & i & 0 \end{pmatrix}, \lambda_8 = \frac{1}{\sqrt{3}} \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & -2 \end{pmatrix}. \tag{2.17}$$

The Gell-Mann generators are found to satisfy $\mathsf{Tr}(\lambda_i \lambda_j) = 2\delta_{ij}$, and adhere to the commutation relation:

$$[\lambda_a, \lambda_b] = 2i\, f^{abc} \lambda^c, \tag{2.18}$$

imposed by the Lie algebra, where the structure constants introduced in Section 1.4.2, are given by f^{abc}.

Having defined the appropriate covariant derivative for this gauge theory, the Lagrangian for QCD may be expressed as:

$$\mathcal{L}_{\mathsf{QCD}} = \sum_f \overline{\psi}_f \left(i\gamma_\mu D^\mu - m_f \right) \psi_f - \frac{1}{4} G^a_{\mu\nu} G_a^{\mu\nu}, \tag{2.19}$$

where the sum has been taken over each flavour, f, in the quark sector. The quark mass and Dirac matrices are respectively given by m and γ^μ, as before. The quark fields whose dynamics are governed by QCD are denoted, for a given flavour, by ψ_f. The adjoint field, $\overline{\psi}_f$, is set to $\overline{\psi}_f = (\psi^*)^\mathsf{T}$. The components of

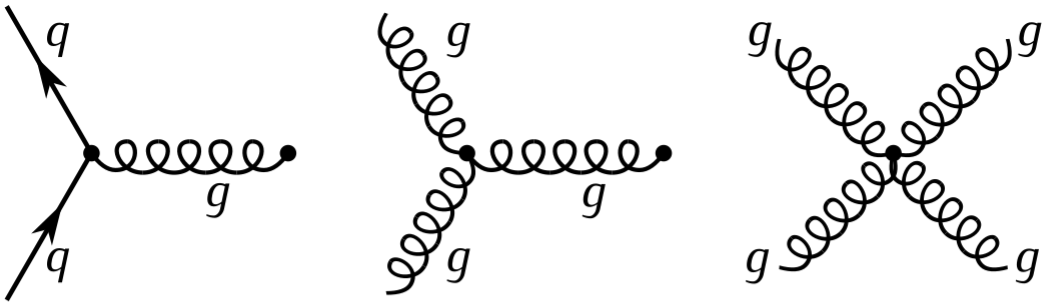

Figure 2.2 Feynman diagrams depicting the gluon coupling to (*left*) quarks, together with the gluon tri-linear and quartic self-interactions.

each quark field are given by the colour spinors, ψ_i, for a given quark flavour, f. An additional term has been introduced for the gluonic field-strength tensor, defined as:

$$G^a_{\mu\nu} = \partial_\mu G^a_\nu - \partial_\nu G^a_\mu - g_s f^{abc} G^b_\mu G^c_\nu, \tag{2.20}$$

for the structure constants of the SU(3) Lie algebra, f^{abc}, and the strong coupling strength, g_s. The non-Abelian nature of the SU(3) group gives rise to triple and quartic gluonic self-interations — shown in Figure 2.2. This marks a further contrast with QED, in which photons do not interact with one-another. While the photon is not in possession of an electric charge, the combination of colour and anti-colour charge carried by each gluon facilitates the interaction of gluons with one-another via the strong force.

2.3.2 Renormalisation and Asymptotic Freedom

As is true for all QFTs, the calculation of measurable quantities such as cross-sections in QCD involves contributions from higher-order quantum loops. While the four-momenta must be conserved at every vertex in processes which involve such loops, the intermediate momenta are free to assume any value — quantum loop contributions are comprised of virtual particles, and therefore enjoy the additional liberty afforded by the Heisenberg Uncertainty Principle.

The computation of physical observables with the inclusion of quantum corrections results in the problem of divergent integrals, and the prediction of seemingly untameable, unphysical infinities. The solution to this mathematical impasse was found to lie in the *renormalisation* of coupling constants and fermion masses.

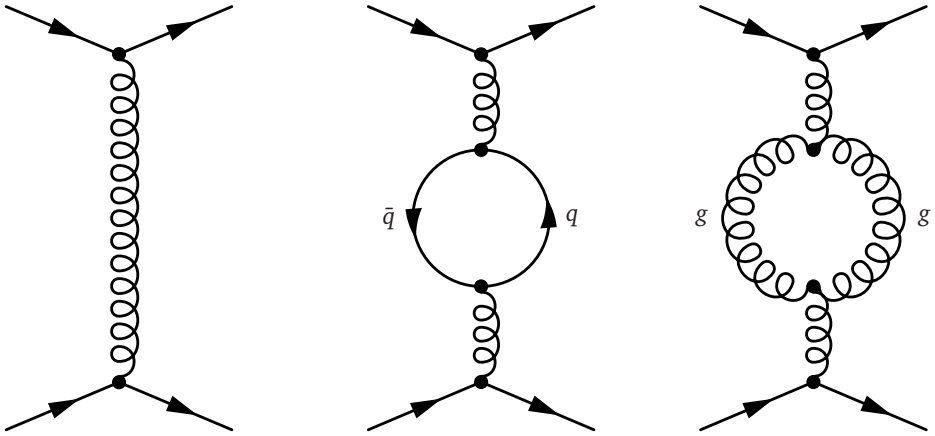

Figure 2.3 Feynman diagrams of contributions to the quark-quark scattering process, which feature in the perturbative expansion associated with the corresponding cross-section. These include a virtual fermion loop and a virtual boson loop — the latter of which is rendered permissible by the gluon self-interaction terms of QCD.

The need for renormalisation is not unique to QCD, but is required of all QFTs — indeed, the ability to renormalise a candidate theory is regarded as an indication of its viability. As a result of gluonic self-interactions, however, the renormalisability of QCD gives rise to unique physics which merits its discussion here.

An intuition for the physics of renormalisation may be developed through first giving consideration to the electromagnetic force. Polarisation of the quantum vacuum leads to the production of virtual e^+e^- pairs; in the presence of a real electron, the positron component of this pair is attracted towards it on account of its positive charge, while the negatively charged virtual electron is repelled. This phenomena induces a 'cloud' of virtual positrons which are said to 'shield', or 'screen', the *bare* charge of the real electron. When observed from a distance, the *effective* charge of the electron is reduced from its original bare value due to the presence of positive charge from the virtual positrons. This is termed the 'dressed' charge. The corrections resulting from quantum mechanical screening can be evaluated through determination of the probabilities corresponding to each interaction with virtual particle states emerging from the vacuum — an infinite series of such interactions.

The trick of renormalisation lies in the recognition that such processes were always present and interacting with the underlying particle, influencing their measured (dressed) properties. As such, many of the virtual loop diagrams introduced in

the perturbative computation of observables such as the cross-section for a given process were already present in the definition of the underlying particle — that is to say, they have already been considered implicitly in the computation. Their inclusion in the computation of the quantity of interest therefore results in a double-count of the effects of such virtual processes. A further normalisation of the quantity such that the double-counting of loop processes is addressed — the renormalisation step — results in the cancelation of the infinities produced by such processes, yielding a finite prediction. Theoretical calculations of quantities which are performed using renormalisation are found to agree precisely with measurement [4, 6, 10].

Further to virtual electrons and positrons, quantum fluctuations of the vacuum also produce virtual $q\overline{q}$ pairs and virtual gluons. This leads the colour charge of quarks to experience quantum-mechanical screening in the same manner as the electric charge of the electron. The colour screening effect of this can be calculated by considering the probabilities of each virtual contribution, as before. The renormalisation of QCD processes is distinguished, however, by the inclusion of gluonic self-interaction terms (as shown in Figure 2.3) — for which there is no photonic counterpart in QED.

While the screening of electric charge from vacuum polarisation processes leads to a decrease in effective electric charge, the inclusion of gluonic terms in the screening of colour charge is found to increase the strength of effective colour charge. Equivalently, it may be stated that the effective strength of colour charge on a quark is found to decrease as the distance between quarks decreases. For this reason, the screening of colour charge is sometimes termed 'anti-screening'. This can be mapped to the afore-stated observation that the strong force increases with separation — a phenomenon first observed in experiments investigating the nature of deep inelastic scattering.

This may be understood as follows. As the charge of a given quark is probed at small distances, the cloud of virtual particles which inflate the charge has been penetrated — their influence is not felt, as they exist at a greater radial distance from the quark. The charge at this distance is therefore the bare charge. As the distance from which the quark charge is measured increases, so too does the number of virtual particles which may contribute to the effective charge. The strength of this charge — and the strong force it induces — is therefore measured to increase with separation.

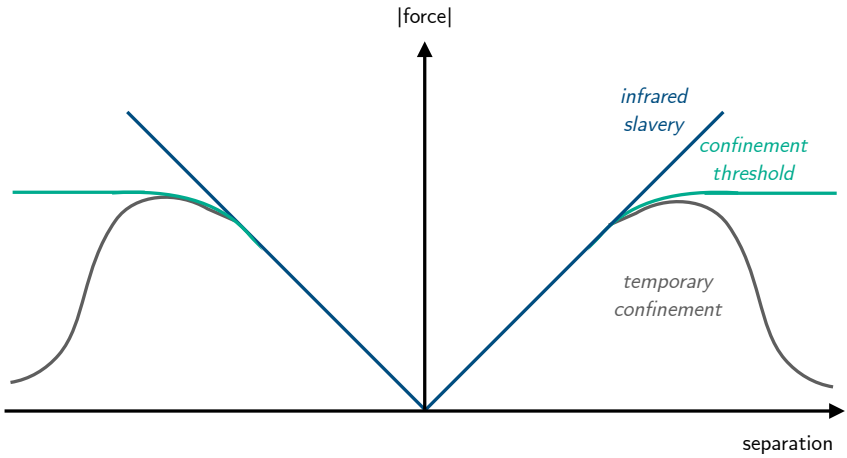

Figure 2.4 sketch of possible colour confinement scenarios, illustrating the possibility that the force with which quarks are retained in bound states may converge to a constant value, or decrease, beyond a given threshold. The scenario in which colour confinement extends to all energy regimes — shown as 'infrared slavery' — is thought to be an intrinsic property of QCD.

This phenomenon is known as 'asymptotic freedom' [20]. The terminology alludes to the fact that, as the inter-quark distances become asymptotically small, so too do the chromodynamic forces between them — allowing the quarks to behave as though they are free particles. The discovery of the asymptotically-free behaviour of QCD not only accounts for the experimental observation of a strong force which increases in magnitude with distance (often termed 'infrared slavery' — see Figure 2.4); moreover, it enables perturbative techniques to be employed in the computation of quantum-chromodynamic properties [10]. This is possible as a result of the strong coupling strength assuming a sufficiently small value in certain conditions such that a perturbative approach becomes analytically valid [4, 6, 7]. The significance of this development led to David J. Gross, H. David Politzer, and Frank Wilczek being awarded the Nobel prize in physics in 2004 [21].

The phenomenon of asymptotic freedom is intimately connected to another unique feature of the gauge theory of strong interactions: *colour confinement*. This feature of QCD is of immutable importance to the study of fundamental physics at hadron colliders, as will be investigated in the sections which follow.

2.3.3 Colour Confinement

Throughout the history of experimental high energy physics, no observation of a single, free quark has ever been made. Instead, quarks are always observed to be *confined* within hadrons of an overall colour-neutral charge.

While the phenomenon of infrared slavery introduced in the discussion of asymptotic freedom seems to account for this behaviour, it is not known whether such confinement extends indefinitely at all energies, or if an energy threshold, beyond which quarks are liberated, may exist beyond the energetic reach of current collider facilities [10]. At the energies achieved at particle colliders to-date, no observation of free quarks has been made. Such a threshold must therefore extend beyond the current energy regime — on the order of 13 TeV — if it exists at all.

It is hypothesised that colour confinement of this nature is, in fact, an intrinsic property of QCD at all energy scales — a property thought to arise as a result of the non-Abelian nature of the gauge theory. While this question is outwith the regime in which perturbative techniques may be utilised, this hypothesis is consistent with calculations performed using lattice methods [22].

The property of colour confinement in QCD can be understood as follows. As quarks produced in high energy collisons begin to separate in space, the force between them — as discussed in Section 2.3.2 — will increase. The chromodynamic field lines which emanate between the quarks do not tend to infinity, as is true of the field lines associated with electrostatic force which extend spherically outward. Instead, due to the non-Abelian nature of QCD, self-interactions of the gluonic gauge field cause the chromodynamic field lines to converge back together [10] — this is visualised in the simulation presented in Figure 2.5.

This leads to the formation of 'flux-tubes' between the quarks. As more work is expended by the separation of the quarks, the energy stored in the flux-tube becomes sufficient to promote a virtual $q\overline{q}$ pair from the quantum vacuum. This process reiterates as the new quark and anti-quark separate, resulting in a spray of bound hadrons — a process known as *hadronisation*. The phenomenon of hadronisation is consistent with all experimental data recorded at the time of writing.

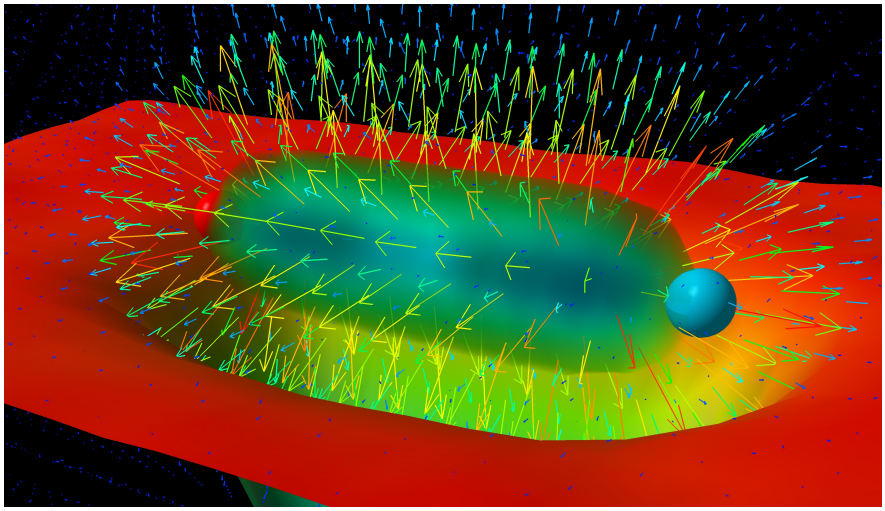

Figure 2.5 numerical simulation of chromodynamic flux-tube formation between a pair of quarks. The potential stored in such flux-tubes creates $q\overline{q}$ pairs from the quantum vacuum, leading to the phenomenon of hadronisation. Image source: [23].

2.4 Electroweak Unification

The first theory of the weak interaction was proposed by Fermi in 1933 as an attempt to explain radioactive β-decay [24]. Fermi proposed that a four-fermion interaction was responsible for the β-decay of the neutron, postulating that the neutron couples directly to a proton, electron, and electron-neutrino (later understood to be an anti-neutrino). Fermi's theory of the weak interaction evolved over subsequent years to further account for the decay of the muon.

With the successful development of QED, efforts were made to reconcile Fermi's theory with the more modern quantum field-theoretic view of fundamental physics. It was found that QFT could not accommodate Fermi's theory in a renormalisable manner [10]. This, together with aesthetic objections, lead the theory to be usurped by a more compelling model of the weak interaction.

The remarkable unification of electricity and magnetism by Maxwell into a single fundamental force fired the starting gun on the pursuit of ever-greater unification, perhaps culminating in a grand unified theory of everything at a high energy scale. This logic was applied with success to the unification of the quantum theory of electromagnetism with the weak force by Glashow, Weinberg, and Salam in the 1960s, resulting in the *The Glashow-Weinberg-Salam (GWS) Model* [25, 26]. This achievement represents one of the most significant milestones in the development

of the SM, earning each of the three founders the Nobel prize in 1979.

The GWS model lead to the prediction of three new bosons: the charged $W^{\pm}$ bosons, and the neutral Z boson. With their discovery in 1983 at CERN, the GWS was confirmed to be consistent with experimental observation [10].

2.4.1 The Weak Interaction

The weak interaction is described by the $\mathsf{SU}(2)_{\mathsf{L}}$ gauge symmetry group. Gauge transformations are generated by $N^2 - 1 = 3$ generator matrices,

$$t^i = \frac{1}{2}\sigma^i, \tag{2.21}$$

for the Pauli matrices, σ^i, with $i = 1, 2, 3$. The generators of the $\mathsf{SU}(2)_{\mathsf{L}}$ group obey the commutation relation:

$$\left[t^i, t^j\right] = i\epsilon^{ijk}t^k, \tag{2.22}$$

where ϵ^{ijk} is the anti-symmetric Levi-Civita tensor defined according to the symmetrical properties of the index permutations:

$$\epsilon^{ijk} = \begin{cases} +1 & : \text{symmetric permutation} \\ -1 & : \text{anti-symmetric permutation} \\ 0 & : \text{otherwise} \end{cases}. \tag{2.23}$$

This commutation relation serves as the Lie algebra of the $\mathsf{SU}(2)_{\mathsf{L}}$ group, defining it uniquely.

The weak interaction is experienced by states which have a non-zero value of *weak isospin*, I_3, the conserved quantum number which serves an analogous role to the electric charge or colour charge in the electromagnetic and chromodynamic gauge theories.

Local gauge invariance is assured by the introduction of three gauge fields, analogously to the inclusion of the photonic and gluonic gauge fields in the afore-reviewed theories of QED and QCD.

Similarly to QCD, the gauge theory of the weak interaction is non-Abelian in nature. The electroweak field-strength tensor is given by:

$$W^i_{\mu\nu} \equiv \frac{i}{g_W}[D_\mu, D_\nu] = \partial_\mu W^i_\nu - \partial_\nu W^i_\mu + g_W \epsilon^{ijk} W^j_\mu W^k_\nu, \tag{2.24}$$

where g_W is the coupling constant of the weak force, and D_μ is the covariant derivative.

Particles do not, in general, travel along the axis of their spin. The direction of spin along the axis of travel for a given particle is quantified by the quantum number of *helicity* [3], defined as:

$$\widehat{h} \equiv \frac{\vec{S} \cdot \vec{p}}{|\vec{p}|}, \tag{2.25}$$

where $\vec{p}$ is the momentum three-vector of the particle.

For a fermion of spin $S = \frac{1}{2}$, its associated helicity is instantiated to $h = +\frac{1}{2}$ should its spin point in the direction of travel, or $h = -\frac{1}{2}$, should its spin point in the opposite direction to the direction of travel. Fermions of with positive helicity are said to be "right-handed", whereas those with negative helicity are labelled "left-handed".

It is convenient to define an additional quantum number: *chirality* [3]. This quantum number is defined as a function of the $\gamma^5 = -i\gamma^0\gamma^1\gamma^2\gamma^3$ matrix which satisfies the anti-commutation relation: $\{\gamma^5, \gamma^\mu\} = 0$. Spinors which are eigenstates of the γ^5 matrix are said to be chiral, and are either left-handed or right-handed.

Parity symmetry (as introduced in Section 1.4.3) is known empirically to be broken under the weak interaction, as discovered by C.S. Wu in 1957 [27]. The property of parity violation dictates that the weak force must interact differently with states of differing chirality [3]. The fermion spinors may be decomposed into their left-handed and right-anded components, ψ_L and ψ_R, through appeal to the chiral projection operators:

$$\psi_L = \widehat{P}_L \psi, \qquad \psi_R = \widehat{P}_R \psi, \tag{2.26}$$

where the projection operators are defined as:

$$\widehat{P}_{\mathsf{L}} = \frac{1}{2}\left(1-\gamma^5\right), \qquad \widehat{P}_{\mathsf{R}} = \frac{1}{2}\left(1+\gamma^5\right). \tag{2.27}$$

The behaviour of particles which experience the weak force may now be understood in terms of their chirality. Left-handed chiral fields form weak isospin doublets which correspond to fermions of the same flavour, differing by one unit of electric charge. This is realised by doublets containing leptons of each generation together with their associated neutrino:

$$\psi_{\mathsf{L}}^{\mathsf{lepton}} = \begin{pmatrix} \nu_{e,\mathsf{L}} \\ e_{\mathsf{L}} \end{pmatrix}, \quad \begin{pmatrix} \nu_{\mu,\mathsf{L}} \\ \mu_{\mathsf{L}} \end{pmatrix}, \quad \begin{pmatrix} \nu_{\tau,\mathsf{L}} \\ \tau_{\mathsf{L}} \end{pmatrix}. \tag{2.28}$$

Quarks are similarly described by doublets of left-handed up-type and down-type quarks of each generation:

$$\psi_{\mathsf{L}}^{\mathsf{quark}} = \begin{pmatrix} u_{\mathsf{L}} \\ d_{\mathsf{L}} \end{pmatrix}, \quad \begin{pmatrix} c_{\mathsf{L}} \\ s_{\mathsf{L}} \end{pmatrix}, \quad \begin{pmatrix} t_{\mathsf{L}} \\ b_{\mathsf{L}} \end{pmatrix}. \tag{2.29}$$

Right-handed fermion fields which do not partake in the weak interaction are then described by singlet states of the form:

$$\psi_{\mathsf{R}}^{\mathsf{lepton}} = \nu_{e,\mathsf{R}}, \nu_{\mu,\mathsf{R}}, \nu_{\tau,\mathsf{R}}, e_{\mathsf{R}}, \mu_{\mathsf{R}}, \tau_{\mathsf{R}}, \qquad \psi_{\mathsf{R}}^{\mathsf{quark}} = u_{\mathsf{R}}, d_{\mathsf{R}}, c_{\mathsf{R}}, s_{\mathsf{R}}, t_{\mathsf{R}}, b_{\mathsf{R}}. \tag{2.30}$$

The components of weak isospin may now be defined as:

$$I_3 = \begin{cases} +\frac{1}{2} & : \text{up-type } \psi_{\mathsf{L}} \\ -\frac{1}{2} & : \text{down-type } \psi_{\mathsf{L}} \\ 0 & : \psi_{\mathsf{R}} \end{cases}. \tag{2.31}$$

Anti-fermions experience the same interaction with inverse chirality. As the weak interaction couples only to left-handed (right-handed) fermions (anti-fermions), it may be described as a 'maximally chiral' theory.

2.4.2 The Glashow-Weinberg-Salam Model

In the 1960s, Glashow, Weinberg, and Salam developed a theory which successfully unifies QED with the weak interaction [25, 26]. The newly-unified *electroweak* interaction is represented by the $\mathsf{SU}(2)_\mathsf{L} \times \mathsf{U}(1)_Y$ group, where the subscript L denotes the fact that the weak interaction is experienced only by states of left-handed chirality, while Y represents a new conserved quantum number introduced by the theory: *hypercharge*.

Hypercharge is related to the generators of the electroweak group by the Gell-Mann-Nishijima relation:

$$Q = t_3 + \frac{Y}{2}, \tag{2.32}$$

where Q is the electric charge familiar from QED.

The unification of QED with the weak force requires a modification to the relevant covariant derivative, defined within the GWS model as:

$$D_\mu = \partial_\mu - i\, g_W \sigma_a W_\mu^a - \frac{i}{2} g^{'} Y\, B_\mu, \tag{2.33}$$

such that local gauge invariance of the unified theory is preserved. Here, σ_a are the usual Pauli matrices, $g^{'}$ is a dimensionless coupling constant, and B_μ is a new gauge field associated with the $\mathsf{U}(1)_Y$ symmetry group. Index a is taken to run from 1 to 3.

The gauge fields of the GWS model may be parameterised as:

$$W_\mu^\pm = \frac{1}{\sqrt{2}} \left(W_\mu^1 \mp i\, W_\mu^2 \right), \tag{2.34}$$

,

$$A_\mu = \cos\theta_W B_\mu + \sin\theta_W W_\mu^3, \tag{2.35}$$

$$Z_\mu = -\sin\theta_W B_\mu + \cos\theta_W W_\mu^3, \tag{2.36}$$

where A_μ is the photon field previously introduced in QED, and θ_W is the *electroweak mixing angle*. The degree of gauge mixing between the $\mathsf{SU}(2)_\mathsf{L}$ and

$\mathsf{U}(1)_Y$ groups is quantified by the electroweak mixing angle in accordance with:

$$\sin\theta_W = \frac{g'}{\sqrt{g_W^2 + g'^2}}, \qquad \cos\theta_W = \frac{g}{\sqrt{g_W^2 + g'^2}}. \tag{2.37}$$

The mixing of the $W_\mu^\pm$ and B_μ gauge fields thereby accounts for the existence of the photon, while predicting the existence of a new bosonic state: the neutral Z boson.

The gauge term of the electroweak Lagrangian may be expressed as:

$$\mathcal{L}_{\text{gauge}} = -\frac{1}{4} W_{\mu\nu}^i W^{\mu\nu,\, i} - \frac{1}{4} B_{\mu\nu} B^{\mu\nu}, \tag{2.38}$$

where $B_{\mu\nu} = \partial_\mu B_\nu - \partial_\nu B_\mu$ is the field-strength tensor for the gauge field B_μ, analogous to the $F_{\mu\nu}$ field-strength tensor of QED. Similarly, the Dirac term of the electroweak Lagrangian is given by:

$$\mathcal{L}_{\text{Dirac}} = i \sum_f \overline{\psi}_{\text{L}} \gamma^\mu D_\mu \psi_{\text{L}} + i \sum_f \overline{\psi}_{\text{R}} \gamma^\mu D_\mu \psi_{\text{R}}, \tag{2.39}$$

where f iterates over each charged fermion. From this expression, together with Equation 2.33 and the the expressions for the electroweak gauge fields, Lagrangians for the charged and neutral currents carried by the electroweak bosons may be derived. Expressing the Lagrangian for the neutral current in terms of a photonic component and a component carried by the Z boson, the Lagrangian for QED may be recovered: $\mathcal{L}_{\text{NC}} = \mathcal{L}_{\text{QED}} + \mathcal{L}_Z$. The electromagnetic interaction may therefore be regarded as fully accommodated by the more fundamental framework provided by electroweak unification.

The GWS model came to be incorporated as a core pilar of the SM, whose predictions of the W and Z boson awaited discovery such that electroweak unification may be experimentally vindicated.

However, the discoveries of the electroweak bosons in 1983 were accompanied by a problem. The bosons were measured to be massive, thereby breaking the local gauge invariance of the electroweak Lagrangian. This revelation posed an existential threat to the entirety of the SM, until rescued by the Higgs mechanism.

2.5 Brout-Englert-Higgs Mechanism

The Brout-Englert-Higgs mechanism (termed 'Higgs mechanism' hereinafter for brevity) is an integral component of the SM. Without it, the SM would not stand as an experimentally-supported theory. The conceptual machinery provided by the Higgs mechanism presents a means by which the local gauge symmetry of the SM may be spontaneously broken in nature, thereby providing a theoretical explanation for the origin of the masses of elementary particles.

The Higgs mechanism was first postulated by Peter Higgs of The University of Edinburgh in 1964 [28], and independently by Robert Brout and François Englert of L'Université Libre de Bruxelles [29], as an explanation for the finite masses of the electroweak W and Z gauge bosons. Once fully incorporated into the SM, the consequential prediction to arise from the Higgs mechanism — namely, the existence of the Higgs boson — came to represent the Holy grail of modern science, the discovery of which would render the SM consistent with empirical observation.

The experimental effort to discover the boson predicted by Higgs would thereafter unfold over a number of decades, encompassing scientists from around the world. This effort culminated in the reported discovery of a scalar boson consistent with that predicted by Higgs by the ATLAS and CMS collaborations in 2012 [16, 17] — a discovery widely regarded as one of the greatest scientific accomplishments and intellectual achievements of mankind. With the discovery of the Higgs boson, the Higgs mechanism and associated scalar Higgs field was accepted as the primary origin of elementary mass in the Universe.

2.5.1 Spontaneous Symmetry Breaking

There exist in physics many asymmetric solutions to problems which adopt an otherwise symmetric form. This can be understood to arise as a consequence of the fact that the symmetric state is not necessarily the state in which the energy of the system is minimised. In such cases, as the system seeks to return to the state of minimum energy, the symmetry once manifest in the system is lost — the symmetry is said to be 'broken'.

One manner in which this may be realised is that of *spontaneous symmetry*

breaking. The symmetry of a physical system may be said to be spontaneously broken when the ground-state solution (or the state of thermal equilibrium) does not respect the same symmetries as the Lagrangian which describes the system. An intuition for this may be developed through the visualisation of a marble in a wine bottle. As the marble sits precariously atop the bulge at its base, the bottle can be said to exhibit rotational symmetry. The system is, however, unstable — only a small disturbance is required to displace the marble from its position. This is because the marble is in possession of potential energy on account of its elevation; the marble is not in its state of minimum energy. Upon the return of the marble to its state of minimum energy, or 'ground state', the rotational symmetry of the wine bottle is broken. Spontaneous symmetry breaking in fundamental particle physics may be considered to be somewhat analogous to this classical example.

Spontaneous Breaking of Global Gauge Symmetries

Spontaneous symmetry breaking in the context of a globally gauge-invariant theory corresponds to the scenario in which, although the Lagrangian of the theory is symmetric, the ground state of system described by the Lagrangian is asymmetric.

The ground state of a quantum field — *i.e.* its vacuum state — is that in which the field assumes its lowest possible energy-level, in the absence of field excitations (or particles). For the majority of quantum fields, the energy is minimised when the average value of the field is zero. For some fields, however, the energy is minimised upon the field assuming a non-zero value — a concept to which the subsequent discussion of the Higgs field will return.

Where a global gauge symmetry is broken, it may be shown that a massless, spin-0 'Goldstone' boson is admitted. No such massless, spin-0 particle is observed in nature. Further, the spontaneous breaking of global symmetries renders the theory at hand non-renormalisable [10, 3], consequentially leading to unphysical infinities as was introduced in Section 2.3.2.

Spontaneous Breaking of Local Gauge Symmetries

Spontaneous symmetry breaking may also be considered for Lagrangians which prove to be invariant under local gauge transformations. In demanding that local gauge invariance be respected, a gauge particle is introduced — this may be understood in the same manner in which the photon field, A_μ, was shown to emerge from the locally gauge invariant construction of QED in Section 2.2.

In the process of spontaneously breaking the local gauge symmetry, the massless gauge particle introduced by the enforcement of local invariance (analogous to the massless photon in QED) *mixes* with the vacuum value of the field such that it acquires a mass term. The emergent mathematics which gave rise to the Goldstone bosons in the case of global symmetry breaking becomes absorbed by the massive gauge term, such that it no longer corresponds to the prediction of a new physical state. It is colloquially said that the gauge particle 'eats' the Goldstone boson and thereby becomes massive.

2.5.2 Electroweak Symmetry Breaking

Local gauge invariance is one of the foundational corner-stones upon which the SM is built. As was shown in the previous introduction to QED in Section 2.2, the enforcement of local gauge invariance naturally gives rise to a gauge boson without an associated mass term. In the case of QED and QCD, the $\mathsf{U}(1)$ and $\mathsf{SU}(3)$ local symmetries demand that the mediators of the electromagnetic and strong forces be massless. Similarly, the $\mathsf{SU}(2)$ group symmetries of the electroweak sector necessitate that the W and Z electroweak mediators be massless. In the case of the photon and the gluon, this is in agreement with experimental measurement (hence why light is able to travel at the maximum relativistically-permissible velocity which does not violate causality). The discovery of the electroweak bosons, by contrast, revealed the W and Z to be of finite, non-zero mass. Should a reconciliation between the local gauge principle and the masses of the electroweak mediators fail to materialise, the core conceptual undergirding of the SM would be rendered invalid by nature.

The solution was found by Higgs to lie in spontaneous symmetry breaking — an idea borrowed from solid state physics. The mechanism proposed by Higgs extends the SM Lagrangian by:

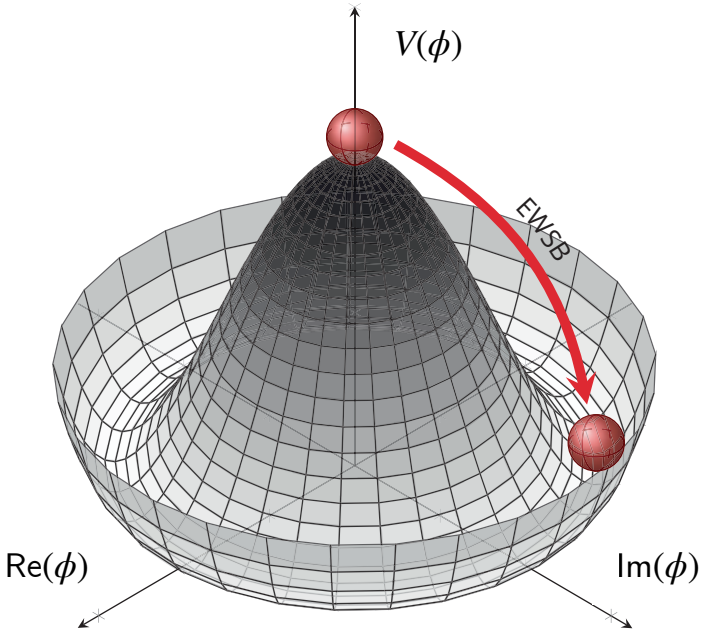

Figure 2.6 graphical depiction of the Higgs potential, $V(\phi)$, as a function of the real and imaginary components of the complex field, ϕ. The red arrow illustrates the field transitioning to its state of minimum energy through the process of (spontaneous) electroweak symmetry breaking (EWSB), thereby losing its symmetrical form. Figure adapted from [30].

$$\begin{aligned}\mathcal{L}_{\text{Higgs}} &= (D_\mu \phi)^\dagger (D^\mu \phi) - V(\phi) \\ &= (D_\mu \phi)^\dagger (D^\mu \phi) - \mu^2 (\phi^\dagger \phi) - \lambda (\phi^\dagger \phi)^2 ,\end{aligned} \tag{2.40}$$

where ϕ is an $\mathrm{SU}(2)$ doublet of complex-valued scalar fields. The components of ϕ may be expressed as:

$$\phi = \begin{pmatrix} \phi^+ \\ \phi^0 \end{pmatrix} = \frac{1}{\sqrt{2}} \begin{pmatrix} \phi_1 + i\phi_2 \\ \phi_3 + i\phi_4 \end{pmatrix}, \tag{2.41}$$

where the field has a hypercharge value of $Y = 1$. Equation 2.40 introduces the terms μ and λ: two free parameters of the Higgs potential. The free parameters are instantiated such that $\mu^2 < 0$ and $\lambda > 0$, yielding the 'Sombrero shaped' potential illustrated in Figure 2.6. This may be compared to the example of the wine bottle provided in the introductory remarks. The term λ may be interpreted as the coupling strength of the Higgs boson to itself.

While the potential illustrated in Figure 2.6 represents the simplest geometrical

form required to facilitate spontaneous symmetry breaking within the SM, the potential geometry favoured by nature has not, at the time of writing, been measured empirically. The Higgs potential has many physical and cosmological implications, and its measurement is accordingly a priority for the ATLAS collaboration.

As previously alluded to, the minimum, ϕ_0, of $V(\phi)$ occurs at non-zero values of ϕ, where such values satisfy:

$$\phi^\dagger\phi = \frac{1}{2}\left(\phi_1^2 + \phi_2^2 + \phi_3^2 + \phi_4^2\right) = -\frac{\mu^2}{\lambda}. \tag{2.42}$$

The minimum of the field can be taken as any value which satisfies Equation 2.42. Letting the minimum, ϕ_0, assume the value $\phi_0 = 2^{-1/2}(0, v)^T$, the Higgs potential may be expanded about the minimum by a small field excitation, $h(x)$:

$$\phi(x) = \frac{1}{\sqrt{2}}\begin{pmatrix} 0 \\ v + h(x) \end{pmatrix}, \tag{2.43}$$

such that the expression for the Higgs potential becomes

$$\begin{aligned} V\left(\phi(x)\right) &= \mu^2\frac{1}{2}\left(v + h(x)\right)^2 + \lambda\frac{1}{4}\left(v + h(x)\right)^4 \\ &= \underbrace{-\frac{1}{4}\lambda v^4}_{\text{VEV}} + \underbrace{\lambda v^2 h^2}_{\text{mass term}} + \underbrace{\lambda v h^3 + \frac{1}{4}\lambda h^4}_{\text{interaction terms}}, \end{aligned} \tag{2.44}$$

where we have used the fact that $\mu^2 = -\lambda v^2$. The field excitation, $h(x)$, takes the physical form of a spin-0 boson, just as an excitation of the photon field produces a photon. This is the prediction of the famed Higgs boson. By arguments of dimensionality, the mass term for the Higgs boson may be extracted from Equation 2.44 as:

$$m_H = \sqrt{2\lambda}v. \tag{2.45}$$

Further, an expression for the vacuum expectation value (VEV) of the Higgs field is given by: $v_0 = |\mu|/\sqrt{\lambda} = 2m_W/g = 246\,\text{GeV}$. This value defines the electroweak

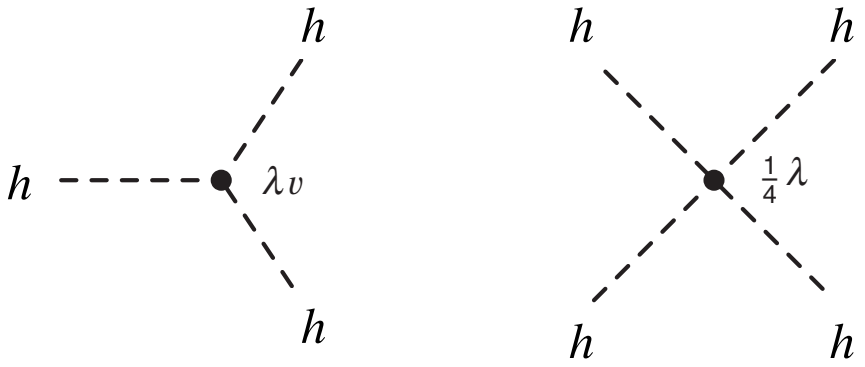

Figure 2.7 Feynman diagrams corresponding to the tri-linear and quartic self-interactions of the Higgs boson, each with a respective coupling strength of λv and $\lambda/4$. Figure adapted from [3].

scale, determined indirectly through the measurement of the W mass, m_W. It is a priority of high energy physics to measure directly the self-coupling strength λ, such that the VEV together with the shape of the Higgs potential may be determined from direct measurement — empirically probing for the potential influence of new physics.

The latter two terms of the expanded Higgs potential describe the nature of the interactions between the newly-predicted Higgs boson. The term $\lambda v h^3$ gives rise to a tri-linear interaction between three Higgs bosons with an interaction strength of λ, while $\lambda h^4/4$ induces a quartic interaction between four Higgs bosons with an interaction strength of $\lambda/4$. The Feynman diagrams for such self-interaction processes are shown in Figure 2.7.

Gauge Boson Masses

The manner in which the Higgs mechanism endows gauge bosons with finite mass may be illuminated through consideration of the covariant derivative. The construction of covariant derivatives was previously shown to secure the preservation of local gauge symmetries. In this spirit, a covariant derivative may be defined for the electroweak sector:

$$\partial_\mu \rightarrow D_\mu = \partial_\mu + i g_W \sum_{i=1}^{3} \frac{1}{2} \sigma_i \, W_\mu^{(i)} + i g' \frac{Y}{2} B_\mu, \tag{2.46}$$

for the three Pauli matrices, σ_i, weak interaction strength, g_w, and the W-boson

field, $W_\mu^{(i)}$. Recalling that the Lagrangian for the Higgs field takes the form:

$$\mathcal{L}_{\text{Higgs}} = (\partial_\mu \phi)^\dagger (\partial^\mu \phi) - V(\phi),$$

the substitution $\partial_\mu \rightarrow D_\mu$ requires the term $D_\mu \phi$ to be computed. The hypercharge for field ϕ is $Y = 2(Q - I_W^{(3)}) = 1$, therefore:

$$D_\mu \phi = \left[\partial_\mu + \frac{1}{2} \left(i g_W \sum_{i=1}^{3} \sigma_i \, W_\mu^{(i)} + i g^{'} B_\mu \right) \right] \phi. \tag{2.47}$$

To obtain the mass terms for the gauge bosons within the Higgs Lagrangian, the product of the contravariant derivative of the Higgs field with the Hermitian conjugate of the covariant derivative of the Higgs field must be evaluated. Through algebraic spade-work, the following expression may be obtained:

$$\begin{aligned} (D_\mu \phi)^\dagger (D^\mu \phi) = & \frac{1}{8} g_W^2 \left(W_\mu^{(1)} W^{\mu(1)} + W_\mu^{(2)} W^{\mu(2)} \right) (v + h)^2 \\ & + \frac{1}{8} \left(-g_W W_\mu^{(3)} + g^{'} B_\mu \right) \left(-g_W W^{\mu(3)} + g^{'} B^\mu \right) (v + h)^2 \\ & + \frac{1}{2} \partial_\mu h \, \partial^\mu h. \end{aligned} \tag{2.48}$$

Post-expansion of the $(v + h)^2$ factors, Equation 2.48 may be contrasted with the form and dimensionality expected of Lagrangian mass terms. By such a comparison, it may be deduced that the mass of the electroweak W boson is given by:

$$m_W = \frac{1}{2} g_W \, v. \tag{2.49}$$

Through the introduction of a new scalar, spin-0 quantum field whose symmetry is broken spontaneously, the Higgs mechanism has successfully attributed non-zero mass to the electroweak W boson without violating the local gauge invariance of the parent Lagrangian. Moreover, the testable prediction of a new scalar, spin-0 boson — the Higgs boson — has arisen naturally from the mechanism. The coupling of the Higgs boson to elementary particles may be deduced from the multiplicative factors of h and h^2 in the expanded Higgs Lagrangian in a manner

similar to the deduction of the W mass.

The term of Equation 2.48 which contains the third component of the W-spinor may be expressed in terms of a mixing matrix:

$$\begin{aligned}
&\frac{1}{8}v^2\left(g_W W_\mu^{(3)} - g' B_\mu\right)\left(g_W W^{\mu(3)} - g' B^\mu\right) \\
&= \frac{1}{8}v^2\begin{pmatrix} W_\mu^{(3)} & B_\mu \end{pmatrix}\begin{pmatrix} g_W^2 & -g_W g' \\ -g_W g' & g'^2 \end{pmatrix}\begin{pmatrix} W^{\mu(3)} \\ B^\mu \end{pmatrix} \\
&= \frac{1}{8}v^2\begin{pmatrix} A_\mu & Z_\mu \end{pmatrix}\begin{pmatrix} 0 & 0 \\ 0 & g_W^2 + g'^2 \end{pmatrix}\begin{pmatrix} W^{\mu(3)} \\ B^\mu \end{pmatrix}.
\end{aligned} \tag{2.50}$$

Applying the same logic as was employed to arrive at the mass of the W boson, mass terms for two bosons which couple to the A_μ and Z_μ fields — *i.e.* the photon and the Z boson — may be identified as:

$$m_\gamma = 0, \qquad m_Z = \frac{1}{2}v\sqrt{g_W^2 + g'^2}. \tag{2.51}$$

As required, the Higgs mechanism has endowed both the W and Z electroweak bosons with finite mass while recovering the massless nature of the photon and respecting the local gauge invariance of the SM Lagrangian.

The mass of the W and Z bosons is found to be related by the *electroweak mixing angle*, θ_W:

$$\frac{m_W}{m_Z} = \frac{g_W}{\sqrt{g_W^2 + g'^2}} = \cos\theta_W. \tag{2.52}$$

The relations introduced in this section allow the Higgs vacuum expectation energy to be determined experimentally. Electroweak measurements imply that:

$$v = 246\ \text{GeV}. \tag{2.53}$$

.

It is in this manner that the value of the electroweak scale previously stated is obtained.

Fermion Masses

In the Dirac Lagrangian, fermionic mass terms arise in the form of 'Dirac bilinears': $m\,\overline{\psi}\psi$. Such terms may be expanded in terms of spinor components, yielding terms of the form:

$$m\left(\overline{\psi}_{\mathsf{L}}\psi_{\mathsf{R}}+\overline{\psi}_{\mathsf{R}}\psi_{\mathsf{L}}\right). \tag{2.54}$$

Left-handed and right-handed chiral fields experience transformations differently under the $\mathsf{SU}(2)_{\mathsf{L}}$ symmetry group — for this reason, the coupling of left-handed and right-handed fields in the above term violate the gauge invariance of the Dirac Lagrangian.

The Higgs mechanism addresses this problem in the same manner by which the symmetrical issues with bosonic mass terms were resolved — by the introduction of the scalar, spin-0 Higgs field, ϕ. Restricting consideration of leptonic $\mathsf{SU}(2)_{\mathsf{L}}$ doublet terms, without loss of generality, to those which correspond to the electron, the relevant *Yukawa* term after spontaneous symmetry breaking is given by:

$$\mathcal{L}^{e}_{\mathsf{Yukawa}}=-\frac{1}{\sqrt{2}}\lambda_{e}\nu\overline{e}_{\mathsf{L}}e_{\mathsf{R}}+\mathsf{h.c.}, \tag{2.55}$$

where λ_e is the electron-Higgs coupling strength. From the electron Yukawa term, a mass term for the electron may be identified as:

$$m_{e}=\frac{1}{\sqrt{2}}\lambda_{e}\,\nu. \tag{2.56}$$

Returning to a broader field of view, this result may be generalised for all generations of fermion:

$$m_{\mathsf{f}}=\frac{1}{\sqrt{2}}\lambda_{f}\nu, \tag{2.57}$$

which also holds true for the masses of quarks.

Physics of the Higgs Boson

The existence of the Higgs boson gives rise to rich phenomena at high-energy hadron colliders. The phenomenology of the Higgs sector enabled the decay signatures of the Higgs boson to be understood, which paved the way for its discovery — solidifying the role of the Higgs mechanism in our understanding of nature. A thorough understanding of the production and decay modes of the Higgs boson will similarly allow many outstanding questions in high energy physics to be probed.

The production of Higgs bosons from pp-collisions can occur via a number of channels. The most dominant production mode at hadron colliders is that of gluon-gluon fusion (ggF), where partonic gluons within the proton interact (as permitted by the non-Abelian $\mathsf{SU}(3)$ symmetry of QCD) such that an energetic excitation of the Higgs field is produced. This is made possible through a virtual t-quark (or b-quark) loop. As the protons are accelerated to higher energies, more mass is accrued in the form of additional partons within the proton, resulting in a higher probability of a ggF interaction. This logic is reflected by the energy dependence of the cross-section.

The following most-commonly analysed mode of production is vector-boson fusion (VBF). In the case of VBF-production, two quarks produced in the underlying pp-collision initiate a scattering event, where the resultant W/Z bosons interact to produce a Higgs boson. While sub-dominant, this topology boasts the experimental advantage of providing an additional trigger handle through the hadronisation of the parent quarks.

Additional production modes in hadron collisions include top/anti-top quark fusion, $t\bar{t}$F, and the 'Higgs-Strahlung' process (from the German for 'radiation', hence the name of 'Bremsstrahlung' radiation). The four most dominant Higgs production processes are shown in Figure 2.8.

The mass of the Higgs boson is intrinsically connected to the vacuum expectation value of the Higgs field:

$$\nu^2 = -\frac{\mu^2}{\lambda}, \qquad m_H^2 = 2\lambda\nu^2. \tag{2.58}$$

It may therefore be said that the Higgs gives mass to itself. Upon its discovery, the mass of the Higgs boson was measured to be $\sim 125\ \mathsf{GeV}$.

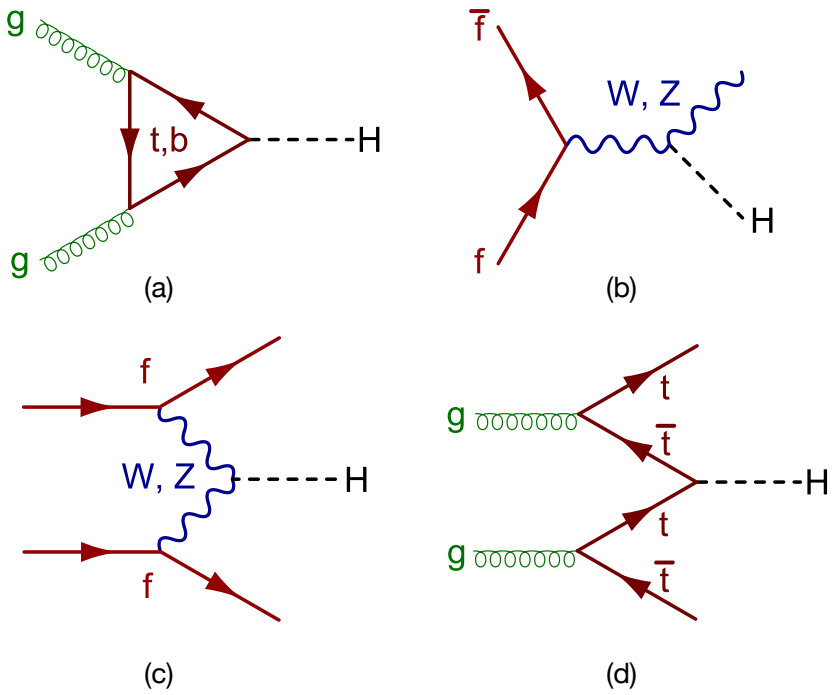

Figure 2.8 Feynman diagrams for the most prominent Higgs boson production modes for *pp*-collisions, shown to leading order. These are (a) gluon-gluon fusion, (b) Higgstrahlung, (c) vector-boson fusion, and (d) top-quark fusion. Source: [31].

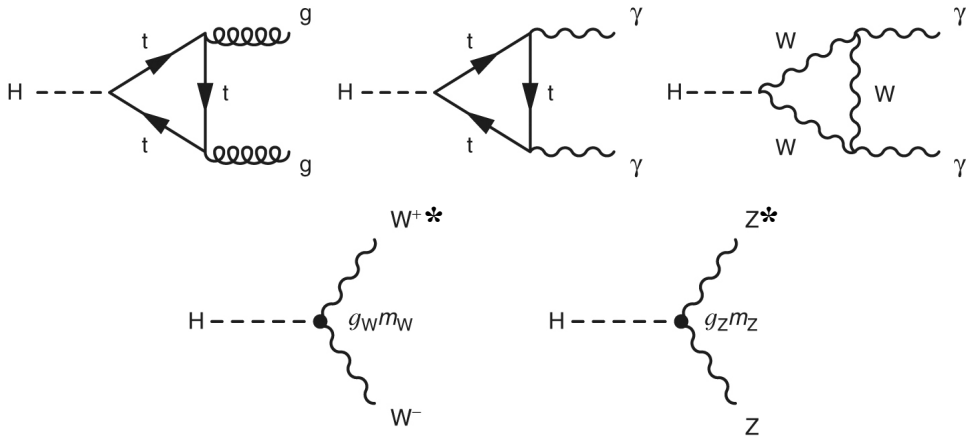

Figure 2.9 leading order Feynman diagrams for the bosonic decay of the Higgs boson. Quantum loops containing virtual particles allow the Higgs to decay to the gluon and photon, while the decay into WW^*-pairs and ZZ^*-pairs is permissible where one of the child bosons is off mass-shell. Source: [3].

By conservation, the Higgs boson may decay through any topology where $m_H > 2m_{\text{decay}}$ (to lowest order). Accordingly, the Higgs boson, H, may decay leptonically via $H \to e^+e^-$, $H \to \mu^+\mu^-$, and $H \to \tau^+\tau^-$, and to quarks via $H \to u\overline{u}$, $H \to d\overline{d}$, $H \to c\overline{c}$, $H \to s\overline{s}$, and $H \to b\overline{b}$. The magnitude of the couplings between the Higgs boson and the W, Z, and virtual t quarks are large. This allows the decays $H \to gg$ and $H \to \gamma\gamma$ to occur via virtual particles with significant branching fraction — as may be observed in Figure 2.10 and Figure 2.11.

Additionally, the Higgs boson may decay to pairs of W and Z bosons. As the mass of the Higgs boson is measured to be $m_H = 125$ GeV, this requires one of each bosonic pair to be off mass-sell (*i.e.* a virtual particle). Such decay topologies are shown in Figure 2.9, with their associated production and decay rates quantified in Figure 2.10.

2.6 Particle Phenomenology at Hadron Colliders

The particulate constituents of the SM and the interactions between them give rise to complex and evolving signatures at particle detectors, such as those situated at the LHC. In order to probe the fundamental physics which drives the evolution of such signatures form the initial pp-collision, the phenomenology of such signatures must first be examined. Once equipped with a clear understanding of the high-

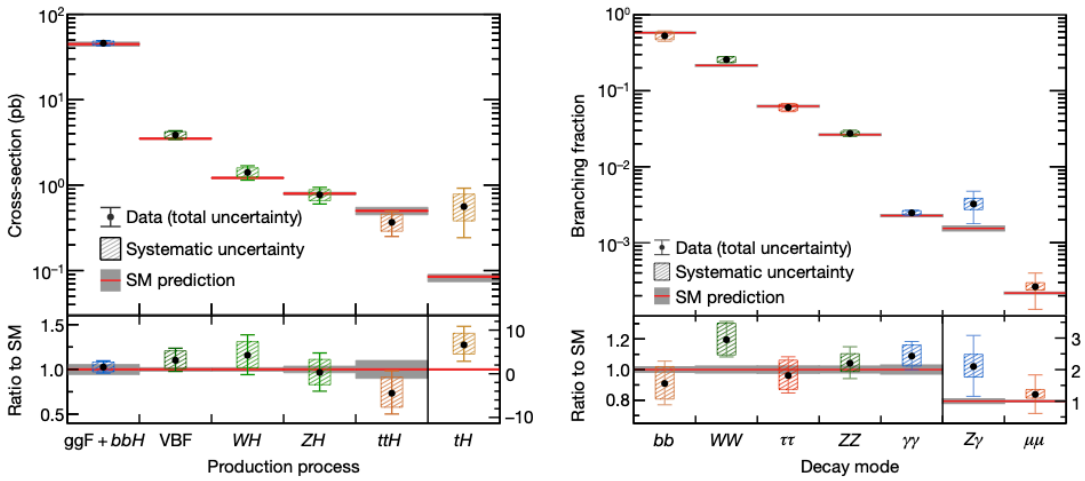

Figure 2.10 (*left*) SM predicted cross-sections for various Higgs boson production modes, and; (*right*) SM branching fraction for a range of Higgs boson decay topologies, compared to data collected by the ATLAS experiment in Run-2 of the LHC. Ratios of predicted to measured values are shown at the bottom of each plot. Source: [32].

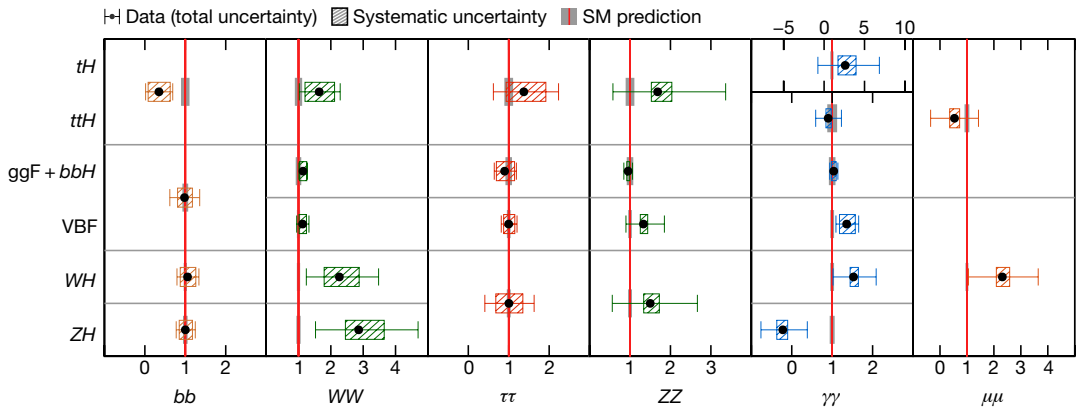

Figure 2.11 ratio of predicted event-rate to measured event-rate as observed by the ATLAS experiment during Run-2 of the LHC, shown for different combinations of production and decay topologies. Source: [32].

energy phenomena produced at hadron colliders, studies of the underlying physics may be performed.

The particles produced in high-energy pp-collisions at the LHC may be crudely delineated into a number of categories: stable particles (such as the electron); quasi-stable particles (such as the pion); short-lived particles (including the top quark, which decays before the QCD hadronisation scale), and; (inferred) missing particles (such as neutrinos). The research put forth by this thesis seeks to add the category of 'long-lived particles' to this list.

Particles which are not stable will decay into different species of particle until the final end-states of the decay-chain are stable particle states. Particles traverse the many layers of the detectors at the LHC, depositing their energy in the sensor systems housed therein — it is from such energy deposits that the species and kinematic properties of a given particle may be identified. From the picture which emerges as such decay processes unfold, it is possible to extract information concerning the fundamental laws by which such processes are governed. The physics presented in this section is of utility in connecting the phenomenology observed at the LHC with the fundamental laws of nature.

2.6.1 Parton Distributions

Protons are composite particles comprised of quarks and gluons. When protons are brought to collision, it is the constituent quarks and gluons, or 'partons', which interact, giving rise to the phenomena observed at hadron colliders. This is possible due to the weak strength of the confining force between the partons at high energies, as prescribed by the asymptotic freedom introduced in Section 2.3.2.

For a proton moving at relativistic energies, a quark or gluon emerging from a proton will posses some fraction, x_i, of the parent hadron's momentum, P_i, such that the momentum of each parton is given by: $p_i = x_i P_i$ for $0 \leqslant x_i \leqslant 1$. Here, x_i is said to be the *momentum fraction* of the proton carried by the parton. As each parton is a fundamental particle — unlike the parent proton — perturbative QCD may be employed to compute the cross-section for each respective parton.

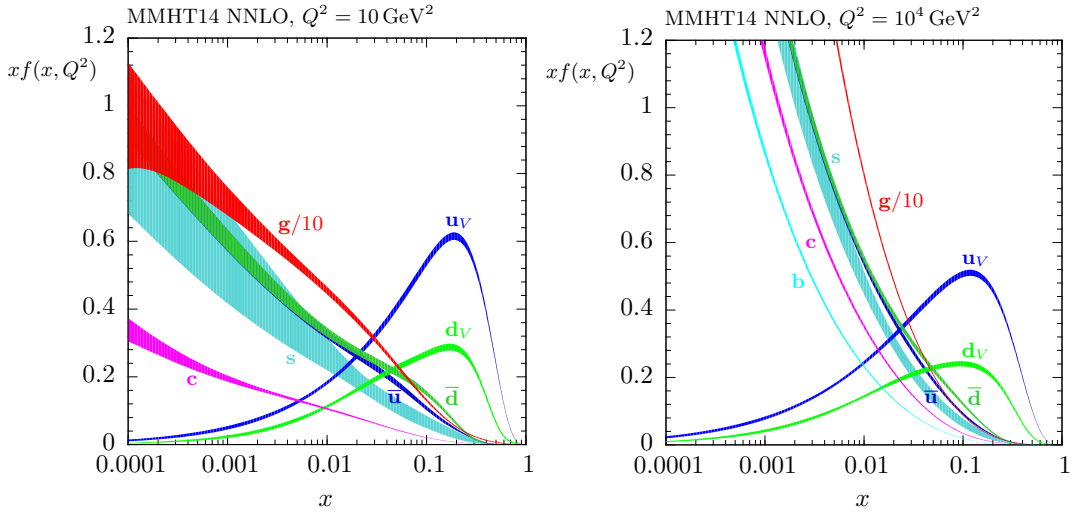

Figure 2.12 the MMHT14 PDFs [33] shown for momentum fraction, x, at the energy scale of (*left*) $Q^2 = 10$ GeV and (*right*) $Q^2 = 10^4$ GeV, calculated at next-to-next-to-leading order (NNLO).

From such components, the cross-section for the proton may be expressed as:

$$\sigma = \sum_{i,j \in \{q,\bar{q},g\}} \int_0^1 \mathrm{d}x_1 \int_0^1 \mathrm{d}x_2 \; f_i\left(x_1, \mu_\mathrm{F}\right) \; f_j\left(x_2, \mu_\mathrm{F}\right) \; \sigma_{ij}\left(x_1, x_2, \mu_\mathrm{F}\right), \tag{2.59}$$

where σ_{ij} is the cross-section for the colliding partons i and j, and μ_F sets the *factorisation scale* — to be discussed shortly. Functions $f_i(x_k)$ quantify the probability of encountering a parton of species i with momentum fraction x_k within an incoming proton, k — such functions are termed *parton distribution functions* (PDFs). Taking the sum of the relevant PDFs over all possible momentum fractions yields the total cross-section for the proton collision. Equation 2.59 is known as the *QCD factorisation theorem*, which may be used in order to generalise the computation of cross-sections for elementary particles, as introduced in Section 1.5, to those of composite hadrons such as the proton.

The PDFs introduced in Equation 2.59 cannot be computed directly via perturbation theory. Instead, they are measured experimentally. The distributions obtained from such measurements are used extensively in simulation in order to predict observable measurements at hadron colliders, and to aide the determination of the theoretical uncertainties associated with such predictions. Recent measurements of PDFs as obtained in [33] are shown in Figure 2.12.

While PDFs may not themselves be calculated from the first principles of QCD,

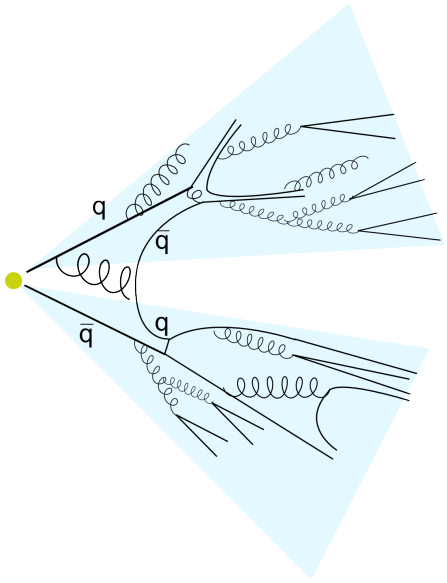

Figure 2.13 illustration of jet formation as a result of quark-induced hadronisation. Additional quark pairs are produced in conjunction with soft gluon radiation such that hadronic jets are observed. Image sourced from [34].

the manner in which they evolve with the *factorisation scale*, μ_F, is, in fact, calculable with perturbative techniques. The factorisation scale specifies the energy threshold above which radiative effects are included in the calculation of the hard scatter process. Radiative contributions which occur below this scale are incorporated into the parton distributions themselves. Having set the factorisation scale, the behaviour of PDFs may be obtained by:

$$\mu_F^2 \frac{\partial f_i\left(x_i, \mu_F^2\right)}{\partial \log_e\left(\mu_F^2\right)} = \sum_j \int_{x_i}^1 \frac{dz}{z} P_{ij}(z)\, f_i\left(\frac{x_i}{z},\, \mu_F^2\right). \tag{2.60}$$

Equations of this form are known as Dokshitzer-Gribov-Lipatov-Altarelli-Parisi (DGLAP) equations [9]. Here, z denotes the fraction of momentum carried by an outgoing quark after the incoming quark has radiated a gluon, for radiative processes which occur above μ_F, while P_{ij} describes the parton splitting (in the collinear limit) which emits the radiated gluon in question.

2.6.2 Hadronisation and Jets

In Section 2.3.3, the confinement of quarks within colour-neutral states was introduced. As quarks produced in collisions at the LHC separate, the

chromodynamic flux-tubes which extend between them become sufficiently energetic, such that new $q\overline{q}$ pairs are produced. This process unfolds iteratively.

At the LHC, this process of hadronisation leads to the observation of *jets*. Quarks and protons produced in the collisions are not observed as individual particles, but instead as hadronic cascades — showers of hadrons, as illustrated in Figure 2.13, which deposit their energy in the sub-systems of particle detectors. Reconstruction algorithms are tuned to identify such jets and, where possible, infer the flavour of the quark which initiated the cascade. The methodology with which this is achieved shall be the subject of Section 8.3.3.

The quark- and gluon-initiated jets which arise from QCD processes form a significant statistical background in the search for displaced τs presented in this thesis. The statistical methodologies by which background effects are estimated and suppressed are discussed in Chapter 11.

2.7 Experimental Coverage

It is elucidating to step back from the detailed workings of the SM and examine the broader landscape of SM measurements. Figure 2.14 depicts the broad consensus between theoretical prediction and experimental measurement with regards to the production cross-section of various high-energy processes. This serves to illustrate the predictive power of following a small number of core axiomatic principles — such as the preservation of symmetry — to their logical limit, and provides a reminder of the profound success of the scientific method.

This chapter has introduced the standard model of particle physics and summarised the theoretical concepts which have proven to yield great power in explaining, and predicting, the observations made at collider experiments. It has been shown that, by insisting upon the preservation of a small number of key concepts, a wealth of physical observations may be accounted for. It has further been shown that the fundamental physics underpinning the standard model may be encapsulated in a Lagrangian of the form $\mathcal{L}_{\mathsf{SM}} = \mathcal{L}_{\mathsf{QCD}} + \mathcal{L}_{\mathsf{EW}} + \mathcal{L}_{\mathsf{Higgs}} + \mathcal{L}_{\mathsf{Yukawa}}$, which yields, in full:

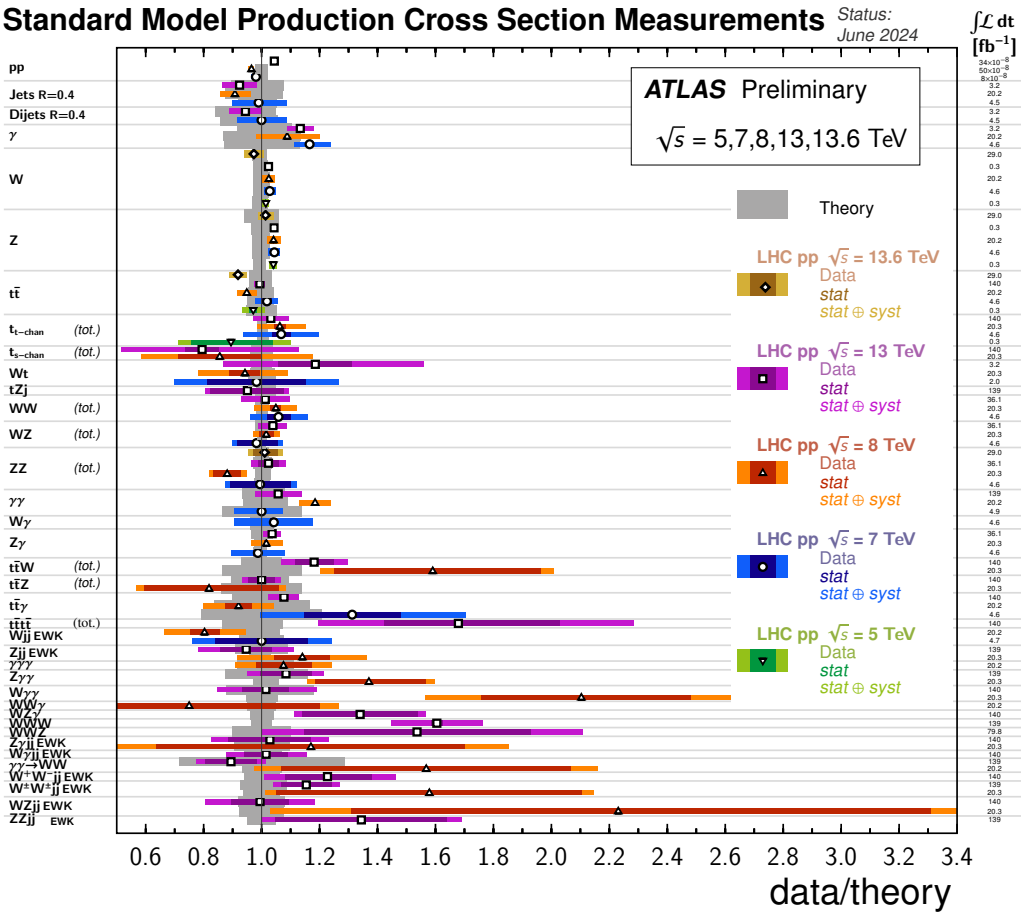

Figure 2.14 summary of ratios with respect to theoretical prediction for several measurements of total and fudicial cross-sections. Reported total cross-sections are corrected for their respective branching fractions. Figure extracted from [35].

$$
\begin{aligned}
\mathcal{L}_{\text{SM}} = & -\frac{1}{4} W^{j}_{\mu\nu} W^{j}_{\mu\nu} - \frac{1}{4} B_{\mu\nu} B^{\mu\nu} - \frac{1}{4} G^{a}_{\mu\nu} G^{\mu\nu}_{a} \\
& + \sum_{f} i \overline{\psi}_{f} \gamma^{\mu} D_{\mu} \psi_{f} \\
& + (D_{\mu}\phi)^{\dagger} (D^{\mu}\phi) - \mu^{2} (\phi^{\dagger}\phi) - \lambda (\phi^{\dagger}\phi)^{2} \\
& - \frac{1}{\sqrt{2}} \sum_{f} \overline{u}_{f,\text{L}} \lambda_{f} \nu u_{f,\text{R}} \\
& + \overline{\psi}_{\text{L}} \lambda \phi \psi_{\text{R}} + \text{h.c.}
\end{aligned}
\tag{2.61}
$$

This is a testament to the success of the scientific method, as pursued by many scientists over many decades, before arriving at the development of the standard model. While the surface of theoretical particle physics has in this chapter barely

been scratched, the concepts presented provide a suitable backdrop against which the research put forth in this thesis may be motivated.

