

UiO : **University of Oslo**

Simen Hellesund

**Searching for New Physics in Resonant
and Non-Resonant Dilepton Final States
With the ATLAS Detector at the LHC**

Thesis submitted for the degree of Philosophiae Doctor

Department of Physics
Faculty of Mathematics and Natural Sciences



2021

© **Simen Hellesund, 2022**

*Series of dissertations submitted to the
Faculty of Mathematics and Natural Sciences, University of Oslo
No. 2486*

ISSN 1501-7710

All rights reserved. No part of this publication may be
reproduced or transmitted, in any form or by any means, without permission.

Cover: Hanne Baadsgaard Utigard.
Print production: Representralen, University of Oslo.

Abstract

The Standard Model (SM) of particle physics has proven to be a very successful framework for describing the interactions of elementary particles. So far, no significant deviations have been measured from its predictions. However, some open questions remain about nature, that the SM in its current form cannot answer. What is the nature of dark matter? Can a quantum theory of gravity be constructed? Do the SM interactions all originate from a common unified interaction? Numerous extensions to the SM have been proposed, aiming to solve these questions and others. Extensions to the SM may give different predictions about nature than the SM itself. Physicists look for evidence of such new physics by finding discrepancies in measurements from the SM predictions, for example in particle collision experiments.

This thesis presents two searches for new physics in dilepton final states, analysing 139 fb^{-1} of proton-proton collisions collected at $\sqrt{s} = 13 \text{ TeV}$ using the ATLAS detector at the Large Hadron Collider. The first search is for a resonant dielectron or dimuon signal of invariant mass between 250 GeV and 6,000 GeV. A functional form is used to model the SM background. No significant deviation from the background is observed. The upper limit on the cross-section of a model-independent resonance is set, and lower mass limits are set at 5.1 TeV, 4.5 TeV, and 4.8 TeV for the Z'_{SSM} , Z'_{ψ} , and Z'_{χ} benchmark heavy boson models, respectively.

The second search is for non-resonant excesses at high dilepton invariant mass. A novel method is used for estimating the SM background, in which a functional form is fitted to the data at low invariant mass, and extrapolated into a high invariant mass signal region where the search is carried out as a single bin counting experiment. No significant deviation from the background is observed. Upper limits are set on the model independent number of signal events. The results are also interpreted in terms of lower limits on the string scale parameter M_S of the Arkani-Hamed, Dimopoulos, and Dvali (ADD) model of large extra dimensions. The lower limit on M_S is set to 6.6 TeV in the dielectron channel, and 6.5 TeV in the dimuon channel, using the Giudice, Rattazzi, and Wells (GRW) convention for summing over Kaluza-Klein (KK) excitations. These limits are the first ADD model limits set by the ATLAS Collaboration in dilepton final states using the full Run 2 data set.

Acknowledgements

I am writing these words standing on the shoulders of giants. First and foremost amongst which are my two supervisors, Farid Ould-Saada and Magnar Kopangen Bugge. To Farid, for giving me the opportunity to work on such an interesting and challenging project, but also for allowing me freedom to pursue my own curiosity. And to Magnar: your knowledge of ATLAS physics and statistical concepts are rivalled only by your willingness to share it. From the bottom of my heart, I thank you both. I have enjoyed our collaboration immensely.

A special thank you goes out to Alexander Lincoln Read. I still would not know what a spurious signal is without our discussions. And also to James Richard Catmore, for your guidance during my ATLAS authorship qualification task.

I want to thank all the members of the ATLAS exotic dilepton working group for four years of great collaboration. In particular, I am indebted to Etienne Dreyer, Peter Johannes Falke, Aaron White, Deshan Abhayasinghe, Daniel Hayden, and Tracey Berry.

To all my friends and colleagues: Andreas, Eli, Even, Inga, Jeriek, Knut, Steffen, and everyone at the high energy particle physics section at the department of physics. Thank you for all the lunch break quizzes, the ski trips, the schools, the conferences, and the after-work pub visits. This work may have been possible without you, but it would not have been enjoyable without you.

To my parents, Wenche and Kjell Rune, and my brother Vegard. Thank you for your support and encouragement. Our weekly Skype calls have truly helped me get through some tough times while writing this thesis.

And finally, to Kristin. For your love, support, and patience. I love you.

Simen Hellesund
Oslo, November 2021

Contents

Abstract	iii
Acknowledgements	v
Introduction	1
1. The Standard Model and Beyond	3
1.1. The Standard Model of Particle Physics	3
1.1.1. The Particle Zoo	3
1.1.2. Standard Model Mathematical Formalism	6
1.1.3. The Interactions of the Standard Model	9
1.1.4. Electroweak Symmetry Breaking	18
1.2. Physics Beyond the Standard Model	22
1.2.1. Heavy Neutral Bosons	26
1.2.2. Large Extra Dimensions	27
1.2.3. Contact Interactions	29
2. Experimental Setup	31
2.1. Introduction to Particle Collider Experiments	31
2.2. CERN and the Large Hadron Collider	33
2.3. The ATLAS Detector	34
2.3.1. The ATLAS Detector Coordinate System	36
2.3.2. ATLAS Detector Layout	36
2.3.3. Triggers and Data Acquisition	43
2.3.4. Event Reconstruction and Particle Identification	46
3. Dilepton Standard Model Backgrounds	51
3.1. Drell-Yan	53
3.2. Top Quark	54
3.3. Diboson	55
3.4. Fake Electrons	56
3.5. Monte Carlo Background Samples	58
3.6. Transfer Functions	58
4. Event Selection for Dilepton Analyses	61
4.1. Data Set	61
4.2. Object Selection	61
4.2.1. Electrons	62
4.2.2. Muons	62
4.3. Calibration and Corrections	63
4.4. Object Overlap Removal	64
4.5. Event Cleaning	64

4.6. Triggers	65
4.7. Event Selection	65
4.8. Data and Monte Carlo Comparison Plots	65
4.9. Highest Dilepton Invariant Mass Event Displays	66
5. Resonant Dilepton Analysis	73
5.1. Modelling Exotic Dilepton Resonances	73
5.1.1. Signal Re-Weighting	76
5.2. Determining Acceptance Times Efficiency	77
5.3. Data-Driven Background Estimation	78
5.3.1. Spurious Signal	79
5.3.2. Standard Model Background Estimation	81
5.4. Parameterising the Spurious Signal Uncertainty	82
5.5. Systematic Uncertainties	83
5.6. Statistical Analysis	85
5.6.1. Discovery	85
5.6.2. Exclusion	88
5.7. Results	89
6. Non-Resonant Dilepton Analysis	95
6.1. Background Estimation	95
6.2. Processing ADD Theory Signal Samples	97
6.2.1. Kaluza-Klein Formalism Conversion	99
6.3. Systematic Uncertainties	100
6.3.1. Background Uncertainties	100
6.3.2. Signal Uncertainties	101
6.4. Statistical Analysis	102
6.4.1. Discovery	104
6.4.2. Exclusion	105
6.5. Results	107
6.6. Setting Limits on the String Scale Parameter	110
6.6.1. Reinterpreting Model-Independent Limits	111
6.6.2. Direct M_S Limits	113
6.7. Signal Region Re-Optimisation	114
7. Conclusions and Outlook	119
Appendices	121
A. Testing the Sliding Window Fit Method	123
A.1. The Sliding Window Fit Method	123
A.2. Optimising the Sliding Window Fit Method	123
A.2.1. Function Choice Optimisation for Sliding Window Fits	125
A.2.2. Window Shape Optimisation for Sliding Window Fits	127
A.3. Comparing The Global Fit and Sliding Window Fit Methods	136
A.4. Conclusion	139
B. Auxiliary ADD Model Material	141
B.1. Generating ADD Model Signal Samples	141

B.2. Kaluza-Klein Graviton Angular Distribution	147
C. ATLAS Authorship Qualification Report	151
C.1. Introduction	151
C.2. File Transfer Anomalies	152
C.3. Correlating Anomalies and Shifter Tickets	152
C.4. Seasonal Trend Decomposition	154
C.5. Machine Learning Approach	155
C.5.1. A Brief Introduction to Machine Learning	156
C.5.2. Boosted Decision Trees	158
C.5.3. Collecting File Transfer Metadata	160
C.5.4. Building the Classifier and Analysing Hyperparameters	162
C.5.5. Results and Performance	165
C.6. Conclusion	167
D. Physics Education Projects	169
D.1. Introduction	169
D.2. Measuring the Acceleration of Gravity	169
D.3. Measuring the Speed of Sound in Air	170
D.4. Measuring the Strength of Earth’s Magnetic Field	170
References	187
Acronyms and Abbreviations	197

Introduction

The Standard Model (SM) of particle physics has proven to be an extremely accurate framework for classifying the elementary particles of the universe and their interactions. Yet, some phenomena, such as the nature of dark matter or the non-zero mass of neutrinos, are not currently explained in the SM. In Chapter 1 of this thesis, I briefly outline the mathematical framework of the SM, as well as some of its shortcomings. Chapter 1 also introduces some possible extensions to the SM that could possibly rectify some of these issues. Emphasis is placed on theories predicting enhanced production of dilepton (dielectrons or dimuons of opposite electric charge) final states in particle collisions, which is the main focus of this thesis. One such class of models are grand unification theory (GUT) models based on the E_6 symmetry group, predicting the existence of at least one heavy neutral boson, Z' , at low energy scales.

The Large Hadron Collider (LHC), located at the European Organization for Nuclear Research (CERN) in Switzerland, delivered 139 fb^{-1} of proton-proton collision data, at a centre of mass (COM) collision energy of $\sqrt{s} = 13 \text{ TeV}$, to the ATLAS (A Toroidal LHC ApparatuS) experiment during its Run 2 collision campaign. In this thesis, I present two analyses of the Run 2 data set, searching for new physics in dilepton final states. The so-called dilepton channel has a long history of discovering new physics in particle collider experiments, with the discovery of the J/Ψ meson in 1974 [1, 2], the Υ meson in 1977 [3], and the Z boson in 1983 [4, 5]. Therefore, the search for new physics in the dilepton channel is considered one of the "flagship" analyses of the ATLAS Collaboration. The LHC and the ATLAS detector are described in Chapter 2, with an emphasis on how electrons and muons are measured and reconstructed in ATLAS.

The SM processes that form the background to the new physics searched for in the analysis, as well as the tools and methods we use to describe them, are described in Chapter 3. We generate Monte Carlo (MC) samples for all relevant SM dilepton backgrounds. However, in our two analyses, the SM background is modelled using parametric fits to the data. The MC background samples are used for validation of the data-driven background modelling, and for calculating systematic uncertainties in the analyses. The object definitions and event selection criteria for constructing the dilepton data set used in the analyses are given in Chapter 4.

Chapter 5 presents a search for resonant new physics phenomena in dilepton final states in the full Run 2 dilepton dataset. The analysis uses a parametric fit to the dilepton invariant mass distribution to describe the SM background in the search, and a generic parametric signal shape, based on a non-relativistic Breit-Wigner distribution, to describe the new physics signal. We consider pole mass hypotheses of the new signal between 250 GeV and 6 TeV, and widths of the signal between 0% and 10% of the resonance pole mass. Upper limits are set on the cross-section of the generic dilepton signal, as well as for three benchmark Z' models. A sliding window fit (SWiFt) method is considered for the background estimation in the analysis. More information about the SWiFt method, and its comparison to the global fit method used in the analysis can be found in Appendix A. The results from the search are published in Ref. [6].

Chapter 6 details a search for new phenomena manifesting as broad, non-resonant, deviations at high dilepton invariant masses. This analysis is a reinterpretation of the search for contact interaction (CI) processes presented in Ref. [7] into lower limits on the string scale parameter of the Arkani-Hamed, Dimopoulos, and Dvali (ADD) theory of large extra dimensions [8]. A novel method for estimating the SM background is used in the analysis, in which a parametric background function is fitted to the data in a low invariant mass control region, and then extrapolated into the high invariant mass signal regions. String scale limits are produced in the Giudice, Rattazzi, and Wells (GRW), Hewett, and Han, Lykken, and Zhang (HLZ) conventions for summing Kaluza-Klein (KK) modes. The results of the analysis are presented in Ref. [9], a so-called ATLAS public note.

During my time as a PhD student at the University of Oslo, I have been working in the ATLAS exotic dilepton analysis group. I joined the group in the closing stages of a search for new physics in 36 fb^{-1} of Run 2 ATLAS data [10]. I contributed to this analysis by calculating p -values for the discovery of a new physics signal. In 2017, the dilepton group started working on an analysis of the full Run 2 ATLAS dataset, searching for resonant signals above the SM prediction in dilepton final states. The results from this search are presented in Chapter 5. I had an active role in the analysis team throughout this period. The emphasis of my work in the analysis was the development and testing of the SWiFt method for performing background fits, as well as the production of data samples and figures for the publication. I also developed a system based on machine learning for tagging events originating from top quark decays. The idea being that this system could reduce the top quark contribution to the overall SM background. However, this system was not used in the final analysis. In 2019, after publishing our results from the resonant analysis, focus shifted in the working group to exotic non-resonant dilepton signatures. In the resulting analysis, presented in Chapter 6, I was responsible for the interpretation of the search in the context of large extra dimensions in the framework of an ADD model. Unfortunately, due to time constraints, my results were not published with the rest of the analysis, in Ref. [7], but rather as a stand-alone ATLAS public note [9].

During my time at the University of Oslo, I have also worked on some projects outside of particle physics searches. To qualify for an ATLAS authorship, I made an anomaly detection system, using machine learning methods, for file transfers on the Worldwide LHC Computing Grid (Grid). A report from this project is included in Appendix C. While working as a teaching assistant in an experimental physics course at the university, I was inspired to write several articles demonstrating how a modern smartphone can be used as a physics laboratory. These articles are included in Appendix D.

1. The Standard Model and Beyond

1.1. The Standard Model of Particle Physics

The ancient Greek philosopher Democritus is credited with the invention of the atomic theory; the idea that all matter is made up of some indivisible constituents of nature. Democritus named these fundamental quantities *atoms*.¹ Later it has become clear that atoms are not elementary building blocks of nature, but are themselves composite objects.

Today, we talk of elementary particles as being the fundamental constituents of nature. Particle physics is the study of these elementary particles. The Standard Model (SM) is the mathematical framework used to describe the elementary particles and their interactions. Or rather some of their interactions, as a quantum theory of gravity has yet to be described. This section presents a cursory review of the SM elementary particle content and mathematical framework.

1.1.1. The Particle Zoo

Spin is an intrinsic property of elementary particles, somewhat comparable to angular momentum. However, the elementary particles of the SM are point particles, and as such cannot rotate around their own axes. Therefore, this macroscopic analogy breaks down, and spin must be considered a purely quantum mechanical property of particles. Spin is quantised, and may take half-integer or integer values, in units of the reduced Planck constant \hbar . The particles of the SM are divided into two classes based on their spin: *fermions*, with half-integer spin, and integer spin *bosons*.

Fermions are matter particles, making up all the visible matter in the universe. The fermions come in 12 flavours, split into two categories based on their properties: six leptons and six quarks.

Three of the leptons carry electric charge: the electron e , the muon μ , and the tau lepton τ . Each charged lepton has an electrically neutral companion particle called a neutrino. These are the electron neutrino ν_e , the muon neutrino ν_μ , and the tau neutrino ν_τ . Each charged and uncharged lepton pair is said to belong to a family or generation, ordered by increasing mass of the charged leptons. This grouping is shown in Table 1.1.

Neutrinos are very feebly interacting particles, and are usually assumed to be massless in the SM. However, neutrino oscillation experiments have shown that neutrinos do in fact have a small, but non-zero, mass. Neutrino oscillation is a phenomenon in which neutrinos change flavour after being produced as a particular flavour, and is a consequence of the fact that the interaction quantum states of the neutrinos are different from their propagating mass quantum states. The masses of the three neutrino flavours, ν_e , ν_μ , and ν_τ , are therefore determined by linear combinations of three neutrino mass states ν_1 , ν_2 , and ν_3 . Direct measurements of the neutrino masses are difficult. Current upper limits for the individual neutrino masses are shown in Table 1.1. In addition to the laboratory

¹From the Greek *atomos*, meaning "indivisible".

measurements of neutrino masses, cosmological measurements put the upper limit on the sum of the three neutrino masses to be $m_\nu^{\text{tot}} < 0.26$ eV [11].

The six quarks are arranged in the same way as the leptons, in three generations ranked by increasing mass. Each quark generation consists of an up-type and a down-type quark. The up-type quarks are, in order of increasing mass: the *up* quark, the *charm* quark, and the *top* quark. The down-type quarks are: the *down* quark, the *strange* quark, and the *bottom* quark. The quarks all have fractional electric charge in units of the elementary charge: up-type quarks have charge $2/3e$, while the down-type quarks have charge $-1/3e$.

Quarks form bound states, called hadrons. The proton, for example, is made up of two up quarks and a down quark. Similarly, the neutron is made up of two down quarks and one up quark. The fractional electric charges of the quarks add up to give the proton charge e , and make the neutron electrically neutral. In fact, all hadrons have integer electric charge, despite the fractional charge of the quarks themselves. Protons and neutrons are collectively known as nucleons. Different numbers and configurations of protons and neutrons are all that is needed to make the nuclei of all atoms. Protons and neutrons are just two examples of so-called *baryons*, bound states of three quarks, although many more exist. Two quarks, more precisely a quark-anti-quark pair, may also go together to form mesons.

Fermions obey the Pauli exclusion principle, which states that two identical fermions may not occupy the same quantum state. With the discoveries of the first baryons consisting of three same-flavour quarks, such as the 1964 discovery of the Ω^- baryon, which consists of three strange quarks [12], it became clear that a new quark quantum number was needed. Two same-flavour spin-1/2 fermions may occupy the same state by different ordering of their spin projections ($+1/2$ and $-1/2$), but it is impossible to introduce a third fermion without violating the Pauli exclusion principle, unless one introduces a new quantum number. In addition to the quantum numbers electric charge and spin, quarks also carry so-called *colour charge*. A colour charged elementary particle interacts through the strong force. The strong force is what binds quarks together to form hadrons. The three colour charges are called red, green and blue. Colour charge is so named for its analogy to the primary colours of light. Like a mixture of red, green, and blue light mixes to make white light, the colour charges red, green and blue add to form a colour neutral state. All composite particles are "white", or colour singlets. The three valence quarks² of a baryon will therefore always be charged blue, green and red (or anti-blue, anti-green, and anti-red for anti-quarks). Similarly, mesons consist of a two-quark colour singlet system. This effect, that no colour charged particles is observed in nature, is called *confinement*, as quarks and gluons seem to be confined to the interior of hadrons.

In particle physics, a force acting is taken to mean the exchange of a force mediating boson between two particles charged under said interaction. Bosons are the second main class of elementary particles. Bosons have integer spin, and unlike fermions they do not obey the Pauli exclusion principle. Each fundamental force of the SM has one or more associated mediator particles. These force mediating particles are bosons. The massless photon is responsible for communicating the electromagnetic force between electrically charged particles. The strong force is mediated by eight gluons, which are also massless. Unlike the electrically neutral photon, the gluons themselves carry colour charge. There-

²The image presented thus far, that a baryon consists of only three quarks is somewhat simplistic. The three quarks that determine the quantum numbers of the baryon are called *valence quarks*, but the baryon also consists of a so-called *sea* of virtual quarks, anti-quarks, and gluons. These virtual particles may carry a significant fraction of the total momentum of the baryon.

Table 1.1.: The fermions of the Standard Model.

Class	Charge [e]	Generation		
		I	II	III
Quarks	2/3	$m = 2.16$ MeV u <i>Up Quark</i>	$m = 1.27$ GeV c <i>Charm Quark</i>	$m = 172.9$ GeV t <i>Top Quark</i>
	-1/3	$m = 4.67$ MeV d <i>Down Quark</i>	$m = 93$ MeV s <i>Strange Quark</i>	$m = 4.18$ GeV b <i>Bottom Quark</i>
Leptons	0	$m < 1.1$ eV ν_e <i>Electron Neutrino</i>	$m < 0.19$ MeV ν_μ <i>Muon Neutrino</i>	$m < 18.2$ MeV ν_τ <i>Tau Neutrino</i>
	-1	$m = 511$ keV e <i>Electron</i>	$m = 105.6$ MeV μ <i>Muon</i>	$m = 1,777$ GeV τ <i>Tau</i>

fore, gluons self-interact. This property is what leads to the property known as *colour confinement*, which is covered in more detail in Section 1.1.3.

Three massive bosons carry the weak force. The neutral Z^0 carry neutral current weak interactions, while the W^+ and W^- bosons carry charged current weak interactions. Weak interactions, those involving W bosons specifically, are unique in that they are the only interactions capable of changing fermion flavour. This property is necessary, for example, in beta decay, in which a down quark in a neutron is transformed into an up quark through the exchange of a W boson, turning the neutron into a proton and emitting an electron and an anti-electron neutrino. Unlike the photon and gluon, the weak bosons are massive particles. The Z^0 has a mass of 91.2 GeV and the W bosons each have a mass of 80.4 GeV [11]. Historically, this high mass made the weak bosons hard to produce in a laboratory. As a result, the Z and W bosons were not directly discovered until 1983 [13]. The force-carrying bosons of the SM are listed in Table 1.2.

Gravity is not included in the SM, but the hypothesised mediator boson of the gravitational interaction is known as the *graviton*. Graviton excitations are covered in more detail in Section 1.2.2.

Mesons are also bosons. The strong nuclear force, responsible for binding nucleons in nuclei, can be expressed as an effective theory of meson exchange.

Each fundamental force has a range. This range of interaction is determined by the Compton wavelength of the boson mediating it [14]. The Compton wavelength of a particle is the wavelength of a photon if this photon had the same energy as the rest mass of the particle in question. This means that heavier force mediators lead to shorter interaction lengths. The forces communicated by massless bosons have infinite range. This is not the case for the strong force. Although gluons are massless, the range of the strong interaction is effectively limited to 10^{-15} m by confinement. The large mass of the Z and W bosons

Table 1.2.: The force mediating bosons of the Standard Model.

<i>Force</i>	<i>Boson</i>	<i>Charge</i> [e]	<i>Mass</i> [GeV]	<i>Colour Charge</i>	<i>Weak Isospin</i>
Electromagnetism	γ <i>Photon</i>	0	0	0	0
Weak Force	W^{\pm} <i>W Bosons</i>	± 1	80.4	0	$\mp \frac{1}{2}$
	Z^0 <i>Z Boson</i>	0	91.2	0	0
Strong Force	g <i>Gluon</i>	0	0	r, g, b	0

give the weak interaction an effective range of the order of 10^{-18} m or less. At low energies, well below the mass of the weak bosons, the weak interaction can be considered a zero distance *contact interaction*. In fact, in the first proposal for a theoretical explanation of beta decay by Fermi in 1933, what came to be known as the weak interaction was assumed to be a four fermion contact interaction [15].

The spin-0 Higgs boson is the latest addition to the SM particle menagerie. It is named after Peter Higgs, one of the scientists behind the theory of its existence. The Higgs boson is not a force mediator like the other bosons of the SM. Instead, it is needed in the SM mathematical framework to spontaneously break the electroweak symmetry. This mechanism, which is responsible for giving mass to the fermions and the weak bosons, is explained in more detail in Section 1.1.4. The so-called Brout-Englert-Higgs (BEH) mechanism was developed in the 1960s. The discovery of the weak bosons at masses predicted by spontaneous symmetry breaking was a strong indicator for the existence of the so-called Higgs field, but it was not until 2012 that a 125 GeV resonance consistent with the Higgs boson was discovered [16, 17].

Most elementary particles have anti-particle partners. Anti-particles have the same mass as their ordinary matter counterparts, but with opposite-sign charges. Truly neutral particles are particles with no anti-particle counterparts. Rather, they are regarded as their own anti-particles. The photon, the Z^0 boson, and the Higgs boson are truly neutral particles. The postulation of the positron, the antimatter counterpart of the electron, is attributed to Paul Dirac, as an attempt to explain the issue of negative energy solutions to his famous Dirac equation.

1.1.2. Standard Model Mathematical Formalism

The elementary particles of the Standard Model (SM) are represented by *quantum fields*. Quantum field theory (QFT) can be considered a successful union of Einstein's theory of special relativity and quantum mechanics. Einstein demonstrated through his famous equation $E = mc^2$ that energy and mass are interchangeable quantities. An important consequence of this is that matter can be created from energy. Quantum mechanics on its own is ill-equipped to handle such systems with a variable number of particles. The infinite degrees of freedom in a quantum field does not have such limitations. Even at energies too

low for particle anti-particle pair creation, intermediate multi-particle states are allowed and necessary in higher-order perturbation theory of elementary particle interactions [18].

Lagrangian mechanics is the language of QFT. The core component of Lagrangian mechanics is the *action* S , which is the time integral of the *Lagrangian* L . The Lagrangian of a system is defined as $L = T - V$, where T and V are the system's kinetic and potential energy, respectively. In a local theory, such as QFT, rather than using L , it is often conducive to use the *Lagrangian density* \mathcal{L} . In high energy physics, \mathcal{L} is often referred to as simply the *Lagrangian*. This text will also adhere to this convention.

For one or more fields $\phi_i(x)$, \mathcal{L} is dependent on the fields themselves and their derivatives $\partial_\mu\phi_i(x)$. When using \mathcal{L} , S becomes the integral

$$S = \int_{\Omega} d^4x \mathcal{L}(\phi, \partial_\mu\phi), \quad (1.1)$$

over the space-time volume Ω . From this point we suppress the indices and dependencies of ϕ to aid readability.

Hamilton's principle of least action states that the trajectory of a mechanical system is such that the action is stationary. For this trajectory, an infinitesimal variation of the action will vanish, $\delta S = 0$. By differentiation of the Lagrangian, this becomes

$$\delta S = \int_{\Omega} d^4x \left\{ \frac{\partial \mathcal{L}}{\partial \phi} \delta \phi + \frac{\partial \mathcal{L}}{\partial (\partial_\mu \phi)} \delta (\partial_\mu \phi) \right\} = 0. \quad (1.2)$$

Using the product rule of differentiation, we see that

$$\partial_\mu \left(\frac{\partial \mathcal{L}}{\partial (\partial_\mu \phi)} \delta \phi \right) = \partial_\mu \left(\frac{\partial \mathcal{L}}{\partial (\partial_\mu \phi)} \right) \delta \phi + \frac{\partial \mathcal{L}}{\partial (\partial_\mu \phi)} \delta (\partial_\mu \phi). \quad (1.3)$$

In the last term we use the fact that $\partial_\mu(\delta\phi) = \delta(\partial_\mu\phi)$. Inserting Equation (1.3) into Equation (1.2) yields

$$\delta S = \int_{\Omega} d^4x \left\{ \frac{\partial \mathcal{L}}{\partial \phi} \delta \phi - \partial_\mu \left(\frac{\partial \mathcal{L}}{\partial (\partial_\mu \phi)} \right) \delta \phi + \partial_\mu \left(\frac{\partial \mathcal{L}}{\partial (\partial_\mu \phi)} \delta \phi \right) \right\} = 0. \quad (1.4)$$

Using Gauss' divergence theorem the final term of this integral can be written as a surface integral over the boundary of Ω , and shown to disappear. We see that to satisfy the Hamiltonian principle, the following equation must hold:

$$\frac{\partial \mathcal{L}}{\partial \phi} - \partial_\mu \left(\frac{\partial \mathcal{L}}{\partial (\partial_\mu \phi)} \right) = 0. \quad (1.5)$$

This is the *Euler-Lagrange equation* [14]. For a system described by a Lagrangian \mathcal{L} , the Euler-Lagrange equation defines the system's *equations of motion*.

As an example, take the Lagrangian of a spin-0 scalar field ϕ :

$$\mathcal{L} = \frac{1}{2} (\partial_\mu \phi \partial^\mu \phi - m^2 \phi^2). \quad (1.6)$$

It should be noted that two or more Lagrangians may lead to the same equations of motion. The choice of this particular Lagrangian is motivated by the classical Lagrangian of a harmonic oscillator. Applying the Euler-Lagrange equation to this Lagrangian, the first term becomes

$$\frac{\partial \mathcal{L}}{\partial \phi} = -m^2 \phi. \quad (1.7)$$

Differentiating \mathcal{L} with respect to $\partial_\mu\phi$ gives

$$\begin{aligned}
 \frac{\partial\mathcal{L}}{\partial(\partial_\mu\phi)} &= \frac{\partial}{\partial(\partial_\mu\phi)} \frac{1}{2} (\partial_\nu\phi\partial^\nu\phi) = \frac{1}{2} \frac{\partial}{\partial(\partial_\mu\phi)} (\partial_\nu\phi g^{\nu\alpha}\partial_\alpha\phi) \\
 &= \frac{1}{2} g^{\nu\alpha} \left\{ \left(\frac{\partial}{\partial(\partial_\mu\phi)} \partial_\nu\phi \right) \partial_\alpha\phi + \partial_\nu\phi \left(\frac{\partial}{\partial(\partial_\mu\phi)} \partial_\alpha\phi \right) \right\} \\
 &= \frac{1}{2} g^{\nu\alpha} (\delta_\nu^\mu \partial_\alpha\phi + \partial_\nu\phi \delta_\alpha^\mu) = \frac{1}{2} (\partial^\mu\phi + \partial^\mu\phi) \\
 &= \partial^\mu\phi.
 \end{aligned} \tag{1.8}$$

Here we introduce the dummy variables ν and α , and use the metric tensor $g^{\nu\alpha}$ to lower the index of ∂^ν . The second term of the Euler-Lagrange equation now becomes $-\partial_\mu\partial^\mu\phi$. Putting this together with the result from Equation (1.7) yields the *Klein-Gordon equation*

$$(\partial_\mu\partial^\mu + m^2)\phi = 0. \tag{1.9}$$

Similarly, we can start with the Lagrangian of a spin-1/2 fermionic field ψ ,

$$\mathcal{L} = \bar{\psi} (i\gamma^\mu\partial_\mu - m) \psi, \tag{1.10}$$

where $\bar{\psi}$ is the *adjoint field*, defined as $\bar{\psi} = \psi^\dagger\gamma^0$. Applying the Euler-Lagrange equation yields the Dirac equation

$$(i\gamma^\mu\partial_\mu - m)\psi = 0. \tag{1.11}$$

The analogue of momentum for a field ϕ is known as the *conjugate momentum field* π , and is defined as

$$\pi = \frac{\partial\mathcal{L}}{\partial\dot{\phi}}. \tag{1.12}$$

Here, $\dot{\phi}$ denotes the partial time derivative of ϕ . Having defined the conjugate field, we can now define the *Hamiltonian density* \mathcal{H} of a field ϕ

$$\mathcal{H} = \pi\dot{\phi} - \mathcal{L}(\phi, \partial_\mu\phi). \tag{1.13}$$

which defines the local energy density of the field.

So far the fields described in this section have been classical. Quantisation of the fields is achieved by promoting the fields themselves to operators and imposing certain commutation and anti-commutation relations on them. For a scalar field and its conjugate field, the following same-time relations must hold:

$$\begin{aligned}
 [\phi(\mathbf{x}, t), \pi(\mathbf{x}', t)] &= i\delta(\mathbf{x} - \mathbf{x}'), \\
 [\phi(\mathbf{x}, t), \phi(\mathbf{x}', t)] &= [\pi(\mathbf{x}, t), \pi(\mathbf{x}', t)] = 0.
 \end{aligned} \tag{1.14}$$

Fermionic fields have similar requirements hoisted upon them when promoted to operators. But rather than commuting, these fields are required to anti-commute:

$$\begin{aligned}
 \{\psi(\mathbf{x}, t), \psi^\dagger(\mathbf{x}', t)\} &= \delta(\mathbf{x} - \mathbf{x}'), \\
 \{\psi(\mathbf{x}, t), \psi(\mathbf{x}', t)\} &= \{\psi^\dagger(\mathbf{x}, t), \psi^\dagger(\mathbf{x}', t)\} = 0.
 \end{aligned} \tag{1.15}$$

Here we use that the conjugate field of ψ is $\pi = i\psi^\dagger$. The commutation relations of equations (1.14) and (1.15) restrict the possible values of certain observables of the fields to discrete, quantised, steps.

1.1.3. The Interactions of the Standard Model

In the previous section, we introduced free fields and their equations of motion. Next, we would like to include the fundamental interactions of nature to our model.

Symmetries are important in the SM. A symmetry is a feature or characteristic of some system that is unchanged, or *invariant*, under some mathematical transformation of the system. In QFT, symmetries are understood as operations performed on fields that leave their Lagrangian invariant.