The standard model stands as one of the most successful theories in scientific history. It is known, however, that it cannot be the ultimate explanation. Many natural phenomena cannot be accounted for within the framework of the standard model, which necessitate, at least, an extension thereof, or perhaps a new theoretical architecture through which to view the natural world. A light will be cast upon the short-comings of the standard model in the following chapter.

> *There are known knowns, and known unknowns...*
>
> — Donald Rumsfeld

3

Physics Beyond the Standard Model

Despite its remarkable success and predictive power, the Standard Model cannot be the final word in particle physics. Many observed physical phenomena stand at variance with SM predictions, and cannot be reconciled with our current understanding of the Universe within the theoretical apparatus that the SM provides. This chapter presents key instances where the SM is in tension with observation, in addition to issues of a more theoretical nature which engender a sense of philosophical discomfort. A set of possible mechanisms with which to extend the SM in an effort to address such short-comings is then summarised.

3.1 Overview

The success of the SM is beyond contestation. The values of physical parameters predicted by the SM agree with measurement to the greatest degree of precision in the history of scientific inquiry. Such success, however, is finite in scope.

Many physical phenomena have been observed which cannot be accounted for within the theoretical framework provided by the SM. Further, many such observations present direct tension with the predictions of the SM. While the SM boasts a profound command of the natural phenomena which lies within its purview, it is apparent that many physical phenomena fall outwith the scope of its explanatory power. It is therefore necessary to consider the SM either as an incomplete theory to which an extension may be attached, or as an approximation of a more general truth to be revealed at higher energies, analogous to classical physics.

This chapter surveys the shortcomings of the SM, together with the conceptual reasons to expect new physics beyond that harboured by the SM to exist. Thereafter, selected attempts to address such problems through novel extensions of the SM are presented.

3.2 The Need for New Physics

The quest to uncover new physics is rooted in a variety of SM shortcomings. The SM presents a number of conceptual issues, a sub-set of which — while not concretely in disagreement with observation — engender a sense of dissatisfaction of a more philosophical nature. Many issues, however, pertain to the empirical observation of phenomena which stand starkly at variance with the description of nature proffered by the SM. An overview of such theoretical shortfalls is now presented.

3.2.1 The Hierarchy Problem

In Section 2.3.2, it was shown that the process of renormalisation furnishes the bare quantities which feature in a given QFT with quantum corrections, resulting in a new effective value for the quantity in question. Such quantum corrections

include the dressed mass of the electron, and the anti-screened charge of quarks — leading to the asymptotically-free behaviour which is characteristic of QCD.

Typical quantum corrections of this nature are small in magnitude, with the resultant effective quantities similar in scale to those of their bare uncorrected counterparts. In some instances, however, the renormalisation process results in quantities which differ vastly in scale to their original values. This leads in some circumstances to the appearance of delicately balanced cancellations in terms.

The *hierarchy problem* [36] is said to be present in a theory which gives rise to quantities at seemingly unnaturally different scales from one-another. Where a natural explanation for the difference in scale cannot be provided, this is often taken as a suggestion that the theory in question is not yet the complete picture of reality.

An example of the hierarchy problem may be found in the mass of the Higgs boson. The SM regards the mass of the Higgs boson as a free parameter, and does not offer a precise prediction of the value it ought to assume. Similarly to other particles in the SM, the Higgs boson is subject to quantum corrections from virtual processes — corrections which influence its mass. In the case of the Higgs, such corrections are found to be very large: on the order of 10^{15} GeV [10]. It is logical to therefore expect the mass of the Higgs boson to similarly be on the order of 10^{15} GeV. Instead, it is measured to be ~ 125 GeV. This implies a very precise, delicate, and unnatural cancellation of virtual terms — a fine tuning of hierarchies [37].

Some have attempted to address the hierarchy problem by invoking what is known as the *anthropic principle* [38]. This is the principle that, however finely-balanced a set of parameters must be to give rise to the Universe as it is observed, this will *perforce* be the set of values measured by an observer — else, the observer would not exist to perform the measurement. Alternatively put, if a hierarchy of scales is required to produce the current state of the natural world, it follows that, however unlikely or unnatural such values may appear, those are the only values that could have been measured by an observer in the Universe to which such values gave rise. The appearance of fine tuning is removed by arguments of probabilistic necessity.

While the logic of the anthropic principle is undeniably without fault, it is not widely regarded as a satisfactory solution. Discarding the prospect of a technical explanation on philosophical grounds is considered by some as a dereliction of

scientific inquisition. It is desirable to instead uncover a scientifically tangible explanation for the hierarchy problem.

The philosophical distaste provoked by the apparent need for fine tuning in the Higgs mass may be evaded by a more aesthetically satisfying theory of new physics. The desire to obtain a more natural explanation for the mass of the Higgs is accordingly a strong motivation for physics beyond the SM.

3.2.2 Matter/Anti-Matter Asymmetry

The baryonic matter content of the Universe is observed to dominate over anti-matter, with vanishingly small proportions of anti-matter naturally present in the cosmos [39]. This poses a number of conceptual questions. At the time of the Big Bang, matter and anti-matter particles should have been produced in equal number. Had nature subscribed to this expectation and evolved accordingly, the matter and anti-matter produced in the early Universe would have mutually annihilated — leaving a Universe devoid of structure and comprised solely of radiation.

The survival of the matter observed today, of which we ourselves are comprised, imposes a number of restrictions on the dynamics of the early Universe. It was shown in 1967 by Andrei Sakharov that an excess of baryonic matter over anti-matter can only arise to the satisfaction of the following conditions: (1) violation of baryon number; (2) the violation of charge (C) and charge-parity (CP) symmetry, and; (3) the departure from thermal equilibrium — the so-called *Sakharov conditions* [40]. While all three conditions can be satisfied naturally by the SM, the degree of CP violation present is insufficient to account for the observed excess of ordinary matter. A further source of CP-violation beyond the SM is therefore required.

3.2.3 Neutrino Mass and Mixing

Neutrinos were first postulated by Wolfgang Pauli in 1930 as an attempt to resolve the apparent violation of energy conservation in the beta decay of radioactive nuclei. The term 'neutrino', meaning "little neutral one", first appeared in Fermi's theory [24] of the weak interaction; today, neutrinos are placed on an equal footing with other fermions in the SM.

The weakly-interacting nature of the neutrino renders their detection and study a profoundly difficult endeavour. The vast majority of neutrinos will refrain from interactions with the material of a given detector, depositing no energy from which to deduce their presence. Neutrinos were discovered in 1956 by Clyde Cowan and Frederick Reines [41], who used the large flux of anti-neutrinos emanating from the Savanna River nuclear reactor in South Carolina to facilitate their discovery. The project was named 'Poltergeist', in reference to the elusive nature of the neutrino, and resulted in the Nobel prize in physics being awarded to Reines in 1995 (by which time Cowan had deceased).

Subsequent experiments were devised to ascertain the properties of the neutrino. One such experiment — known as the 'Homestake' experiment [42] — was developed to detect the flux of neutrinos produced in the core of the Sun. It was found that the measured solar neutrino flux was smaller than that expected from the standard solar model by a factor of one third. This observation came to be known as the *solar neutrino problem* [43]. The explanation was found to lie in the quantum mechanical oscillation, or mixing, of neutrino flavours. Mixing between the three generations of neutrino accounts for the diminished number of solar electron-neutrinos observed: a proportion of electron-neutrinos will have assumed a different flavour by the point of detection.

The neutrino oscillation hypothesis was confirmed in 2003 by the Sudbury Neutrino Oscillation experiment (SNO) [44]. The SNO detector consisted of 9,600 photo-multiplier tubes immersed in 1,000 tonnes of Deuterium, achieving sensitivity to both charged-current and neutral-current nuclear processes. This enabled the detection of all neutrino flavours by the experiment, such that the solar electron-neutrino deficit may be compared to the total solar neutrino flux. The total neutrino flux was measured to be consistent with the predictions of the standard solar model, rendering the hypothesis of neutrino oscillations consistent with empirical observation. The director of the experiment, Art McDonald, was awarded the Nobel prize in physics in 2015.

Quantum mechanical mixing relates to the rotation of flavour states and mass states with respect to one-another. The discovery of flavour mixing in the neutrino sector therefore necessitates the non-zero mass of neutrinos. This is in direct contradiction with the SM, which dictates that neutrinos are massless fermions. While quark mixing in accordance with the CKM matrix can be accommodated by the SM on account of their QCD colour charge and double-handed chirality, neutrinos exist only in left-anded, colourless states. In the SM, the absence of a

right-handed neutrino state prevents the Lagrangian from containing a neutrino mass-generating term. The observation of flavour mixing in the neutrino sector is therefore a concrete observation of physics beyond the SM.

3.2.4 Dark Matter

Early observational evidence for the existence of dark matter was uncovered by Fritz Zwicky in 1933 [45]. Zwicky observed a larger spread in the velocities of galaxies in the Comma galaxy cluster than could be attributed to the effects of the luminous matter present in the cluster. Such an observation was an early indication of the presence of additional matter in the Universe which does not interact with light.

This observation was later corroborated by measurements of the rotational velocity of galaxies performed by Vera Rubin in 1970 [45]. The acceleration required to sustain a circular trajectory at constant velocity, v, is given in classical dynamics by $a = v^2/r$ at a given fixed radius of r. This may then be equated to the acceleration which is induced by a Newtonian gravitational potential at radius r, given for the total mass, $M(r)$, enclosed within r by $a = GM(r)/r^2$, where G is the Newtonian gravitational constant. It is then trivial to show that:

$$v = \sqrt{\frac{GM(r)}{r}}. \tag{3.1}$$

The distribution of visible matter within typical galaxies is sufficiently uniform as to result in values of $M(r)$ which are approximately constant, exhibiting little-to-no dependence on the radial distance from the galactic centre. It therefore follows from Equation 3.1 that the rotational velocity of galaxies should follow $v \propto r^{-1/2}$.

In disagreement with this expectation, the galactic rotation curves obtained through astronomical observation are observed to remain flat at increasing radial distance. This cannot be accommodated through variations in luminous matter content, such as stars, gas, and dust.

It is hypothesised that this observation is the result of galaxy halos, comprised of additional matter which does not interact with light. In the context of astronomical observation, such *dark matter* is only discernible by the gravitational

effects it exerts on the ordinary baryonic matter in its vicinity. In order to agree with the observed rotational velocity data, dark matter halos must extend to appreciably larger distances than the baryonic content of the host galaxy, distributing spherically through space in contrast to the planar disk assumed by ordinary galactic matter. This may be satisfied by a dark matter halo with a density profile given by:

$$\rho_{DM}(r) = \rho_0 \left[1 + \left(\frac{r}{r_c} \right)^2 \right]^{-1}, \tag{3.2}$$

where ρ_0 and r_c respectively denote the overall density scale and the characteristic radius of the halo. From such requirements it may be inferred that dark matter must interact gravitationally, while only interacting weakly with light and with itself.

An alternative hypothesis for the rotational behaviour of galaxies postulates the conjecture that gravitational dynamics behaves differently at large scales — known as *Modified Newtonian Dynamics* (MOND) [46]. This theory is largely disfavoured, however, largely as a function of its failure to account for myriad additional observations which lend support to the dark matter hypothesis. Further evidence to suggest the presence of dark matter in the Universe can be found in a number of different forms, a sub-set of which are now discussed.

Gravitational Lensing

Additional evidence for the presence of dark matter can be found in the *gravitational lensing* [47] which occurs in the vicinity of massive galaxy clusters, as observed by the Hubble Space Telescope.

Einstein's General Theory of Relativity dictates that sufficiently massive objects will warp the geometry of space-time to such an extent that matter and light in their presence will pursue a deflected path — in agreement with extensive astronomical observation. As is true for the optical analogue, the degree of distortion is proportional to the strength of the lens: in this case, the density of mass. This phenomena provides a powerful tool with which to probe the mass density of the observable Universe.

Galaxy clusters are observed to exhibit gravitational lensing. Light originating

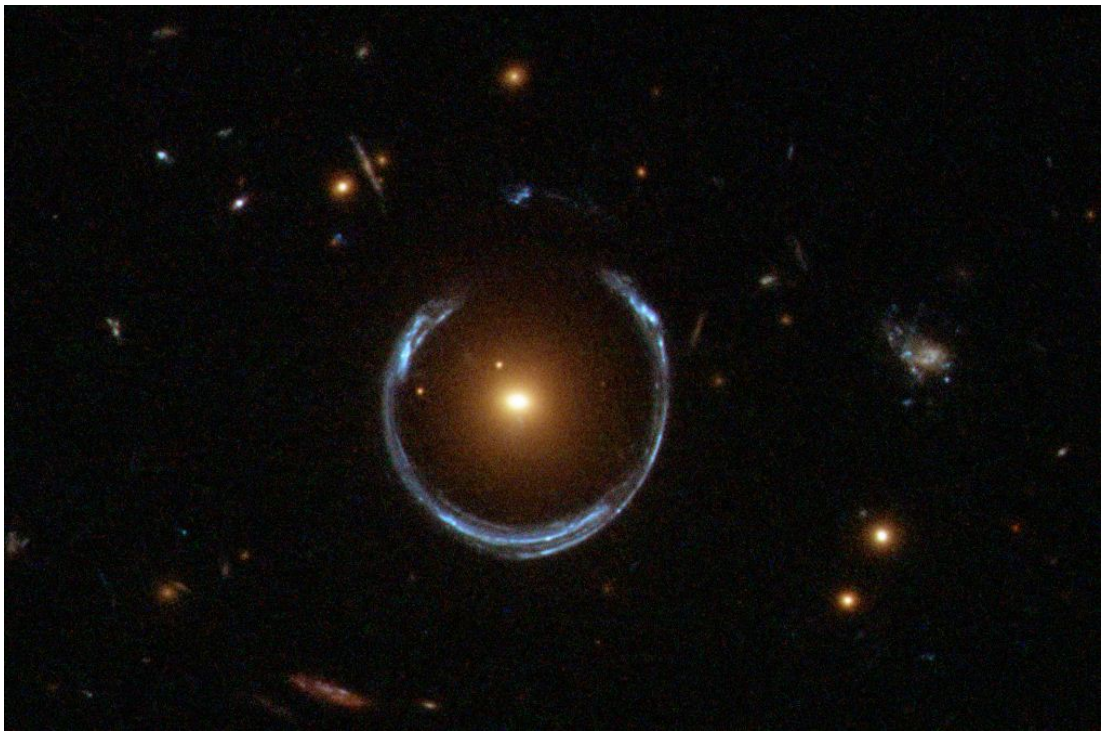

Figure 3.1 weak gravitational lensing induced by the mass present in a galaxy cluster, as captured by the Hubble Space Telescope.

from behind the cluster will experience maximal distortion around the greatest mass concentration — it would therefore be expected to observe distinct optical rings encircling each individual galaxy within the cluster. Contrastingly, massive galaxy clusters are found to exhibit a single large ring around the cluster as a whole, as is captured in Figure 3.1. This is suggestive of a more uniform mass distribution than that which would arise from a cluster whose total mass is chiefly attributable to the stellar mass of each constituent galaxy. The uniformity of gravitational lensing around galaxy clusters therefore indicates the presence of intermediate massive matter within clusters which does not interact with light: dark matter.

Cosmic Microwave Background

Approximately 380,000 Earth years after the Big Bang, the plasma of the early Universe became sufficiently cool — due to expansion — for electrons and nuclear ions to combine into atomic formations. From this point onwards, the Universe became transparent to radiation. Light from before this period in cosmic time cannot be observed; the light which radiates from this*recombination* epoch is therefore the earliest observable light in the Universe, and is known as the *cosmic*

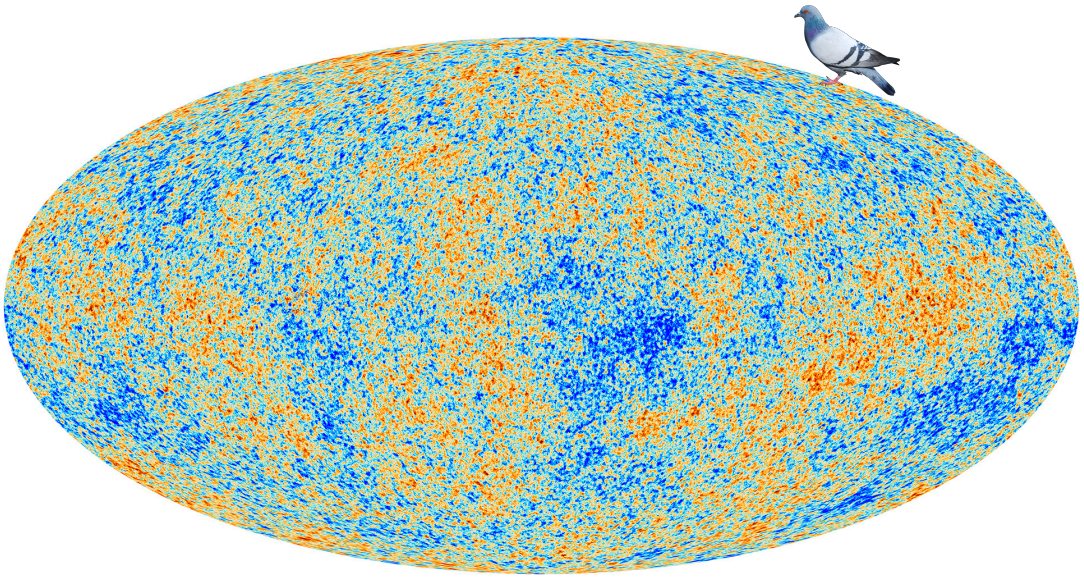

Figure 3.2 cosmic microwave background (CMB) as measured by the Planck Space Telescope. Variations in the heat-map depict temperature anisotropies in the CMB, arising from density fluctuations in the matter present at the time of recombination. The power sepctra of such temperature anisotropies can be used to deduce the energy density of dark matter in the early Universe. Image modified from [48].

microwave background (CMB).

The CMB was discovered accidentally in 1964 by Robert Wilson and Arno Penzias as they sought an explanation for the background noise which plagued the Holmdel Horn Antenna [49]. The background was initially attributed to the presence of pigeons nesting in the apparatus. Analysis of CMB data now constitutes a key pilar of cosmological research, from which anisotropies in cosmic temperature provide information regarding the seeds of cosmic structure formation, and illuminate many other outstanding questions regarding the early Universe.

The power spectrum of CMB temperature anisotropies is sensitive to the energy density of dark matter during the recombination epoch, and can thus be used to derive the abundance of dark matter in the early Universe — which is expected to remain constant. In this manner, the cosmic echo of the CMB allows the relic abundance of dark matter to be measured as $\Omega_{\mathrm{DM}} = \rho_{\mathrm{DM}}/\rho_{\mathrm{critical}} = 0.261 \pm 0.004$, where ρ_{DM} is the dark matter energy density and ρ_{critical} is the total critical density required to obtain a geometrically flat Universe, consistent with observation. This compares to a measured baryonic energy density of approximately 5%. This implies that dark matter comprises approximately 26% of the matter in

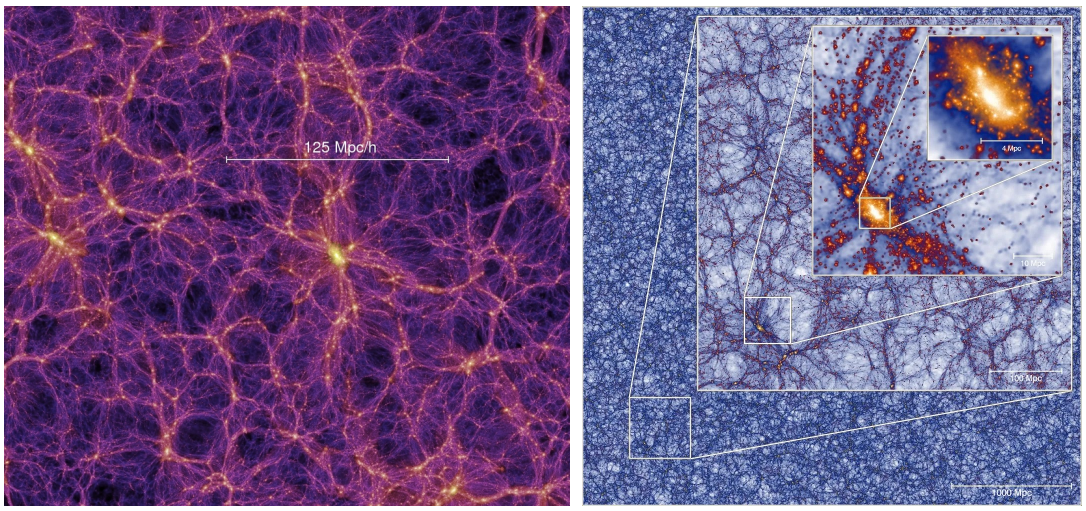

Figure 3.3 distribution of dark matter filaments shown for different scales of cosmological distance, as numerically simulated by the Millennium Simulation Project [50]. The inclusion of dark matter is required to obtain a result consistent with observation.

the Universe, with ordinary baryonic matter consituting a mere 5% of the total matter content of the Universe. The remaining difference is attributed to dark energy, which is investigated overleaf.

Cosmological Simulation

Large hydrodynamical *N*-body simulations are used in cosmology to obtain the predicted evolution of matter distributions and structure formation in the Universe. Simulated distributions may be compared to direct observation, such that the values of the simulated input parameters may be tested.

As a consequence of the observational and numerical evidence presented in this section, the field of physics has largely converged upon the consensus that an as-yet unknown form of massive, weakly-interacting matter pervades the Universe. Dark matter need not necessarily assume a particulate nature — alternative models, including primordial black holes [51], have been proposed. It is expected, however, that dark matter is comprised of weakly-interacting massive particles. As no such particle within the SM can serve as such a candidate, a new particle beyond the SM is required.

Many models of new physics have been proposed as candidate explanations of the dark matter problem. Candidate dark matter particles must satisfy the

requirements listed herein, while also boasting a cross-section which can account for the dark matter abundance in the early Universe. The ATLAS collaboration performs extensive searches for the production of viable dark matter candidate particles in the proton collisions which occur at the LHC.

3.2.5 Dark Energy

The expansion of the Universe was first measured by Edwin Hubble in 1929 [52]. By measuring the changing redshift of their light, Hubble deduced that galaxies were receding from the Earth's view at a rate proportional to their distance — in so doing, uncovering the expansion of the Universe itself.

The rate of expansion was widely expected to decrease under the gravitational attraction of matter in the Universe. To ubiquitous surprise, physicists were disabused of this consensus in 1998 with the discovery that the expansion of the Universe is *accelerating*. This revelation arose from the study of distant supernovae redshift by two independent groups, earning each group leader the Nobel prize in physics in 2011.

The accelerated expansion of the Universe requires an energetic source. Studies of the CMB suggest that such energy, known as *dark energy*, constitutes approximately 70% of the content of the Universe. The nature of dark energy is unknown; no fundamental field within the SM satisfies the properties required. The observed existence of dark energy therefore implies the existence of as-yet unknown physics, awaiting discovery.

In the formulation of his general theory of relativity, Einstein introduced a term known as the *cosmological constant* such that the field equations were tuned to describe a static Universe, as was thought to be accurate at the time. In light of Hubble's discovery of an expanding Universe, Einstein removed the cosmological constant, later referring to it as his "biggest blunder".

With the more recent discovery that the expansion of the Universe is accelerating, however, the cosmological constant has assumed a new significance. When instantiated to a positive value, the cosmological constant provides a strikingly accurate description of the accelerated expansion of the Universe. The cosmological constant has the effect of ascribing a non-zero energy to the vacuum of space-time itself. This concept is familiar to modern physics, and may be mapped

to the quantum mechanical fluctuation of quantum fields in their ground state — the so-called 'zero point energy' of the Universe, as has been discussed in the preceeding chapters. The energy value of the vacuum according to the quantum field-theoretic framework, however, exceeds that which is required to account for the cosmological constant by many orders of magnitude — a problem known as the 'vacuum catastrophe' [53]. While it appears that the source of dark energy is intimately connected to the vacuum energy of the Universe, an account of this energy is not provided by present theoretical thinking. An explanation for dark energy will perhaps require new physics at the fundamental level, deep within the QFT rafters of the SM.

3.2.6 Additional Puzzles in Physics

To exhaustively survey the outstanding problems in the field of physics would perhaps be worthy of a thesis in itself. The current best theoretical explanations of the Universe are peppered with holes; problems afforded extended consideration in this chapter are but a selected sub-set. As is true of the problems afore-discussed in this chapter, those not examined herein come in varieties of both an aesthetic and a substantive nature.

Problems not discussed at length in this chapter include the origins of cosmological inflation [47], an explanation of the strong-CP problem [54], and the question of why the SM should be composed of three flavour generations [55]. Perhaps the most fundamental of all remaining questions pertains to the quantum mechanical nature of gravitation. If Einstein's general theory of relativity is to be regarded as a semi-classical approximation of the true explanation of gravity, the more fundamental quantum mechanical explanatory framework for gravity must be uncovered. The reconciliation of general relativity with quantum mechanics has proven to be an extremely difficult task [56].

3.3 Models of New Physics

The motivation to uncover and understand physics beyond that which is contained within the SM is well-founded. From the existence of dark matter to the asymmetry of matter and anti-matter in the observable Universe, the necessity for a new or extended theoretical view of the natural world is clear.

To the theoretical shortcomings of the SM, a vast array of candidate explanations have been proposed. The nature which new physics may assume has been modelled in many powerful and elegant theories. A selected set of well-motivated models of physics beyond the SM is now presented.

3.3.1 Super-Symmetry

Symmetries of various forms are characteristic of the most successful fruits of twentieth century physics. From the Lorentz symmetry which underpins the relativistic treatment of the Universe, to the gauge symmetries which naturally gave rise to the prediction of many SM particles, the mathematical presence of symmetry in a given theory is a powerful watermark of its potential to capture the behaviour of the natural world.

In 1971, the ability to construct a new symmetry within QFTs was discovered independently by Yuri Gol'fand and Evgeny Likhtman [57], Andre Neveu and John Schwarz [58], and Pierre Ramond [59, 60]. The symmetry in question relates the particles of a given QFT with differing spin to one-another, and came to be known as *supersymmetry* (SUSY). SUSY currents transform particles of integer spin into particles of half-integer spin, and *vice-versa*. Given that integer spin states describe bosons which are symmetric under interchange, and that half-integer states describe fermions which are anti-symmetric under interchange, this new symmetry implies the existence of a partner particle for each particle which respects the symmetry.

It is for this reason that in a supersymmetric QFT, every fermionic (bosonic) particle is accompanied by a bosonic (fermionic) partner particle — such particles are known as *super-particles*, or *sparticles*. Sparticles share the same quantum numbers as their ordinary counterparts, while differing in the value of their spin. An additional quantum number by the name of R*-parity* is introduced in order to secure the prediction of a stable proton. This is defined as $R = (-1)^{2j+3B+L}$ for the spin, j, baryon number, B, and lepton number, L, of each particle in the supersymmetric theory. SM particles are ascribed an R-parity value of $R = 1$, while their super-partners are described by $R = -1$. Conventionally, the super-partners of fermions share the name of the regular fermion with a prefixed 's' (for instance, the supersymmetric electron is known as the 'selectron'), while the super-partners of bosons take the name formed by appending '*ino*' to the end of the name associated with the regular boson (such that the supersymmetric

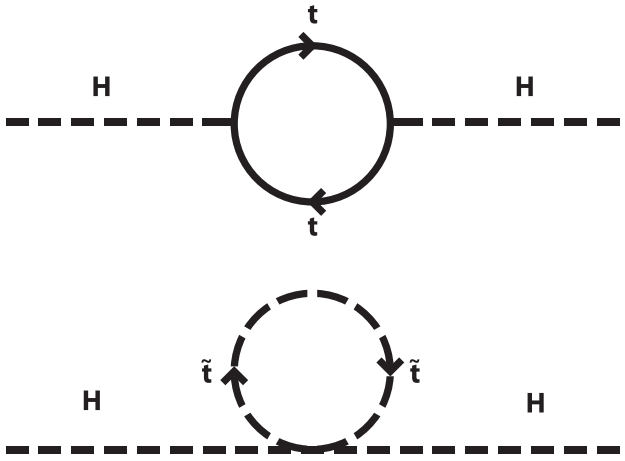

Figure 3.4 virtual loop contributions to the mass of the Higgs boson from (*top*) the SM top quark and (*bottom*) the supersymmetric stop quark. Pairwise cancellation of loop contributions is achieved on the basis of the equal mass and opposite spin statistics of each SM particle and their associated super-partner. SUSY thereby provides a natural solution to the hierarchy problem.

Higgs becomes the 'Higgsino'). The prediction that the SM — if formulated as a supersymmetric theory — should be accompanied by twice as many particles as are currently known presents an excellent, if daunting, signature with which to test its viability.

Of the many features to emerge from a supersymmetric formulation of the SM, the most important concerns the natural cancellation of higher-order quantum corrections. It was shown in Section 3.2.1 that the contributions to the mass of the Higgs boson which arise from virtual quantum loop processes lead to a predicted mass many orders of magnitude greater than that of the empirically measured value. The inclusion of sparticle terms in such loops achieves the precise cancellation of this contribution on account of the differing spin statistics and otherwise identical quantum numbers of particles and their super-partners — a process captured by the Feynman diagram shown in Figure 3.4. Cancellation of virtual loop terms in this pairwise manner results in no virtual contribution to the mass of the Higgs boson. The hierarchy problem is thereby resolved in a natural fashion, evading the philosophically uncomfortable fine tuning of scales.

The lack of observational evidence for sparticles at energies currently probed by colliders implies that, if it is indeed realised in nature, SUSY is a broken symmetry. Precedent for broken symmetries may be found in various quarters of theoretical physics, including the electroweak sector of the SM — the breaking of

SUSY in nature at low energies does not, therefore, diminish the viability of the theory.

The cancellation of virtual loop processes becomes partially truncated in the case of broken supersymmetric theories. If a broken supersymmetric extension of the SM is realised in nature, particles and their super-partners are free to differ in mass. The cancellation of quantum loops which arise from a broken SUSY may nonetheless result in the correct value of mass for the Higgs boson, provided that each sparticle is no heaver than ~ 1 TeV.

Further interrogation of the properties manifested by supersymmetric theories reveals additional attractive features. It can be shown that imposing the mandate of R-parity conservation requires that the lightest super-particle of a supersymmetric theory is stable. In the high-temperature conditions of the early Universe, the production of such particles ought to have occurred in abundance. On account of their stable nature, such super-particles would persist to the present day. Moreover, if comprised of the Higgsino, Wino, or Bino, the lightest stable SUSY sparticle would be weakly-interacting in nature. Solely through arguments of conservation, SUSY has provided the ideal dark matter particle candidate. If the mass of such a particle is taken to be on the order of the weak scale — as is required of a broken SUSY if it is to solve the hierarchy problem — its thermal relic abundance is calculated to agree with the dark matter density obtained from measurements of the CMB. This is known as the 'Weakly-Interacting Massive Particle (WIMP) Miracle' [61].

The desire to unify the seemingly discrete forces of nature into a single, more fundamental description of reality has long been a pursuit of the theoretical physicist. Among the earliest such examples is that of electricity, magnetism, and light — revealed by Maxwell in 1865 as manifestations of the single, more fundamental force of electromagnetism. This logic was extended by the unification of the quantum field theoretic description of electromagnetism with the weak force, yielding the electroweak cornerstone of the SM.

SUSY represents a further advance in the unification of the fundamental forces. Attempts to unify the component forces of the SM into a *grand unified theory* encounter an impasse: the gauge couplings of the SM, when projected forward in energy, fail to meet at a single point. It is accordingly impossible to describe the forces of the SM as originating from a single grand unified theory at a higher scale. Strikingly, this impasse is overcome by the inclusion of SUSY sparticles,

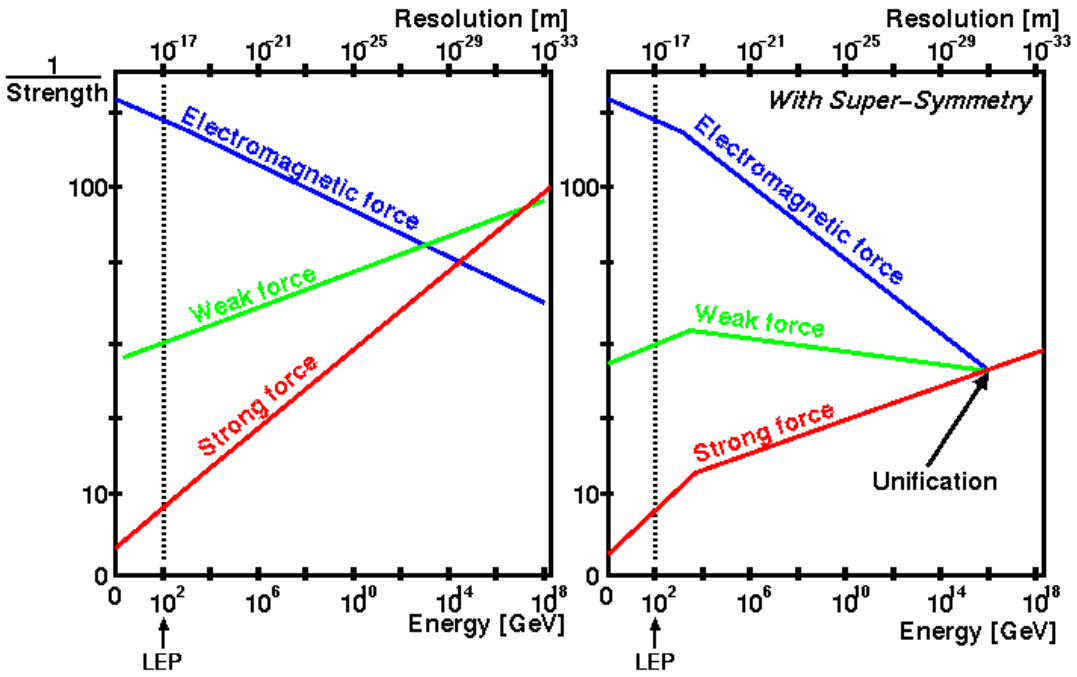

Figure 3.5 running of gauge coupling strengths as a function of energy, shown with constraints placed by data recorded at the Large Electron Positron collider (LEP). SM forces fail to converge to a single point at high energy scales. This is rectified by the inclusion of supersymmetric terms (*right*), which render a grand unified theory feasible. Original source unknown.

via virtual loops, in the determination of the fundamental force gauge couplings. With the inclusion of super-partner terms, supersymmetric formulations of the SM provide a natural unification of the fundamental forces into a single grand unified theory at higher energy scales — as is shown in Figure 3.5. The sense of philosophical satisfaction attached to this feature renders SUSY a yet more appealing theory of physics beyond the SM.

The remarkable elegance and unforced ease with which SUSY addresses an eclectic range of SM shortcomings earned it prominence among the most promising theories of new physics. With the prediction of many new particles states within the energetic reach of current collider technologies, the expectation of uncovering evidence in support of SUSY in the early stages of the LHC programme was widely held. Such discoveries, however, were not forthcoming. Extensive searches for evidence of various supersymmetric models have been undertaken to date, of which none have borne fruit. Current SUSY exclusion limits are shown in Figure 3.6. While by no means excluded, the phase-space in which the most attractive SUSY models may reside is fast diminishing.

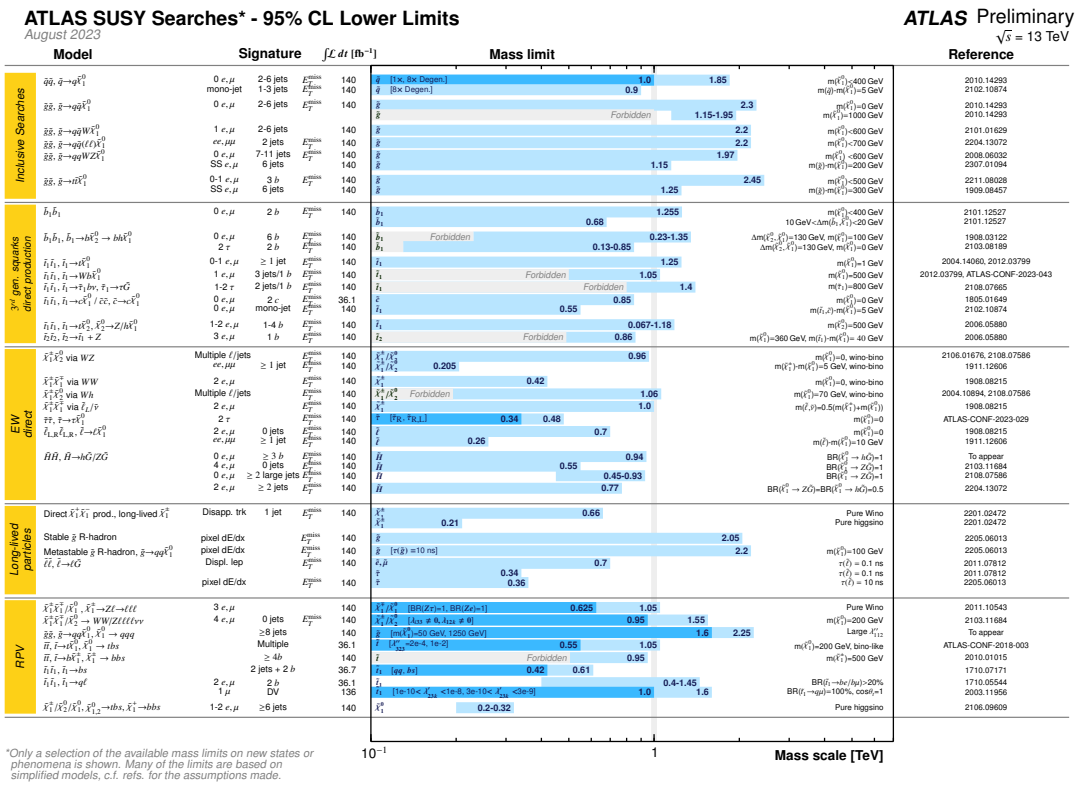

Figure 3.6 current 95% confidence-level (CL) exclusion limits on the possible masses of predicted SUSY states, as recorded by the ATLAS experiment. While many of the simplest incarnations of the supersymmetric SM have been excluded by LHC data, a significant space of well-motivated SUSY models remain viable.

Experimental efforts to expand the expanse of well-motivated parameter space to which particle searches are sensitive is therefore a priority for the field.

3.3.2 Heavy Neutral Leptons

As described in Section 2.4, the SM is home to neutrinos of only left-handed chirality. The omission of right-handed neutrino states in the construction of the SM was borne of two chief considerations: that left-handed neutrinos are naturally massless, and that the exclusion of right-handed neutrinos leads to the conservation of lepton flavour. It is now known that neutrinos possess non-zero mass, and that lepton flavour states may mix analogously to the quark sector. Had such observational evidence been available at the time of its construction, it is likely that chiral symmetry of the SM would not have been manually removed by the exclusion of right-handed neutrinos.

The inclusion of right-handed neutrinos in the SM gives rise to a number of consequences. If the lack of direct observational evidence for their existence is to be addressed, any right-handed neutrino state inserted into the SM must be uncharged with respect to each fundamental gauge symmetry. Unlike their left-handed counterparts, such states do not, therefore, experience the electroweak force — consequentially, right-handed neutrinos are said to be 'sterile'. The inclusion of sterile neutrinos with right-handed chirality within the SM enables neutrinos to acquire a Dirac mass term within the SM Lagrangian. The observation of massive neutrinos, and by extension their flavour oscillations, is thereby explained. It is conventional in the field of high energy physics to refer to massive, right-handed sterile neutrinos as 'heavy neutral leptons' (HNLs) [62].

There remains, however, an outstanding question: given upper-bound constrains on the mass of SM neutrinos on the order of $1\ \text{eV}$, why should their mass be smaller than those of other SM particles to such a significant extent? If neutrinos are to acquire their mass via the Higgs mechanism, mass parameters which differ so vastly from other particles whose mass is acquired by the same phenomenon of electroweak symmetry breaking appear unnatural.

An elegant answer to this question may be found in the *see-saw* mechanism [63]. Owing to their uncharged nature, HNLs are free to acquire mass in an alternative manner. A gauge-invariant *Majorana* mass term arises in the extended Lagrangian for uncharged particles which does *not* depend on the vacuum expectation value, and is therefore independent of the Higgs mechanism. It may be shown that, for 'type-I' see-saw mechanisms [64], the physical mass of the neutrino is dependent upon the ratio of the square of its familiar Dirac mass term to its afore-introduced Majorana mass term [63]. The physical mass of the neutrino is then free to assume a small value, while its associated Dirac mass term remains comparable in magnitude to those of other fermions in the SM.

Extending the neutrino sector of the SM with HNL states provides a compelling solution to the incompatibility of the SM with the experimental discovery of neutrino masses. Moreover, that the solution is achieved through the removal of restrictions manually placed upon the natural chiral symmetry of the SM, rather than by the conjecture of new fundamental fields, enhances its aesthetic appeal. It is further postulated in some physics models that, should it prove to be stable, a given HNL generation may serve as an appealing candidate for dark matter [65].

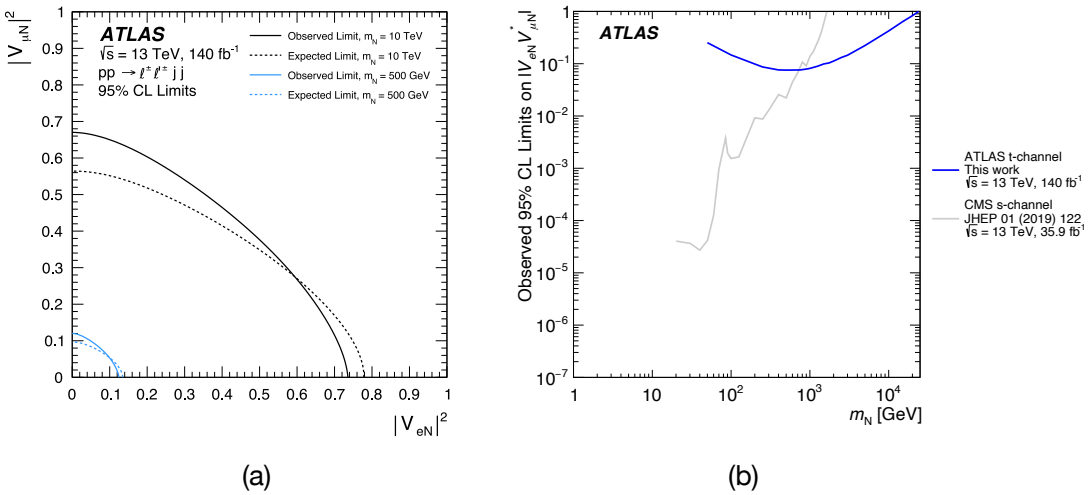

Figure 3.7 95% confidence level (CL) limits on HNL coupling strength to SM electrons ($|V_{eN}|$) and SM muons ($|V_{\mu N}|$) as observed by *(a)* ATLAS and *(b)* ATLAS & CMS search campaigns [66].

Searches for HNL signatures have been performed with the ATLAS experiment, leveraging the predicted HNL mixing with SM constituents. W^+W^+ scattering events, where the production of charged (e or μ) leptons may be mediated by heavy neutral leptons, are scrutinised for evidence of coupling between SM leptons and HNL states. To date, no evidence to support the existence of HNLs has been observed — as is depicted in Figure 3.7.

3.3.3 Extended Electroweak Gauge Sector

A common prediction is made across a number of different physics models. Many models which extend the SM, including Little Higgs models [67], technicolour models [68], composite Higgs models [69], models of extra dimensions [70], and a number of grand unified theories [71], admit a consistent prediction: the existence of additional heavy gauge bosons. While the theoretical machinery with which such models address the problems associated with the SM differs, the predicted existence of new heavy gauge bosons arises from each.

The new particles predicted by such theories are proposed to share the couplings and decay topologies of the SM electroweak bosons, and are accordingly labelled the W' and Z' bosons. The expected mass of the new gauge bosons varies with each model of new physics; it is broadly expected, however, that the mass of the W' and Z' may extend into the TeV regime. It is therefore possible to test

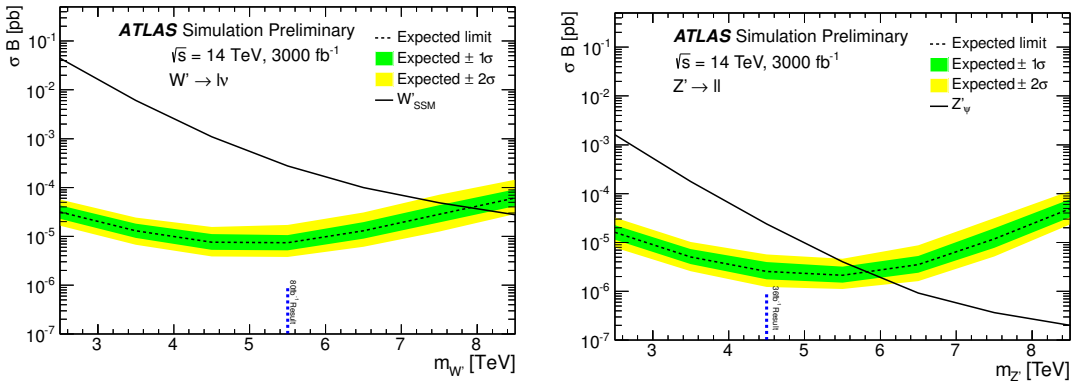

Figure 3.8 simulated cross-section exclusions as a function of mass for the extended (*left*) W' and (*right*) Z' gauge bosons at 95% confidence-level, predicted for future measurements for the ATLAS experiment at a centre-of-mass energy of $\sqrt{s} = 14$ TeV. Decay topologies $W' \to l\nu$ and $Z' \to ll$ are considered, where $l \in \{e, \mu\}$. Error bands at the 1σ and 2σ level are also shown [72].

the prediction of heavy gauge bosons with the existing suite of collider-based experiments.

Projected search sensitivity will place constraints on the permissible mass of heavy gauge bosons, as is shown in Figure 3.8. Additionally, while not excluded from the realms of possibility, existing coverage suggests that the W' and Z' bosons, if realised in nature, are likely to be restricted in the phase-space they may occupy.

Despite its remarkably successful ability to describe the natural world, the success of the SM is limited. Observations spanning across the cosmological and sub-atomic scales dictate that the SM cannot be the final word in high energy physics. New physics beyond the SM is required.

To this end, a plethora of theories have been proposed. Numerous elegant and compelling solutions have been conjectured to address the shortcomings of the SM in ingenious and elegant ways. To date, however, no model of new physics has emerged victorious from experimental assault. Nature is yet to reveal her secrets.

The absence of a clear signature of new physics obscures the path ahead; as the space of experimentally excluded ideas expands, it is unclear which direction will yield the fruits of discovery. In the endeavour to overcome the drought

of discovery, this thesis attempts to advance the exploration of a new frontier: particle lifetime.

> *There are also unknown unknowns.*
>
> — Donald Rumsfeld

4

Long-Lived Particles

As the experimental quest to uncover the nature of physics beyond the SM progresses, the theories proposed to address SM shortcomings increasingly appear to stand at variance with reality. The likelihood that the ideas underpinning current BSM models are aligned with those of nature is severely diminished. Long-lived particles (LLPs) beyond the SM offer a possible remedy to this situation. In concert with strong experimental motivation, many models of new physics appear to naturally give rise to particles with appreciable lifetime without manipulation. This chapter introduces the experimental and theoretical foundations upon which the hypothesis of LLPs is constructed, before discussing the privileged role that SM τ-leptons may be expected to play in their detection.

4.1 The Life-Time Frontier

The necessity for new physics beyond that formulated in the SM has been exposed to a level beyond dispute. Despite extensive, sweeping experimental campaigns aimed at uncovering its nature, no observation of a physical state beyond the SM has been made to date. Presented with such an impasse, it is natural to question the presumptions previously made on the form new physics may assume, and to re-assess the experimental techniques devised to uncover the finger-prints which may be deposited by the production of BSM resonances.

Novel models of new physics, in addition to unconventional topological signatures which break with conventional expectation, are therefore an area of active interest. Proposals of this kind account for the lack of experimental observation of BSM resonances by introducing new physics within new regions of phase-space, charting new territory and advancing the frontiers of parameter-space. One such frontier is that of particle *lifetime*.

The lifetime of a given particle species is defined as the mean duration of time elapsed before it decays into a different, lighter species. This was shown in Chapter 1 to arise in a quantum field theory as:

$$\tau^{-1} \equiv \Gamma = \frac{1}{2m_{\mathsf{x}}} \int \mathrm{d}\Pi_{\mathsf{f}} \, |\mathcal{M}\left(m_{\mathsf{x}} \to \{p_{\mathsf{f}}\}\right)|^2,$$

for a particle x with mass m_{x} decaying into the Lorentz-invariant phase-space $\mathrm{d}\Pi_{\mathsf{f}}$ associated with final states $\{p_{\mathsf{f}}\}$, as was introduced in Section 1.5. This may be viewed relativistically as the *proper* lifetime.

Existing search coverage of BSM scenarios is optimised for the discovery of particles which are expected to be of short lifetime. Such particles decay quickly into SM particles, whose final states may be measured — particles of this description are termed *promptly-decaying* particles. Promptly-decaying, or short-lived, particles are not, however, the only possible incarnation of BSM states.

Rather, many logical avenues of thought cast favour on the existence of new, *long-lived* particles beyond the SM. Such particles, however copiously produced in collider experiments, would evade efficient detection under experimental programmes optimised for the detection of promptly-decaying particles. The motivation to support the conjecture of new, long-lived states assumes many

forms, of which a sub-set is now presented.

4.2 Experimental Motivation

The absence of a concrete resonance beyond the SM in conventional topologies and parameter-space lends credence to the case for new physics manifesting itself in a hitherto ill-investigated signature. As the space of experimentally-excluded models corresponding to conventional signatures expands, so too does the imperative to explore more exotic and unusual possibilities. Promptly-decaying particles beyond the SM are increasingly disfavoured by measurements performed at the current energy frontier, rendering the investigation of novel BSM modalities more attractive. The phase-space associated with LLPs is largely unexplored, harbouring many mass and lifetime points where new physics may reside hidden from our experimental gaze.

High energy physics experiments, such as the ATLAS experiment, are largely optimised for the detection of promptly-decaying particles. Nominal analytical procedures, reconstruction techniques, and trigger algorithms are typically designed to efficiently capture the production of particles with short lifetimes — a consequence of their primary scientific design-objective: discovering the promptly-decaying Higgs boson. It is therefore possible that LLPs are produced in pp-collisions in the current energy regime, yet evading efficient detection due to the design choices which underpin current experimental methods. LLPs would traverse a more significant distance within the ATLAS detector than SM constituents, before decaying into visible end-states in an atypical region of the structure — accounting for the current absence of new resonances detected in more central regions of the detector.

4.3 Theoretical Motivation

While many decay promptly, the particles of the SM are characterised by a wide range of lifetimes. This lifetime axis may broadly be separated into three categories. Particles with a proper lifetime within $c\tau < 10\ \mu m$ may be labelled 'detector prompt', and comprise the majority of particles which feature in collider searches. At the opposite extremity of the axis, particles whose lifetime exceeds

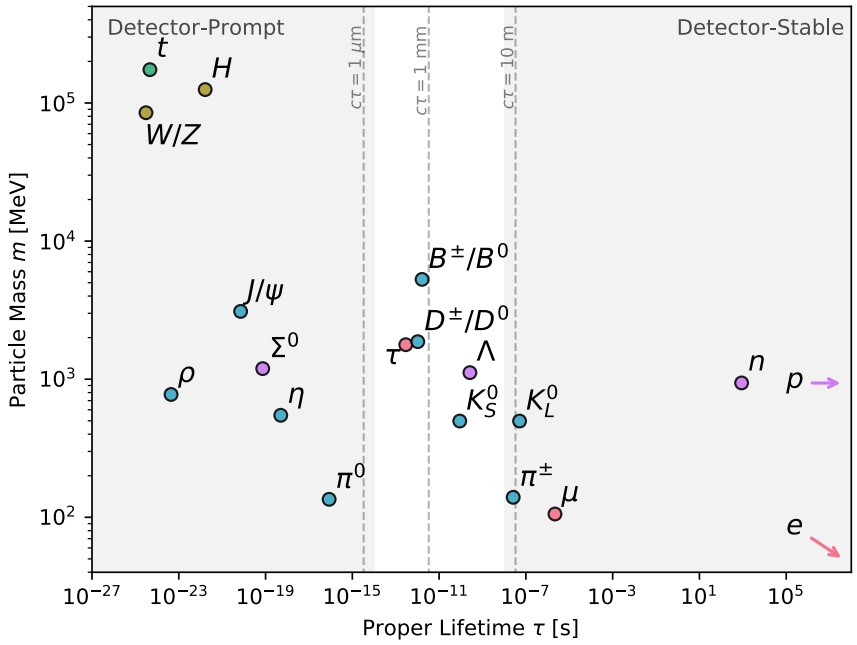

Figure 4.1 subset of SM particle content shown on the mass-lifetime plane. A spectrum of lifetimes may be observed, with a band of meta-stable particles shown in white. The existence of particles in the SM which may reasonably be classified as long-lived lends credibility to the hypothesis of new long-lived states beyond those harboured by the SM. Image extracted from [73].

$c\tau > 10$ m may be termed 'detector stable', a category which includes the muon, charged pions, and the K_L^0 meson. The intermediate lifetime window is home to particles such as the τ, K_S^0, and $B^\pm$. This category of particles may reasonably be described as long-lived. Particles within this window of proper lifetime traverse a distance within the detector before decaying at a secondary vertex, producing a signature analogous to those expected from BSM LLP states. In light of the precedent for long-lived particles in the SM, the existence of LLPs beyond the SM assumes a less radical hue.

Long-lived states can emerge in theories with a number of kinematical properties. Theories which feature decays *via* heavy states, for instance, can lead to particles with large lifetimes. In the SM, this is realised in the decay of the muon. In the process $\mu^- \to e^- \overline{\nu}_e \nu_\mu$, the muon decays *via* a heavy off-shell W^* boson. The lifetime of the muon may be calculated by:

$$\tau = \left[\frac{m_W}{m_\mu g_W}\right]^4 \cdot \frac{12\,(8\pi)^3}{m_\mu}. \tag{4.1}$$

As $m_W >> m_\mu$, the proper lifetime of the muon is determined as $\tau_\mu = 2 \times 10^6\,\mathrm{s}$. By the same logic, BSM scenarios which include decay processes via heavy intermediate states are free to produce long-lived particles.

Long lifetimes may also arise in theories where the available phase-space in which to decay is limited. Decay suppression is found to occur in the SM as a result of phase-space limitations. It is this phenomenon which is responsible for the difference in lifetime found between the K_S^0 and K_L^0 mesons in the SM. For the process $K_S^0 \to \pi^+\pi^-$, the mass splitting of $m_K - 2m_\pi = 220\,\mathrm{MeV}$ results in a kaon lifetime of $\tau = 9 \times 10^{-11}\,\mathrm{s}$. This compares to a more restricted mass splitting of $m_k - 3m_\pi = 80\,\mathrm{MeV}$ for the process K_L^0, producing a kaon lifetime of $\tau = 5 \times 10^{-8}\,\mathrm{s}$. Similarly small mass splittings in a BSM scenario would reduce the available decay phase-space, resulting in long-lived particles.

If dark matter is to be accounted for by a new particle, that particle must be stable on the order of the lifetime of the Universe (~ 13.8 billion years). By definition, this requires the acceptance of new particle states which may extend in lifetime beyond the remit of promptly-decaying SM particles. From this observation, the leap required to become comfortable with a spectrum of lifetimes in the BSM sector, akin to that of SM particles, is small.

4.4 Displaced τ-Leptons

Abreast of the arguments in support of long-lived particles beyond the SM, a search for LLPs in pp-collision data may be viewed as a well-motivated endeavour. The production of LLPs in hadronic collisions would be distinguished by the displacement of their child states within the detector — a discriminating feature which may be exploited by an informed search strategy. The specific properties of this decay signature is dependent upon the identity of the particle to which the LLP decays. The search described herein gives consideration to LLPs which decay preferentially to τ-leptons, for the reasons which follow.

4.4.1 The Third Generation

LLPs must couple to particles within the SM if they are to decay into visible, measurable end-states. Many models of physics beyond the SM feature mass-

dependent couplings, where the strength of the interaction between the BSM sector and SM particles is dependent upon the mass of the particles in question. This has given rise to many 'Higgs portal' models, where new particles are produced in various decay channels of the Higgs boson on account of the hypothesised non-zero Yukawa-style coupling between the Higgs and the new particles featured in the theory. Many other forms of mass-dependent couplings are further proposed.

The third generation fermions are the heaviest constituents of the SM. It is reasonable therefore to expect that the third generation may assume a privileged role in the decay of BSM states. As the τ is the most massive of the SM leptons, it is reasonable still further to predict that massive LLPs beyond the SM may decay preferentially to the SM τ lepton. This would induce the detector signature of a 'displaced' τ-lepton, where a SM τ is produced at the LLP decay vertex at a greater radial distance within the detector than is typical.

The long-lived BSM particles which have been introduced in this chapter are all expected to couple to the SM τ-lepton. Their production in a collision environment would therefore engender the signature of a displaced τ, whose decay products would emerge at an appreciable displacement within the detector when compared with those associated with typical τ production and decay.

Despite this strong theoretical and experimental motivation, experimental coverage of the phase-space corresponding to displaced τ production has been largely neglected — as may be seen in Figure 4.2. This is particularly true for the hadronic decay channels of the τ-lepton, which constitute approximately 65% of the associated branching fraction. The hadronic decay of displaced τ leptons is therefore ripe for further exploration.

This section and those which proceed it have collectively presented the case supporting a search for long-lived particles beyond the SM. The analysis designed and conducted in this thesis is such a search, focusing on the preferential decay of LLPs to SM τ-leptons. The philosophy which guides the strategic and analytical decisions underpinning this search is now discussed.

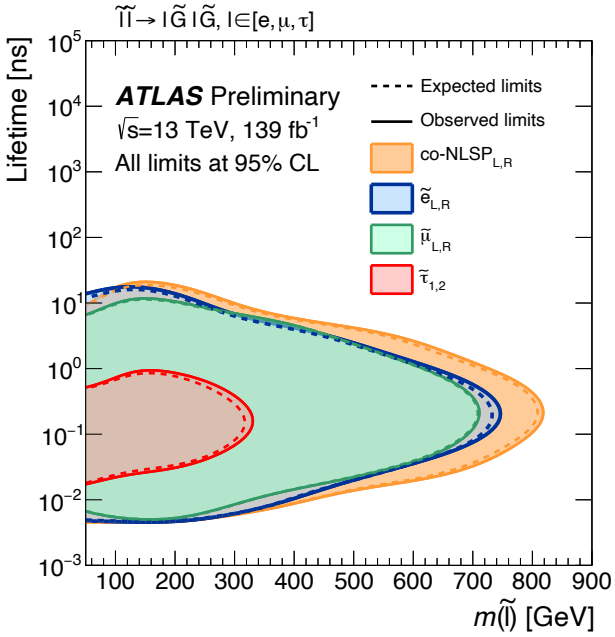

Figure 4.2 mass-lifetime coverage of the ATLAS search [74] for displaced leptons emerging from long-lived sparticles. Coverage of the phase-space associated with the τ-lepton does not extend far in lifetime.

4.4.2 Model Independence

The preceding chapters have examined the most promising theories proposed to-date to address myriad problems in theoretical physics. Having failed to deliver experimentalists to the discovery of new physics, the best theoretical instincts encapsulated by such models appear to have navigated researchers to a dead end.

It is for this reason that the thesis presented herein advocates an alternative approach: that of *model independence*. With this approach, minimal assumptions are imposed regarding the form new physics may assume. An analysis of collision data conducted in this manner is sensitive to BSM states not envisaged by theoretical prediction, with the widest possible aperture to the unknown. "There are known knowns, ...and there are known unknowns", as exclaimed by Rumsfeld. But there are also "unknown unknowns". The philosophy of model independence is one which is cognisant of this observation, and one which seeks to maximise sensitivity to particle messengers of new physics which lie beyond the current predictive power of the field.

Results obtained in a model-independent manner offer a further advantage: ease of re-interpretation. New models of physics beyond the SM which are proposed

in years subsequent to a model-independent search may be tested for survival against obtained exclusions with ease, as limits obtained in a model-independent manner are not restricted to the specific parameters of a given model.

The history of scientific progress is littered with examples of discoveries which blind-sided the research community. The discovery of the muon is such an example, prompting Nobel Laureate Isidor Isaac Rabi to proclaim "who ordered that?" A similar example is found in the Ultraviolet Catastrophe [75] at the turn of the twentieth century. This is a healthy and necessary component of the scientific method — when theoretical progress begins to abate, science turns to experimentalism for disruption. This cyclic process is observed throughout the history of science, where serendipitous and unexpected discoveries are found to play a critical role in the extension of our knowledge concerning the natural world. To bias an experimental search towards a single concrete prediction is to exclude the possible discovery of the unexpected. Model independence approaches discovery with a sense of humility, inherently open to new possibilities.

It is in this spirit that the search for displaced τ-leptons presented herein is conducted. The analysis presented in the pages overleaf is not tuned to the properties of a specific LLP prediction, and is thereby sensitive to the widest possible scope of possibilities. Effort is exerted at the design level to ensure the analysis is devoid of bias towards any particular BSM scenario.

This chapter has examined the case for long-lived particles as an explanation for the elusiveness of new physics. Moreover, motivation of both a theoretical and experimental persuasion for the active search for their presence in collision data has been presented. In light of the logic put forth in this chapter, the production of displaced τ-leptons has been shown to merit particular attention.

A search for displaced τ production may only be designed with the aide of simulated data samples. Such samples are required in order to understand the appearance of LLP signatures as perceived by detector systems, and to ascertain search sensitivity to signals of interest. Suppression of background processes of disinterest to such a search is also dependent upon the efficient production of large simulated samples which correspond to SM processes. The manner in which high energy processes are simulated is now considered.

> *A computer is only as smart as the person using it.*
>
> — Edward A. Murphy

5

Simulation of Physical Processes

The study of high energy physics at particle colliders is dependent upon the accurate simulation of particle interactions. Simulated data is used in every domain of collider physics: the calibration of measurements, the estimation of background processes, and the design of analyses which aim to uncover new physics. It is for this reason that Monte Carlo event generators may be regarded as the work-horses of high energy physics. This chapter introduces the mathematical machinery which underpins such event generators, before examining the methods by which they facilitate the simulation of fundamental particle interactions and the phenomenological signatures thereof.

5.1 The Role of Simulation in High Energy Physics

In the search for evidence of new physics in high energy collisions, a thorough understanding of the manner in which both SM and BSM processes manifest themselves in the environment of a collider is of paramount importance. The precise manner in which SM background processes are likely to evolve in repeated high energy collisions must be understood such that their effects on the observables of interest may be mitigated. Further, simulated data samples which represent the physical signature of new physical models — including those which correspond to LLPs — are required if an analysis which is sensitive to their production is to be designed. The responsibility of such tasks falls upon a common tool: the Monte Carlo (MC) *event generator*.

Event generators leverage efficient means of numerical integration such that the expressions which describe the interactions of interest may be solved computationally. The numerical techniques which serve as the computational engine of MC event generators are introduced in this chapter.

MC event generation represents the interface between high-energy theory and experiment, serving as a bridge between each domain. An intimate command of QFT and the SM is required in the development of event generators, in conjunction with a strong appreciation of the practical considerations imposed by the remarkable complexity of modern collider experiments. Three of the most ubiquitous software implementations used for MC simulation are `PYTHIA` [76], `HERWIG/HERWIG++` [77], and `SHERPA` [78]. The three main generators serve as the work-horses of most experimental analyses performed at colliders, facilitating the design of analytical techniques and enabling the sensitivity of measurements to be assessed.

Event generators typically separate the evolution of a hadronic collision into distinct components such that each component may be treated in isolation, before being concatenated into a realistic representation of a full event. Beyond the initial stages of the event, the strong coupling strength is such that perturbative methods are rendered unsuitable — phenomenological models are therefore employed to provide an accurate representation of the underlying physics, often tuned to data obtained from the experimental measurement of well-understood processes. Each event generator makes different choices with respect to such models, as will be highlighted throughout this chapter.

Before the physical models which underpin the simulation of hadronic collisions may be introduced, and understanding of the key numerical techniques which enable their computation is required. An elementary overview of such methods is now presented.

5.2 Numerical Integration

The generation of synthetic data samples which correspond to particle interactions requires the computation of various integrals. Numerical integration techniques allow the integrals which describe various physical quantities to be estimated computationally. A range of numerical integration methods are available; the *Monte Carlo* method is the method of choice for event generation.

5.2.1 Monte Carlo Method

A number of alternative options to the MC method of numerical integration exist. The trapezoidal method [79] and Simpson's method [79] feature prominently among this list. The Monte Carlo (MC) method offers a number of advantages over alternative methods, making it the most suitable choice for the task at hand.

The MC method converges upon a stable estimation of an integral of dimensionality d at a rate proportional to $1/N^{1/2}$. This compares to a convergence rate $\propto 1/N^{2/d}$ for the trapezoidal method, and a rate $\propto 1/N^{4/d}$ for Simpson's method. MC techniques therefore facilitate a more efficient estimate for high-dimensional integrals than alternative methods. Furthermore, the MC method enables the estimation of integrals of arbitrary complexity, with an error rate which may be determined with ease.

An estimation may be obtained via the MC method for an integral of the form:

$$I = \int_a^b \mathrm{d}x\, f(x) = (b-a)\ \langle f(x) \rangle \tag{5.1}$$

for a well-behaved function f, where $\langle f(x) \rangle \equiv \mathbb{E}[f(x)] = (b-a)^{-1} \int_a^b \mathrm{d}x\, f(x)$ is the

expectation value of $f(x)$. For such an integrand, the MC estimate is given by:

$$I_N = \lim_{N\to\infty} (b-a)\frac{1}{N}\sum_{i=1}^{N} f\left(x_i\right), \tag{5.2}$$

for N random numbers, x_i, drawn uniformly from $[a, b]$. The estimation is determined with an associated standard deviation of $\sigma = \sqrt{V(f)/N}$, where $V(f) \equiv \mathbb{E}[(f - \mathbb{E}(f))^2] = \mathbb{E}[f^2] - (\mathbb{E}[f])^2$ is the variance of function f. This is known as the *mean value theorem* of integration.

Equation 5.2 may be viewed as a uniform distribution of random variables, x_i, weighted by multiplicative factors $w_i \equiv f(x_i)$. Accordingly, a solution to integrals of the form given for I may be estimated through the generation of random numbers. For this purpose the *hit-and-miss* method may be employed.

5.2.2 Hit-and-Miss Method

A MC estimate of an integral may be obtained with the hit-and-miss method. This method proceeds with the generation of a set of random variables, of which a sub-sample which falls within the area cast by the function $f(x)$ is accepted.

The algorithm unfolds as follows:

1 **algorithm** *Hit-and-Miss*
2 define over-estimate, $f_{\max}$, such that $f_{\max} \geqslant f(x) \forall x \sim [x_{\min}, x_{\max}]$;
3 draw random variable, x, from the uniform distribution in $[x_{\min}, x_{\max}]$;
4 draw second random variable, y, such that $y \leqslant f_{max}$;
5 **if** $y \leqslant f(x)$ **then**
6 accept x;
7 **end**
8 **else**
9 reject x;
10 return to step 1.
11 **end**
12 **end**

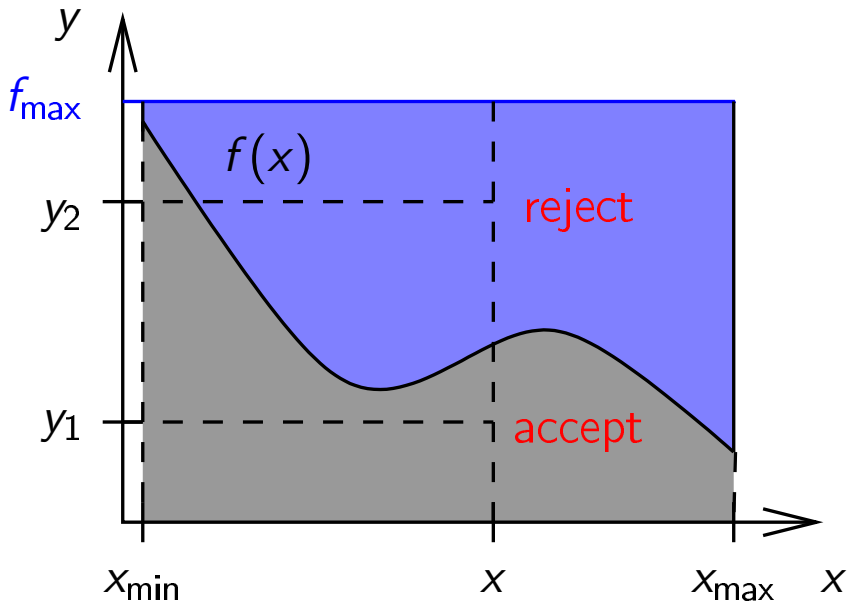

Figure 5.1 illustration of MC integration of function $f(x)$ with the *hit-and-miss* method. Figure extracted from [80].

With this procedure, an estimate of the integral of $f(x)$ is obtained by:

$$\int_{x_{min}}^{x_{max}} df\ f(x) = f_{max}\,(x_{max} - x_{min}) \cdot \frac{N_{accept}}{N_{total}} = A_{total}\frac{N_{accept}}{N_{total}}, \tag{5.3}$$

where N_{accept} and N_{total} respectively denote the number of accepted and total number of random points, and A_{total} is the total area of the (x, y) plane on which the numbers were evaluated. An illustration of this procedure is provided in Figure 5.1.

The hit-and-miss method is capable of integrating any function $f(x)$ provided it is bound from above, and boasts the further advantage of converging more quickly than alternative methods. Despite offering the fastest convergence of the main numerical integration techniques, this method proves inefficient when the variance of f is large. For this reason, a number of *variance reduction* techniques exist which may be employed to improve the rate of convergence while also reducing the error on the result.

Variance reduction Methods

Numerous methods of variance reduction are available, with which the efficiency of MC integration may be improved. A prominent example is found in the method

of *importance sampling*.

This techniques requires a 'friendly function', $g(x)$, to be determined such that $g(x) \geqslant f(x) \forall x \in [x_{\mathsf{min}}, x_{\mathsf{max}}]$. The friendly function serves as an upper-bound on the plane on which random numbers are generated, such that a smaller space of numbers is computed.

Having determined a suitable instantiation for $g(x)$, the integral of interest may be expressed as:

$$\int \mathsf{d}x\, f(x) = \int \mathsf{d}x\, g(x) \frac{f(x)}{g(x)} = \int \mathsf{d}G(x) \frac{f(x)}{g(x)}, \tag{5.4}$$

where $\mathsf{d}G(x) = \mathsf{d}x\, g(x)$. Having reduced the sampling space, the modified algorithm proceeds as follows:

1 **algorithm** *ImportanceSampling*
2 draw random variable, x, from $g(x)$;
3 draw second random variable, y, such that $y \leqslant g(x)$;
4 **if** $y \leqslant f(x)$ **then**
5 accept x;
6 **end**
7 **else**
8 reject x;
9 return to step 1.
10 **end**
11 **end**

This procedure is illustrated in Figure 5.2.

The efficiency boost provided by importance sampling may be further improved with *multi-channel sampling* [80]. In a similar fashion to importance sampling, this method invokes a function $g(x)$ to serve as a ceiling to the integrand on the (x, y)-plane.

Multi-channel sampling differs from importance sampling in the construction of

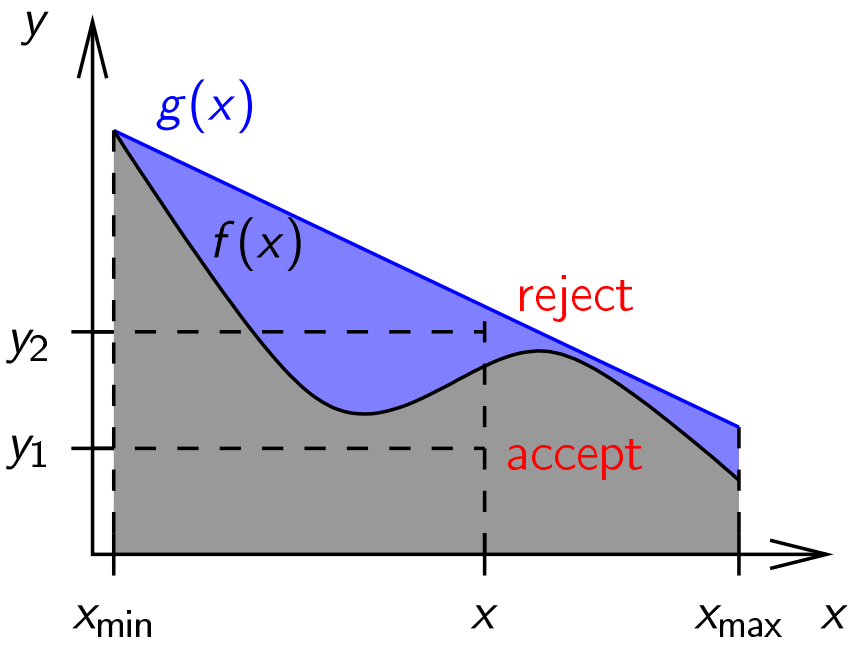

Figure 5.2 illustration of the reduced sampling space associated with the importance sampling variance reduction technique. Image extracted from [80].

the ceiling function, defining $g(x) \geqslant f(x)$ as a sum of friendly functions:

$$g(x) = \sum_i g_i(x), \quad \text{with } A_i = \int_{x_{\min}}^{x_{\max}} \mathrm{d}x\, g_i(x), \tag{5.5}$$

where A_i is the area encompassed by the i^{th} friendly function on the (x, y)-plane. The component friendly functions which sum to define $g(x)$ may be chosen wisely if *a-priori* knowledge of the peak-content of $f(x)$ is available from the physics at hand. Once a function $g(x)$ has been suitably defined, the variance reduction procedure may proceed in the same manner as was presented for importance sampling. An illustration of the corresponding (x, y)-plane is shown in Figure 5.3 (a).

The use of multi-channel sampling significantly reduces the space on which random numbers must be 'thrown' in order to compute an estimate for the integral in question, thereby significantly improving the efficiency of the MC method. The reduction in the error associated with the integral computation for the generation of N random numbers is shown in Figure 5.3 (b).

The methods of efficient numerical integration introduced in this section provide

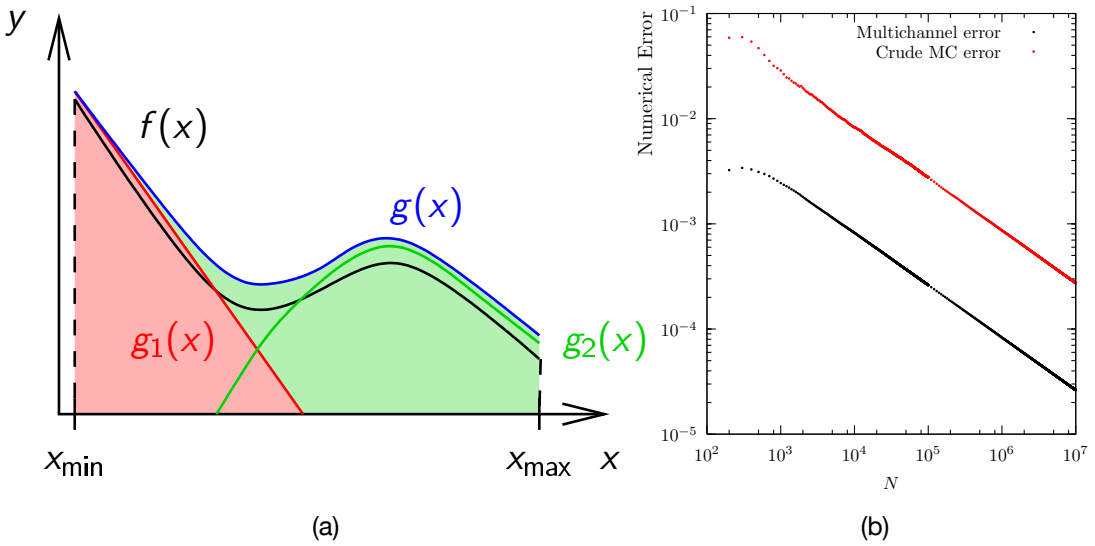

Figure 5.3 presents (*a*) the plane on which integrals are evaluated under the multi-channel sampling scheme, together with (*b*) the performance improvement attained with this method in comparison to the nominal ('crude') MC method. Part (a) sourced from [80], whereas part (b) is sourced from [81].

the computational toolkit with which the simulation of physical processes may be performed. The manner in which this is implemented in MC event generators is now examined.

5.3 Generation of Physical Events

The simulation of high energy particle interactions is underpinned by the methods of numerical integration introduced in the proceeding section. Collision events may be decomposed into distinct elements, whose computation is distilled into integrals which may be solved with MC methods.

Hadronic collisions may be broadly delineated into three different stages: the underlying hard-scatter event which produces the subsequent interactions, the showering process driven by QCD radiation, and the process by which the partons produced in the collision hadronise into the particles observed as final states. MC event generators typically address each step of this process in turn.

Before the primary hard-scatter event may be computed, a term for the cross-section of the process at hand must be determined. This necessitates the

computation of the corresponding matrix element for the process.

5.3.1 Matrix Element Computation

The S-matrix was shown in Chapter 1 to quantify the probability that a given quantum state will transition to another. A given element of this matrix therefore represents the probability that such a transition will occur *via* a particular process. For this reason, the computation of the matrix element which corresponds to the process of interest typically marks the initial step of the event generation sequence.

Matrix elements are calculated to LO in dedicated matrix element generators, such as `MadGraph` [82] and `COMIX` [83] (which is used in `SHERPA`). To LO approximation, matrix elements may be calculated with Feynman rules, which provides a suitable starting point upon which to build subsequent elements of the event simulation. NLO solutions are available for many processes, with $\mathrm{N^2LO}$ increasingly avilable for a more limited number of processes. Matrix element solutions provide accurate results for a small number of hard, well-seperated partons.

The cross-section of a given process — upon which subsequent elements of the simulation depend — is calculated as a function of the corresponding matrix element. In general, a cross-section for two partons, a and b, producing N final state particles may be expressed as:

$$\sigma_{ab\to N} = \int f_a\left(x_1, \mu_\mathrm{F}^2\right)\ f_b\left(x_2, \mu_\mathrm{F}^2\right)\ |\mathcal{M}|^2\,\Theta(\text{cuts})\,\mathrm{d}x_1\,\mathrm{d}x_2\,\mathrm{d}\phi_N, \tag{5.6}$$

for given kinematic cuts set by Θ, where the N-particle differential phase-space is given by:

$$\begin{aligned}\mathrm{d}\phi_N &= \left[\prod_{i=1}^{N}\frac{\mathrm{d}^4 p_i}{(2\pi)^4}\delta\left(p_i^2 - m_i^2\right)\Theta\left(E_i\right)\right](2\pi)^4\delta^{(4)}\left(p_a + p_b - \sum_{i=1}^{N} p_i\right)\\ &= \left[\prod_{i=1}^{N}\frac{\mathrm{d}^3 p_i}{(2\pi)^3 2E_i}\right](2\pi)^4\delta^{(4)}\left(p_a + p_b - \sum_{i=1}^{N} p_i\right).\end{aligned} \tag{5.7}$$

Here, μ_F sets the factorisation scale, with the four-momentum, mass, and energy of the i^{th} particle is given respectively by p_i, m_i, and E_i, and x_i representing the momentum fraction carried by the corresponding parton. The determination of the appropriate matrix element therefore allows the cross-section of a given process to be determined with the MC techniques previously introduced.

5.3.2 hard-scatter Event

The primary *hard-scatter* event is the first element of the physical proton interactions to be considered in the simulation chain. This corresponds to the main collision of interest, from which particles of importance to a given analysis emerge. The simulation of the hard-scatter event is performed through the evaluation of the production cross-section which corresponds to an N-particle state emerging from a proton-proton interaction.

Owing to the composite nature of the proton, the probability of a given parton-parton interaction must be considered. The momenta of incoming protons is distributed among the constituent partons in a non-trivial, energy-dependent distribution. As this distribution is not known *a-priori*, it is necessary to invoke the parton distribution functions (PDFs) previously introduced in Section 2.6.1.

The cross-section which corresponds to a $pp \to N$ process is then given by:

$$\sigma_{pp\to N} = \sum_{a,b} \int \mathrm{d}x_1 \, \mathrm{d}x_2 \, f_a\left(x_1, \mu_F^2\right) \, f_b\left(x_2, \mu_F^2\right) \, \widehat{\sigma}_N^{ab}, \tag{5.8}$$

where the cross-section for the interaction between two given partons, a and b, yielding N final states is given by $\widehat{\sigma}_N^{ab}$, and the probability distribution which describes the probability of parton a carrying momentum fraction x_1 is given by $f_a\left(x_1, \mu_F^2\right)$. This is determined at the factorisation scale, μ_F, as introduced in Section 2.6.1. Radiative effects which occur below this scale are folded into the definitions of the PDFs, with radiative corrections above this threshold included in the computation.

The procedure of renormalisation was introduced in Section 2.3.2 as a means by which to address the unphysical infinities which arise in perturbative expansions. It is necessary to set the scale at which such divergencies are absorbed into the definition of the underlying parton — the cross-section for a given hard-scatter

event is calculated as a function of this scale.

The hard-scatter process may then be simulated in accordance with the steps which follow:

1. obtain the expression for the partonic cross-section, $\widehat{\sigma}_N^{ab}$, corresponding to the process or model of interest (this may be determined with the matrix element generators introduced previously);
2. resolve divergences through the application of renormalisation and set the renormalisation scale, μ_R;
3. set the factorisation scale, μ_F, and consider soft radiative effects beyond this scale;
4. having absorbed the radiative effects below μ_F, combine PDFs in conjunction with $\widehat{\sigma}_N^{ab}$ such that the interaction cross-section may be determined (via Equation 5.8).

Having successfully simulated data which accurately captures the physical behaviour of the primary hard-scatter event, execution proceeds to the subsequent step in the simulation chain: the generation of *parton showers*.

5.3.3 Parton Showers

Just as electrically-charged particles emit photonic radiation under acceleration, accelerating partons which carry a QCD charge radiate gluons. Unlike the photons emitted by the electromagnetic Bremsstrahlung process, gluons carry colour charge under the SU(3) group — as such, the gluons produced in this process radiate further gluons themselves. This initiates an iterative cascade event which produces a shower of partons. Such cascades may be approximated by MC event generators *via* algorithms known as *parton showers*.

The gluonic radiation emergent from partons may be delineated into two categories: initial-state radiation (ISR), and final-state radiation (FSR). Initial-state radiation refers to the gluonic emissions which originate from the incoming partons which form the hard-scatter event. Final-state radiation contrastingly refers to the soft gluons which radiate from the outgoing partons which are produced in the primary hard interaction.

The parton shower algorithm may be understood with reference to the following example. Consider a final-state parton, i, which splits into two approximately collinear partons, j and k, each respectively carrying a momentum fraction z and $(1-z)$ of the parent parton momentum. The partonic differential cross-section for a general $(n+1)$-parton process is given by:

$$\mathrm{d}\widehat{\sigma}_{n+1} = \mathrm{d}\widehat{\sigma}_n \sum_i \int_{k_{\mathrm{min}}}^{k_{\mathrm{max}}} \frac{\mathrm{d}|k_{\mathrm{T}}|^2}{|k_{\mathrm{T}}|^2} \int_{z_{\mathrm{min}}}^{z_{\mathrm{max}}} \frac{\mathrm{d}z}{z} P_{ji}(z), \tag{5.9}$$

where k_{T} denotes the transverse four-momenta of parton k and $P_{ji}(z)$ is the DGLAP splitting function corresponding to the process $i \to jk$. The sum is taken over all possible partons, i, which may produce the final state parton, j. Lower limits imposed upon the possible momenta of k ensure the convergence of the integral, while upper momentum limits are set by the physical scale relevant to the process in question. Similarly, kinematic considerations dictate the limits placed upon the possible values of momentum fraction, z_{min} and z_{max}.

From Equation 5.9 it may be observed that the differential cross-section corresponding to the $(n+1)$-parton process, $\mathrm{d}\widehat{\sigma}_{n+1}$, factorises with the cross-section associated with the n-parton process, $\mathrm{d}\widehat{\sigma}_n$. The $(n+1)^{\mathrm{th}}$ parton, k, may therefore be treated independently. This result may be generalised to an arbitrary number of collinear partonic emissions, which will similarly exhibit mutual independence. It is therefore possible to approximate cross-sections for a large number of final-state partons. The partonic radiation which emanates from a set of well-separated partons produced by a given underlying hard-scatter event may be described by this factorisation in the collinear limit.

The term in Equation 5.9 which serves as a coefficient to the n-parton cross-section,

$$\sum_i \int_{k_{\mathrm{min}}}^{k_{\mathrm{max}}} \frac{\mathrm{d}|k_{\mathrm{T}}|^2}{|k_{\mathrm{T}}|^2} \int_{z_{\mathrm{min}}}^{z_{\mathrm{max}}} \frac{\mathrm{d}z}{z} P_{ji}(z), \tag{5.10}$$

may be interpreted as the probability of a parton emission with momentum fraction $z_{\mathrm{min}} \leqslant z \leqslant z_{\mathrm{max}}$ of parton i and transverse momentum $k_{\mathrm{min}} \leqslant |k_{\mathrm{T}}| \leqslant k_{\mathrm{max}}$. It is convenient to introduce the Lorentz-invariant quantity *virtuality*, defined as the squared momenta of a given particle, such that the integrands which

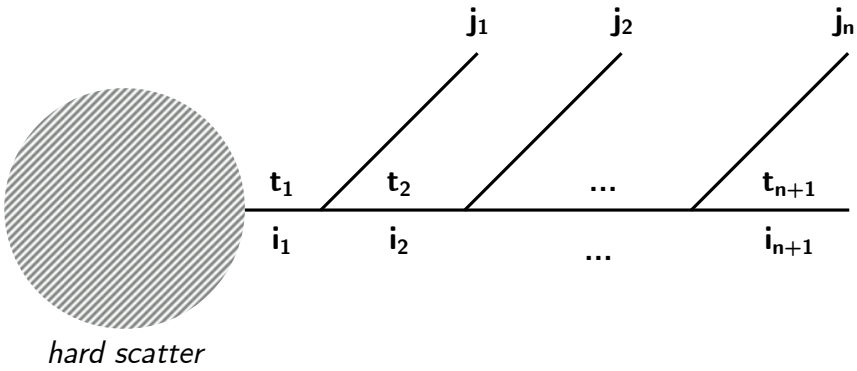

Figure 5.4 illustration of parton cascade emerging from a final-state parton produced in an $(n{+}1)$ hard-scatter process. Parton splittings are labelled with inidices and corresponding virtualities. Inspiration drawn from a similar graphic presented in [9].

correspond to probabilities of this form may be alternatively expressed as:

$$\int \frac{\mathrm{d}|k_{\mathrm{T}}|^2}{|k_{\mathrm{T}}|^2} = \int \frac{\mathrm{d}t}{t}. \tag{5.11}$$

Virtuality may be interpreted as a quantification of the degree to which a given particle is off mass-shell. In this sense, the more positive the metric of virtuality, the more closely the particle in question resembles a virtual particle.

For the collinear splitting of a parton, i, into an arbitrarily large number, $(n+1)$, of partons, a momentum fraction z_m and virtuality t_m may be defined before each soft emission. The initial momentum fraction carried by the parent parton is given by $z_1 = 1$. The species of the parent parton may change with each emission as the conservation of quantum numbers is preserved. By similar logic, the virtuality of the parent parton will decrease with each soft emission, as each radiated parton carries with it a non-zero fraction of momentum: $t_1 > t_2 > ... > t_{n+1}$. The parent parton may therefore be considered to become increasingly 'real' as the shower evolves. This process is illustrated in Figure 5.4.

Between any two values of virtuality, t_0 and t_1, there is a non-zero probability that the emitted parton is *non-resolvable* — soft partons of this nature escape measurement due to the finite resolution of particle detectors. Denoted by $\Delta(t_0, t_1)$, this probability may be expressed as:

$$\Delta_i\left(t_0, t+\delta t\right) = \Delta_i\left(t_0, t\right)\left[1 - \sum_j \int_t^{t+\delta t} \frac{\mathrm{d}t}{t} \int \mathrm{d}z\, P_{ji}(z)\right], \tag{5.12}$$

for a given parent parton, i, where appeal has been made to Equation 5.9. Here, $\Delta_i(t_0, t+\delta t)$ is the probability that no resolvable emissions will radiate from parton i between the virtualities t_0 and $t+\delta t$, while $\Delta_i(t_0, t)$ represents the probability of the same phenomena occurring between the virtualities t_0 and t. Unitarity is imposed on Equation 5.12 such that probability is correctly conserved. Given that terms of the form summed over index j were shown to correspond to the probability of a single emission occurring between the virtuality values of t and $t+\delta t$, the full term contained within the square brackets can therefore be understood, by probability conservation, to give the probability of *no resolvable emissions* occurring within the virtuality window set by t and $t+\delta t$ (assuming small values of δt).

Equation 5.12 can therefore be interpreted as equating the probability of a non-resolvable emission occurring within the range $(t_0, t+\delta t)$ to the product of non-resolvable emission probabilities for the ranges (t_0, t) and $(t, t+\delta t)$ — a statement which passes the test of intuition. Solving this equation for $\Delta_i(t_0, t)$, a Taylor-expansion may be taken around δt:

$$\Delta_i(t_0, t) + \delta t \frac{\mathrm{d}}{\mathrm{d}t}\Delta_i(t_0, t) + \cdots = \Delta_i(t_0, t)\left[1 - \sum_j \frac{\delta t}{t}\int \mathrm{d}z\, P_{ji}(z)\right], \tag{5.13}$$

yielding a differential equation in Δ_i:

$$\frac{\mathrm{d}}{\mathrm{d}t}\Delta_i(t_0, t) = -\frac{1}{t}\left(\sum_j \int \mathrm{d}z\, P_{ji}(z)\right)\Delta_i(t_0, t)\,. \tag{5.14}$$

A solution to this differential equation is given by:

$$\Delta_i(t_0, t) = \exp\left[-\sum_j \int_{t_0}^{t} \frac{\mathrm{d}t'}{t'} \int_{z_{\min}}^{z_{\max}} \mathrm{d}z\, P_{ji}(z)\right]. \tag{5.15}$$

This expression is known as the *Sudakov factor*. The Sudakov factor, as has been shown, quantifies the probability that a soft, non-resolvable partonic emission will radiate from a given parton, i, within the virtuality range cast by (t_0, t). The derivation of this quantity invoked the concept of probability conservation — as unitarity can only be achieved within a QFT with the full consideration of virtual corrections, both real and virtual emissions are thereby accounted for by

this factor.

From the derivation presented thus far, it is apparent that a given partonic cascade may be characterised by the set of values formed by $\{(z_i, t_i)\}$, denoting the momentum fraction and virtuality as they are valued before each radiative emission. Equipped with an expression for the Sudakov factor, the probability associated with a given partonic cascade is obtained by:

$$P^{j_1,j_2,\cdots,j_n}_{i_1,i_2,\cdots,i_n}\left(\{z_i\},\{t_i\}\right) = \left[\prod_{m=1}^{n}\Delta_i(t_m, t_m+1)\right]\left[\prod_{m=1}^{n}\int \mathrm{d}z_{m+1}\; P_{i_{m+1}i_m}\left(z_{m+1}\right)\right], \quad (5.16)$$

where $z_{m+1} = z_{m+1}/z_m$ gives the fraction of momentum carried from the m^{th} parton by the $(m+1)^{\text{th}}$ parton.

Equation 5.16 enables the numerical generation of partonic cascades, provided that a set of numbers can be generated for $(\{z_i\},\{t_i\})$. This may be accomplished with the Monte Carlo method introduced in Section 5.2.1, while the integral may similarly be solved with numerical methods. The parton shower algorithm proceeds in this manner, in accordance with the steps outlined in Algorithm 1.

While this discussion of the parton shower algorithm is centred on final-state radiation, the same logic may also be applied to the soft emissions radiating from initial-state partons. In the case of initial-state radiation, PDFs must be included into the definition of the Sudakov factor such that the distribution of momentum among incoming partons is considered. Further, the treatment presented in this chapter may be extended beyond the collinear limit. Soft, wide-angle radiation may also be simulated with the parton shower algorithm — this is implemented with a procedure known as *angular ordering*, as described in [9].

5.3.4 Matching and Merging

It has been shown that the parton shower algorithm may be used to estimate the effects of partonic radiation. An alternative approach is found in the calculation of higher-order terms in the relevant matrix element.

The two approaches to partonic emission estimation are largely orthogonal in their qualities, as is summarised in Table 5.1. Fixed-order matrix element generators

algorithm *PartonShower*

define parent parton with virtuality t_1 and momentum fraction z_1

do

compute t_{i+1} via the numerical solution to Sudakov factor $\Delta_i(t_{i+1}, t_i) = r_i$, for random number $r_i \sim [0, 1]$

generate momentum fraction $z_{m+1} = z_{m+1}/z_m$ by solving

$$\frac{\int_{z_1}^{z_{m+1}} \mathrm{d}z' \, P_{i_{m+1} i_m}(z')}{\int \mathrm{d}z' P_{i_{m+1} i_m}(z')} = r_z,$$

for random number $r_z \sim [0, 1]$

generate azimuthal angle $\phi \sim [0, 2\pi]$ for parton $(m + 1)$

increment parton number m by 1

while $t_{i+1} \geqslant t_0$;

end

Algorithm 1: sequence of steps by which a parton cascade may be simulated numerically.

provide accurate results for a small number of hard, well-separated partons, yet perform poorly for the generation of many soft, collinear emissions. Parton shower algorithms, meanwhile, attain accurate results when generating many soft, collinear partons, but perform poorly in the generation of a small number of hard, well-separated partons.

Parton Shower Method	Matrix Element Method
Can simulate large № partons	Limited to small № partons
Poor approx. of widely-separated partons	Exact solut'n for well-separated partons
~Exact solution for adjacent particles	Incapable of simulating close particles
Limited incl. of quantum interference	Incl. all orders of quantum interference

Table 5.1 comparison of advantages and disadvantages offered by the parton shower and matrix element methods of simulating partonic radiation.

It is natural to ask whether the respective advantages of each method may be combined. This logic motivates the technique of *matching and merging*.

This procedure operates by correcting for the first few terms of the parton shower

solution with the solution obtained from the matrix element generator to (N)LO precision. Subsequently, having removed the overlapping terms (such that double-counting is avoided), the corrected parton shower terms may be merged with the solutions obtained from matrix element computation. Generated events must be re-weighted such that they may be combined in a sensible manner. The merging scale must be chosen such that the matrix element and parton shower solutions may be clearly delineated. Care must be taken to ensure that the property of unitarity is preserved by merging procedure.

For terms computed at NLO or above, further care must be taken to avoid the double-counting of virtual contributions. This issue arises due to the Sudakov factor: as was previously shown, this factor naturally accounts for virtual effects due to its preservation of unitarity. Some systems, such as `POWHEG` [84], overcome this issue by implementing a small modification to the definition of the Sudakov factor. Alternative solutions to this issue are found in [9].

5.3.5 Hadronisation

The manner in which the quarks and gluons produced in the parton showering process combine to form baryons and mesons must be simulated by Monte Carlo event generators. This is implemented in the *hadronisation* stage of the simulation.

The property of colour confinement introduced in Section 2.3.3 dictates that the partons produced in parton showers must combine to form colour-neutral states. The manner by which particles charged under SU(3) form colour-singlets cannot be described from the first-principles of QCD — the storng coupling strength lies outwith the regime in which perturbative techniques are applicable. It is therefore necessary to rely upon phenomenological models of the hadronisation process.

Two main phenomenological models of hadronisation exist: the *Lund string model* [85], and the *cluster* [86] model.

Lund String Model

The Lund string model is constructed on the basis that the strong potential between a quark and an anti-quark increases as they separate — in accord with

the discussion on asymptotic freedom in QCD presented in Section 2.3.2. The gluon field intensifies along the line which joins the quark and anti-quark as their separation increases. This is the flux tube introduced in Section 2.6, also known as the QCD string.

The strong potential is found to increase linearly with quark/anti-quark separation, as is verified by Lattice QCD solutions. For this reason, the QCD string is constructed such that it possesses uniform energy per unit length. As the quark/anti-quark pair separate, the strong potential increases to such an extent as to render the production of a new quark/anti-quark pair from the vacuum energetically economical. The QCD string is thereby split in two, yielding two new strings.

The Lund string model constructs QCD strings between the quarks produced in the parton showering process. The strings which extend between each quark/anti-quark pair break iteratively until further breaking is kinematically prohibited. Hadrons may then be formed by grouping adjacent quarks and anti-quarks together.

The probability that a QCD string should break is provided by the *fragmentation function*. This function may be expressed as:

$$f\left(z, p_{\mathsf{T}}\right) \sim \frac{1}{2}\left(1-z\right)^{a} \exp\left[-\frac{b\left(m_{\mathsf{h}}^{2}+p_{\mathsf{T}}^{2}\right)}{z}\right], \tag{5.17}$$

where a, b, and the hadron mass, m_{h}, are tunable parameters. Quarks are assigned a p_{T} value drawn from a Gaussian distribution at each juncture in the string-splitting process.

This process alone, of course, is only capable of producing mesons. To model baryonic formation, the concept of 'diquarks' is introduced. Diquark (anti-diquark) states are those which correspond to composite systems of two quarks (anti-quarks). The process of grouping a quark (anti-quark) with a diquark (anti-diquark) then yields (anti-) baryons.

Gluonic radiation from quarks must also be considered. Gluons carry both colour charge and anti-colour charge, such that QCD strings may be constructed from a quark to a gluon, and from the same gluon to an anti-quark. This appears as a regular QCD string with a kink where the colour-lines meet at the gluon. Gluons

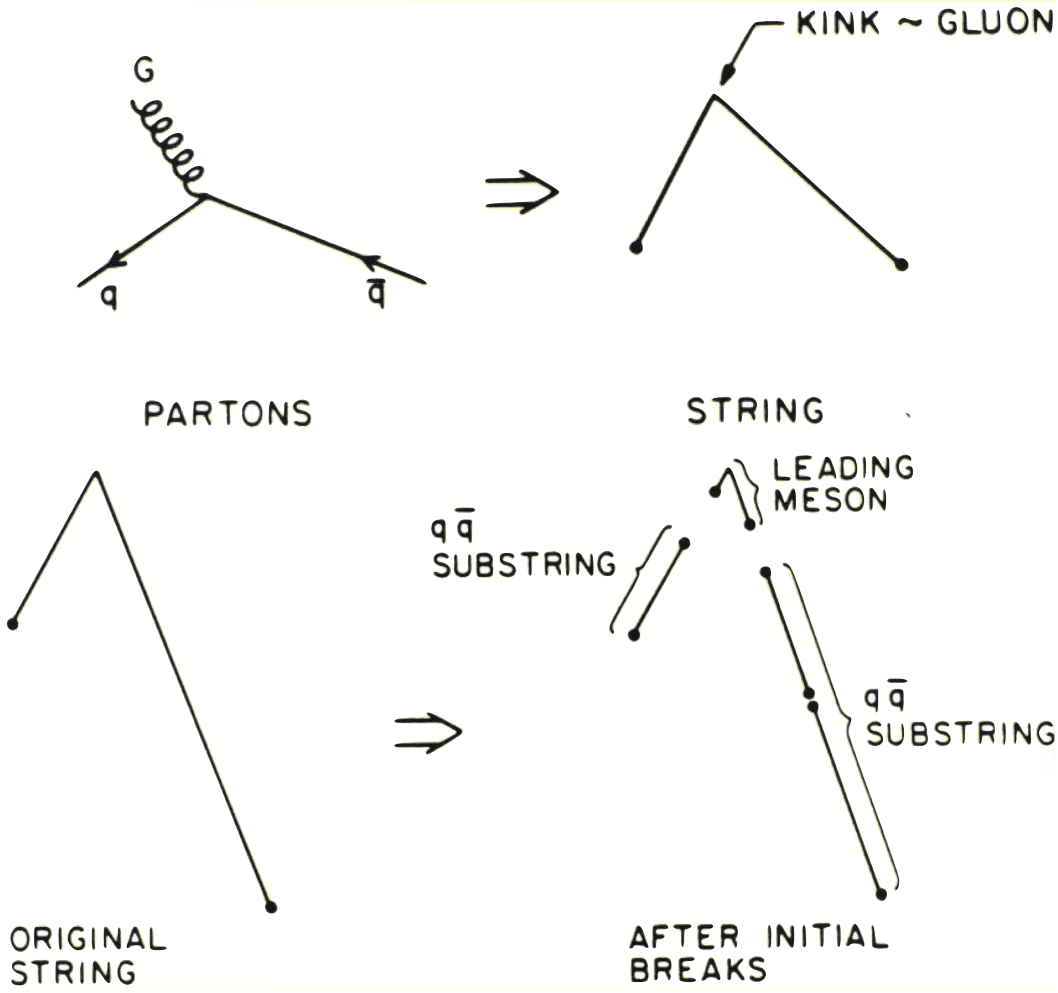

Figure 5.5 illustration of the kinked string representation of gluons in the hadronisation process. Lund strings may subsequently be broken in an iterative fashion, yielding colour-singlet hadronic states. Source: [87].

may therefore be modelled with kinked QCD strings, as is illustrated in Figure ??.

The Lund string model of hadronisation features a number of free parameters. The value of these parameters may be tuned such that the predictions of the model agree with experimental data. Data which correspond to well-understood processes thereby enable the Lund string model to produce physically-accurate simulations of the hadronisation process in the absence of an analytical solution derived from the first-principles of QCD.

This model of hadronisation boasts the advantages of providing an excellent description of experimental data, and of doing so from a phenomenological basis

with strong physical motivation. It does, however, require the optimisation of many free parameters. A model which incorporates additional features of QCD is therefore desirable — this can be found in the *cluster* hadronisation model.

Cluster Model

The cluster model of hadronisation is based upon the concept of colour *pre-confinement*. The principle of pre-confinement states that, subsequent to the parton showering process, it is always possible to form clusters of colour-singlet hadronic states from the radiated partons. It may be shown that the mass distribution of such clusters is calculable from the first-principles of QCD — for this reason, the cluster model may be considered as more theoretically complete than the Lund string model. Fewer tunable parameters are required as a result.

Clusters are formed by first splitting gluon and diquark content into quark/anti-quark pairs in the large $N_c \to \infty$ limit. Adjacent (anti-) quarks and diquarks will therefore be of matching (anti-) colour quantum numbers, and may then be combined in a simplistic manner.

The clusters which result from this process may then be matched to known hadronic states by their invariant mass. If no matching hadronic partner is found, the cluster will decay into pairs of hadrons until a matching hadron may be associated with the sub-clusters. This results in a set of 'proto-hadrons' which decay into the final-state hadrons observed in the simulated event.

Although fewer in number than those found in alternative models, the cluster model is configurable with a set of tuneable parameters. The value of these parameter allows the uncertainty associated with the description of this process to be off-set, having been tuned with experimental data.

Both the Lund string model and the cluster model of hadronisation allow Monte Carlo event generators to simulate data in excellent agreement with experimental measurement. The most common event generators in high energy physics each make a different choice of hadronisation model. `PYTHIA` implements hadronisation with the Lund string model, whereas `SHERPA` and `HERWIG++` each use their own bespoke variants of the cluster model.

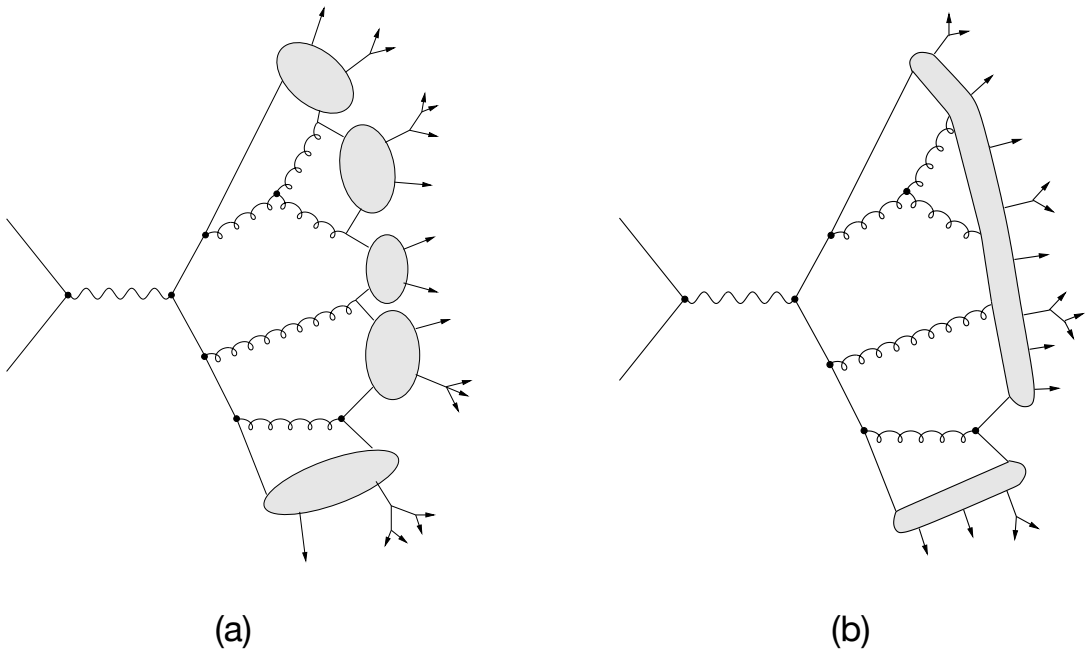

Figure 5.6 Feynman-style depiction of the hadronisation process as it unfolds in the (*a*) Lund string and (*b*) cluster model. Source: [81].

5.4 Underlying Event

The hard-scatter, parton shower, and hadronisation process form the main elements of a MC event generator. Additional elements, however, must be considered if generated events are to provide a fuller description of real experimental data — as conveyed by the complexity shown in Figure 5.7 and Figure 5.8. Such considerations are said to pertain to describing the *underlying event*, which, in this context, includes all phenomena not directly relating the hard-scattering, showering, or hadronisation processes.

In seeking to provide a closer description of experimental data, the formal, concrete concepts provided by QFT are rendered less applicable — the number of practical considerations which enter the scenario produce intractable theoretical calculations. It is therefore typically necessary to rely on further phenomenological models in the simulation of the underlying event.

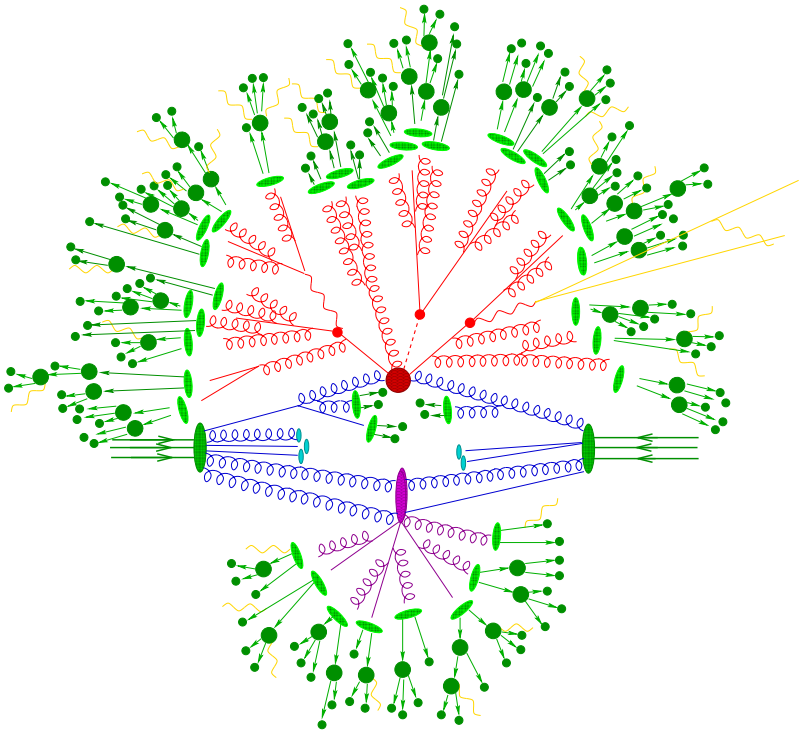

Figure 5.7 illustration of typical simulated hadron-hadron collision, depicting each stage of the Monte Carlo simulation process. The central red blob depicts the primary hard-scatter event, from which the Bremsstrahlung radiation which forms the parton shower is shown in red gluon and quark lines. Gluons which originate from the colliding parent hadrons are shown in blue, while a secondary hard-scatter event is depicted by the purple blob. Light-green blobs illustrate the hadronisation of child partons, while dark-green blobs indicate the decay of hadrons. Soft photonic radiation is shown in yellow. Image extracted from [80].

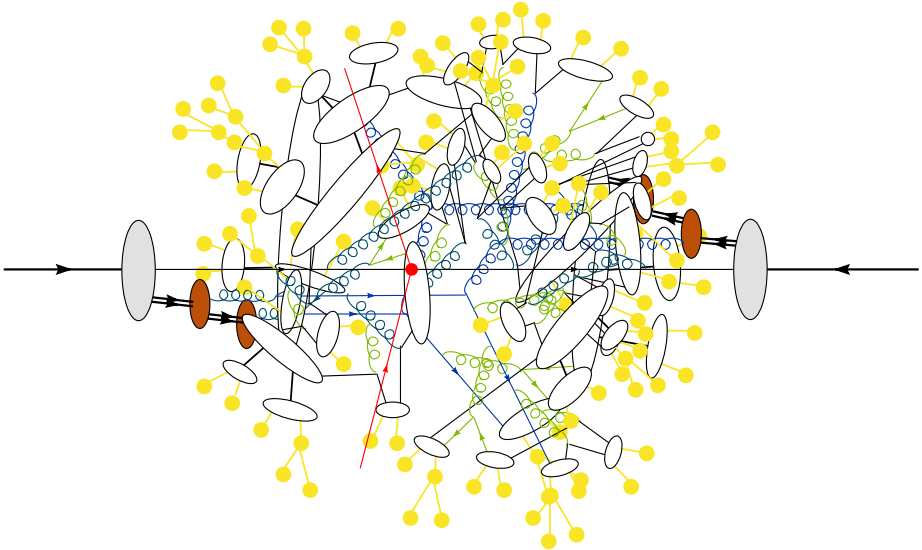

Figure 5.8 alternative depiction of a hadron-hadron collision, as simulated by a Monte Carlo event generator. The stages of event evolution cited above are similarly represented here. Image extracted from [81].

5.4.1 Beam Remnants

One element of the underlying event which must be addressed is the modelling of *beam remnants*. This refers to partons which belong to the incoming hadrons and do not partake in the primary hard-scatter interaction. The behaviour of the remaining partons which do not form part of the hard-scatter event must be modelled if an accurate representation of a hadronic collision is to be simulated.

Remnant quarks may be combined by the generator into diquarks, whose flavour is determined by the flavour of the interacting (*i.e.* non-remnant) quark. The interacting quark and the diquark must collectively form a colour-singlet system on the grounds of conservation. Similarly, total energy and momenta must also be conserved across the collective system formed by the interacting quark and remnant diquark.

Beam remnants must be included in the hadronisation process, else the condition of colour conservation is not satisfied. This stage of the event generation process is visualised in Figure 5.9.

5.4.2 Multi-Parton Interactions

The QCD factorisation theorem (Equation 2.59), introduced in Section 2.6.1, suffers from the deficiency of only giving consideration to single-parton exchanges. Only events in which a single parton of one incoming hadron interacts with a single parton of the other incoming hadron are described by the theorem.

In reality, multiple partonic interactions may occur — a scenario which is particularly likely where the overlap of incoming hadrons is large. The inclusion of multiple partonic interactions (MPIs) is required if generated samples are to agree with experimental data. Despite this, no factorisation theorem which includes MPI effects is known — phenomenological models are therefore relied upon to simulate effects of the kind illustrated in Figure 5.9.

5.4.3 Pileup Interactions

Further consideration must be given to interactions which arise from the collision of other protons included in the same bunch-crossing [9] as those which produce

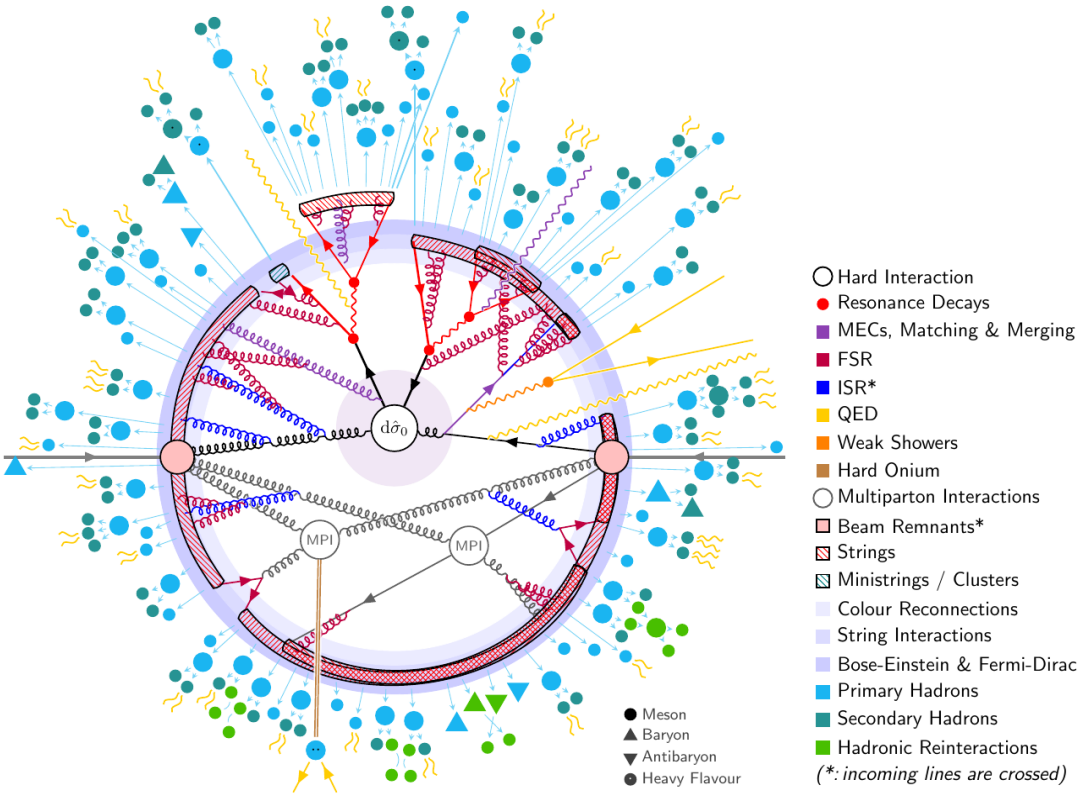

Figure 5.9 illustration of each stage in the Monte Carlo event generation sequence. The stage at which colour reconnection effects become pertinent is shown in light-lilac, with other simulation stages shown as denopted by the legend, *right*. Image extracted from [76].

the primary hard-scatter event. At the Large Hadron Collider, an average of ~ 33 interactions are produced in every bunch-crossing, with 40 million bunch-crossings produced every second (resulting in approximately 1 billion *pp* collisions per second) — a matter to which the following chapter will return.

To replicate this environment, generated hard-scatter events are overlayed with simulated *minimum-bias* events: soft, inelastic collisions which are modelled on experimental measurements of similar events. Overlapped minimum-bias events are then re-weighted with a scheme known as pile-up re-weighting [76, 77, 78].

The main elements of the software event generators used to simulate *pp*-collision data have now been presented. The computational technology discussed in this chapter was used extensively in the design of the analysis presented in this thesis, as will be shown in Chapter 10.

This chapter has introduced the computational mechanisms by which particle collisions are simulated — an essential technology which enables the search for long-lived particles put forth in this thesis. Phenomenological models of parton showering and hadronisation are combined with fixed-order perturbative matrix element solutions to provide an accurate representation of the full evolution of a proton-proton collision. The full sequential integration of each simulation component described in this chapter, including the role of colour reconnection, is depicted in Figure 5.9. Any number of such events for a given process may be generated in this manner, providing a simulated data-sample with which to design experimental analyses.

The discussion of MC event generation marks the final topic in the introduction to the theoretical framework of this thesis. Equipped with an understanding of the core theoretical concepts and computational methods with which a search for new physics may be conducted, an introduction to the experimental methods by which high-energy physics measurements are performed is now required. An introduction to the manner in which measurements are undertaken with the ATLAS experiment is now presented.

Part II

Experimental Apparatus & Techniques

> *Anyone who thinks the LHC will destroy the world is a twat.*
>
> — Brian Cox

6

The Large Hadron Collider

The Large Hadron Collider is the flagship particle collider of our era. At almost 27 km in circumference, it constitutes the largest and (in conjunction with its associated detectors) most complex machine built by mankind to date. It exists to service four major high energy physics experiments — ATLAS, CMS, LHCb, and ALICE — by providing them with ultra-relativistic protons travelling both clockwise and counter-clockwise around the ring of the accelerator. Protons are brought to collision at each experiment, allowing the most fundamental quantum fields in the Universe and the interactions between them to be probed. This chapter reviews the enabling technologies which underpin the operation of the Large Hadron Collider, before evaluating the performance it attains.

6.1 Overview

The preceding chapters have depicted the current landscape of fundamental physics, and the need for new empirical disruption to challenge our existing theoretical models. Particle colliders are the most promising tool with which to stress-test the predictions of the SM and search for resonances beyond its predictive foresight. The energies attainable with modern colliders allow the partonic interactions produced in their collisions to access heavy mass states, providing the deepest window into the fundamental nature of reality in humanity's possession. The energy generated by particle accelerators is dispersed into the fundamental quantum fields of the vacuum, giving rise to exotic particle phenomena — phenomena which may act as a sign-post to new physics.

The Large Hadron Collider (LHC) is a circular particle accelerator located at the European laboratory for fundamental particle physics, CERN, near Geneva, Switzerland. Extending to 26.7 km in circumference, the LHC occupies the tunnel which previously housed the former Large Electron-Positron (LEP) collider [88], having ceased operation at the turn of the millennium. The distance below ground at which the collider sits varies as a function of a number of geological considerations — at its deepest point, it can be found at approximately 170 m below ground. Under nominal operation, the LHC accelerates proton bunches to the highest energies it can attain. The initial run of collisions (run-1) [89] operated at a maximum energy of 4 TeV per beam, taking the total centre-of-mass energy to 8 TeV. During its second run (run-2) [90] a maximum energy of 13 TeV was attained, increasing yet further to 13.6 TeV in the subsequent (and ongoing) run-3 [91]. An additional option exists to accelerate lead ions for a proportion of the year [92], facilitating the study of heavy-ion collisions and their associated phenomena, such as quark-gluon plasma. Focus will be restricted to proton acceleration hereinafter, as is pertinent to the research undertaken in this thesis.

Proton beams are accelerated along two separate beam-pipes in opposite directions, brought to collision at each experimental detector along the collider ring. The suite of LHC experiments exists to cover the full confluence of high energy phenomena produced at the collider in a complimentary fashion. The LHCb [93] and ALICE [94] experiments are tasked with studying the respective domains of flavour physics and heavy-ion physics. ATLAS [95] and CMS [96], conversely, are general purpose detectors, designed to provide the widest possible aperture with

which to study high energy physics. The duality of general purpose experiments ensures the re-producibility of results: a cornerstone of the scientific method. At full design capacity, the LHC will operate at a total centre-of-mass energy of $\sqrt{s} = 14$ TeV, corresponding to 7 TeV per proton beam. (Here, s denotes the Lorentz-invariant Mandelstam variable [3].) Before maximum beam energy is reached, protons will reach the maximal-limit of their velocity imposed by special-relativistic restrictions as their velocity converges upon the limit of $v \approx c$. Beyond this point, additional energy contributes to the valence quark and gluon content of the proton — thereby increasing the likelihood of partonic interactions between the colliding protons at each interaction point. Protons therefore gain mass as they are accelerated to their maximum energy.

6.2 Proton Source and Injector Chain

In a manner analogous to a first-stage rocket booster providing initial thrust before main engines carry a spacecraft to orbit, protons are first pre-accelerated before entering the LHC. This initial acceleration stage is achieved by the *injector chain* [97], a series of pre-accelerators each accelerating protons to successively higher energies.

Protons are sourced from hydrogen gas which is deprived of its electrons by an electric field and fed to the injector chain. The accelerator complex begins with *Linac2* [98], a 37 m linear accelerator which takes protons from rest to an energy of 50 MeV. From here they are ejected into the *Proton Synchrotron Booster* (PSB) [99], a circular accelerator with a circumference of ~ 157 m. It is within the PSB that the protons are further accelerated to an energy of 1.4 GeV. The next step of the accelerator complex is the *Proton Synchrotron* (PS) [97], a ~ 628 m circular accelerator providing a further energetic kick of ~ 24.6 GeV before feeding the protons to the final link in the injection chain: the *Super Proton Synchrotron* (SPS) [100]. With a circumference of 7 km, SPS takes the protons to an energy of 450 GeV before injecting them into the main LHC ring via two tunnels each 2 km in length. A schematic of this accelerator network is shown in Figure 6.1.

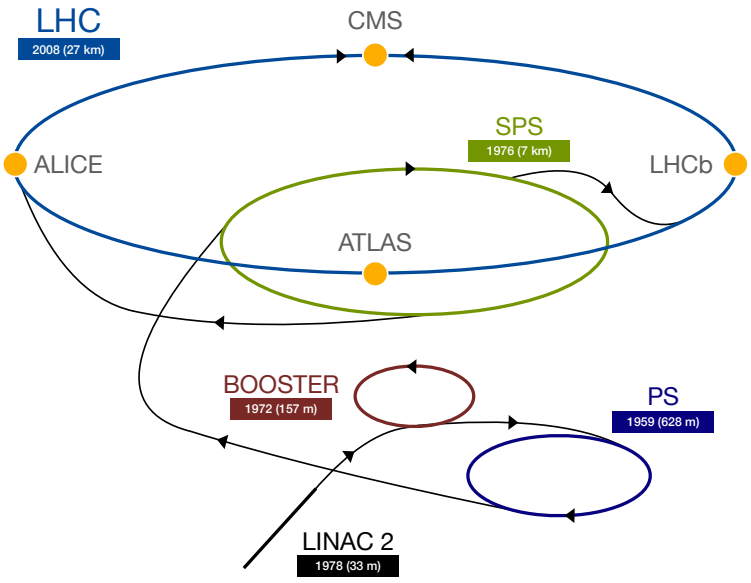

Figure 6.1 schematic of the full accelerator complex required for the operation of the LHC. A network of smaller accelerators form an injector chain which supplies protons to the main ring of the LHC where they will reach their maximum energy. Figure sourced from [97].

6.3 LHC Accelerator Technology

The ring formed by the LHC is not perfectly circular but is instead formed by eight octants, each of which can be further decomposed into an arced segment and a straight segment. Taken collectively, the arc segments of the LHC are home to approximately $1,200$ dipole magnets responsible for deflecting the beam along the correct path. Straight segments of the LHC ring serve a number of purposes. Four of the eight linear segments are used to perform tasks pertaining to beam hygiene, including beam cleaning and beam dumping. Such segments are also home to the radio-frequency cavities used to accelerate the beam with each circumnavigation. The remaining four straight segments of the ring are home to the 'interaction points' which house the experiments.

6.3.1 Radio-Frequency Cavities

Proton beams are brought from their pre-accelerated energy of 450 GeV to their target energy by the superconducting radio-frequency (RF) cavities [101] of the LHC. Eight RF cavities can be found in each direction of the ring in the form of metallic chambers, inside which an electromagnetic field provides an energetic

boost to the protons as they traverse the relevant segment of the collider. RF cavities are housed within cryogenic modules ('cryomodules') which refrigerate the cavities to a temperature converging upon $1.9° \text{ K}$, required to sustain their superconducting state. The LHC is thus one of the coldest places in the known Universe.

The active component responsible for the accelerative power of each RF cavity is known as the 'klystron' [101]: a series of high-power electron beams in cylindrical housing. Electron beams are oscillated at a frequency of 400 MHz, thereby inducing an oscillating magnetic field. The resultant electromagnetic energy is guided to the cavities by conducting wave-guides, whose geometry is optimised such that a resonant amplification of the wave is achieved. This can be viewed as reminiscent of the operation of a guitar chamber. Charged particles experience the focused electromagnetic field as an accelerating force.

Maximum beam energy is attained after a duration of approximately 20 minutes, after which the particles will have circulated through the 27 km accelerator ring over 10 million times — or, equivalently, over 11,000 times per second.

6.3.2 Superconducting Magnet System

Main Dipole Magnets

The large dipole electromagnets [102] which comprise the arcs of the LHC ring are responsible for deflecting the charged particle beam in such a manner as to maintain the precise trajectory of the beam through the accelerator. This is achieved by inducing a powerful 8.3 T magnetic field — more than $100,000$ times more powerful than the magnetic field of the Earth. The induction of such a large field necessitates a colossal electric current. For this reason, power is supplied via superconducting coils in order that energy loss due to electrical impedance is mitigated. (It may be perceived as somewhat poetic that the spontaneous symmetry breaking responsible for super-conduction played an integral role in the discovery of spontaneous symmetry breaking in the electroweak sector of the Standard Model.)

Insertion Magnets

Insertion magnets [103] are responsible for guiding the proton beam to the experimental areas where the particle detectors are located. Systems known as 'inner triplets', consisting of three quadrupole magnets, are used to tightly collimate particle beams before injection into the detector structures; tightly-compacted proton bunches secure a greater probability of collision. Each of the four large experiments on the LHC is home to a pair of inner triplets with which the beams are collimated before being brought to collision.

6.3.3 Thermal Cooling

The superconducting state required to minimise the electrical resistance of the magnet systems can only be sustained in an ultra-cold thermal environment. The main electromagnets of the LHC are kept at a temperature of 1.9° K (−271.3° C) — by contract, the vacuum of outer-space has an average temperature of 2.7° K (−270.5° C).

Sustaining such low temperatures is the task of the LHC cryogenic system. A supply of 40 MW of electricity and 120 tonnes of Helium is consumed to ensure the temperature of the Niobium-Titanium (NbTi) magnet coils can sustain their superconducting state.

The cooling process unfolds in three stages over a number of weeks (to avoid subjecting the systems to excessive mechanical stress). The initial stage of the process sees the liquid Helium cooled to a provisional temperature of 4.5° K. The second state entails the injection of the liquid Helium into the 'cold mass' components of the magnet systems, before being subsequently cooled further to a 1.9° K super-fluid state during the third and final stage of the cooling procedure.

6.4 Proton Beam Formation

This chapter has thus far examined the methodologies of beam acceleration and guidance in place at the LHC. In concert with such considerations, the means by which proton beams are focused and collimated must also be considered.

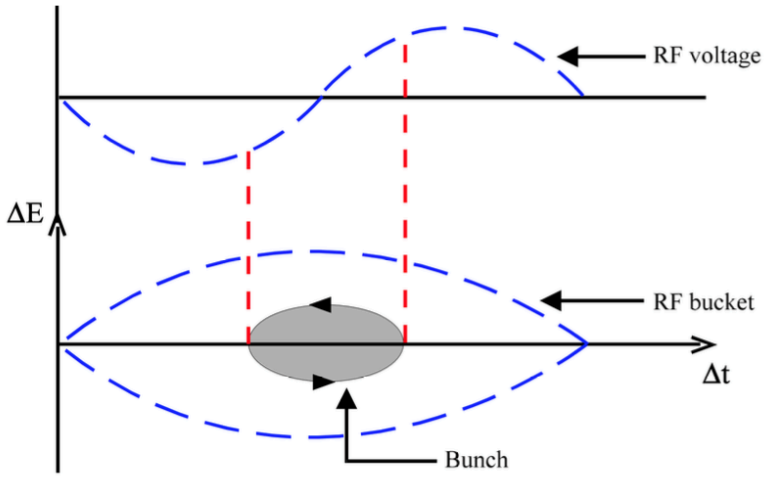

Figure 6.2 pictorial representation of the 'RF buckets' invoked within the RF cavities present within the LHC. Such buckets impose an envelope of allowable velocities within which protons are placed as they traverse the cavity. Image sourced from [104].

6.4.1 Beam Formation and Focusing

The manner in which RF cavities are employed to accelerate protons has been discussed. The RF cavities of the LHC, however, serve a further purpose: bringing protons together into collimated bunches. Protons which traverse a given RF cavity in perfect phase with the 400 MHz oscillation of the field receive no additional acceleration. Such protons are termed *synchronous protons*. Protons which arrive at the RF cavity out-of-phase with respect to the synchronous proton are either accelerated or decelerated to match the phase of the 400 MHz oscillation.

Such phase-selection behaviour of the RF cavities results in protons forming bunches around the synchronous protons. Bunched protons result in a higher probability of collision as they cross at experimental interaction points. Each proton bunch harbours on the order of 10^{11} protons.

The oscillatory patterns exhibited by the RF fields invoke an envelope of permissible positions in velocity-space in which protons must reside. Such envelopes are known as *RF buckets*. The negative-feedback behaviour of the envelope ensures that protons iteratively become successively more collimated with every pass through an RF cavity. This phenomenon is illustrated in Figure 6.2.

Having been brought into a bunch formation, the proton bunches must then be focused. Quadrupole magnets are employed in order to maintain a tight collimation of protons per bunch, thereby increasing the probability of collision per bunch crossing at the interaction points. The four poles of each magnet are arranged symmetrically around the beam pipe such that the separation of passing protons can be constrained on both the vertical and horizontal axis.

Additional sextupole, octupole, and decapole magnets are placed at the extremities of each dipole, correcting for small imperfections in the dipole field and any beam misalignment or stray particles which may otherwise arise from such imperfections.

Still further to the collimation provided by the quadrupole magnet systems, beam cleaning services are situated at points 3 and 7 of the LHC. Beam cleaning aims to mitigate the risk of LHC components sustaining radiative or thermal damage as a consequence of charged particles diverging from their prescribed path. This is achieved via devices known as *collimators*. Collimators absorb the energy of stray particles before they come into contact with sensitive LHC components, thereby circumventing the damage that would otherwise occur.

6.4.2 Beam-Dump Procedure

A given proton-fill circulating the LHC will become increasingly depleted and sparse as experimental collisions consume the particles it supplies. It will therefore become inefficient to continue operating with the current fill at some time-stamp during the run. At this point, the protons are ejected from the LHC through the execution of a procedure known as a 'beam-dump', after which the next supply of protons is injected into the ring in order to sustain the experiments with collisions more efficiently.

Depleted beams are extracted from the accelerator with devices known as 'pulsed extraction kickers' [105] and subsequently ejected. Few materials could withstand the colossal energy-densities of the ejected proton beams. Graphite was selected as the optimal material of choice for the construction of the 8 m long 'dump blocks' into which the beams are deposited. During run-2, LHC dump blocks were required to withstand temperatures in excess of 900° C (projected to double when the forthcoming High-Luminosity LHC commences operation) and the mechanical stress such temperatures entail. The trajectory of each ejected beam is deflected

by a number of 'dilution kickers' [105] such that an irregular pattern is swept over the surface of the dump block — this ensures that no single region of the dump block is subjected to an excessive density of energy. Numerous protective sub-systems are in place to prevent the beam-dump technology from sustaining damage, should the extracted proton beam stray off-course.

6.5 Luminosity and Pile-Up

In collider physics, the number of potential collisions which may be produced per unit time is quantified by *luminosity*, $\mathcal{L}$ [9]. This term forms the constant of proportionality which relates the number of interactions per second to the probabilistic cross-section of interaction, σ_{pp}, between the colliding particles:

$$\frac{\mathrm{d}N}{\mathrm{d}t} = \mathcal{L} \cdot \sigma_{pp}. \tag{6.1}$$

This quantity is a key metric by which the performance and physics-reach of a collider is gauged. It can be read from the differential equation above that the dimensionality of luminosity is measured in $\mathsf{cm}^{-2}\,\mathsf{s}^{-1}$. It is conventional in the field of high energy physics, however, to measure probabilistic cross-sections in dimensions of *inverse femto-barns*, $\mathsf{fb}^{-1} = 10^{-15}\,\mathsf{b}^{-1}$. The term 'barn', defined as $1\mathsf{b} = 10^{-28}\,\mathsf{m}^2 = 10^{-24}\,\mathsf{cm}^2$ has etymological roots in the Manhattan Project, where a code-name was required to obscure the nature of investigations into the cross-sections of various nuclear reactions. 'Barn' was chosen as a nod to the differences in scale pertinent to the objects of study, with the area of their collision target considered comparable in size to a barn relative to the nucleus of an atom. The term has pervaded in the field to this day.

Taking the particles within the beam to be distributed in accordance with a Gaussian distribution in the transverse plane, the luminosity is given by:

$$\mathcal{L} = \frac{N_1\ N_2\ n_\mathsf{b}\ f_\mathsf{rev}\ \gamma}{4\pi\ \epsilon_\mathsf{n}\ \beta^*} \cdot F, \tag{6.2}$$

where N_i is the number of particles contained within the bunch from beam i, n_b denotes the number of particle bunches per beam, f_rev denotes the frequency of revolution, and γ represents the relativistic Lorentz factor. The denominator is

composed of terms ϵ_n and β^* for the beam emittance and degree of collimation, respectively. The additional term, F, corrects for the angle at which the particle beams cross one-another.

The LHC is required to operate at a sufficiently high luminosity such that rare processes can be observed within a practical runtime. This is one of the factors which informed the decision to collide two proton beams (pp) rather than a proton and anti-proton ($p\overline{p}$) beam, given the engineering challenges associated with producing and sustaining a high-intensity beam of anti-protons.

The quantification of $\mathcal{L}$ discussed thus far provides a measure of luminosity at a given moment in time, known as the *instantaneous luminosity* [9, 3]. The total luminosity attained over a given time window — the *integrated luminosity* — is then simply given by the temporal integral of the instantaneous luminosity:

$$\mathcal{L}_{\text{int}} = \int_{t_0}^{t_{\text{total}}} \mathrm{d}t \; \mathcal{L}. \tag{6.3}$$

A total integrated luminosity of $\int \mathrm{d}t \; \mathcal{L} = 147 \; \text{fb}^{-1}$ was recorded by the ATLAS experiment over the course of the second run of the LHC, of which $\sim 140\,\text{fb}^{-1}$ was of sufficiently high quality to serve as the subject of physics analyses — see Figure 6.3. The peak instantaneous luminosity recorded by the experiment during this period was approximately double the design luminosity of the collider.

Each proton bunch accelerated by the LHC contains on the order of 100 billion protons. Consequentially, multiple pp-interactions typically occur at the interaction points for a given bunch crossing. The *hard-scatter vetex*, or *primary vertex* [9], is defined as the interaction vertex with the highest associated $\sum p_{\text{T}}^2$ and is taken to be the event with the greatest probability of producing physics of interest.

Interactions which occur outwith the collision associated with the primary vertex are termed *pile-up* interactions. These predominantly consist of soft inelastic scattering events, and thus give rise to less interesting phenomena. 'In-time' pile-up is the label ascribed to pile-up interactions which arise from the same bunch-crossing as the primary vertex; pile-up interactions which originate from previous or subsequent bunch crossings are termed 'out-of-time' pile-up.