Noether's theorem states that for every continuous symmetry of the Lagrangian \mathcal{L} , there exists a conserved quantity, or more accurately a conserved current j^μ . A spatial translation, for example, leaves the Lagrangian invariant. Momentum is the conserved quantity resulting from this symmetry. Similarly, invariance of the Lagrangian in time leads to the conservation of energy. Rotational symmetry leads to spin conservation. These symmetries, based on transformations of the space-time coordinates of the fields, are known as *external* symmetries of the Lagrangian.

Group theory provides a useful mathematical tool when describing symmetries, since continuous symmetries can be expressed using Lie groups. The external symmetries mentioned above belong to the so-called *Poincaré group*. In addition to these external symmetries, the SM also contains three *internal* symmetries, based on the groups $U(1)$, $SU(2)$, and $SU(3)$. The unitary group of n dimensions, $U(n)$, is the group of unitary $n \times n$ matrices with matrix multiplication being its so-called group operation. The special unitary group $SU(n)$ is the subgroup of $U(n)$ where its component matrices have determinant 1. Internal symmetries are based on local symmetries, meaning transformations with space-time dependence.

Quantum Electrodynamics

Quantum electrodynamics (QED) is a quantum field theory of the electromagnetic interaction, based on a local $U(1)$ symmetry. A $U(1)$ transformation of a field ψ is a rotation of the field in the complex plane, and can be written as

$$\begin{aligned}\psi(x) &\rightarrow \psi'(x) = \psi(x)e^{iqf(x)}, \\ \bar{\psi}(x) &\rightarrow \bar{\psi}'(x) = \bar{\psi}(x)e^{-iqf(x)},\end{aligned}\tag{1.16}$$

where q is electric charge and $f(x)$ is any sufficiently well-behaved space-time dependent function. We see that this transformation will not leave the Dirac Lagrangian, introduced in the previous section, invariant since

$$\partial_\mu \psi' = e^{iqf(x)} \partial_\mu \psi + iq \partial_\mu f(x) e^{iqf(x)} \psi.\tag{1.17}$$

Invariance of the Lagrangian can be recovered by replacing the differentiation operator ∂_μ with the *covariant derivative*

$$D_\mu = \partial_\mu + iqA_\mu,\tag{1.18}$$

where we introduce a new electromagnetic field A_μ , also known as the photon field. The substitution from ∂_μ to D_μ is known as *minimal substitution*. The covariant derivative of the field ψ undergoing a $U(1)$ transformation, becomes

$$D_\mu \psi' = (\partial_\mu + iqA_\mu) \psi e^{iqf(x)} = e^{iqf(x)} [\partial_\mu + iq \partial_\mu f(x) + iqA_\mu] \psi.\tag{1.19}$$

We see that if we also require that the new field A_μ transforms as

$$A_\mu \rightarrow A'_\mu = A_\mu - \partial_\mu f(x),\tag{1.20}$$

the Lagrangian remains invariant. The transformation introduced in Equation (1.20) is known as a coupled transformation.

The free fermion Lagrangian we started with now becomes

$$\mathcal{L} = \bar{\psi}(i\cancel{D} - m)\psi - \frac{1}{4}F_{\mu\nu}F^{\mu\nu}. \quad (1.21)$$

Here we have used $\cancel{D} = \gamma^\mu D_\mu$, so-called Feynman slash notation. Also included in this equation is the kinetic term describing the propagation of the free photon field. The electromagnetic field strength tensor $F_{\mu\nu}$ is defined as

$$F_{\mu\nu} = \partial_\mu A_\nu - \partial_\nu A_\mu. \quad (1.22)$$

It may not be immediately clear from looking at the Lagrangian in Equation (1.21) that this now includes a description of the electromagnetic interaction. However, by substituting the full equation of the covariant derivative into Equation (1.21),

$$\begin{aligned} \mathcal{L} &= \bar{\psi}(i\cancel{D} - m)\psi - q\bar{\psi}\cancel{A}\psi - \frac{1}{4}F_{\mu\nu}F^{\mu\nu} \\ &= \mathcal{L}_{\text{Fermion}} + \mathcal{L}_{\text{Interaction}} + \mathcal{L}_{\text{Photon}}, \end{aligned} \quad (1.23)$$

it becomes clear that a new term $\mathcal{L}_{\text{interaction}}$ has appeared in the Lagrangian, in addition to the terms for the free fermion and photon fields. This new term describes the interaction between charged fermions and photons.

The transformation of the field A_μ given in Equation (1.20) is identical to the gauge freedom of Maxwell's equations of classical electromagnetism, where the electromagnetic potential can be modified by the derivative of some arbitrary function $f(x)$ without altering the observable electric and magnetic fields \mathbf{E} and \mathbf{B} . In fact, we can show that by starting from the QED Lagrangian and applying the Euler-Lagrange equation with respect to the field A_μ , we re-derive the inhomogeneous Maxwell equations

$$\partial_\nu F^{\nu\mu} = -q\bar{\psi}\gamma^\mu\psi = J^\mu. \quad (1.24)$$

J^μ is the so-called *charge-current density*. Having imposed a continuous U(1) symmetry on the Lagrangian, by Noether's theorem there must exist some conserved current due to the symmetry. It can be shown that J^μ is exactly this conserved current. The conserved quantity associated with this current is electric charge q . A particle having an electric charge means that the particle interacts electromagnetically.

QED is a very successful theory. The measured value of the electron anomalous magnetic moment $g - 2$ agrees with QED predictions to within one part in a trillion [19].

Feynman Rules

A physical theory on its own is useless if it cannot be tested. We need to use the QED Lagrangian, arrived at in Equation (1.21), to make some predictions about nature. These predictions can then be compared to observation. In the field of high energy physics, this typically means calculating cross-sections³ of scattering and annihilation processes.

No interacting quantum field theory is exactly solvable. Instead, perturbation theory must be employed. Here, the Hamiltonian of the theory is divided into a free term governing the free propagating fields, and an interaction term. This interaction term is then

³The cross-section of a process is closely related to the probability that said process will take place in a particle collision. See Section 2.1.

treated as a perturbation. This approach is justified if the interaction is sufficiently weak, which is the case for QED, where the strength of interactions at the atomic scale is characterised by the fine structure constant $\alpha \approx 1/137$.

The *S-matrix* S , also known as the scattering matrix, contains the probability amplitudes of every possible transition between an initial and a final state of particles. The transition amplitude from some given initial state $|i\rangle$ to some final state $|f\rangle$ is defined as $S_{fi} = \langle f|S|i\rangle$. The Dyson expansion of S , expressed in terms of the interaction Hamiltonian densities \mathcal{H} is

$$S = \sum_{n=0}^{\infty} \frac{(-i)^n}{n!} \int \cdots \int d^4x_1 d^4x_2 \cdots d^4x_n T\{\mathcal{H}(x_1)\mathcal{H}(x_2)\cdots\mathcal{H}(x_n)\}. \quad (1.25)$$

Here, $T\{\}$ indicates time-ordering, such that later times are ordered to the left of earlier times in the product of Hamiltonian densities. Calculating S exactly is impossible, the sum must be truncated at some order n . Calculating cross-sections from Equation (1.25) directly is impractical, and physical interpretation of the underlying physical process of scattering is obfuscated by multiple complicated integrals. Therefore, when calculating cross-sections in high energy physics, it is common to use *Feynman diagrams* and *Feynman rules*.

Feynman diagrams, also known as Feynman graphs, are used to depict S -matrix elements of some order n in a graphical way. From equations (1.13) and (1.23) we see that the interaction Hamiltonian density of QED is

$$\mathcal{H} = -q : \bar{\psi} \not{A} \psi : . \quad (1.26)$$


Here, the $::$ notation indicate normal ordering, meaning that creation operators are to be placed to the left of annihilation operators in the products. We saw earlier that during quantisation, fields are promoted to operators. A field operator ψ can be decomposed into the sum of a *creation* ψ^+ and an *annihilation* operator ψ^- , responsible as the name suggest for creating and annihilating particles in the field. The interaction Hamiltonian density can then be written as

$$\mathcal{H} = -q : \left(\bar{\psi}^+ + \bar{\psi}^- \right) \left(\not{A}^+ + \not{A}^- \right) \left(\psi^+ + \psi^- \right) : . \quad (1.27)$$

This Hamiltonian density has eight terms, representing eight different permutations of a fermion, an anti-fermion, and a photon being created or annihilated in a point. One of these eight cases is drawn as a Feynman diagram in Figure 1.1. In a Feynman diagram, the horizontal axis indicates the flow of time from left to right, while the vertical axis denotes spatial movement. Photons are drawn as sinusoidal lines, while fermions are drawn as arrowed solid lines. Fermion arrows point towards increasing time, while anti-fermion lines point towards decreasing time. This is accordant with the Feynman-Stückelberg interpretation of an anti-particle as a particle running backwards in time. The point where the lines of a Feynman diagram meet is called a *vertex*. The number of vertices in a Feynman diagram exactly equals the order in perturbation theory of its corresponding S -matrix element.

The lowest order QED vertex shown in Figure 1.1 and its eight permutations are the building blocks for every higher-order Feynman diagram. Higher-order diagrams are constructed by joining the fermion or photon lines of two or more of the basic vertex. In this case, fermion lines must be connected to anti-fermion lines and vice versa, in such a way


Figure 1.1: The fundamental QED interaction vertex.



that only lines that "point" in the same direction meet; a fermion line cannot suddenly become an anti-fermion line between vertices.

In QED, the lowest order vertices, like the one shown in Figure 1.1 do not correspond to any real physical processes, as they would all violate energy and momentum conservation. Real physical processes must instead be constructed using diagrams with two vertices or more, such as the one drawn in Figure 1.2. In this process, an electron-positron pair annihilates into a photon, which in turn creates a muon-anti-muon pair. This diagram also illustrates the physical interpretation of the photon as the mediator of the electromagnetic interaction.

Figure 1.2: A Feynman diagram of the process $e^+e^- \rightarrow \mu^+\mu^-$.



Feynman diagrams have become the industry standard for visualising particle interactions, but they are also useful beyond their illustrative purpose. Feynman rules are mathematical tools, used for calculating cross-sections, based on Feynman diagrams. Feynman rules are used to calculate the Feynman amplitude \mathcal{M} of a process, whose absolute square $|\mathcal{M}|^2$ is proportional to the cross-section of said process. The Feynman rules of QED will not be given here⁴, but in essence, they assign a factor to the term for \mathcal{M} for each line and vertex in a Feynman diagram. So, to calculate \mathcal{M} of some process to a certain order of perturbation theory, first, write down every possible topologically distinct Feynman diagram of the chosen order and lower. Next, use the Feynman rules to calculate \mathcal{M} for each diagram, add the terms, and take the absolute square.

Quantum Chromodynamics

We learnt in Section 1.1.1 that hadrons, bound states of two or three quarks, are held together by the strong force. Quantum chromodynamics (QCD) is a quantum field theory of the strong interaction, somewhat analogous to QED. Whereas electromagnetism only has one electric charge, the strong interaction contains three colour charges. The field of each colour charge can be combined into a three-component field

$$\Psi = \begin{pmatrix} \psi_r \\ \psi_g \\ \psi_b \end{pmatrix}. \quad (1.28)$$

⁴See for example in Appendix B of [14]

The free Dirac Lagrangian can then be written

$$\mathcal{L} = \bar{\Psi}^f (i\gamma^\mu \partial_\mu - m) \Psi^f . \quad (1.29)$$

The index f indicates the quark flavour of the field.

The theory of the strong interaction is arrived at in much the same way as QED, only now the Lagrangian is required to be invariant under local SU(3) transformations of the field Ψ :

$$\Psi(x) \rightarrow \Psi'(x) = e^{ig_s \alpha_i(x) \lambda^i / 2} \Psi(x) , \quad (1.30)$$

where $\alpha_i(x)$ are eight arbitrary real space-time dependent angles, and λ_i are the eight Gell-Mann matrices, the *generators* of the SU(3) symmetry group. The covariant derivative is defined analogously to the one in Equation (1.18),

$$D_\mu = \partial_\mu + i \frac{g_s}{2} \lambda^i A_{i\mu} . \quad (1.31)$$

only now there are eight new *gluon fields* $A_{i\mu}$ instead of the single photon field in QED. The transformations of $A_{i\mu}$ themselves are complicated somewhat, due to the fact that SU(3) is a *non-Abelian* symmetry group, meaning that its group elements do not commute. For small α_i , $A_{i\mu}$ transforms as

$$A_{i\mu} \rightarrow A'_{i\mu} = A_{i\mu} - \partial_\mu \alpha_i - g_s f_{ijk} \alpha^j A_\mu^k , \quad (1.32)$$

where f_{ijk} are totally anti-symmetric so-called structure constants. We see here that a new term $-g_s f_{ijk} \alpha_j A_{k\mu}$ has appeared, which has no QED analogue. The full QCD Lagrangian now becomes

$$\mathcal{L} = \bar{\Psi} (i\not{D} - m) \Psi - \frac{1}{4} G_{i\mu\nu} G_i^{\mu\nu} . \quad (1.33)$$

$G_i^{\mu\nu}$ is the gluon field strength tensor, defined as

$$G_{i\mu\nu} \equiv \partial_\mu A_{i\nu} - \partial_\nu A_{i\mu} - g_s f_{ijk} A_\mu^j A_\nu^k . \quad (1.34)$$

The final term in this expression is needed to make the Lagrangian invariant, due to the new term added in the transformation of the gluon fields in Equation (1.32). When the dust has settled after imposing a local SU(3) symmetry to the theory, we can see that three new terms have appeared in the Lagrangian that can be interpreted as interaction terms:

$$\mathcal{L}_{\text{Interaction}} = -\frac{g_s}{2} \bar{\Psi} \lambda_i \not{A}^i \Psi + \frac{g_s}{2} f^{ijk} (\partial^\mu A_i^\nu - \partial^\nu A_i^\mu) A_\mu^j A_\nu^k - \frac{g_s^2}{4} f^{ijk} f_{ilm} A_j^\mu A_k^\nu A_\mu^l A_\nu^m . \quad (1.35)$$

The first of these terms is a quark-gluon interaction term analogous to the QED interaction, while the other two are *gluon self-interaction terms*. These interaction terms are illustrated as Feynman graphs in Figure 1.3. The gluon self-interaction terms are a consequence of the non-Abelian nature of SU(3). Gluon self-interaction means that the gluons themselves must carry colour charge. In fact, each gluon carries two units of colour charge, one colour- and one anti-colour charge.

The strong interaction is an asymptotically free theory. This means that the strength of the interaction grows with increasing length scales. This is directly opposite to QED, where the interaction wanes with distance, and is a direct result of the gluons carrying colour charge. If one tries to rip apart the quark-anti-quark pair in a meson into a free




Figure 1.3.: The fundamental QCD vertices.

quark and anti-quark, the force required to separate the quarks increases like in a rubber band. At some point this rubber band "snaps" as it becomes more energetically favourable to spontaneously create a quark-anti-quark pair from the vacuum, to pair up the colour charges in the meson. This effect, known as colour confinement, is the reason why no colour charged objects are seen in nature. It should be pointed out here that although asymptotic freedom is a predicted trait of QCD, colour confinement is not. Instead, colour charge is observed in experiments to be confined to hadrons.

The Weak Interaction

The weak interaction is the third fundamental interaction described in the SM mathematical framework. The weak interaction was first postulated in the 1930s by Enrico Fermi to explain the process of beta decay. In a beta decay, a down quark is transformed into an up quark, an electron, and an anti-electron neutrino. A new interaction theory is necessary to explain these flavour changing properties, as they cannot occur in strong or electromagnetic interactions. A new fermion also has to be introduced to the SM, the neutrino, to explain the apparent missing energy in beta decays.

Before presenting the theory of weak interactions, we have to introduce the concept of particle *handedness*. The *helicity* of a particle is the projection of its spin onto the direction of its momentum. A particle is said to be of *right-handed* helicity if this spin projection is parallel with its direction of motion, and of *left-handed* helicity if the spin projection is anti-parallel with its direction of motion. The *chirality* of a particle is a more abstract property, and is an intrinsic property of particles like their spin or their mass. A particle that transforms in a left- or right-handed representation of the Poincaré group is said to be left- or right-handed, respectively. The Poincaré group is the symmetry group of special relativity, ensuring Lorentz invariance under translations, rotations and boosts. The chirality and helicity of a massless particle coincide. The helicity of a massive particle, on the other hand, is dependent on the frame of the observer. However, for an ultra-relativistic particle, helicity and chirality can be treated as equal.

Charged weak interactions only couple to left-handed fermions. The left-handed ψ^L and right-handed ψ^R components of a fermion field ψ can be projected out using the *chirality operators* L and R, defined as

$$\left. \begin{aligned} \psi^L &= L\psi \\ \psi^R &= R\psi \end{aligned} \right\} = \frac{1 \mp \gamma^5}{2} \psi. \quad (1.36)$$

The matrix γ^5 is defined, using the Dirac gamma matrices, as $\gamma^5 = i\gamma^0\gamma^1\gamma^2\gamma^3$. The left- and right-handed fermion fields must be treated differently in the theory of weak interactions. A new quantum number is given to left-handed fermions, known as *weak*

isospin I. Left-handed fermions are paired in *weak isospin doublets*, while right-handed fermions form isospin singlets of $I = 0$. The left-handed charged lepton and neutrino pair of each lepton generation, for example, form the weak isospin doublet

$$\Psi_l^L = \begin{pmatrix} \psi_{\nu_l}^L \\ \psi_l^L \end{pmatrix}, \quad \bar{\Psi}_l^L = \left(\bar{\psi}_{\nu_l}^L \bar{\psi}_l^L \right), \quad (1.37)$$

where the index l indicates the lepton flavour. The value of the third component of isospin, I_3 , is $+1/2$ for the neutrino field $\psi_{\nu_l}^L$ and $-1/2$ for the charged lepton field ψ_l^L . The weak isospin doublets are analogous to the strong colour triplets introduced in Section 1.1.3. Weak isospin doublets also exist for left-handed quark fields.

The free Lagrangian of the left- and right-handed lepton fields can now be written

$$\mathcal{L} = \bar{\Psi}_l^L i \not{\partial} \Psi_l^L + \bar{\psi}_l^R i \not{\partial} \psi_l^R + \bar{\psi}_{\nu_l}^R i \not{\partial} \psi_{\nu_l}^R. \quad (1.38)$$

Like for QED and QCD, we arrive at the weak interaction by first imposing a local transformation on the weak isospin doublet. This time we require our Lagrangian to be invariant under an $SU(2)$ transformation

$$\Psi^L(x) \rightarrow \Psi^{L'}(x) = e^{ig\omega_i(x)\sigma_i/2} \Psi^L(x). \quad (1.39)$$

Here, g is the weak coupling constant, ω_i are three arbitrary real functions, and σ_i are the three Pauli matrices. $SU(2)$ transformations when applied to the weak isospin doublets are often denoted with a subscript L as $SU(2)_L$. Right-handed fermion fields ψ^R do not transform under this transformation, these fields are so-called $SU(2)_L$ *singlets*.

There are no mass terms in the free Lagrangian given in Equation (1.38). Any such terms, of the form $-m\bar{\psi}\psi$ would mix fields of left- and right-handed chirality, and as these fields transform differently under $SU(2)_L$, such terms would not be invariant.

As before, the Lagrangian can be made invariant under these new transformations by substituting the covariant derivative

$$D_\mu = \partial_\mu + i\frac{g}{2}\omega_i W_{i\mu}, \quad (1.40)$$

where $W_{i\mu}$ are three new gauge fields, required to transform as

$$W_{i\mu} \rightarrow W'_{i\mu} = W_{i\mu} - \partial_\mu \omega_i - g\epsilon_{ijk}\omega^j W_\mu^k, \quad (1.41)$$

for small ω_i .

The fields associated with the physical charged weak bosons W^+ and W^- are found by combining the fields $W_{1\mu}$ and $W_{2\mu}$ as

$$W_\mu^\pm = \frac{1}{\sqrt{2}} (W_{1\mu} \mp iW_{2\mu}) \quad (1.42)$$

One could be forgiven for assuming then, that the third field $W_{3\mu}$ is associated with the neutral weak boson Z^0 . However, things are not quite as simple, as we shall see in the next section.

For leptons, disregarding flavour changing neutrino oscillations, weak interactions only involve one of the three lepton families at a time. For quarks, however, weak interactions may be *flavour changing*. A W^- boson may, for example, decay to an anti-down quark and a charm quark, both belonging to different quark families. This is due to the fact that weak interaction quark states are not equal to the physical quark states. The relationship between the interaction states and physical mass states of the quarks is given by the *Cabbibo-Kobayashi-Maskawa* (CKM) matrix. Flavour changing weak interactions are only seen in interactions of the charged weak bosons.

Electroweak Unification

Requiring invariance of the free Lagrangian in Equation (1.37) under $SU(2)_L$ transformations, yields the conserved *weak isospin current*

$$J_i^\mu = \frac{1}{2} \bar{\Psi}^L \gamma^\mu \sigma_i \Psi^L. \quad (1.43)$$

The third component J_3^μ of this current, which is an electrically neutral current, takes the form

$$J_3^\mu = \frac{1}{2} \bar{\Psi}^L \gamma^\mu \sigma_3 \Psi^L = -\frac{1}{2} [\bar{\psi}_l^L \gamma^\mu \psi_l^L - \bar{\psi}_{\nu_l}^L \gamma^\mu \psi_{\nu_l}^L]. \quad (1.44)$$

Save for a factor e , the elementary electric charge, J_3^μ resembles the conserved current for charged leptons,

$$J_{EM}^\mu = -e \psi_l \gamma^\mu \psi_l. \quad (1.45)$$

This motivates the idea that the electromagnetic and weak forces can be described in a common mathematical framework. A modification to the weak interaction theory is further necessitated by the fact that the weak boson Z^0 , unlike its charged cousins W^\pm , can couple to right-handed fermions. Combining electromagnetism and the weak force into a common framework, aptly named the *electroweak interaction*, was first proposed by Glashow in 1961, and later expanded by Weinberg and Salam [20–22].

Combining J_3^μ and J_{EM}^μ , dividing out the factor e , yields the *hypercharge current*

$$J_Y^\mu \equiv J_{EM}^\mu / e - J_3^\mu = -\frac{1}{2} \bar{\Psi}^L \gamma^\mu \Psi^L - \bar{\psi}_l^R \gamma^\mu \psi_l^R. \quad (1.46)$$

The conserved charge corresponding to this current is called hypercharge Y , and is found by taking $Y = \int d^3x J_Y^0$. Since electric charge Q and the third component of weak isospin I_3 are both individually conserved, hypercharge Y must also be a conserved quantity. Hypercharge can be calculated using the *Gell-Mann-Nishijima* formula

$$\frac{Y}{2} = \frac{Q}{e} - I_3. \quad (1.47)$$

To marry the existing electromagnetic interaction with the weak interaction we exchange the local $U(1)_{EM}$ symmetry introduced in Section 1.1.3 with a new local hypercharge $U(1)_Y$ transformation

$$\psi(x) \rightarrow \psi'(x) = e^{ig'Yf(x)/2} \psi(x). \quad (1.48)$$

Here, ψ can be both a left-handed weak isospin doublet Ψ^L or a right-handed fermion field ψ^R . Once again the covariant derivative D_μ and a field B_μ must be introduced to ensure invariance of the Lagrangian. The field B_μ transforms analogously to the photon field A_μ of QED

$$B_\mu \rightarrow B'_\mu = B_\mu + \partial_\mu f(x). \quad (1.49)$$

The fields A_μ and Z_μ , associated with the physical photon and Z^0 particles, are linear combinations of B_μ and $W_{3\mu}$,

$$\begin{aligned} A_\mu &= B_\mu \cos \theta_W + W_\mu^3 \sin \theta_W, \\ Z_\mu &= -B_\mu \sin \theta_W + W_\mu^3 \cos \theta_W. \end{aligned} \quad (1.50)$$

Here, θ_W is the *weak mixing angle*, also known as the Weinberg angle. The weak mixing angle can be defined both as a ratio of weak boson masses

$$\cos \theta_W = \frac{m_W}{m_Z}, \quad (1.51)$$

or using the coupling constants g and g'

$$\cos \theta_W = \frac{g}{\sqrt{g^2 + g'^2}}. \quad (1.52)$$

The total symmetry group of the electroweak interaction, then, is $SU(2)_L \otimes U(1)_Y$. We can now write down the full interacting electroweak Lagrangian

$$\mathcal{L}_{EW} = \bar{\Psi}_l^L i \not{D} \Psi_l^L + \bar{\psi}_l^R i \not{D} \psi_l^R - \frac{1}{4} B_{\mu\nu} B^{\mu\nu} - \frac{1}{4} W_{i\mu\nu} W_i^{\mu\nu}. \quad (1.53)$$

The gauge field tensors $B_{\mu\nu}$ and $W_{\mu\nu}$ are analogous to those given in equations (1.22) and (1.34), respectively. The form of the covariant derivative D_μ is dependent on what field it operates on:

$$\begin{aligned} D_\mu \Psi_l^L &= \left[\partial_\mu + i \frac{g}{2} \sigma_i W_{i\mu} - i \frac{g'}{2} B_\mu \right] \Psi_l^L, \\ D_\mu \psi_l^R &= [\partial_\mu - i g' B_\mu] \psi_l^R. \end{aligned} \quad (1.54)$$

There is a factor-2 difference in the B_μ -terms due to the difference in hypercharge between left- and right-handed fermions.

There are no right-handed neutrino fields in the electroweak Lagrangian given in Equation (1.53). From Equation (1.47) we see that right-handed neutrinos have $Y = 0$. Having neither hypercharge, weak isospin, electric charge, or colour charge, right-handed neutrinos do not partake in any of the SM interactions. They are therefore omitted from the Lagrangian.

The basic interaction vertices of the electroweak interaction are shown as Feynman graphs in Figure 1.4 and Figure 1.5. Figure 1.4 shows the possible interactions between fermions and electroweak bosons, while Figure 1.5 shows the possible electroweak boson self-interactions.




Figure 1.4.: The fundamental electroweak fermion interaction vertices. Note that for diagram (a), only the Z^0 boson may annihilate into a neutrino-anti-neutrino pair.

One glaring issue remains: electromagnetism and the weak interaction are manifestly different interactions. Constructing the electroweak theory, all fermions and bosons are assumed to be massless. In nature, fermions and the weak bosons have mass. The unification of electromagnetism and the weak interaction is only occurring at high energy scales. In physics parlance, $SU(2)_L \otimes U(1)_Y$ must be a *broken* symmetry. The mechanism for breaking electroweak symmetry and introducing particle masses is covered in the next section.




Figure 1.5.: The fundamental electroweak self-interaction vertices.

1.1.4. Electroweak Symmetry Breaking

Consider a system described by some Lagrangian, invariant under some symmetry transformation. If the lowest energy state of the system is degenerate, there is no unique eigenstate to represent the ground state of the system. If one of the degenerate states is chosen to represent the ground state, this ground state is no longer invariant under the symmetry transformation. This phenomenon is known as spontaneous symmetry breaking (SSB). A spinning top is a simple example of a spontaneously broken symmetry. As long as the top is spinning, it has a spatial rotational symmetry; it looks the same from all directions. When the spinning top reaches its ground state, meaning when it loses its momentum and falls down, it comes to rest on the table pointing in a random direction. The rotational symmetry is now lost, or broken.

In QFT, the ground state of the Lagrangian is called the *vacuum*. If the electroweak symmetry is spontaneously broken, the lowest energy state of the SM Lagrangian must be degenerate. The vacuum is found by selecting one of these lowest energy states.

The Goldstone Model

A model exhibiting SSB is the *Goldstone model*, which is described by the Lagrangian

$$\mathcal{L} = \partial^\mu \phi^* \partial_\mu \phi - \mu^2 |\phi|^2 - \lambda |\phi|^4. \quad (1.55)$$

Here, ϕ is a complex scalar field, while μ and λ are arbitrary model parameters. This Lagrangian is invariant under a global U(1) symmetry transformation $\phi \rightarrow \phi' = \phi e^{i\alpha}$. The Goldstone potential is given by $V = \mu^2 |\phi|^2 + \lambda |\phi|^4$. To ensure that V is bounded from below, meaning that its ground state is a stable minimum, we require λ to be positive. If μ^2 is positive, the vacuum is unique at $\phi = 0$, prohibiting SSB. If μ^2 is negative, however, a local *maximum* occurs when $\phi = 0$, and the ground state of the Lagrangian, ϕ_0 , is a ring of degenerate states in the complex plane of ϕ ,

$$\phi_0 = \sqrt{\frac{-\mu^2}{2\lambda}} e^{i\theta} = \frac{v}{\sqrt{2}} e^{i\theta}, \quad 0 < \theta < 2\pi, \quad (1.56)$$

where θ is the angle of the minimum in this complex plane. This choice of $\mu^2 < 0$ and $\lambda > 0$ is drawn in Figure 1.6. Due to its shape, this potential is sometimes known as the *Mexican hat-* or *wine bottle potential*.

We can now select a value of θ in Equation (1.56) to represent the vacuum. For simplicity, take $\theta = 0$. The field ϕ can now be rewritten in terms of deviations from the chosen vacuum, as

$$\phi(x) = \frac{1}{\sqrt{2}} [v + \sigma(x) + i\eta(x)], \quad (1.57)$$




Figure 1.6: The Goldstone potential [23].

where σ and η are two real fields. The Goldstone Lagrangian now becomes

$$\mathcal{L} = \frac{1}{2} [\partial^\mu \sigma \partial_\mu \sigma] + \frac{1}{2} [\partial^\mu \eta \partial_\mu \eta] - \lambda v^2 \sigma^2 - \lambda v \sigma [\sigma^2 + \eta^2] - \frac{1}{4} \lambda [\sigma^2 + \eta^2]^2, \quad (1.58)$$

omitting an irrelevant constant term. This Lagrangian is no longer invariant under global U(1) transformations of the field ϕ , due to the term $-\lambda v \sigma [\sigma^2 + \eta^2]$.

In the Goldstone theory, the fields σ and η can be quantised and associated with physical scalar spin-0 particles σ and η . The mass of σ is $\sqrt{2\lambda v^2}$, while η is massless. Such massless particles are known as *Goldstone bosons*. No such particles have been observed in nature. *Goldstone's theorem* states that for any spontaneously broken global continuous symmetry, the theory must contain at least one massless particle [24]. However, in the next section, we shall see how we can avoid unphysical particles appearing in the theory after spontaneous symmetry breaking.

The Brout-Englert-Higgs Mechanism

We start once again with the Goldstone Lagrangian in Equation (1.55), with the constraints $\mu^2 < 0$ and $\lambda > 0$. Having chosen the real vacuum $\theta = 0$ in Equation (1.56), the field ϕ can again be expressed as a variation around the vacuum:

$$\phi(x) = \frac{1}{\sqrt{2}} [v + H(x)] e^{i\alpha(x)}, \quad (1.59)$$

where H and α are real scalar fields. Here, we have used polar coordinates rather than Cartesian coordinates like in Equation (1.57), to simplify the following calculation.

Rather than having it be invariant under global U(1) transformations, we now impose local U(1) transformations to the theory, meaning that the complex field ϕ transforms as $\phi \rightarrow \phi' = \phi e^{if(x)}$. To achieve this we proceed in the same manner as for QED in Section 1.1.3, by promoting the derivative ∂_μ to the covariant derivative $D_\mu = \partial_\mu + iqA_\mu$, and adding a gauge field A_μ that transforms as $A_\mu \rightarrow A'_\mu = A_\mu + \partial_\mu f(x)$.

The field ϕ can be made entirely real in every space-time point by requiring that the U(1) transformation function $f(x)$ takes the form $f(x) = -\alpha(x)$. This procedure of choosing a form of $f(x)$ like this is known as *gauge fixing*. The gauge chosen here is called the *unitary gauge*. The field A_μ now transforms as $A_\mu \rightarrow A'_\mu = A_\mu - \partial_\mu \alpha(x)$.

Dropping all constant and inconsequential terms, the Lagrangian of the transformed field ϕ' , which we will relabel ϕ for simplicity, now becomes

$$\begin{aligned} \mathcal{L} = & \frac{1}{2} \partial^\mu H \partial_\mu H - \frac{1}{2} (2\lambda v^2) H^2 - \frac{1}{4} F_{\mu\nu} F^{\mu\nu} + \frac{1}{2} (qv^2) A_\mu A^\mu \\ & - \lambda v H^3 - \frac{1}{4} \lambda H^4 + \frac{1}{2} q^2 A_\mu A^\mu (2vH + H^2). \end{aligned} \quad (1.60)$$

Here, the first line describes the free, non-interacting, Lagrangian, while the bottom line contains interaction terms. Mass terms have appeared in the Lagrangian for both fields H and A_μ .

Unlike the Lagrangian of Equation (1.58), Equation (1.60) contains no massless Goldstone bosons. The α degree of freedom has not simply vanished. Rather, it has been absorbed by the gauge field A_μ as a longitudinal mode of polarisation. The number of degrees of freedom in the theory has not changed. The process outlined here, where we start with a complex scalar field and a massless gauge field and end up with a real scalar field and a massive gauge field is known as the Brout-Englert-Higgs (BEH) mechanism. The massive spin-0 scalar particle associated with the field H is called a *Higgs boson*.

Spontaneous breaking of $SU(2)_L \otimes U(1)_Y$ through the BEH mechanism is called electroweak symmetry breaking (EWSB), and is analogous to the $U(1)$ breaking introduced above. First, promote the field ϕ to a weak isospin doublet

$$\Phi = \begin{pmatrix} \phi^+ \\ \phi^0 \end{pmatrix}, \quad (1.61)$$

where ϕ^+ and ϕ^0 are complex scalar fields. The Goldstone Lagrangian for a weak isospin doublet field becomes

$$\mathcal{L} = D^\mu \Phi^\dagger D_\mu \Phi - \mu^2 \Phi^\dagger \Phi - \lambda (\Phi^\dagger \Phi)^2. \quad (1.62)$$

The covariant derivative for a weak isospin doublet is given in Equation (1.54)⁵. This Goldstone Lagrangian is then added to the preexisting electroweak Lagrangian of Equation (1.53).

By requiring that $\mu^2 < 0$ and $\lambda > 0$, the ground state of the theory is once again degenerate. Selecting a value for the vacuum,

$$\Phi_0 = \frac{1}{\sqrt{2}} \begin{pmatrix} 0 \\ v \end{pmatrix}, \quad (1.63)$$

allows us to write Φ as variations around this ground state:

$$\Phi = \frac{1}{\sqrt{2}} \begin{pmatrix} \phi_1^+ + i\phi_2^+ \\ v + H + i\phi_2^0 \end{pmatrix}. \quad (1.64)$$

Selecting the unitary gauge, Φ becomes

$$\Phi = \frac{1}{\sqrt{2}} \begin{pmatrix} 0 \\ v + H \end{pmatrix}. \quad (1.65)$$

Once again a massive scalar particle H appears in the theory. The fields ϕ_1^+ , ϕ_2^+ , and ϕ_2^0 have not disappeared completely; the total number of degrees of freedom in the fields of the electroweak theory does not change before and after EWSB. Table 1.3 illustrates this migration of degrees of freedom from the Higgs weak isospin doublet Φ to the gauge fields $W_{i\mu}$.

After EWSB, mass terms for the W^\pm and Z^0 bosons have appeared in the Lagrangian, while the photon γ remains massless. Interaction terms for the new Higgs boson have also appeared in the Lagrangian. The term $D^\mu \Phi^\dagger D_\mu \Phi$ contains interactions between the Higgs boson and the weak bosons, while the Goldstone potential $\mu^2 \Phi^\dagger \Phi - \lambda (\Phi^\dagger \Phi)^2$ contains Higgs boson self-interaction vertices. The basic interaction vertices of these interactions are shown in Figures 1.7 and 1.8.

Table 1.3.: Counting degrees of freedom in the electroweak and Higgs fields before and after electroweak symmetry breaking. The absolute number of degrees of freedom is the same before and after symmetry breaking.

(a) Before EWSB				(b) After EWSB			
<i>Fields</i>	<i>Bosons</i>	<i>Type</i>	<i>DoF</i>	<i>Fields</i>	<i>Bosons</i>	<i>Type</i>	<i>DoF</i>
W^1, W^2	W^+, W^-	Massless	4	$W^{1,2}, \phi_{1,2}^+$	W^+, W^-	Massive	6
W^3, B^0	Z^0, γ	Massless	4	W^3, ϕ_2^0	Z^0	Massive	3
ϕ_1^+, ϕ_2^+	ϕ_1^+, ϕ_2^+	Scalar	2	W^3, B^0	γ	Massless	2
ϕ_1^0, ϕ_2^0	ϕ_1^0, ϕ_2^0	Scalar	2	H	H^0	Scalar	1
<i>Total</i>			12	<i>Total</i>			12




Figure 1.7.: The fundamental vertices of Higgs boson interactions with the electroweak bosons.




Figure 1.8.: The fundamental Higgs boson self-interaction vertices.

So far, the weak bosons have acquired mass terms in the Lagrangian, but the fermions remain massless. Fermion mass terms enter the SM Lagrangian through so-called *Yukawa terms*. For charged leptons, for example, the Yukawa terms take the form

$$\mathcal{L}_l^{\text{Yuk}} = -g_l \bar{\Psi}_l^L \Phi \psi_l^R - g_l \bar{\psi}_l^R \Phi^\dagger \Psi_l^L, \quad (1.66)$$

where the indices l run over the three lepton flavours, and g_l are free model parameters. The second term is the Hermitian conjugate (h.c.) of the first term. The Yukawa terms are invariant under $\text{SU}(2)_L \otimes \text{U}(1)_Y$ transformations. After EWSB, the Yukawa terms become

$$\mathcal{L}_l^{\text{Yuk}} = \frac{-g_l v}{\sqrt{2}} \bar{\psi}_l^L \psi_l^R - g_l \bar{\psi}_l^L \psi_l^R H + \text{h.c.} \quad (1.67)$$


The first of these terms is a mass term for the charged lepton l . The mass of the charged leptons is given by $m_l = g_l v / \sqrt{2}$. The second term of Equation (1.67) is an interaction term between the charged lepton field ψ_l and the massive scalar Higgs field H . We see that since the Yukawa coupling parameter g_l appears in both terms of Equation (1.67), the strength of this interaction is directly proportional to the mass of the charged lepton. The fundamental vertex of the Higgs interactions to fermions is shown as a Feynman graph in Figure 1.9.

Yukawa terms for down-type quarks can be written in the same form as the term in Equation (1.66). For up-type quarks we have to make a slight modification to the form of the Yukawa terms. Up-type quarks have weak isospin $+1/2$, but we chose in Equation (1.63) to place our vacuum expectation value for the Higgs weak isospin doublet in the isospin $-1/2$ position. The $+1/2$ and $-1/2$ positions of the Higgs weak isospin doublet can be rotated into each other by the transformation $\tilde{\Phi} = i\sigma_2 \Phi$, where σ_2 is the second Pauli matrix. The SM Yukawa terms for up-type quarks then become

$$\mathcal{L}_u^{\text{Yuk}} = -g_u \bar{\Psi}_q^L i\sigma_2 \Phi \psi_q^R + \text{h.c.} \quad (1.68)$$

Since right-handed neutrino fields $\psi_{\nu_l}^R$ have been dropped from the theory, no Yukawa terms exist for neutrinos. Neutrinos are therefore assumed to be massless in the SM.

Figure 1.9: The fundamental Higgs-fermion interaction vertex.



1.2. Physics Beyond the Standard Model

The SM is in excellent agreement with experimental results across a range of observables. Nonetheless, the SM has its shortcomings. Some experimental results are in tension with SM predictions. We mentioned in Section 1.1.3 that the observed anomalous magnetic moment $g - 2$ of the electron agrees with its SM prediction to very high precision. The muon anomalous magnetic moment on the other hand, which has also been measured to

⁵Equation (1.54) shows the covariant derivative of the weak isospin doublet of a lepton, which has opposite hypercharge to Φ .

very high precision, shows a discrepancy of 3.3 standard deviations, or 3.3 "sigma", between theory and observation [11].

In electroweak theory, the strength of the coupling of the weak bosons to leptons does not distinguish between lepton flavours. This phenomenon is known as lepton universality. Therefore, the branching ratios of the decay processes $B^+ \rightarrow K^+ e^+ e^-$ and $B^+ \rightarrow K^+ \mu^+ \mu^-$ are predicted to be equal in the SM. However, a recent measurement of these branching ratios by the LHCb Collaboration shows a 3.1 sigma disagreement with theory [25].

The aforementioned experimental results could be due to statistical fluctuations, but they hint at the possibility that the SM is an incomplete theory. In fact, we know that the SM cannot be a complete model of the interactions of elementary particles. Such a theory would of course have to include the fourth fundamental interaction, gravity. We assume that the SM is valid up to some energy scale, at which it can no longer make accurate predictions. The scale at which the quantum effects of gravity can no longer be ignored is called the *Planck Scale* $E_P \approx 10^{19}$ GeV. Such energies are well beyond the reach of particle colliders.

In addition to the experimental disparities mentioned above, some physical phenomena are not described by the SM at all. Listed below are some outstanding theoretical issues of the SM.

Neutrino Mass

This problem was hinted at in the previous section. Neutrinos are assumed to be massless in the SM, but this assumption is wrong. Neutrinos are shown to flavour-oscillate, a feature not possible without mixing between the interaction- and mass eigenstates of the neutrinos, and non-zero neutrino mass. Technically, there is nothing keeping us from adding the right-handed sterile neutrino fields back into the SM and adding neutrino mass terms to the SM Lagrangian through the Higgs mechanism. However, in doing so the issue of neutrino lightness, why is the mass of the neutrinos so small compared to the other SM particles, would remain.

A possible solution to the neutrino mass conundrum is found in the *seesaw mechanism*. By introducing a very heavy right-handed neutrino of mass M_R , the mass of the left-handed SM neutrinos can be shown in the seesaw theory to be of the order $m_{\nu_l} \approx m_l^2/M_R$ [26]. We see that if the mass of the heavy right-handed neutrino grows, the mass of the left-handed neutrino decreases, like the two sides of a playground seesaw. A requirement of the seesaw theory is for neutrinos to be Majorana particles. In other words, that they are their own anti-particles. Being electrically neutral, the neutrinos are the only fermions that could satisfy this requirement. However, no experimental evidence that neutrinos are Majorana particles has so far been observed.

The Hierarchy Problem

Higher-order calculations in QFT involve loop diagrams. These loops introduce divergent four-momentum integrals. Finite calculations are only achieved after *renormalisation* of the theory. The divergent integrals are *regularised* by introducing a regularisation parameter, often denoted as Λ . This dimensionful parameter can be regarded as the energy scale at which the theory (i.e. the SM) loses its validity. We often set this parameter to be equal to the Planck scale.

The bare Higgs mass is given in Equation (1.60) as $m_{0,H} = v\sqrt{2\lambda}$. The observable Higgs mass has been measured at the LHC to be 125 GeV. This mass is given as $m_H^2 =$

$m_{0,H}^2 + \Delta m_H^2$. Here, Δm_H^2 denotes a correction to the Higgs mass from loop diagrams, exemplified in Figure 1.10. Every elementary particle interacting with the Higgs boson adds to these corrections. Whereas most SM loop corrections are of the order $\mathcal{O}(\ln(\Lambda))$, corrections to the Higgs mass are of the order $\mathcal{O}(\Lambda^2)$. This means that for a regularisation parameter Λ set to the value of the Planck scale, Δm_H^2 becomes huge. The contribution to the Higgs mass from loop corrections are orders of magnitude greater than the Higgs mass itself. Achieving almost perfect cancellation between the bare Higgs mass and the loop corrections requires extreme fine tuning of the SM particle masses.




Figure 1.10.: Higgs correction diagrams from fermion (a) and boson (b) loops.

Baryon Asymmetry

There does not seem to be enough antimatter in the universe. In the Big Bang, baryons and anti-baryons, which is to say matter and antimatter, will have been produced in equal quantities. Why is it then, that today the universe only seems to consist of ordinary matter? We say that there is a *baryon asymmetry* in the universe. The SM does not conserve CP-symmetry, which is to say that the probability of a transition from a matter particle to an antimatter particle is not necessarily the same as the probability of the reverse process. CP-violation in the SM has been observed in matter-antimatter oscillations between neutral kaons, and B⁰- and D⁰-mesons. CP-violation in the SM can account for parts of the matter-antimatter asymmetry in the universe, but it cannot account for all of it. We do not know the mechanism through which ordinary matter "won out" over antimatter in the early universe.

Dark Matter

In the 1930s, Fritz Zwicky calculated that the luminous matter in the Coma galaxy cluster was only about 1 per cent of the mass needed to bind the galaxies to the cluster by gravitation. Zwicky's calculation has later turned out to be off by about an order of magnitude, but his observation is still valid: there seems to be more gravitational matter in the Coma cluster than there is luminous matter [27]. Later, in the 1970s, Vera Rubin and W. Kent Ford faced a similar conundrum when observing the rotation curves of galaxies; the luminous matter in the galaxies they observed were only a fraction of that required to keep the stars in the galaxies in their orbits [28]. The leading theory explaining these phenomena is that a new kind of matter, meaning a new class of elementary particle or particles, exist in the universe in addition to baryonic matter. This new matter is called *dark matter*, as it does not appear to interact through the electromagnetic interaction, and therefore emits no light. Dark matter is so far not explained by the SM. Neutrinos are ruled out as dark matter candidates, as their tiny mass cannot account for the abundance of dark matter in the universe. If dark matter does turn out to be a particle, and the dark

matter particles interact through any of the SM interactions, then we should expect to be able to create dark matter particles in particle collision experiments such as at the LHC.

Grand Unification

The coupling parameters of the SM interactions are not constant, but energy scale dependent. This is caused by contributions from higher-order loop effects in elementary particle interactions. At higher energy transfers, loop corrections become more influential, altering the observed coupling strengths of the strong, weak, and electromagnetic interactions. In the case of the strong and weak force, the self-interaction of gluons and weak bosons lead to a so-called anti-screening effect on the couplings. This means that the strong and weak couplings both *decrease* as a function of energy scale. Photons do not self-interact. Rather than anti-screening, *screening* is seen in the electromagnetic interaction. The electromagnetic coupling strength *increases* as a function of energy scale.




Figure 1.11.: The inverse coupling strengths of the strong-, weak-, and electromagnetic interactions as a function of interaction energy. The couplings are extrapolated, from values measured in experiments at lower energies, to high energies, where their curves intersect [29].

If we extrapolate the strong, weak, and electromagnetic coupling strengths, as measured at energies attainable at colliders, to higher energies, the couplings are almost seen to meet in a single point. This is shown in Figure 1.11. The not-quite crossing point of the three coupling strengths happens at an energy scale of about $\mathcal{O}(10^{15})$ GeV, far beyond the current reach of particle colliders. The apparent convergence of the coupling strengths hints at the possibility of unifying the three forces into a single combined interaction, analogous to the electroweak unification. This concept is known as *grand unification*. Grand unified theories (GUTs) are usually built around a common symmetry group including all three SM forces. A common example of a GUT symmetry group is SU(5). This symmetry must then be broken for us to be left with the SM symmetry groups at low energy scales. The SU(5) symmetry group is excluded as a GUT candidate by proton lifetime experiments and by precision measurements of the Weinberg angle.

1.2.1. Heavy Neutral Bosons

Beyond standard model (BSM) is the field of study seeking to extend and ameliorate the SM. This section, and the next, will cover some BSM models that may predict increased production of dilepton events in collider experiments.

The SM is based on the symmetry $SU(3)_C \otimes SU(2)_L \otimes U(1)_Y$. Adding a new $U(1)$ symmetry to the SM predicts the existence of a new neutral gauge boson. This boson is often called Z' , named for its similarity to the SM Z boson. No additional $U(1)$ symmetries are observed in experiments, so any such additional new symmetry must be broken at low energy.

Perhaps the simplest extension to the SM is in the form of an additional $U(1)$ symmetry where the couplings of the predicted Z' boson are identical to those of the SM Z boson [30]. This model is known as the Sequential Standard Model (SSM). The SSM model cannot be realised in nature in this fundamental form, as it is non-renormalisable. However, it is still commonly used as a benchmark model.

In the SM, left-handed fermion fields are grouped together in $SU(2)_L$ isospin doublets. This means that only left-handed fields carry weak isospin. In a *left-right symmetric model* nature does not distinguish between left- and right-handed fermions. Such a left-right symmetry occurs for example in the GUT based on the $SO(10)$ symmetry group [30]. This symmetry may be broken as $SO(10) \rightarrow SU(3)_C \otimes SU(2)_L \otimes SU(2)_R \otimes U(1)_{B-L}$. Here, right-handed fermions are also grouped together in $SU(2)_R$ isospin doublets. The electroweak symmetry group is now extended to become $SU(2)_L \otimes SU(2)_R \otimes U(1)_{B-L}$. In certain left-right symmetric models, hypercharge Y is no longer needed to determine electromagnetic charge Q . Instead, Q becomes a function of baryon and lepton numbers B and L , as

$$Q = I_3^L + I_3^R + \frac{B - L}{2}, \quad (1.69)$$

instead of the familiar Gell-Mann-Nishijima formula of Equation (1.47).

The new symmetry group $SU(2)_R$ introduces three new gauge bosons: a neutral boson Z' , and two charged bosons W'^{\pm}_R . No symmetry $SU(2)_R$ has been observed in nature, meaning that if it exists at all, the left-right symmetry must be a broken symmetry. Left-right symmetry breaking happens in a mechanism analogous to the electroweak symmetry breaking of the SM, covered in Section 1.1.4. The Higgs sector of left-right symmetric models must be more complex than the SM Higgs model, containing for example charged Higgs bosons. During symmetry breaking, the left-right symmetric symmetry group is first broken to produce the electroweak symmetry group, as $SU(2)_L \otimes SU(2)_R \otimes U(1)_{B-L} \rightarrow SU(2)_L \otimes U(1)_Y$. This symmetry breaking step gives mass to the "primed" gauge bosons of the right-handed symmetry group. These masses must necessarily be greater than the electroweak scale. Next, the electroweak symmetry group is broken through the Higgs mechanism, as explained in Section 1.1.4.

Another class of Z' models is based on GUTs involving the E_6 symmetry group [30, 31]. An interesting feature of these theories is that they predict the existence of at least one additional heavy neutral boson at low energies. In many models, the E_6 symmetry group is first broken to a $SO(10) \otimes U(1)_\psi$ symmetry. The $SO(10)$ symmetry group is further broken down, yielding a $SU(5) \otimes U(1)_\chi$ symmetry. The SM symmetry groups are produced by symmetry breaking of the $SU(5)$ symmetry. All three aforementioned symmetry breaking steps happen around the GUT scale. The $U(1)_\psi$ and $U(1)_\chi$ symmetries give rise to two new neutral bosons, Z'_ψ and Z'_χ . These bosons are usually allowed to mix, leading to a

single observable Z' state, given as

$$Z' = Z'_\psi \cos \theta_{E_6} + Z'_\chi \sin \theta_{E_6}, \quad (1.70)$$

where $0 < \theta_{E_6} < \pi$. The value of the mixing angle θ_{E_6} determines the coupling strength of the Z' boson to SM fermions, as well as the width of its resonance. The SSM model Z' has an intrinsic width of around 3% of its mass. The E_6 model Z' is narrower, depending on the choice of θ_{E_6} . A pure Z'_ψ ($\theta_{E_6} = 0$) has a width of 0.5%, while a pure Z'_χ ($\theta_{E_6} = \pi/2$) has a width of 1.2%. These two models bookend the possible decay widths of the E_6 Z' resonance.

Three Z' benchmark models are considered in the search for resonant dilepton phenomena covered in Chapter 5: Z'_ψ , Z'_χ , and Z'_{SSM} .

In a heavy vector triplet (HVT) theory, the SM is extended by a triplet in the so-called adjoint representation of $SU(2)_L$ [32]. Such triplets appear for example in composite- or little Higgs models, and in certain GUT theories. The HVT triplet contains a neutral Z'_{HVT} boson and two W'_{HVT}^\pm bosons. The new bosons are mass degenerate, and have zero hyper- and colour charges. HVT theory contains three free model parameters: g_l , g_q , and g_H , the coupling of the HVT triplet to leptons, quarks, and the Higgs boson, respectively. The coupling to fermions is usually taken to be universal, or $g_f = g_l = g_q$.

So far, we have only considered BSM theories predicting spin-1 bosons, but the SM may also be extended by spin-0 and spin-2 bosons. Spin-2 theories will be covered in the next section. New bosons of spin-0, so-called scalar bosons, may arise for example in the minimal supersymmetric standard model (MSSM) [33].

1.2.2. Large Extra Dimensions

Perhaps the largest issue facing the SM is that it is not a theory of the gravitational interaction. Another conundrum facing physicists is the question of why gravity is so feeble compared to the other fundamental forces. This question is sometimes referred to as the big hierarchy problem. A possible solution to this issue is the existence of one or more extra spatial dimensions. The spread of gravity into these additional dimensions could explain its apparent weakness in the four observable space-time dimensions [34].

To be consistent with observations, the new dimensions are considered to be small. The process of folding dimensions in on themselves is known as *compactification*. A scalar field in a higher-dimensional room will manifest in ordinary space-time as an infinite collection of massive scalar fields. These fields can be considered quantised four-dimensional excitations of the higher-dimensional scalar field's momentum. This is known as a Kaluza-Klein (KK) *tower* of excitations. The mass spacing of the KK modes is determined by the size of the compact extra dimensions. The KK tower can be interpreted as space-time excitations of a massive spin-2 graviton.

The Randall-Sundrum 1 (RS1) model introduces one new spatial dimension [35]. In RS1, five-dimensional space-time, also known as the *bulk*, is warped. The added spatial dimension is compactified on a so-called S^1/\mathbb{Z}_2 *orbifold*. Two four-dimensional "surfaces", so-called *3-branes*, on the bulk are connected by this orbifold, the SM brane and the UV or Planck brane. The SM is located on the SM brane, while gravity is located at or near the Planck brane. Only the gravitational field propagates in the bulk, all other fields are restricted to the SM brane.

In Arkani-Hamed, Dimopoulos, and Dvali (ADD) theory, ordinary space-time is ex-

tended by two or more flat spatial dimensions⁶ [8]. The n extra dimensions are compactified on an n -dimensional torus. The size parameter of the torus is R . As in the RS1 theory, the SM fields are located on a four-dimensional hypersurface of the $4-n$ -dimensional bulk, while gravity is allowed to propagate in the bulk itself. The *fundamental Planck scale* in $4+n$ dimensions M_D is related to the Planck scale M_{Pl} through Gauss' law

$$M_{\text{Pl}}^2 \approx M_D^{n+2} R^n. \quad (1.71)$$

For sufficiently large values of R , the fundamental scale of gravity can be as low as a few TeV [36], which could lead to a solution of the hierarchy problem. For this reason, the new spatial dimensions introduced in the ADD model are sometimes referred to as "large" extra dimensions.

Compactifying the extra dimensions in ADD theory leads to a discrete set of KK graviton excitation states. The mass of the i -th individual excitation is given by

$$m_i^2 = \frac{i^2}{R^2}. \quad (1.72)$$

In the ADD model, R is large. So large in fact, that the spacing between neighbouring KK modes can be considered negligible. In ADD theory, therefore, the graviton spectrum can be considered continuous and non-resonant. In the RS1 model, on the other hand, the graviton spectrum is expected to be more discrete, with definite resonances at each KK mode.

The sum of all KK modes must be regulated, so as not to run to infinity. This is done with a UV cutoff at some energy scale. This cutoff is set on the *string scale* M_S of the theory, given by [37]

$$M_S = 2\sqrt{\pi} \left[\Gamma\left(\frac{n}{2}\right) \right]^{1/(n+2)} M_D, \quad (1.73)$$

where Γ is the gamma function.

The total cross-section σ_{Tot} for the process $q\bar{q} \rightarrow l^+l^-$ in the ADD model can be described by modifying the SM Drell-Yan (DY) cross-section σ_{DY} of the process:

$$\sigma_{\text{ADD}}^{\text{Tot}} = \sigma_{\text{DY}} + \mathcal{F} \frac{F_{\text{Int}}}{M_S^4} + \mathcal{F}^2 \frac{F_G}{M_S^8}. \quad (1.74)$$

The cross-section signal contribution is divided into an interference term and a pure graviton term, represented by F_{Int} and F_G , respectively. Giudice, Rattazzi, and Wells (GRW) [38], Hewett [39] and Han, Lykken, and Zhang (HLZ) [40] define the following conventions for calculating the form factor \mathcal{F} :

$$\begin{aligned} \mathcal{F} &= 1 \quad (\text{GRW}). \\ \mathcal{F} &= \frac{2\lambda}{\pi} = \frac{\pm 2}{\pi} \quad (\text{Hewett}). \\ \mathcal{F} &= \log\left(\frac{M_S^2}{s}\right) \text{ for } n = 2 \quad (\text{HLZ}). \\ \mathcal{F} &= \frac{2}{n-2} \text{ for } n > 2 \quad (\text{HLZ}). \end{aligned} \quad (1.75)$$

Only the Hewett convention allows for a negative sign (destructive) interference term with the SM. However, for practical reasons, only the positive sign convention will be considered in the analysis of Chapter 6.

⁶The ADD case of exactly one extra spatial dimension is excluded by observations. The existence of such a theory predicts modified Newtonian gravity on the solar system scale

1.2.3. Contact Interactions

The ADD model is not the only model predicting non-resonant high-mass dilepton signatures. Contact interactions (CIs) are examples of effective field theories (EFTs). An EFT can be considered a low energy approximation of a more fundamental, or "full", quantum theory [41]. An advantage of EFTs is that they allow us to ignore, or "integrate out", any new degrees of freedom introduced by a high-energy fundamental theory, and rather construct a theory using only SM fields.

Assume that the SM is only valid up to some energy scale Λ , above which new physics is needed to accurately describe nature. New interactions can be constructed with coefficients proportional to the inverse power of Λ . A general effective Lagrangian can be written

$$\mathcal{L} = \mathcal{L}_{\text{SM}} + \sum_{i=1}^{\infty} \frac{c_i}{\Lambda^i} \mathcal{O}_i. \quad (1.76)$$

Here, c_i are dimensionless coefficients, and the operators \mathcal{O}_i are functions of SM fields. The above expansion may be calculated to any order in Λ^{-1} , but the lowest order operators will be dominant. We see that the SM is recovered when taking the limit $\Lambda \rightarrow \infty$. The action S of a Lagrangian is a dimensionless quantity. To satisfy this condition, every SM operator has dimension $[\mathcal{O}_{\text{SM}}] = E^4$. We say that they are dimension four operators. For S to remain dimensionless, the new operators added in Equation (1.76) must have dimensions $[\mathcal{O}] = E^{4+i}$. The effective theory is only valid at energies far below the Λ scale [42]. The most dominant new interactions are the lowest order terms of Equation (1.76). In this work, we will only consider dimension six operators. One gauge-invariant dimension five operator can be constructed, the so-called Weinberg operator. However, this operator does not give rise to four-fermion contact interactions, and will therefore be ignored.

Perhaps the most famous example of an effective field theory is Fermi's model of beta decay [15]. In the 1930s, when Fermi developed his theory, the energy reach of particle physics experiments was far lower than that required to produce weak bosons. Fermi was able to model the weak interaction, not through the exchange of W or Z bosons, but as a four fermion contact interaction. Fermi's model is an effective theory of the weak interaction at energy scales far less than the mass of the weak bosons. This model-building approach can be used today, in the search for new heavy bosons like the ones mentioned in Section 1.2.1. If a Z' boson exists with a mass beyond the reach of modern colliders, we would not be able to observe this particle directly, but we could observe new CIs. A Feynman diagram of a CI scattering process is shown in Figure 1.12.




Figure 1.12: Feynman diagram of a four fermion contact interaction, where a quark pair annihilates to a lepton final state in a single point.

Another high-energy theory that could manifest itself through CIs at lower energies is the theory of quark and lepton compositeness; the idea that quarks and leptons are not elementary particles, but are made up of *preons*. In this theory, quarks and leptons are preons bound together by some force, analogous to the strong force [43].

Fermion fields have mass dimension 3/2, so a four fermion interaction has to be represented by a dimension six operator. 63 dimension six operators can be constructed from

fermion fields [44]. The number of operators reduces to 59 if baryon number conservation is imposed in the interaction. In the analysis presented in Chapter 6, we only consider the operator

$$\mathcal{O}_{lq} = (\bar{q}\gamma_{\mu}q)(\bar{l}\gamma^{\mu}l) , \quad (1.77)$$

which may appear in quark- and lepton compositeness scenarios. When separating the left and right components of the quark and lepton fields, this operator leads to the following Lagrangian for the CI:

$$\begin{aligned} \mathcal{L}_{\text{CI}} = \frac{g^2}{\Lambda^2} [& \eta_{\text{LL}}(\bar{q}_{\text{L}}\gamma_{\mu}q_{\text{L}})(\bar{l}_{\text{L}}\gamma^{\mu}l_{\text{L}}) \\ & + \eta_{\text{RR}}(\bar{q}_{\text{R}}\gamma_{\mu}q_{\text{R}})(\bar{l}_{\text{R}}\gamma^{\mu}l_{\text{R}}) \\ & + \eta_{\text{LR}}(\bar{q}_{\text{L}}\gamma_{\mu}q_{\text{L}})(\bar{l}_{\text{R}}\gamma^{\mu}l_{\text{R}}) \\ & + \eta_{\text{RL}}(\bar{q}_{\text{R}}\gamma_{\mu}q_{\text{R}})(\bar{l}_{\text{L}}\gamma^{\mu}l_{\text{L}})] . \end{aligned} \quad (1.78)$$

Here, the chosen convention for the coupling g is $g^2/4\pi = 1$ [45]. q_{L} (q_{R}) and l_{L} (l_{R}) are left-handed (right-handed) quark- and lepton fields, respectively. The chiral structure of the new interaction is determined by η_{ij} , which can take the values -1, 0, or 1. In the analysis of Chapter 6, each term in Equation (1.78) is considered separately, setting the remaining three terms to zero.

The total cross-section for the process $q\bar{q} \rightarrow l^+l^-$ is found by adding terms for the CI and its interference with the SM to the SM Drell-Yan cross-section:

$$\sigma_{\text{CI}}^{\text{Tot}} = \sigma_{\text{DY}} + \eta_{ij} \frac{F_{\text{Int}}}{\Lambda^2} + \frac{F_{\text{CI}}}{\Lambda^4} . \quad (1.79)$$

The sign of η_{ij} determines the interference of the CI with the SM; a negative sign gives negative interference, while a positive sign gives positive interference. F_{Int} and F_{CI} are functions of invariant mass m_{ll} , and represent the interference- and "pure CI" terms, respectively.

2. Experimental Setup

2.1. Introduction to Particle Collider Experiments

Particle colliders are the main tools by which physicists study the properties of subatomic particles. Only a few subatomic particles are stable in nature: electrons, photons, neutrinos, protons and neutrons (when bound inside nuclei). All other particles will decay, when produced, into stable decay products. To study such transient particles, they must be produced in a laboratory by colliders, or in atmospheric cosmic ray showers.

Particle colliders create particles by exploiting the mass-energy equivalence principle $E = mc^2$. A consequence of this principle is that the energy of two particles colliding can be converted to mass in the form of new particles. The mass of the produced particles is limited by the energy of the initial colliding particles. More specifically, only the energy in the COM frame of reference of the colliding particle system is available to produce new particle states. This has implications for particle collider design. For a so-called fixed-target collider, in which a single beam of particles impacts a stationary target, the collision energy in the COM system of the particle collisions is proportional to the square root of the energy of the beam. For a colliding beam accelerator, on the other hand, the COM energy grows linearly with the beam energy. Two proton beams of 21.6 GeV colliding have the same COM collision energy as a single 1 TeV proton beam striking a fixed-target [26].

Colliding beam experiments benefit from having circular accelerators. When two particle beams collide, only a small fraction of the particles contained in the beams interact. The rest of the beam, the so-called beam remnants, continue past the collision point unaffected. In a linear collider experiment these beam remnants, and the energy spent to accelerate these particles, would be wasted. In a circular collider, on the other hand, the beam remnants of one particle collision can be recovered and brought back around the ring to participate in new collisions. This has the added benefit that the particle beams in a circular accelerator do not have to be accelerated to their target energies in one pass around the detector.

The most common circular accelerator design in use today is the so-called *synchrotron*. In a synchrotron, the particle beams are held in circular orbits by magnetic fields. Acceleration is achieved by electric fields oscillating in so-called radiofrequency cavities. The oscillation of these electric fields is synchronised (hence the name synchrotron) to the orbit of the particle beams in such a way that the force they exert on the beams is always in the same direction, accelerating the particles. A consequence of this setup is that the particle beams cannot be continuous streams of particles, but have to be grouped in so-called bunches. If this was not the case then some particles would inevitably experience a braking force and not an accelerating force when the electric field in the accelerator cavities changes direction. The recycling of the beams means that only a section of the accelerator ring needs to be dedicated to acceleration.