Pile-up interactions represent a key challenge when conducting studies at hadron colliders. The additional energetic signatures such interactions register in detector

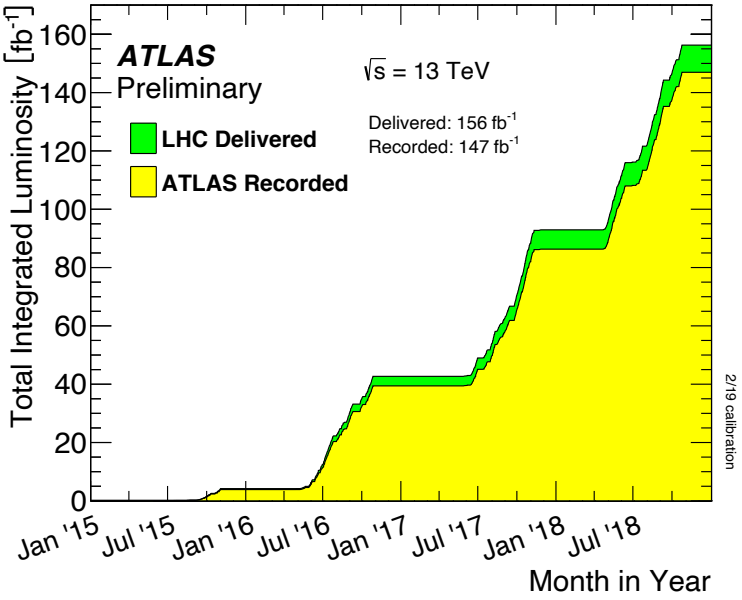

Figure 6.3 attained integrated luminosity over run-2. Integrated luminosity of collisions provided by the LHC is shown in green, with that which was successfully recorded by the ATLAS experiment shown in yellow. The difference between delivered and recorded luminosity can be attributed to detector down-time. Figure sourced from [106].

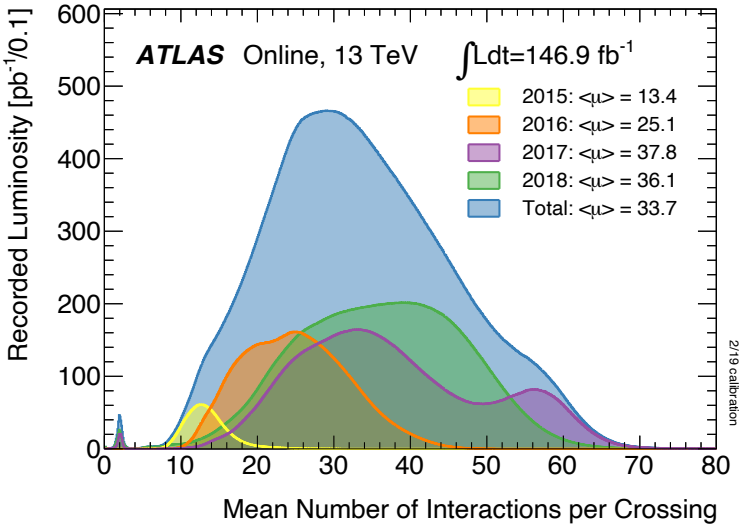

Figure 6.4 luminosity-weighted distributions of the mean number of interactions per bunch crossing for run-2 *pp*-collision data recorded at $\sqrt{s} = 13$ TeV at the LHC. Figure sourced from [106].

sub-systems must be distinguished from those deposited by primary events, such that efficient identification of physical events can be accomplished. As higher luminosities are pursued, the pressure imposed by the corresponding growth in the number of pile-up interactions will pose a yet-greater challenge on the experimental methods employed. The average number of interactions per bunch crossing can act as a useful metric with which to ascertain the number of pile-up interactions which must be accommodated — this quantity is shown for different periods of the second run of the LHC in Figure 6.4, weighted by the luminosity attained at the time of data-recording.

This chapter has examined the enabling technology which underpins the ability of the LHC to provide high energy physics experiments with proton-proton collisions. Before the data produced from such collisions can be analysed, the apparatus with which it is recorded must first be understood. A description of the ATLAS experiment now follows.

> *Big machines are awe-inspiring cathedrals of the 20th century.*

— Daniel Kleppner

7

The ATLAS Experiment

If big machines are the cathedrals of the 20th century, the ATLAS experiment is surely a 21st century Vatican. The largest of the four major LHC experiments, ATLAS stands at 25 m in height, spans 44 m in length, and weighs approximately 7,000 tonnes (broadly equivalent in weight to the Eiffel Tower). The ATLAS detector is comprised of multiple layers, each of which performing a specific purpose — taken all together, the ATLAS system is undoubtably one of the most complex scientific instruments ever constructed. This chapter examines each detector layer in turn, describing the technology which facilitates their operation and the systems which process the data they collect.

7.1 Overview

The methodologies by which protons are accelerated to ultra-relativistic energies has been discussed in the previous chapter. As protons are brought to collision at the various interaction points of the LHC, the various particulate by-products which are produced in the collisions must be measured — it is the job of general purpose particle detectors to measure the products of proton collisions such that an understanding of the underlying physics which governs their production and behaviour can be studied. Particle detectors represent some of the largest technological structures in existence: incredibly large machines are required to study the world of the unimaginably small, to probe reality at the finest possible resolution. The ATLAS experiment [95] is one such detector.

ATLAS (A Toroidal LHC ApparatuS) is a cylindrical general-purpose particle detector located at interaction point 1 of the LHC ring. Spanning $\sim 4\pi$ rad in solid-angle coverage around the beam-pipe, the ATLAS detector provides excellent acceptance of a wide variety of final states. As a general-purpose detector, ATLAS has enabled myriad measurements to be performed in a wide range of physics domains, from precision SM measurements [107] to searches for physics beyond the SM [108].

The detector takes the geometrical form of a central barrel region in which a number of sub-systems are housed, with further sub-detectors found in end-cap regions at each extremity of the barrel structure — see Figure 7.1. This achieves as wide an acceptance of final scatter states as possible. Key parameters concerning the operation of the ATLAS experiment are tabulated in Table 7.1.

Parameter	Run-1	Run-2	Run-3
Centre-of-mass energy $\sqrt{s}$ [TeV]	$7-8$	13	13.6
Min. bunch spacing [ns]	50	25	25
Recorded integrated luminosity [fb^{-1}]	5.08	147	183 (*expected:* 250)
Mean num. interactions / crossing $\langle\mu\rangle$	20	52	62

Table 7.1 ATLAS operational parameters and statistics for *pp*-collisions during each run period of the LHC.

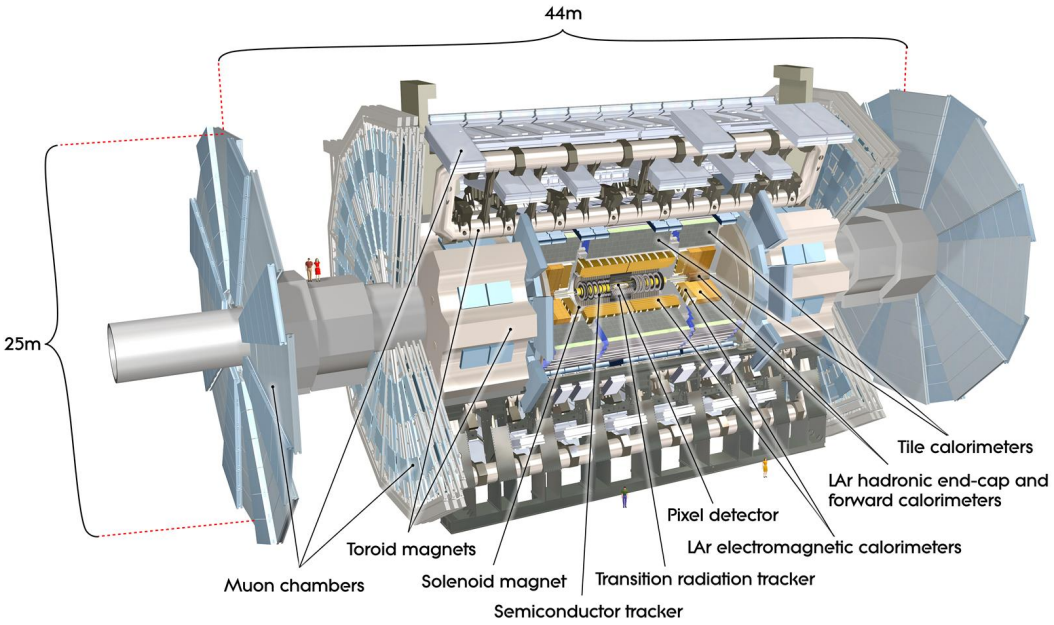

Figure 7.1 cut-away illustration of the ATLAS detector and a decomposition of major ATLAS sub-systems. Human figures are depicted for scale. Image sourced from [95].

7.1.1 Superconducting Magnets

The superconducting magnet system [109] of the ATLAS detector forms a visually-imposing and iconic component of the structure. The system can be decomposed into two sub-systems: the central solenoid and the toroidal magnet system. Both magnet systems are comprised of the same NbTi alloy used in the construction of the LHC magnet system, and are cooled to similarly low temperatures with liquid Helium. The central solenoid extends to 5.1 m in length with a diameter of 2.4 m, whereas the four superconducting toroidal magnets each extend to 22 m in diameter and 26 m in length. The large magnetic fields of the ATLAS experiment are used to deflect the trajectories of particles produced in pp-collisions to a degree proportional to the momentum of the particle.

The force experienced by a charged particle of charge q moving at velocity $\vec{v}$ is given by:

$$\vec{F} = q\,\vec{v} \wedge \vec{B}, \tag{7.1}$$

where $\vec{B}$ denotes the magnetic field vector. $\vec{F}$ acts in a direction perpendicular

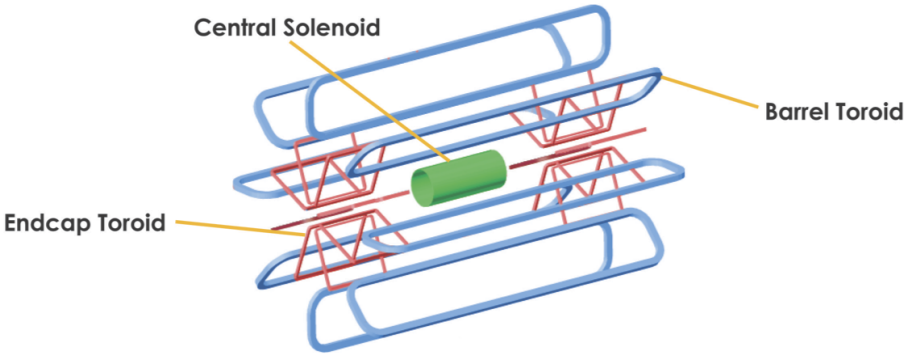

Figure 7.2 schematic of the superconducting magnet systems of the ATLAS experiment. Figure extracted from [110].

to the motion of the particle; when $\vec{B}$ is also mutually orthogonal to $\vec{F}$ and $\vec{\nu}$, the particle is deflected such that it pursues a circular trajectory. The radius of the trajectory curvature is related to the momentum of the deflected particle via $p = Bqr$ (where $B = |\vec{B}|$). The superconducting magnet system of the ATLAS detector therefore facilitates the measurement of charged particle momenta.

The geometrical configuration of the central solenoid system was chosen such that the radiative interference of the magnetic field with the electromagnetic calorimeter is minimised. The solenoid magnet system provides the inner detector with a 2 T axial magnetic field.

The toroidal magnet system is composed of one barrel region magnet and a complimentary magnet system in each end-cap, as illustrated in Figure 7.2. The barrel toroid extends to 25.3 m in length at an inner (outer) diameter of 9.4 m (20.1 m). The end-cap toroids are responsible for providing the muonic spectrometer with a field of 0.5 T (1.0 T) via 8 air-core coils per toroid. This allows the momenta of muons incident with the spectrometer to be determined. This is achieved via toroids of 10.7 m in diameter and 5 m in axial length per end-cap magnet.

7.1.2 ATLAS Coordinate System

A right-handed Cartesian coordinate system is defined for the ATLAS detector, where the x-axis points towards the centre of the LHC ring, the y-axis vertically upwards, and the z-axis along the beam-pipe. The origin of this coordinate system is defined as the centre of the detector, as shown in Figure 7.3. Owing to

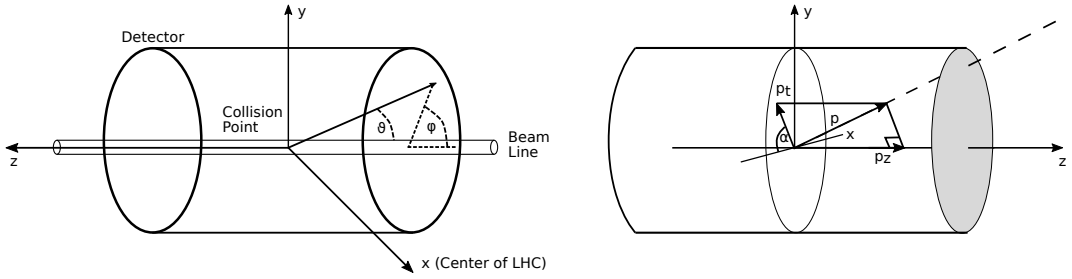

Figure 7.3 illustration of the ATLAS coordinate system in the *(left)* nominal cylindrical basis and *(right)* transverse momentum plane. The (η, ϕ)-plane can be visualised as the resultant plane obtained having unfurled the curved barrel volume of the detector into a flat sheet. Figure extracted from [95].

the approximate rotational symmetry of the detector around the beam-pipe, it is often convenient to employ a cylindrical coordinate frame to describe particle phenomena in the transverse plane. In this cylindrical system, the azimuthal angle, ϕ, is taken to be the angular separation with the x-axis in the transverse plane. The polar angle is denoted by θ in adherence with convention. The cylindrical coordinate system is also depicted in Figure 7.3.

The metric of *rapidity* [3] provides a relativistic measure of velocity in the longitudinal direction, and exhibits the attractive quality of Lorentz-invariance with regards to linear transformations along the beam-axis. Given that interacting partons are initially of unknown momentum, the Lorentz boost experienced along the z-direction cannot be known *a priori*. For this reason, rapidity is a convenient metric in collider physics. The quantity of *pseudo-rapidity* [3] provides an accurate and easily computed approximation of rapidity in the ultra-relativistic regime where $\beta \to 1$, defined as $\eta = -\log_e\left(\tan\frac{\theta}{2}\right)$. This quantity shall be invoked routinely hereinafter. Equipped with the afore-given definitions, the geometric distance metric used to capture particle separation, ΔR, may now be defined as $\Delta R = \sqrt{(\Delta\phi)^2 + (\Delta\eta)^2}$.

7.1.3 Design Criteria

In order to achieve the scientific objectives which motivated its construction, the ATLAS detector is required to exhibit the following qualities.

- **Radiation-hardened electronics:** the large particle fluxes and energies

which constitute the LHC collision environment impose a high toll which the sensor and read-out systems must withstand.

- **Responsive electronics:** the rate at which recorded data can be read out from sensor systems must keep pace with the challenging event-rate of the LHC.
- **Large η and ϕ coverage:** high acceptance of a wide variety of final state geometries is necessary for efficient and effective particle reconstruction.
- **Excellent electromagnetic resolution:** efficient detection of electromagnetic showers produced by electrons, positrons, and photons requires excellent electromagnetic calorimetry capabilities.
- **Excellent hadronic resolution:** efficient reconstruction of hadronic processes and jets requires excellent hadronic calorimetry capabilities.
- **Strong tracking capabilities:** a tracking system with good momentum resolution is essential to efficiently reconstruct the trajectories of charged particles.
- **Muon Identification capabilities:** efficient muon identification with strong momentum resolution is required for many analyses, and provides additional input for trigger systems.
- **Efficient trigger system:** the large rate of collisions at the interaction points of the LHC necessitates the efficient reduction of data-rates without discarding events of interest.

The afore-listed specifications were key considerations which informed the design of the ATLAS detector and its various sub-systems. A description of each ATLAS sub-system and an examination of the degree to which they satisfy such requirements now follows.

7.2 Detector Sub-Assemblies

7.2.1 Inner Detector

With the closest proximity to the beam-pipe, the Inner Detector (ID) [111, 95] serves to identify particle tracks emanating from primary vertices. The ID

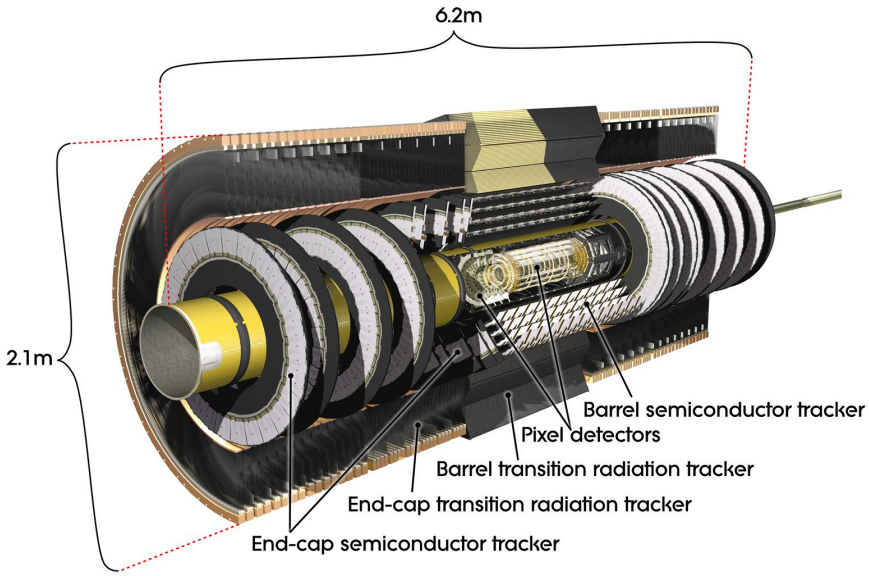

Figure 7.4 cut-away render of the ATLAS Inner Detector and its sub-systems. Figure sourced from [95].

is comprised of a set of distinct modules each in a different spatial location, such that particle 'hits' can be recorded in a given location as charged particles deposit energy in the modules they traverse. Particle hits may subsequently be reconstructed into tracks, allowing the trajectories of the particles produced in pp-collisions to be obtained. Tracking information enables the momentum of a passing charged particle, in conjunction with the sign of its charge, to be determined. The extraction of secondary and tertiary vertex information is also facilitated by the ID.

The ID stands at a total combined length of 6.2 m with a diameter of 2.1 m. Coverage of $|\eta| < 2.5$ is achieved for particles whose transverse momenta is in excess of $p_{\mathrm{T}} > 0.5$ GeV. The system is immersed in the 2 T axial magnetic field produced by the superconducting solenoid. The azimuthal deflection of charged particles allows the momentum and charge of charged particles to be measured, as described in Section 7.1.1.

ID Resolution: A Note on Impact Parameter Measurement

The ID boasts a design momentum resolution of $\sigma_{p_{\mathrm{T}}}/p_{\mathrm{T}} = 0.05\%\ p_{\mathrm{T}}$ [GeV] [3]. The resolution of the ID can alternatively be parameterised in terms of the longitudinal impact parameter, z_0, and the transverse impact parameter, d_0. The value of d_0 can be considered as the distance of closest approach of a given

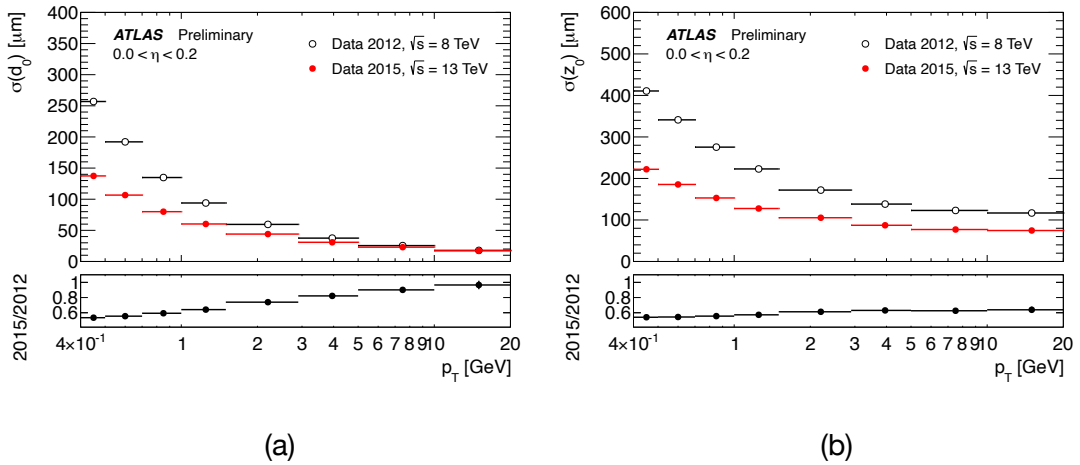

Figure 7.5 unfolded (a) transverse and (b) longitudinal impact parameter resolution as recorded in run-2 (red, filled) and run-1 (black, hollow) data as a function of transverse momentum. Lower sub-plots show the ratio of the run-2 / run-1 data-sets. Source: [112].

charged particle to the beam-pipe. The impact parameter resolution [3] can be written as:

$$\sigma(z_0) = 12 \oplus \frac{88}{p_T\sqrt{\sin\theta}}\ [\mu\text{m}], \tag{7.2}$$

$$\sigma(d_0) = 95 \oplus \frac{160}{p_T\sqrt{\sin^3\theta}}\ [\mu\text{m}]. \tag{7.3}$$

The searched for displaced τ-leptons conducted in this thesis will rely heavily on the measurement of d_0, which will serve as a proxy for the lifetime of the parent BSM particle responsible for any displaced signature which may be observed.

The ATLAS inner detector consists of three sub-systems: the *silicon 'pixel' detector*, the *silicon-strip semiconductor tracker* (SCT), and the *transition radiation tracker* (TRT), as shown in Figure 7.4. Each of these sub-systems is now considered in turn.

Pixel System

The pixel sub-system [95] is the inner-most layer of the inner detector and provides coverage of $|\eta| < 2.5$. The pixel system is composed of 1,744 modules, within which n-type silicon semiconductor material is employed as the sensing

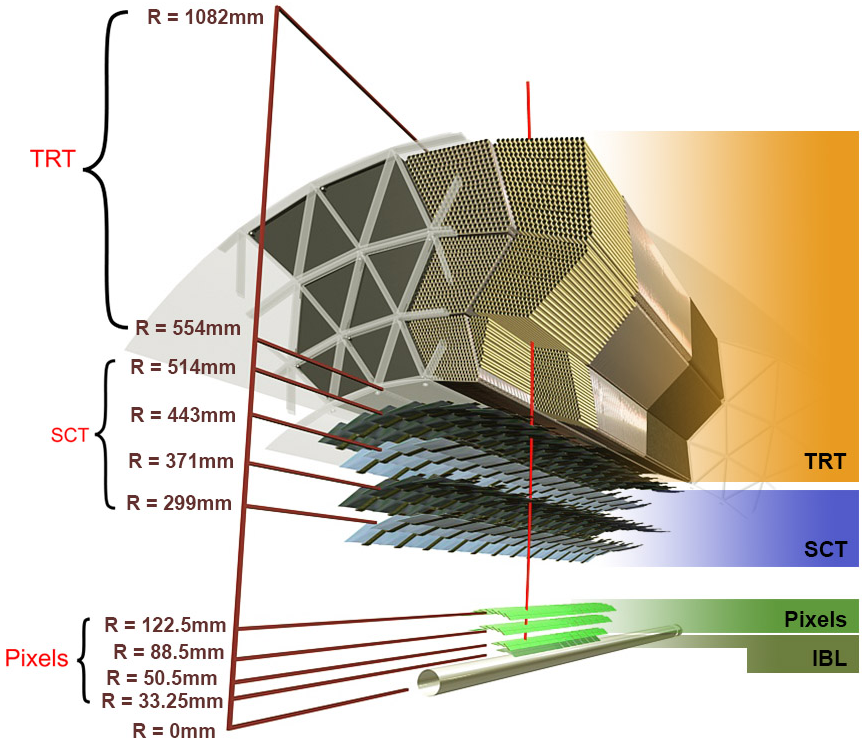

Figure 7.6 depiction of the barrel structure of the ATLAS Inner Detector with the Insertable B-Layer. Figure sourced from [95].

component. Incident charged particles excite electrons within the n-type doped region into the conducting band, resulting in 'holes' in the valence band of the semiconductor [113]. Liberated electrons travel towards dedicated electronic read-out systems under an applied electric field, from which a hit signal can be recorded.

As is the case for most ATLAS sub-systems, the pixel detector has distinct assemblies within the barrel region of the detector and each of the detector end-caps. Within the barrel region, pixel modules are distributed across three concentric layers, each of which stretching to 4 m in length and positioned at 50.5 mm, 88.5 mm, and 122.5 mm radial distance from the centre of the beam-pipe. This is illustrated in Figure 7.6. Modules within the end-caps are organised into disk formations, with three disks situated at $|z| = 495$ mm, $|z| = 580$ mm, and $|z| = 650$ mm per end-cap.

Each of the 1,744 modules in the pixel system is home to 47,232 sensing elements known as *pixels*. Each pixel spans an area of 50 μm × 400 μm, and achieves a spatial resolution of 10 μm in the (x, y)-plane and 115 μm in the z-plane. It is therefore possible to extract a three-dimensional position of each hit encountered

along the trajectory of a charged particle passing through the inner detector. The three layers of the pixel system in the barrel region harbour $1,456$ modules, corresponding to 67 million pixels, whereas the end-cap systems collectively contain 288 modules with 13 million pixels.

Before the commencement of run-2, an additional layer of the pixel system was installed as the inner-most system of the inner detector. At a radial distance of $\sim 33\ \text{mm}$ from the beam-pipe (as shown in Figure 7.6), the *insertable B-layer* (IBL) significantly improves the resolution at which impact parameter measurements can be made. The size of modules within the IBL are smaller than those found in other components of the pixel system, with an area of $115 \times 10\ \mu\text{m}^2$ — this leads to a finer measurement resolution in the z-direction, while the additional hit points closer to the beam-pipe improve the resolution of low-p_T particle tracking in the (r, ϕ)-plane. Altogether, this results in a greater capacity to accurately reconstruct secondary vertices — something of particular import with regards to the identification of b-hadrons, hence the name of the component.

Semiconductor Tracker

Moving radially outward, the next layer of the inner detector is the *semiconductor tracker* (SCT) [95]. The SCT collects particle hits at a greater radial displacement than the pixel system via $4,088$ silicon micro-strip modules. Each module features two sensor layers tilted at an angle to $40\ \text{mrad}$ with respect to one-another. This allows two-dimensional particle hit reconstruction to be achieved in the plane of the module.

Within the barrel region of the detector, SCT modules are arranged into four coaxial cylindrical layers. Each of the $64.0 \times 63.6\ \text{mm}^2$ sensors are mounted at a pitch of $80\ \mu\text{m}$ and protrude at a thickness of $285\ \mu\text{m}$. Sensors are organised into 768 strips of silicon wafers, which constitute the tiles arranged into the cylindrical geometry of the barrel region of the SCT system. SCT modules within the end-caps are trapezoidal in shape and are organised into 9 disks per end-cap. SCT layers within the barrel region are situated at radial distance $299\ \text{mm}$, $371\ \text{mm}$, $443\ \text{mm}$, and $514\ \text{mm}$ from the beam-pipe, whereas layers within the end-cap sectors range in position from $|z| = 854\ \text{mm}$ to $|z| = 2,720\ \text{mm}$. A total of 6.3 million read-out channels are present in the SCT system.

Transition Radiation Tracker

The outer-most layer of the inner detector is the *transition radiation tracker* (TRT) [95], which provides coverage up to $|\eta| < 2.0$. Standing at variance with the pixel and SCT systems, the TRT is a gaseous detector which consists of components known as *straw tubes*. Each straw tube within the TRT system contains an admixture of $70\% : 27\% : 3\%$ $\mathsf{Xe} : \mathsf{CO_2} : \mathsf{O_2}$, which is ionised by incident charged particles. A grounded gold-plated Tungsten wire is present in the centre of each straw tube, whose walls are held at a potential of -1.5 kV. Electrons liberated from the gas by incident particles therefore drift towards the central wire, inducing an electrical signal which can subsequently be read-out and interpreted as a particle hit.

The TRT system within the ATLAS barrel region contains $52,544$ straw tubes which collectively cover $|\eta| < 1.0$. Each straw tube extends to 144 cm in length parallel to the beam-pipe, collectively projecting from $R = 544$ mm to $R = 1,082$ mm radially from the centre of the detector. Straw tubes within the end-cap regions of the detector provide coverage of the area $1.0 < |\eta| < 2.0$ and $0.8\ \mathsf{m} < |z| < 2.7\ \mathsf{m}$. The detector end-caps collectively contain a total of $122,880$ straw tubes, each of which extending to 39 cm in length.

Each straw tube is interlaced with polypropylene fibres which induce transition radiation (to which this component owes its name) as charged particle pass through. Photons emitted in this process are absorbed by the Xe gas within the tube, producing a stronger readout signal which can further discriminate between electrons and pions. Due to an irreparable gas leak which occurred in the system, some Xe gas was replaced with Argon for run-2 — the ability to detect transition radiation was therefore somewhat diminished by the lower absorption rate of Argon in comparison to Xenon.

7.2.2 Calorimetry

Calorimeter (think 'calorie-meter') systems are used to measure particle energy. Particles of a given species are brought to rest by the material of the appropriate calorimeter, thereby depositing their kinetic energy into the detector. This energy is subsequently measured by the electronic systems of the calorimeter and used to inform the particle identification process.

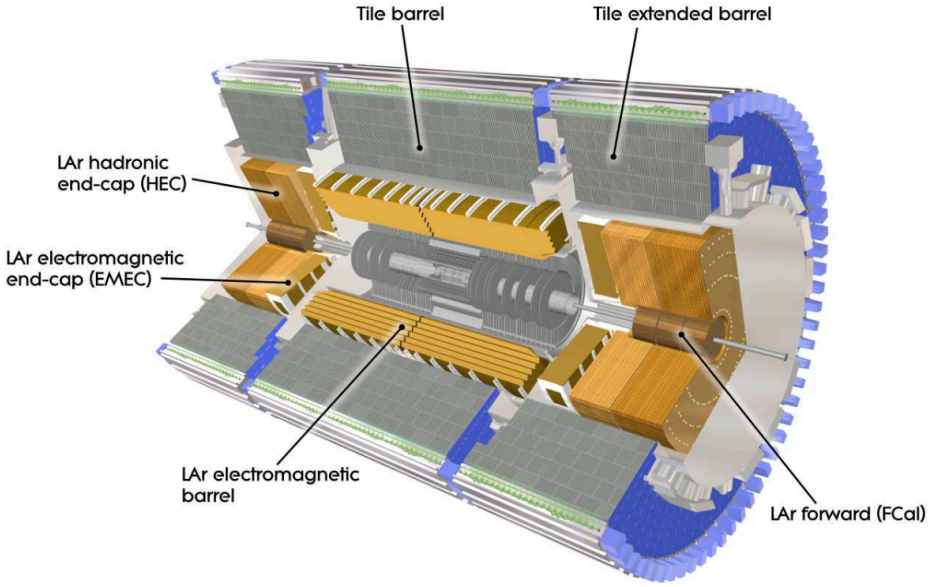

Figure 7.7 cut-away view of the ATLAS calorimetry system in the barrel region of the detector, with annotations for the various calorimeter sub-systems. Image extracted from [95].

Beyond the measurement of particle energy, calorimeters provide additional utility through the measurement of particle geometrical position within the detector, assisting with the computation of missing transverse energy (see subsequent chapter), and providing input to the detector trigger systems.

Given the differing interaction properties of elementary particles and the wide geometrical area which they may traverse, various calorimeter systems feature in the ATLAS detector. The key elements of the ATLAS calorimetry system are depicted in Figure 7.7, each of which is now examined in turn.

Electromagnetic Calorimeter

The electromagnetic calorimeter (ECAL) [114] is the first layer of the detector a particle will encounter as it exits the inner tracking detector. Electrons, positrons, and photons are impeded by the ECAL, depositing their energy in the layers of the system.

As an incoming electromagnetically-interacting particle interacts with the detector material of the ECAL, it produces further, less energetic particles which subsequently interact with the detector themselves. The iteration of this process results in a shower of particles known as an *electromagnetic shower*. It is the

energy deposited by the shower within the calorimeter that is measured, serving as a proxy for the energy of the parent particle which initiated the cascade. High energy electrons and positrons experience energy loss chiefly due to the emission of Bremsstrahlung radiation, whereas the production of e^+e^- pairs is largely responsible for the energy losses of high energy photons.

The expectation value for the energy $E(x)$ of an electron at penetration depth x within the detector material is given by:

$$\langle E(x) \rangle = E_0 \, e^{-\frac{x}{X_0}}, \tag{7.4}$$

where E_0 is the incident energy of the electron and X_0 is the radiation length of the material being traversed. In a similar fashion, the expectation value for the intensity of a photon beam is given by:

$$\langle I(x) \rangle = E_0 \, e^{-\frac{7}{9}\frac{x}{X_0}}, \tag{7.5}$$

with the multiplicative coefficient of $-7/9$ determined empirically to account for the mean free path pursued by a photon before undergoing pair conversion.

The ATLAS ECAL is a *sampling calorimeter*, where different materials are used to induce the showering process and to measure the energy deposited by the shower. The ECAL is thus comprised of alternating layers of dense, passive material with which to impede the traversing particles and layers of sampling material. The ATLAS ECAL is comprised of a mix of lead and stainless steel in its passive layers, while Liquid Argon (LAr) is used as the active sampling material in which kapton electrodes are submersed.

The geometrical form assumed by the ATLAS ECAL can be decomposed into a barrel region and two end-cap regions. The barrel region of the ECAL comprises two coaxial cylinders separated by a small 6 mm gap at $z = 0$, situated behind the superconducting solenoid. Each at 3.2 m in length, the barrel cylinders have an inner (outer) radius of 1.4 m (2 m). Collectively, they provide calorimetry coverage within $|\eta| < 1.475$. The shape of the active and absorptive material in this region is reminiscent of an accordion, thereby achieving complete ϕ-symmetry with no gaps in azimuthal acceptance.

Each of the two end-cap ECAL systems are composed of two coaxial wheels,

where each wheel has a thickness of 63 cm and an inner (outer) radius of 330 mm (2098 mm). End-cap ECAL systems provide strong energy measurements of particles which exit the detector at a trajectory with small angular deviation from the beam-pipe: the outer wheel provides coverage of the region $1.375 < |\eta| < 2.5$, while the inner wheel has an acceptance region of $2.5 < |\eta| < 3.2$.

Active sampling is accomplished by placing a large voltage across the sampling electrodes — ionisation charges produced by incident showers in the LAr are thereby accelerated towards the electrodes, inducing a current which can be amplified and read as an electronic signal proportional in strength to the energy deposited by the shower. The read-out system is split into three radial layers at varying depths such that shower shape information may be determined.

The active LAr system is composed of modules, with each module comprising three layers. The first layer of each module is finely segmented in η with a resolution of $\Delta\eta \times \Delta\phi = 0.003 \times 0.1$. This layer is $\sim 4X_0$ thick and largely serves to provide positional measurements of photons which, if unconverted, do not leave a track in the ID. The second layer of each LAr module is $\sim 16X_0$ thick, with a resolution of $\Delta\eta \times \Delta\phi = 0.025 \times 0.025$. This layer captures the majority of the energy deposited by an incident electromagnetic shower. With a thickness of $\sim 2X_0$, the third and final layer of each module provides a coarser resolution of $\Delta\eta \times \Delta\phi = 0.05 \times 0.025$. This provides a means by which to capture the tail of a given shower, while additionally helping to ascertain the extent of energy leakage which may be present. A schematic of an ECAL LAr module is presented in Figure 7.8. In total, 173,312 read-out channels from such modules are present in the system.

A gap in coverage is present at $|\eta| = 1.5$ as a consequence of the boundary between the barrel region and end-cap regions of the ECAL. An additional gap in coverage arises at $z = 0$ due to the 6 mm gap between the two identical cylindrical blocks which constitute the barrel region of the calorimeter. Limitations in coverage must be accounted for in the data quality and selection strategy at analysis level in order to exclude unreliable regions of the detector.

As a given particle traverses through the inner detector towards the calorimeters it will encounter detector material in its path, absorbing a fraction of its energy. The total material budget encountered by a particle will increase proportionally with η as a consequence of the angle of the trajectory pursued. For this reason, a pre-sampler is installed before the ECAL at various locations within the barrel

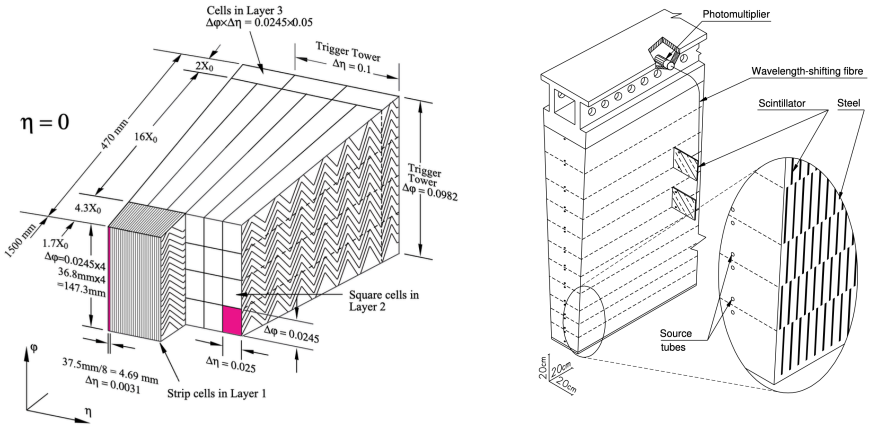

Figure 7.8 shows *(left)* an ECAL barrel module annotated to depict the cell structure and the accordion-like geometry; *(right)* a tile calorimeter module, complete with alternating active and absorptive material and optical read-out. Image taken from [115].

and end-cap regions, providing an additional energy measurement with which to correct for energy losses incurred *en route* towards the calorimeter systems.

The precision with which the four-momenta of incoming particles may be determined by the calorimeter is determined by the *energy resolution*. Many factors can perturb the measured value from its true physical value, including detector noise, the quantum fluctuations which manifest as photon $\leftrightarrow$ electron conversions, and natural statistical fluctuations. Further, interactions with the sampling material can induce additional showering interactions which 'leak' into subsequent layers of the system and skew the measurement. For such reasons, the shower energy measured by calorimetry systems can be observed to fluctuate around a central true value. The energy resolution can be obtained via:

$$\frac{\sigma(E)}{E} = \frac{a}{\sqrt{E}} \oplus \frac{b}{E} \oplus c\,\%, \tag{7.6}$$

where a denotes the *stochastic term*, quantifying the contribution from fluctuations in signal-generating processes, b accounts for the contributions arising from both electronic noise and pile-up interactions, and c provides a metric by which to account for imperfections in the construction of the calorimeter system — these may include non-uniformity of detector response, errors in calibration between detector read-out channels, and obstructions to the measurement of the shower (arising, for instance, from other material in the detector). In the above equation, $\oplus$ denotes the sum in quadrature operation: $a \oplus b = \sqrt{a^2 + b^2}$.

Hadronic Calorimeter

Hadronic cascades are measured by the hadronic calorimeter (HCAL) [116, 95], which resides immediately behind the ECAL assembly. The strong force gives rise to more complex interactions between the hadronic shower constituents and the detector material than is exhibited by electromagnetic interactions, with hadron showers typically penetrating appreciably deeper into the calorimeter. Additionally, the fraction of shower energy which can be measured is lower than that associated with electromagnetic cascades, which *perforce* leads to a lower energy resolution than is attainable with the ECAL.

Hadronic showers are characterised by the nuclear interaction length, λ. This parameter is typically an order of magnitude greater than the radiation length, X_0, associated with electromagnetic showers — hadronic cascades are therefore longer than their electromagnetic counterparts and penetrate a greater distance into the calorimeter. The depth required to contain a given energy fraction of a given shower is also logarithmically dependent on the energy of the incident particle. Consequentially, hadronic calorimeters must typically be greater in size than those tasked with electromagnetic measurements. The sub-systems of the ATLAS HCAL are designed accordingly, with a maximum depth of $\sim 10\,\lambda$.

The barrel component of the HCAL is known as the *tile calorimeter* [95]. A central 5.8 m long barrel and two 2.6 m extended barrels at each side comprise the tile system, with each segment sharing a respective inner and outer radius of $\sim$ 2.3 m and $\sim$ 4.3 m. The central barrel provides coverage of the region within $|\eta| < 1.0$, while the extended barrels cover the region within $0.8 < |\eta| < 1.7$. In common with its electromagnetic counterpart, the tile calorimeter is a sampling calorimeter. Scintillating tiles house the active material, with steel plates serving as the absorbing material — 64 modules in total, containing an active and absorptive component, are present in the tile system. Plastic scintillator was chosen to act as the active material of each tile. The tile calorimeter is segmented into three layers of varying depths and decreasing granularity of measurement, in a further similarity with the ECAL system. As previously alluded to, the cells which comprise the layers of the HCAL are, however, substantially larger than those found in the ECAL, at $\Delta\phi \times \Delta\eta \approx 0.1 \times 0.1$.

The LAr hadronic end-cap calorimeter is situated directly behind the end-cap detectors of the ECAL, operating in a similar manner. A front and rear wheel constitute each end-cap. Wedge-shaped modules of copper absorption material

are interwoven with active LAr gaps, with a total of 32 units of each component installed. Read-out units of size $\Delta\phi \times \Delta\eta = 0.1 \times 0.1$ are installed in the central region of the hadronic end-caps, decreasing in granularity at larger radial distance from the centre. Coverage extends to $1.5 < |\eta| < 3.2$, minimising 'cracks' in acceptance by overlapping with other HCAL components.

A high density LAr *forward calorimeter* (FCAL) extends hadronic coverage to $3.1 < |\eta| < 4.9$. The FCAL is composed of three separate layers, each extending to 45 cm in depth and employing LAr as the active medium. The passive material chosen differs between the layers of the FCAL: the first layer utilises Copper and is optimised for the measurement of electromagnetic interactions arising from the hadronic shower, whereas the subsequent layers employ Tungsten and are predominantly optimised for the measurement of the energy of hadronic interactions between the shower constituents and the HCAL.

7.2.3 Muonic Spectrometer

With the notable exception of neutrinos, muons are the only SM particle to traverse the afore-discussed layers of the ATLAS detector unimpeded. It is for this reason that the Muonic Spectrometer (MS) [95] is situated at the outer-most layer of the ATLAS detector.

As with the afore-discussed components of the ATLAS experiment, the MS is comprised of multiple sub-systems — as may be observed in Figure 7.9. An examination of such sub-systems now follows.

Tracking Chambers

The dedicated MS serves primarily to provide accurate four-momenta measurements of incoming muons, while also providing additional information with which to assist the ATLAS trigger system. Accurate four-momenta measurements are provided for the majority of the pseudo-rapidity coverage by the *monitored drift tubes* (MDT): aluminium tubes encircling Tungsten-Rhenium (W-Re) wires, arranged into chambers (known as 'muon chambers'). MDTs are arranged in three cylindrical layers in the barrel region, extending radially to 5 m, 7.5 m, and 10 m, with a further three end-cap layers, located at $|z| = 7.4$ m, $|z| = 14.0$ m, and $|z| = 21.5$ m. Within $|\eta| < 2.0$, muon tracks are measured by MDT layers and the

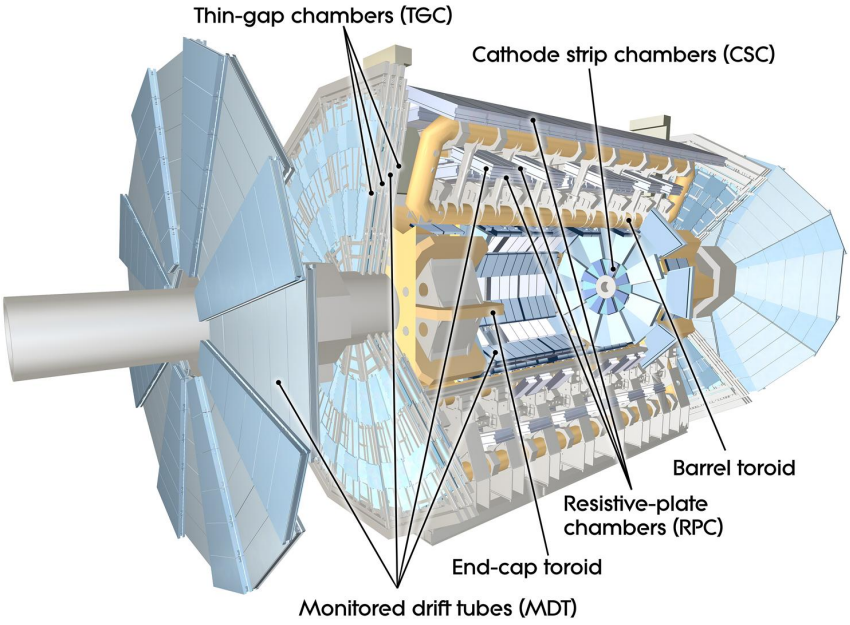

Figure 7.9 cut-away view of the ATLAS muon spectrometer, showing the location of each sub-component of the system. Image sourced from [95].

four-momenta of the associated muon determined through appeal to the sagitta of the deflected track [117].

Each MDT extends to 30 mm in diameter and contains a gaseous admixture of 93% Ar and 7% CO_2. The W-Re anode wire within each MDT is held at a 3 kV potential — ions produced within the gaseous admixture by passing charged particles will therefore experience an acceleration towards the anode, from which an electronic signal can be read. The duration of time taken for a given ionic charge to drift towards the anode is measured, providing a metric with which to compute the displacement of the track from the centre of the MDT — thereby providing spatial tracking information. A schematic depiction of the MDT system is presented in Figure 7.10. Altogether, the MS contains 1,150 muon chambers and 354,000 MDTs.

Tracking capabilities in the large pseudo-rapidity regime of $2.0 < |\eta| < 2.7$ are provided by *cathode strip chambers* (CSC). Particle occupancy in this region is greater than the design capacity of the MDT chambers; CSCs were developed to withstand the conditions of this environment. Each CSC contains two cathode planes segmented into strips, placed in an orthogonal orientation with respect to one-another such that a two-dimensional read-out is facilitated. In contrast to the single anode wire present in each MDT, CSCs employ multi-wire anodes with

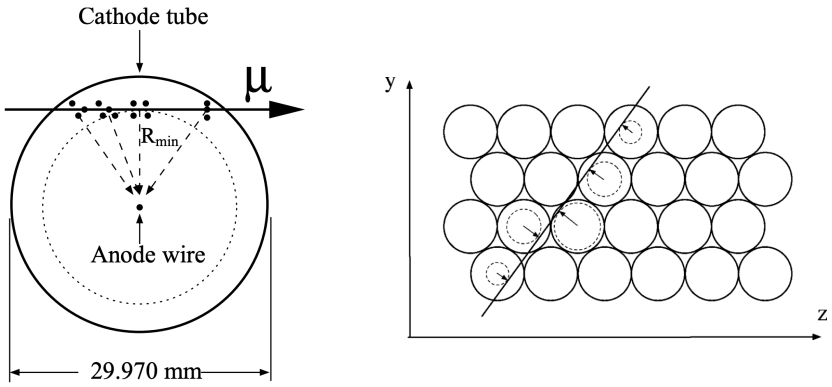

Figure 7.10 *(left)* [118] shows the cross-sectional view of an MDT tube, where ionisation is induced by an incident muon; *(right)* [119] depicts the view of a MDT multi-layer in the (y, z)-plane, with a track fit shown across multiple MDTs.

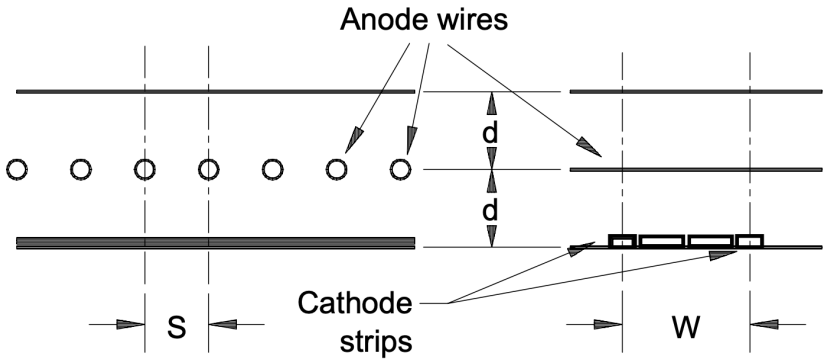

Figure 7.11 diagramatic depiction of an ATLAS cathode strip chamber, 32 of which comprise the full ATLAS CSC system of the muon spectrometer. Wire separation, s, and the cathode-anode separation, d, are both equal to 2.54 mm. W denotes the cathode read-out pitch, which stands at 5.08 mm Image extracted from [120].

which to collect accelerated charges, held at a potential of 1.9 kV. The gaseous volume within each unit is a $80\% : 20\%$ combination of $Ar : CO_2$, acting as the ion source. A total of $31,000$ read-out channels are present in the CSC system across 32 chambers.

Trigger Chambers

Dedicated trigger chambers [120] are installed in the space surrounding the monitored drift tubes. These serve to provide fast signalling which aides the performance of the muon trigger system and assist the muon tracking process

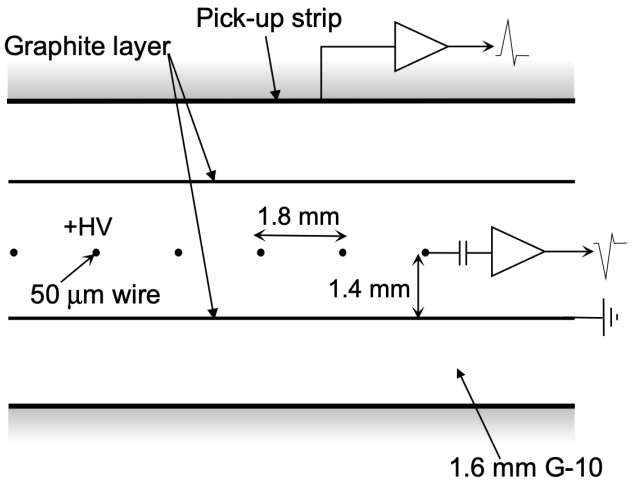

Figure 7.12 schematic depiction of an ATLAS thin gap chamber, 3,588 of which may be found in the ATLAS detector. Image taken from [120].

to attain high precision measurements. Muon trigger chambers come in two varieties: *resistive plate chambers* (RPC) and *thin gap chambers* (TGC).

Resistive plate chambers provide trigger signalling within the range $|\eta| < 1.05$. Each RPC is comprised of two parallel electrode plates, between which a potential of $9.8\ \mathsf{kV}$ is applied. The volume between each pair of resistive plates is filled with Tetrafluorethane ($C_2H_2F_4$) gas which ionises in response to an incident muon, providing a signal in (z, ϕ)-space which is relayed by dedicated read-out strips. This achieves a trigger timing resolution of $2\ \mathsf{ns}$.

In the forward regions, coverage of $1.05 < |\eta| < 2.4$ is provided by thin gap chambers. TGCs are multi-wire proportional chambers and operate in a similar capacity to the CSCs, with an internal potential of $2.9\ \mathsf{kV}$ across the anode and a combination of $55\% : 45\%\ CO_2 : n-C_5H_{12}$ providing the gaseous ion source. The acceleration provided by the potential and the geometry of the chamber design together enable the TGC system to achieve a timing resolution of $4\ \mathsf{ns}$.

The combination of the afore-introduced sub-systems enables the muonic spectrometer to attain a momentum resolution of $\sigma_{p_T}/p_T = 10\%$ at $p_T = 1\ \mathsf{TeV}$ [121]. Tracking information from the inner detector may be taken in conjunction with MS data to provide improved momentum resolution for low-p_T muons — this will be investigated further in the chapter which follows.

7.2.4 Trigger System

The 40 MHz bunch-crossing rate of the LHC, corresponding to an average of 1-billion pp-collisions per second, poses a sizeable challenge for the data acquisition and storage strategy of the ATLAS experiment. The volume of data collected from the raw detector read-out at such a collision rate exceeds that which is able to be practically and affordably stored for analysis. The majority of collisions produced at the LHC do not give rise to physical phenomena of interest — it is therefore possible to filter and selectively store events of interest for subsequent analysis without detriment to the scientific investigations performed on the reduced data-set.

The task of identifying physical events of interest such that they can be stored for analysis falls to the ATLAS trigger system [122]. The trigger system is required to reduce the raw event-rate of 40 MHz to a stream of selected events at a rate of 100 Hz — crucially, this must be accomplished without discarding rare physical processes of interest to the frontier of fundamental physics research. The system responsible for this task can be decomposed into two components.

Level-1 Trigger

The initial stage of the triggering system is the *Level-1 trigger*. This is a hardware-level system which consists of a *central trigger processor*, largely tasked with processing input from the calorimeter systems and the muonic spectrometer. At this stage of the trigger system the event-rate is reduced from 40 MHz to $\sim$ 100 kHz.

The central trigger processor is used to crudely define *regions of interest* (RoI): areas in (η, ϕ)-space which 'seed' subsequent HLT processing. These are used to isolate the location of likely high-p_{T} objects, including muons, electromagnetic clusters, jets, and τ-leptons.

The time taken from the original proton collision until level-1 trigger information is made available to the wider system must not exceed 2.5 µs. Prior to processing, detector read-out information is stored in 'pipe-line' memory located close to the detector. Pipe-line memory systems are radiation hardened such that they can withstand the radiation environment of the immediate detector surroundings.

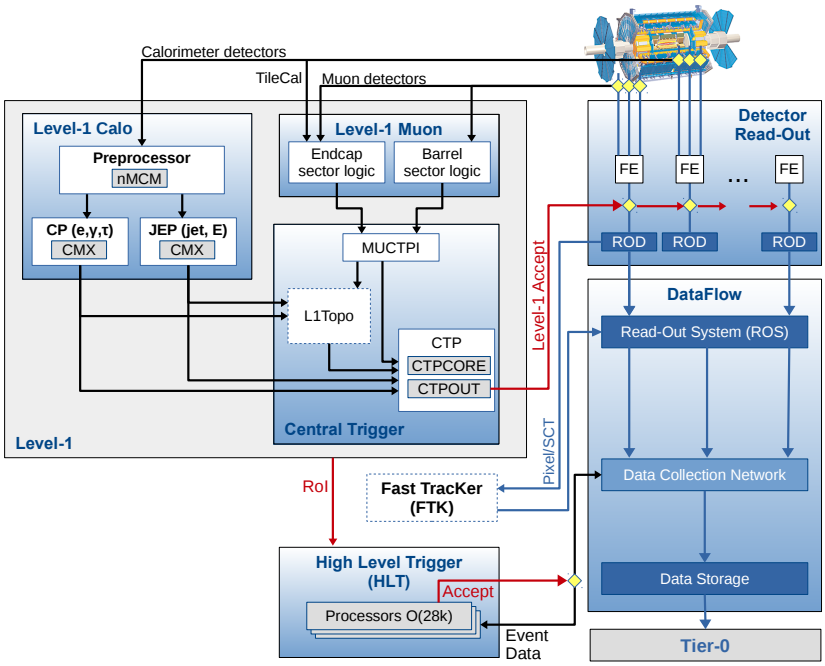

Figure 7.13 shcematic overview of the ATLAS trigger system, depicting the dataflow from low-level ATLAS components to the level-1 and HLT trigger components. 'Tier-0' refers to the CERN compute farm where accepted event data are stored for reconstruction. The 'FTK' system depicted above was not active during the data-taking process relevant to this thesis. Image taken from [123].

High-Level Trigger

Once the initial data-rate is reduced by the level-1 system, it is fed to the *high-level trigger* (HLT). The HLT system executes a range of offline identification algorithms on a dedicated processor farm to perform a refined filtration of the level-1 output within 300 ms. Trigger identification processes are initially run within the level-1 RoIs, before a fuller event-wide view is considered. The event-rate is further reduced by the HLT to ~ 1 kHz.

Events which pass HLT selection are written to various streams and later made available via a set of trigger 'menus'. Trigger menus make relevant detector measurements available for subsequent analysis at the physics level.

The ATLAS detector is among the most impressive and complex technological feats yet accomplished by mankind. Largely optimised by design to be a

discovery experiment, ATLAS has overcome the QCD-induced complexity of high-energy hadronic interactions to accomplish numerous leading high-precision measurements, rendering the LHC both a discovery and precision machine.

While the detector is capable of such precision, the results obtained from ATLAS data are only as accurate as the methodologies utilised to *reconstruct* the signals collected from collision remnants into a fuller picture of the original event. An evaluation of the algorithms used to extract physical insight from low-level ATLAS data now follows.

> *One accurate measurement is worth a thousand expert opinions.*
>
> — Grace Hopper

8

Reconstruction of Physical Objects

LHC collisions occurring in the ATLAS detector create a complex and evolving spray of subatomic particles, depositing their energy in the various sub-systems of the experiment as they traverse its many layers. Identifying the nature of the particles whose energy deposits are measured, before determining the nature of the event to which they owe their production — all within the time constraints imposed by a 40 MHz collision rate — is a conceptual and computational challenge of daunting proportions. It is the role of ATLAS *reconstruction algorithms* to undertake this task. An examination of the algorithms and processes which are employed to reconstruct the remnants of LHC collisions and enable the original underlying physical event to be investigated now follows.

8.1 Overview

The many overlapping *pp*-collisions which occur at the centre of the ATLAS detector can each give rise to rich particle phenomena, with particles decaying into other species before detection, producing secondary interactions with the material budget of the detector, and hadronising to produce complex cascades. Identifying the nature of the physical events which produce such a complex and evolving signatures before the arrival of subsequent proton bunches is therefore a challenging undertaking.

The suite of ATLAS reconstruction algorithms seek to facilitate this by accurately and efficiently ascertaining the energy, four-momenta, charge, identity, and other key attributes of the particles which interact with the detector for subsequent analysis.

Before high-level defined objects pertaining to particles and their properties can be reconstructed, objects of a low-level nature must first be computed. The reconstruction of elementary particles and jets will largely be a function of such low-level objects.

8.2 Low-Level Objects

Low-level objects are those which capture the raw detector response to various incident particles, and form the basis upon which the high-level objects of interest to analyses — such as electrons, muons, and jets — are computed. Accurate and dependable reconstruction of such objects is therefore of critical import.

8.2.1 Track Reconstruction

As has been examined in the previous chapter, charged particles leave 'hits' in the inner detector as they traverse its layers. The variables associated with particle hits can be used to form tracks, such that the trajectory of charged particles within the detector can be determined.

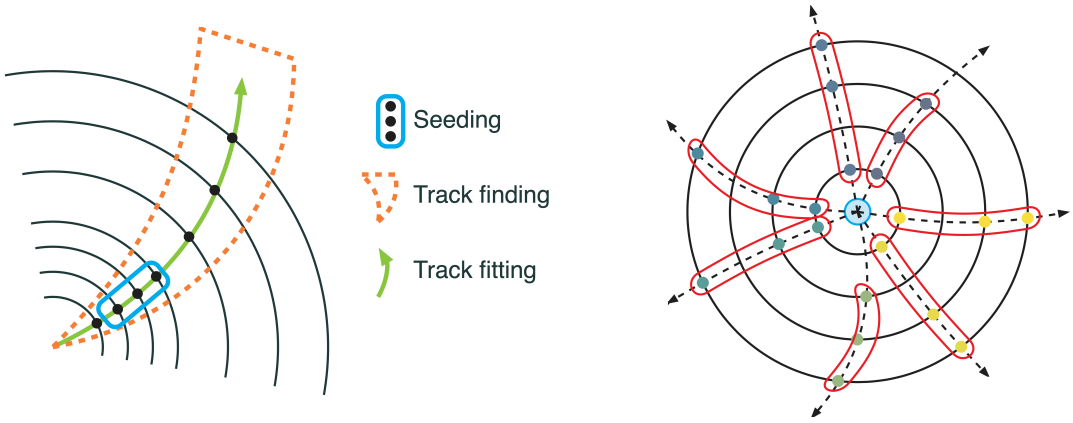

Figure 8.1 illustration of track reconstruction with the ATLAS experiment. (*Left*) shows the various stages of the tracking sequence; (*right*) depicts collections of track seeds forming tracks on the transverse plane. Figure extracted from [124].

Tracks formed in the inner detector are parameterised as:

$$\tau = \tau\left(d_0, z_0, \phi_0, \theta, \frac{q}{p}\right). \tag{8.1}$$

Here, d_0 denotes the transverse impact parameter — defined as the distance of closest approach in the transverse plane of the track candidate to the primary vertex, with z_0 denoting the longitudinal impact parameter — instantiated to the z-coordinate of the perigee point. Addition spatial parameters ϕ_0 and θ are respectively defined as the azimuth and polar angle. Further reconstructive leverage is sourced from the charge-to-momentum ratio, q/p. Three-dimensional space-points are constructed from Pixel and SCT information, in conjunction with the response of drift tubes within the TRT.

Tracks are reconstructed with the *New Tracking* (NEWT) algorithm [125], which undertakes two complementary reconstruction processes. An 'inside-out' procedure initiates from the Pixel layers and constructs track objects in an outward sweep of the detector. This approach secures efficient reconstruction of the trajectories of particles produced in the original hard-scatter event and the decay of short-lived particles. Tracks reconstructed with this process are seeded from triplets of particle hits in the Pixel and SCT sub-detectors. Tracks are then extrapolated from said track seeds to the extremity of the SCT, with the aid of data concerning the detector material budget and magnetic field configuration. This is achieved with the use of a *combinatorial Kalman filter* [126]. An ambiguity

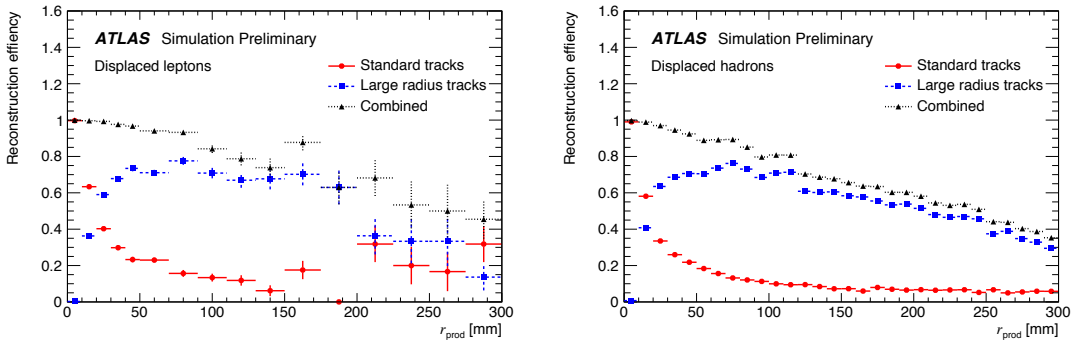

Figure 8.2 performance of the ATLAS large radius tracking algorithm with regards to the reconstruction of (*left*) displaced leptons and (*right*) displaced hadrons. LRT is observed to attain appreciably stronger reconstruction of particles produced at a greater displacement ('radius of production', r_{prod}) than standard tracking. Taking the logical combination of both tracking methods can be seen to attain maximal reconstruction efficiency by leveraging the stronger performance of standard tracking at low values of displacement. Figure extracted from [127].

solving process is executed to resolve any ambiguities which may arise by ranking multiple track candidates with respect to variables of interest (such as quality of fit). Missing particle hits are appropriately penalised in order to suppress the likelihood of track mis-reconstruction. Tracks which survive the strict selection imposed by this process are subsequently extended to the TRT to be combined with the particle hits residing in this layer of the detector.

In concert, a complementary 'outside-in' procedure performs a pass from the outer layers of the detector inwards, seeding TRT track segments from the electromagnetic calorimeter. Compatible hits from the SCT unused by the inside-out process are used to extend the TRT segment to a full track when possible. This secondary pass is better equipped to reconstruct tracks emanating from secondary vertices which may evade reconstruction under the inside-out pass, having deposited a smaller number of silicon hits. An illustration of the ATLAS tracking procedure is provided in Figure 8.1.

Large Radius Tracking

The inclusion of *Large Radius Tracking* (LRT) [127] in the ATLAS tracking reconstruction sequence is partially the result of the sizeable and growing interest in long-lived signatures. LRT seeks to increase acceptance of long-lived particles by performing an additional pass over particle hits registered in the SCT which

Parameter	Standard	Large Radius
Maximum d_0	10 mm	300 mm
Maximum z_0	250 mm	1,500 mm
Maximum $\|\eta\|$	2.7	5
Maximum number of shared hits in Silicon modules	1	2
Minimum number of hits in Silicon modules	7	7

Table 8.1 comparison of parameter instantiations between the standard ATLAS tracking algorithm and the ATLAS large radius tracking algorithm.

were discarded by previous passes of the traditional tracking algorithm.

The algorithm operates in a similar manner to the inside-out process undertaken in standard track reconstruction. Tracks are seeded in the Silicon sub-detectors before an extension into the TRT is attempted. Furthermore, the maximum value of impact parameter to which the algorithm is sensitive is increased from $d_0 = 10$ mm to $d_0 = 300$ mm, thereby improving the acceptance to particles which decay at appreciable displacement from the interaction point. A comparison of the standard and LRT tracking algorithms is presented in Table 8.1.

8.2.2 Primary Vertex Determination

With the availability of fully reconstructed trajectories of charged particles, the original event which gave rise to the charged particles can now be isolated. To do so, the *primary vertex* (PV) of the *underlying event* must be determined.

The PV is the vertex which is associated with the parent pp-collision, known as the underlying event. Identification of the primary vertex is essential if events of interest are to be disentangled from the many secondary and pile-up interactions present in the detector environment.

The ATLAS primary vertex reconstruction algorithm [128] accomplishes this in two steps: *vertex finding*, where candidate vertices are constructed from a selection of tracks, and *vertex fitting*, where the position of each vertex and the associated covariance matrix are computed.

Tracks used to determine the primary vertex must satisfy the following conditions:

- $p_{\mathrm{T}} > 0.5$ GeV;
- $|d_0| < 4$ mm, $\sigma(d_0) < 5$ mm, $\sigma(z_0) < 10$ mm;
- num. silicon hits $\geqslant 9\,(11) \vee |\eta| < 1.65\ (> 1.65)$;
- at least 1 hit in the first two layers of the pixel system;
- no missing hits with respect to the fitted particle trajectory ('holes') in the pixel system, and no more than one hole in the SCT system;
- a maximum of one shared module (*i.e.* 1 shared pixel hit or 2 shared SCT hits).

Tracks which survive selection are processed by the vertex reconstruction algorithm in an iterative manner. The process initiates by computing the seed position for the first vertex in the set of vertex candidates, where the position is determined with respect to the centre of the beam pipe. The position of the vertex is subsequently fitted as a function of the afore-computed seed position and the position of the input tracks.

The position of the vertex is fitted via a process known as *iterative annealing* [128], whereby each input track is ascribed a weight proportional to its compatibility with the vertex candidate under consideration.

It is at this point in the execution that the process reiterates, with a new seed position calculated with the weighted tracks and a new position fitted.

At the completion of a suitable number of iterations, tracks with small weightings are interpreted as exerting minimal influence on the vertex position and regarded as incompatible with the vertex in question. Tracks which are deemed incompatible with the vertex under consideration are ejected from the set of tracks such that they may be considered for the reconstruction of the subsequent vertex. This process repeats until such a time as the set of input tracks is depleted or no further vertices are identified.

Having identified all possible vertices, those which are associated with two or more tracks are considered to be candidates for the PV. The vertex with the maximum sum of track momenta ($\sum p_{\mathrm{T}}^2$) is designated as the PV, with remaining vertices attributed to pile-up interactions.

8.3 High-Level Objects

The computation of low-level variables, such as those which pertain to tracking and vertex information, enables the reconstruction of the types of particles produced in pp-collisions and their kinematic properties. The methods by which this is achieved are now examined.

8.3.1 Electrons and Photons

The presence of electrons and photons in a given event is indicated by the shower deposits in the electromagnetic calorimeter. This signature alone, however, is not a sufficient discriminator with which to differentiate between electrons and photons — electrons will emit photons under the process of Brehmsstrahlung radiation, while quantum field theory dictates that photons will produce electron-positron pairs. Calorimetry energy deposits must therefore be taken in conjunction with tracking information to accurately resolve the different signatures which originate from electrons and photons.

Electrons and positrons produce charged tracks emanating from their production vertex which point towards their respective energy deposits in the electromagnetic calorimeter, as visualised in Figure 8.3. Photons, meanwhile, do not form a track as they traverse the detector owing to their status as an electrically neutral particle.

The reconstruction of electrons with the ATLAS detector is undertaken in accordance with the procedure which follows. The process initiates with the identification of shower seeds in the electroamgnetic calorimeter via the *slding-window algorithm*. Fixed-size clusters of dimension 3×5 cells are constructed. The magnitude of energy deposits in all three calorimeter layers which fall within this window are subsequently summed, yielding the total transverse energy within the window. The position of this window in relation to the calorimeter coverage is then moved iteratively until the window position which maximises the total transverse energy is determined — should the total energy at this position exceed the given threshold of $2.5\ \mathrm{GeV}$, the position is recorded as a calorimeter seed. The value of this energy threshold is selected such that the trade-off between reconstruction efficiency and fake signal acceptance is optimised.

By means of comparison with tracks formed in the inner detector, calorimeter seeds are subsequently classified as having originated from a given type of particle on the basis which follows:

1. **electrons** — should a match be made between the seed and a track which points to the *primary* vertex of the event (positrons are distinguished from electrons as they experience the deflection of the magnetic field in the opposite direction);
2. **converted photons** — should a match be made between the seed and a track which points to the *secondary* vertex (having been produced by a trackless photon emanating from the primary vertex);
3. **unconverted photons** — should a match to a track within the inner detector be absent.

Approximately 20% of photons produced at low values of η convert in the inner detector [9]. The larger detector material quota encountered by photons within $|\eta| \approx 2.3$ increases the probability of electromagnetic showering events, thereby increasing the conversation rate to $\sim 65\%$ in this region [9]. The efficiency for each stage of the reconstruction process is shown in Figure 8.4, as determined with the aide of MC simulation.

A discriminant, $d_L = L_S/(L_S + L_B)$, is then determined for each electron candidate obtained through sliding-window procedure. This is calculated as a function of the likelihood functions given by:

$$L_{S(B)}(\vec{x}) = \prod_{i=1}^{n} P_{S(B),i}(x_i), \tag{8.2}$$

where $\vec{x}$ vectorises the n discriminating variables which correspond to the relevant tracking and calorimetry information, and $P_{S,i}(x_i)$ and $P_{B,i}(x_i)$ respectively denote the probability density functions corresponding to the i^{th} variable at value x_i for signal and background processes. The signal process is taken to be a promptly-produced electron, while background processes include jets which mimic the signature of a prompt electron, photon-conversions within the material budget of the detector, and non-prompt electrons emanating from the decay of heavy-flavour hadrons. Probability density functions which feature in this discriminant are derived from MC simulation.

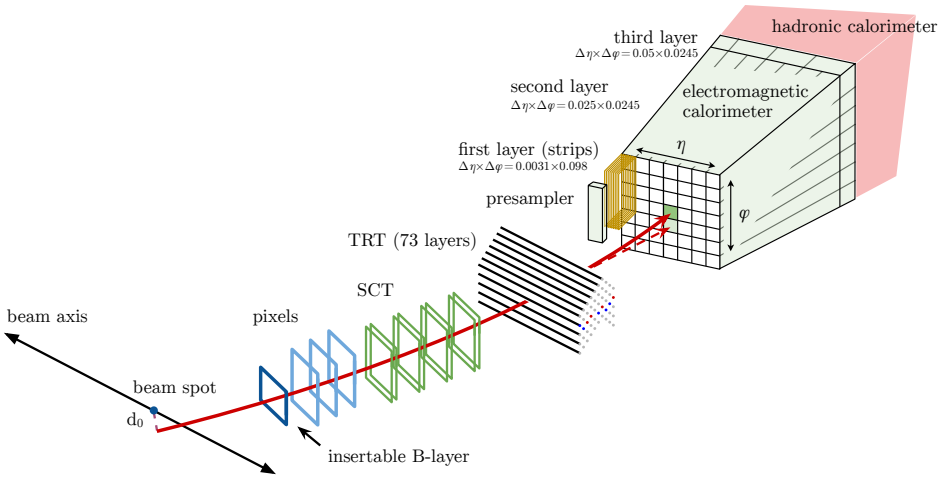

Figure 8.3 schematic illustrating the trajectory of an electron through the various layers of the ATLAS detector. Interactions with the material budget of the detector produce a photon, depicted by the dashed line. Image extracted from [129].

The distribution of electron discriminant values may then be employed to select objects for subsequent analysis. *Working points* may be defined for various 'cuts' along this distribution which quantify the desired balance of signal acceptance and signal purity, for which differing appetites will exist depending on the nature of each given analysis. Three working points by the names of `loose`, `medium`, and `tight` are defined — referring to the degree of stringency exhibited towards background processes by each selection. The electron identification efficiency attained for each working point as a function of transverse energy and pseudorapidity is respectively provided in Figure 8.5 (a) and Figure 8.5 (b) overleaf.

8.3.2 Muons

With the exception of neutrinos — which leave no direct trace in the detector — muons are the only SM particle to traverse to the outer-most layers of ATLAS unimpeded. No other particles leave energy deposits in the muon spectrometer. Taken in concert with tracking information from the inner detector, data from the ATLAS muon sub-systems provide strong detection and reconstruction capabilities. Despite the unique nature of this combined signature, pile-up interactions in addition to gaps in spectrometer coverage necessitate that information from electromagnetic and hadronic calorimeters also be incorporated into the ATLAS identification procedure for second-generation leptons.

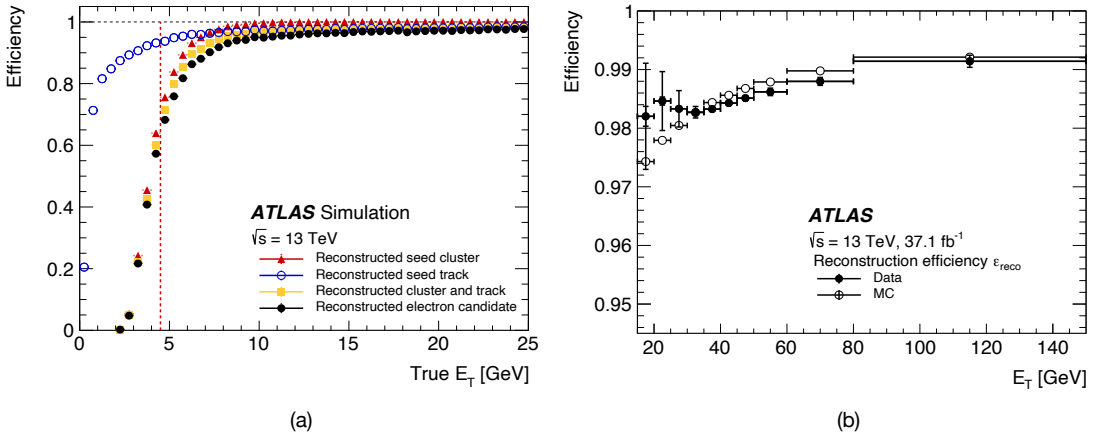

Figure 8.4 electron reconstruction efficiency shown for *(a)* simulated single electron samples as a function of MC transverse energy, as attained at each stage of the reconstruction sequence, and *(b)* simulated $Z \to ee$ events with respect to reconstructed clusters, as a function of MC transverse energy, for measured and simulated data. Figures extracted from [129].

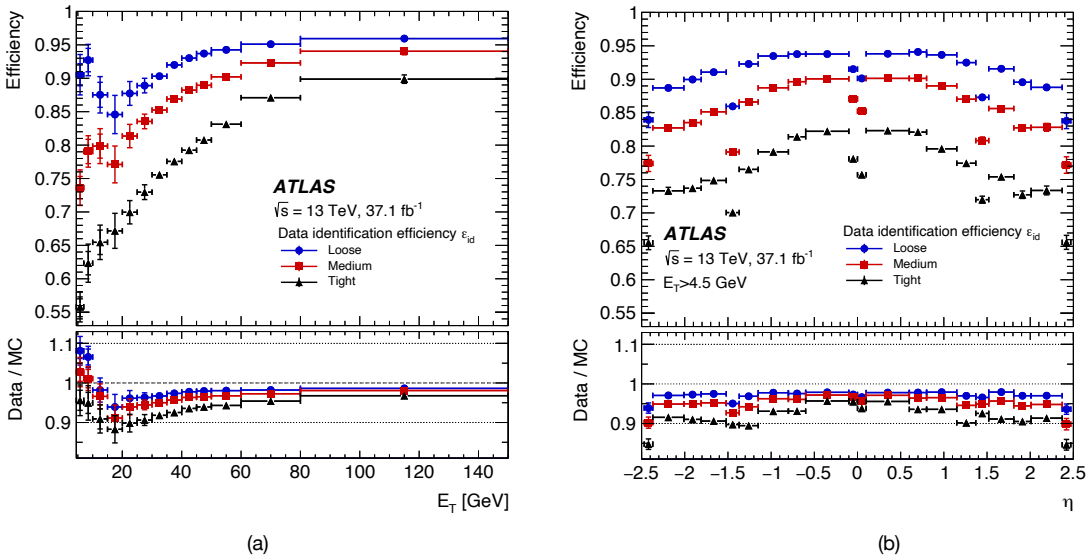

Figure 8.5 electron identification efficiency for `loose`, `medium`, and `tight` working points as a function of *(a)* electron transverse energy and *(b)* pseudo-rapidity, as determined with $J/\psi \to ee$ and $Z \to ee$ simulated samples. Ratios of measured data to simulated MC samples are shown at the bottom of each plot. Figures extracted from [129].

For muons which fall within the pseudo-rapidity window of $|\eta| < 2.5$, measurements of position and momenta are obtained from the inner detector to a high degree of precision. Stringent conditions are placed upon the number of registered hits recorded by the various inner detector sub-systems and the number of traversed silicon layers where a hit is absent allow high-quality muon tracks to be selected.

Pseudo-rapidity measurements can be made with information from the muon subsystems to which such tracks point, in which three layers of drift-tube detectors are equipped to perform six to eight high-precision measurements of particle η-position per muon within $|\eta| < 2.7$. Muon trigger chambers provide ϕ-coordinate information to a more coarse degree of granularity.

Hits registered within each layer of the muon spectrometer are taken in turn to form a *local* track object, each of which is concatenated with similar objects formed in other layers of the spectrometer — this combined object is termed a *global* muon spectrometer track. Four categories of reconstructed muons can then be defined thusly:

1. **combined muons** — where a track in both the inner detector and muon spectrometer have been recored, and a satisfactory match associating one with the other can been asserted;

2. **segment-tagged muons** — where an inner detector track can be matched to local tracks formed in segments of the muon spectrometer, in the absence of a global muon spectrometer track;

3. **stand-alone muons** — where a global muon spectrometer track has been recorded without an accompanying inner detector counter-part (this can arise as a result of limitations in detector coverage);

4. **calorimeter-tagged muons** — where an inner detector track is recorded in the absence of a muon spectrometer track, in conjunction with matching energy deposits in the calorimeter systems consistent with a minimally-ionising particle (such objects can only be reliably called upon in regions of known limitations in muon spectrometer coverage due to their low purity).

In a similar spirit to the procedure employed for electron identification, a discriminating variable is determined for each candidate muon reconstructed via

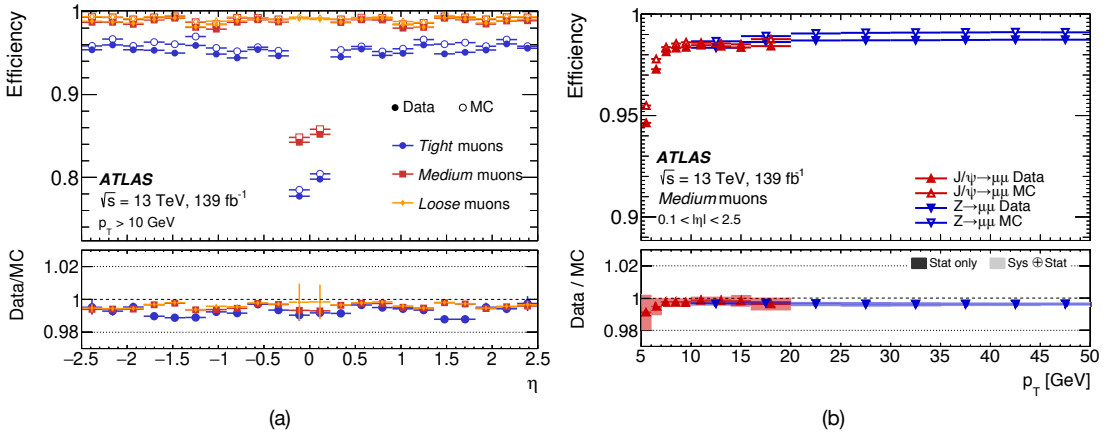

Figure 8.6 muon reconstruction efficiency as a function of *(a)* pseudo-rapidity, shown for `loose`, `medium`, and `tight` measured and simulated muons, and *(b)* transverse momentum, as determined for measured and simulated muons of `medium` quality. Figures extracted from [130].

the afore-described process. The variable of choice is the *charge/momentum* (q/p) *significance*, defined as:

$$(q/p)_{\text{sig.}} = \frac{|q/p|_{\text{MS}} - |q/p|_{\text{ID}}}{\sigma_{q/p}}, \tag{8.3}$$

where $|q/p|_{\text{ID}}$ is the muon charge-to-momentum ratio as measured by the inner detector, $|q/p|_{\text{MS}}$ is the same ratio as determined by the muonic spectrometer, and $\sigma_{q/p}$ is the sum in quadrature of the uncertainties associated with each respective measurement. This variable provides strong discriminating power between muons and background processes which mimic their signature within the ATLAS detector.

Quality working points are ascribed to the distribution of this variable, as was performed for the likelihood-based electron identification discriminant. The reconstruction efficiency attained for each muon working point is provided in Figure 8.6, as determined by the tag-and-probe method [9] using $J/\psi \to \mu\mu$ and $Z \to \mu\mu$ MC samples.

8.3.3 Jets

Quarks and gluons produced in LHC collisions are not observed as discrete individual particles, but rather as sprays or showers of hadrons known as *jets*. Jets

are observed as narrow, collimated cones of partons traversing collinearly through the ATLAS detector. Such hadronisation occurs as a consequence of the strong potential between the partons increasing with separation, as discussed in Section 2.3. The potential energy between the partons — well-described by the model of *flux tubes* [23] — continues to increase as the partons diverge until sufficient in magnitude to produce new $q\overline{q}$ pairs from the quantum vacuum. QCD colour-confinement dictates that colour-charged particles can only exist in colour-neutral quantum states, thus the hadronisation process continues iteratively, giving rise to jet-like phenomena.

Jets of various origin form an integral component of the vast majority of analyses performed with the ATLAS experiment: from precision measurements of Standard Model processes to searches for physics beyond the Standard Model (including the search presented in this thesis). Moreover, jets represent an interesting physical object of study in their own right, providing an insightful and elucidating tool with which to probe the fundamental nature of the strong force. Efficient and performant reconstruction of hadronic jets is therefore a crucial prerequisite for undertaking physics research with the ATLAS experiment.

ATLAS jet reconstruction processes can accept both inner detector tracks and calorimeter energy deposits as algorithmic inputs. Jets which are reconstructed with the aide of inner detector tracking information — termed *track-jets* — prove less sensitive to pile-up effects, given that only tracks which point to the primary vertex are considered. Track-jets must, however, fall within the narrow pseudo-rapidity acceptance of $|\eta| < 2.5$, rendering jets reconstructed from calorimetry deposits more readily usable to most analyses performed with ATLAS.

Two main categories of jets exist: *small-R* jets and *large-R* jets. The difference between the two classifications is largely concerned with the physical process from which they originate and the manner in which they cascade through the detector. Jets which arise from a quark or gluon due to colour-confinement are typically tightly collimated, and can therefore be reconstructed within a smaller angular volume — such jets fall under the small-R categorisation. Calorimeter energy deposits resulting from large-R jets tend to be more diffuse in nature. These typically originate from the hadronic decay of massive particles, such as the Higgs and electroweak bosons, whose decay products have a larger angular separation which is propagated through the parton shower.

The process of reconstructing jets from energy deposits in the calorimeter systems

begins by collating individual cells in the calorimeter which harbour energy above a given threshold into super-sets of clusters. The energy threshold is determined by taking the quadrature sum of the background contribution expected to arise from pile-up effects in conjunction with electronic noise. This threshold is typically denoted by σ.

A *jet clustering algorithm* is subsequently employed to scan over seed cells whose energy exceeds 4σ, iteratively adding neighbouring cells to the cluster should their energy exceed 2σ. The addition of all cells adjacent to the working cluster then follows, with the resultant object known as a topological cluster, or *topo-cluster*. Should more than one local energy maxima be present in a given topo-cluster with at least four neighbouring clusters whose energy is less than the local maximal cell, the topo-cluster is split. This is performed in an effort to avoid the occurrence of overlapping jets.

Owing to the different response of the calorimeters to hadronic and electromagnetic energy, the total energy of the topo-cluster cannot simply be taken to be equal to the energy of the showering event. Instead, the energy of the calorimeter cells is measured at the electromagnetic (EM) scale, with a correction applied to clusters whose origins are classified as hadronic in nature. Topo-clusters originating from hadronic jets are scaled by a factor called the Local Cell Weighting (LCW), derived from Monte Carlo samples of single pion decays. Jets can then be reconstructed from EM topo-clusters or LCW topo-clusters.

While alternative algorithms exist [131, 132, 133], jets produced by collisions at ATLAS are typically reconstructed using the sequential recombination `anti-kt` algorithm [134]. The algorithm commences by treating each jet candidate in the event as a 'pseudo-jet'. A distance metric, d_{iB}, is computed to define the distance between the pseudo-jet object and the proton beam. A further distance metric, d_{ij}, is computed to measure the separation of each pair of pseudo-jets i and j. This metric is determined by the transverse momenta of the pseudo-jets and their angular separation in angle-energy space:

$$d_{iB} = p_{\mathrm{T}i}^{2P}, \tag{8.4}$$

$$d_{ij} = \min\left(p_{\mathrm{T}i}^{2P}, p_{\mathrm{T}j}^{2P}\right) \cdot \frac{\Delta_{ij}^2}{R^2}, \tag{8.5}$$

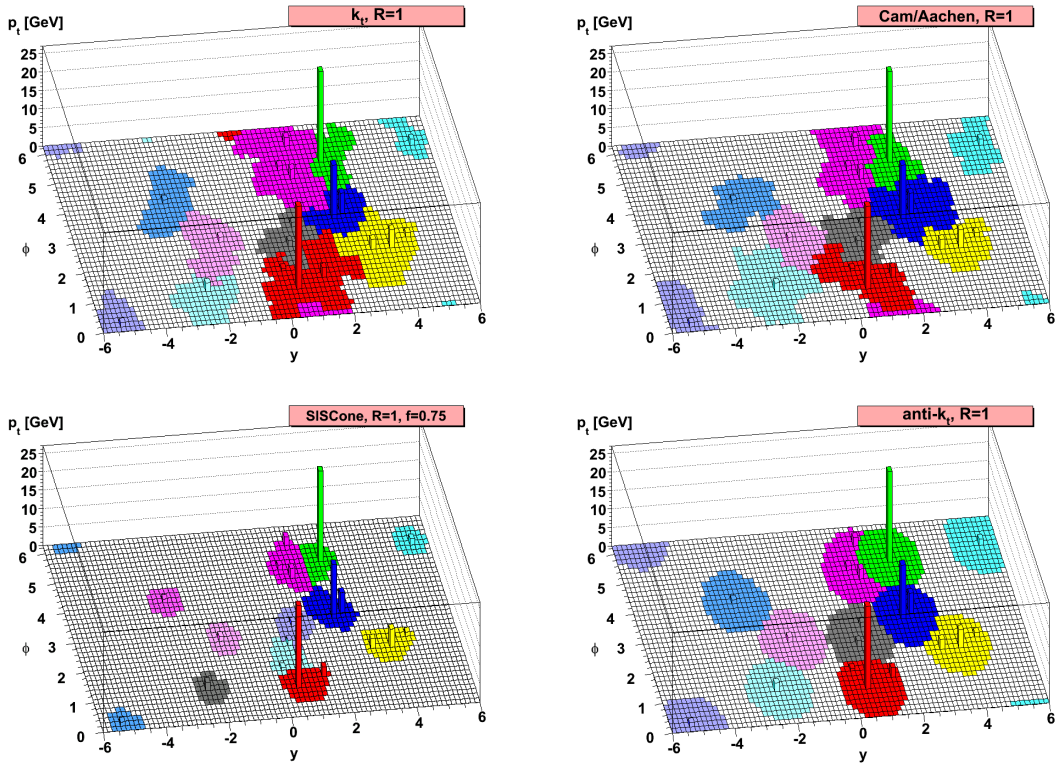

Figure 8.7 cluster formations produced by the `anti-kt` [134] jet clustering algorithm (*bottom-right*) as compared to the output of alternative clustering algorithms [131, 132], performed on a simulated sample of parton-level events, in conjunction with random soft 'ghosts', generated with the `Herwig` [77] event generator. Source: [132].

where $\Delta_{ij}^2 = (y_i - y_j)^2 + (\phi_i - \phi_j)^2$ for rapidity y and azimuth ϕ.

The parameters P and R in the afore-listed equations are configurable. In the case of the `anti-kt` algorithm, P is taken to be $P = -1$, resulting in an inverse-square power-law dependence on the transverse momenta of pseudo-jets. The hardest particles in the event are therefore treated first, with the inclusion of soft radiation occurring at the termination of the algorithm. Alternative reconstruction algorithms instantiate this value differently. The parameter R is used to set the size of the jet. For most analyses this is taken to be $R = 0.4$, increasing to $R = 1.0$ for large-R jets.

Each iteration of the algorithm selects the distance metric with the smallest value. In the case where d_{ij} is the smaller of the two metrics, the two pseudo-jets in question are usurped by a new pseudo-jet with momentum equal to the vector sum of the discarded pseudo-jet pair, i and j. Should d_{iB} prove to be the smaller distance, the pseudo-jet cannot be merged and is removed from the set

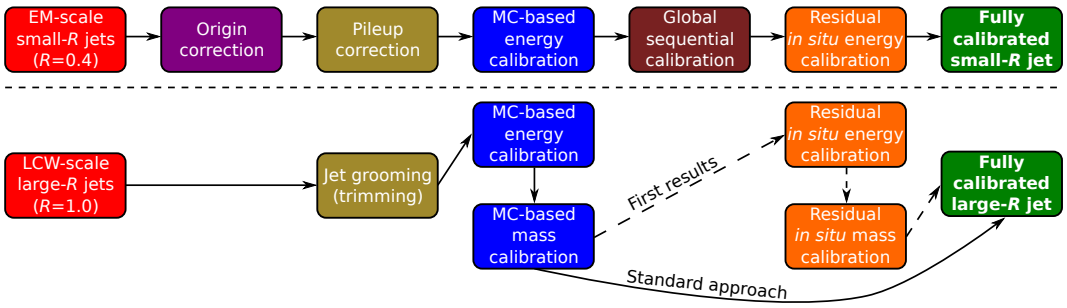

Figure 8.8 Flowchart depicting the steps undertaken in the jet calibration process for both EM-scale and LCW-scale `anti-kt` jets. Source: [135].

— it is thereafter considered to be a full jet. This process of merging and ejecting pseudo-jets continues iteratively until the set of pseudo-jets is depleted. The resultant cluster formations obtained via the `anti-kt` algorithm is contrasted to those attained via alternative procedures in Figure 8.7.

To be of dependable practical utility, a jet reconstruction algorithm must meet the tests set by both infrared safety and collinear safety. An infrared safe algorithm is one whose output is invariant under the emission of extra soft particle radiation which does not originate from the original hard scatter process. Collinear safety is assured if the algorithm is insensitive to the collinear splitting of partons — *i.e.* to the distribution of transverse momenta among the decay products of the parton shower.

The energy scale of the jets is then restored to that of simulated 'truth' jets using multiplicative corrective factors known as *Jet Energy Scale* (JES) calibration factors, derived from MC. Such factors are defined individually for EM and LCW jets. The systematic uncertainty associated with the JES can be a leading source of uncertainty in analyses whose events of interest feature multiple jets in the topology of their final state. This uncertainty is shown in Figure 8.9.

Jet calibration accounts for many factors which can disturb the measured energy of a jet. Origin correction, for instance, ensures that the four-momentum of a jet points towards the primary vertex of the event while leaving the energy of the jet unperturbed. Mitigation of pile-up effects which may artificially inflate, or 'bloat', the measured energy of a jet is also achieved through appropriate calibration, in addition to potential sensitivity to the QCD colour of the partons which initiate the hadronic cascade (through a process known as *global sequential calibration* [136]).

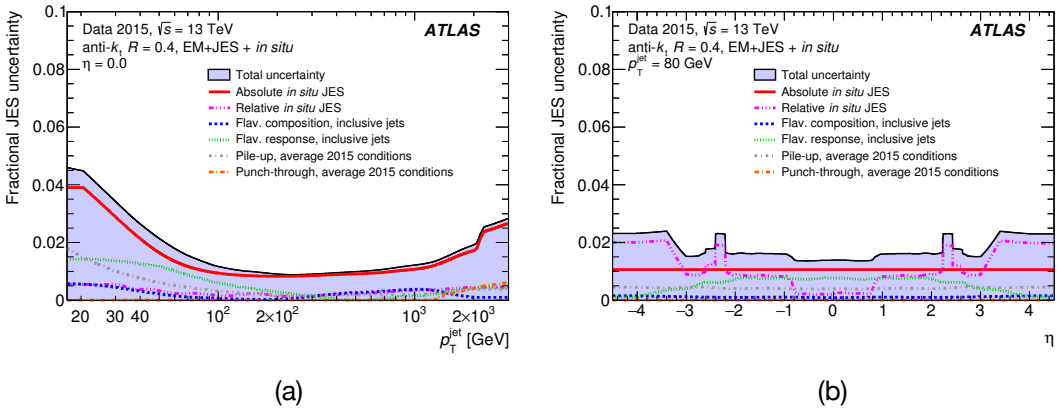

Figure 8.9 shows the combined uncertainty in the JES of fully calibrated jets as a function of (a) jet p_T at constant $\eta = 0$, and (b) η at constant $p_T = 80$ GeV. Total systematic uncertainty includes contributions from pile-up, punch-through, and *in-situ* calibration uncertainties. Source: [137].

Correction of pile-up effects proceeds in two steps. The initial step of the sequence applies an *area-based* correction to the measured value of p_T, where the contribution from pile-up events is subtracted:

$$p_T^{\text{corrected}} = p_T^{\text{measured}} - \rho\, A_T. \tag{8.6}$$

Here, ρ denotes the median p_T density of all jets within $|\eta| < 2.0$, while A_T is the area of the jet under calibration. Inclusion of the jet area accounts for jet-by-jet variations in pile-up sensitivity, while ρ captures event-by-event fluctuations in pile-up intensity. Post-application of the area-based correction, there remains a residual perturbation from the effects of pile-up interactions on the measured value of jet p_T. This is visible as a dependence of the reconstructed p_T upon the number of primary vertices, N_{PV}, and the number of pile-up interactions, $\langle\mu\rangle$, present. A second corrective step in the form of a *residual* correction is therefore applied, counteracting this dependence. Taken together, the effects of both in-time and out-of-time pile-up are adequately mitigated by such steps, as may be observed in Figure 8.12.

A further corrective quantity known as the *jet energy resolution* is computed to account for fluctuations in measured jet energy at a fixed physical energy as a function of detector location. This quantity is determined by measuring the asymmetry of QCD di-jet events in the detector, which must, per force, be attributable to mis-reconstruction.

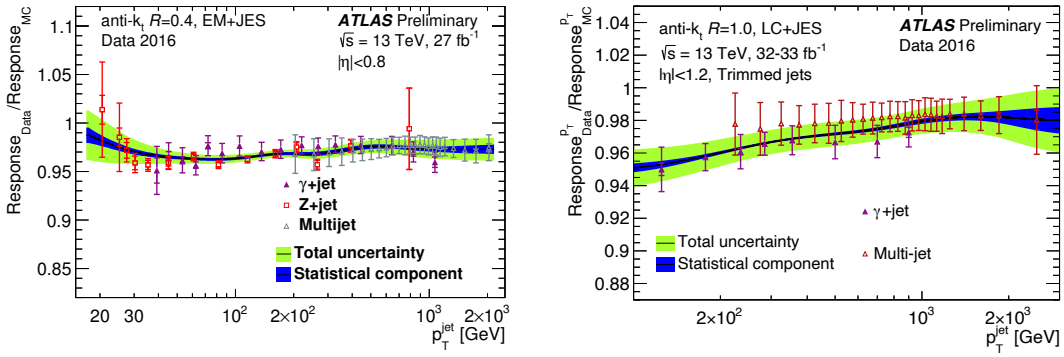

(a) (b)

Figure 8.10 ratio of EM jet response to 'true' MC-level data for *(a)* small-R jets and *(b)* large-R jets after jet energy scale (JES) calibration for $\gamma+$jet, $Z+$jet, and multi-jet processes, shown with associated uncertainties. Small-R jet response is shown to agree strongly with MC data after calibration, with large-R jet reconstruction performing more accurately for higher values of p_{T}. Figure *(a)* and *(b)* respetively extracted from [138, 139].

The larger radius required to reconstruct large-R jets has the unfortunate consequence of rendering the reconstruction more susceptible to interference from pile-up interactions, simply due to the greater detector volume encompassed by the jet cone. An additional correction is therefore required for large-R jets in the form of *jet grooming*. The grooming process seeks to remove contributions to the jet which likely originate from pile-up interactions rather than from the underlying event. This is achieved by re-clustering the constituents of large-R jets using the alternative `kt` clustering algorithm [133]. A cut is imposed on the transverse momentum of the original `anti-kt` jet, typically of the value f_{cut} : 5% $p_{\mathrm{T}}^{\text{anti-kt}}$. Jet constituents whose momentum falls below this cut are considered to have originated from pile-up contributions and are removed from the original `anti-kt` parent jet. The full calbiration sequence applied to small-R and large-R jets is shown in Figure 8.8, with the post-calibration performance achieved for each category shown in Figure 8.10.

8.3.4 Heavy Flavour Tagging

Unlike gluons or light quarks, the flavour of heavy quarks which initiate the hadronic cascade of an observed jet can be ascertained. This is an important feature for many analyses — largely due to the mass-dependent nature of many SM and BSM couplings and preferential decays to heavier generations of fermions.

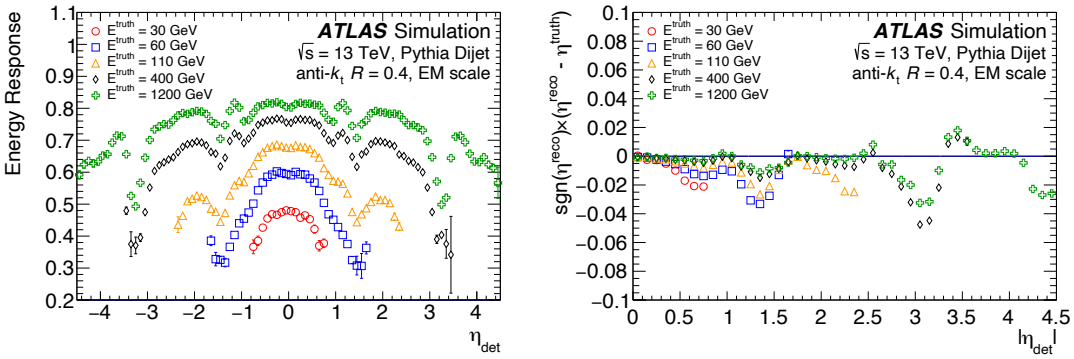

Figure 8.11 shows (*left*) the average energy response as a function of η_{det}, the η vector pointing from the geometric centre of the ATLAS detector. This is shown for jets whose truth energy is drawn from $p_{\text{T}}^{\text{truth}} \in \{30, 60, 110, 400, 1200\}$ GeV post-application of origin and pile-up corrections. The *right* plot shows the difference in reconstruction and truth pseudo-rapidity multiplied by the sign of the reconstruction pseudo-rapidity. The signed difference is due to biases in jet reconstruction which are addressed in calibration. Source: [136].

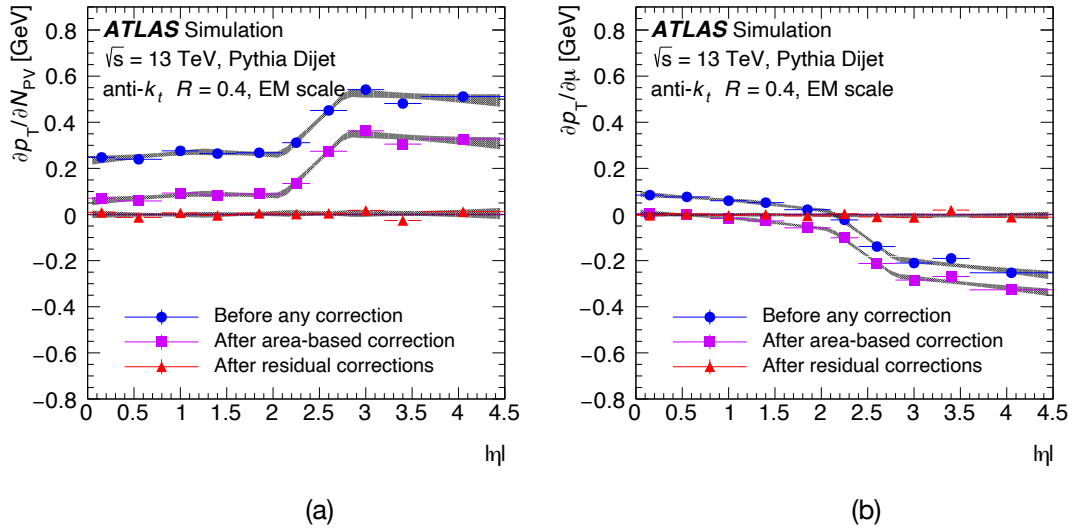

Figure 8.12 shows the dependence of `anti-kt` EM jet p_{T} on (a) in-time pile-up — defined as the number of primary vertices, N_{PV}, averaged over pile-up, μ — and (b) out-of-time pile-up — defined as μ averaged over N_{PV} — as a function of absolute pseudo-rapidity. Shaded bands show the 98 % confidence intervals of the linear fits in four regions of $|\eta|$. The p_{T} spectrum is seen to remain approximately constant under exposure to pile-up interactions when all appropriate corrections are applied. Source: [136].

The top quark does not directly produce a jet-like signature, with a lifetime below that of the QCD hadronisation scale. The next two heaviest quarks, b and c, possess differentiating properties over jets initiated by lighter quarks and thus receive tailored treatment.

The primary hadron associated with a heavy-flavour jet will decay having first traversed some distance (on the order of a few millimetres) within the detector. In the case of b-hadrons, this is primarily driven by the large CKM suppression imposed on third-generation quarks and the kinematic prohibition of decaying to a child top-quark. Moreover, the fragmentation function of the the b-quark is comparatively large, consequentially leading to a large proportion of its momentum being carried by the b-hadron it progenerates before decaying weakly. This subsequently generates a large relativistic Lorentz boost for the b-hadron, rendering the signature associated therewith yet more distinctive. The c-quark exhibits a similarly long decay-length, which, in conjunction with the greater relativistic time dilation experienced by a particle of lighter mass, also render its experimental signature discriminable.

Such distinct properties allow for the jets initiated by b- and c-quarks to be differentiated from those engendered by gluons or quarks of lighter mass. The act of assigning a flavour of origin to a given jet is referred to as 'flavour tagging'.

b-Jet Tagging

The process of tagging jets produced from b-quarks begins with the execution of some basic algorithms — these use tracking information to produce a set of discriminating variables with which to delineate jets of varying flavour origins. Such algorithms fall under three categories: *impact parameter based algorithms*, *secondary vertex reconstruction algorithms*, and *decay chain multi-vertex reconstruction algorithms*. The tracks which constitute the input to such algorithms are each matched to a single jet using the measured angular separation between the two reconstructed objects, ΔR — itself a function of p_{T} (jets with higher p_{T} produce narrower, more collimated jet cones). Track selection criteria vary with each category of algorithm.

Impact parameter based algorithms — such as `IPD2` [140] and `IPD3` [140] — calculate the signed impact factor significance, $\mathsf{sig}\,(x) = x/\sigma_x$, of the jet-matched tracks for respective impact factor and associated uncertainty, x and σ_x. A

positive (negative) sign is ascribed to this quantity if the point of closest approach of the track and primary jet vertex lies geometrically ahead (behind) the primary vertex with respect to the direction of the jet in $\eta-\phi$ space. The algorithm returns a log likelihood ratio discriminant quantifying whether the b-jet hypothesis or light flavour hypothesis is more likely for the jet under consideration.

The secondary vertex approach reconstructs displaced secondary vertices which reside within jet substructure. Tracks forming secondary vertices are rejected if flagged as having likely originated from composite particles with appreciable lifetime (*e.g.* K_S, Λ), photon conversion, or nuclear interactions with the detector material budget. The surviving track population is used for the reconstruction of a single vertex, with outliers iteratively removed. The resultant secondary vertex serves as a potent discriminator of jet flavour, as can be observed in Figure 8.13.

The ATLAS decay chain multi-vertex algorithm, `JetFitter` [141], exploits the topological structure of the weak decays of the parent b- and c-hadrons, which unfold within the jet. The reconstruction of the full decay chain ($\mathrm{PV} \to b \to c$) within the jet is the objective of this algorithm. A Kalman filter [126] is used to project a trajectory through the parent b-hadron and the subsequent decay of any child hadrons. The properties of tracks associated with each vertex in this trajectory provide the necessary discriminating power to serve as useful variable in evaluating the likely flavour of genesis.

The variables described (or a selected subset thereof) are then fed to a pre-trained multi-variate classifier, the output of which can then be employed by analyses. The classifier used to produce the b-jet score is a *boosted decision tree* (BDT), to which an introductory explanation can be found in [142]. Two variants of the BDT are available: `MV2c20`, trained on a MC sample of 80% light-flavour jets and 20% c-jets, and `MV2c00`, trained on a MC sample consisting exclusively of light-flavour jets. The performance of each classifier with regards to b-jet identification and the rejection of background jets is depicted in Figure 8.14.

8.3.5 Missing Transverse Energy

Conservation of momentum necessitates that the transverse momenta of particles produced in LHC collisions must sum to zero. Neglecting the small contribution which may arise from momentum mis-reconstruction, imbalances in the summed final transverse momenta can therefore indicate the presence of an additional

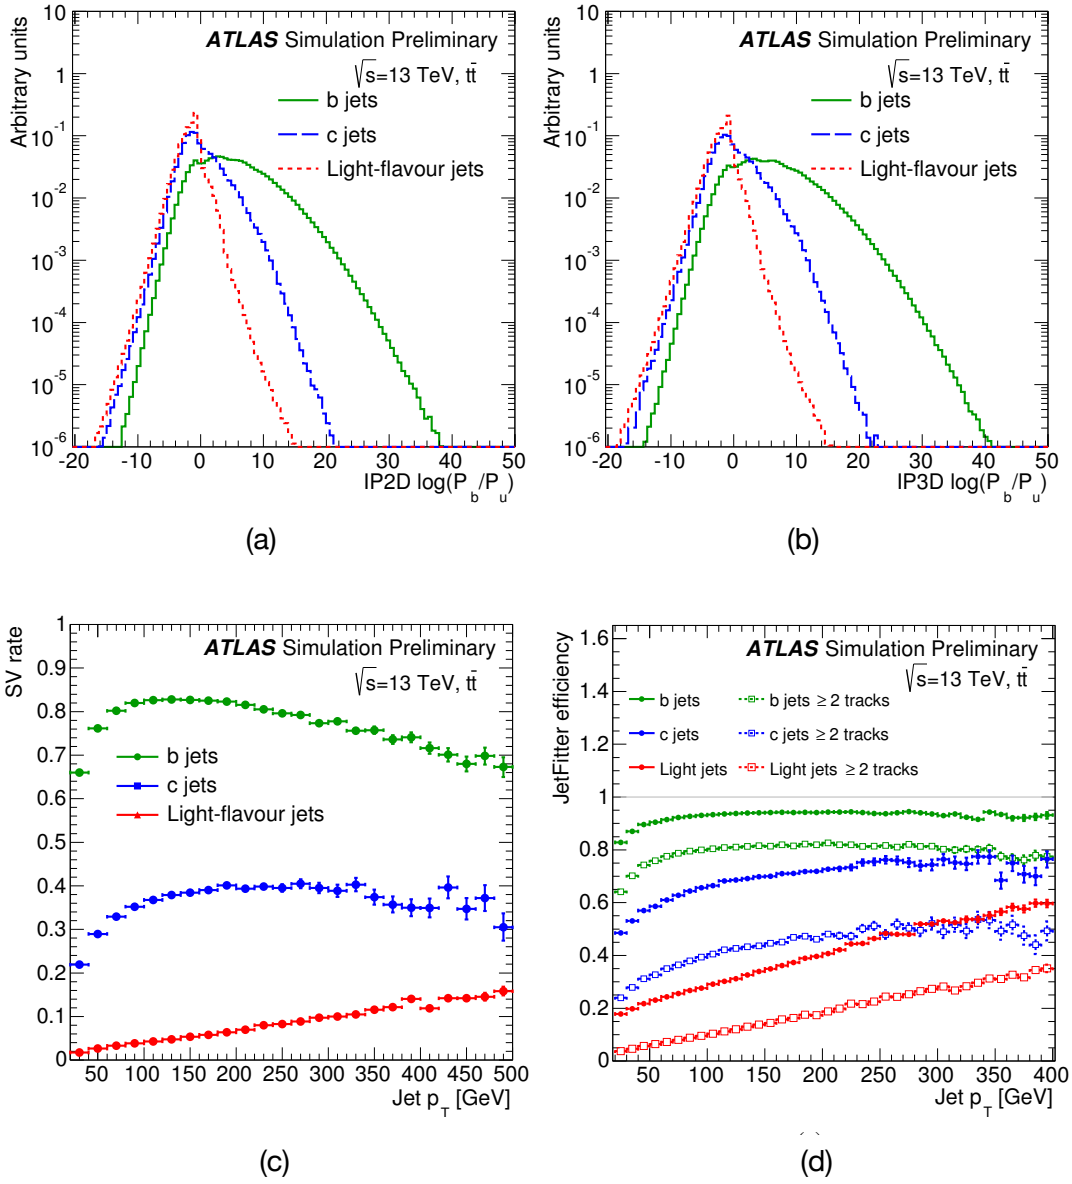

Figure 8.13 shows the algorithmic output of *(a)* `IP2D`, *(b)* `IP3D`, *(c)* secondary-vertex algorithm, and the *(d)* `JetFitter` algorithm for jets of various origin. Figures obtained from [140].

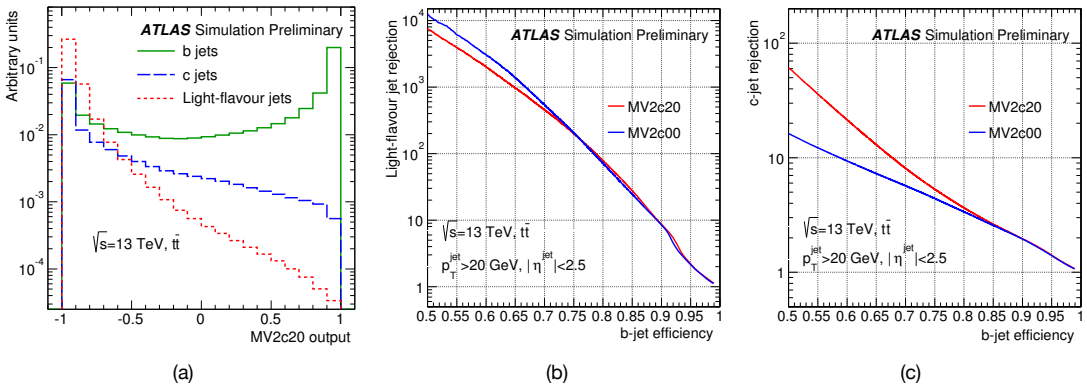

Figure 8.14 output of the flavour-tagging MVA classifier presented in terms of *(a)* jet classification, and *b*-jet classification efficiency vs. the rejection of *(b)* light-flavour jets and *(c)* *c*-jets. Figures extracted from [140].

particle in the final event state which has evaded direct detection.

Such 'missing' transverse momenta (or missing transverse energy, $E_{\mathrm{T}}^{\mathrm{miss}}$: 'MET') therefore proves to be a powerful experimental probe and discriminating variable. SM neutrinos, for instance, deposit little-to-no energy in the ATLAS detector, traversing its layers unimpeded owing to its weakly-interacting nature. The presence of neutrinos produced in pp-collisions is identified by the measurement of MET in events where neutrino production is expected. Moreover, measurement of MET is a central pilar of myriad BSM searches for super-symmetric particles or weakly-interacting dark matter candidates which are not predicted to interact strongly with the detector systems.

While not formally of the same dimensionality as energy, $E_{\mathrm{T}}^{\mathrm{miss}}$ is taken to be equal to the magnitude of the vector formed by missing transverse momentum:

$$E_{\mathrm{T}}^{\mathrm{miss}} = \sqrt{\left(p_x^{\mathrm{miss}}\right)^2 + \left(p_y^{\mathrm{miss}}\right)^2}, \tag{8.7}$$

with the direction of the MET given by $\phi^{\mathrm{miss}} = \arctan\left(p_y^{\mathrm{miss}}/p_x^{\mathrm{miss}}\right)$. Non-physical sources of MET, including finite reconstruction resolution and non-collision backgrounds, must be accurately quantified to avoid attributing false discoveries to their influence. The performance of MET identification with ATLAS is carefully studied with MC data.

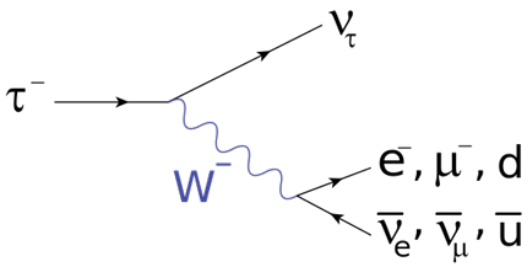

Figure 8.15 Feynman diagram of common τ-lepton decays, mediated by an off-shell electroweak *W* boson.

8.3.6 τ-Leptons

τ-leptons exhibit unique characteristics by which they are distinguished from their first and second generation leptonic counterparts. With a mass of 1.78 GeV and a lifetime of 2.9×10^{-13} s, τ-leptons are the only leptons sufficiently massive to decay hadronically, and decay shortly after production at the interaction point. Their lifetime corresponds to a proper decay length of 87 μm — τ-leptons do not, therefore, reach the inner detector.

τ Decay Modes

The decay process of the τ is primarily mediated by an off-shell *W* boson, as shown in Figure 8.15. Each of the various decay modes can broadly be characterised as either leptonic or hadronic in nature, with respective branching ratios of ~ 36% and ~ 64%. Hadronic decay modes can be further categorised by their track multiplicity, or 'prongness'. This refers to the number of charged tracks present in the decay signature, owing to the number of charged pions produced in the final state. Primary consideration will be given to the categories of 1-prong (1 p) and 3-prong (3 p) in the analysis which follows. As the electrons and muons produced in the final state of fully-leptonic τ-decays cannot be distinguished from those produced directly by the underlying event, such τ-decays cannot be reconstructed. A summary of kinematically-accessible decay modes for the τ-lepton, together with their associated branching ratio, is provided in Table 8.2.

Every decay topology of the τ-lepton features the production of a τ-neutrino (ν_τ). This leads, perforce, to the presence of missing transverse energy in every τ-decay which must be considered in the reconstruction process; the full energy of the τ cannot be reconstructed.

Topology		Final State	Branching Ratio [%]
Leptonic Decay		$\tau^- \to e^- \overline{\nu}_e \nu_\tau$	17.82
		$\tau^- \to \mu^- \overline{\nu}_\mu \nu_\tau$	18.39
Total:			36.21
Hadronic Decay	1-*prong*	$\tau^- \to \pi^- \pi^0 \nu_\tau$	25.49
		$\tau^- \to \pi^- \nu_\tau$	10.82
		$\tau^- \to \pi^- 2\pi^0 \nu_\tau$	9.26
		$\tau^- \to \pi^- 3\pi^0 \nu_\tau$	1.04
	3-*prong*	$\tau^- \to 3\pi^\pm \nu_\tau$	8.99
		$\tau^- \to 3\pi^\pm \pi^0 \nu_\tau$	2.74
Total:			63.79

Table 8.2 available leptonic and hadronic decay modes of the τ-lepton and their associated branching ratios.

The presence of pions in the final state of hadronically-decaying τ-leptons leads to a jet-like signature, complete with either 1 or 3 charged tracks. A jet initiated by a quark or gluon, by comparison, will typically produce between 5 and 20 tracks. A further distinguishing property of the 'τ-jet' is the greater degree of collimation exhibited in comparison to a jet of QCD origin. Such jet-like properties dictate the reconstruction approach adopted in the identification of τ-leptons.

Hadronic τ Reconstruction

Hadronic τ reconstruction commences with the execution of the standard `anti-kt` jet reconstruction algorithm, taking the jet size to be $R = 0.4$. Jets are seeded from three-dimensional calorimeter topo-clusters, to which local hadronic calibration is applied. Jets are required to satisfy the selection of $p_\mathrm{T} > 10$ GeV and $|\eta| < 2.5$ to be considered a τ-jet candidate.

A matching primary vertex must be identified for an event to be considered for τ candidacy. This is required to have three or more tracks. Where more than one possible primary vertex has been identified, the vertex with maximum transverse momenta across the associated tracks ($\max \sum p_\mathrm{T}^\mathrm{track}$) is selected.

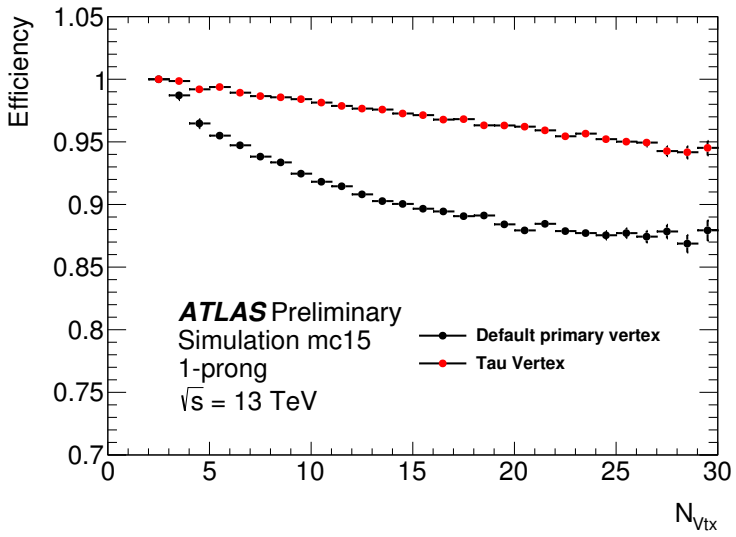

Figure 8.16 τ-lepton reconstruction efficiency before (black) and after (red) execution of the Tau Vertex Association Algorithm. Source: [144].

Given its finite lifetime, the τ will travel some distance from the primary vertex before decaying into visible decay products. The selected primary vertex therefore does not necessarily correspond to the vertex at which the τ was produced. The *Tau Vertex Association Algorithm* [143] attempts to address this by taking the set of tracks within a $\Delta R = 0.2$ cone around the direction of the jet seed, and assigning the vertex with the largest fraction of the total track p_T within this set to be the τ vertex. The effect of pile-up interactions is thereby minimised, resulting in greater reconstruction efficiency as shown in figure 8.16. This definition of τ vertex is subsequently used to determine the direction of the visible τ decay products, perform track selection, and define the coordinate system used to re-compute impact parameters and variables used in the identification process.

Tracks originating from pile-up interactions or the underlying event, or those reconstructed from background processes, may fall within the cone of the reconstructed τ jet. The presence of such tracks can degrade the accuracy with which the correct charged track multiplicity of the τ jet is determined. For this reason, a track classification algorithm is applied to the tracks within the jet cone in order to select those which likely originate from the decay of the τ, such that the track charge and multiplicity of the τ jet can be accurately measured.

Having determined the τ vertex, the variables associated with the τ jet tracks are re-calculated. The following selection criteria are then applied:

- tracks must fall within a jet cone of $\Delta R < 0.25$;

- transverse momenta of each track must exceed $p_T > 1.0$ GeV;
- a minimum of 2 hits within the pixel system of the inner detector must be recorded;
- at least 7 hits across the pixel and SCT systems must be recorded, and;
- the impact parameters computed with respect to the τ vertex must satisfy $d_0 < 1.0$ mm and $z_0 \sin\theta < 1.5$ mm.

Such selection criteria are reflective of the fact that tracks which originate from pile-up interactions are typically lower in p_T than those which originate from the decay of a τ-lepton and are produced at a larger longitudinal distance with respect to the τ vertex. Additional tracks which fall within an 'outer-cone' region defined as $0.25 < \Delta R < 0.4$ may be considered for selection if they constitute the closest track to the jet seed, or if they satisfy the conditions for ghost particle association [145, 146, 147].

Tracks which satisfy the afore-listed criteria are subsequently subjected to the track classifier [148]. The previous classifier, based on a *boosted decision tree* (BDT) ensemble [142], was replaced in run-2 with a novel *recurrent neural network* architecture [149]. The RNN was trained on a selection of track variables found to exhibit discriminating power, such that it can effectively distinguish between tracks which originate from the charged hadrons produced in the decay of τ-leptons, and those which arise from background processes. The track classifier assigns each track to one of the following categories:

- **tau tracks (TT):** charged tracks which emanate from the $\pi^{\pm}$ produced in the decay of the τ;
- **conversion tracks (CT):** tracks produced by electrons or positrons which originate from the photonic decay products of the neutral π^0 within the τ jet-cone. Such tracks enable the momentum fraction of the π^0 to be determined;
- **isolation tracks (IT):** tracks which reside within the outer-cone region of the τ jet;
- **fake tracks (FT):** tracks which arise from processes other than hadronic τ decays.

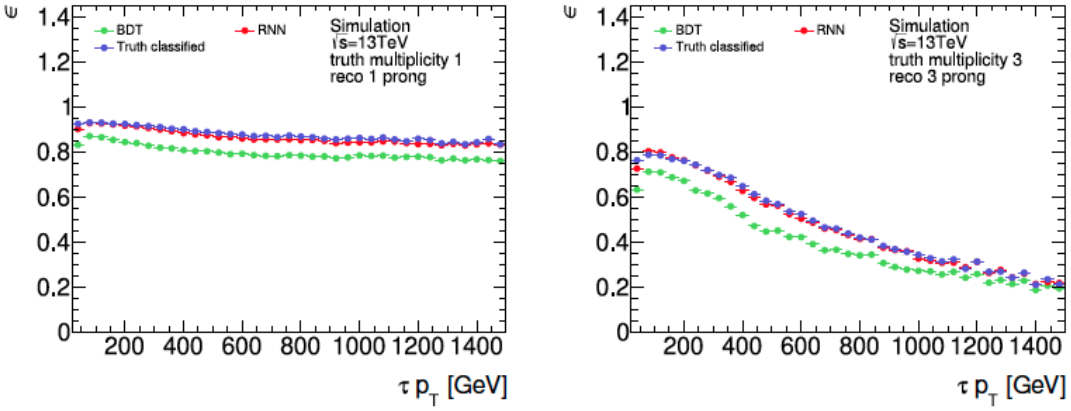

Figure 8.17 comparison of track classification performance attained by BDT and RNN classifiers for *(left)* 1-prong and *(right)* 3-prong simulated τ-lepton decays. The RNN classifier is found to attain appreciably better performance across values of p_T for all track multiplicities, agreeing strongly with 'true' MC performance. Figures extracted from [150].

Furthermore, the RNN track classifier assigns the track multiplicity of the τ jet candidate as either 1 p or 3 p as a function of the number of TTs present in the jet-cone. The performance of the BDT-based and RNN-based track classifiers are compared in Figure 8.17.

Given that τ jets are necessarily hadronic in nature, local cell weighting (LCW) is applied to τ candidates during the seeding step. An additional calibration accounting for the differences in QCD jets and jets originating from τ decays is applied in the form of the *Tau energy scale* (TES) calibration [151]. The TES aims to improve the energy resolution of τ candidates by scaling their energy back to the 'true' visible energy. This calibration factor is derived from MC samples which account for the different jet energy responses between τ jets and QCD jets. Additional pile-up corrections are also applied.

The final step of the τ reconstruction chain is that of jet identification. Tracks which survive the pre-selection process described previously are presented to a further RNN [152], trained to distinguish τ jet candidates from jets initiated from quarks or gluons. This RNN is trained on a combination of low-level variables, including calorimeter clusters and charged inner detector tracks, and high-level variables with significant discriminating power — replacing the BDT previously used for the same purpose.

This chapter has introduced the algorithmic methods utilised in the reconstruction of raw ATLAS system data into low-level and high-level physics objects, such that physical events of interest may be selected for analysis.

If evidence for the decay of LLPs to τ leptons is to be rendered detectable, additional reconstruction methods tailored to the signature of displaced τ production are required. A presentation of such methods now follows.

amsmath mathtools accents algorithm2e

“*Exploration is wired into our brains. If we can see the horizon, we want to know what's beyond.*”

— Buzz Aldrin

9

Displaced τ-Lepton Identification

The search for displaced τ-leptons necessitates the development of bespoke reconstruction tools with which to expose new phase-space to the light of experimental scrutiny. This chapter reviews the algorithmic methods employed in the identification of detector signatures associated with displaced-τ production, before introducing the high-level trigger chain developed in this thesis with which displaced-τ tracks may be captured more efficiently. It is with this experimental toolkit that the search for new physics presented herein may venture beyond the horizon of existing experimental coverage.

9.1 Overview

If a search for displaced τ production is to be conducted, bespoke identification and trigger algorithms must be developed such that the relevant phase-space is rendered accessible. This chapter introduces the displaced τ track classifier and identification algorithms — optimised for displaced signatures in the thesis [150] submitted by Estifa'a Zaid — in conjunction with the LRT-based displaced τ trigger developed in the course of this thesis. Each algorithm is now discussed in turn.

9.2 Displaced τ Classification

Section 8.3.6 introduced the track classification and identification algorithms employed at the ATLAS experiment for the purpose of τ-lepton reconstruction. Such algorithms were designed to identify the presence of τs among the decay products produced by promptly-decaying particles, such as the Higgs boson. As such, nominal classification algorithms are ill-equipped to efficiently identify the presence of displaced τ-leptons originating from the decay of LLPs.

This observation motivates the optimisation of τ identification algorithms for the efficient detection of displaced signatures. The displaced τ track and ID classifiers — developed outwith this thesis in [150] — are used extensively in the analysis described in subsequent chapters, and are incorporated into the trigger algorithm developed in this thesis. Their operation and performance are described below.

9.2.1 Displaced τ Track Classifier

The τ track classifier introduced in Section 8.3.6 was developed with a view towards the classification of promptly-produced τ tracks. Its performance with regards to track content produced in the decay of τ-leptons at some displacement within the detector is therefore sub-optimal.

The τ track classifier was accordingly re-developed in an effort to improve the efficiency of displaced track classification. Accurate track classification enables more efficient overall displaced τ reconstruction, as classified tracks serve as input to the downstream τ identification classifier.

The dedicated displaced track classifier was developed with an identical architecture to the RNN used for 'nominal' track classification. This was trained on a long-lived $Z' \rightarrow \tau\tau$ MC signal sample, with a QCD di-jet MC sample serving as background training data. The instantiation of mass and lifetime parameters was varied in the signal sample such that the RNN was exposed to a variety of possible signal signatures.

The classification efficiency with respect to displaced charged track multiplicity was derived with MC simulations of the $Z' \rightarrow \tau\tau$ and $\bar{\tau}\bar{\tau} \rightarrow \tau\tau$ processes, and is presented in Figure 9.1. Sub-figures (a) and (b) show the efficiency attained by both the nominal and displaced track classifier for true 1-prong and 3-prong $Z' \rightarrow \tau\tau$ displaced tracks, while sub-figures (c) and (d) present the same comparison of performance for true 1-prong and 3-prong $\bar{\tau}\bar{\tau} \rightarrow \tau\tau$ displaced tracks. Each efficiency is derived for MC tracks which meet the track quality criteria imposed by $p_T > 10$ GeV and $|\eta| > 2.5$. It may be observed that, under the prompt track classifier, the majority of 1-prong τ tracks are misclassified as belonging to the category of 0-prong; such erroneous behaviour is shown to be yet more acute for the true 3-prong scenario. This inefficient classification of displaced track content would severely deplete the statistical power of future analyses, diminishing the source of charged track decays with which to perform a search. The newly-developed displaced classifier is found to migrate such mis-classifications to the appropriate true category, attaining significantly higher classification efficiency in both 1-prong and 3 prong τ decays. Future analyses therefore stand to profit from the higher statistics of correctly-classified displaced tracks with visible, non-zero track multiplicity.

It may be observed from Figure 9.1 that the performance attained on the Z' MC sample, which served as the training set for the RNN, generalises well to the $\bar{\tau}\bar{\tau}$ sample, to which the RNN was not exposed during the training process. This stands to provide reassurance that the underlying neural network was not overtrained, verifying that the power of the RNN to discriminate between tracks of varying classification is independent of the τ production process.

The efficiency with which displaced τ tracks may be classified is significantly improved by the development of a dedicated RNN for this task. Consideration of the subsequent step in the τ reconstruction chain is now due: displaced τ identification (ID).

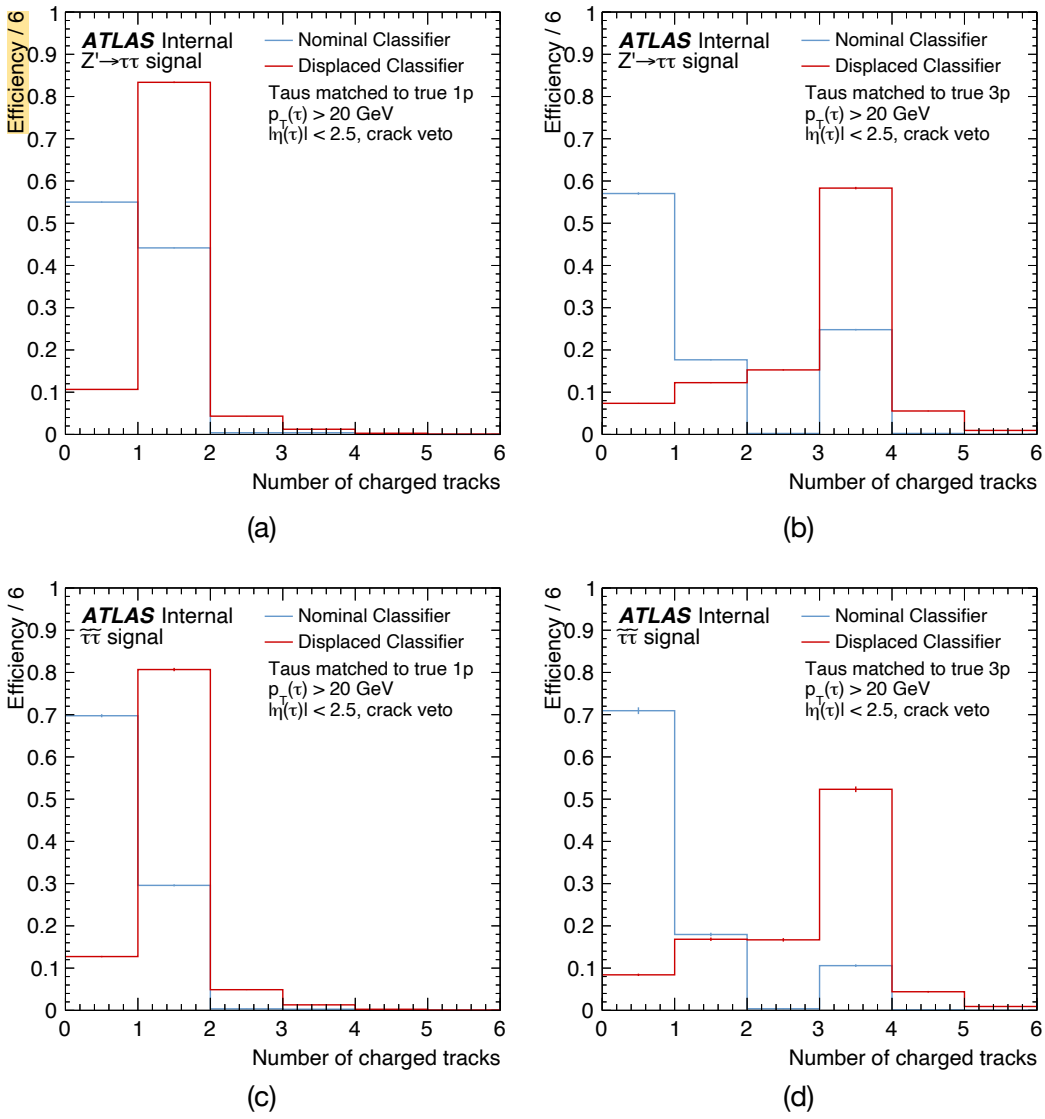

Figure 9.1 track classification efficiencies for the nominal (*blue*) and displaced (*red*) hadronic-τ track classifiers. Efficiences are derived for MC samples corresponding to the (*top*) $Z' \to \tau\tau$ and (*bottom*) $\widetilde{\tau}\widetilde{\tau} \to \tau\tau$ processes, for (*a*,*c*) true 1p and (*b*,*d*) 3p τ-decays. The newly-developed displaced classifier is observed to out-perform the nominal classifier in the classification of all displaced multi-prong τ-decays, regardless of the process of origin. The re-trained RNN is migrating tracks which were previously mis-classified as 0p to their correct multi-prong classification — enabling future analyses to profit from the greater acceptance of multi-pronged displaced decays.

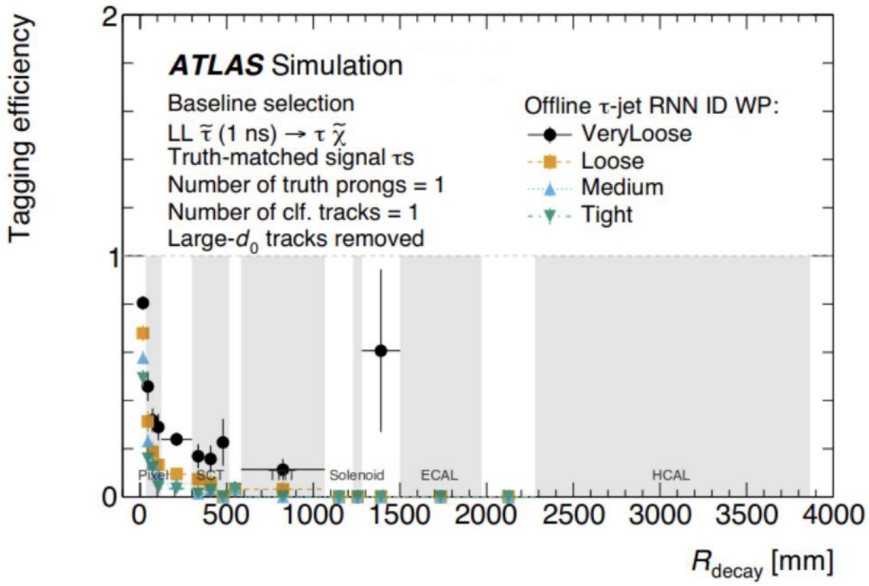

Figure 9.2 identification efficiency of nominal RNN with regards to single-prong displaced τ-leptons produced in MC τ̃ decays. Large radius tracks are removed from the MC sample data such that the efficiency of the RNN may be determined in isolation, without influence from the track classifier. Efficiency is reported as a function of the decay radius of the parent long-lived τ̃, equivalent to the production radius of the displaced τ. It may be observed that poor τ classification efficiency is attained across all working points for all non-negligible values of displacement. Source: [150].