Circular colliders also have their limitations. For one, in colliding beam experiments, the rate of interaction between the particles in the colliding beams is small, compared

to interaction rates in fixed-target experiments. This necessitates high beam intensities, which means adding more particles to the beams or focusing the beams into smaller areas before the collisions. The particles in a circular collider are also constantly accelerating inwards due to the bending magnets keeping the particles in their circular orbits. According to the laws of electrodynamics, this means that they radiate energy. The power radiated by one particle can be found using Liénard's relativistic generalisation of the Larmor formula

$$P = \frac{\mu_0 q^2 \gamma^6}{6\pi c} \left(a^2 - \left| \frac{\mathbf{v} \times \mathbf{a}}{c} \right|^2 \right), \quad (2.1)$$

where μ_0 is the vacuum permeability, q is the electric charge of the particle, and γ is the relativistic Lorentz factor $\gamma \equiv 1/\sqrt{1-v^2/c^2}$ [46]. When the acceleration experienced by the particle is normal to the direction of motion, this formula becomes

$$P = \frac{\mu_0 q^2 \gamma^4 a^2}{6\pi c}. \quad (2.2)$$

The acceleration of the particles in a circular accelerator is given by $a = v^2/r$, where r is the radius of the accelerator. The energy of the particle in the laboratory reference frame is $E = \gamma mc^2$. For ultra-relativistic particles $v \approx c$. Using these relations, Equation (2.2) may be written

$$P = \frac{\mu_0 q^2}{6\pi c^5 r^2} \left(\frac{E}{m} \right)^4. \quad (2.3)$$

This type of radiation is known as synchrotron radiation. We see that the power of this radiation depends strongly on the beam energy E . This means that for a fixed radius r of a circular accelerator, the beam energy cannot be made arbitrarily large. At some point, the energy loss to synchrotron radiation becomes greater than the energy put in to accelerate the particles in the beams. To get around this limitation, and reach higher beam energies, one can either decrease the curvature of the accelerator, or increase the mass m of the particles being accelerated. The factor m^{-4} in Equation (2.3) means that light particles lose far more energy to synchrotron radiation than heavy particles.

The general rule when constructing accelerators is to make them as compact as possible to minimise material costs. This introduces another limitation on the beam energy, namely the strength of the bending magnets keeping the particles in orbit around the accelerator. A smaller accelerator means a shorter radius of curvature for the particles, and having a shorter radius of curvature means having a stronger magnetic field in the bending magnets.

In addition to beam energy, *luminosity* \mathcal{L} is an important figure of merit for a particle collider. Luminosity is closely related, but not equal, to the particle collision rate of the collider. Rather, it is a measure of the number of particles colliding per area per second. Luminosity integrated over time is also a much-used quantity in particle physics, and is often simply called luminosity L . The symbol \mathcal{L} is sometimes used for both instantaneous and time-integrated luminosity, exacerbating the chance of confusion. In this text, luminosity will be referred to as *instantaneous luminosity*, while time-integrated luminosity will be called luminosity or integrated luminosity.

The reason for using luminosity is closely related to the measure known as *cross-section* σ , which is a measure of the probability that a given process will occur in a particle collision. Cross-section is defined as having unit area. On the macroscopic scale, for example for two balls being tossed in the air towards each other, the probability of a collision occurring

between the balls depends on their size. More precisely, it depends on an area transverse to the relative direction of motion of the balls. If the areas of the projections of the two balls onto the plane perpendicular to the relative direction of motion overlap, the balls will collide. For two balls of equal radius r , the cross-section becomes $\sigma = 4\pi r^2$. When we move to the microscopic world of subatomic physics, this analogy of hard spheres colliding breaks down somewhat. For example, cross-sections on the microscopic scale may be energy dependent. In the macroscopic example of flying balls, this would be analogous to the balls growing or shrinking depending on their kinetic energy. A useful property of cross-sections is that they are independent of beam conditions, such as the intensity and focus of the beams. A cross-section obtained at one experiment can therefore be directly compared to theory and other experiments. When given the instantaneous luminosity of a beam experiment, and the cross-section of some scattering process, the production rate n in the experiment of said process is simply found by multiplying these two values:

$$n = \sigma \mathcal{L}. \quad (2.4)$$

Similarly, finding the total number of events N occurring for the process of interest in a certain data-taking period of the experiment, is found by multiplying the cross-section of the process and the integrated luminosity:

$$N = \sigma L. \quad (2.5)$$

The unit of cross-section most commonly used in high energy physics is *barn* b. $1 \text{ b} = 10^{-24} \text{ cm}^2 = 100 \text{ fm}^2$. Luminosity is often given in inverse barn b^{-1} .

2.2. CERN and the Large Hadron Collider

The European Organization for Nuclear Research (CERN) is a multi-national organisation for particle and nuclear physics research. The organisation operates the world's largest particle physics laboratory, situated just outside of Geneva, Switzerland, on the French-Swiss border. CERN was founded in 1954 and has been instrumental in many physics discoveries since. A few examples are the discovery of the W and Z bosons at the Super Proton-Antiproton Synchrotron (Sp \bar{p} S) in 1983, the first synthesis of anti-hydrogen in the Low Energy Antiproton Ring (LEAR) in 1995, and the discovery of a 125 GeV neutral boson consistent with the Higgs boson at the LHC in 2012.

The Large Hadron Collider (LHC) is a circular hadron collider located at the CERN site. It is the current world record holder for beam energy in an accelerator. The collider is located in an approximately 27 kilometre long tunnel, around 100 metres underground, originally dug to house the Large Electron Positron Collider (LEP). The LHC is set up for both proton-proton (pp) collisions, and for collisions of heavier nuclei, but this text will only consider results from pp collisions. 1,232, 8 tesla superconducting dipole magnets bend the proton beams to their circular orbits. Before being brought to collide in four points along the accelerator ring, the beams of the LHC are focused and squeezed by thousands of quadrupole magnets and magnetic fields of higher order still. Some important design specifications of the LHC are listed in Table 2.1.

The LHC is not capable of accelerating protons from a standstill. The proton beams already have an energy of 450 GeV when injected into the LHC. Ramping up the beams to this energy takes place in stages, in accelerators from CERN's earlier history. The so-called *LHC injector chain* begins in an unassuming canister of hydrogen gas. The hydrogen gas

is ionised and the protons are accelerated to 50 MeV in the Linac2 linear accelerator. From Linac2, the protons are injected into the Proton Synchrotron Booster (PSB), known simply as the Booster. The PSB consists of four individual beam pipes, stacked vertically, which means that it can accelerate four proton bunches in parallel. In the Booster the proton bunches are accelerated to 1.4 GeV before being injected into the Proton Synchrotron (PS). Next, when the energy of the beams reach 25 GeV, they are injected into the Super Proton Synchrotron (SPS). This is the final stage of ramp-up of the beams, where the beams are accelerated to the required 450 GeV for the LHC [47]. Filling each beam of the LHC and accelerating them to collision energies takes of the order of one hour. One fill of the machine can produce particle collisions for several tens of hours. When not tasked with filling the LHC, the injector accelerators provide particle beams for other CERN experiments.

Table 2.1.: Design specifications of the LHC beams [48]. Also included in the table is the maximum achieved values for the same parameters during the LHC Run 2 data-taking campaign [49].

<i>Feature</i>	<i>Design (Nominal)</i>	<i>Run 2 Delivered (Max)</i>
Peak instantaneous luminosity [$10^{34}\text{cm}^{-2}\text{s}^{-1}$]	1.0	2.1
Beam energy [GeV]	7,000	6,500
Bunches per beam	2,808	2,556
Protons per bunch [10^{11}]	1.15	1.25

The LHC started up in 2008, but had to be shut down for 14 months for repairs after an accident caused by a defective electrical connection. Between 2010 and 2013 the machine delivered proton and heavy ion collisions to its experiments before being shut down for upgrades and maintenance. This period of data-taking was named *Run 1*. The LHC started back up in 2015 and ran until 2018. *Run 2* is the label given to data collected in this period. During Run 2 the LHC delivered over 150fb^{-1} integrated luminosity to its main experiments, exceeding projected goals for the period. Table 2.1 also lists the maximum values delivered in Run 2 for some important accelerator parameters. The searches covered in Chapters 5 and 6 are performed on pp data collected during Run 2.

2.3. The ATLAS Detector

Creating high-energy particle collisions in the LHC is only half the story. For the collisions to have any kind of scientific benefit, we need some way of observing them. For this we need *particle detectors*. ATLAS (A Toroidal LHC ApparatuS) is one the four main physics experiments at the LHC. A general-purpose detector, ATLAS is designed to study a broad swath of phenomena, from precision measurements of the SM to searches for BSM physics in a wide range of final states. ATLAS is a so-called hermetic detector, aiming to have as close to a perfect solid angle acceptance of 4π steradians as possible.

The integrated luminosity recorded in Run 2 by the ATLAS detector is shown as a function of time in Figure 2.2. Although the LHC delivered over 150fb^{-1} of collision events during this time, some of this data is not usable, for various reasons such as detector downtime or beam quality issues. The amount of Run 2 ATLAS data deemed "good for physics" is around 140fb^{-1} .




Figure 2.1.: The LHC injection chain. Also depicted on the schematic are the other physics experiments being filled by this interconnected system of accelerators [50].




Figure 2.2.: LHC delivered and ATLAS recorded Run 2 total integrated luminosity [51].

2.3.1. The ATLAS Detector Coordinate System

The origin of the ATLAS coordinate system is at the collision point of the LHC beams, also known as the interaction point, at the centre of the detector. Using the standard Cartesian coordinate system, the x -axis is defined to point towards the centre of the LHC ring, the y -axis points straight up, and the z -axis is parallel to the LHC beams in the counterclockwise direction [52]. Spherical coordinates are also utilised, where the azimuthal angle ϕ lies in the x,y coordinate plane, and the polar angle θ gives the angle of the position vector to the z -axis. Instead of the polar angle θ , it is common to use the so-called *pseudorapidity*

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

Position vectors perpendicular to the beam direction have $\eta = 0$. Pseudorapidity grows exponentially with decreasing θ , becoming infinite when the position vector is parallel to the LHC beamline, at $\theta = 0$. One reason for this convention is that differences in pseudorapidity $\eta_1 - \eta_2$ are approximately invariant under Lorentz boosts along the beam direction, for ultra-relativistic particles.

The x,y -coordinate plane is referred to as the transverse plane, because any vector in this plane is normal to the LHC beams. The transverse plane is important in the collider experiments. In proton collisions it is not the protons themselves that collide. Rather, it is the partons that make up the colliding protons that interact. This has the effect that we cannot know the momentum of the colliding particles, since the partons only carry a fraction of the total proton momentum. What we do know, however, is that the momentum of the colliding partons is zero in the transverse plane. This means that the sum of the transverse momenta of the final state particles, the particles being created in the collision, must also be zero in the transverse plane. The presence of particles invisible to the detector, such as neutrinos, in the final state of the proton collisions is inferred by observing a deficit in the sum of the transverse momenta of all the final state particles. In the SM, only neutrinos manifest in the detector as missing energy, but many BSM models predict feebly interacting particles with missing energy signatures similar to that of neutrinos.

ΔR is an often used parameter in ATLAS, defined as angular distance in the space of the azimuthal angle and pseudorapidity: $\Delta R = \sqrt{\Delta\eta^2 + \Delta\phi^2}$.

2.3.2. ATLAS Detector Layout

Like most modern particle detector experiments, ATLAS consists of several subdetector systems. These systems are tasked with measuring different properties of the particles created in the proton collisions delivered by the LHC in the centre of the detector. The subdetectors are arranged in layers around the interaction point, and are generally arranged in such a way as to minimise the amount of material having to be traversed by the incoming particles: particles being deflected or absorbed by the material in the detector leads to measurement inefficiencies and inaccuracies. A three-dimensional schematic of the ATLAS detector can be seen in Figure 2.3. The design resolution and pseudorapidity range of the ATLAS detector subsystems are summarised in Table 2.2.

ATLAS can be divided into three main components: the cylindrical central barrel, covering a pseudorapidity range of around $|\eta| < 1.37$, bookended by two end-caps, extending the pseudorapidity range of the detector to around $|\eta| < 2.5$. The transition region between the barrel and end-caps is sometimes called the *crack region*. For the electromagnetic

Table 2.2.: Design resolution and pseudorapidity coverage of the ATLAS subdetector systems [53].

<i>Subdetector</i>	<i>Resolution</i>	<i>Coverage</i>
Inner Detector	$\sigma_{p_T}/p_T = 0.05\% \times p_T/(1 \text{ GeV}) \oplus 1\%$	$ \eta < 2.5$
EM Calorimeter	$\sigma_E/E = 10\%/\sqrt{E/(1 \text{ GeV})} \oplus 0.7\%$	$ \eta < 3.2$ (trigger: $ \eta < 2.5$)
Hadronic Calorimeter		
Barrel and End-Cap	$\sigma_E/E = 50\%/\sqrt{E/(1 \text{ GeV})} \oplus 3\%$	$ \eta < 3.2$
Forward Region	$\sigma_E/E = 100\%/\sqrt{E/(1 \text{ GeV})} \oplus 10\%$	$3.1 < \eta < 4.9$
Muon Spectrometer	$\sigma_{p_T}/p_T = 10\%$ at $p_T = 1 \text{ TeV}$	$ \eta < 2.7$ (trigger: $ \eta < 2.4$)

calorimeter, described below, this region occurs around $1.37 < |\eta| < 1.52$. This region of the detector suffers lower detection efficiency, due to the small gaps between the barrel and end-cap subdetectors.




Figure 2.3.: A three-dimensional schematic of the ATLAS detector. A wedge shaped piece is cut away to reveal the inner layers of the detector [54].

Inner Detector

The three innermost detector systems of the ATLAS detector, closest to the interaction point, are the Pixel Detector, the SemiConductor Tracker (SCT) and the Transition Radiation Tracker (TRT). These detector systems are collectively known as the *Inner Detector*. The main purpose of the Inner Detector is to track the path of charged particles as they flow from the interaction point, and use these so-called tracks to determine the vertices, meaning the collision points, of the proton collisions. The point in space where two protons from the LHC collide is known as a primary vertex. Any vertices stemming from the subsequent decay of particles produced in these central interactions are known as secondary vertices. Schematics of the inner detector can be seen in Figures 2.4 and 2.5.

The innermost layer of the ATLAS detector is the *Pixel Detector* [55]. It consists of four layers of silicon-based semiconductor sensors segmented into small two-dimensional portions called pixels, analogous to the pixels in a digital camera. The four layers are arranged as concentric cylinders in the barrel, and as parallel disks in the two end-caps. The pixels vary in size, but most have dimensions $50 \mu\text{m}$ times $400 \mu\text{m}$. Pixel detectors

give excellent spatial resolution, but they are very costly, which is why the Pixel Detector only makes up the innermost part of the ATLAS detector. In total there are some 80 million pixels in the ATLAS Pixel Detector, covering a total area of around 1.7 m^2 .

In 2014, during the shutdown of the LHC between Run 1 and Run 2, a new smaller radius beam pipe was fitted in the ATLAS detector. This allowed for the installation of a new barrel layer of pixel detectors, the so-called Insertable B-Layer (IBL) [56]. This new layer of pixel sensors is located very close to the interaction point, at an average radius of only 33 mm. This proximity to the interaction point gives good resolution when placing vertices. Pinpointing vertices is in turn important for identifying, or *tagging*, particle showers originating from bottom quarks. Bottom quark hadrons have lifetimes of order picoseconds. This is short enough for the b-hadrons to decay within the beam pipe after being produced, but long enough for the secondary vertices from their decay to be measurably displaced from the primary vertex.

The layer succeeding the pixel detector is the SCT. Like the pixel detector, the SCT is a silicon-based semiconductor detector. But, instead of consisting of pixels, the SCT sensors are shaped like long thin microstrips. Each strip is $80 \text{ }\mu\text{m}$ wide and 12 cm long. The SCT is made up of four cylindrical layers of microstrip sensors in the barrel of the Inner Detector, and 18 end-cap disk stations, constructed in such a way that any track will intersect at least four stations, giving four spatial measurements [57]. In principle, the silicon microstrip sensors can only determine the position of a track in one dimension, perpendicular to the length of the strip; it is impossible to determine where along its 12 cm length a hit has occurred. To solve this problem, each SCT station consists of two overlapping microstrip layers, mounted back to back, rotated by a 40 mrad stereo angle. This 40 mrad angle, although small, enables the 2D position of each track to be determined at each SCT station. One advantage of having such a small angle separating the two microstrip layers compared to, say, a 90° rotation, is that it reduces the density of so-called ghost hits. Ghost hits are ambiguous hit positions and may stem from simultaneous close proximity hits in the sensor [58]. Furthermore, in the barrel region, both microstrip layers are approximately parallel to the beamline. This maximises resolution in the transverse plane, where charged particles are bent by a magnetic field to measure their momenta.

The outermost subsystem of the Inner Detector is the TRT. The TRT consists of around 300,000 hollow kapton tubes, or "straws". The straws have a diameter of 4 mm. A $31 \text{ }\mu\text{m}$ radius gold plated tungsten wire runs down the centre of each straw. The straws are kept at a -1.5 kV voltage with respect to the grounded central wire. The straws are filled by a mixture of xenon or argon (70%), carbon dioxide (27%) and molecular oxygen (3%) [59]. The straws operate as drift chambers; when ionising radiation traverses the straws, it ionises the gas mixture contained in the straws. The freed electrons drift to the central electrode where an electrical signal is registered. The strong electric field close to the central wire causes the drifting electrons to further ionise the gas, liberating yet more charge in so-called cascades. This effect further amplifies the signal.

In the barrel portion of the Inner Detector, the TRT straws are arranged in 36 layers, parallel to the beam axis. In the Inner Detector end-caps, the straws are mounted radially, perpendicular to the beam axis. The central wires of the straw tubes in the barrel section are separated in the middle by a section of glass wire and each end of the straws are read out independently. This is done to better cope with high occupancy rates in the detector. The area between straws is filled with polypropylene. The task of this material is to induce transition radiation in the TRT. This radiation is then later used for particle identification, a topic covered in Section 2.3.4.

The spatial resolution for each measurement in the TRT is low compared to the semiconductor detectors in the Inner Detector. But where the pixel and microstrip detectors typically have three and four measurements per track, respectively, the TRT performs 36 measurements. As a result of this, no individual measurement from the three Inner Detector subsystems dominates the position measurement of the charged tracks.




Figure 2.4.: A radial cross-sectional view of the ATLAS Inner Detector, showing the Pixel Detector and IBL closest to the beamline, followed by the SCT, and finally the TRT [60]. The view in the schematic is from the centre of the barrel, where semiconductor modules and TRT straws are mounted parallel with the beamline.




Figure 2.5.: A plan of the inner detector, seen in the r - z plane [59]. Here, we see that sensor modules are mounted parallel to the beam pipe in the central region of the barrel. Further out in the barrel, sensors are mounted radially. Also visible here is the ATLAS Central Solenoid magnet, which will be covered in the next section.

ATLAS Magnet Systems

The reason for putting a magnet inside a particle detector is to enable momentum measurement of charged particles. A charged particle travelling through a magnetic field is

deflected by the Lorentz force:

$$F = q\mathbf{v} \times \mathbf{B}. \quad (2.7)$$

If the particle is moving perpendicular to the magnetic field, and if it is not acted upon by any other forces, it follows a circular trajectory of radius r . Using the expression for the acceleration of a body in a constant circular motion, $a = v^2/r$, the momentum of the particle can be expressed as

$$p = |q|Br. \quad (2.8)$$

This means that if the strength of the magnetic field and charge of the particle is known, the momentum of the particle can be inferred from measuring the radius of its path.

It is usually only possible to measure a short arc of the circular trajectory of the charged particles. Instead of measuring the radius of the tracks in the detector, it is more common to use the *sagitta* of the tracks through the detector. The sagitta s of an arc is defined as the distance from the centre of the arc to the centre of its baseline L (see Figure 2.6). The sagitta is therefore a measure of how far the curved trajectory is from being a straight line. The shortest distance h from the origin of the arc to L , can be written, using the Pythagorean theorem, as

$$\begin{aligned} h^2 &= r^2 - \frac{L^2}{4}. \\ h &= r\sqrt{1 - \frac{L^2}{4r^2}} \\ &= r\left(1 - \frac{1}{2}\frac{L^2}{4r^2} + \dots\right) \\ &\approx r\left(1 - \frac{L^2}{8r^2}\right), \end{aligned} \quad (2.9)$$

where the final approximation holds when $L \ll r$. The sagitta then becomes

$$s = r - h = \frac{L^2}{8r}. \quad (2.10)$$

Substituting the sagitta for the radius in Equation (2.8) now yields

$$p = \frac{L^2|q|B}{8s}. \quad (2.11)$$

Momentum smearing from particles interacting with detector material, as well as imperfections in the magnetic field of the detector, makes momentum calculation in ATLAS a bit more involved than simply applying Equation (2.11) directly. However, the principle remains the same.

In ATLAS, the bending of charged particles for momentum measurement is done by four separate magnet systems: one solenoid magnet around the inner detector and three toroidal magnets further out from the interaction point. The four magnets are all superconducting, with an operating temperature of 4.6 K [53].

The Central Solenoid surrounds the Inner Detector. It is a 2.3 m diameter, 5.3 m long solenoid magnet with an average field strength of 2 T, parallel to the beam axis [61]. The




Figure 2.6.: A geometric representation of a circular arc of radius r and baseline L . The sagitta s of the arc is the distance from the centre of the arc to the baseline. The height h is the shortest distance from the origin to L .

Central Solenoid is designed to be as thin as possible. This is done to minimise the amount of material in the detector, thereby decreasing the risk of particles scattering off of, or being absorbed by, the detector material. The Central Solenoid has an average thickness of 0.66 radiation lengths. One radiation length in a material is the mean free path traversed in the material by a high-energy photon before annihilating into an electron-positron pair [11].

Further out radially from the Central Solenoid, ATLAS contains three toroidal magnets: one magnet mounted in the barrel portion of the detector, and one in each of the end-caps. Each of these three magnets consists of 8 air core superconducting coils, arranged in a discrete azimuthal symmetry around the beam axis (see Figure 2.7). This configuration of the magnetic coils results in a magnetic field tangential to a circle in the transverse plane of the detector. The barrel toroid and end-cap toroids have an average magnetic flux density of 0.5 T and 1 T, respectively. The task of the toroidal magnets is to provide a magnetic field to the Muon Spectrometer, which is described below.

Calorimeters

In high energy physics, the main task of a *calorimeter* is to measure the energy of a particle by stopping it and absorbing its energy. Calorimeters are also used for position measurement, although they usually lack the spatial resolution of dedicated tracking detectors.

Calorimeters are separated into two main classes: electromagnetic (EM) and hadronic calorimeters. EM calorimeters measure the energy of particles that lose most of their energy through the electromagnetic interaction, mainly electrons and photons. An average high-energy photon will traverse one radiation length of material before annihilating into an electron-positron pair. After traversing another radiation length of detector material each of these particles in turn create a new photon through bremsstrahlung. This means that after traversing n radiation lengths of material, 2^n new particles are created. This cascading process is known as a *shower*, and continues until ionisation becomes the primary mode of energy loss rather than pair production and bremsstrahlung [26]. High-energy hadrons, on the other hand, lose their energy through nuclear interactions in collisions with atomic nuclei. Despite the different mechanisms for energy loss, hadrons travelling through




Figure 2.7.: Schematic of the ATLAS toroid magnet system. The calorimeter system is also shown [53].

matter also create particle showers. Hadronic showers usually take longer to develop than electromagnetic showers and need denser and thicker material. Hadrons typically travel unimpeded through EM calorimeters and hadronic calorimeters are therefore usually placed outside of EM calorimeters.

The ATLAS calorimeters are sampling calorimeters. A sampling calorimeter consists of alternating layers of active and passive material. The passive layers are usually made up of some dense material tasked with braking the incoming particles, inciting cascades of secondary particles. The energy of the incoming particle is inferred by measuring the particle activity in the active layers. Spatial resolution in the calorimeters is achieved by dividing the active calorimeter material into *cells* or *towers* in η , ϕ and r . Each of these cells can then be read out separately.

The ATLAS EM calorimeter is made up of lead absorber plates arranged in an accordion structure. The space between the lead plates, the active layer, is filled with liquid argon. Kapton electrodes measure ionisation in the argon, analogously to the gas-filled straw tubes of the Transition Radiation Tracker. The EM calorimeter is made up of three segments: one barrel section and two end-cap sections. To save material in the detector, the Central Solenoid magnet is housed in the cryostat of the barrel EM calorimeter. A thin liquid argon layer is placed in front of the central solenoid to account for energy lost by the particles traversing the magnet. The thickness of the EM calorimeter is greater than 22 radiation lengths in the barrel and greater than 24 radiation lengths in the end-caps.

The ATLAS hadronic calorimeter is located immediately outside of the EM calorimeter. The hadronic calorimeter is made up of three subsystems: A barrel tile calorimeter ($|\eta| < 1.7$), a liquid argon end-cap calorimeter ($1.5 < |\eta| < 3.2$), and a liquid argon forward calorimeter ($3.1 < |\eta| < 4.9$). The overlap in pseudorapidity between these calorimeter subsystems ensures coverage in the transition regions between them. The tile calorimeter is made up of plastic scintillators sandwiched between steel absorber plates. Fibre optic cables carry the scintillation light from the scintillators to photomultiplier tubes mounted on the outside of the calorimeter. The steel of the tile calorimeter acts as a return yoke

for the Central Solenoid magnetic field. Liquid argon is used instead of plastic in the end-cap and forward calorimeters. This is due to the intrinsic radiation hardness of liquid argon. Copper plates are used as passive material in the end-cap hadronic calorimeters. The copper plates are mounted in a parallel stack, perpendicular to the beam axis. The forward liquid argon calorimeter is built in a rod-in-cylinder geometry embedded in a solid metal matrix. The forward calorimeter is separated into three sections. The innermost section is made from copper material, and is designed for electromagnetic measurements. The outermost two sections are made from tungsten and are intended for hadronic measurements. Hadronic calorimeter thickness is usually measured in *interaction lengths* λ , rather than in radiation lengths. The interaction length of a medium is the mean path length required to reduce the number of relativistic charged particles traversing it to $1/e$. The ATLAS hadronic calorimeter has a thickness of 9.7λ in the barrel section, and 10λ in the end-caps.

Muon Spectrometer

The muon has a lifetime of $2.2 \mu\text{s}$, making it a relatively long-lived particle. Also, at the energies typical in the ATLAS detector, muons are near minimum ionising particles, meaning that their energy loss from interactions with detector material is close to the minimum. Muons are over 200 times heavier than electrons and therefore lose much less energy due to bremsstrahlung. These properties give muons a distinct detector signature. Muons interact with all detector subsystems, but are not absorbed by the calorimeters. For this reason, the ATLAS Muon Spectrometer (MS), tasked with identifying and tracking muons, is the outermost detector subsystem of the ATLAS detector. The MS utilises four different detector types: Monitored drift tubes (MDTs) and Cathode strip chambers (CSCs) are used for precision position measurement of the muon tracks, while resistive plate chambers (RPCs) and thin gap chambers (TGCs) are used for triggering muon events.

The MS detectors are located in and around the barrel and end-cap toroid magnets. The toroidal magnetic field bends the muon trajectories, enabling measurement of the muon momenta, in the same way as in the Inner Detector. The Muon Spectrometer's ability to measure the momentum spectrum of muons is what gives it its name. In addition to the detectors embedded in the toroid magnet systems, two large muon detector "wheels" are situated outside each end-cap of the detector. These wheels can be seen in Figure 2.3.

2.3.3. Triggers and Data Acquisition

When operating with a bunch spacing of 25 ns, the peak bunch crossing frequency of the LHC is 40 MHz. Each proton bunch contains of the order of a hundred billion protons (see Table 2.1). The probability of two individual protons colliding at each bunch crossing is exceedingly small, but due to the high number of protons in each bunch, each bunch crossing usually sees a few tens of hard proton collisions. Having more than one concurrent scattering event happening at a time in the detector is known as *pileup*. The integrated luminosity recorded as a function of the mean number of interactions per bunch crossings is shown, for each year of Run 2, in Figure 2.8. From the figure we see that the mean number of collisions per bunch crossing is just under 34 events.

Given a mean number of interactions per bunch crossing of 34, the event rate in the ATLAS detector is in excess of 1 GHz. Although the storage needed to record a single event is quite modest, on the order of 1 MB, it would be physically impossible to store every




Figure 2.8.: Recoded luminosity in ATLAS during Run 2 as a function of pileup [51].

event for physics analysis. And we would not necessarily want to. Proton collisions are not created equal in the eyes of a physicist; some events are simply not interesting to the searches or measurements being performed by the ATLAS Collaboration. To select only the most interesting collision events and to reduce the event rate down to a manageable size to be read out, ATLAS uses a system of *triggers*. Triggers are hardware and software systems designed to quickly recognise noteworthy events, and mark them for storage and later physics analysis. During Run 1, the ATLAS trigger system consisted of three stages: the *Level 1* trigger, the *Level 2* trigger, and the *Event Filter*. For Run 2, a new two-stage trigger system was implemented, consisting of a hardware level trigger, still known as the Level 1 trigger, and a software trigger, known as the *High-Level* trigger [62]. A diagram of the current ATLAS trigger system is shown in Figure 2.9.

The Level 1 trigger is a hardware trigger using information from the Muon Spectrometer and calorimeters. The RPCs and TGCs in the MS have good timing characteristics. These detectors are used to identify high- p_T muons. High- E_T photons, electrons, and jets, as well as events with large missing transverse energy, are identified using reduced granularity information from the calorimeters. Information from the Level 1 trigger is used to define $\eta - \phi$ Regions of Interest (ROIs) in the detector. While the Level 1 trigger processes an event, the output signals from the detector are stored in pipeline memories. The Level 1 trigger decides to keep or discard an event within $2.5 \mu\text{s}$, and reduces the rate of candidate collision events to around 100 kHz.

The High-Level trigger is a software trigger, housed in a dedicated computer farm at the ATLAS site at CERN. This trigger is seeded by the ROIs defined by the Level 1 trigger, but also has access to the full detector readout, including information from the Inner Detector. The High-Level trigger has enough time to make a decision to perform track and vertex reconstruction, and fully reconstruct events with algorithms similar to those employed in offline analyses, without events accumulating in the queue to be processed. The High-Level trigger also has access to alignment information and magnetic field maps of the detector. Events selected by the High-Level trigger are written to mass storage for later analysis. The output rate of the High-Level trigger is on average 1.2 kHz. This means that about three events in every 100,000 bunch crossings are kept for analysis.




Figure 2.9.: Schematic of the ATLAS data readout and trigger system [62].

2.3.4. Event Reconstruction and Particle Identification

The ATLAS detector does not measure particles directly. Instead, the various subdetectors register electric signals and energy depositions as particles traverse the detector. The existence of physical particles and their properties must be inferred from these signals during event reconstruction. First, charge and energy depositions in the detector are identified as particle hits in the subdetectors. Next, these hits are used to construct objects such as charged tracks, collision vertices, and calorimeter clusters. These objects are then used to construct physical objects like particles and jets. In the final physics analyses performed on the data, physical objects are organised in *events*. An event is a time-independent snapshot of a particle collision, analogous to a long exposure photograph of a fireworks explosion, where one can see both the initial explosion and the tracks of the firework all in the same photo.

Track reconstruction in the Inner Detector takes place in two stages: First, using a so-called *inside-out* algorithm, and then an *outside-in* algorithm [63]. The inside-out method starts by identifying track seeds made up of three hits in different stations of the pixel and SCT detectors. Pixel and SCT hits that loosely fall in line with the curve extended by the seed hits are added to the candidate track. Various quality criteria are used to remove bad track candidates. Overlapping tracks are vetoed against each other using an ambiguity solver algorithm. The candidate tracks are extended into the TRT to check if there is a valid set of hits matching the track there. Finally, accepted tracks are fitted utilising all available information.

The outside-in method for track reconstruction uses a Hough transform to identify track segments in the TRT. Tracks already found in the inside-out procedure are removed in this step to improve efficiency. Track candidates are extended towards the interaction point to test whether they match any hits in the silicon subdetectors. The outside-in algorithm can identify tracks from converted photons, which the inside-out algorithm may miss. Photon conversion is explained in Chapter 3.

Track reconstruction in the MS also relies on a Hough transform to reconstruct tracks [64]. In certain areas of the MS, straight-line fits are constructed from hits in individual MDT and CSC stations. These local straight-line segments are then used as seeds for the full track finding algorithm.