9.2.2 Displaced τ ID

The recurrent neural network (RNN) used for the purpose of τ-lepton identification was introduced in Section 8.3.6. This classifier was re-trained such that its performance with respect to the identification of displaced τ leptons may be optimised.

Figure 9.2 visualises the classification, or tagging, efficiency attained by the pre-existing RNN ('nominal RNN' hereinafter) as a function of the decay radius of the τ's parent particle. This is evaluated on a long-lived τ̃ MC sample, with a proper τ̃ lifetime of 1 ns. Consideration is restricted to τ-leptons of single charged track multiplicity, with the exclusion of large radius tracks (such that the inefficiency of the nominal track classifier does not skew the observed distribution). It may be observed that the efficiency of the nominal RNN classifier drops sharply with τ displacement across all working points.

The poor performance exposed in Figure 9.2 is not unexpected. The nominal

Displaced τ-ID Working Point	$Z' \to \tau\tau$ (1-Prong)	$Z' \to \tau\tau$ (3-Prong)
`VeryLoose`	95%	95%
`Loose`	85%	75%
`Medium`	75%	60%
`Tight`	60%	40%

Table 9.1 working point definitions in terms of signal selection efficiency, as determined on 1-prong and 3-prong $Z' \to \tau\tau$ MC samples. Tighter selections apply a more stringent background rejection at the expense of signal acceptance.

RNN was trained on a sample of MC events corresponding to the $\gamma^* \to \tau\tau$ process, in which τ-leptons are produced in the prompt decay of (virtual) photons. The machine learning classifier was therefore exclusively exposed to promptly-produced τ signatures in the training process. The optimisation of the RNN for the identification of displaced τ signatures is centred upon the modification of this training process — the 'displaced RNN' is trained on MC samples which contain displaced τ content.

The architecture of the RNN [153], in addition to the variables which constitute the input to the network [153], were unchanged in this process. The displaced RNN was trained on track and cluster variables, in addition to derived high-level variables deemed to be of discriminating utility [153]. The training scheme for the displaced neural network, however, employed the long-lived $Z' \to \tau\tau$ MC sample as its signal training set, and was therefore exposed to large radius tracks and larger values of displacement excluded from the $\gamma^* \to \tau\tau$ training sample. Training data-sets were flattened in variables of pile-up, $\langle\mu\rangle$, and transverse momenta, $p_T(\tau_{had})$, such that a uniform response is achieved across different phase-space. raining is performed separately per τ track multiplicity, owing to the differing detector signatures produced by 1-prong and 3-prong τ-leptons. Both standard track and large radius track content was present in the training data-set.

Figure 9.3 presents a Receiver Operator Characteristic (ROC) curve [142] for the displaced RNN against competing ID taggers. This quantifies the signal acceptance efficiency against the background rejection achieved by each algorithm. A comparison is provided for (*blue*) the RNN trained on the displaced $Z' \to \tau\tau$ process — as deployed in the ATLAS experiment, against (*green*) a RNN trained on the ~~$\tau\tau$~~ $\to \tau\tau$ process, (*purple*) a RNN trained on the original $\gamma^* \to \tau\tau$

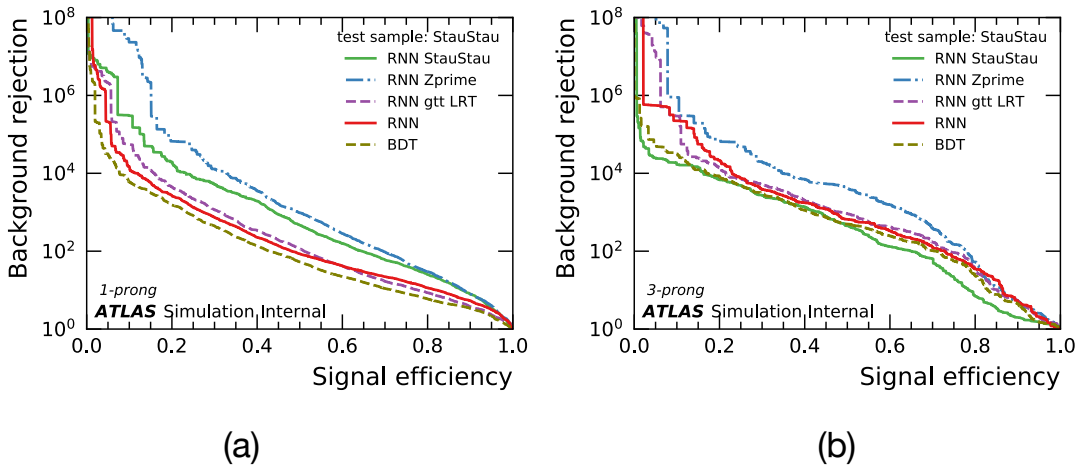

Figure 9.3 displaced τ identification efficiency *vs.* background rejection with respect to (*a*) 1-prong and (*b*) 3-prong displaced τ track multiplicities. Performance is shown with respect to simulated $\tilde{\tau}\tilde{\tau} \to \tau\tau$ MC samples. *Green*, *blue*, and *pruple* lines show the performance achieved by the RNN classifier respectively trained on the displaced track content of $\tilde{\tau}\tilde{\tau} \to \tau\tau$, $Z' \to \tau\tau$, and (LRT) $\gamma^* \to \tau\tau$ MC samples. Nominal RNN and previous generation BDT performances are shown for comparison. Figure extracted from [150].

process enhanced with LRT content, and the respective (*red*) RNN and (*brown*) BDT nominal algorithms. Across all values of track multiplicity, the displaced RNN τ-ID trained on the long-lived Z' sample attains better performance than both nominal and alternative displaced algorithms.

To be of analytic utility to the search in question, the performance shown in Figure 9.3 must be translated into a signal acceptance efficiency as a function of displacement. This is presented in Figure 9.4. Shown for all τ-track multiplicities, Figure 9.4 (a) shows the performance of the Z'-trained displaced RNN on the Z' MC sample, whereas sub-figure (b) shows the same performance of the same displaced RNN on the $\tilde{\tau}\tilde{\tau}$ MC sample. For each respective selection working point — as defined in Table 9.1 — the displaced RNN is shown to achieve significantly greater signal acceptance efficiency at all values of displacement. Working points with more stringent appetites for background events are observed to exclude a greater degree of displaced-τ signal, as expected. It is clear from the performance captured in this figure that the displaced τ-ID will afford the search conducted in this thesis additional analytic leverage which would not have otherwise been available under with the nominal ID tagger. This is most true for high values of R_{decay}, the very phase-space of interest to this analysis.

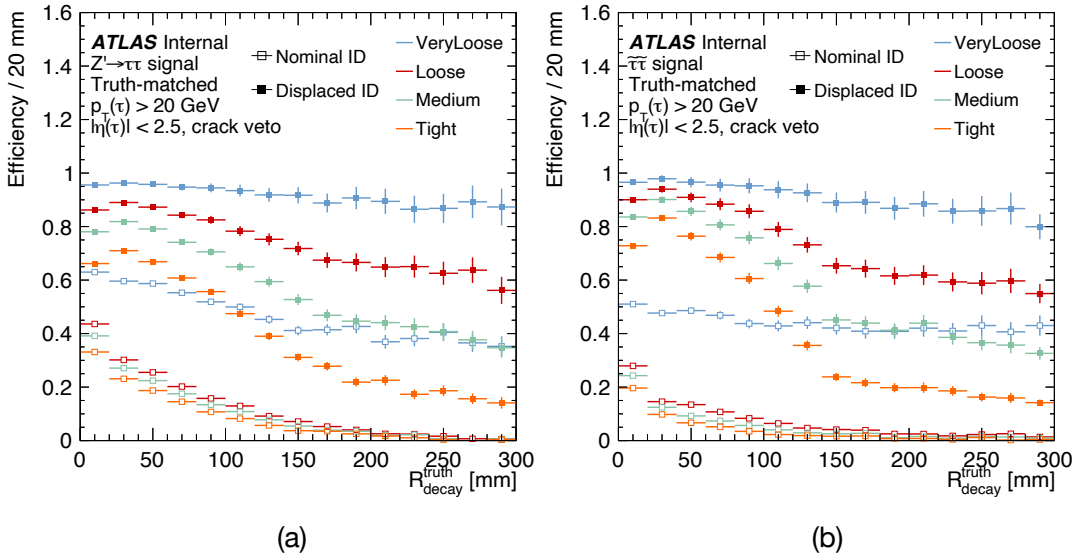

Figure 9.4 signal efficiency attained for each working point selection by both the nominal and displaced τ-ID taggers. Efficiencies are derived for the (*a*) $Z' \rightarrow \tau\tau$ and (*b*) $t\bar{t} \rightarrow \tau\tau$ MC samples (inclusive of all track multiplicities). The displaced ID is found to achieve significantly higher signal acceptance across all MC decay displacements, with starker improvements visible at greater values of R_{decay}. Nominal and displaced efficiencies begin to converge as R_{decay} approaches zero, as expected.

The efficient performance of the displaced RNN at large values of τ displacement must be secured if the classifier is to provide utility to a search for displaced τ-leptons. Figure 9.4 depicts the efficiency gains attained by the displaced RNN trained on long-lived Z' sample events ('LLZ RNN') in comparison to the performance of the nominal RNN. Significantly improved efficiency is observed across all working points, with improvements sustained at high values of τ production radius.

The RNN used for displaced-τ identification will feature in the analysis of ATLAS data presented in subsequent chapters. Moreover, the displaced RNN also features in the HLT developed in this thesis for the efficient triggering of displaced-τ production. A discussion of this trigger chain now follows.

9.3 Displaced τ Trigger Development

Equipped with τ-lepton identification algorithms optimised for the detection of displaced signatures, trigger algorithms tasked with the data acquisition of third-

generation leptons may be optimised. Nominal high-level triggers (HLTs) are optimised for the detection of τs produced in the prompt decay of particles, such as the Higgs boson. Such algorithms are not well-equipped to facilitate the efficient detection of τs produced in the decay of LLPs.

To this end, a suite of HLT algorithms optimised for the signature of displaced τ-leptons have been developed. The efficient detection and capture of events in which such a signature is present enables the downstream analysis of displaced-τ events to be undertaken with greater statistical precision. Outwith this thesis, a trigger chain was developed which employed the newly-optimised identification algorithm presented in Section 9.2 in conjunction with the execution of conventional tracking techniques. This thesis presents an additional HLT algorithm, configured to similarly employ the τ identification classifier optimised for displaced signatures, while performing a modified tracking sequence in which the ATLAS large radius tracking (LRT) algorithm is executed. The improved acceptance of displaced tracks provided by the LRT sequence enhances the ability of the HLT to efficiently capture events in which a displaced τ-lepton is produced. While its development does not benefit the search conducted on the existing run-2 LHC data-set presented herein, the LRT-based trigger chain will equip future run-3 searches to expand their sensitivity to new phase-space.

9.3.1 LRT Trigger Chain

A bespoke high-level displaced-τ trigger algorithm for the ATLAS experiment was developed in the course of this thesis. Existing trigger chains optimised for displaced-τ signatures, developed outwith this thesis, leverage improved displaced-τ identification techniques in conjunction with standard track reconstruction practises. The trigger developed in this thesis seeks to further advance the efficiency with which displaced τ-leptons may be acquired at trigger-level by invoking the improved reconstruction performance provided by the ATLAS large radius tracking (LRT) algorithm — expanding the acceptance of tracks emanating from displaced vertices.

A Monte Carlo study was undertaken to ascertain the extent to which a search for displaced τs may profit from a trigger equipped with large radius tracking capabilities. A MC sample was generated corresponding to the decay of stau-stau sparticle pairs ($\tilde{\tau}\tilde{\tau}$) into pairs of SM τ-leptons. Stau particles were generated with a mass of $m_{\mathsf{LLP}} = 100\ \mathsf{GeV}$ and a proper lifetime of $\tau_{\mathsf{LLP}} = 1.0\ \mathsf{ns}$, such that

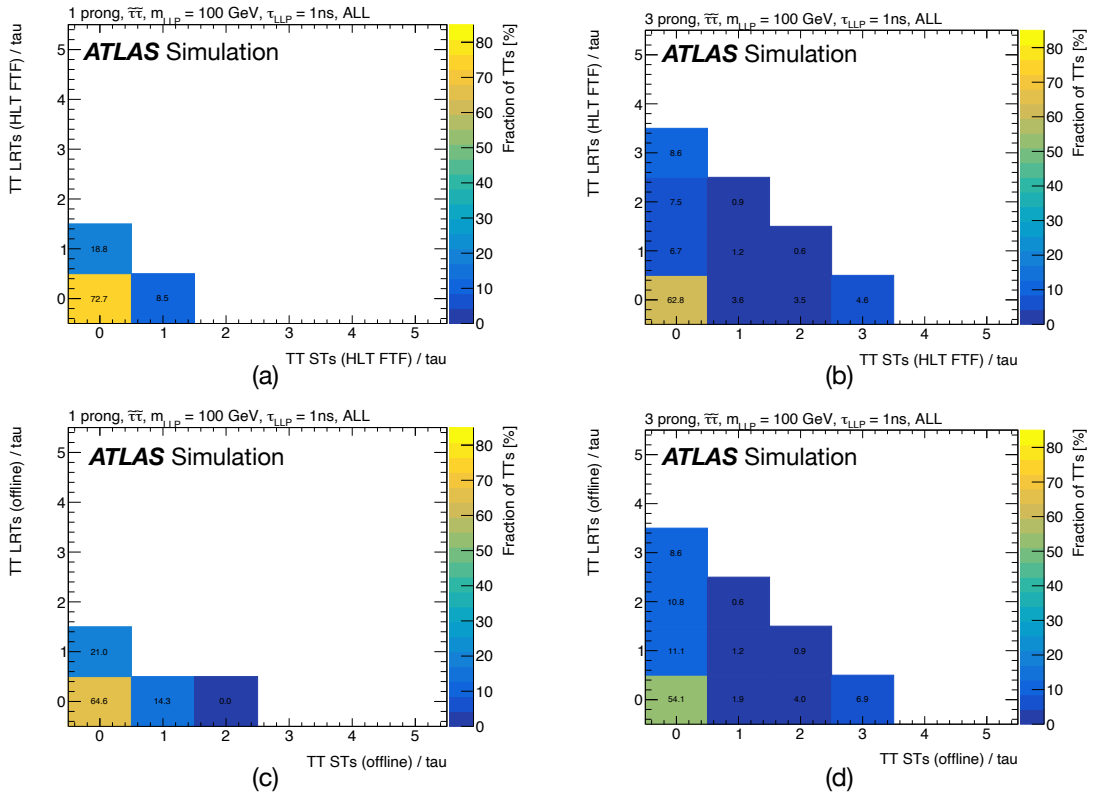

Figure 9.5 Monte Carlo study of track content produced by simulated displaced τ production. Standard track (ST) and large radius track (LRT) counts produced by charged truth τ (TT) content are compared on an event-by-event basis for both 1-prong (*a*,*c*) and 3-prong (*b*,*d*) charged track multiplicities. Tracks obtained via ATLAS Fast Track Finder (FTF) (*a*,*b*) and offline (*c*,*d*) reconstruction are considered. Approximately 30% of simulated events are found to produce no STs (0^{th} bin of bottom-axis) while producing one or more charged tracks of a large radius nature. Such statistics motivate the development of an LRT-based trigger chain with which to capture said events.

each τ may serve as a representative LLP and a source of displaced τ content. The number of standard tracks (STs) and large radius tracks (LRTs) produced by each parent LLP decay was determined for both 1 prong and 3 prong child τ-leptons. Track multiplicities were obtained for two different track reconstruction methodologies: ATLAS 'fast tracks', and offline tracks.

It may be observed from Figure 9.5 that approximately 30% of simulated events containing displaced τ content produce one or more large radius track while producing no standard tracks. In the case pertaining to the offline tracks originating from 3-pronged displaced τs, the number of events which exclusively produce large radius tracks is found to exceed 30% — as is quantified in Figure 9.5 (d). Trigger systems which are limited to the reconstruction of standard tracks

are likely blind to the vast majority of such events. The availability of a high-level trigger chain equipped with large radius tracking capabilities is therefore projected to furnish a run-3 displaced τ search with a data-set containing ~ 30% more events, substantially expanding the statistical resources with which such a search may be conducted. Figure 9.5 further quantifies the scenario where LLP decays produce one or more standard tracks in the absence of large radius tracks — this scenario may be addressed at the analysis level by taking the logical-OR combination of both the nominal and LRT-enabled trigger, thereby achieving acceptance of the full τ lifetime spectrum. The development of a LRT trigger chain is accordingly well-motivated.

The developed LRT-based trigger commences execution with the identification and extraction of τ candidate seeds from calorimeter deposits [154], after which a tracking sequence is subsequently performed. The ATLAS L1 trigger system accepts hardware information from the electromagnetic and hadronic calorimeters together with the muon spectrometer. The software-based HLT system, contrastingly, accepts additional input information from the inner detector. Efficiencies of HLT trigger chains may therefore be computed with respect to such 'online' tracks in the inner detector.

Online tracking is typically the most computationally expensive element of HLT processes. The online tracking sequence conventionally proceeds in two stages: (1) ATLAS Fast Tracking [155], and (2) Precision Tracking [155].

The ATLAS Fast Track Finder [155] coordinates the seeding and formation of tracks at HLT level using information from the inner detector. This procedure is optimised to prioritise speed of execution and computational efficiency such that it may be executed within the timing constraints and CPU budget of the HLT system. In so doing, a degree of precision is sacraficed when compared to the full offline tracking procedure described in Section 8.2.1. Tracks obtained with the Fast Track Finder are termed *FTF tracks*.

The second tracking step — Precision Tracking — seeks to reconstruct tracks in a manner much closer to full offline reconstruction. Precision tracks are seeded from the spatial coordinates of the previously-determined FTF tracks, which are then subjected to ambiguity removal and TRT-extension in the same fashion as described in 8.2.1. The execution of Precision Tracking may in this manner increase the purity of trigger-level tracks, where the CPU budget of a given HLT chain may accommodate the increased complexity of this additional step.

Given the computational intensity of the ATLAS LRT algorithm, the LRT-based displaced-τ trigger was designed to undertake a single-step tracking sequence. In the context of large radius tracking, fast tracking may be performed in one of two modes: 'RoI-based' fast tracking, or 'full-scan' fast tracking. The former method is a single-step process where RoIs are utilised for every tracking hit. The full-scan method unfolds in two steps: initially executing the standard non-LRT tracking sequence, before executing the LRT algorithm exclusively on unused hits. While the execution of the LRT algorithm solely on hits discarded by the standard tracking algorithm — as is done in offline reconstruction — has a number of advantages, the single-step RoI-based approach is implemented in this trigger. In this manner the computational complexity — and, by extension, the execution time — of the HLT chain may be minimised.

This design decision has implications for the choice of other trigger design parameters: most notably, the designated size of the RoI. As is illustrated in Figure 9.6, a typical two-step tracking sequence at trigger level is undertaken with two RoIs of differing size. This reflects the coarseness and resolution of the respective Fast Tracking and Precision Tracking sequences. It was determined that, in the absence of a second Precision Tracking stage, the LRT-enabled displaced-τ trigger could accommodate the specification of a large RoI without resulting in a prohibitive computational load. For this reason, an RoI-size of $\Delta\eta \times \Delta\phi = 0.4 \times 0.8$ was implemented — equal in magnitude to the larger of the two RoIs used in the two-step LRT trigger chains detailed in [1]. The precedent for this parameterisation of RoI supports the assessment that such a catchment area should fall within the computational budget of the HLT system. The choice of a larger RoI enables the acceptance of displaced tracks to be maximised.

Upon the completion of the HLT tracking stage, the displaced-τ trigger chain is then designed to schedule the execution of the *ID stage*. At this point in the execution, tracks obtained from the previous step of the chain are fed to the RNN classifier optimised for the detection of displaced τ-leptons, as introduced in the previous section. τ decays to which the RNN ID algorithm ascribes a score in excess of the chosen `Medium` working point are accepted by the trigger and stored for subsequent analysis.

The efficiency of the HLT chain may be expressed with respect to both FTF and offline tracks. Figure 9.7 presents the trigger signal efficiency as a function of various track kinematic parameters, computed for a reference 'truth' (*i.e.* MC) FTF track from the $\tau\tau \to \tau\tau$ process, where the reference track falls within

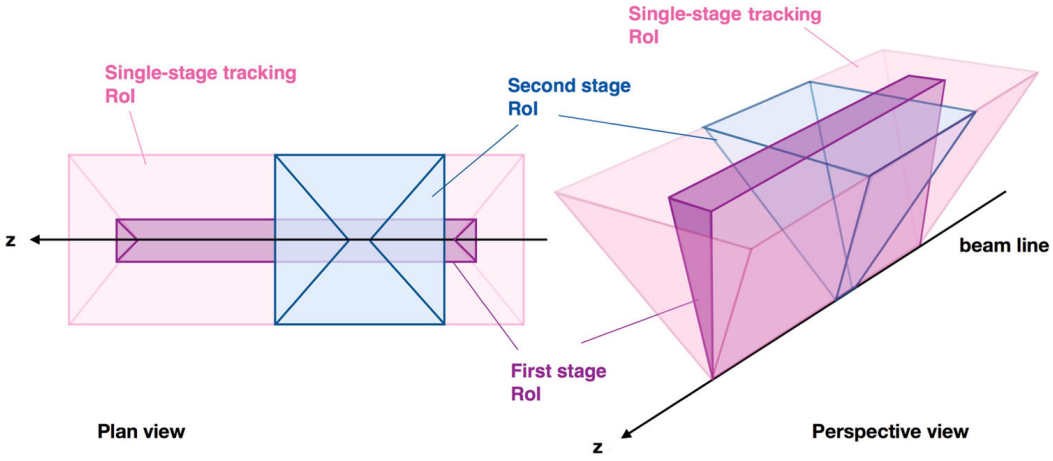

Figure 9.6 diagram depicting the geomatrical form and size of trigger regions of interest (RoIs) with respect to the z-axis of the ATLAS coordinate system. Typical HLT chains equipped with a two-step tracking sequence define two RoIs of differing sizes, whereas triggers which employ single-step tracking procedures operate with a single RoI of fixed size. Source: [156].

the inner detector. Efficiencies are determined for a τ p_T threshold of 25 GeV — this threashold was chosen for testing purposes, such that the available test statistics were maximised. A threshold of 160 GeV was envisaged as the target for analysis deployment at the time of design. It may be observed in Figure 9.7 that the attained efficiency with respect to FTF tracks is approximately constant across the pseudo-rapidity coverage of the ATLAS detector, with an approximately flat distribution in η. A similarly flat distribution is attained with respect to ϕ, suggesting strong trigger acceptance across the geometry of the detector. The τ p_T threshold is reflected in the 'turn-on' curve of the efficiency distribution as a function of transverse momentum, demonstrating expected sub-optimal acceptance of τ-leptons below this threshold. The performance of the new LRT-enabled displaced τ trigger is found to outperform the non-LRT LLP trigger by approximately 30%, with notably greater efficiency attained at high values of track impact parameter, d_0 and z_0. This result exemplifies the extend to which a future analysis may profit from greater acceptance of LLPs which produce displaced tracks.

The *offline* efficiency with which the LRT-based trigger chain acquires displaced τ tracks is shown in Figure 9.8 and Figure 9.9. the offline track efficiency is shown in Figure 9.8 against the decay radius of long-lived MC τs on various scales, in addition to the transverse momenta of the τ. It may be observed that the efficiency of the LRT-based trigger, shown in blue, exceeds that of

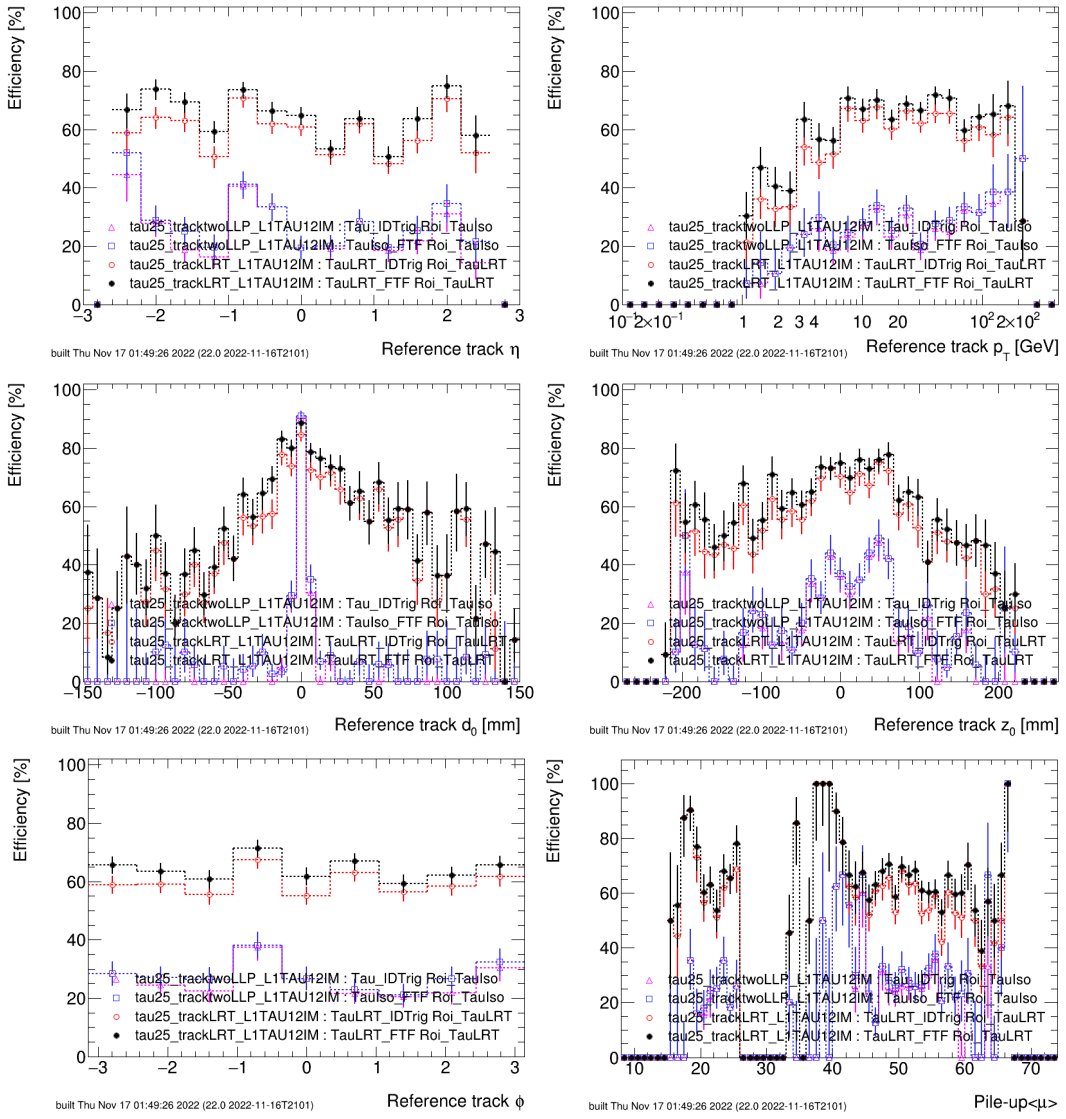

Figure 9.7 online efficiency with respect to various kinematic quantities, as determined for a high-quality 'reference' track in the inner detector. Pink and blue points show the performance of the LLP trigger with standard tracking with respect to inner detector and FTF and tracks, respectively; whereas the red and black points show the LRT-equipped trigger performance, also with respect to inner detector and FTF tracks, respectively. It may be observed that the inclusion of large radius tracking achieves significantly greater online efficiency across all kinematic variables.

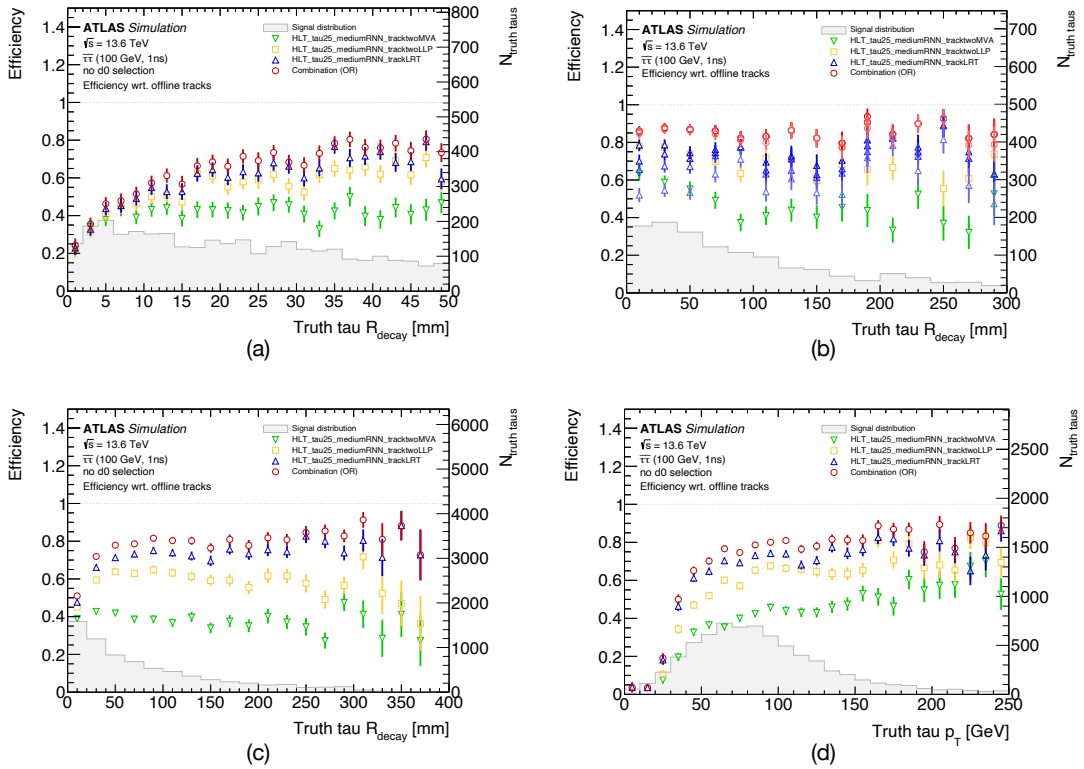

Figure 9.8 efficieny of LRT-based HLT chain (blue triangle) with respect to offline-level simulated tracks, compared to LLP trigger with standard tracking (yellow square) and nominal τ trigger chain (green triangle). Results obtained from an MC sample of $x\,\tilde{\tau}\tilde{\tau} \rightarrow \tau\tau$ events, where parent $\tilde{\tau}$ particles are generated with a mass of 100 GeV and a proper lifetime of 1 ns. Efficiencies are presented as a function of (*a-c*) τ decay radius on various distance scales, and (*d*) τ transverse momentum. LRT-chain efficiency is observed to exceed alternative chains at greater displacement and momentum, with the logical combination of all three triggers (red circle) achieving best performance.

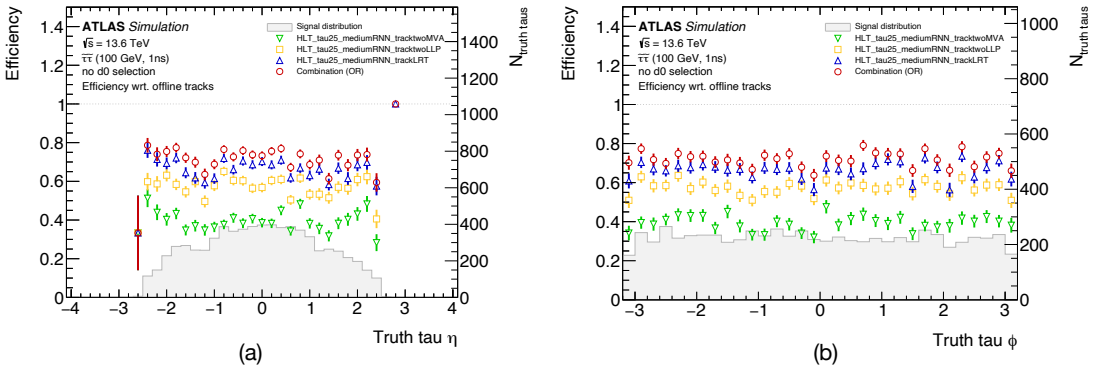

Figure 9.9 LRT-based HLT chain (blue triangle) with respect to offline-level simulated tracks, compared to LLP trigger with standard tracking (yellow square) and nominal τ trigger chain (green triangle), as a function of τ (*a*) pseudo-rapidity and (*b*) azimuthal position. Efficiencies computed with respect to the same MC sample as specified above. LRT-capabilities achieve greater efficiency at trigger-level across the detector geometry, with the logical combination of all three trigger options attaining optimal performance.

Trigger	p_T Cut	I/P Rate [Hz]	Express Rate [Hz]	Event Time [ms]
Prompt Chain	25 GeV	740,503.0	93,589.0	604.4
(non-LRT)	160 GeV	158.0	200.0	2,252.0
LLP Chain	25 GeV	740,503.0	134,087.1	299.7
(LRT)	160 GeV	1,996.1	243.8	1,025.0

Table 9.2 input rate, express rate, and execution time per event for (*top*) the nominal (prompt) τ trigger chain, and (*bottom*) the LRT-equipped displaced τ trigger chain. Performance is shown for two p_T thresholds of 25 GeV and 160 GeV for each trigger.

both the nominal trigger chain and the LLP chain equipped with standard tracking capabilities. This observation is most acutely expressed at high values of displacement — a strong endorsement of the approach adopted. Figure 9.8 (b) further conveys the efficiency attained by each ID working point, with lighter variations in hue corresponding to looser ID selection criteria. The logical `OR` combination of all three chains is shown to achieve the most optimal efficiency, with a ~ 25%-30% improvement over the nominal chain alone.

Figure 9.9 similarly depicts the same efficiency as a function of the kinematic distributions formed by the τ pseudo-rapidity and azimuthal coordinates. The observed efficiency comparison of trigger chains is consistent with that presented in Figure 9.8. Further, efficient coverage across the geometry of the detector is confirmed.

In addition to the time consumed in its execution, an important metric of consideration in the development of a trigger is the rate at which it outputs accepted event data. This is known as the *express rate*. Table 9.2 compares the express rate and HLT execution time for the newly developed LRT trigger chain, and that exhibited by the nominal τ chain without LRT capabilities. Such figures are shown for two $\tau\, p_T$ thresholds — a test threshold of 25 GeV and a more physically reasonably threshold of 160 GeV — together with the input ('I/P') event rate for each given threshold. The execution time is found to be competitive with the exiting nominal chain, executing in approximately half the time of the non-LRT counter-part. The express rate, while not unreasonable, may be judged to be on the higher end of what can reasonably be expected of the HLT system. Future work to reduce this rate may include the imposition of a more stringent

p_T threshold, or the reduction of the RoI size.

The LRT-based trigger chain developed in the course of this thesis enables displaced τ-lepton production to be triggered upon with approximately 30% greater accuracy. The additional events rendered accessible by this trigger will profit future analyses performed on the ATLAS run-3 data-set in the search for new physics, helping to cast light upon new regions of lifetime phase-space.

This chapter has introduced the suite of identification and trigger algorithms with which the search for displaced-τs presented in this thesis may be conducted. The incorporation of large radius tracking at the ATLAS trigger level has been shown in this thesis to trigger on the production of displaced τ-leptons with approximately 30% greater efficiency than is attainable with existing trigger algorithms.

Having examined the theoretical and experimental basis of this thesis, it is timely to provide an overview of the sequence in which an ATLAS search for new physics may be conducted. The manner in which the theoretical and experimental material discussed herein is interwoven such that ATLAS data may be analysed for evidence of new physics is encapsulated overleaf.

Part III

Displaced τ-Lepton Analysis

“*All we can do is search for the falsity content of our best theory.*”

— Karl Popper

10

Analysis Overview

Previous chapters have surveyed the theoretical and experimental underpinnings of the search for new physics advanced in this thesis. With an understanding of such topics now in place, the analytical and statistical design of the scientific investigation may be examined. With reference to the theoretical simulation procedures, experimental detection and reconstruction methods, and bespoke displaced signature identification algorithms introduced in the preceding chapters, this chapter provides an overview of the analysis strategy employed in the search for displaced τ-lepton production.

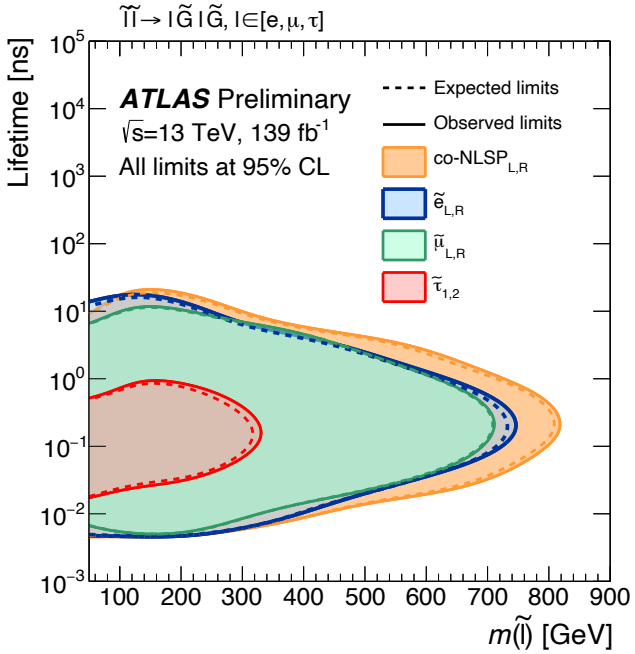

Figure 10.1 existing experimental coverage of displaced-τ phase-space does not extend far in parent LLP lifetime. Source: [157].

10.1 Search for Displaced τ-Leptons

This thesis has thus far explored both the theoretical and experimental motivation for the existence of long-lived particles beyond the SM. It has been argued herein that the large mass of third-generation fermions may afford τ-leptons a privileged role in the kinematics of LLP decays, which may constitute a favoured child-state. It is in accordance with such reasoning that this thesis advances the search for long-lived particles decaying to τ-leptons with the ATLAS experiment.

Despite the uniquely strong motivation associated with their large mass, existing coverage of signatures corresponding to displaced leptons has largely neglected the third generation. This may be observed clearly in Figure 10.1. The gap in existing coverage which corresponds to this phase-space is therefore ripe for further exploration.

The decay of LLPs to τ-leptons would give rise to the signature of τ production and subsequent decay occurring at appreciable displacement within the detector. The production of τ-leptons at high displacement is manifested by the signature of displaced vertices, which are characterised by large values of longitudinal impact parameter relative to the associated uncertainty: d_0/σ_0. The subsequent decay of

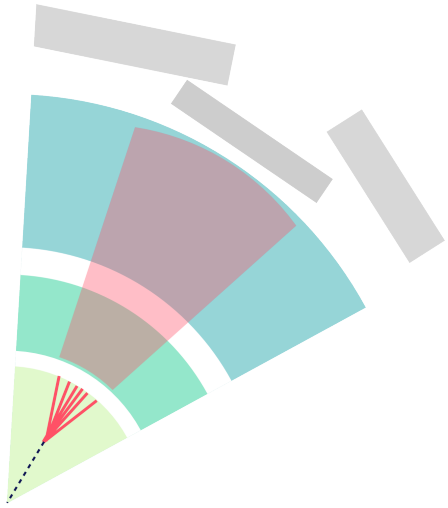

Figure 10.2 illustration of a long-lived particle decaying hadronically within the ATLAS inner detector, giving rise to the signature of a displaced jet in the calorimeters. Figure taken from [158].

τ-leptons may unfold leptonically or hadronically, via the decay modes tabulated in Table 8.2. While decay topologies which include child leptonic states feature in the broader analysis, the scope of consideration is restricted in this thesis to the doubly-hadronic decay channel: $\tau_{had}\tau_{had}$. Figure 10.2 illustrates the hadronic decay of a generic LLP state within the inner detector of the ATLAS experiment, depositing energy in the calorimeter systems which may be recognised as a displaced hadronic jet.

Signatures of this nature present unique experimental challenges. This necessitates the development of bespoke reconstruction and identification algorithms with which to conduct this search, as were introduced in Chapter 9. The search is further complicated by the model-independent philosophy to which it subscribes. The design decisions pertinent to the analysis — such as those which concern object selections and background estimation — are not informed by the properties of any particular model of new physics, but rather seek to achieve acceptance of a broad range of possible LLP states. This approach casts the widest possible net with which to capture evidence of new physics, without bias. In this context, MC simulated samples serve as a benchmark with which the sensitivity of the search to well-motivated signatures may be gauged. The $\tilde{\tau}\tilde{\tau} \rightarrow \tau\tau$ process assumes a significant role to this effect.

This search considers events which contain two displaced τ-leptons, emanating from the decay of a long-lived particle beyond the Standard Model. The search presented in this thesis is designed to target the $\tau_{had}\tau_{had}$ channel, where both τs

decay hadronically. The particular design considerations which pertain to this channel are discussed in this chapter.

Expected limits are placed at the 95% confidence level on the number of signal events observed in the phase-space which corresponds to displaced-τ production in the $\tau_{had}\tau_{had}$ decay channel. Models which predict the existence of LLPs may then be compared against this observed limit such that their compatibility with observation may be assessed.

10.2 Analysis Inputs

Data employed in the course of this analysis may be separated into simulated data, and those which are collected by the ATLAS experiment from proton collisions at the LHC. A description of each category now follows.

10.2.1 ATLAS Data

Real collision data recorded in the second proton-proton run of the LHC is analysed in the course of this search, where raw collision events are subjected to the event reconstruction chain described in Chapter 8.

Owing to the computational expense of the algorithm, large radius tracking (LRT) content may only be computed for a sub-set of the full run-2 ATLAS data-set. (The computational load of the ATLAS LRT algorithm was successfully reduced in preparation for run-3 such that this restriction may be avoided.) The full run-2 *pp* data-set is therefore reduced through the application of successive filters such that the execution of the LRT algorithm may be performed within the appropriate reconstruction CPU budget.

The data filtration process is performed on raw detector output (*i.e.* before subsequent re-processing) with the `RPVLL` (*R*-Parity Violating Long-Lived [supersymmetry]) filter, originally developed to aide the search for various SUSY scenarios which incorporate the violation of *R*-parity. The resultant data-stream achieves a reduced event-rate of 50 Hz — which may be regarded as computationally-economical when contrasted with the 1 kHz HLT event-rate. The `RPVLL` stream thereby enables the execution of the ATLAS LRT reconstruction within the allocated CPU time quota.

As LRT content forms an integral component of the LLP search described herein, the ATLAS pp data studied for the purposes of this analysis are reconstructed from the `RPVLL` stream. Object reconstruction is performed on `RPVLL` event data in accordance with the methodologies presented in Chapter 8.

The single-τ and di-τ triggers employed in this search were included in the `RPVLL` stream in 2018. The search for displaced τ-leptons therefore considers data collected from this period onward. This corresponds to a total data-set across all analysis channels of $58.8\,\text{fb}^{-1}$, of which $37.5\,\text{fb}^{-1}$ falls into the $\tau_{\text{had}}\tau_{\text{had}}$ channel — the subject of this thesis.

10.2.2 Simulated Data

Simulated data assumes an important role in the search for displaced-τ production. The development of identification algorithms optimised for the detection of displaced signatures required a source of displaced τ-lepton data with which to train the associated multivariate classifiers. This was achieved with the aide of MC simulation.

More pertinently to the analytical component of the search, simulated data is of critical import in the determination of background process contribution to the measured event yield. While the estimation of the primary background source is obtained in a data-driven manner (as will be shown shorlty), the effect of non-dijet SM processes on the total event yield is directly accounted for with simulated samples. Further, simulation facilitates the determination of the expected sensitivity of the analysis to new physics by serving as a benchmark model against which to gauge the exclusion power of the search strategy.

Simulated data corresponding to signal processes include a $\tilde{\tau}\tilde{\tau} \to \tau\tau$ sample and a $Z' \to \tau\tau$ sample — details of which are tabulated in Table 10.1 and Table 10.2. Supersymmetric $\tilde{\tau}$ decays in this channel were simulated via the MC methods introduced in Chapter 5 with the `MadGraph` (`MadGraph5_aMC@NLO`) general-purpose event generator to leading-order precision. Events were simulated to emerge from pp-collisions at a centre-of-mass energy of $\sqrt{s} = 13\,\text{TeV}$, corresponding to the collision energy of the LHC during the run-2 data-taking period. Samples were generated for various mass and lifetime values in order to assess our sensitivity to each possible instantiation of such values as may be favoured by nature — unknown, of course, *a priori*. Events produced

in accordance with this model were simulated with the NNPDF23LO parton distribution functions.

The p_T spectrum of τ-leptons produced by MC $\tau\tau$ simulation is presented in Figure 10.3. This shows the transverse-momentum for (*top*) leading-p_T τ-leptons and (*bottom*) sub-leading-p_T τ-leptons for a variety of lifetime values. For each value of lifetime, the p_T distribution is shown for a range of mass points.

Samples corresponding to the $pp \to Z' \to \tau\tau$ process are described in Table 10.2. This process was simulated with a combination of the POWHEG and Pythia event generators using the NNPDF23LO parton distribution function values. All signal MC samples are ascribed a numerical weight proportional to the collider luminosity and the cross-section of the corresponding process, such that their event yield matches that which would be expected for the operating conditions of the second run of the LHC. Additional weighting is performed at the event generator level such that predictions remain accurate across different regions of phase-space.

Background processes simulated for analysis include $Z \to \tau\tau$, $t\bar{t} \to \tau\tau$, and di-boson decays: $\{WW, WZ, ZZ\} \to \tau\tau$. While QCD di-jet events represent the primary background process in this search, Monte Carlo representations of such events are not sufficiently accurate for the purposes of this analysis — as such, the contribution of such events to the search background was evaluated in a data-driven manner. This will be examined in the following chapter.

Details of di-boson, Z, and $t\bar{t}$ SM background samples are respectively tabulated in Table 10.3, Table 10.4, and Table 10.5. Di-boson events are produced with the SHERPA MC event generator to next-to-leading-order (NNLO) precision with the NNPDF30NNLO parton distribution set. In keeping with the generation of signal samples, background MC samples are weighted by luminosity, process cross-section, and generator weightings such that the predicted event yield and kinematics reflect expected observations during run-2 LHC conditions.

Having examined the data sources which comprise the subject of this analysis, it is pertinent to define the strategy by which a sub-set of events will be selected. Such considerations begin with the designation of a suitable trigger strategy.

DSID	Process	$(m, c\tau)$	Generator	Prec.	$\sqrt{s}$ [TeV]
399001	$\tilde{\tau}\tilde{\tau} \to \tau\tau$	$(50, 0.01)$	MadGraph	LO	13
399002	$\tilde{\tau}\tilde{\tau} \to \tau\tau$	$(50, 0.1)$	MadGraph	LO	13
399003	$\tilde{\tau}\tilde{\tau} \to \tau\tau$	$(50, 1.0)$	MadGraph	LO	13
399005	$\tilde{\tau}\tilde{\tau} \to \tau\tau$	$(100, 0.01)$	MadGraph	LO	13
399006	$\tilde{\tau}\tilde{\tau} \to \tau\tau$	$(100, 0.1)$	MadGraph	LO	13
399007	$\tilde{\tau}\tilde{\tau} \to \tau\tau$	$(100, 1.0)$	MadGraph	LO	13
399009	$\tilde{\tau}\tilde{\tau} \to \tau\tau$	$(200, 0.01)$	MadGraph	LO	13
399010	$\tilde{\tau}\tilde{\tau} \to \tau\tau$	$(200, 0.1)$	MadGraph	LO	13
399011	$\tilde{\tau}\tilde{\tau} \to \tau\tau$	$(200, 1.0)$	MadGraph	LO	13
399013	$\tilde{\tau}\tilde{\tau} \to \tau\tau$	$(300, 0.01)$	MadGraph	LO	13
399014	$\tilde{\tau}\tilde{\tau} \to \tau\tau$	$(300, 0.1)$	MadGraph	LO	13
399015	$\tilde{\tau}\tilde{\tau} \to \tau\tau$	$(300, 1.0)$	MadGraph	LO	13
399017	$\tilde{\tau}\tilde{\tau} \to \tau\tau$	$(400, 0.01)$	MadGraph	LO	13
399018	$\tilde{\tau}\tilde{\tau} \to \tau\tau$	$(400, 0.1)$	MadGraph	LO	13
399019	$\tilde{\tau}\tilde{\tau} \to \tau\tau$	$(400, 1.0)$	MadGraph	LO	13
399021	$\tilde{\tau}\tilde{\tau} \to \tau\tau$	$(500, 0.01)$	MadGraph	LO	13
399022	$\tilde{\tau}\tilde{\tau} \to \tau\tau$	$(500, 0.1)$	MadGraph	LO	13
399023	$\tilde{\tau}\tilde{\tau} \to \tau\tau$	$(500, 1.0)$	MadGraph	LO	13

Table 10.1 details of simulated data samples corresponding to the long-lived $pp \to \tilde{\tau}\tilde{\tau} \to \tau\tau$ process at a centre-of-mass energy of $\sqrt{s} = 13$ TeV. The MadGraph event generator was used to produce MC samples for various mass ([GeV]) and lifeltime ([ns]) points to leading-order (LO) precision.

DSID	Process	$(m, c\tau)$	Generator	Prec.	$\sqrt{s}$ [TeV]
312565	$Z' \to \tau\tau$	$(200, 0.01)$	POWHEG + Pythia	LO	13
312566	$Z' \to \tau\tau$	$(500, 0.01)$	POWHEG + Pythia	LO	13
312567	$Z' \to \tau\tau$	$(1000, 0.01)$	POWHEG + Pythia	LO	13
312568	$Z' \to \tau\tau$	$(200, 0.1)$	POWHEG + Pythia	LO	13
312569	$Z' \to \tau\tau$	$(500, 0.1)$	POWHEG + Pythia	LO	13
312570	$Z' \to \tau\tau$	$(1000, 0.1)$	POWHEG + Pythia	LO	13
312571	$Z' \to \tau\tau$	$(200, 1.0)$	POWHEG + Pythia	LO	13
312572	$Z' \to \tau\tau$	$(500, 1.0)$	POWHEG + Pythia	LO	13
312573	$Z' \to \tau\tau$	$(1000, 1.0)$	POWHEG + Pythia	LO	13
312574	$Z' \to \tau\tau$	$(200, 10.0)$	POWHEG + Pythia	LO	13
312575	$Z' \to \tau\tau$	$(500, 10.0)$	POWHEG + Pythia	LO	13
312576	$Z' \to \tau\tau$	$(1000, 10.0)$	POWHEG + Pythia	LO	13
312577	$Z' \to \tau\tau$	$(200, 100.0)$	POWHEG + Pythia	LO	13
312578	$Z' \to \tau\tau$	$(500, 100.0)$	POWHEG + Pythia	LO	13
312579	$Z' \to \tau\tau$	$(1000, 100.0)$	POWHEG + Pythia	LO	13

Table 10.2 details of simulated data samples corresponding to the long-lived $pp \to Z' \to \tau\tau$ process at a centre-of-mass energy of $\sqrt{s} = 13$ TeV. The POWHEG and Pythia event generators were used to produce MC samples for various mass ([GeV]) and lifeltime ([ns]) points to leading-order (LO) precision.

DSID	Process	Generator	Prec.	$\sqrt{s}$ [TeV]
363355	$ZZ \to qq\,\nu\nu$	SHERPA	NLO	13
363356	$ZZ \to qq\,ll$	SHERPA	NLO	13
363357	$WZ \to qq\,\nu\nu$	SHERPA	NLO	13
363358	$WZ \to\ ll$	SHERPA	NLO	13
363359	$WW \to qq\,l\nu$	SHERPA	NLO	13
363360	$WW \to l\nu\,qq$	SHERPA	NLO	13
363489	$WZ \to l\nu\,qq$	SHERPA	NLO	13
364250	$\{WW, WZ, ZZ\} \to ll\,ll$	SHERPA	NLO	13
364253	$\{WW, WZ, ZZ\} \to ll\,l\nu$	SHERPA	NLO	13
364254	$\{WW, WZ, ZZ\} \to ll\,\nu\nu$	SHERPA	NLO	13
364255	$\{WW, WZ, ZZ\} \to l\nu\,\nu\nu$	SHERPA	NLO	13

Table 10.3 details of simulated data samples corresponding to SM diboson processes. Samples were simulated at $\sqrt{s} = 13$ TeV to NNLO precision with the SHERPA MC event generator.

DSID	Process	Generator	Prec.	$\sqrt{s}$ [TeV]
700326	$Z \to \tau\tau$	SHERPA	NLO	13
700327	$Z \to \tau\tau$	SHERPA	NLO	13
700328	$Z \to \tau\tau$	SHERPA	NLO	13
700329	$Z \to \tau\tau$	SHERPA	NLO	13
700330	$Z \to \tau\tau$	SHERPA	NLO	13
700331	$Z \to \tau\tau$	SHERPA	NLO	13
700332	$Z \to \tau\tau$	SHERPA	NLO	13
700333	$Z \to \tau\tau$	SHERPA	NLO	13
700334	$Z \to \tau\tau$	SHERPA	NLO	13

Table 10.4 details of simulated data samples corresponding to SM $Z \to \tau\tau$ process. Samples were simulated at $\sqrt{s} = 13$ TeV to NNLO precision with the SHERPA MC event generator.

DSID	Process	Generator	Prec.	$\sqrt{s}$ [TeV]
410470	$t\bar{t} \to \tau\tau$	SHERPA	NLO	13
410471	$t\bar{t} \to \tau\tau$	SHERPA	NLO	13

Table 10.5 details of simulated data samples corresponding to SM $t\bar{t} \to \tau\tau$ process. Samples were simulated at $\sqrt{s} = 13$ TeV to NNLO precision with the SHERPA MC event generator.

10.3 Trigger Strategy

The search for displaced τ-leptons utilises an admixture of single-τ and di-τ triggers. The available prompt trigger chains, despite being developed without a view towards displaced signatures, were found to achieve satisfactory performance for this run-2 analysis. This is a consequence of the fact that such triggers are inclusive of trackless τ-leptons — including those which may emanate from the decay of neutral LLPs. Such events, while triggered by nominal chains, may then benefit from the enhancement in performance attained by the optimised track classifier and offline τ-ID RNN. Subsequent analyses performed on run-3 data may profit from the full efficiency benefits of the displaced LRT-based trigger chain developed in this thesis.

Single-τ triggers employed in this analysis include the HLT_tau160_mediumRNN and HLT_tau200_mediumRNN_tracktwoMVA chains. The former imposes a $p_{\mathrm{T}}(\tau)$ requirement of 160 GeV and schedules the regular τ-ID step. The latter chain requires a transverse momentum value in excess of 200 GeV, and is therefore more susceptible to high-p_{T} background jets — for this reason it schedules the more stringent BDT-based tracktwoMVA classifier, which imposes a stricter track multiplicity and quality requirement to further mitgate background acceptance. The di-τ trigger HLT_tau80_mediumRNN_tau60_mediumRNN is used, with a p_{T} threshold of 80 GeV (60 GeV) applied to the leading-p_{T} (sub-leading-p_{T}) τ-lepton. An additional two chains are included in an effort to exploit the trigger handles provided by the missing transverse energy necessarily produced in the decay of τ-leptons, on account of the child neutrino states. These are HLT_tau35_mediumRNN_xe70 and HLT_tau60_mediumRNN_tau25_mediumRNN_xe50, where the former is a single-τ trigger with a p_{T} threshold of 35 GeV and a minimum MET requirement of 70 GeV, and the latter is a di-τ trigger with respective leading and sub-leading p_{T}-threaholds of 60GeV and 25GeV,

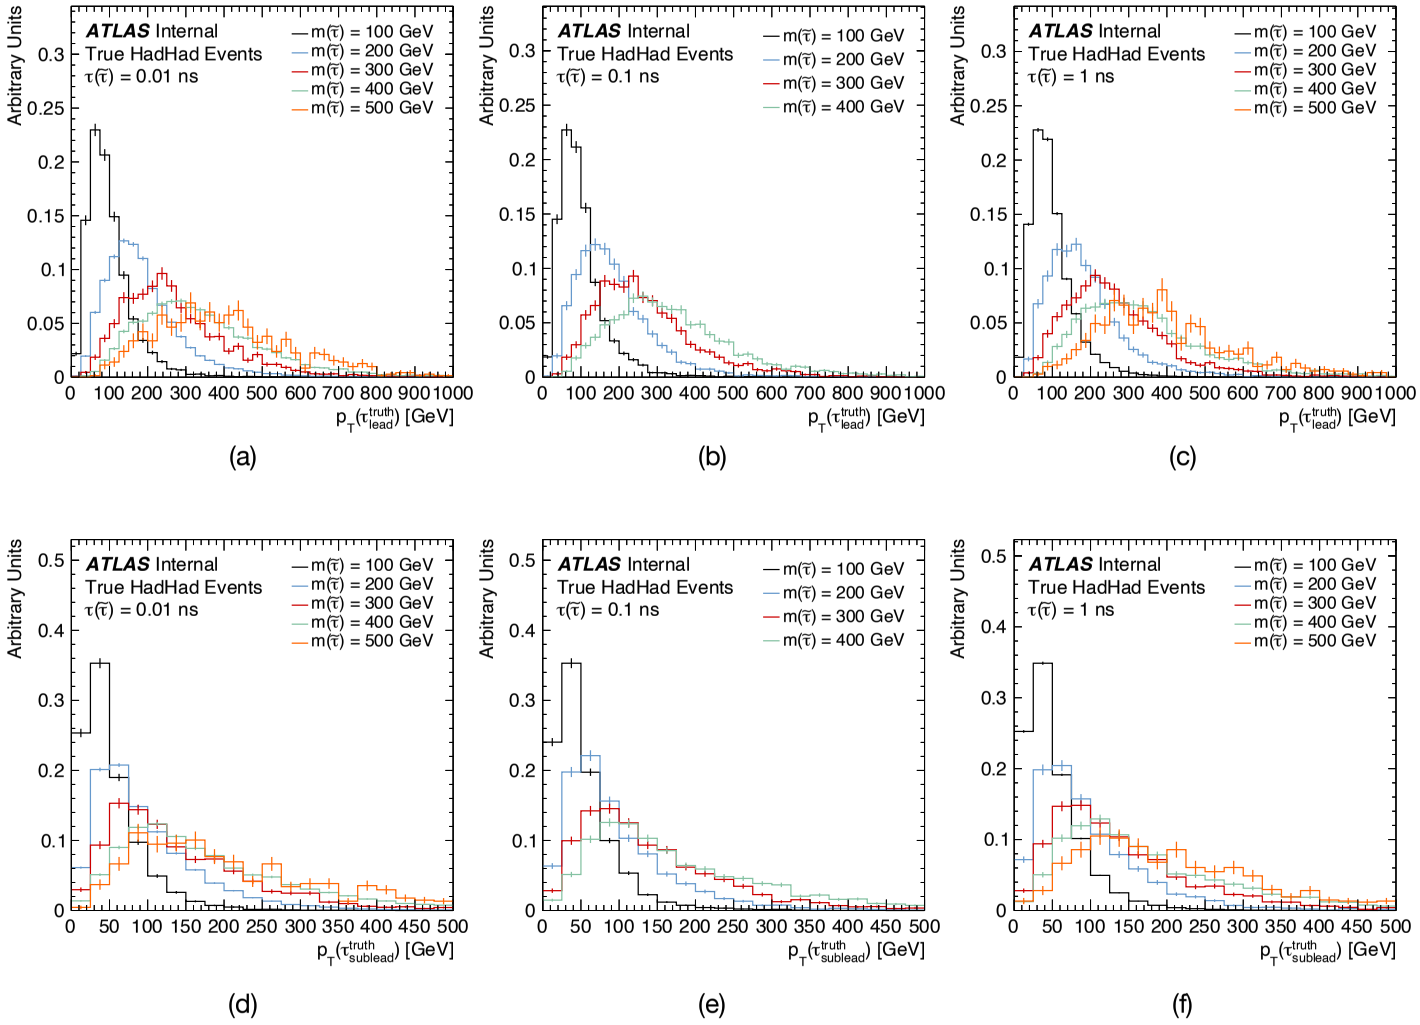

Figure 10.3 Truth p_{T} distributions for the MC $\tilde{\tau}\tilde{\tau} \to \tau\tau$ process. Distributions are computed for the (*top*) leading-p_{T} and (*bottom*) sub-leading-p_{T} τ-leptons for various mass and lifetime points.

Filter Name	Trigger Chains (OR Selection)	Description
DTau_SiTau	HLT_tau160_mediumRNN	Single-τ
	HLT_tau200_mediumRNN_tracktwoMVA	
DTau_DiTau	HLT_tau80_mediumRNN_tau60_mediumRNN	Di-τ
DTau_TauMET	HLT_tau35_mediumRNN_xe70	$\tau + E_{\mathrm{T}}^{\mathrm{miss}}$
	HLT_tau60_mediumRNN_tau25_mediumRNN_xe50	

Table 10.6 trigger filters as used in the selection of τ_{had} objects for each analysis channel. Suitable τ_{had}-leptons may then feature in events selected for subsequent analysis.

with a MET requirement of $E_{\mathrm{T}}^{\mathrm{miss}} \geqslant 50$ GeV. In all such chains, the mediumRNN denotes the application of the RNN-based prompt τ-ID classifier at the medium working point.

The triggers introduced in this section are organised into *filters*, wherein a logical combination (logical OR) is taken of a sub-set of τ-HLT chains such that they may be applied collectively as a single selection. Table 10.6 introduces the DTau_SiTau, DTau_DiTau, and DTau_TauMET filters, respectively defined as the logical-OR of each trigger belonging to the appropriate category. Such filters may then be used to apply the required trigger selections for each analysis channel. The filters defined in Table 10.6 are used to define objects suitable for reconstruction as a hadronically-decaying τ_{had}-lepton. Once reconstructed, τ_{had} objects may then be selected for analysis if deemed to be of sufficient quality.

Table 10.7 presents the signal acceptance attained by the afore-described trigger filters with respect to τ_{had} objects. The percentage of τ_{had}-leptons accepted by the trigger filters is shown for various mass and lifetime values of the parent $\tilde{\tau}$ in the $\tilde{\tau}\tilde{\tau}$ MC sample, as accepted by the logical-OR of the DTau_SiTau and DTau_DiTau filters. The MET-based filter is not employed in the definition of the $\tau_{\mathrm{had}}\tau_{\mathrm{had}}$ channel and is therefore neglected for this study. It is found that trigger acceptance is greater for LLPs of higher mass points, with the highest lifetime points for each mass value typically achieving the greatest acceptance. This trigger strategy therefore provides suitable performance with which to conduct a search for displaced τ-leptons.

$\tilde{\tau}$ Mass [GeV]	$\tilde{\tau}$ Lifetime [ns]	DSID Acceptance [%]
100	0.01	12.7
	0.1	12.2
	1.0	13.1
200	0.01	38.5
	0.1	37.5
	1.0	38.5
300	0.01	55.3
	0.1	56.8
	1.0	57.4
400	0.01	69.9
	0.1	68.7
	1.0	69.0
500	0.01	74.5
	1.0	74.6

Table 10.7 MC signal acceptance for various $\tilde{\tau}$ mass and lifetime points in the $\tau_{had}\tau_{had}$ channel. Acceptance is reported as the percentage of the number of events in the DSID which corresponds to the $\tilde{\tau}\tilde{\tau}$ MC sample with the given $(m, c\tau)$ instantiation.

10.4 Object Definitions

A subset of the 'defined objects' introduced in Chapter 8 are selected from the data inputs for analysis in the search for displaced τ-leptons. Whether measured from LHC pp-interactions or generated through MC simulation, all physics data are processed by the ATLAS reconstruction chain such that defined objects may be made available for subsequent analysis. This section defines the experimentally reconstructed objects which qualify for analysis in this search for displaced-τ production.

Hadronic τ-lepton reconstruction is seeded by calorimeter topological clusters, as introduced in Section 8.3.6. Such clusters are calibrated with Local Cell Weighting (LCW), before processing with the $\texttt{anti-}k_{\mathsf{T}}$ algorithm. τ_{had} jet seeds are processed with a distance parameter of $R = 0.4$, after which they must survive the quality selection criteria of $p_{\mathsf{T}} > 10\,\mathsf{GeV}$ and $|\eta| < 2.5$ to be considered for selection as a τ_{had} candidate.

The Tau Vertex Association algorithm (TVA) [144] is used to determine the vertex from which the τ-lepton decays, which may not correspond to the primary vertex of the event on account of the lifetime of the SM τ. High-p_{T} τ-leptons, in particular, may traverse some distance from their production vertex before decaying into reconstructable products. The vertex correction performed by the TVA therefore improves τ_{had} reconstruction efficiency while simultaneously improving resilience to the effects of pile-up. The vertex selected by the TVA is that which exhibits the highest p_{T}-weighted fraction of all tracks associated therewith, which satisfy $p_{\mathsf{T}}\,(\mathsf{track}) > 0.5\,\mathsf{GeV}$ within $\Delta R = 0.2$ of the jet axis. Consideration is not extended to LRT track content by the TVA.

Tracks are then associated with the τ_{had} candidate should they fall within $\Delta R = 0.4$ of the jet seed axis. Large radius tracks, as introduced in 8.2.1, are considered for τ_{had} association. It is at this point that the services of the displaced track classifier, introduced in the previous chapter, are employed. The displaced τ-ID RNN is then applied to the classified tracks, profiting from the improved classification of τ_{had} decays which are characterised by a non-zero charged track multiplicity. A score-flattening procedure is applied to the output of the displaced τ-ID RNN such that the ID response is uniform in p_{T} and $\langle\mu\rangle$.

τ_{had} candidates are required to pass the displaced τ-ID under the application of

Quality Cut	Description
Acceptance	$p_T > 20$ GeV && $\|\eta\| < 1.32 \,\|\|\, (1.52 < \|\eta\| < 2.5)$
Displaced-ID Score	pass `VeryLoose` working point selection
Track Multiplicity	> 0 *w.r.t.* displaced track classifier

Table 10.8 baseline τ_{had} quality requirements for analysis consideration in the $\tau_{had}\tau_{had}$ channel.

the `VeryLoose` RNN working point, as defined in Table 9.1. Further quality cuts of $p_T(\tau_{had}) > 20$ GeV and $|\eta| < 2.5$ must be met if the τ_{had} candidate is to be considered for analysis. Should at least one charged track be associated with the lepton, as determined by the displaced track classifier, the candidate is determined to be a *baseline* τ_{had}-lepton — that which satisfies the minimum quality criteria, but is yet to be selected for analysis. The requirements for baseline τ_{had} selection are summarised in Table 10.8.

The efficiency with which offline baseline τ_{had} objects are accepted by the displaced τ-ID RNN is presented 10.4. A markedly more efficient acceptance of truth-matched MC τ_{had}-leptons is observed for the displaced-ID (solid markers) than that attained by the nominal-ID (hollow markers), where both acceptance efficiencies are derived for the `medium` working point as defined for each RNN. This indicates that, given a suitable event selection strategy, a performant signal acceptance of displaced τ-leptons is within reach.

Equipped with a baseline selection for hadronically-decaying τ-leptons, events which harbour two τ_{had}-leptons may be selected for the $\tau_{had}\tau_{had}$ analysis channel.

10.5 $\tau_{had}\tau_{had}$ Event Selection

The experimental objects defined in the previous section, having passed initial trigger filtration and baseline quality selection, may be used to define the selection process with which to accept particle interaction measurements for statistical analysis.

Particle decay events are selected for the $\tau_{had}\tau_{had}$ channel if they pass the trigger selection imposed by either the single-τ chain `HLT_tau160_mediumRNN`, or the di-τ chain `HLT_tau80_mediumRNN_tau60_mediumRNN`. Events are further required to

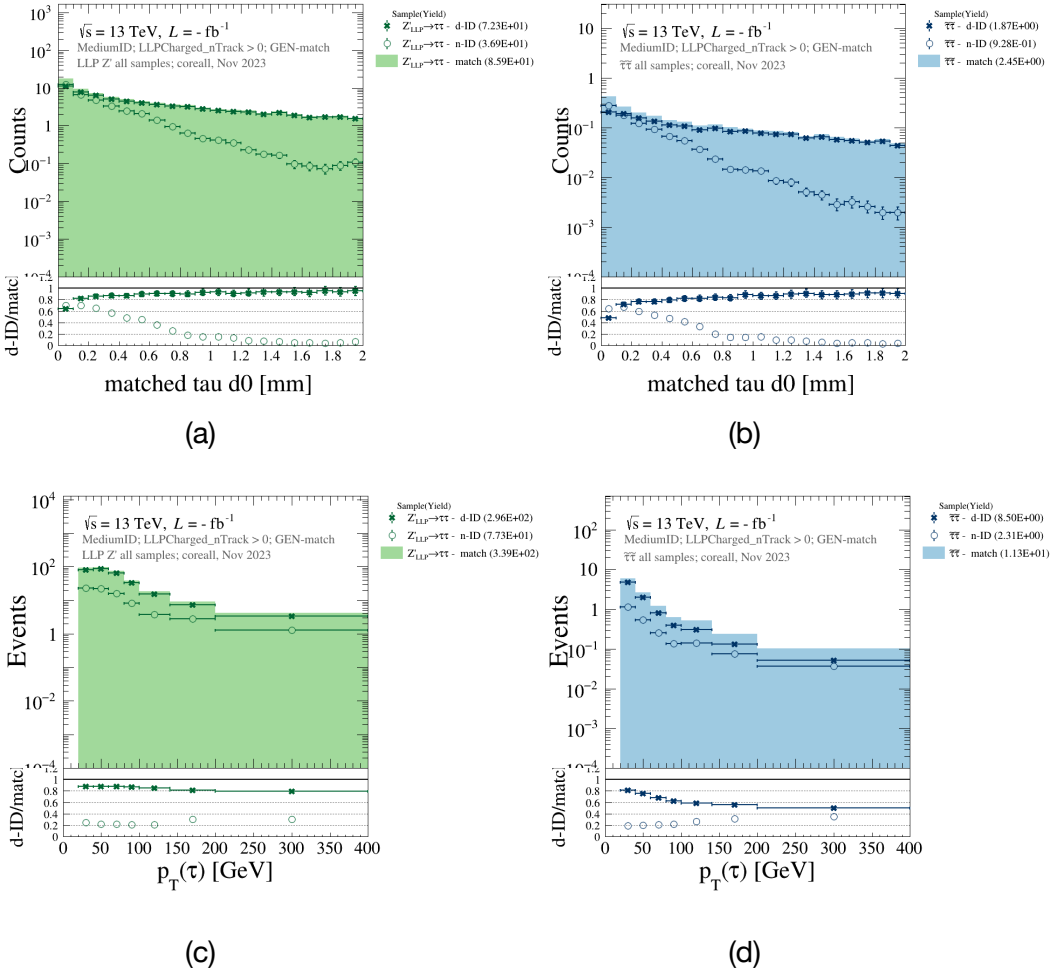

Figure 10.4 displaced τ-ID signal acceptance efficiency with respect to baseline τ_{had} objects. The acceptance efficiency is shown for truth-matched (*left*) $Z' \to \tau\tau$ and (*right*) $t\bar{t} \to \tau\tau$ MC samples, parameterised as a function of (*top*) the impact parameter, d_0, of the truth-matched τ_{had}, and (*bottom*) the p_T of the truth-matched τ_{had}. The acceptance efficiency attained by the displaced-ID (*solid markers*) is compared to that attained by the nominal-ID (*hollow markers*). Markedly more efficient acceptance of baseline τ_{had}-leptons is observed for the displaced-ID RNN.

Selection Type	Description
Filters	`DTau_SiTau` ||
	`DTau_DiTau`
Triggers	`HLT_tau160_mediumRNN` ||
	`HLT_tau80_mediumRNN_tau60_mediumRNN`
τ_{had} Multiplicity	$\tau_{had} \geqslant 2$ baseline
e Multiplicity	0
μ Multiplicity	0
Kinematics	$p_T(\tau_{lead}) > 165$ GeV ||
	$p_T(\tau_{lead}) > 85$ GeV && $p_T(\tau_{sub\text{-}lead}) > 65$ GeV

Table 10.9 summary of event selection criteria for the $\tau_{had}\tau_{had}$ channel, including trigger-level and kinematic-level requirements.

contain a minimum of 2 τ-leptons which pass the τ_{had} baseline selection introduced in the previous section.

Subsequent to such requirements, a leptonic veto is enforced — this ensures that no reconstructed electrons or muons are present in the events selected for the $\tau_{had}\tau_{had}$ channel, thereby ensuring an appropriate level of data hygiene with respect to the to the two $\tau_{lep}\tau_{had}$ channels. Electrons and muons are reconstructed in accordance with the procedures outlined in Section 8.3.

A final kinematic selection requires that either the leading-p_T τ-lepton has a transverse momentum in excess of 165 GeV, or that collectively the leading-p_T and sub-leading-p_T τ-leptons respectively possess a transverse momentum in excess of 85 GeV and 65 GeV. Collectively, such selections define the events which are taken to comprise the $\tau_{had}\tau_{had}$ analysis channel, and are summarised in Table 10.9.

The pre-selections (that is, the quality selections made in advance of analysis-level cuts) defined in this section achieve a material suppression of the background process population. This may be observed in the *cut-flow* results presented in Table 10.10. Such numbers quantify the event yield attributable to each tabulated MC process, as measured at each successive stage of selection defined thus far. It is clear that signal models with a higher value of $c\tau$ are favoured by the selections which define the $\tau_{had}\tau_{had}$ channel — as is appropriate for the target of this search.

The numbers derived from QCD di-jet MC samples — the primary background source in this search — cannot be relied upon in this phase-space, and do not feature in this cut-flow tabulation. For this reason, the statistical background associated with this process is evaluated in a data-driven manner — as will be discussed in the subsequent chapter. Event yield from the $t\bar{t} \to \tau\tau$ process was of negligible significance and accordingly excluded from the table. The effect of the displaced-ID application for various working points will be examined at the analysis level.

Having established a series of object and event selections with which to define the analysis channel of interest, this channel may now serve as the basis of a statistical analysis. To conduct such an analysis, a statistical strategy must be specified.

10.6 Analysis Strategy

The object and event selections defined in this chapter provide the basis on which to conduct a search for displaced-τ production in the $\tau_{\text{had}}\tau_{\text{had}}$ analysis channel.

While they provide the required data-set of the quality necessary to conduct a search, primitive selections of this nature do little to discriminate between the physics processes from which such events originate. The number of events which arise from background processes, rather than from the LLP states of interest, must therefore be estimated.

This is achieved by the definition of suitable data *regions*. Such regions delineate the observed event yield as a function of variables which serve to discriminate between processes of origin. A *signal region* is defined such that the background contamination in this portion of phase-space is minimised. The proportion of the event yield in this region which may be attributed to background processes is estimated with the aide of an additional region, enriched with background events by design. The manner in which this is achieved in this thesis is the subject of the following chapter.

Equipped with an estimate of the expected background yield in the signal region, the observed yield in this region constrains the values of signal production cross-section which are compatible with observation. This is achieved with a statistical limit-setting procedure, as will be discussed in Chapter 12. Limits are derived

Selection Cut	Z' (500 GeV, 1.0 ns)	$\tau\tau$ (200 GeV, 0.1 ns)	$Z \to \tau\tau$	Di-Boson
No cut	$432,983 \pm 2,078$	$34,272 \pm 270$	$83,347,882 \pm 98,017$	$3,723,041 \pm 10,595$
`DTau_SiTau`	$88,233 \pm 939$	$1,495 \pm 58$	$52,684 \pm 928$	$41,593 \pm 579$
`||DTau_DiTa` trigger filter				
τ_{had} multiplicity $\geqslant 2$	$54,162 \pm 735$	848 ± 44	$26,927 \pm 653$	$10,394 \pm 261$
Track multiplicity $\geqslant 1$	$28,041 \pm 530$	473 ± 32	$14,540 \pm 565$	$2,666 \pm 131$
Kinematic selection	$27,115 \pm 522$	445 ± 31	$12,119 \pm 550$	$1,456 \pm 96$

Table 10.10 event yield observed after each successive cut in the $\tau_{had}\tau_{had}$-channel selection criteria. The 'cut-flow' results highlight the background suppression achieved at the level of selection, while signal acceptance is illustrated with a given (mass, lifetime)-point of each signal MC sample.

in a model-independent manner in the $\tau_{had}\tau_{had}$ channel, together with a specific model interpretation in the form of cross-section limits on $\tilde{\tau}\tilde{\tau}$ production. Model-specific interpretations allow the sensitivity of the search to be gauged, and may be viewed as a performance benchmark of the model-independent search.

Limits of this kind represent the scientific output of this thesis. Models which predict the existence of long-lived particles beyond the Standard Model must do so in a manner which is compatible with the observed limits obtained in the research presented herein. It is with such limits that this thesis constrains the phase-space in which the secrets of nature may hide.

This chapter has introduced the overall strategy deployed in this thesis in the search for displaced-τ production, and the selections with which the $\tau_{had}\tau_{had}$ analysis channel is defined. If the production of LLP states beyond the SM is to be probed, a suitable signal region accompanied by an estimate of the background content it contains must be obtained. This work is presented overleaf.

“*Science is perhaps the only human activity in which errors are systematically criticised and, in time, corrected.*”

— Karl Popper

11

Region Design and Background Estimation

While the development of bespoke trigger, reconstruction, and identification technologies enables the effective extraction of displaced τ signatures, such systems do not reject background processes with 100% efficiency. Moreover, the object selection criteria imposed at the offline analysis-level cannot filter residual background events in their entirety — an irrepressible population of background events will therefore contaminate the selected data. Consequentially, the number of events produced via background process which are expected to satisfy signal selection criteria must be estimated such that their effects may be understood. This chapter presents the statistical methodologies and analytical design decisions which facilitate the estimation of background effects in the displaced $\tau_{\text{had}}\tau_{\text{had}}$ search channel.

11.1 Overview

The irrepressible nature of the residual background event population, which evade the rejective power of identification algorithms and signal selection criteria, necessitates the estimation of their contribution to the measured event yield in the search for displaced τ-leptons. In the $\tau_{\text{had}}\tau_{\text{had}}$ decay channel, the overwhelming majority of such events arise due to QCD di-jet production. Di-jet signatures which mimic the detector signature of a hadronic-τ decay can be erroneously reconstructed as a τ-lepton — such objects are termed 'fake' τs.

At the time of writing, QCD processes are not sufficiently well-modelled as implemented in MC event generators — simulated data samples do not, therefore, provide an accurate representation of the number of di-jet events which are likely to satisfy the signal selection criteria of the analysis. Moreover, the reliability of MC di-jet data is further diminished when extrapolated to high values of displacement. It is therefore necessary to obtain a *data-driven* estimation of the fake τ content in the $\tau_{\text{had}}\tau_{\text{had}}$ decay channel.

To enable the estimation of its fake τ content, the selected data-set is divided into various 'regions'. A *signal region* (SR) is defined in terms of ID requirements which maximise sensitivity to signal processes, and secure a data sample depleted of background entries to the greatest degree possible. Supporting *side-band regions* are constructed, and are shown to aide the determination of background yields which correspond to a single fake τ or two fake τs in the $\tau_{\text{had}}\tau_{\text{had}}$ SR.

The fake τ content in the SR is then determined with the data-driven *hadronic fake-factor method* [159]. A quantification of the expected fake content in the $\tau_{\text{had}}\tau_{\text{had}}$ channel is thereby obtained without reference to MC-simulated di-jet content.

Sub-dominant sources of background events which arise from alternative SM processes do not suffer from the modelling limitations experienced by QCD di-jet events at the simulation level. It is therefore possible to observed the number of events produced by such processes — as introduced in the previous chapter — which fall within the signal region, and proceed accordingly. This is the manner in which the background processes other than QCD di-jet are estimated in this analysis.

11.2 Hadronic Fake-Factor Method

The hadronic fake-factor method is centred upon the derivation of a representative ratio for the number of fake τ-jets which satisfy signal selection criteria to those which fail. This 'fake-factor' may then serve as a multiplicative corrective factor with which to re-weight the event yield observed in the signal region.

The fake-factors are derived in a dedicated region, termed the *control region*. The selection criteria of the control region are defined with a view to attaining a background-enriched and signal-depleted data sample, such that the fake factor extracted from the control region provides a representative quantification exclusively for the *fake* τ content which satisfy signal selection.

With the definition of a suitable control region, the ratio of the fake factor may be computed. In order to obtain an accurate estimate of the QCD-initiated fake τ population, contributions from alternative background processes are removed through MC simulation. In contrast with dijet events, the contribution of sub-dominant SM background processes may accurately be determined through examination of simulated sample yields in the signal region, and are accordingly excluded from the fake jet estimation. For a given ID working-point, the fake factor, F, may therefore be defined as:

$$F = \frac{N^{\mathsf{CR}}_{\tau-\mathsf{ID}}(\mathsf{data}) - N^{\mathsf{CR}}_{\tau-\mathsf{ID}}(\mathsf{MC\ jet} \neq \mathsf{QCD})}{N^{\mathsf{CR}}_{\mathsf{anti}\ \tau-\mathsf{ID}}(\mathsf{data}) - N^{\mathsf{CR}}_{\mathsf{anti}\ \tau-\mathsf{ID}}(\mathsf{MC\ jet} \neq \mathsf{QCD})}, \tag{11.1}$$

where the number of fake τ-leptons attributable to non-QCD background processes, $N^{\mathsf{CR}}_{\tau-\mathsf{ID}}(\mathsf{MC\ jet} \neq \mathsf{QCD})$ and $N^{\mathsf{CR}}_{\mathsf{anti}\ \tau-\mathsf{ID}}(\mathsf{MC\ jet} \neq \mathsf{QCD})$, are respectively subtracted from the total number of τ-leptons which pass, $N^{\mathsf{CR}}_{\tau-\mathsf{ID}}(\mathsf{data})$, and fail, $N^{\mathsf{CR}}_{\mathsf{anti}\ \tau-\mathsf{ID}}(\mathsf{data})$, τ-ID as applied within the control region.

The numerator of the fake-factor therefore quantifies the number of fake τs attributable to mis-identified QCD-jets which *pass* τ-ID, while the denominator quantifies the QCD fake population which *fail* τ-ID. The multiplication of this factor with the number of τ-leptons which fall within a given data-region therefore scales the corresponding event-yield such that the contribution which likely originates from QCD-jets may be obtained.

It is of import to recognise that the fake factor ratio is determined *per* τ, not per event — this must be accounted for in the application of the fake factor such that

processes which produce differing fake τ multiplicities are treated appropriately.

11.2.1 Region Design

It is necessary to delineate the data obtained from the displaced-τ reconstruction and selection process into various regions, such that the hadronic fake factor method may be tailored to the requirements posed by a search for displaced τ-lepton production. This is required such that a dedicated, fake-enriched control region may be defined for the purpose of fake factor extraction, while additionally providing a distinct data-sample with which to estimate the number of background events which fake a single or double hadronic τ decay signature.

The construction of data regions which correspond to the $\tau_{\text{had}}\tau_{\text{had}}$ decay channel may broadly be decomposed into a *signal region* (SR) and supporting *side-band* regions, in addition to a FF control region. Such regions are constructed in terms of the score produced by the RNN developed for displaced-τ identification, as was introduced in Section 9.2.2. A data plane is defined as a function of the binary ID pass/fail score obtained with the `Medium` working point, where the score ascribed to the leading p_{T} τ (τ_{lead}) and sub-leading p_{T} τ ($\tau_{\text{sub-lead}}$) of each event respectively form the axes of the plane. All τ-leptons entering the plane are required to pass the `VeryLoose` displaced-ID working point such that a floor is imposed on τ reconstruction quality. The following regions of data are defined on this plane:

- **region *A* (SR):** τ_{lead} *pass* `medium`-ID and $\tau_{\text{sub-lead}}$ *pass* `medium`-ID;
- **region *B*:** τ_{lead} *pass* `medium`-ID and $\tau_{\text{sub-lead}}$ *fail* `medium`-ID;
- **region *C*:** τ_{lead} *fail* `medium`-ID and $\tau_{\text{sub-lead}}$ *pass* `medium`-ID;
- **region *D*:** τ_{lead} *fail* `medium`-ID and $\tau_{\text{sub-lead}}$ *fail* `medium`-ID.

Regions B and C therefore provide representative data for the estimation of individual fake τ_{lead} and fake $\tau_{\text{sub-lead}}$ contributions to the background event population. The matrix comprised by regions A, B, C, and D is termed the $ABCD$-plane, and is visualised in Figure 11.1.

Additionally, a dedicated control region is defined for the extraction of fake-factors. The control region imposes the selection criteria requiring the leading

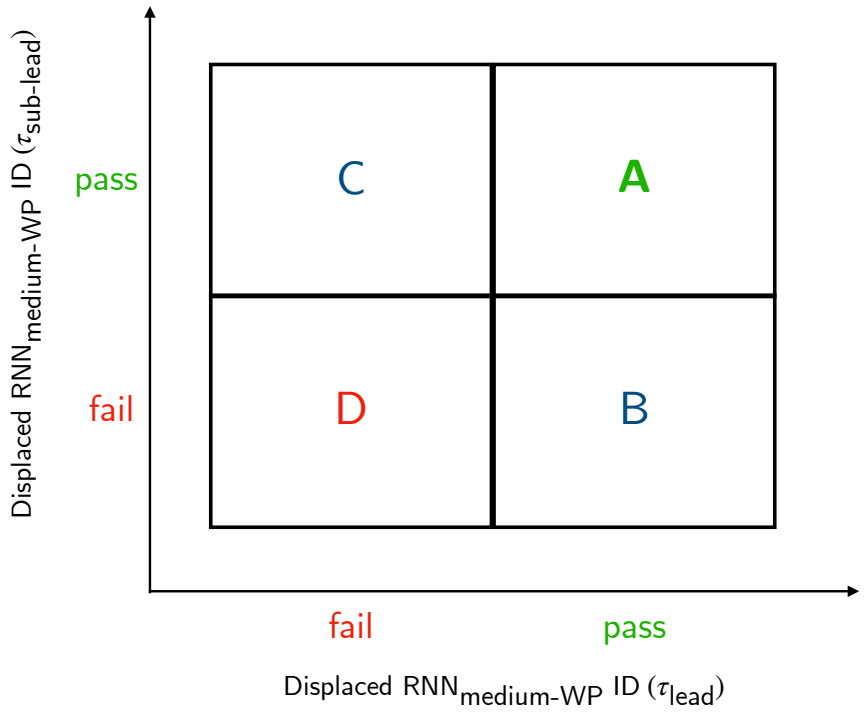

Figure 11.1 plane constructed from the leading p_{T} and sub-leading p_{T} τ ID score, as obtained with the `Medium` working point of the RNN classifier optimised for displaced-τ signatures. Regions with which to estimate fake τ_{lead} and $\tau_{\text{sub-lead}}$ content are defined, in addition to a signal region (A).

τ of each event to pass the `VeryLoose` ID requirement, while the sub-leading τ is required to pass the more stringent `Loose` requirement. The fake-factor ratio is then determined with the sub-leading τ such that sufficient statistics for low-momentum leptons is are available to inform the value of the ratio.

The hadronic fake-factor may then be applied to regions B, C, and D, such that an estimate of the number of events containing a fake leading τ, a fake sub-leading τ, or two fake τs may be ascertained. This estimated fake background may then be accounted for in the final event yield observed in the signal region, A.

The MC signal event yield in each region of the $ABCD$-plane is presented in Figure 11.2. Figure 11.2 (a) and (b) show the composition of signal acceptance of each region across the plane for the Z' and *tautau* processes, where it may be observed that the vast majority of signal events which enter the $ABCD$-plane are selected for the signal region, region A. It may further be noted that regions which are intended by design be depleted of signal event harbour a negligible yield of such events. Figure 11.2 (c) shows the absolute event count for the $\tau\tau$*process* exclusively within signal region A. All results are quantified for a variety of mass

and lifetime points — it may be observed that strong signal acceptance is attained in the signal region across all phase-space considered.

Before the hadronic fake-factor method may be applied to the afore-defined regions, the performance of the background estimation it provides must be validated. This ensures that the estimation method is effective within tolerable error bounds *before* the signal region is exposed, or 'unblinded', to the analysis. The method by which the $\tau_{\text{had}}\tau_{\text{had}}$ background estimation procedure is validated is now introduced.

11.2.2 Validation Strategy

The background estimation procedure is validated on regions designed to verify the effectiveness of the procedure without the need to 'unblind' the main signal region, as the content of which could bias the design of the analysis. It is critical that the content of the signal region does not influence the design decisions pertinent to the analysis in a manner which may inflate the significance of the findings — specific validation regions are therefore defined such that the performance of the background estimation procedure may be ascertained without exposure to the primary signal region which is sensitive to new physics.

For the purposes of performing this validation, dedicated regions are obtained by sliding the pass/fail matrix formed by the *ABCD* regions down the plane of leading and sub-leading τ ID quality criteria, as is shown in Figure 11.3. This results in the definition of two new regions:

- **region *E*:** (τ_{lead} *pass* `Loose`-ID & *fail* `Medium`-ID) and ($\tau_{\text{sub-lead}}$ *pass* `VeryLoose`-ID & *fail* `Loose`-ID);
- **region *F*:** τ_{lead} *pass* `Medium`-ID and ($\tau_{\text{sub-lead}}$ *pass* `VeryLoose`-ID & *fail* `Loose`-ID).

As a result of sliding one quality working point down the ID plane, it should be noted that a new minimum quality cut has been imposed on the additional regions: requiring τ-leptons to pass the displaced ID with the `VeryLoose` working point, rather than the `Loose` working point required for the minimum quality of τ in the *ABCD* matrix.

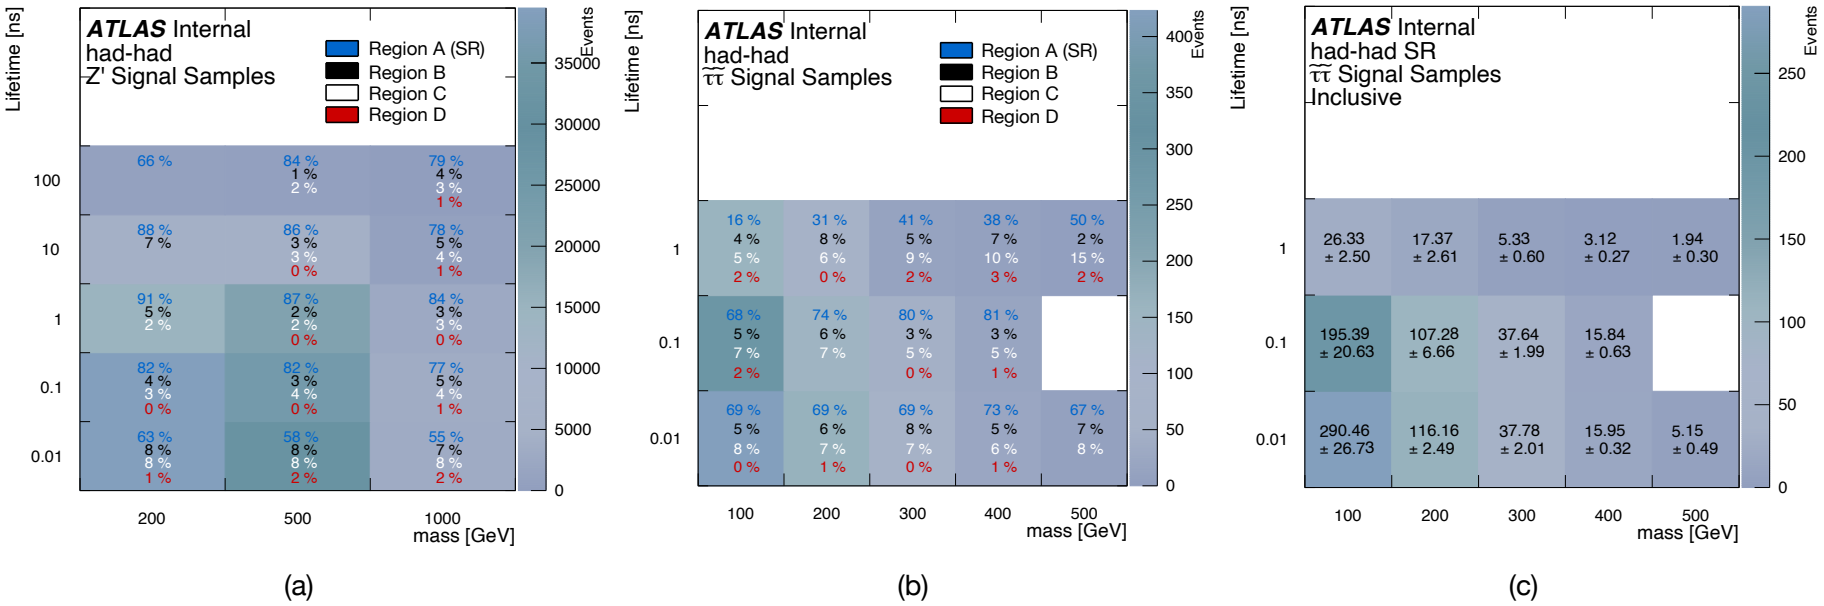

Figure 11.2 MC signal yield in the $\tau_{had}\tau_{had}$ *ABCD*-plane for (*a*) $Z' \to \tau\tau$ and (*b*) $\widetilde{\tau\tau} \to \tau\tau$ processes, quantified for different mass and lifetime points. Figure (*c*) restricts consideration to the signal region, presenting the absolute $\widetilde{\tau\tau}$ event yield in region *A*. It is observed that the majority of true signal events in the *ABCD*-plane fall into the signal region, with negligible signal contamination into regions intended to be signal depleted by design.

The validation procedure then unfolds in the same manner as the primary estimation procedure, applying the appropriate fake-factors to the leading and sub-leading τs found in regions *D*, *E*, and *F* such that an estimate of the number of events containing a fake leading τ, a fake sub-leading τ, or two fake τs may be determined. This result may then be compared to the data content in region *B*, having first accounted for the background content expected to arise from other SM processes with the aide of MC simulation. As this region is, by design, background enriched, the content of this region should be dominated by background processes — a close agreement between the expected fake count in this region and the data yield in this region, subtracting SM MC contributions, therefore indicates strong background estimation performance.

As is illustrated in Figure 11.3, the decreased quality cut required by the validation *BDEF* matrix with respect to the primary *ABCD* matrix requires the derivation of an additional fake-factor, such that the transfer from `VeryLoose` regions to `Loose` regions may be determined. An additional control region must therefore be defined, such that the $\tau_{\text{sub-lead}}$ fake content may be correctly estimated for the validation SR, *B*, given that the quality requirements for the leading and sub-leading τs are mis-aligned in the validation *BDEF* matrix.

The additional validation control region is obtained analogously to the construction of regions *E* and *F*, decreasing the quality cut on the sub-leading τ and thereby sliding down the ID quality plane by one working point bin. An understanding of the quality thresholds applied to each τ-lepton in the plane is best gained with reference to Figure 11.3. The fake-factor extracted from this region may then be used to determine the fake τ content which corresponds to the sub-leading τ in regions *E* and *F*.

11.3 Validation of Background Estimation

The validation of the hadronic fake-factor method was undertaken in accordance with the afore-outlined procedure. This initiates with the extraction of the appropriate fake-factor ratios for the validation *BDEF* region matrix. As was highlighted in the previous section, this process differs from that which corresponds to the main SR in its need for *two* fake-factors — due to the differing quality requirements on the leading and sub-leading τs in the validation plane.

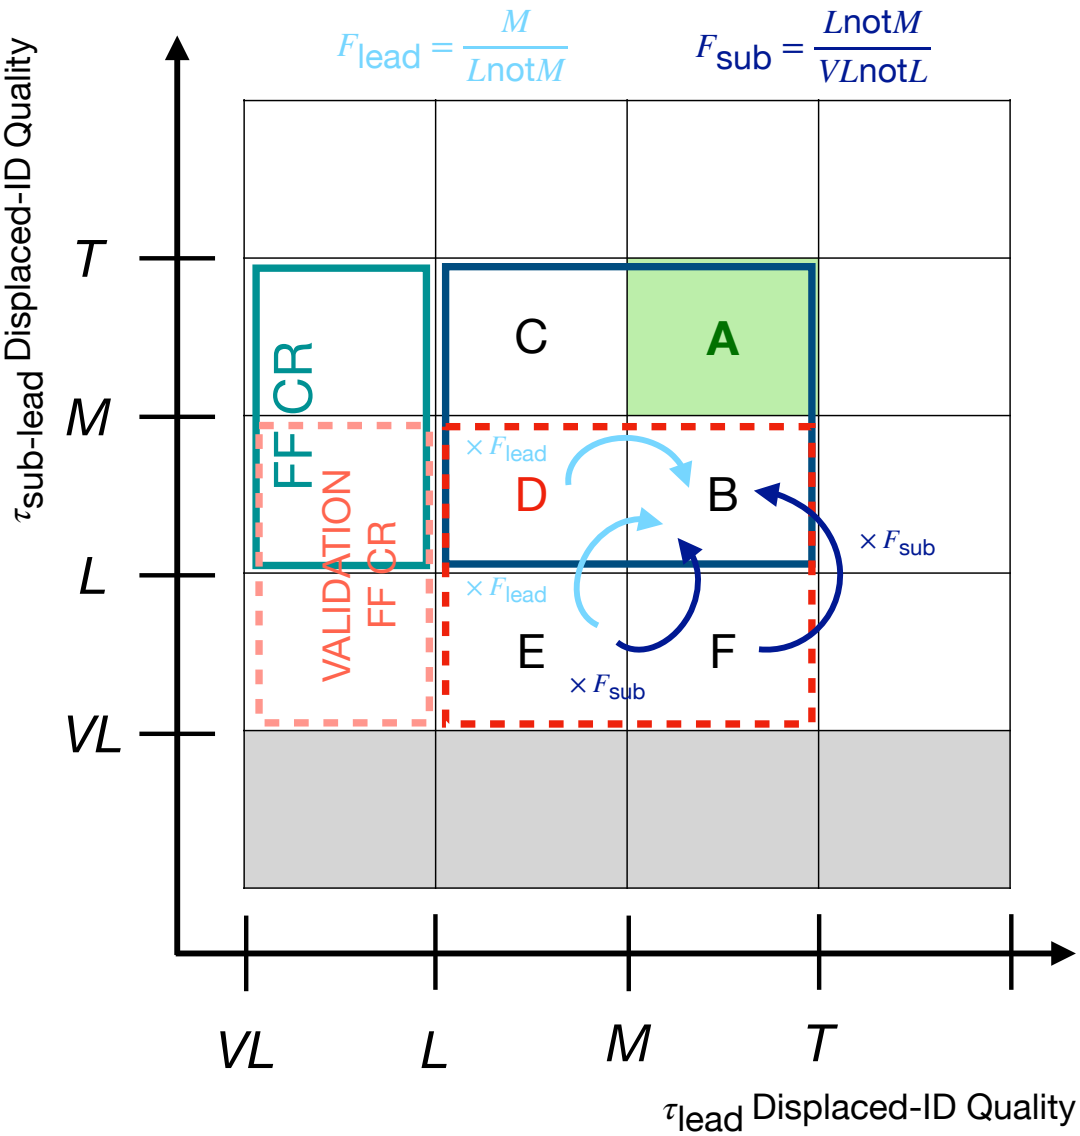

Figure 11.3 application of fake-factors to the validation SR (B), where the validation SR and supporting side-band regions are obtained by decrementing the working point of the ID used for the selection criteria of each of the original regions. This provides a distinct τ population with which to validate the effectiveness of the fake-factor method without unblinding regions sensitive to new physics. Axis bin labels abbreviate the the working point selections imposed by the ID classifier on the leading and sub-leading τs.

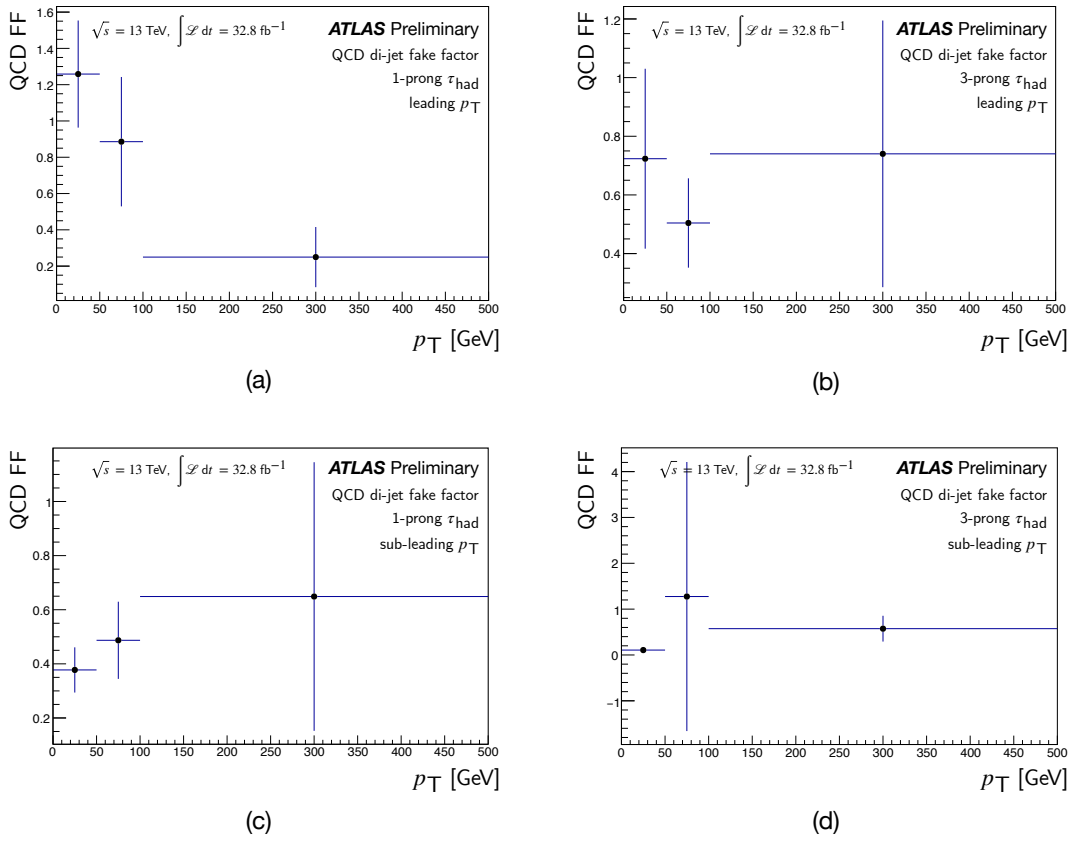

Figure 11.4 validation fake-factor values for *(a,b)* leading p_T τ and *(c,d)* sub-leading p_T τ, as determined for 1-prong and 3-prong reconstructed τ_{had}-leptons. As a consequence of the sliding validation scheme, the leading and sub-leading τs require individual fake-factors to account for their differing ID quality requirements in the validation plane. Fake-factor values are binned in the p_T of the subject τ such that variations in background population as a function of τ kinematics are accommodated by the estimate.

11.3.1 Validation Region Fake-Factors

The fake-factor definition provided in Equation 11.1 may now be expressed in terms of the validation pass/fail criteria imposed on each leading and sub-leading τ. This yields two fake-factor expressions:

$$F_{\text{lead}} = \frac{N^{\texttt{Medium}}_{\text{data}} - \sum_{\text{MC}\neq\text{dijet}} N^{\texttt{Medium}}_{\text{MC}}}{N^{\texttt{L!M}}_{\text{data}} - \sum_{\text{MC}\neq\text{dijet}} N^{\texttt{L!M}}_{\text{data}}} \qquad F_{\text{sub-lead}} = \frac{N^{\texttt{L!M}}_{\text{data}} - \sum_{\text{MC}\neq\text{dijet}} N^{\texttt{L!M}}_{\text{MC}}}{N^{\texttt{VL!L}}_{\text{data}} - \sum_{\text{MC}\neq\text{dijet}} N^{\texttt{VL!L}}_{\text{data}}}, \quad (11.2)$$

where F_{lead} and $F_{\text{sub-lead}}$ are respectively extracted from the primary and validation

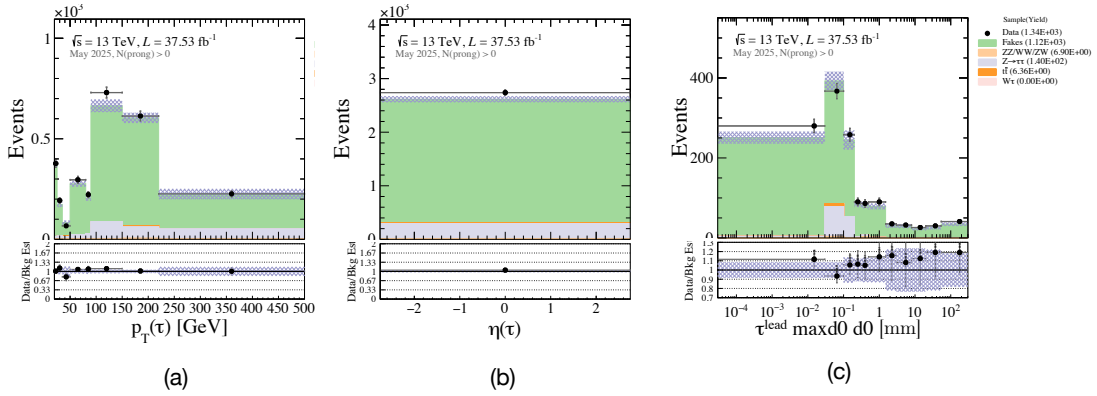

Figure 11.5 kinematic distributions for events observed in the validation SR, region B. Distributions are shown for the (*a*) p_T and (*b*) η values of all τ-leptons in the region, in addition to (*c*) largest impact parameter of the tracks associated with the leading-p_T τ in the region. Events observed in real data (*black*) are compared to the fake estimate obtained from the fake factor method (*green*), together with non-QCD SM MC samples.

CRs and applied to the τ of the corresponding p_T-rank. Each fake-factor is binned in values of τ p_T in order to capture variations in background composition as a function of τ kinematics. Additional binning strategies which leverage alternative kinematic variables, such as d_0, η, and φ were investigated and found to attain less accurate background estimation results than are achieved with p_T-binned fake-factors. Further, fake-factors are derived selectively as a function of τ track multiplicity — thereby accounting for variations between fake 1-prong and fake 3-prong τs. A variable bin-width is designed to reduce the statistical uncertainty in the higher-p_T bin where statistics would otherwise be limited. The fake-factors obtained for validation purposes are shown in Figure 11.4, together with their associated statistical uncertainties.

11.3.2 Background Estimation in Validation Region

The fake-factors derived from the primary and validation control regions may now be applied. The total number of fake events in our validation SR, B, is given by the sum of the number of events which produce a single fake τ (either leading or sub-leading in p_T) with the number of events which produce two fake τs:

$$\begin{aligned} N^{\mathrm{B}}_{\mathrm{fake}}(\mathrm{events}) &= N^{\mathrm{B}}_{\mathrm{fake}}(1\tau) + N^{\mathrm{B}}_{\mathrm{fake}}(2\tau) \\ &= N^{\mathrm{B}}_{\mathrm{fake}}(\tau_{\mathrm{lead}}) + N^{\mathrm{B}}_{\mathrm{fake}}(\tau_{\text{sub-lead}}) + N^{\mathrm{B}}_{\mathrm{fake}}(\tau_{\mathrm{lead}}, \tau_{\text{sub-lead}}). \end{aligned} \tag{11.3}$$

This may be expressed in terms of the *BDEF* validation plane definitions as:

$$\begin{aligned} N^{\mathrm{B}}_{\mathrm{fakes}} = F_{\mathrm{lead}} \times \left[N^{\mathrm{D}}_{\mathrm{data}} - \sum_{\mathrm{MC}\neq\mathrm{dijet}} N^{\mathrm{D}}_{\mathrm{MC}} \right] + F_{\text{sub-lead}} \times \left[N^{\mathrm{F}}_{\mathrm{data}} - \sum_{\mathrm{MC}\neq\mathrm{dijet}} N^{\mathrm{F}}_{\mathrm{MC}} \right] \\ + (F_{\mathrm{lead}} \times F_{\text{sub-lead}}) \times \left[N^{\mathrm{E}}_{\mathrm{data}} - \sum_{\mathrm{MC}\neq\mathrm{dijet}} N^{\mathrm{E}}_{\mathrm{MC}} \right], \end{aligned} \tag{11.4}$$

where the appropriate fake-factor as been applied to each validation region, serving as a multiplicative weighting which scales each region yield by a factor proportional to the expected fake τ event fraction. MC content corresponding to sub-dominant background processes is subtracted from each validation region, such that the final estimate pertains exclusively to QCD di-jet events.

Having verified through simulated event samples that the region is sufficiently signal-depleted, region B was determined safe to unblind such that a background estimate may be determined. The non-dijet SM MC content was subtracted from region B such that the corresponding event-yield may fairly be compared to the estimated dijet content of the region — non-dijet SM processes are suitably addressed with the standard ABCD method using MC samples. A total of 3.06×10^2 non-QCD MC events were observed in region B, comprised of 2.80×10^2 $Z \to \tau\tau$ events, 13.8 di-boson events, and 12.7 $t\bar{t}$ events. Fraction event counts arise from the weighting applied at the event generation level. The resulting validation of the estimation procedure is tabulated in Table 11.1.

Region B Yield	Region B MC	Region B (Yield − MC)	Region B Estimate
2.74×10^3	3.06×10^2	2.43×10^3	2.31×10^3

Table 11.1 results obtained from the closure-test of the hadronic fake-factor method of background estimation, as performed on dedicated validation data-regions.

An estimation of the fake τ content in region B may accordingly be reported as:

$$N^{B}_{\text{fakes}} = 2.31 \times 10^{3}, \tag{11.5}$$

This is found to agree with the unblinded region B yield of 2.43×10^{3} (having accounted for non-QCD background processes) to within 120 events. This difference is termed the *non-closure* uncertainty, and must be quantified such that the final background estimation is reported to within the correct degree of systematic error. Kinematic distributions for the estimated fake τ-jet content in this region are shown in Figure 11.5, together with those corresponding to real data observations and MC samples of alternative background sources. The close agreement of observed data to the estimated fake-τ content is a positive outcome for a region which is expected to be largely comprised of background events.

11.4 Systematic Uncertainties on Fake-τ Estimate

Systematic uncertainties are those inherently associated with a given apparatus or experimental procedure, rather than those attributable to random statistical fluctuation. By definition, therefore, a systematic uncertainty must be quantified for the fake-factor method of background estimation.

While a number of sources of systematic uncertainty exist in the fake-factor procedure — including those associated with the MC theoretical modelling of non-QCD processes — the non-closure uncertainty obtained in the validation process is found to be the most dominant. Accordingly, this will constitute the systematic error associated with the background estimation presented in this thesis.

The estimated background calculated in this section is found to agree with the 'true' background content in the validation region within 19%. As this is viewed as an uncertainty inherent to the estimation procedure itself, the equivalent percentage may be migrated to the final estimation in the SR such that a systematic error-band may be reported.

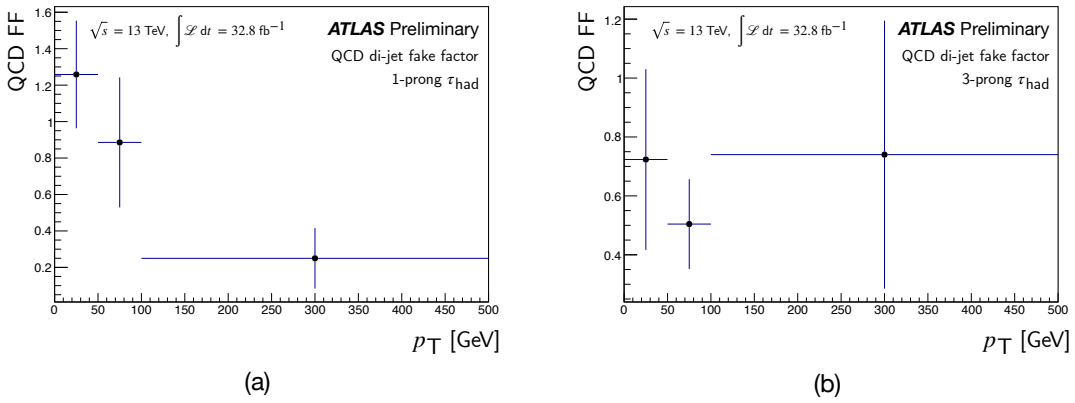

Figure 11.6 fake-factor values for *(a)* 1-prong and *(b)* 3-prong τ-leptons, as measured in the primary τ_{had} control region for the `Medium` working point displaced ID selection. Fake-factor values are binned in the p_T of the τ such that variations in background population as a function of τ kinematics are absorbed into the estimate.

11.5 Background Estimation in Signal Region

With the assurance that the fake-factor method will attain a background estimate within a tolerable degree of error, the procedure was applied to obtain an estimate of fake τ content in the primary SR. When extracted from the control region which corresponds to the `Medium` ID working point, as is used in the primary *ABCD* matrix, the fake-factor ratio is defined as:

$$F = \frac{N^{\texttt{Medium}}_{\text{data}} - \sum_{\text{MC}\neq\text{dijet}} N^{\texttt{Medium}}_{\text{truth-MC}}}{N^{\texttt{L!M}}_{\text{data}} - \sum_{\text{MC}\neq\text{dijet}} N^{\texttt{L!M}}_{\text{truth-MC}}}, \tag{11.6}$$

where $N^{\texttt{Medium}}_{\text{data}}$ denotes the number of τ-leptons within the CR which pass the `Medium` ID selection, and $N^{\texttt{L!M}}_{\text{data}}$ is the number of τ-leptons within the CR which pass the `Loose` ID selection while failing the `Medium` selection. From both counts the MC yield attributable to non-QCD SM processes is subtracted, such that an estimate for the number of dijet processes mimicking hadronically-decaying τs may be obtained. It should be noted that this definition is equivalent to those applied to the leading τ during the validation procedure — which may now be applied to both the leading and sub-leading τs, given the matching quality cuts imposed by each axis of the *ABCD* plane.

It may then be stated that the number of fake τs which pass the `Medium` working

point displaced ID selection is given by:

$$N_{\text{fake}}^{\texttt{pass-Medium}} = F \times \left(N_{\text{data}}^{\texttt{L!M}} - \sum_{\text{MC}\neq\text{dijet}} N_{\text{truth-MC}}^{\texttt{L!M}} \right). \tag{11.7}$$

In order to obtain an estimate for the number of fake *events* which pass signal selection, consideration must be given to each of the scenarios where the leading-p_{T} τ, the sub-leading-p_{T} τ, or both τs are fake:

$$N_{\text{fake}}^{\text{A}}(\text{events}) = N_{\text{fake}}^{\text{A}}(1\tau) + N_{\text{fake}}^{\text{A}}(2\tau). \tag{11.8}$$

In terms of the A, B, C, and D region definitions, this is given by:

$$\begin{aligned} N_{\text{fake}}^{\text{A}}(\text{events}) = F \times \left\{ \left[N_{\text{data}}^{\text{B}} + N_{\text{data}}^{\text{C}} \right] - \sum_{\text{MC}\neq\text{dijet}} \left[N_{\text{truth-MC}}^{\text{B}} + N_{\text{truth-MC}}^{\text{C}} \right] \right\} \\ + F^2 \times \left\{ N_{\text{data}}^{\text{D}} - \sum_{\text{MC}\neq\text{dijet}} N_{\text{truth-MC}}^{\text{D}} \right\}. \end{aligned} \tag{11.9}$$

The fake-factor is applied twice to region D such that an estimation of the contribution from events which produce two fake τs is obtained.

The fake-factors extracted from the primary CR are binned in p_{T}, such that variations in the background population as a function of momentum are accounted for. A variable bin-width strategy is employed such that the constraints imposed by statistical limitations are mitigated as far as possible. Further, individual fake-factors are derived for each τ-track multiplicity, in light of their differing detector signatures. The fake-factors applied in the background estimation process are shown in Figure 11.6.

Having performed adequately in validation, the afore-described estimation procedure was applied to the primary $ABCD$ matrix and an estimate of the background content in region A obtained. The true content of the signal region A remains blinded, pending the development of the full analysis into a more mature state. Statistical and systematic uncertainties (as derived from the validation non-closure error) are quoted. The expected number of QCD background events

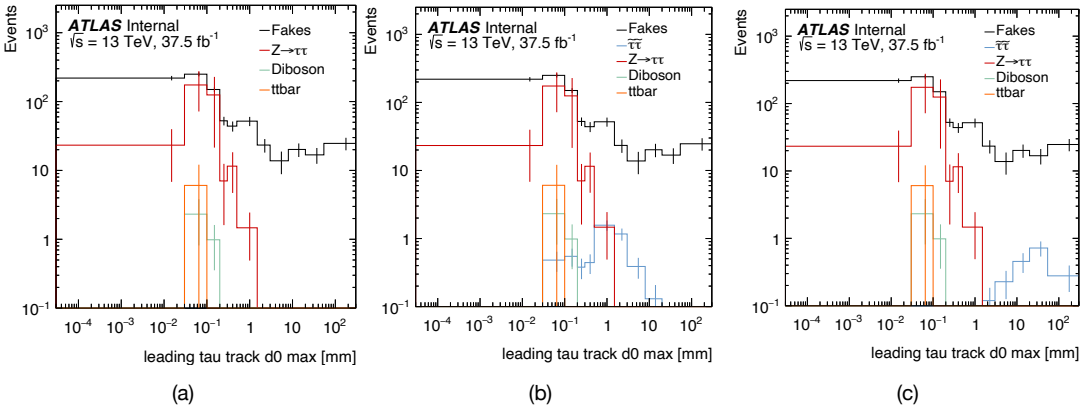

Figure 11.7 distributions of the largest impact parameter, d_0^{max}, of all tracks associated with each leading-p_{T} τ_{had} in the signal region, region A. Distributions are shown for all background sources, including the fake τ-jet estimate calculated in this chapter. Figures (*b*) and (*c*) additionally present the same distributions for the (500 GeV, 0.01 ns) and (500 GeV, 1.0 ns) signal points, respectively.

which mimic the τ_{had} decay signature is reported to be:

$$N_{\text{fakes}}^{\text{A(SR)}} = 866 \pm 19\% \text{ (sys.)}. \tag{11.10}$$

This result constitutes the estimated number events which produce quark-initiated or gluon-initiated jets expected to satisfy the signal requirements imposed on the $\tau_{\text{had}}\tau_{\text{had}}$ decay channel of the search presented in this thesis. This fake τ-jet estimation is compared to the background yield within the signal region which is attributable to other SM background processes in Table 11.2, where the background yield in the signal region is delineated by process of origin. Background distributions within the $\tau_{\text{had}}\tau_{\text{had}}$ signal region, including those which correspond to the fake τ-jet estimate, are shown in Figure 11.7 parameterised in a number of kinematic variables, together with MC signal samples. Distributions for additional MC signal mass and lifetime points are shown in 11.8. A complete estimate of the background content in the signal region of the $\tau_{\text{had}}\tau_{\text{had}}$ analysis channel, for all background processes under consideration, is obtained.

This chapter has introduced the analytical machinery with which the number of QCD-initiated background events expected to pollute the search for displaced τ-leptons in the fully-hadronic decay channel is estimated. The expected 'fake' τ

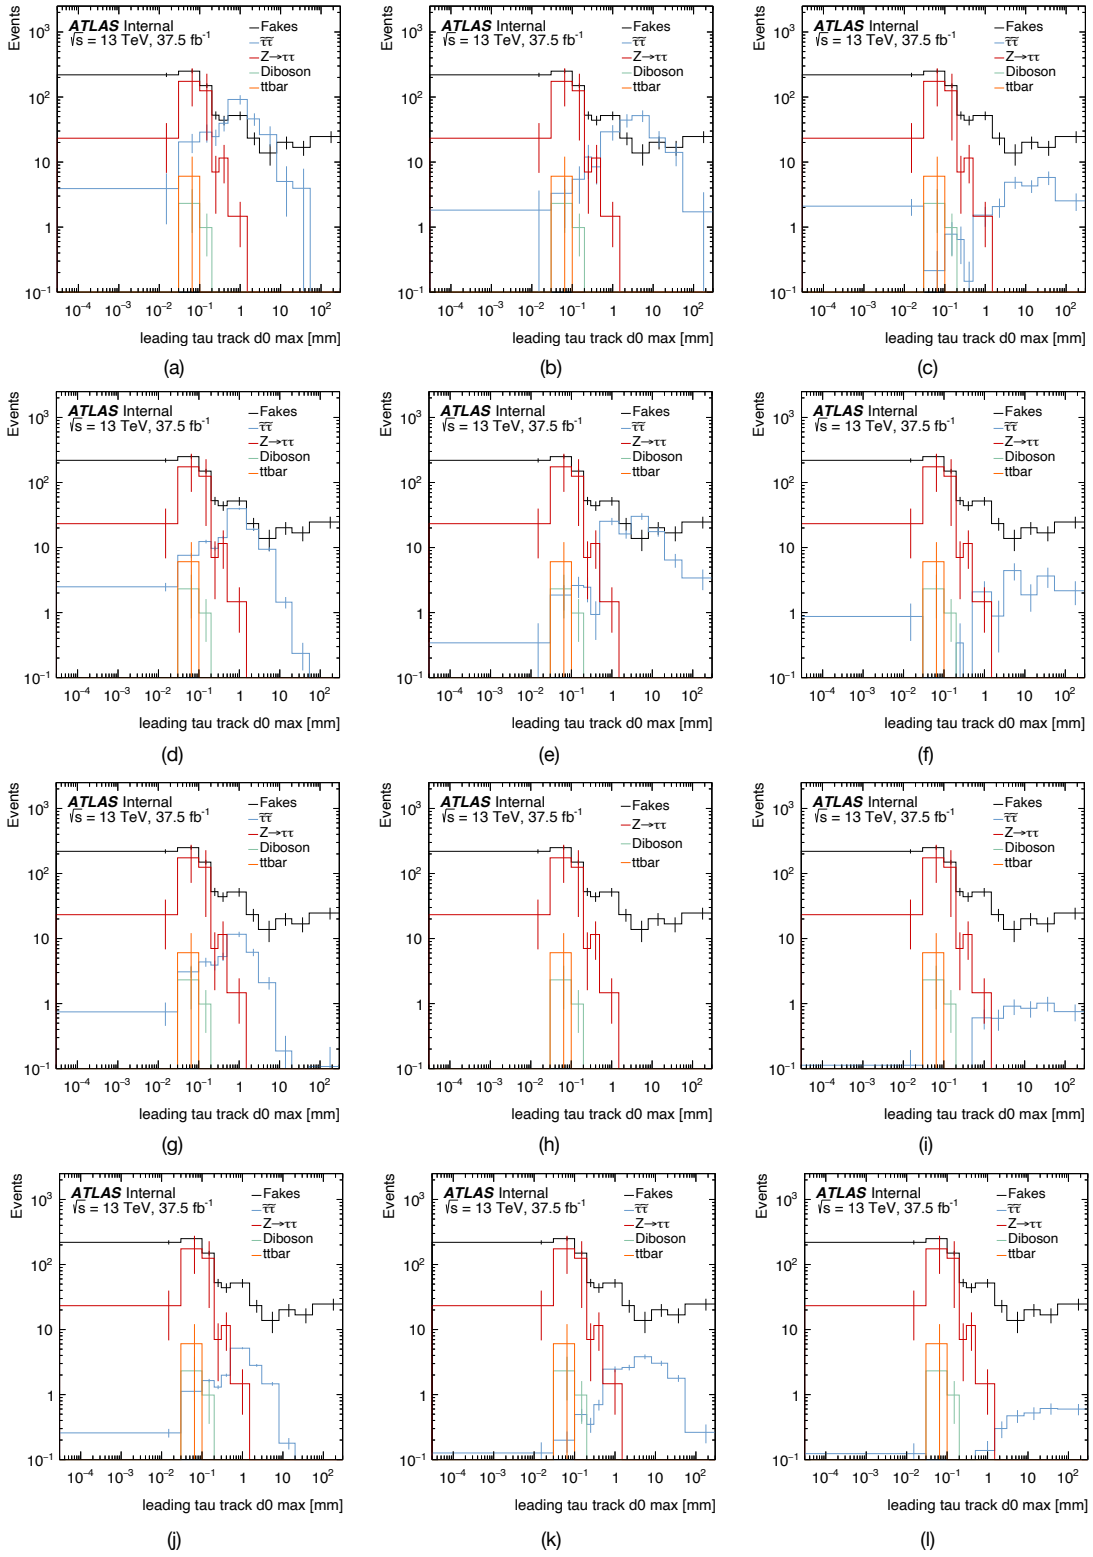

Figure 11.8 shows the same d_0^{max} distributions as Figure 11.7 for additional MC signal $(m, c\tau)$-points. Distributions are shown for signal events within the signal region with lifetime values of (*left*) 0.01 ns, (*middle*) 0.1 ns and (*right* 1.0ns). From top to bottom, each row respectively presents signal d_0^{max} distributions for 100 GeV, 200 GeV, 300 GeV, and 400 GeV mass values.

Background Process	Event Yield
Fake τ-jet	866
$Z \to \tau\tau$	342
$t\bar{t} \to \tau\tau$	6
Di-boson	3

Table 11.2 Event yield by background process in the $\tau_{had}\tau_{had}$ signal region.

yield is reported in conjunction with associated systematic uncertainties, where the latter is derived from the non-closure error obtained in the validation of the estimation procedure.

A reliable quantification of background effects is required if accurate limits on the production of displaced τ-leptons are to be derived. The statistical analysis with which such limits are obtained, together with the manner in which the expected background event-yield is accommodated, is now presented.

> *Every genuine test of a theory is an attempt to falsify it.*
>
> — Karl Popper

12

Statistical Analysis and Limit Setting

With the appropriate theoretical understanding and experimental methodology set in place, the search for displaced τ production is equipped to make new statements concerning fundamental physics. This information is conveyed through results in two primary channels: (1) exclusion limits on the $\tau\bar{\tau} \to \tau\tau$ process, and (2) model-independent constraints on the cross-section of states present in the phase-space under investigation, including those yet to be proposed by theorists or phenomenologists. This chapter presents provisional model-dependent results in the $\tau_{\text{had}}\tau_{\text{had}}$ channel, before discussing the advancements set to improve the exclusion power of the analysis and place more stringent constraints on the elusive physics beyond the SM.

12.1 Overview

A discovery may only be reported in the search for displaced τ-leptons should an excess of 5 standard deviations (5σ) above the expected SM background be observed. In the absence of such an excess, it is desirable to place a limit upon the number of events observed in the phase-space associated displaced τ signatures, such that the viability of a given model of new physics may be assessed against this observation. In this manner, a model-independent constraint on the cross-section of LLP models beyond the SM may be achieved. The design of a statistical analytical procedure by which to obtain such constraints begins with the specification of a 'statistical estimator'.

12.2 Statistical Estimators

Measurements performed with collision data aim to reveal the nature of the 'true' distribution of the physical observable under investigation as it is valued in nature. The finitude of both the collision luminosity and the run-time of the experiment prohibits the act of doing so with infinite precision. It is therefore necessary to *estimate* the value of observables with a metric which is sensitive to their true distribution.

By their very nature, statistical estimators are subject to statistical fluctuations. For this reason, it is desirable to determine the value of estimators with the largest data-set available, such that the effect of fluctuations is minimised and an accurate estimate may be obtained.

The choice of statistical estimator is an example of a design decision which must be made at the conception of the analysis at hand. The estimator of choice should be *unbiased*: that is to say, the metric should asymptotically converge upon the true value of the observable, without skewing in a manner which is unreflective of the variation of the underlying parameter of measurement. A popular form of unbiased estimator are those based upon the metric of *likelihood*.

12.2.1 Likelihood-Based Estimators

The likelihood provides a metric with which to encode the compatibility of data with a given statistical model , expressed as the product of individual probability density functions (PDFs). A general expression for the likelihood is given by:

$$L\left(\vec{\Theta}\middle|\vec{x}\right) = \prod_i p_i\left(x_i\middle|\vec{\Theta}\right), \tag{12.1}$$

where the likelihood is expressed as a function of a (vectorised) set of parameters, $\vec{\Theta}$, given the observed data, $\vec{x}$. This is equated to a product of i PDFs, each of which is a function of statistically-independent measurements x_i given the parameters encoded by $\vec{\Theta}$. The parameters encapsulated by $\vec{\Theta}$ may be decomposed into a parameter(s) of interest, denoted by $\vec{\mu}$, and the remaining *nuisance parameters*, $\vec{\theta}$.

The ability of the likelihood to capture the degree of agreement between data and a given model may be exemplified through comparison to the χ^2 metric. The χ^2-fit is a common measure of the 'goodness of fit' between two distributions:

$$\chi^2 = \sum_{x\in\{x\}} \frac{\left(y(x) - f(x)\right)^2}{f(x)}, \tag{12.2}$$

where $y(x)$ gives the distribution of x as measured in data, and $f(x)$ describes the distribution of the same variable in accordance with the predictions of a given model. In the limit of large statistics, the statistical variance is approximately given by $\sigma^2(x) \approx f(x)$. In the context of data collected at the LHC, it is therefore possible to reduce the χ^2 metric to:

$$\chi^2 = \sum_{x\in\{x\}} \frac{\left(y(x) - f(x)\right)^2}{\sigma^2(x)}. \tag{12.3}$$

The quantity χ^2 can accordingly be viewed as providing a distance metric between the predictions of a given model and the measured data, scaled by the variance of each statistical bin in which the measurement is performed.

The likelihood of observing a set of bin values $\{y(x)\}$ given the model $\mathcal{M}$ may be

expressed as:

$$L = \prod_{x \in \{x\}} p(y|\mathcal{M}). \tag{12.4}$$

Under the assumption that the random variable $y(x)$ in each bin is described by a Gaussian probability density function, $\mathcal{N}(y; f(x), \sigma(x))$, the logarithm of the likelihood is:

$$\begin{aligned} \log_e L &= \log_e \left[\prod_{x \in \{x\}} p(y|\mathcal{M}) \right], \\ &= \sum_{x \in \{x\}} \log_e \left[\mathcal{N}(y(x); f(x), \sigma(x)) \right], \\ &\approx \sum_{x \in \{x\}} -\frac{(y(x) - f(x))^2}{2\sigma^2(x)}, \\ &= -\frac{1}{2}\chi^2. \end{aligned} \tag{12.5}$$

As the χ^2 goodness of fit metric may be recovered therefrom, the definition of the likelihood can be viewed intuitively as providing a measure of agreement between a given model and the data collected by experiment. For this quantity to serve ..., a suitable instantiation for the likelihood in the context of displaced τ production must be set.

12.2.2 Likelihood Instantiation

This analysis expresses the likelihood function in terms of the various regions defined for each search channel, taking the relevant search regions to form the bins of the expression. As the measurement of the event-yield in each region may be regarded as an independent counting experiment, the binned likelihood may be defined as:

$$L\left(\mu, \vec{\theta} \middle| \vec{x}\right) = \prod_{i=1}^{N} \text{Pois}\left(x_i \middle| \mu \cdot s_i\left(\vec{\theta}\right) + b_i\left(\vec{\theta}\right)\right) \times \prod_j C\left(\alpha_j | \theta_j\right), \tag{12.6}$$

where it is assumed that the number of observed events, n, in each of the search regions included in the likelihood estimator will follow a Poisson distribution of the form:

$$\text{Pois}\left(n|\mu s+b\right)=\frac{\left(\mu s+b\right)^{n}}{n!}e^{-\left(\mu s+b\right)}. \tag{12.7}$$

The Poisson distribution for each likelihood bin, i, is parameterised in terms of the expected signal yield, $s=s_i\left(\vec{\theta}\right)$ and expected background yield, $b=b_i\left(\vec{\theta}\right)$, where assumptions regarding signal content may scale with the parameter μ.

The additional term appended to Equation 12.6 serves as a normalisation factor which imposes external constraints on each of the nuisance parameters harboured by $\vec{\theta}$. In practise, this enables pre-determined systematic uncertainties associated with background estimation procedures and theoretical uncertainties to be incorporated into the likelihood. This analysis instantiates this term to a Gaussian distribution, thereby accommodating small fluctuations in the systematic error-rate where favoured by the observed data. The inclusion of such Gaussian constraints results in a final likelihood expression of the form:

$$\begin{aligned}L\left(\mu,\vec{\theta}\middle|\vec{x}\right)=&\prod_{i=1}^{N}\frac{\left(\mu s+b\right)^{x}}{x!}e^{-\left(\mu s+b\right)}\times\prod_{\theta_{\text{sig}}\in\vec{\theta}_{\text{sig}}}\text{Gauss}\left(\theta_{\text{sig}}|0,\sigma_{\text{sig}}\right)\\&\times\prod_{\theta_{\text{bkg}}\in\vec{\theta}_{\text{bkg}}}\text{Gauss}\left(\theta_{\text{bkg}}|0,\sigma_{\text{bkg}}\right),\end{aligned} \tag{12.8}$$

where each Gaussian is constructed with a mean value equal to the relevant nuisance parameter, and a standard deviation equal to the uncertainty associated with the parameter in question:

$$\prod_{\theta\in\vec{\theta}}\text{Gauss}\left(\theta|0,\sigma\right)=\prod_{\theta\in\vec{\theta}}\frac{1}{\sqrt{2\pi}}e^{-\theta^{2}/2\sigma^{2}}. \tag{12.9}$$

12.2.3 Incorporating Region Event-Yields

With the appropriate definition for the likelihood estimator specified, the manner in which the relevant search channels are incorporated into the likelihood binning

scheme is now considered.

The double-hadronic decay channel considered in this search is incorporated into the profile likelihood expression as a single bin. A Poisson distribution is configured for the $\tau_{\mathsf{had}}\tau_{\mathsf{had}}$ signal region, as defined in the preceding chapter.

The expected event-yield attributable to background processes is taken to be the fake-τ estimation obtained via the fake-factor method, together with additional SM background events obtained through MC simulation and the standard ABCD method [160]. The expected background event-yield is therefore given by:

$$b_{\tau_{\mathsf{had}}\tau_{\mathsf{had}}}(\mathsf{total}) = b_{\tau_{\mathsf{had}}\tau_{\mathsf{had}}}(\mathsf{fake}\text{-}\tau) + b_{\tau_{\mathsf{had}}\tau_{\mathsf{had}}}(\mathsf{SM})\,. \tag{12.10}$$

The appropriate systematic uncertainties will be incorporated into the likelihood as nuisance parameters. In the case of the fake-τ estimation, the non-closure uncertainty was found to be the dominant source of error. Theoretical uncertainties associated with the MC simulation of SM processes, together with the statistical error associated with the ABCD method, will be determined for the non-fake background content.

12.3 Limit Setting

Equipped with a suitable estimator in the form of the binned profile likelihood, the data may now be examined for evidence of displaced τ production. In practise, this corresponds to setting a limit on the number of events which may arise from the decay of a long-lived state to τ-leptons in the phase-space under consideration — and, indirectly, the cross-section of the corresponding process. For a given model of LLPs beyond the SM to be deemed compatible with the derived limits, it must predict a cross-section which corresponds to an observable signal-yield within this limit. BSM models whose predicted signal-yield fall outwith this limit may be considered *excluded* from the realms of possibility (with respect to their predicted decay signatures in the relevant phase-space).

The process of obtaining such limits commences with *hypothesis testing*.

12.3.1 Hypothesis Testing

Hypothesis testing determines the extent to which data favour one scenario over another. This requires two hypotheses to be defined:

1. **the null hypothesis,** H_0: corresponding to the scenario where the data are compatible with the SM, without the presence of additional BSM states;
2. **the alternative hypothesis,** H_1: where the statistical model which describes the SM together with an additional BSM presence is favoured.

An additional quantity known as the *test statistic* must further be defined. This should be chosen such that the value of this quantity, as measured in data, enables the strength of agreement with data to be determined for each hypothesis. The process of hypothesis testing then assesses whether the null hypothesis may be rejected — science is an exercise in falsification, as argued by Popper.

The *Neyman-Peasrson lemma* states that the optimal choice of test statistic for a simple test of two hypotheses is given by the likelihood ratio:

$$\lambda = \frac{L\left(x|H_1\right)}{L\left(x|H_0\right)}. \tag{12.11}$$

This, however, does not apply in the scenario where free parameters are present. For the purposes of this analysis, the test statistic may instead be amended to:

$$\lambda = -2\log_e\left(\frac{L\left(x|H_1\right)}{L\left(x|H_0\right)}\right) = -2\log_e\left(\frac{L\left(x|\mu_1, \hat{\hat{\vec{\theta}}}_1\right)}{L\left(x\middle|\hat{\mu}_0, \hat{\vec{\theta}}_0\right)}\right), \tag{12.12}$$

where the denominator is the global maximised likelihood — denoting respectively the PoI and nuisance parameters which maximise the likelihood with $\hat{\mu}$ and $\hat{\vec{\theta}}$ — and the numerator is the likelihood calculated for a fixed value of μ, where $\hat{\hat{\vec{\theta}}}$ are the values of nuisance parameters for which this likelihood is maximised.