When locating the primary collision vertices of an event, only the highest quality charged tracks are used. An iterative vertex finding algorithm is used. The steps of this algorithm are outlined below [65]:

- A seed vertex position is constructed. The x - and y -coordinates of the vertex seed are chosen to be in the centre of the luminous collision region, also called the *beam spot*. The z -coordinate of the vertex seed is the mode of the z -coordinate of closest approach to the beam spot of all the charged tracks used in the vertex finding process.
- The charged tracks and vertex seed are used to find the best vertex position in an iterative fit. Before each iteration of the fit, less compatible tracks are weighted down.
- After the vertex position has been determined, tracks that are incompatible with the vertex by more than seven standard deviations are disassociated with the vertex. A valid vertex candidate must retain at least two charged tracks in this procedure.
- The procedure is repeated. Any charged track not associated with a vertex is used in the next iteration.

Bare quarks and gluons cannot be observed in nature due to confinement. When quarks, anti-quarks and gluons are produced in a collider, more quarks and anti-quarks are produced, through gluon radiation, to form colour neutral objects. The resulting particle cascades in the detector are known as *jets*. Jets typically manifest in the detector as collimated tracks in the inner detector pointing to broad energy depositions in the electromagnetic and hadronic calorimeters. The main jet identification algorithm used in ATLAS event reconstruction is the so-called anti- k_t algorithm [66]. The inputs for the algorithm are calorimeter depositions, Inner Detector tracks, or a combination of both. Energy depositions in neighbouring calorimeter cells are grouped by a clustering algorithm before being fed into the jet reconstruction algorithm.

Missing transverse momentum E_T^{miss} is reconstructed as the negative vector sum of the transverse momenta of the visible particles and objects in a collision. There are two main contributions to this calculation, *hard terms* and *soft terms*. Hard terms stem from the reconstructed objects in the event: lepton, jets, etc. Soft terms stem from charged tracks or calorimeter activity that are not associated with any of the reconstructed hard objects. Different techniques are used in ATLAS to calculate the soft terms, but the most common algorithm relies on track information.

Having found the hard objects and soft terms in an event, the missing transverse momentum components $E_{x(y)}^{\text{miss}}$ are calculated as [67]

$$E_{x(y)}^{\text{miss}} = - \sum_{\text{hard objects}} p_{x(y)} - \sum_{\text{soft terms}} p_{x(y)}. \quad (2.12)$$

The magnitude of the transverse missing momentum is given by $E_T^{\text{miss}} = \sqrt{E_x^{\text{miss},2} + E_y^{\text{miss},2}}$. The reconstructed momenta going into the E_T^{miss} are largely taken from calorimeter energy measurements. Therefore, E_T^{miss} is often called missing transverse energy. For relativistic particles, energies and momenta are approximately equal.

Particle properties, such as mass and electric charge, lead to the particles interacting differently in the different subdetectors. This is used to distinguish and identify particles in events. Charged particles leave tracks in, and are bent by the magnetic fields of, the Inner Detector and the Muon Spectrometer. Most particles, regardless of electric charge, leave depositions in the calorimeters. Photons and electrons will deposit most of their energy in the electromagnetic calorimeter, while hadrons, like protons and neutrons, will deposit their energy in the hadronic calorimeter. Pions are an exception to this. They are hadrons, but being very light they will deposit their energy in the electromagnetic calorimeters. Muons tend to traverse the whole detector volume, interacting with every subdetector along the way, though not depositing a lot of energy in the calorimeters. Neutrinos are too feebly interacting to be detected, and will therefore manifest as missing transverse energy in events. A schematic of how different particles interact with ATLAS detector material is depicted in Figure 2.10.

Misidentification of particles does occur. To name some examples: photons may interact with detector material, annihilating to a collimated electron-positron pair; calorimeter depositions from neutral particles may falsely be associated with a charged track; or two particle classes may have similar detector signatures, making it difficult to accurately label them in event reconstruction. Identifying charged pions as electrons is an example of the latter.

Electrons and muons are the most important particles in this work. The reconstruction and identification of these objects in ATLAS is covered in more detail in the sections below.




Figure 2.10.: Stylised radial cross-section of the ATLAS detector, with the detector signatures of various common particles [68]. A dotted line in the figure signifies that the particle does not interact with the detector, while a solid line signifies that it does interact.

Electron Reconstruction and Identification

Electron reconstruction in ATLAS is based on matching "clusters" of energy deposition in the electromagnetic calorimeter with tracks in the Inner Detector. Electrons are identified using a likelihood-based method, utilising a set of cluster-shape and track variables [69]. Four different identification *working points* are defined using the likelihood discriminant: *Very Loose*, *Loose*, *Medium* and *Tight*. Looser working points accept more electron candidates, at the cost of higher probability of jets being identified as electrons. This will be discussed in Section 3.4. The efficiencies of accepting a 40 GeV E_T electron are 93%, 88%, and 80% for the Loose, Medium, and Tight working points, respectively [69].

The ratio of an electron's energy and momentum E/p , measured by the calorimeter and tracker, respectively, should be close to 1 due to its low mass. This trait can be used to separate electrons from heavier objects, such as jets. Separating electrons from charged pions is helped by information from the TRT. Transition radiation is produced when relativistic charged particles traverse a boundary between two media of different dielectric constants. The space between the kapton tubes of the TRT is filled with a multi-layer radiator to induce transition radiation. The amount of transition radiation produced by a particle in the TRT radiator grows with the particle's Lorentz factor γ , with a turn-on of γ between 10^3 and 10^4 . At the same momentum, electrons will have a larger Lorentz factor than charged pions. The increased photon activity of electrons in the TRT can be measured, and used to discriminate against pions. The TRT has a pion rejection efficiency of $>90\%$ [70].

Muon Reconstruction and Identification

There are several methods employed in ATLAS to reconstruct and identify muons. A "stand-alone muon" is reconstructed from a reconstructed track from the Muon Spectrometer alone. A "combined muon" is reconstructed from an Inner Detector track combined

and refitted with a track in the MS. A "segment-tagged muon" is like a combined muon, but only requiring a partial track from the MS. Similarly, "calorimeter-tagged muons" are identified by extrapolating an Inner Detector track into the calorimeters, and finding a deposit consistent with a minimum ionising particle. However, unlike combined muons, the tracks of segment-tagged and calorimeter-tagged muons are not refitted. Instead, the kinematics of these muons are taken from their Inner Detector tracks.

The High- p_T working point requires candidate muons to be combined muons, with registered "precision hits" in at least three stations of the Muon Spectrometer [71]. Here, a precision hit is defined as one registered in one of the two high-precision tracking systems of the Muon Spectrometer, namely the MDT or CSC detector systems. This requirement increases the muon invariant mass resolution at high p_T , at a cost of muon selection efficiency. The High- p_T working point has been optimised for high invariant mass searches such as the searches carried out in this work.

3. Dilepton Standard Model Backgrounds

When considering a single particle collision event on its own, we can never be sure whether the particles we are observing stem from some new physics process, from a known SM interaction, or simply from measurement errors. It is only by statistical analysis, comparing detector measurements against SM predictions that we claim discoveries. Events that mimic the signal being searched for in an analysis are called *background events*. Background events can be due to SM processes with the same final state as the proposed new physics signal, or from detector effects such as particle misidentification.

To increase the chance of discovery of a potential new physics signal, particle physicists try to select events such as to maximise the signal-to-background ratio of their analyses. Backgrounds that can be removed this way are known as reducible backgrounds. However, some background processes will have identical signatures to the signals being tested. Such backgrounds cannot be removed by making clever cuts and selections, and are called irreducible backgrounds.

A Monte Carlo (MC) algorithm is any mathematical method that relies on random number sampling to solve a problem. Particle scattering processes are inherently stochastic; one can never calculate exact scattering angles or final state momenta of single elementary particle interactions, only the probability distributions of these properties. MC generators are used to create large sets of simulated particle collision events from SM calculations of differential cross-sections to some order of perturbation theory, usually leading order (LO) or next-to-leading order (NLO). These MC samples can then be compared to data recorded by detectors to test SM predictions. Several event generator software packages exist on the market, each with its strengths and weaknesses. In addition to simulating hard scattering processes, parton shower generators can also handle hadronisation and other non-perturbative effects.

If we know the momentum of two colliding quarks or gluons, we can in principle calculate the cross-section of some SM process to arbitrary orders of perturbation theory. However, in proton collisions, like at the LHC, the momentum distributions of the partons inside the protons cannot be calculated from first principles. These properties must instead be measured in experiments. The parton density function (PDF) $f_i(x, Q^2)$ is given as the number density of partons of flavour i and *momentum fraction* x in a proton. In the Breit reference frame, also known as the "brick wall" frame¹, x is the fraction of the total proton momentum carried by the parton. Q^2 is the square of the momentum transfer in the collision. PDFs are Q^2 -dependent. In the field of microscopy, optical resolution grows with decreasing wavelength. Smaller length scales can be probed by increasing photon energy. The same principle holds in deep inelastic proton collisions. At low Q^2 , meaning low scattering energies, only the dominant proton structures are visible, and the PDFs are dominated by the valence quarks of the proton. At higher momentum transfers, on the

¹The Breit frame is boosted such that in a collision between particles A and B , the momentum of particle A is reversed as if bouncing off a brick wall.

other hand, the proton resolution increases, and the sea quarks and gluons become more prominent.

PDF sets may vary from one provider to the next based on the data set, theory assumptions, and error treatment used to infer it. The order of perturbation used in the theory calculations will also affect the results. The uncertainty of PDF values tends to grow as a function of x , making it an important systematic uncertainty of high-mass searches such as in this thesis.

It is not feasible to calculate SM cross-sections to arbitrarily high orders of perturbation theory. With each added term in the expansion of the coupling α , the complexity of the calculation grows exponentially. Most event generators calculate cross-sections to LO or NLO. Sometimes, however, higher-order terms contribute non-negligibly to the cross-section being calculated. In this case, we can apply so-called *K-factors*, short for knowledge-factors, to each event to compensate for higher-order loop corrections. K-factors are dependent on the choice of PDF set used by the generator.

Before a set of MC events can be compared to real data, measured in a detector, the response of the detector itself must be simulated, modelling how the generated particles interact with the detector material. In ATLAS, a complete computer-simulated model of the detector is constructed using the Geant4 toolkit [72]. Using this model, each particle created by a MC generator is propagated through the detector, modelling its interaction with the detector material. These interactions model desired particle-detector interactions, such as calorimeter showers or hits in the semiconductor detector material, and also undesired effects and inefficiencies in the detector, such as particles being deflected or absorbed by passive detector material. An example of such a process is *photon conversion*, in which a photon reacts with detector material, and annihilates into an electron-positron pair before reaching the electromagnetic calorimeter. Therefore, what should have been recorded as a photon in the reconstruction of the event, instead looks like a dielectron event. In the cases where both converted electrons are reconstructed, their invariant mass will be close to zero, because the mass of the photon is zero. Such cases are of no concern to our analyses. However, sometimes only one of the electrons is reconstructed. The second electron may for example be too soft to be reconstructed. This single reconstructed electron may then be paired with an electron from another process, such as a pileup event, another converted photon, or a from a decaying W boson, to form a high invariant mass dielectron pair.

The detector simulation step described above models the energy deposition in the detector in a simulated particle collision event. Next in the ATLAS MC production chain comes the so-called *digitisation* step. Here, the energy deposits from the Geant4 simulation is converted into detector response information, in the same form as the raw data measured by the detector when taking real data. This detector response information now goes through the same event *reconstruction* as real data, where the identity and kinematics of physical objects are inferred.

Performing the full Geant4 detector simulation as described above is very costly in terms of computing resources. Modelling calorimeter response is particularly time consuming. Processing one event in this way takes of the order of ten minutes. Such time constraints limit the number of background events that can be produced to model a background process, increasing the statistical fluctuations of background samples. Therefore, methods exist that model detector energy deposition without going through the complete Geant4 detector simulation, to speed up the MC production. One such tool is Atlfast-II, which models tracking in the Inner Detector and Muon Spectrometer using Geant4, while using

a parametric description of calorimeter response [73]. Another technique for simulating detector response uses analytic parameterisations to model detector resolution. Not just in the calorimeter like Atlfast-II, but in the whole detector. This method, which in this work will be referred to as transfer functions (TFs), is used to create sufficiently large MC samples used in the analyses of this thesis. TFs are covered in more detail in Section 3.6

This chapter covers the most relevant background processes for the resonant and non-resonant dilepton searches performed in this thesis, and the production of MC samples for said processes. These MC samples are not used to estimate the SM background in the analyses, but rather to validate and optimise the data-driven background estimation, which is based on parametric fits to the data. MC samples are also used to estimate the spurious signal uncertainty associated with the parametric fits employed in the analyses.

3.1. Drell-Yan

The Drell-Yan (DY) process is an electroweak interaction in which a quark-anti-quark pair from colliding protons annihilate into a virtual photon or Z^0 boson, subsequently decaying into a fermion-anti-fermion pair. The case where these fermion pairs are electron-positron or muon-anti-muon pairs, is the dominant background process for high invariant mass dilepton searches for new physics at the LHC. The tree-level Feynman diagram of the DY process is shown in Figure 3.1. While the DY dilepton invariant mass spectrum has a sharp peak at 91.2 GeV, the Z^0 boson mass, DY events are created at any invariant mass. Therefore, DY is an irreducible background.




Figure 3.1: Feynman diagram of the DY process.

In this work, DY samples are generated using the POWHEG-BOX [74] event generator, applying the CT10NLO [75] PDF set. The POWHEG-BOX generator calculates the DY cross-section to NLO in QCD. A mass-dependent K-factor is applied to the DY samples to account for higher-order loop effects from NLO electroweak and next-to-next-to-leading order (NNLO) QCD. QCD corrections are calculated using VRAP 0.9 [76], while MCSANC [77] is used to calculate electroweak corrections. The CT14 NNLO PDF set [78] is used in both cases. Final state photon radiation is modelled by the PHOTOS [79] event generator. The DY samples are generated in separate ranges in dilepton invariant mass. This is done to enhance production, and thus lower statistical fluctuations, in the high invariant mass region.

In our analyses, we aim for the statistical uncertainty on the MC background to be as small as possible. In a previous analysis published by the ATLAS exotic dilepton working group, using 36 fb^{-1} of 2015 and 2016 LHC data, the desired level of MC statistical uncertainty was set to be no larger than 10% of the statistical uncertainty in the data [10]. Due to the computational cost of generating MC events, this requirement was not met. The problem of generating large enough MC samples is even more challenging when analysing the full LHC Run 2 data set, corresponding to 139 fb^{-1} of data. This limitation is one of

the main reasons the ATLAS exotics dilepton working group opted to use a data-driven method for estimating the background over a MC simulation-based background estimation.

No conventional ATLAS detector simulation is performed on the DY MC samples. Instead, TFs are used to model detector effects on the dilepton invariant mass resolution. The TF method is detailed in Section 3.6. Skipping the computationally intensive detector simulation step allows for more events to be generated. The number of generated DY events is equivalent to at least $8,000 \text{ fb}^{-1}$ of LHC data.

3.2. Top Quark

The top quark is, with its 173 GeV mass, the heaviest known elementary particle. It is maybe not surprising then, that top quark signatures show up as background in high-mass analyses. SM processes producing a top quark, or a pair of top quarks, are the second most important background category in our dilepton searches. When produced, the top quark decays predominantly into a W^+ boson and a bottom quark. In the cases where these W^+ bosons decay leptonically, top quark events may show up as background in lepton analyses. The branching ratios of leptonic W boson decays are about 11% per lepton generation. Both the top quark and the W bosons have mean lifetimes of the order of 10^{-25} s, meaning that they do not travel any significant distance before decaying. This means that any lepton tracks stemming from top events originate in the primary vertex of the event.

The top background is separated into two main categories: single top quark production and top-anti-top quark pair production ($t\bar{t}$). $t\bar{t}$ events are produced predominantly through the strong interaction, either through quark-anti-quark annihilation or gluon-gluon fusion. These two production modes are shown as tree-level Feynman diagrams in Figure 3.2.




Figure 3.2.: Three example top-anti-top quark pair production mechanisms: a gluon-gluon fusion process (a), a quark scattering processes (b), and a t-channel top-quark exchange process (c).

Single-top quark production, on the other hand, is only possible through weak interactions. As a result, single-top quark events are less common than $t\bar{t}$ events. In these events, a top or anti-top quark is produced in association with a down- or anti-down-type quark, through the exchange of a W boson; or vice versa, in association with a W boson through the exchange of a down- or anti-down-type quark. Example tree-level Feynman diagrams for down-type associated top quark production and Wt production are shown in Figures 3.3 and 3.4, respectively.

The top quark background MC samples used in this work were generated using the POWHEG-BOX event generator, calculated at NLO and applying the NNPDF3.0 [80] PDF set. The MADSPIN [81] event generator is used to account for the effect of top quark




Figure 3.3.: Two production mechanisms for single-top quark production in association with a bottom quark, through the exchange of a W boson: an s-channel process (a) and a t-channel process (b).




Figure 3.4.: Two production mechanisms for single-top quark production in association with a W boson, through quark exchange: an s-channel process (a) and a t-channel process (b).

spin correlation. For the single top quark samples, ATLAS detector resolution effects are modelled using the full Geant4 detector simulation. TFs are used to model detector resolution in the $t\bar{t}$ samples. A static K-factor, calculated using Top++ 2.0 [82], is applied to the top quark background samples to normalise to NNLO QCD predictions.

3.3. Diboson

Two weak bosons, a ZZ -, WZ -, or WW -boson pair, may be produced in the proton-proton collisions at the LHC. These bosons may then decay leptonically, sometimes resulting in two or more electrons or muons. Such events form the so-called *diboson* background category, the third most important background to the dilepton searches presented here. Two LO diboson production Feynman diagrams are shown in Figure 3.5.

The SHERPA [83] event generator, applying the CT10 NLO PDF set, is used to produce the diboson MC samples. Diboson cross-sections are calculated to NLO. The samples are generated in separate ranges of dilepton invariant mass. This procedure, identical to the one used in the DY production, is performed to increase the size of the diboson MC sample in the high dilepton invariant mass region.

Not all events are created equally. Event generators assign a weight to each event. The event weights of most events in the diboson MC sample are of the order 1, but a few have event weights of the order 10^6 . Events with event weights higher than 10^3 are removed from the sample. This does not change the overall shape of the dilepton invariant mass distribution significantly, but removes some unphysical single bin peaks.




Figure 3.5.: Example diboson production mechanisms. Drell-Yan production of a W boson pair (a), and t-channel production of a Z or W boson pair through quark exchange (b).

3.4. Fake Electrons

As previously mentioned, the ATLAS detector does not detect particles directly. Rather, it registers hits and energy depositions in its detector subsystems. Physical objects are identified later, during event reconstruction. From time to time, the reconstruction algorithms fail to correctly identify a particle. Jets, for example, are sometimes identified as leptons. Such mislabelled jets are called *fake leptons* or *fakes*, colloquially. The electron misidentification rate in ATLAS is high enough for fake electrons and positrons to show up as background in analyses. In this section, electrons and positrons are both referred to as electrons.

We use the *matrix method* to estimate the fake electron background. We cannot measure the number of fake electrons in our data directly. What we can count is how many electrons pass the event selection requirements of our analysis. These requirements are listed in Section 4.2.1 and will be referred to here as the "tight" selection. We can also loosen these requirements, which will accept more fake electrons into the selected data set. In our case, a "loose" selection is defined by loosening the electron identification requirement, and removing the electron isolation requirement.

The number of electrons passing the loose selection is written here as N_L . The subset of N_L that also pass the tight selection requirement is written N_T . The number of electrons passing the loose selection, but failing the tight selection, is denoted as $N_{\bar{T}}$.

The *real efficiency* r and the *fake efficiency* f are defined as

$$r = \frac{N_T^{\text{real}}}{N_L^{\text{real}}}, \quad f = \frac{N_T^{\text{fake}}}{N_L^{\text{fake}}}. \quad (3.1)$$

The real efficiency is measured as the fraction of real electrons that, after having passed the loose selection, also pass the tight selection. The fake efficiency is similarly defined for fake electrons. We measure r using MC simulations, while f is measured using data in a fake-enhanced control region. Both r and f are functions of transverse momentum and pseudorapidity of the fake-candidate electron.

The real and fake efficiencies can be used to build a set of equations relating the number of real and fake electrons in our data to the number of electrons passing or failing the tight selection requirements:

$$\begin{pmatrix} N_{TT} \\ N_{T\bar{T}} \\ N_{\bar{T}T} \\ N_{\bar{T}\bar{T}} \end{pmatrix} = M \begin{pmatrix} N_{RR} \\ N_{RF} \\ N_{FR} \\ N_{FF} \end{pmatrix}, \quad (3.2)$$

where

$$M = \begin{pmatrix} r_1 r_2 & r_1 f_2 & f_1 r_2 & f_1 f_2 \\ r_1(1-r_2) & r_1(1-f_2) & f_1(1-r_2) & f_1(1-f_2) \\ (1-r_1)r_2 & (1-r_1)f_2 & (1-f_1)r_2 & (1-f_1)f_2 \\ (1-r_1)(1-r_2) & (1-r_1)(1-f_2) & (1-f_1)(1-r_2) & (1-f_1)(1-f_2) \end{pmatrix}. \quad (3.3)$$

Here, we are looking at the two-lepton case. The first index of N in these equations, as well as the efficiencies r_1 and f_1 , are associated with the leading transverse energy electron. The second index of N and the efficiencies r_2 and f_2 are associated with the subleading transverse energy electron. The indices R and F indicate whether the electrons are real or fake, respectively.

We want to know the number of electron pairs passing the tight selection, where one or both electrons are fake. This is given by the equation

$$N_{\text{TT}}^{\text{fake}} = r_1 f_2 N_{\text{RF}} + f_1 r_2 N_{\text{FR}} + f_1 f_2 N_{\text{FF}}. \quad (3.4)$$

The numbers N_{RF} , N_{FR} , and N_{FF} are not measurable. However, these numbers can be calculated in terms of N_{TT} , $N_{\text{T}\bar{\text{T}}}$, $N_{\bar{\text{T}}\text{T}}$, and $N_{\bar{\text{T}}\bar{\text{T}}}$ by inverting the matrix M . This is how the matrix method gets its name.

The above procedure is sometimes referred to as the standard matrix method. This method has its limitations. The estimated number of fake electrons may sometimes become negative, which is hard to interpret physically. The method can also become numerically unstable if r and f take similar values. To circumvent these issues we use a modified version of the matrix method, known as the *likelihood matrix method* [84]. Rather than inverting the matrix M , the likelihood matrix method finds the number of fake electrons using a maximum likelihood fit. The likelihood function used,

$$L = \text{Pois}(N_{\text{TT}}, N_{\text{TT}}^{\text{pred}}) \text{Pois}(N_{\text{T}\bar{\text{T}}}, N_{\text{T}\bar{\text{T}}}^{\text{pred}}) \text{Pois}(N_{\bar{\text{T}}\text{T}}, N_{\bar{\text{T}}\text{T}}^{\text{pred}}) \text{Pois}(N_{\bar{\text{T}}\bar{\text{T}}}, N_{\bar{\text{T}}\bar{\text{T}}}^{\text{pred}}), \quad (3.5)$$

contains four Poisson terms, $\text{Pois}()$. The values $N_{\text{TT}}^{\text{pred}}$, $N_{\text{T}\bar{\text{T}}}^{\text{pred}}$, $N_{\bar{\text{T}}\text{T}}^{\text{pred}}$, and $N_{\bar{\text{T}}\bar{\text{T}}}^{\text{pred}}$ are calculated as functions of N_{RR} , N_{RF} , N_{FR} , and N_{FF} in Equation (3.2). The values of N_{RR} , N_{RF} , N_{FR} , and N_{FF} are chosen such as to maximise the likelihood function.

The fake electron invariant mass distribution is smoothed using a parametric function. The function found to best represent the shape of the fake electron background is given by

$$y(x) = (1 - x^{1/3})^{p_1} x^{p_2 \log(x)}, \quad (3.6)$$

where $x = m/\sqrt{s}$ is the dielectron invariant mass divided by the collision energy.

The fake electron background distribution obtained using the method described here, using the likelihood matrix method and smoothing using a functional fit, was only used in the non-resonant dilepton analysis described in Chapter 6. For the resonant dilepton analysis described in Chapter 5, the fake electron background sample used was the same as that obtained in Ref. [10], scaled up to the Run 2 luminosity. This sample was generated using the standard matrix method. A functional fit to the fake electron invariant mass distribution was used to accurately model the high and low invariant mass range.

A muon is identified in ATLAS by matching a track in the Inner Detector with one in the Muon Spectrometer. The misidentification rate of muons is much smaller than that of electrons. The contribution from fake muons to the dimuon background, at the high invariant masses considered in this work, was found to be negligible [85].

3.5. Monte Carlo Background Samples

For the background MC samples using the full Geant4-based detector simulation, the effect of pileup is taken into account in the simulation. Dedicated simulations of pileup is generated using PYTHIA 8 [86] with the NNPDF2.3 [87] LO PDF set. These pileup samples are weighted to replicate the average number of pileup events during data-taking. For Run 2, the average number of collisions per bunch crossing in ATLAS was around 34 [51]. The generated pileup events are passed through the Geant4 detector simulation. The fully simulated pileup samples are then added to the hard scattering samples to simulate real Run 2 data-taking conditions in ATLAS.

In addition to the aforementioned pileup re-weighting procedure, the events in the MC samples are multiplied by weights to account for observed discrepancies between data and MC simulation. Lepton calibration scale factors account for discrepancies between data and simulation in lepton reconstruction, and lepton trigger-, identification-, and isolation efficiencies. Calibration corrections are also added to the scale and resolution of electron energies and muon momenta.

For every MC simulated background category except the diboson background, the PYTHIA 8 and EVTGEN [88] event generators are used to model hadronisation and other non-perturbative effects after the central hard scattering process. EVTGEN is used for charm- and bottom quark decays, specifically.

When using the SHERPA event generator, there is no need for a dedicated treatment of non-perturbative effects. SHERPA handles the full event generation, from hard scattering down to final state particles.

A summary of the event generators and PDF sets used to generate the MC background samples for this work is shown in Table 3.1.

Table 3.1.: List of MC event generators used in the production of background samples.

<i>Process</i>	<i>Hard Scattering Process with PDFs</i>	<i>Parton Shower, Non-Perturbative Effects with PDFs</i>
Drell-Yan	POWHEG-BOX [74] [89], CT10NLO [75], PHOTOS [79]	PYTHIA v8.186 [90], CTEQ6L1 [91] [92], EVTGEN 1.2.0 [88]
tt	POWHEG-BOX, NNPDF3.0NLO [80]	PYTHIA v8.230 [86], NNPDF23LO [87], EVTGEN 1.6.0
Wt	POWHEG-BOX, NNPDF3.0NLO	PYTHIA v8.230, NNPDF23LO, EVTGEN 1.6.0
Single t	POWHEG-BOX, NNPDF3.04FNLO, MADSPIN [81]	PYTHIA v8.230, NNPDF23LO, EVTGEN 1.6.0
Diboson	SHERPA 2.1.1 [83], CT10NLO	SHERPA 2.1.1, CT10

3.6. Transfer Functions

Transfer functions (TFs) are functional parameterisations of the dilepton invariant mass resolution. The TFs determine the probability density of finding a dilepton event of reconstructed invariant mass m , given a generated invariant mass of m_t . The transfer functions are parameterised with a Gaussian function convoluted with a Crystal Ball function, totalling 7 free parameters. This function fits the invariant mass detector resolution simultaneously in 200 different NLO Drell-Yan MC samples. These MC samples are generated at fixed points of dilepton invariant mass, spaced logarithmically between 130 GeV and 6,000 GeV. The detector resolution in the MC samples is modelled using the full Geant4 detector simulation. To obtain the reconstructed invariant mass distribution $S(m)$ of a generator-level invariant mass spectrum $S_t(m_t)$, the generator-level spectrum is first multiplied with the function of dilepton acceptance and efficiency in the detector before being convoluted with the TFs. More information on the TF method can be found in Ref. [93].

The TFs are used to model the detector resolution of the Drell-Yan and $t\bar{t}$ MC simulated backgrounds. TFs are also used to make a generic signal model shape for the resonant dilepton analysis; If the width of the resonant signal is less than the detector mass resolution, a so-called zero-width signal, the signal shape can be described by the TFs. For signal hypotheses of non-negligible width, the TFs must also be convoluted with the truth-level signal shape.

Once the TFs have been constrained, the dilepton invariant mass resolution is obtained by adding the width parameters of the Gaussian and Crystal Ball functions, with their respective weights, in quadrature. The relative dilepton invariant mass resolutions, σ_m/m , in both the electron and muon channels are shown in Figure 3.6. The muon resolution is seen to degrade with increasing invariant mass, while the electron resolution is seen to improve. The momentum of a muon is inferred from the curvature of its track in the detector. Higher momentum muons are bent less by the magnetic field of the detector, and as a result have straighter tracks. The radius of a near straight track is harder to measure than a sharper curve. Therefore, muon momentum is seen to degrade with increasing momentum. This effect is reflected in the invariant mass resolution, as high invariant mass dilepton events inevitably involve higher momentum muons. The dielectron mass resolution, on the other hand, is driven by the energy resolution of the electromagnetic calorimeter of the detector. This resolution improves as electron momentum increases.




Figure 3.6.: Relative invariant mass dilepton resolution, extracted from TFs.

4. Event Selection for Dilepton Analyses

4.1. Data Set

As mentioned in Section 2.2, the analyses of this work use the full pp collision data set collected in ATLAS during the Run 2 data-taking period of the LHC. The amount of integrated luminosity recorded by the ATLAS detector in each year of Run 2 is shown in Table 4.1. The total luminosity recorded is measured to be $139.0 \pm 2.4 \text{ fb}^{-1}$. The luminosity is determined using a method similar to that detailed in Ref. [94], using measurements from the LUCID-2 detector [95].

Table 4.1.: Total integrated luminosity collected each year in ATLAS during Run 2.

<i>Year</i>	$\int L dt \text{ [fb}^{-1}\text{]}$
2015	3.2
2016	33.0
2017	44.3
2018	58.5
Total	139.0

4.2. Object Selection

The only physical objects of interest in the searches carried out in this thesis are electrons and positrons, and muons and anti-muons. To stay sensitive to as wide a range of new physics signals as possible, we make as few assumptions as possible on the properties of the events selected for the analysis; no selection is made on non-leptonic activity in the events, such as the missing transverse energy or jet activity. We call such an event selection strategy an *inclusive* dilepton selection.

An event is required to have at least one reconstructed proton-proton interaction vertex. If more than one vertex is reconstructed in the event, the primary vertex is defined as the one where the squared sum of the transverse momenta of all tracks with transverse momentum over 0.5 GeV is the largest. All electron and muon candidates must originate from the primary vertex. Deciding whether or not a track is consistent with the primary vertex is done by imposing requirements on the *impact parameters*, d_0 and z_0 , of the lepton tracks.

Reconstructed primary vertices in ATLAS are positioned on the beamline, or z -coordinate axis of the detector, by construction. This is justified by the fact that the beam spot is much smaller than the vertex position resolution. The *longitudinal impact parameter* z_0 ,

is the z -coordinate of the closest approach of a track to the beamline. The difference in z between z_0 and the primary vertex position is denoted as Δz_0 . The shortest distance from the track to the primary vertex is then $|\Delta z_0 \sin \theta|$. The *transverse impact parameter* d_0 is defined as the distance of closest approach of a track to the beamline in the transverse plane of the detector. The d_0 significance $d_0(\sigma)$ is defined as the length of d_0 expressed in terms of its measurement uncertainty. For our analyses, electron and muon candidates are required to satisfy $|\Delta z_0 \sin \theta| < 0.5$ mm. Electron (muon) candidates must also satisfy $d_0(\sigma) < 5$ (3).

4.2.1. Electrons

The main electron selection uses the Medium electron identification working point. The LooseAndBLayer identification is also used when estimating the fake electron background using the matrix method, as described in Section 3.4. The LooseAndBLayer working point is identical to the Loose working point, but requires the reconstructed electron track to have a registered hit in the Insertable B-Layer in the Inner Detector.

A particle being *isolated*, loosely means that the particle is separated from other particle activity in the detector. Determining the isolation of a particle involves drawing a cone in ΔR around the trajectory of the particle, and requiring that the particle activity in the cone be below some threshold. Electron candidates in our analyses are required to pass the *Gradient* isolation working point, as defined in Ref. [69]. This isolation requirement has an energy-dependent variable threshold on the transverse energy and momentum of calorimeter and track activity in a $\Delta R = 0.2$ cone around the electron candidate. This threshold is designed to yield an isolation efficiency of $\epsilon_{\text{iso}} = (92.14 + 0.1143 \times p_T / (1 \text{ GeV}))\%$.

Electrons are required to have transverse energy E_T greater than 30 GeV, and pseudorapidity less than 2.47. Electrons falling in the transition region between the barrel and end-cap of the electromagnetic calorimeter, $1.37 < |\eta| < 1.52$, are excluded.

The above electron selection criteria are summarised in Table 4.2.

Table 4.2.: Electron object selection criteria.

<i>Feature</i>	<i>Criteria</i>
Pseudorapidity	$ \eta < 2.47$ and not $1.37 < \eta < 1.52$
Transverse Energy	$E_T > 30$ GeV
Vertex Quality	$ d_0(\sigma) < 5$ $ \Delta z_0 \sin \theta < 0.5$ mm
Identification	LooseAndBLayer (matrix method), Medium (main selection)
Isolation	Gradient working point

4.2.2. Muons

Muon candidates are required to pass the High- p_T identification working point, and are required to pass the FixedCutTightTrackOnly isolation working point [96]. This isolation criterion requires the summed p_T of tracks, excluding the muon track itself, in a cone of variable size in ΔR around the muon, to be less than 6% of the muon p_T .

Muons are required to have transverse momentum p_T greater than 30 GeV, and pseudorapidity less than 2.5. The reason for loosening the pseudorapidity requirement somewhat

for muons, compared to the electron selection, and for muons to be accepted in the crack region of the detector, is due to the different methods in which electrons and muons are reconstructed and identified. Electron reconstruction relies on information from the electromagnetic calorimeter, which suffers from decreased efficiency in the transition from the barrel calorimeter to the end-cap calorimeter. Combined muons use calorimeter information to account for energy lost in transit between the Inner Detector and the Muon Spectrometer, but are not critically affected by the crack region calorimeter inefficiency.

The uncertainty of the charge-to-momentum ratio q/p is required to be less than a p_T -dependent threshold. This requirement removes muons from the tail end of the σ_{p_T}/p_T distribution. Muons passing this selection are known as "good muons". The muon selection criteria are summarised in Table 4.3.

Table 4.3.: Muon object selection criteria.

<i>Feature</i>	<i>Criteria</i>
Pseudorapidity	$ \eta < 2.5$
Transverse Momentum	$p_T > 30 \text{ GeV}$
Vertex Quality	$ d_0(\sigma) < 3$ $ \Delta z_0 \sin \theta < 0.5 \text{ mm}$
Identification	High- p_T working point
Isolation	FixedCutTightTrackOnly working point
Quality	Good muons

4.3. Calibration and Corrections

Previously, we described how simulated pileup MC samples are added to the background MC samples to accurately model additional proton-proton collisions taking place, in addition to the triggered hard scattering collision of interest, at each LHC bunch crossing. Usually, MC samples are produced at the same time as, or even before, data are recorded in the detector. The pileup MC samples added to the background are generated using a best estimate of the actual pileup conditions in the data. More or fewer pileup events lead to more or less activity in the detector, which affects detection efficiencies. Using the wrong estimate for the pileup in simulations will lead to mismodelling between the data and the MC samples. In the analyses presented in this thesis, we apply a weight to each MC event to re-weight the pileup conditions in the sample to that of the data.

Additional scale factors are used to account for discrepancies between data and MC samples. For electrons, trigger efficiency-, reconstruction efficiency-, isolation efficiency-, and identification efficiency scale factors are applied on an event-by-event basis. Differences in the electron energy scale and resolution between data and simulations are also corrected. The corrections are independent of electron transverse energy and are derived using the method described in Ref. [97]. Here, the absolute electron energy scale is determined using a large dedicated sample of Z boson decays into electron-positron pairs.

For muons, trigger efficiency-, reconstruction efficiency-, isolation efficiency-, and track-to-vertex association (TTVA) efficiency scale factors are applied. The differences in muon p_T scale and resolution between data and simulations are also corrected. These corrections are calculated as a function of p_T for $p_T < 300 \text{ GeV}$, using the method described in

Ref. [96]. Here, a large sample of J/ψ and Z boson decays to muon-anti-muon pairs are used to determine the scale and resolution of ATLAS muon momentum measurements. For transverse momenta over 300 GeV, the correction factors are determined in muon alignment studies.

4.4. Object Overlap Removal

Sometimes, reconstructed particles in the detector are built using the same detector objects. For example, a muon candidate may share the same Inner Detector track as an electron candidate, or two different clusters in the electromagnetic calorimeter may both be combined with the same Inner Detector track to form two electron objects. For our analyses, we perform "overlap removal", where we veto overlapping objects against each other to resolve such reconstruction ambiguities.

The overlap removal criteria are as follows:

1. If two or more electrons share an Inner Detector track, the higher- E_T electron is kept.
2. If a muon and an electron share an Inner Detector track, the muon is kept.
3. If a lepton and a jet share an Inner Detector track, the lepton is kept. However, such leptons are not likely to survive the isolation selection criteria.

As previously mentioned, only combined muons, with tracks both in the Inner Detector and the Muon Spectrometer, are considered for this work. For a muon and electron sharing the same Inner Detector track, one of two things may have occurred: either the electron "punched through" the calorimeter system to appear as a track in the Muon Spectrometer; or the muon caused sufficient activity in the electromagnetic calorimeter to be identified as an electron before entering the Muon Spectrometer, for example by radiating a photon through bremsstrahlung. The second case is much more likely and is why we keep muons over electrons.

Muons do not need to be vetoed against muons in the same way as electrons against electrons or muons. Such overlap removal occurs explicitly in the muon reconstruction. Therefore, no two reconstructed muon objects share Inner Detector or Muon Spectrometer tracks.

4.5. Event Cleaning

We require stable LHC beams; that all magnets are on and at full magnetic field strength, and that all detector subsystems be operational, during data-taking. To ensure this, the ATLAS data-taking is separated into 1-2 minute "luminosity blocks". Each luminosity block where the above requirements are met is added to a so-called Good Runs List (GRL). Only luminosity blocks in a GRL are accepted for our analyses.

Events recorded during noise bursts in the liquid Argon calorimeter are removed, as are incomplete or corrupted events.

4.6. Triggers

Electron events are recorded using a dielectron trigger, discriminating on electron transverse energy. The trigger threshold in E_T is varied from 12 GeV in the 2015 LHC data-taking period, to 17 GeV in 2016, and finally 24 GeV in 2017 and 2018. The electron trigger accepts events passing the loose identification working point in 2015, and the very loose identification working point in the remaining years of the Run 2 campaign.

Muon events are selected using two different single-muon triggers, both with a lower threshold on muon transverse momentum. One muon trigger only requires a muon with a p_T of at least 50 GeV. The second muon trigger lowers this p_T threshold to 26 GeV, but with an added requirement that the muon is sufficiently isolated.

4.7. Event Selection

Candidate dilepton events in the analyses are required to contain at least two same-generation leptons. We divide the event selection into two orthogonal categories, or *channels*, according to the lepton flavour of the selected dilepton pair: the *electron channel*, containing two electrons and/or positrons, and the *muon channel*, containing a muon-anti-muon pair.

Similarly to the previously described object overlap procedure, events are kept from being counted in both channels. In events with more than two same flavour leptons passing the object selection, the two electrons with the highest transverse energy E_T is kept, or the muon pair with the highest transverse momentum p_T . The most energetic lepton of a particular flavour in an event is known as the *leading* lepton of that particular flavour. Similarly, the subleading lepton is the second most energetic. In events containing a candidate pair of both electrons and muons, the electron pair is kept due to the higher efficiency and resolution for electrons. This last step removes only a tiny fraction of dimuon events.

Two selected muons are required to be of opposite electric charge. Although such a requirement is physically motivated, as the postulated Z' boson is an electrically neutral particle, no such requirement is made in the electron channel. The rate of charge misidentification is large for high- E_T electrons. Imposing an opposite charge requirement will significantly impact the electron selection efficiency, and not increase the signal-to-background ratio significantly.

The invariant mass of the selected dilepton pair can be no less than 130 GeV. Neither the resonant nor the non-resonant analyses covered in this work considered the region in invariant mass below 250 GeV. Nevertheless, this low mass region is still used to validate the shape and normalisation of the MC background samples used in the analyses.

4.8. Data and Monte Carlo Comparison Plots

Figures 4.1 and 4.2 show the dilepton invariant mass distributions of the data and MC simulated background¹ samples after event selection, in the electron and muon channel, respectively. Signal shape templates for four different choices of Z'_χ mass are overlaid in

¹Except for the background due to fake electrons in the electron channel, which is derived using a data-driven method.

the figures. The production of these signal histograms will be described in Section 5.1.1. The MC samples are scaled to 139 fb^{-1} to match the measured luminosity in data.

Bear in mind that the MC templates shown here are not used in any statistical analysis directly. Instead, the MC samples are used to optimise the functional fits used to model the SM background, and to calculate the background modelling uncertainty in the form of spurious signal. The figures shown in this section are not using TFs to model detector resolution in the MC background samples. Instead, detector resolution is modelled using Geant4.




Figure 4.1.: Invariant mass data-to-background comparison in the electron channel. The ratio between the two distributions is shown in the lower panel. Data points are drawn with error bars. The relative statistical uncertainty on the total background template is drawn as a shaded grey band in the lower panel. Red arrows indicate data points falling outside the range of the lower panel.

Figures 4.3 and 4.4, and show the comparison between the data and the MC simulated background for the transverse momentum and pseudorapidity, respectively, of the leading and subleading electrons or positrons in the electron channel. Note that the plots do not include the fake electron background category. Figures 4.5 and 4.6 show the comparison between the data and the MC simulated background for the transverse momentum and pseudorapidity, respectively, of the leading and subleading muon or anti-muon in the muon channel. Only a rough agreement between the data and the background is expected, as no systematic variations are applied to the MC background.

4.9. Highest Dilepton Invariant Mass Event Displays

Figures 4.7 and 4.8 show the so-called ATLAS event displays of the highest invariant mass dielectron and dimuon event in the Run 2 data, respectively. An event display is a graphical representation of a recorded ATLAS event. In these event displays, the barrel section of




Figure 4.2.: Invariant mass data-to-background comparison in the muon channel. The ratio between the two distributions is shown in the lower panel. Data points are drawn with error bars showing their statistical uncertainty. The relative statistical uncertainty on the total background template is drawn as a shaded grey band in the lower panel.




Figure 4.3.: Data-to-background comparison of the transverse momentum of the leading (a) and subleading (b) electrons or positrons in the electron channel.




Figure 4.4.: Data-to-background comparison of the pseudorapidity of the leading (a) and subleading (b) electrons or positrons in the electron channel.




Figure 4.5.: Data-to-background comparison of the transverse momentum of the leading (a) and subleading (b) muon or anti-muon in the muon channel.




Figure 4.6.: Data-to-background comparison of the pseudorapidity of the leading (a) and subleading (b) muon or anti-muon in the muon channel.

the ATLAS detector is projected onto the transverse (x - y) plane. The event displays also contain a view of the full detector, as well as a zoomed in view of the interaction region, in the longitudinal (z - y) plane. The electromagnetic and hadronic calorimeters are marked in green and red, respectively, while the components of the Muon Spectrometer are shown in blue. In the event displays we see tracks emanating from multiple vertices in the interaction region, a nice demonstration of the concept of pileup.




Figure 4.7.: ATLAS event display of the highest invariant mass dielectron event recorded in Run 2, with $m_{ee} = 4.06$ TeV. The leading transverse energy electron has $E_T = 2.01$ TeV, $\eta = 0.47$ and $\phi = -0.78$. The subleading transverse energy electron has $E_T = 1.92$ TeV, $\eta = -0.03$ and $\phi = 2.37$. The trajectories, and deposited calorimeter energy, of the electrons show up in the display as light red lines.




Figure 4.8.: ATLAS event display of the highest invariant mass dimuon event recorded in Run 2, with $m_{\mu\mu} = 2.75$ TeV. The leading transverse momentum muon has $p_T = 1.82$ TeV, $\eta = -0.52$ and $\phi = -0.56$. The subleading transverse momentum muon has $p_T = 1.04$ TeV, $\eta = -0.67$ and $\phi = 2.53$. The trajectories of the muons show up in the display as thin dark red lines. The calorimeter activity overlapping the muon trajectories have $p_T < 50$ GeV, consistent with having been induced by the muons themselves.

5. Resonant Dilepton Analysis

This chapter describes a search for resonant new physics phenomena in dilepton final states. Our analysis strategy is to search for a model-independent signal shape, making no assumptions on the spin, width, or mass of any potential new physics models. The results obtained on this generic model can then be reinterpreted into several benchmark physics models, by this or other searches. The discriminating variable of the analysis is the reconstructed dilepton invariant mass, finely binned to allow for a shape analysis. The search is performed for pole masses of a new resonant signal between 250 GeV and 6 TeV. This invariant mass interval is the "search range" of the analysis. We do not know where in this search range to expect a signal, nor the shape the signal might take. Therefore, we "scan" across signal mass hypotheses in steps of one GeV, and across signal width hypotheses, relative to the Z' pole mass, from 0% to 10% in steps of 0.5%, searching for a signal. The one GeV step size is smaller than the invariant mass resolution of the detector. The results presented in this chapter are published in Ref. [6].

5.1. Modelling Exotic Dilepton Resonances

The search strategy in the resonant analysis is to fit the dilepton invariant mass spectrum with a smooth, exponentially falling function and a generic signal shape. In the cases where a BSM dilepton resonance has a decay width Γ much less than the dilepton invariant mass resolution of the ATLAS detector, the shape of its reconstructed signal will be the same as the shape of the invariant mass detector resolution functions, or transfer functions (TFs), described in Section 3.6. Such signatures are referred to as "zero-width" signals in this text. Zero-width signal shapes for three different values of BSM resonance pole mass are shown in Figure 5.1.

For "wide signals", signatures of non-negligible resonance width, we use a non-relativistic Breit-Wigner function to describe the true underlying shape of the signal invariant mass distribution. This function is given as a function of invariant mass m , as

$$f_{\text{BW}}(m) = \frac{\Gamma_{Z'}}{(m_{Z'} - m)^2 + \Gamma_{Z'}^2}, \quad (5.1)$$

where $m_{Z'}$ and $\Gamma_{Z'}$ are the pole mass and width of the resonance, respectively. This choice of signal shape is motivated by the shape of the SM Z boson, which can be described by a Breit-Wigner function. For wider signals, a relativistic Breit-Wigner function would be a more realistic model.

The total signal shape of the wide signals are found by taking the convolution integral of the non-relativistic Breit-Wigner function with the detector resolution TF. This convolution models how the truth level signal shape is modified by detector resolution. The convolution integral for the Breit-Wigner and detector resolution TF is given as

$$f_{\text{wide}}(m) = \int_{-\infty}^{\infty} f_{\text{BW}}(x) f_{\text{TF}}(m - x) dx. \quad (5.2)$$




Figure 5.1.: Zero-width signal probability density functions (PDFs) at pole masses 3.5, 4, and 4.5 TeV in the electron (a) and muon (b) channels. The area under each signal curve has the same arbitrary normalisation.

For the resulting wide signal function to be centred around the pole mass of the Breit-Wigner function, either the Breit-Wigner or the TF must be translated to be centred around zero before calculating the above integral.

The Breit-Wigner width $\Gamma_{Z'}$ is expressed as a percentage of the pole mass $m_{Z'}$. This is referred to here as the relative width of the signal. Relative signal widths between 0% and 10% are considered in the search. Figures 5.2 and 5.3 show the convolution of a 2.5% width Breit-Wigner function with the detector resolution function to obtain the total signal shape for wide signals. The figures also show how the total signal shape changes with the increasing width of the Breit-Wigner functions.

We generate signal MC samples to be able to calculate cross-sections and selection efficiencies for different spin-0, spin-1, and spin-2 BSM models predicting resonant dilepton signatures. Spin-1 LO DY samples are generated and then *re-weighted* to model different Z' benchmark models. The re-weighting procedure is explained briefly in Section 5.1.1. The DY samples are generated using PYTHIA 8, in separate ranges of invariant dilepton mass to enhance production in the high invariant mass region. Detector response to the DY samples is simulated using full Geant4 detector simulation.

Spin-2 RS1 model MC dilepton samples are generated using Pythia 8, for values of graviton G^* mass between 750 GeV and 5,000 GeV. Samples are generated for three different values of the warping parameter k of the extra dimension introduced by the RS1 model: $k/\bar{m}_{\text{Pl}} = 0.1, 0.2, \text{ and } 0.3$, where \bar{m}_{Pl} is the reduced Planck mass. MC samples of a spin-0 MSSM Higgs boson decaying to a dilepton pair are generated using SHERPA, for values of Higgs boson pole masses between 400 GeV and 1,000 GeV, and for values of the relative resonance widths between 0% and 20%. The detector response to the generated RS1 graviton- and MSSM Higgs MC samples are modelled using Atlast-II.

A summary of the event generators and PDF sets used to generate the signal MC samples is shown in Table 5.1. In this work, the generated RS1- and Higgs boson samples are only used to calculate the acceptance and efficiency in the analysis for spin-0 and spin-2




Figure 5.2.: A 2.5% relative width Breit-Wigner distribution, centred at 1 TeV, convoluted with the zero-width detector response function for electrons (a). The total electron signal shape PDFs for five signal hypotheses with pole mass $m_{Z'} = 1$ TeV, and relative widths between 0% and 10% (b). In (b), The area under each signal curve has the same arbitrary normalisation.




Figure 5.3.: A 2.5% relative width Breit-Wigner distribution, centred at 1 TeV, convoluted with the zero-width detector response function for muons (a). The total muon signal shape PDFs for five signal hypotheses between 0% and 10% relative width (b). In (b), the area under each signal curve has the same arbitrary normalisation.

resonant dilepton signals.

Table 5.1.: List of MC generators used in signal sample production.

<i>Process</i>	<i>Hard Scattering Process with PDFs</i>	<i>Parton Shower, Non-Perturbative Effects with PDFs</i>
Drell-Yan	PYTHIA v8.186, NNPDF23LO	PYTHIA v8.186, NNPDF23LO, EVTGEN 1.2.0
Randall-Sundrum $G^* \rightarrow ll$	PYTHIA v8.210, NNPDF23LO	PYTHIA v8.210, NNPDF23LO, EVTGEN 1.2.0
MSSM $gg \rightarrow H \rightarrow ll$	SHERPA 2.1.1, CT10	PYTHIA v8.212, CTEQ6L1, EVTGEN 1.2.0

To reduce the effect of SM interference and off-shell effects in the signal samples, a *fiducial* selection is applied to each signal sample event. Each electron (muon) candidate is required to pass $|\eta| < 2.5$, $E_T(p_T) > 30$ GeV, and $m_t > m_{Z'} - 2\Gamma_{Z'}$. Here, m_t refers to the dilepton invariant mass of the MC sample at generator level, or the "true" mass of the dilepton pair, which, due to detector resolution effects, may be different from the invariant mass m_{ll} found in event reconstruction. Interference effects may change the low mass tail of the signal distribution, having a large effect on the cross-section without altering the shape of the peak of the distribution considerably. Removing the low mass tails of the signal distributions is done to remove cross-section model dependence. The fiducial selection defined above is the same as the one used in the previous resonant search performed by the ATLAS exotic dilepton working group, using 36 fb^{-1} of Run 2 proton-proton collision data [10], with one exception: Instead of removing just the low mass tail of the signal distribution, Ref. [10] employs a two sided mass cut at twice the width of the resonance, $m_{Z'} - 2\Gamma_{Z'} < m_t < m_{Z'} + 2\Gamma_{Z'}$. We find that the cross-section model dependence is smaller in the high-mass tail of the signal distributions under consideration, so this requirement is removed in the analysis.

5.1.1. Signal Re-Weighting

Spin-1 signal model samples at arbitrary pole mass values are generated by *matrix element re-weighting* the DY samples listed in Table 5.1. This allows us to create distributions for spin-1 Z' models at arbitrary pole mass, without the need for dedicated MC sample production. Each event in the DY sample is re-weighted by a factor w . This weight is given by the ratio of the differential cross-section of the BSM process over the corresponding SM differential cross-section, as

$$w = \frac{d\sigma/d\hat{t} (q\bar{q} \rightarrow \gamma^*/Z/Z' \rightarrow l\bar{l})}{d\sigma/d\hat{t} (q\bar{q} \rightarrow \gamma^*/Z \rightarrow l\bar{l})}, \quad (5.3)$$

where \hat{t} is the kinematic Mandelstam variable representing the squared difference of the initial state quark q and final state lepton l four-momenta.

We create signal distributions for the E_6 models Z'_χ and Z'_ψ . The relative resonance widths of Z'_χ and Z'_ψ are 0.5% and 1.2%, respectively. Signal distribution templates are also created for the Sequential Standard Model (SSM) Z' model. The couplings of the Z'_{SSM} boson are equal to the SM Z boson in all respects, but the resonance pole mass of Z'_{SSM} is larger than the mass of the Z boson. Lastly, we create signal samples for three choices of resonance pole mass, $m_{Z'} = 3, 4, \text{ and } 5$ TeV, in the HVT theory Z'_{HVT} model. More information about these spin-1 BSM models can be found in Section 1.2.1.

5.2. Determining Acceptance Times Efficiency

The *acceptance* A of a detector is defined as its solid angle coverage. If the trajectory of a particle in the ATLAS detector falls outside of the solid angle covered by the detector volume, the particle cannot be detected. Moreover, the ATLAS detector sensors are not perfectly efficient. Nor are the ATLAS trigger-, particle reconstruction-, and particle identification systems. Some number of particles will not be observed, despite hitting the active volume of the detector. The fraction of particles passing the acceptance requirement that is actually recorded is referred to as the *efficiency* ϵ of the detector. However, every event recorded by the ATLAS detector is not necessarily used in physics analyses. Particles are trimmed further from the ATLAS data by the event selections. Therefore it is useful to instead define detector efficiency as the fraction of events created in the detector that, passing detector acceptance, ends up in the final data set of an analysis. If so, then the product of the acceptance multiplied by the efficiency, $A\epsilon$, is equal to the probability, for an event produced in the detector, of ending up in the analysis data set.

Knowing the acceptance times efficiency of the event selection in the analysis is necessary to calculate the production cross-section of a signal. For a new physics signal with production cross-section σ , the expected number of signal events in the analysis data set is given by

$$N_{\text{sig}} = LA\epsilon\sigma\text{Br}, \quad (5.4)$$

where L is the integrated luminosity of the data set, and Br is the branching ratio of the process into a dielectron or dimuon final state.

The acceptance and efficiency of zero-width resonances are determined using the 200 NLO DY MC samples used to model the detector resolution TFs. These samples are generated in narrow slices of dilepton invariant mass. $A\epsilon$ as a function of generator level invariant mass m_t is found by dividing the number of events in the reconstructed sample passing the analysis event selection, by the number of generated events in the sample:

$$A\epsilon(m_t) = \frac{\text{Events passing selection}}{\text{Events in MC sample}} \Big|_{m_t \text{ in slice}}. \quad (5.5)$$

To interpolate values of acceptance and efficiency between the points in m_t covered by the MC samples, polynomial functions are fitted to the calculated dielectron and dimuon $A\epsilon$ values. The acceptance and efficiency calculated here is extracted from DY samples and therefore assume the signal to be spin-1. We test the spin-dependence of the acceptance and efficiency by calculating $A\epsilon$, using Equation (5.5), for the spin-0 and spin-2 samples shown in Table 5.1. The continuous spin-1 $A\epsilon$ functions, as well as the single point spin-0 and spin-2 $A\epsilon$ values, are shown in Figure 5.4. In the muon channel, the $A\epsilon$ values for spin-0, spin-1, and spin-2 signals agree to within 1%. In the electron channel, on the other hand, the spin-0 and spin-2 $A\epsilon$ values are higher than the spin-1 values by as much as 4%. Part of this discrepancy can be explained by the different methods used to model detector resolution in the MC samples. MC samples in which detector response is modelled using AtIfast-II, are seen to have a slightly higher acceptance times efficiency compared to samples using the full Geant4 detector simulation. However, the main source of this discrepancy is thought to be due to the fact that the spin-0 and spin-2 processes in the signal MC samples are more central in pseudorapidity than the spin-1 sample. A larger number of the spin-0 and spin-2 events pass the $|\eta| < 2.5$ requirement of the event selection than events in the spin-1 sample.

For wide signals, the m_t of the signal events are no longer all the same value, but are spread out around the pole mass $m_{Z'}$. In this case the Equation (5.5) cannot be used, as $A\epsilon$ may change as a function of m_t . We have to take into account the fact that the acceptance and efficiency of the signal changes from one end of the distribution to the other. For the generator level distribution $S_t^{Z'}$ of a wide signal, in this case a Breit-Wigner distribution, the acceptance and efficiency is found by integrating the product of the wide signal distribution and the zero-width $A\epsilon$ given in Equation (5.5) as a function of m_t , normalised by the integral of the wide signal shape:

$$A\epsilon(m_{Z'}) = \frac{\int S_t^{Z'}(m_t)A\epsilon(m_t)dm_t}{\int S_t^{Z'}(m_t)dm_t}. \quad (5.6)$$

The acceptance times efficiency values vary between 62% (54%) at 225 GeV and 74% (38%) at 6 TeV for zero-width signals in the electron (muon) channel. The difference in $A\epsilon$ between a zero-width signal and a signal of relative width 10% is less than 0.5% across most of the pole mass range; this difference grows as large as 2% above 5 TeV in the muon channel. In the spin-0 and spin-2 samples, acceptance times efficiency varies by less than 1% across all signal width variations.




Figure 5.4.: Acceptance times efficiency as a function of resonance pole mass for physics processes of spin-0, spin-1, and spin-2. Zero width signals are used to produce the spin-1 values. The spin-0 and spin-2 values are averaged over all choices of signal width [6].

5.3. Data-Driven Background Estimation

In the 36 fb^{-1} analysis presented in Ref. [10], which is the previous result from the ATLAS exotic dilepton working group, the SM background is estimated using MC samples. The SM background in this search is estimated using a data-driven method, by fitting a functional form to the data separately in the electron and muon channels.

We explore two methods of fitting the background: a *global fit* method, where the full invariant mass search range is fitted with a single function, and a sliding window fit (SWiFt) method, where consecutive fits are performed in restricted invariant mass intervals around the pole mass hypotheses of the search.

The sliding window fit method is named for the apparent "sliding" motion of these windows, as it scans across invariant mass subranges, searching for deviations from the SM hypothesis. The main strength of the SWiFt method is that the parametric fit functions used to describe the SM background can be kept simple; in this work we only consider functions of one and two free parameters, not counting the normalisation of the functions. The SWiFt method is not constrained in the high invariant mass region in the same way as the global fit method. The SWiFt method is therefore more likely to suffer from fit instabilities in this region due to high statistical uncertainty. Care must also be taken to not make the fit windows too narrow compared to the width of the signal hypothesis under consideration. If the sideband-, or signal-free, region of the fit is too narrow, the sensitivity to the signal can be reduced. Examples of high energy physics analyses using the SWiFt method can be seen for an ATLAS dijet search in Ref. [98], and for a Compact Muon Solenoid (CMS) Collaboration diphoton search in Ref. [99].

There are two main ways of using sliding windows for data-driven background estimation. The first technique is to build a so-called *SWiFt background template*. This histogram is constructed by scanning over pole masses, fitting the chosen parametric signal and background functions in each sliding window. Next, the background template is made by "stitching together" the background normalisation in the central bin of each window. Due to the large overlap between windows, neighbouring bins in the background template are strongly correlated, resulting in a smooth background shape. The SWiFt background template can then take the place of the SM background estimate in the subsequent data analysis. We do not use the SWiFt background template in the statistical analysis.

The second alternative of the SWiFt background method is the so-called *true SWiFt* method. Here, the signal and background components are fitted in each window individually, calculating the significance of the signal and setting upper limits on the signal strength, before sliding the windows to the next pole mass signal hypothesis. This approach is the method used in this work.

The global fit method is presented in this chapter, while the SWiFt method, and a comparison between the two methods, is presented in more detail in Appendix A. The SWiFt approach is found to be a viable option for doing data-driven background estimation, producing similar results to those obtained using the global fit approach. However, the method introduces a lot of complexity, particularly in the optimisation of the sliding windows. Therefore, the global fit method is selected for this search. In the future, at Run 3 and beyond, as data set sizes continue to grow, it may become impractical, or even impossible, to use the global fit approach. At this point, the SWiFt method could prove a viable option for doing data-driven background estimation.

5.3.1. Spurious Signal

When performing a data-driven background estimation there are always uncertainties tied to the choice of fit function; how do we know that the chosen mathematical function accurately models the SM? Including more free parameters in the functional form increases the *flexibility* of the fit. A polynomial of arbitrarily large order will conform perfectly to any distribution. This is not wanted, as this would hide any potential new physics signal in the background estimate. Instead, we want the function to conform to the broad, unlocalised, behaviour of the SM background, while leaving any narrow resonances in the data to be fit by the signal shape functions introduced in Section 5.1. Choosing the correct fit function, therefore, becomes an optimisation procedure, balancing the bias and power

of the function.

If a discrepancy is observed between the background fit and the data itself, we have a potential observation of new physics. However, we need a way of determining whether or not this discrepancy is due to a new physics signal, or simply due to bias in the fit function. If the chosen fit function poorly describes the data, it may under- or overestimate the background shape. Then, when adding a signal component to the fit, the signal shape may compensate for the error in the background fit. This leads to a non-zero signal contribution, not from any new physics signal, but as a result of background mismodelling. Such false signals are known as *spurious signals*. This effect is illustrated in Figure 5.5 (a) and (b). Here, a linear function is a poor choice to describe the toy background function. The background function is perfectly smooth, containing no localised resonances. And yet, when a Gaussian signal hypothesis is included in the fit model, a significant signal is found. It should be pointed out here that a signal contribution arising from a statistical fluctuation of the data is *not* a spurious signal. Particle collision data are stochastic by nature, and will always contain statistical fluctuations.




Figure 5.5.: Illustrating the concept of spurious signal. A signal-free toy background distribution, generated here by a second-order polynomial, is poorly described by a linear background hypothesis (a). This mismodelling leads to a Gaussian signal being extracted where none should be found (b). In (c), we see that when a spurious signal term is added to the background shape, the correct signal normalisation is extracted. Without spurious signal treatment, the signal is greatly overestimated.

To quantify the expected spurious signal of a background function, we fit the function, along with the function describing the signal shape, to the smooth background MC template described in Chapter 3. Seeing as we do not expect to see any signal in the background MC template, the number of signal events extracted from the signal shape is taken to be the expected spurious signal bias of the background function. We want the spurious signal to be as low as possible. However, the absolute value of the spurious signal is not a good measure of the performance of the fit function. At low invariant mass the number of events in the data is much greater than at higher mass (see Figures 4.1 and 4.2). A spurious signal of 10 events could be insignificant compared to the expected size of an observable signal at low mass, but devastatingly large at high mass. Instead, we want the spurious signal to be low compared to the statistical uncertainty in the data. Therefore, we use the spurious signal *significance* Z_{SS} , given as the ratio of spurious signal events to the uncertainty on the spurious signal yield, N_{SS}/σ_{SS} . The spurious signal uncertainty σ_{SS} is driven by the statistical uncertainty in the data. The MC background template is normalised to the integrated luminosity of the Run 2 proton-proton data set. This means that the extracted uncertainty on the spurious signal in the fit is equivalent to the statistical uncertainty of a real signal measurement on the data. Using the spurious signal significance we define a so-called spurious signal test to assess candidate background functions. If $Z_{SS} \ll 1$, then the spurious signal is not expected to have a significant influence on the statistical analysis. If Z_{SS} is less than some threshold value less than one, for all invariant mass points in the signal scan range, then the function passes the test. The threshold value chosen for this analysis is $Z_{SS} < 0.5$. In addition to using spurious signal to discriminate background fit functions, we also use spurious signal to account for the background mismodelling uncertainty in the statistical analysis. This is done by adding a spurious signal term in the signal normalisation parameter in the signal PDF, essentially allowing the background model to flex or bend as the signal shape. This assumption, that the background mismodelling has the same form as the signal hypothesis, is the worst-case scenario. The size of the spurious signal normalisation is determined by a Gaussian constraint in the statistical model, covered in more detail in Section 5.6. This way, if a signal is observed, we can be sure that it is not a spurious signal. This concept is illustrated in Figure 5.5 (c). Here, a Gaussian signal has been injected over the smooth background shown in Figures 5.5 (a) and (b). Without the spurious signal term in the signal normalisation, the size of the extracted signal is considerably overestimated, due to the background mismodelling. However, by adding a spurious signal component to the signal normalisation, equal to the spurious signal found in the signal plus background fit on the signal-free background (shown Figure 5.5 (b)), the extracted signal strength becomes consistent with the injected signal.