The distribution of the test statistic — in this case, the profile likelihood ratio — must now be obtained for each hypothesis, $g\left(\lambda|H_0\right)$ and $g\left(\lambda|H_1\right)$. The statistical limitations of the analysis in the $\tau_{\text{lep}}\tau_{\text{had}}$ channels render the standard asymptotic

approximation [161] inapplicable. A large number of 'pseudo-experiments' are therefore computed, where samples are drawn from the likelihood using MC methods. A value of the test statistic, λ, is then determined for data drawn from each pseudo-experiment: $x' \sim L(\mu'|\hat{\hat{\vec{\theta}}})$.

12.3.2 From Protons to *p*-Values: Expected Confidence Levels

Expected limits on the production of of displaced τ-leptons will be obtained with the $\mathsf{CL_s}$ method.

With this method, the values of μ which are not rejected define a confidence interval with a confidence level of $\mathsf{CL} = 1 - \alpha$. The term α represents the probability that the null hypothesis is erroneously rejected, and is conventionally taken to be $\alpha = 0.05$ such that confidence intervals may be reported at the 95% confidence level (CL).

Given the test statistic of the likelihood ratio, λ, the *p*-values which correspond to each hypothesis may be computed:

$$p_0 = \int_{\lambda_{\mathsf{obs}}}^{+\infty} \mathrm{d}\lambda \, g\,(\lambda|H_0) = 1 - \mathsf{CL_b}, \tag{12.13}$$

$$p_1 = \int_{-\infty}^{\lambda_{\mathsf{obs}}} \mathrm{d}\lambda \, g\,(\lambda|H_1) = \mathsf{CL_{s+b}}. \tag{12.14}$$

Using the $\mathsf{CL_s}$ method, the null (SM-only) hypothesis is rejected if the ratio given by $\mathsf{CL_s}$ is valued at $\mathsf{CL_s} < \alpha$, where:

$$\mathsf{CL_s} = \frac{\mathsf{CL_{s+b}}}{\mathsf{CL_b}}. \tag{12.15}$$

In this manner, limits mat be placed upon the cross-section of displaced-τ production.

12.4 From Lagrangians to Limits: Expected Exclusions

The results presented in this section convey the sensitivity of the search for displaced τ-leptons which is expected upon the full unblinding of collider data in the $\tau_{had}\tau_{had}$ signal region. Sensitivity to such new physics is quantified as upper-bound limits on the production of displaced τ-leptons which would be visible to the search detailed herein. Through the determination of limits of this form, the compatibility of various BSM models of new physics with ATLAS measurement data can be determined — models which are deemed incompatible with observation are said the be *excluded*.

The expected exclusion power of the search for displaced τ-leptons is expressed in terms of model-independent limits on their production, together with a specific interpretation performed in the context of the ~~ττ~~ BSM model. Collectively, such results represent the main scientific outcome of this PhD.

12.4.1 Model-Independent Results

A binned maximum likelihood fit is performed in the manner introduced in Section 12.3.1. This is performed under the background-only hypothesis on the *Asimov* dataset [161, 162] — the dataset constructed by setting the 'observed' data equal to the signal and background expectations derived previously. This is appropriate for the purposes of estimating the expected limits attainable with the search upon the full unblinding of the $\tau_{had}\tau_{had}$ signal region. The Asimov dataset is named after the science fiction writer Isaac Asimov, whose book *Franchise* envisages political elections held not by collective electorates, but rather by a single 'representative' voter. The Asimov dataset is, by construction, the single most representative dataset, and is in this manner representative of the Asimov book.

Likelihood scans of the signal yield parameter are presented in Figure 12.1, where results are shown for two choices of fit variables: the largest value of impact parameter, d_0^{max}, of tracks associated with the (*blue dashed-line*) leading-p_T τ-lepton, and (*black dashed-line*) either the leading-p_T or sub-leading-p_T τ-lepton. Likelihood scans are presented for the 300 GeV ~~ττ~~ signal for the (a) 0.01 ns, (b) 0.1 ns, and (c) 1.0 ns lifetime points.

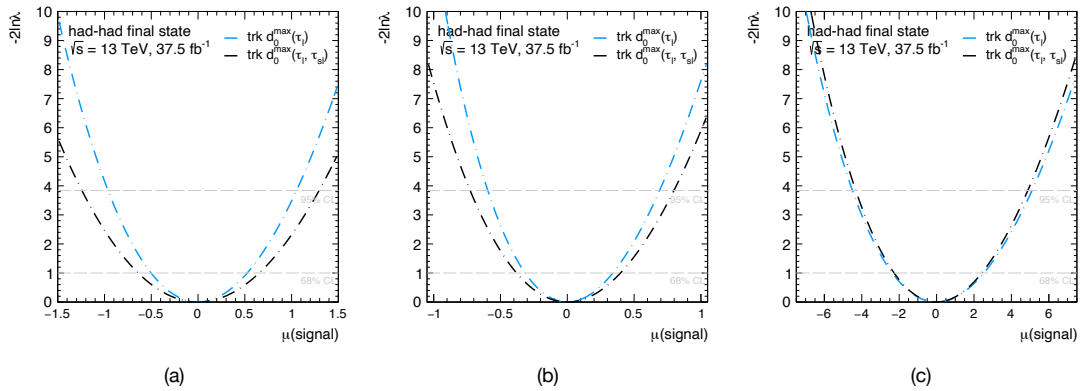

Figure 12.1 likelihood scan of the signal event parameter in the maximised likelihood function (or, minimised negative logarithm of the likelihood function). Fit is performed with the largest impact parameter of tracks associated with the leading-p_{T} τ-lepton, d_0^{max}, which is shown to perform for favourably than that which is performed with the largest value of d_0 taken across the collective pool of tracks associated with both the leading-p_{T} and sub-leading-p_{T} τ. Results are shown for (*a*) 0.01 ns, (*b*) 0.1 ns, and 1.0 ns lifetime points of the 300 GeV $\tilde{\tau}$.

The maximum likelihood fit results in a value of 0 for the signal parameter, μ, while the value of the parameter quantifying the fake τ-jet background converges upon 0 ± 1. This represents the behaviour expected of the fitting procedure. The post-fit results are parameterised in the d_0^{max} of the leading-p_{T} τ-lepton ($d_0^{\mathrm{max}}(\tau_{\mathrm{lead}})$) and presented in Figure 12.2. Distributions are provided for each physics process included in the fitting procedure.

With the maximum likelihood fitting procedure in place, upper-bound limits are placed upon the number of signal events in the $\tau_{\mathrm{had}}\tau_{\mathrm{had}}$ signal region at the 95% CL, derived with the $\mathrm{CL_s}$ method introduced in Section 12.3.2. This procedure was performed with the *asymptotic approximation* [161, 162], which approximates the behaviour of the profile likelihood ratio test statistic in the limit of large datasets (such as that provided by the $\tau_{\mathrm{had}}\tau_{\mathrm{had}}$ analysis channel of this search). The resultant limits are expressed in terms of both signal event yield and signal production cross-section in the phase-space which corresponds to the $\tau_{\mathrm{had}}\tau_{\mathrm{had}}$ signal region, and are tabulated in Table 12.1.

The results tabulated in Table 12.1 constitute the collective outcome of the work detailed in this thesis. Such results highlight the expected sensitivity to new physics of the search described herein, upon the full unblinding of the $\tau_{\mathrm{had}}\tau_{\mathrm{had}}$ signal region. Upon the unblinding of this search to real pp-collision data, entirely new phase-space will be exposed to experimental scrutiny to the level of sensitivity

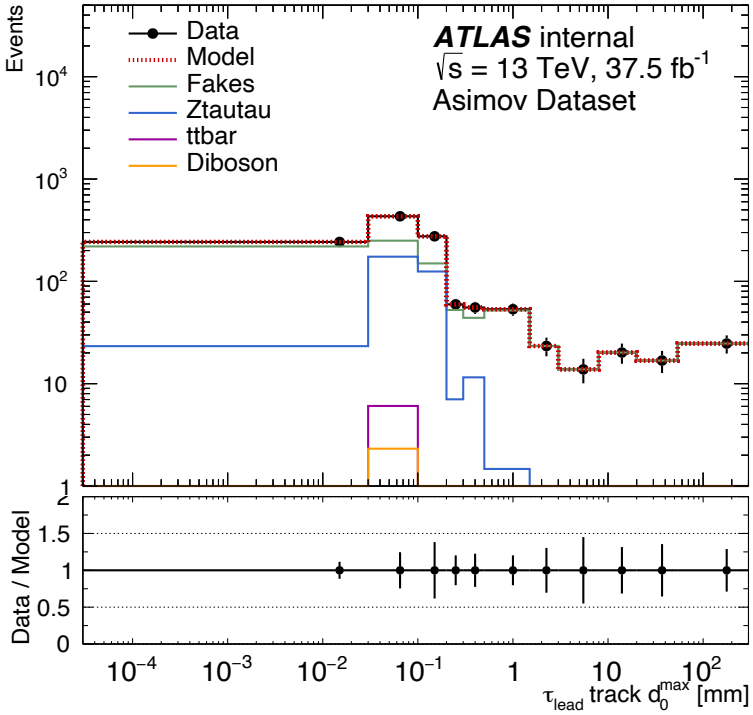

Figure 12.2 post-fit results of the binned maximum likelihood fit to the Asmimov dataset, constructed under the assumption of the background-only hypothesis in the $\tau_{had}\tau_{had}$. The fit is parameterised as a function of $d_0^{max}(\tau_{lead})$, as previously discussed.

Discovery Channel	Signal Yield	$\sigma_{prod.}$ (LLP $\rightarrow \tau\tau$) [fb]
$\tau_{had}\tau_{had}$	109	2.9

Table 12.1 95% CL upper limits on the signal yield and production cross-section of displaced τ-leptons in the phase-space defined by the $\tau_{had}\tau_{had}$ signal region.

expressed in this thesis.

The model-independent results presented in this section represent the expected results of the search for displaced τ-leptons and the fruits of this thesis, constraining the space of possible remedies to the short-falls of the SM in a manner which is decoupled from any given theoretical prediction of new physics. It is important, however, to demonstrate that such results may be interpreted in the context of well-motivated BSM models. This serves not only to benchmark the sensitivity of the search to hypothesised BSM processes of interest, but further highlights the ability for the compatibility of theories (perhaps not yet conceived) with the derived results to be ascertained. This results of this search are accordingly interpreted in the context of the gauge-mediated SUSY-breaking model of $\tilde{\tau}\tilde{\tau}$-production.

12.4.2 Model-Specific Interpretations

An interpretation of the results of the analysis is performed in the context of the super-symmetric model of the $\tilde{\tau}$ discussed throughout this thesis. Limits are derived on the $pp \to \tilde{\tau}\tilde{\tau}$ production cross-section at the 95% CL. As was true for the main model-independent results afore-presented, this is achieved with the $\mathsf{CL_s}$ method under the asymptotic approximation.

Limits on the model-specific signal event-yield in the $\tau_{\mathsf{had}}\tau_{\mathsf{had}}$ signal region are presented in Table 12.2. As before, such results constitute the *expected* limits attainable upon the full unblinding of collision data within the $\tau_{\mathsf{had}}\tau_{\mathsf{had}}$ signal region. The $\tau_{\mathsf{had}}\tau_{\mathsf{had}}$ channel is observed to be most sensitive to $\tilde{\tau}$-signals with high mass, with a particularly strong sensitivity to the signal-points which correspond to a lifetime of 0.01 ns. This was found outwith this thesis to be counter-balanced by the particularly high sensitivity to higher lifetimes achieved by the $\tau_{\mathsf{lep}}\tau_{\mathsf{had}}$ analysis channels.

Figure 12.3 further presents 95 CL limits on the model-specific cross-section $\sigma(pp \to \tilde{\tau}\tilde{\tau})$ as a function of $\tilde{\tau}$-mass, $m_{\tilde{\tau}}$. This result incorporates all decay channels considered in the broader analysis, including the two channels which correspond to $\tau_{\mathsf{lep}}\tau_{\mathsf{had}}$ decays. Semi-leptonic analysis channels were incorporated into the fitting procedure with the 'modified-ABCD method' [163], where the ansatz which underpins the traditional ABCD method is encoded in the likelihood function [161, 162].

$\tilde{\tau}$ Mass [GeV]	$\tilde{\tau}$ Lifetime [ns]	95% CL Limit on $\tau_{had}\tau_{had}$ Yield
100	0.01	0.13
	0.1	0.13
	1.0	1.02
200	0.01	0.31
	0.1	0.22
	1.0	1.41
300	0.01	1.05
	0.1	0.69
	1.0	5.11
400	0.01	2.27
	0.1	1.59
	1.0	8.58
500	0.01	6.79
	1.0	10.33

Table 12.2 95%CL expected limits on $\tilde{\tau}\tilde{\tau}$ signal yield in the $\tau_{had}\tau_{had}$ search channel. Sensitivity is observed to be greatest for high-mass $\tilde{\tau}$ states.

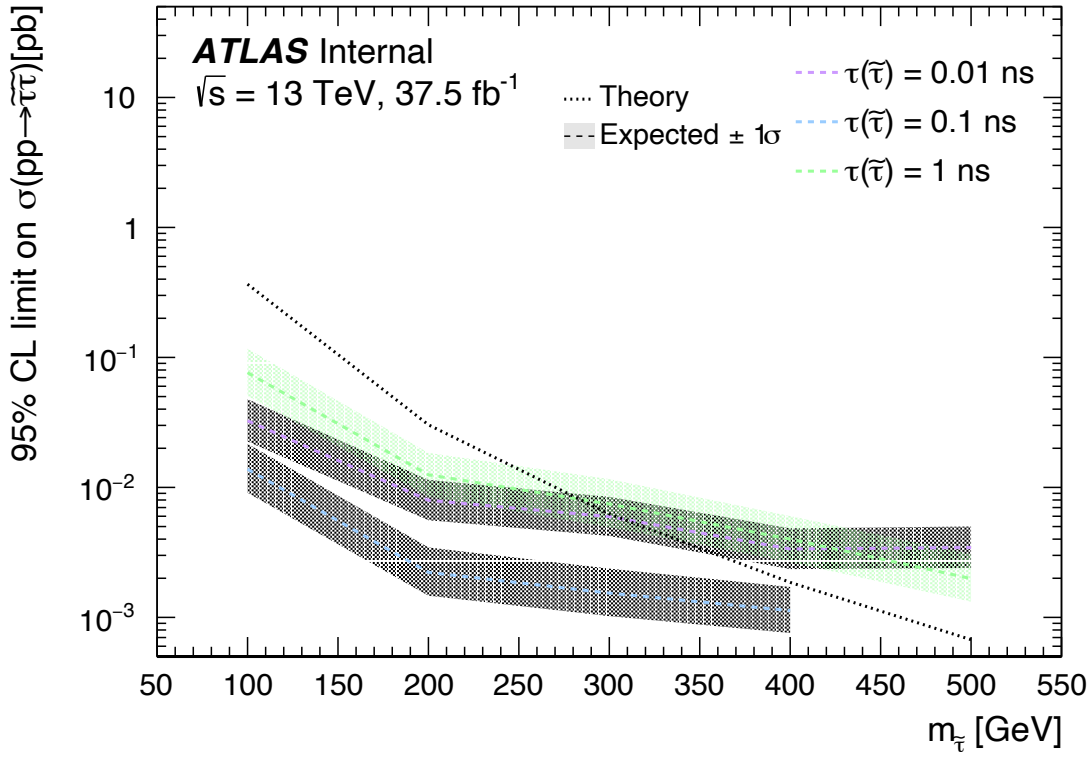

Figure 12.3 95% CL exclusion limits in the $\tilde{\tau}\tilde{\tau}$ model interpretation as a function of $\tilde{\tau}$ mass, $m_{\tilde{\tau}}$. Limits are derived on the $\tilde{\tau}$ production cross-section for three lifetime points: 0.01 ns, 0.1 ns, 1.0 ns. The theoretical cross-section for this signal model (*black dashed-line*) is excluded by the expected limit for lower values of $m_{\tilde{\tau}}$.

The expected limits presented in Figure 12.3 demonstrate particularly strong sensitivity to lower-mass $\tilde{\tau}$-states, once all analysis channels are incorporated into the fitting procedure. Production of higher-mass $\tilde{\tau}$-states is best excluded by the expected limits for lifetime points corresponding to 0.1 ns. The expected limits obtained for the MC $\tilde{\tau}\tilde{\tau}$ model thereby provide a benchmark of the LLP search sensitivity in the manner envisaged in this thesis, while further providing a physically-insightful probe of a well-motivated BSM model.

12.4.3 Treatment of Systematic Uncertainties

The sole source of systematic error currently incorporated into the maximum likelihood fitting procedure is that which corresponds to the non-closure uncertainty associated with the background estimation procedure. A conservative systematic uncertainty of 19% is applied to the fit such that the non-closure error examined in the previous chapter is accommodated by the limit-setting procedure.

While this is likely to comprise the largest systematic effect, other sources of systematic error are, of course, under active investigation. Uncertainties which arise from additional systematic effects will be incorporated into the fitting procedure upon their complete determination.

The uncertainty associated with large radius track reconstruction largely arises due to secondary particle interactions with the material budget of the ATLAS detector. To determine the effect of this error, the material quota of the inner detector is varied by 5% in a `GEANT4` [164] simulation of the ATLAS detector. In accordance with the same procedure, the material budget of the Insertable B-Layer and Inner Detector are respectively scaled by 10% and 25%. The effects induced by such variations on the efficiency of the LRT algorithm (and, by extension, the displaced τ-ID RNN) is under investigation, and will be included as a systematic uncertainty in the maximum likelihood fitting procedure.

The uncertainty associated with the luminosity conditions of the collision environment, as determined by the LUCID-2 detector [165], is measured to be 0.84%. This error will be included as a systematic uncertainty in the limit-setting procedure.

The main uncertainty associated with object reconstruction is that attributed to the Tau Energy Scale calibration, which is determined, via $Z \to \tau\tau$ events, to be $\sim 4\%$. This will similarly be incorporated into the fitting procedure.

Relevant uncertainties of a theoretical nature include the choice of normalisation scale and factorisation scale in the MC calculation of theoretical predictions, as introduced in Chapter 2. To determine their effect on the exclusion power of this search, the value of each scale will be varied in both directions — the resultant deviation from the nominal MC predictions will be applied as a systematic uncertainty. The accommodation of errors of this nature in the fitting procedure will likely result in a small reduction of the upper-bound limits attainable in this search.

This chapter has presented the main scientific results of the research conducted in this thesis. The expected exclusion power of the search for displaced τ-leptons has been presented in a form independent of any given model of new physics, thereby maximising sensitivity to physics beyond the SM without favour or bias. the physics reach of the search is determined with a specific super-symmetric

interpretation, where concrete predictions of new particle states are shown to be excluded by the expected limits on displaced-τ production achieved in this thesis.

The search for physics beyond the Standard Model continues apace.

The Path Continues

This thesis has detailed the theoretical, mathematical, and conceptual underpinnings of the Standard Model of particle physics, before surveying the observations and arguments by which it may be deemed to be incomplete. It was then argued herein that the conjectured existence of long-lived particles beyond those harboured in the Standard Model would provide a compelling and well-motivated solution to the short-comings of current theoretical thinking. The privileged role which the leptons of the third generation ought to assume in the decay of such long-lived states was subsequently examined.

It is against this backdrop that the search for displaced τ-lepton production in proton-proton interactions at the Large Hadron Collider was conducted with the ATLAS experiment.

In service of this search, bespoke event reconstruction capabilities were developed for the ATLAS experiment. This includes the development of neural network classifiers optimised for the identification of τ-leptons which emanate from hitherto-unknown long-lived particles, and the construction of a novel high-level trigger system — a trigger which leverages the enhanced discriminating power of the afore-mentioned neural network, together with new tracking techniques with which the acceptance of highly-displaced τ-production is increased. The latter development represents a personal contribution to this search.

The development of such algorithmic technologies enabled the proton-proton interactions recorded during the second run of the Large Hadron Collider to be probed. The search for displaced τ-leptons in this dataset initiated with the design of suitable strategies for trigger selection, object selection, and statistical region design. This thesis presented the strategic decisions made with respect to such considerations, as they pertain to the doubly-hadronic decay channel of the analysis: the $\tau_{\mathsf{had}}\tau_{\mathsf{had}}$ channel. Analytical design of the the $\tau_{\mathsf{had}}\tau_{\mathsf{had}}$ channel is a

further personal contribution to this research effort.

An irrepressible fraction of background events will, of course, evade the mitigating selections imposed by the definition of such regions, thereby polluting the measured event yield in the corresponding signal region. The estimation of the event yield attributable to background processes in the $\tau_{\text{had}}\tau_{\text{had}}$ channel posed a particularly large challenge, as the lack of a leptonic trigger handle enables QCD-induced di-jet events to mimic the jet-like signature of both τ-leptons in this signature. Further, the simulation-based methods for estimating the expected yield of such events cannot be relied upon in the phase-space associated with this final-state. A data-driven method was accordingly employed in this thesis, with which the background contribution to the $\tau_{\text{had}}\tau_{\text{had}}$ event yield was successfully estimated to an accuracy within 20%. This represents a significant personal contribution to the search for displaced τ-leptons.

The final chapter of this thesis presented the expected sensitivity of this search effort. This was expressed both as a model-independent upper-bound limit on the production cross-section of long-lived particles which is compatible with the observations of this research, and a specific interpretation of this limit in the context of super-symmetric $\tilde{\tau}$-production in the phase-space probed by this investigation. Such limits represent the main outcome of this doctoral research. The expected exclusion power of this analysis of run-2 LHC data is found to be competitive with that attained by a recent displaced-leptons analysis [166], which was performed on a combined run-2 and run-3 LHC dataset approximately $3\times$ larger, and which neglects to consider final states which contain the τ-lepton. Moreover, the exclusion reach of the search presented herein extends to entirely new quantum phase-space — thereby exposing an entirely new space of model parameters to the light of empirical scrutiny.

Upon the full unblinding of the collision data selected by the $\tau_{\text{had}}\tau_{\text{had}}$ signal region, a further constraint will be placed upon the parameter space in which nature may choose to hide. Until that time, the path to discovery continues.

Bibliography

[1] The ATLAS Collaboration. "The ATLAS Trigger System for LHC Run 3 and Trigger Performance in 2022". In: *Journal of Instrumentation* 19.06 (June 2024), P06029. ISSN: 1748-0221. DOI: 10.1088/1748-0221/19/06/p06029. URL: http://dx.doi.org/10.1088/1748-0221/19/06/P06029.

[2] X. Fan et al. "Measurement of the Electron Magnetic Moment". In: *Physical Review Letters* 130 (2023). DOI: 10.1103/physrevlett.130.071801. URL: http://dx.doi.org/10.1103/PhysRevLett.130.071801.

[3] Mark Thomson. *Modern Particle Physics*. Cambridge University Press, 2013. ISBN: 978-1-1070-3426-6.

[4] Matthew D. Schwartz. *Quantum Field Theory and the Standard Model*. Cambridge University Press, 2013. ISBN: 978-1-1070-3473-0.

[5] I. I. Bigi and A. I. Sanda. *CP Violation*. 2nd ed. Cambridge University Press, 2009. ISBN: 978-0-5115-8101-4.

[6] Michael Edward Peskin and Daniel V. Schroeder. *An Introduction to Quantum Field Theory*. Westview Press, 1995. ISBN: 0-201-50397-2.

[7] Lewis H. Ryder. *Quantum Field Theory*. 2nd ed. Cambridge University Press, 1996. ISBN: 978-0-5214-7814-4.

[8] Frank Wilczek. "Quantum Field Theory". In: *Reviews of Modern Physics* (1999). DOI: 10.1103/revmodphys.71.s85. URL: http://dx.doi.org/10.1103/RevModPhys.71.S85.

[9] Andy Buckley, Christopher White, and Martin White. *Practical Collider Physics*. IOP Publishing, 2021. ISBN: 978-0-7503-2444-1. DOI: 10.1088/978-0-7503-2444-1.

[10] James E. Dodd and Ben Gripaios. *The Ideas of Particle Physics*. 4th ed. Cambridge University Press, 2020. ISBN: 978-0-5115-8101-4.

[11] R. L. Jaffe. "Casimir Effect and the Quantum Vacuum". In: *Physical Review D* 72.2 (July 2005). ISSN: 1550-2368. DOI: 10.1103/physrevd.72.021301. URL: http://dx.doi.org/10.1103/PhysRevD.72.021301.

[12] Don N Page. "Hawking Radiation and Black Hole Thermodynamics". In: *New Journal of Physics* 7 (Sept. 2005), pp. 203–203. ISSN: 1367-2630. DOI: 10.1088/1367-2630/7/1/203. URL: http://dx.doi.org/10.1088/1367-2630/7/1/203.

[13] Howard Georgi. *Lie algebras in particle physics: from isospin to unified theories; Second edition*. CRC Press, 2018. ISBN: 978-0-4294-9921-0.

[14] P.C. Kendall and D.E. Bourne. *Vector Analysis and Cartesian Tensors*. 2nd ed. Routledge, 1977. ISBN: 0-748-75460-1.

[15] Léo Morel et al. "Determination of the Fine-Structure Constant with an Accuracy of 81 Parts per Trillion". In: *Nature* 588 (Dec. 2020), pp. 61–65. DOI: 10.1038/s41586-020-2964-7.

[16] The ATLAS Collaboration. "Observation of a New Particle in the Search for the Standard Model Higgs Boson with the ATLAS Detector at the LHC". In: *Physics Letters B* 716.1 (2012), pp. 1–29. ISSN: 0370-2693. DOI: 10.1016/j.physletb.2012.08.020. URL: http://dx.doi.org/10.1016/j.physletb.2012.08.020.

[17] "Observation of a New Boson at a Mass of 125 GeV with the CMS Experiment at the LHC". In: *Physics Letters B* 716.1 (2012), pp. 30–61. ISSN: 0370-2693. DOI: 10.1016/j.physletb.2012.08.021. URL: http://dx.doi.org/10.1016/j.physletb.2012.08.021.

[18] Brian Moser. "The Beauty and the Boost: a Higgs Boson Tale — Measurements of Higgs Boson Production at High Energy in Decays to Bottom Quarks and Their Interpretations with the ATLAS Experiment at the LHC". Presented 04 Mar 2022. PhD thesis. Amsterdam U., 2021. URL: https://cds.cern.ch/record/2803776.

[19] G. ZWEIG. "MEMORIES OF MURRAY AND THE QUARK MODEL". In: *International Journal of Modern Physics A* 25 (2010). ISSN: 1793-656X. DOI: 10.1142/s0217751x10050494. URL: http://dx.doi.org/10.1142/S0217751X10050494.

[20] David J. Gross and Frank Wilczek. "Ultraviolet Behaviour of Non-Abelian Gauge Theories". In: *Phys. Rev. Lett.* 30 (26 1973), pp. 1343–1346. DOI: 10.1103/PhysRevLett.30.1343. URL: https://link.aps.org/doi/10.1103/PhysRevLett.30.1343.

[21] Frank Wilczek. "Nobel Lecture: Asymptotic Freedom: From Paradox to Paradigm". In: *Reviews of Modern Physics* 77 (2005). ISSN: 1539-0756. DOI: 10.1103/revmodphys.77.857. URL: http://dx.doi.org/10.1103/RevModPhys.77.857.

[22] Rajan Gupta. *Introduction to Lattice QCD*. 1998. arXiv: hep-lat/9807028 [hep-lat]. URL: https://arxiv.org/abs/hep-lat/9807028.

[23] F. Bissey et al. "Gluon Flux-Tube Distribution and Linear Confinement in Baryons". In: *Physical Review D* 76.11 (2007). ISSN: 1550-2368. DOI: 10.1103/physrevd.76.114512. URL: http://dx.doi.org/10.1103/PhysRevD.76.114512.

[24] G. Rajasekaran. "Fermi and the Theory of Weak Interactions". In: *Resonance* 19.1 (2014). ISSN: 0973-712X. DOI: 10.1007/s12045-014-0005-2. URL: http://dx.doi.org/10.1007/s12045-014-0005-2.

[25] Sheldon L. Glashow. "Partial-Symmetries of Weak Interactions". In: *Nuclear Physics* 22.4 (1961), pp. 579–588. ISSN: 0029-5582. DOI: https://doi.org/10.1016/0029-5582(61)90469-2. URL: https://www.sciencedirect.com/science/article/pii/0029558261904692.

[26] Steven Weinberg. "A Model of Leptons". In: *Phys. Rev. Lett.* 19 (21 1967), pp. 1264–1266. DOI: 10.1103/PhysRevLett.19.1264. URL: https://link.aps.org/doi/10.1103/PhysRevLett.19.1264.

[27] C. S. Wu et al. "Experimental Test of Parity Conservation in Beta Decay". In: *Phys. Rev.* 105 (4 1957), pp. 1413–1415. DOI: 10.1103/PhysRev.105.1413. URL: https://link.aps.org/doi/10.1103/PhysRev.105.1413.

[28] Peter W. Higgs. "Broken Symmetries and the Masses of Gauge Bosons". In: *Phys. Rev. Lett.* 13 (16 1964), pp. 508–509. DOI: 10.1103/PhysRevLett.13.508. URL: https://link.aps.org/doi/10.1103/PhysRevLett.13.508.

[29] F. Englert and R. Brout. "Broken Symmetry and the Mass of Gauge Vector Mesons". In: *Phys. Rev. Lett.* 13 (9 1964), pp. 321–323. DOI: 10.1103/PhysRevLett.13.321. URL: https://link.aps.org/doi/10.1103/PhysRevLett.13.321.

[30] Stefan Gadatsch. "The Higgs Boson". Report number: CERN-THESIS-2015-113. PhD thesis. Amsterdam U., 2015.

[31] *Feynman Diagrams of Higgs Boson Production Modes*. https://en.wikipedia.org/wiki/Higgs_boson. Accessed: 2024-07-21.

[32] The ATLAS Collaboration. "A detailed Map of Higgs Boson Interactions by the ATLAS Experiment Ten Years After the Discovery". In: *Nature* 607.7917 (2022). ISSN: 1476-4687. DOI: 10.1038/s41586-022-04893-w. URL: http://dx.doi.org/10.1038/s41586-022-04893-w.

[33] L. A. Harland-Lang et al. "Parton Distributions in the LHC Era: MMHT 2014 PDFs". In: *The European Physical Journal C* 75.5 (2015). ISSN: 1434-6052. DOI: 10.1140/epjc/s10052-015-3397-6. URL: http://dx.doi.org/10.1140/epjc/s10052-015-3397-6.

[34] Heather Russell. *Illustration of Hadronic Jet Formation*. https://hrussell.web.cern.ch/hrussell/graphics.html. Accessed: 2024-07-28.

[35] *Standard Model Summary Plots June 2024*. Tech. rep. Report number: ATL-PHYS-PUB-2024-011. Geneva: CERN, 2024. URL: https://atlas.web.cern.ch/Atlas/GROUPS/PHYSICS/PUBNOTES/ATL-PHYS-PUB-2024-011/.

[36] Michael E. Peskin. *What is the Hierarchy Problem?* 2025. arXiv: 2505.00694 [hep-ph]. URL: https://arxiv.org/abs/2505.00694.

[37] P. Grangé et al. "Aspects of Fine-Tuning of the Higgs Mass Within Finite Field Theories". In: *Physical Review D* 88.12 (Dec. 2013). ISSN: 1550-2368. DOI: 10.1103/physrevd.88.125015. URL: http://dx.doi.org/10.1103/PhysRevD.88.125015.

[38] Brandon Carter. *Anthropic Principle in Cosmology*. 2006. arXiv: gr-qc/0606117 [gr-qc]. URL: https://arxiv.org/abs/gr-qc/0606117.

[39] Michael Dine and Alexander Kusenko. "Origin of the Matter-Antimatter Asymmetry". In: *Reviews of Modern Physics* 76.1 (Dec. 2003), pp. 1–30. ISSN: 1539-0756. DOI: 10.1103/revmodphys.76.1. URL: http://dx.doi.org/10.1103/RevModPhys.76.1.

[40] Graham White. *The Sakharov Conditions*. 2053-2563. IOP Publishing, 2022, 2–1 to 2–2. ISBN: 978-0-7503-3571-3. DOI: 10.1088/978-0-7503-3571-3ch2. URL: https://dx.doi.org/10.1088/978-0-7503-3571-3ch2.

[41] Frederick Reines and Clyde L. Cowan. "The Neutrino". In: *Nature* 178.4531 (1956), pp. 446–449.

[42] Paolo Walter Cattaneo. *Critical Review of the Results of the Homestake Solar Neutrino Experiment*. 2002. arXiv: astro-ph/0211534 [astro-ph]. URL: https://arxiv.org/abs/astro-ph/0211534.

[43] John N. Bahcall. "The Solar Neutrino Problem". In: *Solar Physics* 100.1 (1985), pp. 53–63.

[44] A. Bellerive et al. "The Sudbury Neutrino Observatory". In: *Nuclear Physics B* 908 (July 2016), pp. 30–51. ISSN: 0550-3213. DOI: 10.1016/j.nuclphysb.2016.04.035. URL: http://dx.doi.org/10.1016/j.nuclphysb.2016.04.035.

[45] Sidney van den Bergh. "The Early History of Dark Matter". In: *Publications of the Astronomical Society of the Pacific* 111.760 (1999), p. 657. DOI: 10.1086/316369. URL: https://dx.doi.org/10.1086/316369.

[46] Eugene Oks. "Brief Review of Recent Advances in Understanding Dark Matter and Dark Energy". In: *New Astronomy Reviews* 93 (2021), p. 101632. ISSN: 1387-6473. DOI: https://doi.org/10.1016/j.newar.2021.101632. URL: https://www.sciencedirect.com/science/article/pii/S1387647321000191.

[47] J. A. Peacock. *Cosmological Physics*. Cambridge University Press, 1998. ISBN: 978-0-5118-0453-3.

[48] ESA/Planck Collaboration. "Planck-2018 Results: VI. Cosmological Parameters". In: *Astronomy amp; Astrophysics* 641 (Sept. 2020), A6. ISSN: 1432-0746. DOI: 10.1051/0004-6361/201833910. URL: http://dx.doi.org/10.1051/0004-6361/201833910.

[49] Ruth Durrer. "The Cosmic Microwave Background: The History of its Experimental Investigation and its Significance for Cosmology". In: *Classical and Quantum Gravity* 32.12 (June 2015), p. 124007. ISSN: 1361-6382. DOI: 10.1088/0264-9381/32/12/124007. URL: http://dx.doi.org/10.1088/0264-9381/32/12/124007.

[50] Volker Springel et al. "Simulations of the Formation, Evolution and Clustering of Galaxies and Quasars". In: *Nature* 435.7042 (June 2005), pp. 629–636. ISSN: 1476-4687. DOI: 10.1038/nature03597. URL: http://dx.doi.org/10.1038/nature03597.

[51] Anne M. Green. "Primordial Black Holes as a Dark Matter Candidate — a Brief Overview". In: *Nuclear Physics B* 1003 (2024). Special Issue of Nobel Symposium 182 on Dark Matter, p. 116494. ISSN: 0550-3213. DOI: https://doi.org/10.1016/j.nuclphysb.2024.116494. URL: https://www.sciencedirect.com/science/article/pii/S0550321324000609.

[52] Malcolm A H MacCallum. "Milestones of General Relativity: Hubble's Law (1929) and the Expansion of the Universe". In: *Classical and Quantum Gravity* 32.12 (June 2015), p. 124002. ISSN: 1361-6382. DOI: 10.1088/0264-9381/32/12/124002. URL: http://dx.doi.org/10.1088/0264-9381/32/12/124002.

[53] Antonio Padilla. *Lectures on the Cosmological Constant Problem*. 2015. arXiv: 1502.05296 [hep-th]. URL: https://arxiv.org/abs/1502.05296.

[54] Roberto D. Peccei. "The Strong CP Problem and Axions". In: *Axions*. Springer Berlin Heidelberg, 2008, pp. 3–17. ISBN: 9783540735182. DOI: 10.1007/978-3-540-73518-2_1. URL: http://dx.doi.org/10.1007/978-3-540-73518-2_1.

[55] Masahiro Ibe, Alexander Kusenko, and Tsutomu T. Yanagida. "Why Three Generations?" In: *Physics Letters B* 758 (2016), pp. 365–369. ISSN: 0370-2693. DOI: https://doi.org/10.1016/j.physletb.2016.05.025. URL: https://www.sciencedirect.com/science/article/pii/S0370269316301721.

[56] M. V. Takook. "Conceptual and Technical Challenges of Quantum Gravity". In: *International Journal of Theoretical Physics* 59.8 (2020), pp. 2540–2556.

[57] Yu. A. Gol'fand and E. P. Likhtman. "Extension of the Algebra of Poincare Group Generators and Violation of P-Invariance". In: *Supergravities in Diverse Dimensions* (1989). DOI: 10.1142/9789814542340_0001. URL: https://www.worldscientific.com/doi/abs/10.1142/9789814542340_0001.

[58] A. Neveu and J.H. Schwarz. "Factorisable Dual Model of Pions". In: *Nuclear Physics B* 31.1 (1971), pp. 86–112. ISSN: 0550-3213. DOI: https://doi.org/10.1016/0550-3213(71)90448-2. URL: https://www.sciencedirect.com/science/article/pii/0550321371904482.

[59] P. Ramond. "Dual Theory for Free Fermions". In: *Phys. Rev. D* 3 (10 1971), pp. 2415–2418. DOI: 10.1103/PhysRevD.3.2415. URL: https://link.aps.org/doi/10.1103/PhysRevD.3.2415.

[60] P. Ramond. "Boson-Fermion Confusion: The String Path to Supersymmetry". In: *Nuclear Physics B - Proceedings Supplements* 101.1–3 (Aug. 2001), pp. 45–53. ISSN: 0920-5632. DOI: 10.1016/s0920-5632(01)01491-8. URL: http://dx.doi.org/10.1016/S0920-5632(01)01491-8.

[61] Can Kilic, Karoline Kopp, and Takemichi Okui. *LHC Implications of the WIMP Miracle and Grand Unification*. 2010. DOI: https://doi.org/10.1103/PhysRevD.83.015006. arXiv: 1008.2763 [hep-ph]. URL: https://arxiv.org/abs/1008.2763.

[62] Asli M Abdullahi and Pablo et. al. Barham Alzás. "The Present and Future Status of Heavy Neutral Leptons". In: *Journal of Physics G: Nuclear and Particle Physics* 50.2 (2023), p. 020501. DOI: 10.1088/1361-6471/ac98f9. URL: https://dx.doi.org/10.1088/1361-6471/ac98f9.

[63] S F King. "Neutrino Mass Models". In: *Reports on Progress in Physics* 67.2 (Dec. 2003), pp. 107–157. ISSN: 1361-6633. DOI: 10.1088/0034-4885/67/2/r01. URL: http://dx.doi.org/10.1088/0034-4885/67/2/R01.

[64] Vedran Brdar et al. "Type I Seesaw Mechanism as the Common Origin of Neutrino Mass, Baryon Asymmetry, and the Electroweak Scale". In: *Physical Review D* 100.7 (Oct. 2019). ISSN: 2470-0029. DOI: 10.1103/physrevd.100.075029. URL: http://dx.doi.org/10.1103/PhysRevD.100.075029.

[65] Latham Boyle, Kieran Finn, and Neil Turok. "CPT-Symmetric Universe". In: *Physical Review Letters* 121.25 (Dec. 2018). ISSN: 1079-7114. DOI: 10.1103/physrevlett.121.251301. URL: http://dx.doi.org/10.1103/PhysRevLett.121.251301.

[66] The ATLAS Collaboration. "Search for Heavy Majorana Neutrinos in $e\pm e\pm$ and $e\pm\pm$ Final States via WW Scattering in pp Collisions at $\sqrt{s} = 13$ TeV with the ATLAS Detector". In: *Physics Letters B* 856 (Sept. 2024), p. 138865. ISSN: 0370-2693. DOI: 10.1016/j.physletb.2024.138865. URL: http://dx.doi.org/10.1016/j.physletb.2024.138865.

[67] Maxim Perelstein. "Little Higgs Models and Their Phenomenology". In: *Progress in Particle and Nuclear Physics* 58.1 (Jan. 2007), pp. 247–291. ISSN: 0146-6410. DOI: 10.1016/j.ppnp.2006.04.001. URL: http://dx.doi.org/10.1016/j.ppnp.2006.04.001.

[68] Edward Farhi and Leonard Susskind. "Technicolour". In: *Physics Reports* 74.3 (1981), pp. 277–321. ISSN: 0370-1573. DOI: https://doi.org/10.1016/0370-1573(81)90173-3. URL: https://www.sciencedirect.com/science/article/pii/0370157381901733.

[69] Giacomo Cacciapaglia, Aldo Deandrea, and K. Sridhar. "Introduction to Composite Models". In: *The European Physical Journal Special Topics* 231.7 (2022), pp. 1223–1228.

[70] Abdel Pérez-Lorenzana. "An Introduction to Extra Dimensions". In: *Journal of Physics: Conference Series* 18 (Jan. 2005), pp. 224–269. ISSN: 1742-6596. DOI: 10.1088/1742-6596/18/1/006. URL: http://dx.doi.org/10.1088/1742-6596/18/1/006.

[71] John Ellis, D.V. Nanopoulos, and K. Tamvakis. "Grand Unification in Simple Supergravity". In: *Physics Letters B* 121.2 (1983), pp. 123–129. ISSN: 0370-2693. DOI: https://doi.org/10.1016/0370-2693(83)90900-0. URL: https://www.sciencedirect.com/science/article/pii/0370269383909000.

[72] The ATLAS Collaboration. *Prospects for Searches for Heavy Z' and W' Bosons in Fermionic Final States with the ATLAS Experiment at the HL-LHC*. Tech. rep. Report number: ATL-PHYS-PUB-2018-044. Geneva: CERN, 2018. URL: https://cds.cern.ch/record/2650549.

[73] Juliette Alimena et al. "Searching for Long-Lived Particles Beyond the Standard Model at the Large Hadron Collider". In: *Journal of Physics G: Nuclear and Particle Physics* 47.9 (2020), p. 090501. DOI: 10.1088/1361-6471/ab4574. URL: https://dx.doi.org/10.1088/1361-6471/ab4574.

[74] The ATLAS Collaboration. "Search for Displaced Leptons in $\sqrt{s} = 13$ TeV pp Collisions with the ATLAS Detector". In: *Phys. Rev. Lett.* 127 (5 2021), p. 051802. DOI: 10.1103/PhysRevLett.127.051802. URL: https://link.aps.org/doi/10.1103/PhysRevLett.127.051802.

[75] Robert Martin Eisberg and Robert Resnick. *Quantum Physics of Atoms, Molecules, Solids, Nuclei, and Particles*. second edition. Wiley, New York, 1985. ISBN: 0-471-87373-X.

[76] Torbjörn Sjöstrand. "The Pythia Event Generator: Past, Present and Future". In: *Computer Physics Communications* 246 (Jan. 2020), p. 106910. ISSN: 0010-4655. DOI: 10.1016/j.cpc.2019.106910. URL: http://dx.doi.org/10.1016/j.cpc.2019.106910.

[77] Gennaro Corcella et al. "HERWIG 6: An Event Generator for Hadron Emission Reactions with Interfering Gluons (Including Supersymmetric Processes)". In: *Journal of High Energy Physics* 2001.01 (Jan. 2001), pp. 010–010. ISSN: 1029-8479. DOI: 10.1088/1126-6708/2001/01/010. URL: http://dx.doi.org/10.1088/1126-6708/2001/01/010.

[78] Enrico Bothmann et al. *Event Generation with Sherpa 3*. 2024. arXiv: 2410.22148 [hep-ph]. URL: https://arxiv.org/abs/2410.22148.

[79] William H. Press et al. *Numerical Recipes 3rd Edition: The Art of Scientific Computing*. 3rd ed. Cambridge University Press, Sept. 2007. ISBN: 0521880688.

[80] Korinna Zapp. *Introduction to Monte Carlo Event Generators*. 2016. URL: https://www.uni-muenster.de/imperia/md/content/physik_tp/lectures/ws2016-2017/computational_methods_ii/mcgenerators_muenster_1610-i.pdf.

[81] Stefan Gieseke. *Introduction to Monte Carlo Event Generators*. 2019. URL: https://indico.cern.ch/event/796134/contributions/3560243/attachments/1908781/3153323/lect1and2.pdf.

[82] Fabio Maltoni and Tim Stelzer. "MadEvent: Automatic Event Generation with MadGraph". In: *Journal of High Energy Physics* 2003.02 (Feb. 2003), pp. 027–027. ISSN: 1029-8479. DOI: 10.1088/1126-6708/2003/02/027. URL: http://dx.doi.org/10.1088/1126-6708/2003/02/027.

[83] Tanju Gleisberg and Stefan Höche. "Comix, A New Matrix Element Generator". In: *Journal of High Energy Physics* 2008.12 (Dec. 2008), pp. 039–039. ISSN: 1029-8479. DOI: 10.1088/1126-6708/2008/12/039. URL: http://dx.doi.org/10.1088/1126-6708/2008/12/039.

[84] Stefano Frixione, Paolo Nason, and Carlo Oleari. "Matching NLO QCD Computations with Parton Shower Simulations: The POWHEG Method". In: *Journal of High Energy Physics* 2007.11 (Nov. 2007), pp. 070–070. ISSN: 1029-8479. DOI: 10.1088/1126-6708/2007/11/070. URL: http://dx.doi.org/10.1088/1126-6708/2007/11/070.

[85] Silvia Ferreres-Solé and Torbjörn Sjöstrand. "The Space–Time Structure of Hadronization in the Lund Model". In: *The European Physical Journal C* 78.11 (2018), p. 983. DOI: 10.1140/epjc/s10052-018-6459-8. URL: https://doi.org/10.1140/epjc/s10052-018-6459-8.

[86] Stefan Gieseke et al. *Phenomenological Constraints of the Building Blocks of the Cluster Hadronization Model*. 2025. arXiv: 2505.14542 [hep-ph]. URL: https://arxiv.org/abs/2505.14542.

[87] Stephen Mrenna. *Monte Carlo and Event Generators*. 2009. URL: https://www.hep.wisc.edu/cteq09/lectures/mrenna.pdf.

[88] R W Assmann, M Lamont, and S Myers. "A Brief History of the LEP Collider". In: *Nucl. Phys. B, Proc. Suppl.* 109 (2002). Report number: CERN-SL-2002-009-OP, pp. 17–31. DOI: 10.1016/S0920-5632(02)90005-8. URL: https://cds.cern.ch/record/549223.

[89] R Alemany-Fernandez et al. *Operation and Configuration of the LHC in Run 1*. Tech. rep. Report number: CERN-ACC-NOTE-2013-0041. 2013. URL: https://cds.cern.ch/record/1631030.

[90] Jorg Wenninger. *Operation and Configuration of the LHC in Run 2*. Tech. rep. Report number: CERN-ACC-NOTE-2019-0007. 2019. URL: https://cds.cern.ch/record/2668326.

[91] Stephane Fartoukh et al. *LHC Configuration and Operational Scenario for Run 3*. Tech. rep. Report number: CERN-ACC-2021-0007. 2021.

[92] Natalia Triantafyllou, Roderik Bruce, and Stefano Redaelli. "Analysis of the performance in the 2023 LHC Pb-Pb run". In: *JACoW IPAC* 2024 (2024), TUBD2. DOI: 10.18429/JACoW-IPAC2024-TUBD2. URL: https://cds.cern.ch/record/2912703.

[93] *Framework TDR for the LHCb Upgrade II: Opportunities in Flavour Physics, and Beyond, in the HL-LHC Era*. Tech. rep. Report number: CERN-LHCC-2021-012, LHCB-TDR-023. Geneva: CERN, 2021. URL: https://cds.cern.ch/record/2776420.

[94] Antonin Maire. *ALICE Technical Design Report (TDR) for LS2 Upgrades*. Report number: 2021. URL: https://cds.cern.ch/record/2771691.

[95] The ATLAS Collaboration. "The ATLAS Experiment at the CERN Large Hadron Collider". In: *JINST* 3 (2008), S08003. DOI: 10.1088/1748-0221/3/08/S08003.

[96] The CMS Collaboration. "The CMS Experiment at the CERN LHC". In: *Journal of Instrumentation* 3.08 (2008), S08004. DOI: 10.1088/1748-0221/3/08/S08004. URL: https://dx.doi.org/10.1088/1748-0221/3/08/S08004.

[97] Karlheinz Schindl. *The Injector Chain for the LHC*. Tech. rep. Report number: CERN-OPEN-99-052. 1999. URL: https://cds.cern.ch/record/397574.

[98] L et. al. Arnaudon. *Linac4 Technical Design Report*. Tech. rep. Report number: CERN-AB-2006-084, CARE-Note-2006-022-HIPPI. Geneva: CERN, 2006. URL: https://cds.cern.ch/record/1004186.

[99] Kurt Hübner et al. *Fifty years of the CERN Proton Synchrotron: Volume 2*. 2013. DOI: 10.5170/CERN-2013-005. URL: http://cds.cern.ch/record/1597087.

[100] Roberta Arnaldi. *Future Facilities: the CERN SPS*. 2025. arXiv: 2505.10286 [nucl-ex]. URL: https://arxiv.org/abs/2505.10286.

[101] H Padamsee. *Designing Superconducting Cavities for Accelerators*. Tech. rep. 2004. DOI: 10.5170/CERN-2004-008.233. URL: https://cds.cern.ch/record/808364.

[102] L Rossi. *Superconducting Cable and Magnets for the Large Hadron Collider*. Tech. rep. Report number: LHC-Project-Report-694, CERN-LHC-Project-Report-694. 2004. URL: https://cds.cern.ch/record/729430.

[103] R Ostojic. "The LHC Insertion Magnets". In: *IEEE Trans. Appl. Supercond.* 12.1 (2002), pp. 196–201. DOI: 10.1109/TASC.2002.1018382. URL: https://cds.cern.ch/record/544637.

[104] S Baird. *Accelerators for Pedestrians; rev. version*. Tech. rep. Report number: AB-Note-2007-014, CERN-AB-Note-2007-014, PS-OP-Note-95-17-Rev-2, CERN-PS-OP-Note-95-17-Rev-2. Geneva: CERN, 2007. URL: https://cds.cern.ch/record/1017689.

[105] Michael Barnes. "Kicker systems". In: *CERN Yellow Rep. School Proc.* 5 (2018), p. 229. DOI: 10.23730/CYRSP-2018-005.229. URL: https://cds.cern.ch/record/2674119.

[106] *Public ATLAS Luminosity Results for Run-2 of the LHC*. https://twiki.cern.ch/twiki/bin/view/AtlasPublic/LuminosityPublicResultsRun2. Accessed: 2024-12-16.

[107] Yi Yu. *The Latest Highlighted Measurements of the Standard Model with the ATLAS Experiment*. Tech. rep. Report number: ATL-PHYS-PROC-2024-021. Geneva: CERN, 2024. URL: https://cds.cern.ch/record/2892696.

[108] Mitsou, Vasiliki A. and On behalf of the ATLAS Collaboration. "Hunting New Physics with ATLAS". In: *EPJ Web Conf.* 182 (2018), p. 02089. DOI: 10.1051/epjconf/201818202089. URL: https://doi.org/10.1051/epjconf/201818202089.

[109] H H J ten Kate. "Superconducting Magnet System for the ATLAS Detector at CERN". In: *IEEE Trans. Appl. Supercond.* 9.2, pt.1 (1999), pp. 841–846. DOI: 10.1109/77.783428. URL: https://cds.cern.ch/record/409763.

[110] *ATLAS Magnet System: Technical Design Report, 1*. Technical design report. ATLAS. Geneva: CERN, 1997. DOI: 10.17181/CERN.9O5C.VDTM. URL: https://cds.cern.ch/record/338080.

[111] *ATLAS Inner Detector: Technical Design Report, 1*. Technical design report. ATLAS. Geneva: CERN, 1997. URL: https://cds.cern.ch/record/331063.

[112] Nicoletta Garelli. "Performance of the ATLAS Detector in Run-2". In: *EPJ Web Conf.* 164 (2017), p. 01021. DOI: 10.1051/epjconf/201716401021. URL: https://cds.cern.ch/record/2310929.

[113] Lee Sze. *Semiconductor Devices: Physics and Technology*. WILEY. ISBN: 978-0-470-53794-7.

[114] M et. al. Aleksa. "Construction, Assembly and Tests of the ATLAS Electromagnetic End-Cap Calorimeters". In: *JINST* 3 (2008), P06002. DOI: 10.1088/1748-0221/3/06/P06002. URL: https://cds.cern.ch/record/1114196.

[115] *ATLAS Liquid-Argon Calorimeter: Technical Design Report*. Technical design report. ATLAS. Geneva: CERN, 1996. DOI: 10.17181/CERN.FWRW.FOOQ. URL: https://cds.cern.ch/record/331061.

[116] The ATLAS Collaboration. "Operation and Performance of the ATLAS Tile Calorimeter in LHC Run 2". In: *The European Physical Journal C* 84.12 (Dec. 2024). ISSN: 1434-6052. DOI: 10.1140/epjc/s10052-024-13151-4. URL: http://dx.doi.org/10.1140/epjc/s10052-024-13151-4.

[117] S Aefsky. *Alignment of the Muon Spectrometer in ATLAS*. Tech. rep. Report number: ATL-MUON-PROC-2011-003. Geneva: CERN, 2012. DOI: 10.1016/j.phpro.2012.02.355. URL: https://cds.cern.ch/record/1380912.

[118] Y Arai et al. "ATLAS Muon Drift Tube Electronics". In: *JINST* 3 (2008). Report number: ATL-MUON-PUB-2008-008, ATL-COM-MUON-2008-010, P09001. DOI: 10.1088/1748-0221/3/09/P09001. URL: https://cds.cern.ch/record/1100199.

[119] P. Branchini et al. "Global Time Fit for Tracking in an Array of Drift Cells: The Drift Tubes of the ATLAS Experiment". In: *IEEE Transactions on Nuclear Science* 55.1 (2008), pp. 620–627. DOI: 10.1109/TNS.2007.914020.

[120] *ATLAS Muon Spectrometer: Technical Design Report*. Technical Design Report. ATLAS. Geneva: CERN, 1997. URL: https://cds.cern.ch/record/331068.

[121] A Salvucci. *Measurement of Muon Momentum Resolution of the ATLAS Detector*. Tech. rep. Report number: ATL-PHYS-PROC-2012-013. Geneva: CERN, 2012. URL: https://cds.cern.ch/record/1417538.

[122] The ATLAS collaboration. "Operation of the ATLAS trigger System in Run 2". In: *Journal of Instrumentation* 15.10 (Oct. 2020), P10004–P10004. ISSN: 1748-0221. DOI: 10.1088/1748-0221/15/10/p10004. URL: http://dx.doi.org/10.1088/1748-0221/15/10/P10004.

[123] The ATLAS Collaboration. "Performance of the ATLAS Trigger System in 2015". In: *The European Physical Journal C* 77.5 (May 2017). ISSN: 1434-6052. DOI: 10.1140/epjc/s10052-017-4852-3. URL: http://dx.doi.org/10.1140/epjc/s10052-017-4852-3.

[124] Guilherme Almeida et. al. Xiaocong Ai Corentin Allaire. *A Common Tracking Software Project*. 2021. URL: https://acts.readthedocs.io/en/latest/tracking.html.

[125] T Cornelissen et al. *Concepts, Design and Implementation of the ATLAS New Tracking (NEWT)*. Tech. rep. Report number: ATL-SOFT-PUB-2007-007, ATL-COM-SOFT-2007-002. Geneva: CERN, 2007. URL: https://cds.cern.ch/record/1020106.

[126] Yan Pei et al. "An Elementary Introduction to Kalman Filtering". In: *Commun. ACM* 62.11 (Oct. 2019), pp. 122–133. ISSN: 0001-0782. DOI: 10.1145/3363294. URL: https://doi.org/10.1145/3363294.

[127] The ATLAS Collaboration. *Performance of the Reconstruction of Large Impact Parameter Tracks in the ATLAS Inner Detector*. Tech. rep. Report number: ATL-PHYS-PUB-2017-014. Geneva: CERN, 2017. URL: https://cds.cern.ch/record/2275635.

[128] Guennadi Borissov et al. *ATLAS Strategy for Primary Vertex Reconstruction During Run-2 of the LHC*. Tech. rep. 7. Report number: ATL-SOFT-PROC-2015-016. Geneva: CERN, 2015. DOI: 10.1088/1742-6596/664/7/072041. URL: https://cds.cern.ch/record/2015220.

[129] The 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". In: *The European Physical Journal C* 79.8 (). ISSN: 1434-6052. DOI: 10.1140/epjc/s10052-019-7140-6. URL: http://dx.doi.org/10.1140/epjc/s10052-019-7140-6.

[130] The ATLAS Collaboration. "Muon Reconstruction and Identification Efficiency in ATLAS Using the Full Run 2 pp Collision Data-Set at $\sqrt{s} = 13$ TeV". In: *The European Physical Journal C* 81.7 (July 2021). ISSN: 1434-6052. DOI: 10.1140/epjc/s10052-021-09233-2. URL: http://dx.doi.org/10.1140/epjc/s10052-021-09233-2.

[131] Stan Bentvelsen and Irmtraud Meyer. "The Cambridge Jet Algorithm: Features and Applications". In: *The European Physical Journal C - Particles and Fields* 4.4 (1998), pp. 623–629.

[132] Shao-Feng Ge. "The Georgi Algorithms of Jet Clustering". In: *Journal of High Energy Physics* 2015.5 (May 2015). ISSN: 1029-8479. DOI: 10.1007/jhep05(2015)066. URL: http://dx.doi.org/10.1007/JHEP05(2015)066.

[133] J.M. Butterworth et al. "K_t-Jet: A C++ Implementation of the K_t Clustering Algorithm". In: *Computer Physics Communications* 153.1 (June 2003), pp. 85–96. ISSN: 0010-4655. DOI: 10.1016/s0010-4655(03)00156-5. URL: http://dx.doi.org/10.1016/S0010-4655(03)00156-5.

[134] Matteo Cacciari, Gavin P Salam, and Gregory Soyez. "The anti-K_t Jet Clustering Algorithm". In: *Journal of High Energy Physics* 2008.04 (Apr. 2008), pp. 063–063. ISSN: 1029-8479. DOI: 10.1088/1126-6708/2008/04/063. URL: http://dx.doi.org/10.1088/1126-6708/2008/04/063.

[135] Steven Schramm. *ATLAS Jet Reconstruction, Energy Scale Calibration, and Tagging of Lorentz-Boosted Objects*. Report number: ATL-PHYS-SLIDE-2017-781. 2017. URL: https://cds.cern.ch/record/2284807.

[136] The ATLAS Collaboration. *Jet Calibration and Systematic Uncertainties for Jets Reconstructed in the ATLAS Detector at $\sqrt{s} = 13$ TeV*. Tech. rep. Report number: ATL-PHYS-PUB-2015-015. Geneva: CERN, 2015. URL: https://cds.cern.ch/record/2037613.

[137] The ATLAS Collaboration. "Jet Energy Scale Measurements and Their Systematic Uncertainties in Proton-Proton collisions at $\sqrt{s} = 13$ TeV with the ATLAS Detector". In: *Physical Review D* 96.7 (Oct. 2017). ISSN: 2470-0029. DOI: 10.1103/physrevd.96.072002. URL: http://dx.doi.org/10.1103/PhysRevD.96.072002.

[138] Aliaksei Hrynevich. *ATLAS Jet and Missing Energy Reconstruction, Calibration and Performance in LHC Run-2*. Tech. rep. 06. Geneva: CERN, 2017. DOI: 10.1088/1748-0221/12/06/C06038. URL: https://cds.cern.ch/record/2263777.

[139] The ATLAS Collaboration. *In-Situ Measurements of the ATLAS Large-Radius Jet Response in $\sqrt{s} = 13$ TeV pp Collisions*. Tech. rep. Geneva: CERN, 2017. URL: https://cds.cern.ch/record/2275655.

[140] *Optimisation of the ATLAS b-Tagging Performance for the 2016 LHC Run*. Tech. rep. Report number: ATL-PHYS-PUB-2016-012. Geneva: CERN, 2016. URL: https://cds.cern.ch/record/2160731.

[141] Yulia Rodina. *ATLAS b-Tagging Performance During LHC Run 2 with the New Insertable B-Layer*. Tech. rep. Geneva: CERN, 2015. URL: https://cds.cern.ch/record/2112120.

[142] Kevin P. Murphy. *Machine Learning: A Probabilistic Perspective*. The MIT Press, 2012. ISBN: 0-262-01802-0.

[143] *Performance of the Reconstruction and Identification of Hadronic Tau Decays in ATLAS with 2011 Data*. Tech. rep. Report number: ATLAS-CONF-2012-142. Geneva: CERN, 2012. URL: https://cds.cern.ch/record/1485531.

[144] Blake Burghgrave. *Reconstruction of Hadronically Decaying Tau Leptons with ATLAS*. Tech. rep. Report number: ATL-PHYS-PROC-2016-188. Geneva: CERN, 2016. DOI: 10.22323/1.282.1118. URL: https://cds.cern.ch/record/2228946.

[145] The ATLAS Collaboration. "Performance of Jet Substructure Techniques for large-R Jets in Proton-Proton Collisions at $\sqrt{s} = 7$ TeV using the ATLAS Detector". In: *Journal of High Energy Physics* 2013.9 (Sept. 2013). ISSN: 1029-8479. DOI: 10.1007/jhep09(2013)076. URL: http://dx.doi.org/10.1007/JHEP09(2013)076.

[146] Matteo Cacciari, Gavin P Salam, and Gregory Soyez. "The Catchment Area of Jets". In: *Journal of High Energy Physics* 2008.04 (Apr. 2008), pp. 005–005. ISSN: 1029-8479. DOI: 10.1088/1126-6708/2008/04/005. URL: http://dx.doi.org/10.1088/1126-6708/2008/04/005.

[147] Matteo Cacciari and Gavin P. Salam. "Pileup Subtraction using Jet Areas". In: *Physics Letters B* 659.1–2 (Jan. 2008), pp. 119–126. ISSN: 0370-2693. DOI: 10.1016/j.physletb.2007.09.077. URL: http://dx.doi.org/10.1016/j.physletb.2007.09.077.

[148] *Identification of the Hadronic Decays of Tau Leptons in 2012 Data with the ATLAS Detector*. Tech. rep. Report number: ATLAS-CONF-2013-064. Geneva: CERN, 2013. URL: https://cds.cern.ch/record/1562839.

[149] Robin M. Schmidt. *Recurrent Neural Networks (RNNs): A gentle Introduction and Overview*. 2019. arXiv: 1912.05911 [cs.LG]. URL: https://arxiv.org/abs/1912.05911.

[150] Estifa'a Zaid. "Search for Long-Lived Particles Decaying into the Semi-Leptonic Di-Tau Final State with the ATLAS Detector at the LHC". PhD thesis. The University of Edinburgh, 2023.

[151] A Brennan. *Tau energy Calibration in the ATLAS Experiment*. Tech. rep. Report number: ATL-PHYS-PROC-2013-260. Geneva: CERN, 2013. URL: https://cds.cern.ch/record/1609636.

[152] *Reconstruction, Identification, and Calibration of Hadronically Decaying Tau Leptons with the ATLAS Detector for the LHC Run 3 and Reprocessed Run 2 Data*. Tech. rep. Geneva: CERN, 2022. URL: https://cds.cern.ch/record/2827111.

[153] *Identification of Hadronic Tau Lepton Decays Using Neural Networks in the ATLAS Experiment*. Tech. rep. Report number: ATL-PHYS-PUB-2019-033. Geneva: CERN, 2019. URL: https://cds.cern.ch/record/2688062.

[154] Geert-Jan Besjes. *Performance of the ATLAS Tau Trigger in Run 2*. Tech. rep. Report number: ATL-DAQ-PROC-2016-027. Geneva: CERN, 2016. DOI: 10.22323/1.282.0857. URL: https://cds.cern.ch/record/2231874.

[155] Jonathan Long. *The ATLAS Inner Detector Tracking Trigger at 13 TeV in LHC Run-2 and New Developments on Standard and Unconventional Tracking Signatures for the Upcoming LHC Run-3*. Tech. rep. Report number: ATL-DAQ-PROC-2022-005. Geneva: CERN, 2022. URL: https://cds.cern.ch/record/2813981.

[156] The ATLAS Collaboration. "The ATLAS Inner Detector Trigger Performance in pp Collisions at 13 TeV During LHC Run 2". In: *The European Physical Journal C* 82.3 (2022), p. 206.

[157] The ATLAS Collaboration. "Search for Displaced Leptons in $\sqrt{s} = 13$ TeV pp Collisions with the ATLAS Detector". In: *Phys. Rev. Lett.* 127 (5 2021), p. 051802. DOI: 10.1103/PhysRevLett.127.051802. URL: https://link.aps.org/doi/10.1103/PhysRevLett.127.051802.

[158] Heather Russell. *Illustration of Hadronic LLP Decay with the ATLAS Inner Detector*. https://hrussell.web.cern.ch/hrussell/graphics.html. Accessed: 2024-11-28.

[159] Konstantin Lehmann and Bernd Stelzer. "The Fake Factor Method and its Relation to the Matrix Method". In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 1054 (2023), p. 168376. ISSN: 0168-9002. DOI: https://doi.org/10.1016/j.nima.2023.168376. URL: https://www.sciencedirect.com/science/article/pii/S0168900223003662.

[160] The ATLAS Collaboration. "Search for Long-Lived Neutral Particles in pp Collisions at $\sqrt{s} = 13$ TeV that Decay into Displaced Hadronic Jets in the ATLAS Calorimeter". In: *The European Physical Journal C* 79.6 (2019). ISSN: 1434-6052. DOI: 10.1140/epjc/s10052-019-6962-6. URL: http://dx.doi.org/10.1140/epjc/s10052-019-6962-6.

[161] Glen Cowan. *Statistical Data Analysis*. Oxford University Press, Mar. 1998. ISBN: 9780198501565. DOI: 10.1093/oso/9780198501565.001.0001. URL: https://doi.org/10.1093/oso/9780198501565.001.0001.

[162] Olaf Behnke et al. *Data Analysis in High Energy Physics: A Practical Guide to Statistical Methods*. 1st. Wiley-VCH, 2013. ISBN: 3527410589.

[163] *Introduction to the ABCD Background Estimation Method*. https://indico.cern.ch/event/1122790/contributions/4713580/attachments/2381493/4270369/ABCDGuide_draft18Oct18.pdf. Accessed: 2024-11-28.

[164] S. Agostinelli et al. "Geant4—A Simulation Toolkit". In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 506.3 (2003), pp. 250–303. ISSN: 0168-9002. DOI: https://doi.org/10.1016/S0168-9002(03)01368-8. URL: https://www.sciencedirect.com/science/article/pii/S0168900203013688.

[165] C. Sbarra. "The LUCID-2 Detector". In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 936 (2019). Frontier Detectors for Frontier Physics: 14th Pisa Meeting on Advanced Detectors, pp. 152–153. ISSN: 0168-9002. DOI: https://doi.org/10.1016/j.nima.2018.07.058. URL: https://www.sciencedirect.com/science/article/pii/S0168900218308933.

[166] ATLAS Colloaboration. "Search for Displaced Leptons in $\sqrt{s} = 13$ TeV and $\sqrt{s} = 13.6$ TeV pp-Collisions with the ATLAS Detector". In: *Physical Review D* 112.1 (July 2025). ISSN: 2470-0029. DOI: 10.1103/w8hh-xf24. URL: http://dx.doi.org/10.1103/w8hh-xf24.

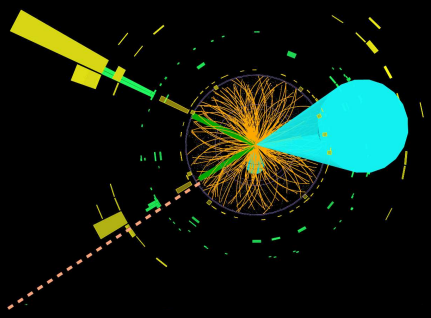

That’s all, folks!