5.3.2. Standard Model Background Estimation

The *global fit method* is the name given to one approach for doing the data-driven background estimation, where the entire invariant mass search range is fitted using a single parametric function. With its large fit range, the global fit method is well suited for fitting broad signal resonances. The method also has the benefit of the low mass, high-statistics, region constraining the background fit in the high-mass, low-statistics, region. However, finding a suitable function to fit the background, balancing the opposing needs of minimising bias and maximising expected significance, can be challenging. The global fit method used in this analysis is similar to the high invariant mass diphoton resonance searches of

Refs. [100] and [101].

The optimal choice of fit function for the global fit method is determined using a spurious signal test, as described in Section 5.3.1: the candidate background fit functions must satisfy $Z_{SS} < 0.5$ for any pole mass when scanning over the signal-free MC background template. If more than one function satisfy this requirement, then the function with the least free parameters is chosen. Functions of fewer parameters are less "flexible", and thus less likely to conform to, and hide, a real signal present in the data. The function found to best describe the dilepton SM background is

$$f_u(m) = f_Z(m)(1 - x^C)^b x^{\sum_{i=0}^3 p_i \log(x)^i}, \quad (5.7)$$

where $x = m/\sqrt{s}$, where $\sqrt{s} = 13$ TeV is the COM collision energy of the LHC during Run 2, and $f_Z(m)$ is the non-relativistic Breit-Wigner shape presented in Equation (5.1), with the Z boson mass and width used in place of those of the Z' . The term $(1 - x^C)^b$ in Equation (5.7) is physically motivated, as it tends to zero as m approaches \sqrt{s} ; it is impossible in a collider to observe collision events more energetic than the available COM collision energy. The parameter C is not a free parameter in the background fits, but a constant, set to $C = 1$ in the electron channel, and $C = 1/3$ in the muon channel. The value of C is chosen so as to optimise the fit to the smooth background template. The spurious signal significance of Equation (5.7) is found to be around 0.3 or less for any pole mass in a spurious signal test for a zero-width signal. For a 10% relative signal, the spurious signal significance grows, but remains less than the 0.5 threshold.

Alternate values for the constant C in Equation (5.7), as well as setting $b = 0$, are tested, without increasing the performance in the spurious signal test. To test that the chosen function have an appropriate number of free parameters we perform a test where an extra degree of freedom is added to Equation (5.7) in the form of the factor $x^{p_4 \log(x)^4}$. This extended function is kept if it improves the likelihood value of a background-only fit by 2σ compared to the non-extended function. This requirement is not met, and the non-extended function is kept. The stability of the chosen global fit function is tested by injecting a signal of varying strength into the smooth MC template at selected pole masses and widths across the full scan range, and testing that the extracted signal yields match those injected.

5.4. Parameterising the Spurious Signal Uncertainty

The size of the spurious signal is measured using signal plus background fits to the signal-free SM background template in a spurious signal test. There is a risk that random fluctuations in the MC background sample will lead to an underestimation of the spurious signal for some signal hypotheses. The underestimation of the size of the spurious signal may lead to overestimation of the sensitivity to new physics signals. Therefore, we make a conservative estimate of the spurious signal by fitting a parameterised function to the local maxima of the spurious signal values from the spurious signal test.

The function chosen to describe the spurious signal envelope is

$$f_{\text{spur}}(x) = (1 - x^{1/3})^a x^{\sum_{i=1}^4 p_i \log(x)^i}, \quad (5.8)$$

where once again, $x = m/\sqrt{s}$. This function is almost identical to the muon channel background function of Equation (5.7), except for a small difference in the exponential logarithmic polynomial, which here includes terms of the first to the fourth degree. The

spurious signals extracted for a zero-width signal in the global fit approach is shown in Figures 5.6. These figures also show the best fit of f_{spur} to the local maxima of the spurious signal distributions. The spurious signal scan and envelope fit are performed for all signal




Figure 5.6.: Spurious signal yields extracted in a scan for a zero-width signal on the smooth MC template in the electron (a) and muon (b) channels [93]. The best fits of the parametrisation functions f_{spur} to the ensemble of spurious signal maxima are also shown.

width assumptions, from 0% to 10% in steps of 0.5% signal width.

5.5. Systematic Uncertainties

All background shape uncertainties are propagated into the spurious signal uncertainty. The background uncertainties under consideration are choice and shape uncertainty in generator PDFs and shape and normalisation uncertainties on the top quark and fake electron backgrounds. These uncertainties are handled by producing smooth dilepton invariant mass templates, using the TF method described in Chapter 3, shifted up and down by one standard deviation of the shape systematic uncertainties. The spurious signal yield is extracted for each shifted template, at every pole mass hypothesis in the scan range. The maximum spurious signal yield at each pole mass of any of the MC templates, nominal or shifted by systematic uncertainty, is taken to be the final spurious signal estimate used in the analysis.

Systematic uncertainties on the signal can be divided into two categories: signal normalisation uncertainties and signal shape uncertainties. Sources of uncertainty are assumed to only belong to one of these two categories. The signal normalisation uncertainties considered in the statistical analysis are lepton identification and isolation uncertainties, uncertainty on the luminosity, and uncertainty on the "good muon" selection requirement. Only systematic uncertainties seen to affect the signal normalisation by more than 0.5% are considered. All the signal normalisation uncertainties are assumed to be fully uncorrelated between the lepton channels. One exception of this is the luminosity scale uncertainty, which is taken to be fully correlated between the electron and muon channels.

Determining the signal shape uncertainties require a refit of the detector response TFs for the up- and down variation of the 200 zero-width NLO Drell-Yan MC samples by each systematic uncertainty. Energy and momentum scale and resolution uncertainties are the only uncertainties that alter the detector response shape. The scale and resolution

uncertainties alter the mean and width, respectively, of the Gaussian and Crystal Ball functions that make up the TFs. These effects are visualised in exaggerated form in Figure 5.7.




Figure 5.7.: Exaggerated impact on the nominal detector resolution functions in the electron channel of up- and down variations of the energy resolution- (a) and scale (b) uncertainties [93]. The dashed lines show the Crystal Ball component of the TFs.

Smooth parameterisations of the systematic uncertainties as a function of dilepton invariant mass are obtained using the TF framework, introduced in Section 3.6, developed to model detector response. Parameterisations of the impact of the shape systematic uncertainties on the mean and width parameters of the TFs are obtained for zero-width signals, and for signals of widths between 0% and 10% in steps of 0.5%. We see no significant difference between the nominal zero-width parameterisations and the parameterisations obtained for wider signals. Therefore, only the zero-width parameterisations are kept and used for all signal width hypotheses.

The relative impact on the signal yield of the systematic uncertainties can be seen, for two pole-mass points, in Table 5.2. Spurious signals are the largest source of uncertainty at low pole masses. At high pole-masses, the electron channel uncertainty is dominated by the electron identification efficiency. In the muon channel, the uncertainty in the high pole-mass region is also dominated by the uncertainty associated with the Good Muon selection requirement.

Systematic uncertainties are included in the statistical analysis as nuisance parameters with log-normal or Gaussian constraints in the statistical model, which will be presented below. Log-normal constraints are used for uncertainties on the resolution, identification and isolation of leptons, as well as on the luminosity scale uncertainty. Gaussian constraints are used for energy and momentum scale uncertainties, and the spurious signal uncertainty.

The signal normalisation at any given pole mass, including nuisance parameters, becomes

$$N_{\text{Sig},ee} = L \times \sigma_{Z' \rightarrow ee} \times A_{ee} \times \epsilon_{ee} \times \kappa_{\text{Lumi}} \times \kappa_{\text{Iso},ee} \times \kappa_{\text{Ident},ee} + N_{\text{SS},ee} \times \theta_{\text{SS},ee}, \quad (5.9)$$

and

$$N_{\text{Sig},\mu\mu} = L \times \sigma_{Z' \rightarrow \mu\mu} \times A_{\mu\mu} \times \epsilon_{\mu\mu} \times \kappa_{\text{Lumi}} \times \kappa_{\text{Iso},\mu\mu} \times \kappa_{\text{Ident}(\text{stat}),\mu\mu} \\ \times \kappa_{\text{Ident}(\text{syst}),\mu\mu} \times \kappa_{\text{GoodMuon},\mu\mu} + N_{\text{SS},\mu\mu} \times \theta_{\text{SS},\mu\mu}, \quad (5.10)$$

in the electron and muon channel, respectively. Here, L is the integrated luminosity of the data sample and $A_l \times \epsilon_l$ is the acceptance times efficiency of the event selection at

Table 5.2.: Relative impact, in per cent, on the number of extracted signal events from a $\pm 1\sigma$ shift of the systematic uncertainties included in the analysis. The shifts are shown for a zero-width (10% relative width) signal at two example pole masses.

<i>Pole mass</i> [GeV]	<i>Electron Channel</i>		<i>Muon Channel</i>	
	300	5,000	300	5,000
Spurious signal	± 12.5 (12.0)	± 0.1 (1.0)	± 11.7 (11.0)	± 2.1 (2.2)
Identification	± 1.6 (1.6)	± 5.6 (5.6)	± 1.8 (1.8)	$^{+25}_{-20}$ ($^{+25}_{-20}$)
Isolation	± 0.3 (0.3)	± 1.1 (1.1)	± 0.4 (0.4)	± 0.4 (0.5)
Luminosity	± 1.7 (1.7)	± 1.7 (1.7)	± 1.7 (1.7)	± 1.7 (1.7)
Electron energy scale	$^{+1.7}_{-4.0}$ ($^{+1.0}_{-1.8}$)	$^{+0.1}_{-0.4}$ (± 0.8)	-	-
Electron energy resolution	$^{+7.9}_{-8.3}$ ($^{+1.1}_{-0.9}$)	$^{+0.4}_{-0.9}$ (± 0.1)	-	-
Muon ID resolution	-	-	$^{+0.8}_{-2.3}$ ($^{+0.3}_{-0.8}$)	$^{+0.6}_{-0.4}$ ($^{+0.5}_{-0.3}$)
Muon MS resolution	-	-	$^{+2.8}_{-3.8}$ ($^{+1.0}_{-1.3}$)	± 2.4 (2.1)
Good Muon requirement	-	-	± 0.6 (0.6)	$^{+55}_{-35}$ ($^{+55}_{-35}$)

the pole mass in question. The fiducial cross-section times branching ratio of a potential new signal, $\sigma_{Z' \rightarrow ll}$, is determined by fitting the signal and background PDFs to the data. The value of the spurious signal nuisance parameter $\theta_{SS, ll}$, is also determined in this fit. $N_{SS, ll}$ is the spurious signal normalisation, determined in Section 5.4. The factors $\kappa_{i, ll}$ control remaining signal normalisation nuisance parameters. The nominal values of these parameters are set to one.

5.6. Statistical Analysis

The RooFit [102] and RooStats [103] data analysis tools are used to perform the statistical analysis, which is based on a frequentist interpretation of statistics and probability. Previous iterations of ATLAS exotic dilepton analyses have employed Bayesian statistics. No discernible difference is expected between the two methodologies.

The statistical analysis takes place in two steps: the *discovery* step and the *exclusion* or limit setting step. In the discovery step, we check the validity of the SM given the observations. If no discrepancy with the SM is seen, we move on to the exclusion step. Here, we quantify the upper limit on the parameter of interest: the cross-section times branching ratio of a new dilepton (e^+e^- , $\mu^+\mu^-$) resonance.

5.6.1. Discovery

To claim the discovery of a new physics signal, the standard model must be shown to be inconsistent with experimental results. The SM in this case can be referred to as the background or null hypothesis, H_0 , of the analysis.

For a set of observables \mathbf{x} and model parameters Θ , the likelihood $\mathcal{L}(\Theta|\mathbf{x})$ of Θ , given the outcome \mathbf{x} , is equal to the probability $P(\mathbf{x}|\Theta)$ of observing the outcome \mathbf{x} , given the values of Θ . In this search we have one parameter of interest, μ , and a set of nuisance parameters θ . The signal strength parameter μ is used here to symbolise the cross-section times branching ratio $\sigma_{Z' \rightarrow ll}$ of Equations (5.9) and (5.10). The symbol μ is kept to stay consistent with literature. The nuisance parameters are related to the uncertainties on

the scale of the luminosity, the spurious signal uncertainty, etc. The likelihood of a set of values of μ and $\boldsymbol{\theta}$, given the binned data \mathbf{n} , is given as a product of the individual bins in the dilepton invariant mass histograms, as

$$\mathcal{L}(\mu, \boldsymbol{\theta}|\mathbf{n}) = \prod_i^{\text{bins}} \text{Pois}(n_i|\mu S_i(\boldsymbol{\theta}) + B_i(\boldsymbol{\theta})) \times \prod_j C_j(\theta_j|\mu_j^\theta, \sigma_j). \quad (5.11)$$

Here, n_i , $S_i(\boldsymbol{\theta})$, and $B_i(\boldsymbol{\theta})$ are the event yield of each bin for the data, the signal function, and the background function, respectively. $\text{Pois}()$ is the probability mass function of a Poisson distribution, giving the probability of drawing n_i from a Poisson distribution of mean $\mu S_i(\boldsymbol{\theta}) + B_i(\boldsymbol{\theta})$. The constraint terms for $\boldsymbol{\theta}$, C_j , are either Gaussian or log-normal probability density functions of mean μ_j^θ and standard deviation σ_j , where σ_j is the systematic uncertainty associated with the nuisance parameter θ_j .

A *test statistic* is a variable quantifying how well measured data match a given hypothesis. A likelihood ratio test, also known as a Wilks test, is used in the analysis. The logarithmic likelihood-ratio test statistic is given as

$$q_0 = \begin{cases} +2 \log \left[\frac{\mathcal{L}(0, \hat{\boldsymbol{\theta}}_0)}{\mathcal{L}(\hat{\mu}, \hat{\boldsymbol{\theta}})} \right] & \hat{\mu} < 0, \\ -2 \log \left[\frac{\mathcal{L}(0, \hat{\boldsymbol{\theta}}_0)}{\mathcal{L}(\hat{\mu}, \hat{\boldsymbol{\theta}})} \right] & \hat{\mu} \geq 0. \end{cases} \quad (5.12)$$

Here, $\hat{\boldsymbol{\theta}}_0$ and $\hat{\boldsymbol{\theta}}$ are the nuisance parameter values that maximise the likelihood, given a signal strength value of 0 or $\hat{\mu}$, respectively. Likewise, $\hat{\mu}$ is itself chosen to maximise the likelihood. Note that we omit the measurements \mathbf{n} from the likelihood expression for clarity. The signal strength parameter is allowed to be both positive and negative, allowing us to test both upward and downward fluctuations of the data.

We can calculate the value of q_0 on an ensemble of toy distributions drawn under the null hypothesis and make a probability density function $f(q_0|H_0)$. The resulting PDF will look something like the distribution shown on Figure 5.8. The probability p_0 that the background fluctuates to yield a signal-like excess equal to or larger than the observed value q_0^{obs} , can be calculated using the $f(q_0|H_0)$ distribution as

$$p_0 = P(q_0 \geq q_0^{\text{obs}}|H_0) = \int_{q_0^{\text{obs}}}^{\infty} f(q_0|H_0) dq_0. \quad (5.13)$$

p_0 is also known as the *p-value* of the experiment. It is common to express p_0 as a significance Z , which has units of standard deviations σ of a unit Gaussian distribution. The *p-value* can be converted to a significance using the inverse of a cumulative unit Gaussian Φ^{-1} , as

$$Z = \Phi^{-1}(1 - p_0). \quad (5.14)$$

Small *p-values*, less than 0.5, correspond to positive values of $\hat{\mu}$, and *p-values* greater than 0.5 correspond to negative values of $\hat{\mu}$. It is customary in the field of high energy physics to only claim discovery if the *p-value* is smaller than 2.87×10^{-7} , which is equivalent to a significance of 5σ [104].

Wilks' theorem states that in the so-called asymptotic limit, when sample size approaches infinity, the distribution $f(q_0|H_0)$ approaches a χ^2 distribution for one degree of freedom [106]. In the cases where this condition is met, such as in the example test statistic distribution of Figure 5.8, p_0 can be calculated without the need for drawing toy distributions. This can greatly improve the computation speed when calculating *p-values*.




Figure 5.8.: Example test statistic distribution (red) for null hypothesis toy data sets [105]. Also shown (blue) is a χ^2 distribution for one degree of freedom. The test statistic distribution is seen to be consistent with Wilks' theorem.

The asymptotic formula for calculating the p -value from the test statistic directly is given as [107]

$$p_0 = 1 - \Phi(\sqrt{q_0}), \quad (5.15)$$

and similarly, the significance can be calculated as

$$Z = \sqrt{q_0}. \quad (5.16)$$

For negative values of q_0 , corresponding to negative values of $\hat{\mu}$, the negative sign must be brought outside of the square root in Equations (5.15) and (5.16).

In this analysis, we do not calculate only one p -value. Instead, we scan across the dilepton invariant mass histogram, calculating the p -value at every one GeV step. For each new p -value checked, the probability grows that a random background fluctuation will mimic a signal. This is known as the look-elsewhere effect. The p -values found at each individual pole mass scan point are known as *local* p -values. A *global* p -value is the probability for the background to fluctuate to yield a signal-like excess equal to or larger than what is observed, at *any* of the tested invariant mass points in the scan, taking into account the look-elsewhere effect.

Global p -values can be found by throwing background hypothesis toy distributions, performing a scan on them for local p -values, and then noting the frequency at which certain p -values show up. The local significances found in the scan on data can then be mapped to global significance values using these frequencies.

An asymptotic approximation, valid for local significances above 2σ , can be used to avoid drawing toy distributions. First, choose a reference significance level Z_{ref} , usually 0σ . Next, count the number of times N_{up} the local significance curve crosses this threshold in the upward direction. A local p -value can now be converted to a global p -value using the equation [108]

$$p_0^{\text{glob}} = p_0^{\text{loc}} + N_{\text{up}} e^{-\frac{1}{2}(Z_{\text{loc}}^2 - Z_{\text{ref}}^2)}. \quad (5.17)$$

The statistical uncertainty on this global p -value can be found by varying N_{up} by $\pm\sqrt{N_{\text{up}}}$ and repeating the above calculation. To increase the precision of the procedure, the value of N_{up} measured on the data can be replaced by the mean number of upward crossings found in an ensemble of background-only toys. Note that this does not require drawing as many pseudoexperiments as the non-asymptotic procedure.

5.6.2. Exclusion

The exclusion step can also be referred to as the limit setting step of the statistical analysis. If we are unable to reject the null hypothesis in the discovery step, then we want to determine how large the parameter of interest μ could be while remaining undetected in the experiment. In other words, we want to set a reasonable upper limit on the cross-section of a generic resonant dilepton signal. The competing hypothesis to the null hypothesis, the alternative hypothesis H_1 , is that such a generic signal exists, with some non-zero value of the cross-section $\sigma_{Z' \rightarrow ll}$.

When setting limits we use a slightly modified test statistic from the one used in the discovery step. The exclusion test statistic q_μ takes the form

$$q_\mu = \begin{cases} -2 \log \left[\frac{\mathcal{L}(\mu, \hat{\theta}_\mu)}{\mathcal{L}(0, \hat{\theta}_0)} \right] & \hat{\mu} < 0, \\ -2 \log \left[\frac{\mathcal{L}(\mu, \hat{\theta}_\mu)}{\mathcal{L}(\hat{\mu}, \hat{\theta})} \right] & 0 \leq \hat{\mu} \leq \mu, \\ 0 & \hat{\mu} > \mu. \end{cases} \quad (5.18)$$

Here, μ is the value of the signal strength parameter associated with the chosen H_1 hypothesis. At this value of μ , $\hat{\theta}_\mu$ are the maximum likelihood values of the nuisance parameters. Note that since this analysis step aims to set upper limits on the new signal cross-section, we are no longer interested in negative values of signal strength μ . When the maximum likelihood value of $\hat{\mu}$ becomes negative, $\hat{\mu} = 0$ enters the q_μ likelihood ratio calculation. Similarly, q_μ is set to zero in cases where $\hat{\mu} > \mu$, to ensure that these do not count as evidence against the alternative hypothesis.

Like in the discovery step, probability density functions can be created for the test statistic q_μ by drawing toy distributions under the null and alternative hypotheses. An example of two such distributions can be seen in Figure 5.9. Similar to the probability p_0 of Equation (5.13), we can now define a probability p_μ , which is the probability of observing a value q_μ as large as, or larger than, that observed:

$$p_\mu = P(q_\mu \geq q_\mu^{\text{obs}} | H_1) = \int_{q_\mu^{\text{obs}}}^{\infty} f(q_\mu | H_1) dq_\mu. \quad (5.19)$$

Just like a small value p_0 is evidence that the data is inconsistent with the null hypothesis, a small value p_μ means that the data is inconsistent with the alternative hypothesis. Again, in the asymptotic limit, Equation (5.15) can be used in place of Equation (5.19).

p_μ is known as the *confidence level* CL_{s+b} of the signal hypothesis. Similarly, the p -value p_μ^0 is known as the confidence level CL_b of the background. This p -value is calculated by substituting in the null hypothesis in place of the alternative hypothesis in Equation (5.19). For a chosen hypothesis H_1 , if CL_{s+b} is less than some predetermined threshold α , we say that the hypothesis is excluded with a $(1 - \alpha)$ *confidence level*. It is customary in high energy physics to choose $\alpha = 0.05$ for exclusion. We say that α is the *false exclusion rate* of the experiment. This means that if we were able to repeat the experiment several times, the H_1 hypothesis would be falsely excluded in one out of twenty experiments.

CL_{s+b} has an undesirable feature. If the null and alternative hypotheses are not well separated, a downward fluctuation of the data may lead to a false exclusion of a signal the analysis should not be sensitive to. A common solution to this problem is to use the so-called *modified frequentist* or CL_s method [109]. CL_s is not a confidence level, but a ratio of confidence levels, given as

$$\text{CL}_s = \frac{\text{CL}_{s+b}}{\text{CL}_b} = \frac{p_\mu}{p_\mu^0}. \quad (5.20)$$




Figure 5.9.: Example exclusion test statistic distributions for toy data sets drawn from the null- (blue) and alternative (red) hypotheses [105].

Finding the upper limit on the parameter of interest, μ_{up} , at 95% CL is an iterative process. We stepwise increase μ , calculating p_μ and p_μ^0 at each step. p_μ and p_μ^0 are used to calculate CL_s using Equation (5.20). μ_{up} is the value for which CL_s crosses the α , or 0.05, threshold.

In the asymptotic approximation, CL_s can be calculated using q_μ directly. Now, instead of drawing an ensemble of toy distributions, we can use the so-called Asimov data set to calculate the limits. The Asimov set is a data set where all observable parameters are set to their expected values. In the asymptotic limit, CL_s becomes

$$\text{CL}_s = \frac{1 - \Phi(\sqrt{q_\mu})}{\Phi(\sqrt{q_{\mu,A}} - \sqrt{q_\mu})}. \quad (5.21)$$

Here, $q_{\mu,A}$ is the test statistic calculated for the Asimov set [110].

To find the expected upper limit on the parameter of interest, given that there is no signal present in the data, generate toy experiments under the null hypothesis. For each of these pseudoexperiments, find the 95% CL upper limit on μ using the CL_s method outlined above and draw the resulting μ_{up} distribution. The expected limit on μ is the median of this distribution. The $\pm 1(2)\sigma$ uncertainty bands on the expected limit are defined by the 16 (2.3)% and 84 (97.7)% quantiles of the distribution. The expected limit μ^{exp} is found using the Asimov data in place of the data and finding the value of μ where $\text{CL}_s = 0.05$. The $N\sigma$ error bands on the expected limit are then calculated as

$$\mu_N^{\text{exp}} = \frac{\mu^{\text{exp}}}{\sqrt{q_{\mu,A}}} (\Phi^{-1}(1 - 0.05\Phi(N)) + N). \quad (5.22)$$

The validity of the asymptotic approximation in the high invariant mass, low statistics, region of the analysis is checked using toy experiments. The limits obtained with the asymptotic formulae are seen to be stronger than those found using toy distributions above 3 TeV. The impact on the Z' pole mass limit from using the asymptotic approximation is found to be less than 100 GeV for the benchmark models considered in the analysis.

5.7. Results

The best-fit background shapes are drawn with the data in Figure 5.10. Note that the data in these plots are not as finely binned as the data on which the analysis is performed. The

best-fit parameters of the background fit function in the electron channel are $b = 1.5 \pm 1.0$, $p_0 = -12.38 \pm 0.09$, $p_1 = -4.295 \pm 0.014$, $p_2 = -0.9191 \pm 0.0027$, and $p_3 = -0.0845 \pm 0.0005$, with a background normalisation of $N_{\text{bgr}} = 178,000 \pm 400$, in a background-only fit to the data. In the muon channel, the equivalent best-fit background-only parameters are $N_{\text{bgr}} = 138,700 \pm 400$, $b = 11.8 \pm 0.5$, $p_0 = -7.38 \pm 0.12$, $p_1 = -4.132 \pm 0.017$, $p_2 = -1.0637 \pm 0.0029$, and $p_3 = -0.1022 \pm 0.0005$. No deviations from the background hypothesis can be seen by eye. We scan the search range, 250 GeV to 6 TeV, in steps of 1 GeV, and widths between




Figure 5.10.: Dielectron (a) and dimuon (b) invariant mass distributions with the best-fit background PDFs. Three generic zero-width signals are also included in the plots at pole masses 1.34, 2 and 3 TeV. For added readability, the signal cross-sections are scaled by 20 times their upper limit values [6].

zero-width and 10% relative width in steps of 0.5%, searching for a signal. The 1 GeV steps are smaller than the mass resolution in both channels. Results are produced separately in the electron and muon channels and combined under the assumption of lepton universality. The local significance values from the scans are shown for a zero-width and a 10% relative width generic signal in Figure 5.11. Local significances are drawn in two dimensions as a




Figure 5.11.: Significance of a zero-width (a) and a 10% relative width (b) signal as a function of pole mass in the electron, muon, and combined channels [6].

function of pole mass and signal width in Figure 5.12. To aid readability of the plots in

Figure 5.12, only positive significances are drawn, with negative significances being set to zero. The global significance of the largest local significance points found in the scan is




Figure 5.12.: Positive significance as a function of pole mass and signal hypothesis width in the electron (a), muon (b), and combined (c) search channels [6].

calculated using the method outlined in the previous section. For a zero-width signal, the largest deviations from the background hypothesis are seen at 774 GeV, 267 GeV and 264 GeV, with local significances 2.9σ , 2.4σ , and 2.3σ in the electron, muon, and combined channels, respectively. In the same channels, the number of upwards crossings of the reference significance threshold, N_{up} , with uncertainty, are found to be 26 ± 5 , 7 ± 3 , and 26 ± 5 . This means that the global significances of the most locally significant scan points are 0.1σ , 0.3σ and 0.0σ in the electron, muon, and combined channels, respectively. For a 10% relative width signal, the highest significance points for the electron, muon, and combined channels are found at 771 GeV, 267 GeV and 2,390 GeV, with local significances 2.5σ , 2.2σ , and 1.4σ , and global significances of 0.5σ , 0.0σ and 0.0σ . Here, the values of N_{up} , with uncertainty, are 7 ± 5 , 5 ± 2 , and 10 ± 3 for the three channels. The largest local significance excesses and deficits for any signal width assumptions are shown in Table 5.3. From these results we conclude that no significant deviation from the SM is seen.

The upper limit on the fiducial cross-section at each pole mass, with the added restriction that the signal decays to a dilepton final state, is found using the asymptotic approximation presented in the previous section. The observed limits, as a function of pole mass, for a zero-width signal is compared to the expected limit with $\pm 1\sigma$ and $\pm 2\sigma$

Table 5.3.: Each dilepton channel’s most locally significant excess and deficit at any signal width hypothesis.

<i>Channel</i>	<i>Excess</i>			<i>Deficit</i>		
	$m_{Z'}$ [GeV]	$\Gamma_{Z'}/m_{Z'}$ [%]	Z [σ]	$m_{Z'}$ [GeV]	$\Gamma_{Z'}/m_{Z'}$ [%]	Z [σ]
ee	773	2.5	3.0	1,957	4.0	−3.2
$\mu\mu$	268	2.5	2.5	349	8.5	−2.8
ll	264	0	2.3	1,958	3.0	−2.9

error bands in Figure 5.13. This figure only shows limits up to pole masses of 2 TeV. Above this threshold, the expected limit error bands obtained in the asymptotic approximation becomes considerably wider than those obtained by toy-based methods.

Upper limit curves on the fiducial cross-section times branching ratio for zero-width signals and relative signal width assumptions 0.5%, 1.2%, 3%, 6% and 10% are shown in Figures 5.14, 5.15, and 5.16, for the electron, muon and combined channel, respectively. Expected limit error bands are omitted from these figures. We see that the zero-width and 10% relative width signal hypotheses bookend the upper limit values of the intermediate signal widths. The observed limits range from 3.6 (13.1) fb at 250 GeV, to 0.014 (0.018) fb at 6 TeV for zero-width (10% relative width) signals in the combined channel. Systematic uncertainties are seen to have a small impact on the limits. The limits are between 4% and 7% weaker when including systematic uncertainties, depending on the channel and pole mass value.

In Figures 5.14, 5.15, and 5.16, theoretical fiducial cross-section curves are drawn for the exotic heavy boson models Z'_{SSM} , Z'_{ψ} , and Z'_{χ} . Uncertainties on the Z'_{SSM} cross-section, due to choice and variations of PDFs, are shown in the limit figures for illustration purposes. From Section 1.2.1 we know that the relative resonance widths of these three models are 3%, 0.5% and 1.2%, respectively. Therefore, the lower mass limits on the Z' in these models, can be calculated as the intersection between the theoretical cross-section curve and the observed and expected limit curves of corresponding signal width. The resulting mass limits are summarised in Table 5.4. These mass limits improve previous ATLAS results for the same models by between 500 GeV and 800 GeV [10].

 Table 5.4.: Lower limits on the Z' mass, in TeV, for the models Z'_{SSM} , Z'_{ψ} , and Z'_{χ} .

<i>Model</i>	ee		$\mu\mu$		ll	
	<i>Expected</i>	<i>Observed</i>	<i>Expected</i>	<i>Observed</i>	<i>Expected</i>	<i>Observed</i>
Z'_{SSM}	4.9	4.9	4.5	4.5	5.1	5.1
Z'_{ψ}	4.3	4.1	4.0	4.0	4.5	4.5
Z'_{χ}	4.6	4.6	4.2	4.2	4.8	4.8




Figure 5.13.: Upper limits at 95% confidence level on the fiducial cross-section times branching ratio for zero-width signals in the electron (a), muon (b), and combined (c) channels [6].




Figure 5.14.: Upper limits at 95% confidence level on the fiducial cross-section times branching ratio for tested signal widths of zero, 3%, and 10% relative width (a), and 0.5%, 1.2%, and 6% relative width (b), in the electron channel. Also included are predicted fiducial cross-section values for the heavy boson benchmark models Z'_{SSM} (a), and Z'_ψ and Z'_χ (b) [6].




Figure 5.15.: Upper limits at 95% confidence level on the fiducial cross-section times branching ratio for tested signal widths of zero, 3%, and 10% relative width (a), and 0.5%, 1.2%, and 6% relative width (b), in the muon channel. Also included are predicted fiducial cross-section values for the heavy boson benchmark models Z'_{SSM} (a), and Z'_ψ and Z'_χ (b) [6].




Figure 5.16.: Upper limits at 95% confidence level on the fiducial cross-section times branching ratio for tested signal widths of zero, 3%, and 10% relative width (a), and 0.5%, 1.2%, and 6% relative width (b), in the combined channel. Also included are predicted fiducial cross-section values for the heavy boson benchmark models Z'_{SSM} (a), and Z'_ψ and Z'_χ (b) [6].

6. Non-Resonant Dilepton Analysis

In the previous chapter, we presented a search for localised resonant new physics phenomena. For theories such as the quark compositeness and large extra dimensions, presented in Chapter 1, dilepton final state signals are predicted to manifest as broad deviations from the SM predictions in the high invariant mass range. The background estimation method used in the resonant analysis is less suited in such searches, as the broad nature of the signal in a low statistics region in data risks biasing the background fit. After publishing the results from the resonant analysis, the ATLAS exotic dilepton working group shifted its focus to the search for non-resonant signals, using the same full Run 2 dataset, and the same object definition and event selection scheme, as the resonant analysis. In the non-resonant analysis, a model-independent search was carried out for high invariant mass dilepton excesses, as well as a search for signatures of contact interactions (CIs) and gravitons originating from the ADD model of large extra dimensions. Due to time constraints, the part of the analysis interpreting the results in the framework of the ADD model, which is my main contribution, was not published with the CI result of the search, in Ref. [7], but as a so-called ATLAS public note, in Ref. [9]. The contact interaction analysis sets a lower limit on the energy scale parameter Λ of an effective field theory (EFT) contact interaction (CI) model. CI processes with both constructive and destructive interference with the SM are considered in a single bin counting experiment on the invariant mass variable. This chapter focuses on the results from the search for non-resonant ADD model graviton production, aiming to set a lower limit on the string scale parameter M_S of the ADD theory of gravity in large extra dimensions. The dilepton invariant mass distributions for three ADD model hypotheses of M_S are shown with the Run 2 data and the smooth background MC histogram estimate in Figure 6.1.

6.1. Background Estimation

This analysis employs a data-driven method to estimate the SM background to the potential new physics signal. Unlike in the resonant analysis, where we can make no prior predictions as to the pole mass of a potential new signal, for the models considered in the non-resonant case, the signal is known to be located at high dilepton invariant mass, as broad deviations from the SM prediction. We can therefore use the low invariant mass region to constrain the background fit.

The background fit procedure is as follows: First, a parametric function is fitted in a low invariant mass control region (CR). The best-fit parameters of the background function are only determined in this CR. Next, the background function is extrapolated into a single bin, high-invariant mass signal region (SR). A schematic of the extrapolation procedure is shown in Figure 6.2, using example values for the CR and SR ranges. The figure shows the invariant mass distribution of a CI signal which destructively interferes with the SM. In the figure, the signal hypothesis predicts fewer events than the background prediction for invariant mass values between 1 and 2 TeV. To avoid having this negative signal prediction




Figure 6.1.: Data-to-background comparison in the electron (a) and muon (b) channels, along with the predicted invariant mass distributions of three string scale hypothesis ADD models in the GRW convention [9].

cancel any potential positive signal yield in the SR, the control and signal regions may be separated by a gap in invariant mass.




Figure 6.2.: Schematic of the background estimation extrapolation procedure [7].

Around 50 parametric functions are considered for the background fit. Each function candidate is fitted to the smooth background template described in Chapter 3 in a set of 15 provisional CR and SR configurations. The residual of the background fit with regards to the smooth background template is used to evaluate the candidate functions. For a function to be considered, the relative residual between the fit and the template in each bin is required to be less than three in both the CR and SR. Five parametric functions pass this requirement, one of which is the function chosen to model the background in the resonant analysis, presented in the previous chapter. This function, given in Equation (5.7), is again chosen to model the SM background of the analysis. However, the parameter b , previously a free parameter determined in the fit, is now fixed. The optimal choice of b

is determined by fitting the background function, given in Equation (5.7), to the smooth MC template. The fixed parameter C is determined in the same way.

Once the background parametric function has been determined, the choice of CR and SR is optimised separately for a CI signal with constructive and destructive interference with the SM. The interference of a CI model is determined by the sign of η_{ij} in Equation (1.79), given in Chapter 1. The optimised choices for the control and signal regions are given in Table 6.1. Destructive interference models are seen to prefer a separation of around 1.3 TeV between the control and signal regions, while constructive interference models prefer adjoining regions. The optimal choices of CRs and SRs are not dependent on the chirality of the CI models. Therefore, only one set of CRs and SRs is used in the constructive and destructive case, for the four possible combinations of left- and right-handed fermion fields present in the Lagrangian presented in Equation (1.78). These will be referred to in this text as the constructive and destructive SRs.

Table 6.1.: Control region and signal region definitions, in units GeV. $SR_{\max} = 6,000$ GeV for each signal region.

<i>Channel</i>	<i>Constructive Interference</i>			<i>Destructive Interference</i>		
	CR_{\min}	CR_{\max}	SR_{\min}	CR_{\min}	CR_{\max}	SR_{\min}
ee	280	2,200	2,200	310	1,450	2,770
$\mu\mu$	310	2,070	2,070	320	1,250	2,570

6.2. Processing ADD Theory Signal Samples

We produce MC signal samples for seven ADD string scale hypotheses: 3 TeV, 4 TeV, 5 TeV, 6 TeV, 7 TeV, 8 TeV, and 50 TeV. Details about the signal sample production are presented in Appendix B. After passing the event selection step, presented in Chapter 4, the signal samples are scaled to the 139 fb^{-1} luminosity of the Run 2 ATLAS dataset, and then binned in invariant mass like the dilepton dataset. The matrix element of the signal samples is calculated at NLO precision. The invariant mass signal templates are scaled by an invariant mass-dependent K -factor, to account for higher-order QCD corrections to the ADD cross-section. The K-factor is calculated using VRAP 0.9 [111] and the CT14NNLO PDF set [78]. The K-factor, plotted as a function of dilepton invariant mass in Figure 6.3, is the same K-factor used to scale our DY background MC samples, introduced in Section 3.1 of Chapter 3. Having applied the K-factor, the ADD model MC samples now correspond to samples generated with the CT14NNLO PDF set. The K-factor is derived on MC samples generated using the POWHEG-BOX and PYTHIA8 event generators, while the ADD model MC samples are generated using the SHERPA event generator. After a side by side comparison between a DY sample generated using POWHEG-BOX+PYTHIA8 and CT10NLO, and the 50 TeV signal sample, generated using SHERPA and CT10NLO, we conclude that the K-factor can be applied to the ADD signal samples. Applying the K-factor to the ADD samples is shown to have less than a 1% impact on the limits set on the string scale parameter of the ADD theory.

The matrix element of the SHERPA generator used to generate the signal samples contains a SM DY component that cannot be disentangled at the generator level. A fraction of the dilepton data set is made up of DY events. The DY component of the signal samples




Figure 6.3.: The K-factor used to account for higher-order cross-section corrections in the ADD theory signal samples, plotted as a function of dilepton invariant mass.

must therefore be removed to avoid double-counting events in the final statistical analysis. When setting the string scale parameter to 50 TeV, the resulting MC sample is virtually entirely made up of DY events, as the signal is moved beyond the energy reach of the LHC. We therefore call the $M_S = 50$ TeV sample the "signal-free" sample. The DY component in the signal samples is removed by subtracting the shape of the invariant mass distribution of the signal-free MC sample from each signal sample. The remaining events should be purely due to KK production of graviton excitations. The invariant mass distributions of the ADD model signal samples, before and after the DY subtraction step, are shown for the electron channel in Figure 6.4, and for the muon channel in Figure 6.5.




Figure 6.4.: Electron channel invariant mass distributions before (a) and after (b) DY subtraction, for all ADD theory string scale hypotheses. The signal histograms are drawn with a 100 GeV bin width.

The signal event yield in the constructive and destructive SRs, for each of the six string scale hypotheses, are given in Table 6.2. Naively, the SR event yield is expected to fall with increasing string scale. However, in the electron channel, more events are seen when




Figure 6.5.: Muon channel invariant mass distributions before (a) and after (b) DY subtraction, for all ADD theory string scale hypotheses. The signal histograms are drawn with a 100 GeV bin width.

integrating the $M_S = 4$ TeV sample in the destructive SR, than for the 3 TeV sample. Calculating the cross-section of the signal samples involves a sum over many KK modes. This sum must be cut off at some value so as not to become infinite. Therefore, the samples have a kinematic cutoff at the string scale hypothesis. The ADD model is not valid above this cutoff. The starting point of the destructive SR is close to the kinematic edge of the 3 TeV signal sample. As a result, a large portion of the events in the 3 TeV sample falls below the destructive SR interval. The $M_S = 3$ TeV signal sample is not used in the statistical analysis for this reason.

Table 6.2.: Event yield in the constructive and destructive SRs, for the six string scale hypotheses considered in this search. The cross-sections of the ADD samples are calculated using the GRW convention of summing over KK states.

<i>String Scale</i> [GeV]	<i>Electron Channel</i>		<i>Muon Channel</i>	
	<i>Const. SR</i>	<i>Dest. SR</i>	<i>Const. SR</i>	<i>Dest. SR</i>
3,000	868.1	125.4	618.9	232.9
4,000	233.0	136.5	148.5	97.2
5,000	55.5	38.7	33.1	24.8
6,000	14.5	10.9	10.5	6.5
7,000	4.1	3.1	3.1	2.0
8,000	1.3	1.0	1.8	0.6

6.2.1. Kaluza-Klein Formalism Conversion

The ADD model signal MC samples are generated in the GRW convention for summing over KK states. We also want to be able to express the results of the search in the Hewett and HLZ conventions, without having to produce a dedicated set of signal MC samples for each convention. The total ADD model production cross-section is given in Equation

(1.74) in Chapter 1. After removing the DY component, this cross-section becomes

$$\sigma_{\text{ADD}} = \mathcal{F} \frac{F_{\text{Int}}}{M_{\text{S}}^4} + \mathcal{F}^2 \frac{F_{\text{G}}}{M_{\text{S}}^8}. \quad (6.1)$$

The first term in this equation is referred to as the "interference" term, of the ADD cross-section. This term controls the amount of interference between the KK graviton and the SM term. The second term is known as the "pure graviton" term. The value of \mathcal{F} , given in Equation (1.75) in Section 1.2.2, is determined by the chosen KK convention.

The interference and pure graviton terms both contribute to the events in the ADD signal production, but the ratio between the terms is unknown. Therefore, we cannot convert the number of events predicted in the GRW convention to that predicted in the Hewett and HLZ conventions directly. This would require separating the contribution to the cross-section from the interference and pure graviton terms. Instead, two cases are considered: one where the cross-section scales "interference-like" between the GRW convention and the other convention, and one where the cross-section scales "graviton-like". To obtain the Hewett and HLZ signal template in the interference cross-section scheme, we scale the GRW invariant mass signal histograms by their respective \mathcal{F} values. Similarly, to convert the signal templates in the pure graviton scheme, we scale the GRW histograms by \mathcal{F}^2 . The true value of the Hewett and HLZ cross-sections are bookended by their respective interference- and pure graviton cross-sections. Note that when $n = 4$, $\mathcal{F} = 1$ in the HLZ convention, and the cross-section becomes identical to that in the GRW convention.

Signal templates for some KK conventions cannot be created using the method described here. In the Hewett convention, if the model parameter $\lambda = -1$, the scaling factor \mathcal{F} becomes negative in the interference cross-section scheme. The same thing happens in the HLZ convention when the number of extra dimensions n is equal to two. These conventions are therefore dropped from the analysis. Astronomical observations rule out $n = 1$, as this would modify Newton's law of gravitation at solar system distance scales [11].

6.3. Systematic Uncertainties

6.3.1. Background Uncertainties

The statistical uncertainty on the extrapolated background shape is determined by drawing pseudo-data sets from the nominal background fit to the data. Each of these toy distributions is refitted in the CR and extrapolated into the SR. The mean of the SR yield distribution of the pseudoexperiments are confirmed to be centred on the nominal background expectation, and the standard deviation of the toy background distribution in the SR is taken as the statistical uncertainty. In the electron channel, the relative statistical uncertainty is found to be 14% in the constructive SR and 35% in the destructive signal region. In the muon channel, the relative statistical uncertainty is 21% in the constructive SR and 58% in the destructive SR.

Like in the resonant search presented in Chapter 5, we run the risk of observing a spurious signal, not due to any real signal being present in the data, but due to background mismodelling. The spurious signal uncertainty is quantified and included in the statistical analysis. An ensemble of 10,000 background templates is constructed from a linear combination of the nominal background and all its variations by systematic uncertainties. The systematic uncertainties are the same as for the resonant search presented in Chapter 5.

The weight of each non-nominal variation in the linear combination is drawn between -1 and 1 from a normal distribution centred at zero, with a standard deviation of one. Each pseudo-background template is fitted using the parametric background function in the CR and extrapolated into the SR. The difference in SR events between the pseudo-background template and the fit is taken as the spurious signal. The mean and standard deviation of the spurious signal distribution for all pseudo-background templates are added in quadrature. This quadratic sum is taken to be the spurious signal uncertainty on the background. The mean of the pseudo-background spurious signal distribution is included to take into account any systematic shift in the spurious signal distribution.

The presence of a signal in the data may bias the background fit in the CR. A CR bias uncertainty is measured by taking the difference in the SR between a fit to data in the CR of the background function alone, and a fit to data using the background function with an added signal CI signal component. The resulting relative uncertainty is less than 4% in both the electron and muon channels, and for both constructive and destructive interference CI signal hypotheses. More information about the background uncertainties used in the analysis can be found in Refs. [112] and [113].

6.3.2. Signal Uncertainties

Uncertainties on the signal event yield are divided into two categories: experimental uncertainties and theoretical uncertainties. Experimental uncertainties, in this case, refers to uncertainties coming from the ATLAS measurements themselves and reconstructing physics objects. These uncertainties include uncertainty on the measurement of the luminosity, lepton scale and resolution uncertainties, tracking uncertainty in the Inner Detector and Muon Spectrometers, reconstruction and isolation uncertainties, etc. Theoretical uncertainties are uncertainties due to the choice and variation of the PDF set used for the signal sample event generator. Only experimental signal uncertainties are used in the analysis. The theoretical PDF uncertainties are calculated and quoted for completeness.

Experimental Uncertainties

The uncertainty on the signal yield is found by taking the sum in quadrature of the difference between the nominal signal yield and the signal yield obtained in all variations of the signal template under each category of uncertainty. In the case of CIs, the impact on the CI signal templates of the experimental uncertainty variations were taken to be sufficiently similar to their effect on NLO DY MC samples. Therefore, the experimental signal uncertainties were found by variations of the DY invariant mass distributions. We assume that this assumption still holds for the ADD model signals, and use the experimental signal uncertainty found for CIs in the ADD analysis. The dominating experimental uncertainties in the electron channel are the electron identification and electron energy scale uncertainties. In the muon channel, muon reconstruction, the "good muon" selection criterion, and sagitta bias in track reconstruction are the dominant sources of experimental uncertainty. The relative experimental signal uncertainty in the electron channel is found to be 8% in both the constructive and destructive SRs. In the muon channel, the relative experimental signal uncertainty is $^{+20\%}_{-17\%}$ in the constructive SR and $^{+27\%}_{-22\%}$ in the destructive SR.

The relative background and signal uncertainties are presented in Table 6.3. The uncertainties are seen to be larger for destructive interference models than for constructive interference models. The destructive interference CR is shorter than the constructive

CR, leading to a larger background uncertainty. Similarly, the destructive SR starts at a higher invariant mass threshold than the constructive SR. This means that the integrated SR event yield is smaller in the destructive SR, leading to a larger relative uncertainty.

Table 6.3.: The relative uncertainty in each SR. Here, $\sigma_{\text{bgr}}^{\text{stat.}}$, $\sigma_{\text{bgr}}^{\text{SS}}$, and $\sigma_{\text{bgr}}^{\text{CRB}}$ is the statistical, spurious signal, and control region bias uncertainties on the background, respectively, while $\sigma_{\text{sig}}^{\text{exp.}}$ denotes the experimental signal uncertainty.

<i>Channel</i>	<i>Signal Region</i>	$\sigma_{\text{bgr}}^{\text{stat.}}$	$\sigma_{\text{bgr}}^{\text{SS}}$	$\sigma_{\text{bgr}}^{\text{CRB}}$	$\sigma_{\text{sig}}^{\text{exp.}}$
ee	Constructive SR	14%	4%	2%	8%
ee	Destructive SR	34%	7%	1%	8%
$\mu\mu$	Constructive SR	21%	6%	2%	+20%
$\mu\mu$	Destructive SR	58%	24%	4%	-17%
					+27%
					-22%

Theoretical Uncertainties

Next, we want to determine the uncertainty on the signal yield due to PDF uncertainties in the generated ADD signal MC samples. For the CT14NNLO PDF set, we have an orthogonal set of seven PDF eigenvectors with uncertainties, which parameterise the degrees of freedom in the fit of the PDF to experimental data. Another uncertainty that needs to be considered is the PDF choice uncertainty. This uncertainty is derived from comparisons between CT14NNLO and other available NNLO PDF sets. If the nominal prediction of any other PDF set falls outside of the CT14NNLO 1σ uncertainty band, the deviation from the CT14NNLO uncertainty envelope is added to the PDF uncertainty. The only PDF set that contributes to the PDF choice uncertainty is the NNPDF3.0 set, which falls outside the CT14NNLO uncertainty band above around 4 TeV invariant mass. The total PDF uncertainty on the ADD signal yield is found by adding in quadrature the relative deviation in SR event yield from all eigenvector- and choice variations of the CT14NNLO PDF. The relative uncertainty of each PDF uncertainty, as well as their sum in quadrature, is listed in Table 6.4, and can be seen in Figure 6.6. The relative PDF uncertainty varies between around 10% at low invariant mass and around 20% at high invariant mass in both the electron and muon channels. The impact on the destructive SR signal yield from a 1σ up and down variation of the PDF uncertainty is shown in Figure 6.7. Here, the PDF uncertainty on the event yield varies between 19 events at $M_S = 4$ TeV and 0.2 events at $M_S = 8$ TeV in the electron channel, and between 24 events at $M_S = 3$ TeV and 0.1 events at $M_S = 8$ TeV in the muon channel.

6.4. Statistical Analysis

Like the resonant analysis of Chapter 5, this search employs a frequentist interpretation of probability. Two statistical models are considered. The first model compares the background-only prediction to a generic, model-independent signal hypothesis. In the second model, the background-only prediction is measured against the ADD model signal hypotheses.




Figure 6.6.: Relative change in the destructive SR event yield for a GRW convention ADD signal from each PDF egevector- and choice variation, as well as their sum in quadrature, in the electron (a) and muon (b) channels. This sum is taken to be the total relative PDF uncertainty of the ADD signals.




Figure 6.7.: Absolute impact on the destructive SR event yield from a $\pm 1\sigma$ variation of the PDF uncertainty of a GRW convention ADD signal in the electron (a) and muon (b) channels.

Table 6.4.: Relative PDF uncertainty of the ADD model samples, in the destructive SR.

<i>Channel</i>	M_S [TeV]	$\sigma_{\text{sig}}^{\text{PDF}}$
ee	4	13.7%
$\mu\mu$		12.9%
ee	5	16.3%
$\mu\mu$		15.3%
ee	6	19.6%
$\mu\mu$		18.1%
ee	7	20.5%
$\mu\mu$		18.7%
ee	8	20.7%
$\mu\mu$		20.6%

6.4.1. Discovery

For a model-independent inference of the signal event yield n_{sig} in the SR, the following likelihood function is constructed:

$$\mathcal{L}(n_{\text{sig}}, \theta_{\text{bgr}} | n) = \text{Pois}(n | n_{\text{sig}} + (1 + \theta_{\text{bgr}})n_{\text{bgr}}) \times \text{Gauss}(\theta_{\text{bgr}} | \mu_{\text{bgr}}^\theta, \sigma_{\text{bgr}}). \quad (6.2)$$

Here, n is the observed number of events in the SR. $\text{Pois}()$ is the probability mass function of a Poisson distribution and $\text{Gauss}()$ is the probability density function of a normal distribution. The expected number of background events in the SR, extracted from the extrapolation procedure outlined in Section 6.1, is denoted by n_{bgr} . The nuisance parameter θ_{bgr} accounts for the relative uncertainty on the background, σ_{bgr} . The value of n_{bgr} is a so-called *auxiliary measurement*, being completely independent of the main measurement n . The value of the auxiliary measurement is parameterised using μ_{bgr}^θ , the mean of the Gaussian nuisance term.

The likelihood function of Equation (6.2) is used in the discovery step, or search phase, of the analysis. The likelihood ratio test statistic of Equation (5.12) is again used to quantify the validity of the null hypothesis. In the resonant analysis of Chapter 5, we use the asymptotic approximation to calculate the local significance at each pole mass hypothesis. Here, we instead use an ensemble of pseudoexperiments to create the probability density function $f(q_0 | H_0)$ of the test statistic under the background-only hypothesis. First, we determine the best-fit value of the background nuisance parameter, $\hat{\theta}_{\text{bgr},0}$, such that it maximises the likelihood under the null hypothesis, $n_{\text{sig}} = 0$. The background-only pseudoexperiments are drawn under this assumption for the background nuisance parameter. However, there is a subtle issue regarding the treatment of nuisance parameters of auxiliary measurements when constructing ensembles of pseudoexperiments; one can use the *conditional ensemble*, where the auxiliary measurement μ_{bgr}^θ is fixed at $\hat{\theta}_{\text{bgr},0}$ when calculating the test statistic, or the *unconditional ensemble*, where μ_{bgr}^θ is drawn from a Gaussian distribution, centred at $\hat{\theta}_{\text{bgr},0}$ and of width σ_{bgr} . Both ensembles are valid constructions in frequentist statistics. However, the result of the asymptotic approximation corresponds to that of the unconditional ensemble. The unconditional ensemble approach is therefore recommended by the ATLAS Statistics Forum and is what we use to obtain the results here [110]. We generate 100,000 pseudoexperiments in each of the four SRs, given in Table 6.1. The p -value of the data is found by integrating $f(q_0 | H_0)$ from the test

statistic measured on the data, q_0^{obs} , to infinity, using Equation (5.13). This p -value is then converted into a significance value using Equation (5.14). If no significant deviation from the SM is seen in the data in the discovery step outlined above, we move on to the exclusion step.

6.4.2. Exclusion

A model-independent upper limit on the number of signal events can be obtained using the likelihood function presented in Equation (6.2). We also aim to set a lower limit on the value of the string scale parameter of the ADD models. For that, we extend the likelihood function from Equation (6.2) to include the ADD signal prediction, as

$$\begin{aligned} \mathcal{L}(M_S, \boldsymbol{\theta}|n) = & \text{Pois}(n|(1 + \theta_{\text{sig}})n_{\text{sig}}(M_S) + (1 + \theta_{\text{bgr}})n_{\text{bgr}}) \\ & \times \text{Gauss}(\theta_{\text{bgr}}|\mu_{\text{bgr}}^\theta, \sigma_{\text{bgr}}) \times \text{Gauss}(\theta_{\text{sig}}|\mu_{\text{sig}}^\theta, \sigma_{\text{sig}}). \end{aligned} \quad (6.3)$$

Here, $n_{\text{sig}}(M_S)$ is the number of integrated events in the signal region predicted at some string scale value M_S . The new nuisance parameter θ_{sig} accounts for the uncertainty on this signal yield, and is constrained, like the background nuisance parameter, by a Gaussian factor in the likelihood function.

Calculating $n_{\text{sig}}(M_S)$ requires a continuous description of the integrated number of signal events in the SR for arbitrary values of M_S . Initially, we only know the value of n_{sig} for the six values of M_S for which we have generated signal template histograms. Intermediate values of n_{sig} between these points are found by performing a bin-by-bin exponential interpolation between neighbouring ADD signal samples. The signal histogram at an arbitrary string scale M_S , is constructed from the signal histograms at $M_{S,a}$ and $M_{S,b}$, the string scale values for which we have existing signal sample histograms above and below M_S , respectively. For example, when constructing the histogram for $M_S = 5.5$ TeV, $M_{S,a} = 6$ TeV and $M_{S,b} = 5$ TeV. The bin content of each bin $h_i(M_S)$ is determined by

$$h_i(M_S) = \exp\left(\log(h_i(M_{S,b})) + \frac{\log(h_i(M_{S,a})) - \log(h_i(M_{S,b}))}{M_{S,a} - M_{S,b}}(M_S - M_{S,b})\right). \quad (6.4)$$

Here, $h_i(M_{S,a})$ and $h_i(M_{S,b})$ are the bin contents of the i -th bin in the $M_{S,a}$ and $M_{S,b}$ histograms, respectively. Finally, $n_{\text{sig}}(M_S)$ is found by integrating the M_S histogram in the SR.

The kinematic cutoff at M_S in the ADD signal samples leads to a possible mismodelling effect in the histogram morphing procedure outlined above. The simple bin-by-bin interpolation is not able to accurately model such a shape effect. The mismodelling will only affect the integrated SR count of interpolated histograms if the kinematic cutoff of $M_{S,a}$ or $M_{S,b}$ is below the upper SR cut. To gauge the severity of the mismodelling effect, the SR integral of the morphed histograms are compared to the SR integral of the generated ADD signal histograms for $M_S = 5$ TeV and $M_S = 6$ TeV. The $M_S = 5$ TeV and $M_S = 6$ TeV morphed histograms are morphed between the generated $M_S = 4$ TeV and $M_S = 6$ TeV, and $M_S = 5$ TeV and $M_S = 7$ TeV histograms, respectively. The morphed and generated $M_S = 6$ TeV histogram templates are compared in the electron and muon channels in Figure 6.8. In the figure, the morphed 6 TeV histogram is shown as a solid orange line, while the MC generated histogram is shown as grey circles. Also included in the figure are the generated histograms for the string scale values above and below $M_S = 6$ TeV. In the electron channel, we see that the morphed 6 TeV histogram fails to accurately follow the

shape of the generated 6 TeV histogram above 5 TeV, the kinematic cutoff of the generated 5 TeV sample. This effect is not as pronounced in the muon channel, where the kinematic cutoffs of the generated ADD signal samples are smeared out by the lower detector mass resolution. By eye, the morphed electron channel histogram seems a worse fit than the muon channel histogram. However, in the electron channel, the SR integrals differ by 0.7% between the morphed and generated histograms. In the muon channel, the SR integrals differ by 7.2%. At $M_S = 5$ TeV, the difference in SR integral between the morphed and generated histograms is 0.2% and 1.2% in the electron and muon channel, respectively. A 7.2% difference in the SR integral may seem like a problematic discrepancy brought on by the kinematic edge mismodelling effect. However, if we move the low SR edge to 3 TeV, and compare the morphed and generated histogram integrals, the SR integral discrepancy becomes less than 1% for both channels. This suggest that the largest source of uncertainty in the histogram morphing procedure is low mass statistical uncertainty, and not mismodelling of the kinematic cutoff. This large statistical uncertainty stems from an insufficient number of events generated in the low mass region, in combination with the DY subtraction procedure. Later, when setting limits on M_S in different KK convention schemes, we will see that most of the limits are larger than 6 TeV, where no mismodelling can occur due to the kinematic cutoffs.




Figure 6.8.: Exponential histogram interpolation of a $M_S = 6$ TeV signal sample in the electron (a) and muon (b) channels. Also included in the plots are the MC simulated $M_S = 5$ TeV, $M_S = 6$ TeV, and $M_S = 7$ TeV signal samples. The samples at 5 TeV and 7 TeV are used to generate the morphed 6 TeV sample, shown as a solid orange line.

To set limits on M_S , we use the likelihood function defined in Equation (6.3) and the q_μ likelihood ratio test statistic defined in Equation (5.18) in Chapter 5. Now, M_S takes the place of the signal strength parameter μ . Limits are set at 95% CL using the CL_s method. We use pseudoexperiments to determine the distribution of q_μ under the null- and alternative hypotheses, rather than using the asymptotic approximation. The lower limit on the string scale, at 95% CL, is taken to be the value M_S at which the ratio of confidence levels CL_s becomes 0.05. To find this limit, we vary M_S between 4 TeV and 8 TeV in 100 steps. At each step, we draw the pseudoexperiment ensemble and calculate CL_s using Equation (5.20). The lower limit on M_S is the value for which CL_s crosses the 0.05 threshold.

Drawing pseudoexperiments can use a lot of computing resources, but having a small pseudoexperiment ensemble may lead to large statistical fluctuations of the calculated

limit value. We want to determine how many pseudoexperiments we need for the limits to converge and to remain stable. To do this we calculate and plot the M_S limit as a function of the number of pseudoexperiments drawn in each step of the limit setting procedure. Here, the limits are calculated in the GRW KK convention. The size of the pseudoexperiment ensemble is varied in 14 steps between 1,000 and 50,000. A new ensemble is drawn at each step. The M_S limit is plotted as a function of pseudoexperiment ensemble size in Figure 6.9. The limit is seen to converge for ensemble sizes of around 10,000 and above. We generate 25,000 pseudoexperiments at each step in M_S to calculate the limits.




Figure 6.9.: 95% CL limit on M_S as a function of pseudoexperiment ensemble size, in the electron (a) and muon (b) channels. The limits are calculated for the GRW KK convention.

6.5. Results

The dilepton invariant mass histogram of the full Run 2 data set is shown for both lepton channels in Figure 6.10. The figures also show the best-fit background estimate in the destructive control and signal regions, and the predictions of three ADD model string scale hypotheses, $M_S = 4$ TeV, $M_S = 6$ TeV, and $M_S = 8$ TeV, in the GRW convention.

The best-fit values for the free parameters in the CR background fit is given for the CR associated with each of the four SRs in Table 6.5. In addition to the parameters floating in the fit, the optimal value of the fixed parameter b is found to be 6.1 in the electron channel and 1.3 in the muon channels. In the muon channel, the fixed parameter C is set to $1/3$ in the background fit. This is the same as in the resonant analysis of Chapter 5. In the electron channel, however, C is set to $1/2$, and not to 1 like in the resonant analysis.

Table 6.5.: Best-fit values of the free parameters in the CR background fit [7].

<i>parameter</i>	<i>ee Constructive</i>	<i>ee Destructive</i>	$\mu\mu$ <i>Constructive</i>	$\mu\mu$ <i>Destructive</i>
Norm.	$(6.17 \pm 0.02) \times 10^{-3}$	$(7.87 \pm 0.03) \times 10^{-3}$	$(6.90 \pm 0.03) \times 10^{-6}$	$(4.39 \pm 0.02) \times 10^{-7}$
p_0	-12.2 ± 0.1	-12.1 ± 0.1	-14.9 ± 0.2	-17.0 ± 0.2
p_1	-4.14 ± 0.02	-4.16 ± 0.03	-4.42 ± 0.04	-4.70 ± 0.04
p_2	-0.948 ± 0.005	-0.945 ± 0.006	-0.927 ± 0.008	-0.846 ± 0.008
p_3	-0.0840 ± 0.0008	-0.082 ± 0.001	-0.081 ± 0.001	-0.064 ± 0.001




Figure 6.10.: The dilepton invariant mass distribution of the data, with the best-fit background shape in the destructive control and signal regions, and three ADD model signal hypotheses, in the electron (a) and muon (b) channels [9].

Figure 6.11 shows the distributions of the q_0 test statistic calculated for 100,000 background hypothesis pseudoexperiments in the constructive and destructive SRs in the electron channel. Figure 6.12 shows the corresponding distributions in the muon channel. The p -value of each observation is calculated as the fraction of pseudoexperiments with test statistic values larger than or equal to the test statistic observed in the data, $q_0 \geq q_0^{\text{obs}}$. The expected and observed number of events in each SR are given, along with the significance of the observation, in Table 6.6. No significant excess above the SM expectation is seen.




Figure 6.11.: Electron channel test statistic distributions of 100,000 background hypothesis pseudoexperiments in the constructive (a) and destructive (b) SRs. The red bars show pseudoexperiments for which $q_0 \geq q_0^{\text{obs}}$.




Figure 6.12.: Muon channel test statistic distributions of 100,000 background hypothesis pseudoexperiments in the constructive (a) and destructive (b) SRs. The red bars show pseudoexperiments for which $q_0 \geq q_0^{\text{obs}}$.

Table 6.6.: Expected and observed events in all SRs and the statistical significance of each observation.

<i>Signal Region</i>	<i>Expected Bgr. Events</i>	<i>Observed Events</i>	<i>Significance</i> [σ]
Electron Constructive	12.4 ± 1.9	19	1.53
Electron Destructive	3.1 ± 1.1	2	-0.50
Muon Constructive	9.6 ± 2.1	6	-1.00
Muon Destructive	1.4 ± 0.9	1	-0.20

6.6. Setting Limits on the String Scale Parameter

The model-independent upper limit on the number of signal events, n_{sig} is drawn in Figure 6.13. The expected limits, with 1σ and 2σ uncertainty bands¹, and the observed limits, in all four SRs, are listed in Table 6.7.

Table 6.7.: Expected and observed 95% CL upper limits on n_{sig} in all SRs. The first set of uncertainties on the expected limit values are their $\pm 1\sigma$ uncertainties, while the second set are their $\pm 2\sigma$ uncertainties [7].

<i>Signal Region</i>	<i>Expected n_{sig} Limit</i>	<i>Observed n_{sig} Limit</i>
Electron Constructive	$9.3^{+3.8+8.5}_{-2.5-3.9}$	16.0
Electron Destructive	$5.0^{+2.0+4.8}_{-1.2-1.9}$	4.4
Muon Constructive	$8.0^{+3.3+7.5}_{-2.2-3.5}$	5.8
Muon Destructive	$4.0^{+1.6+3.7}_{-0.8-1.6}$	3.8




Figure 6.13.: Model-independent upper limits on n_{sig} in all SRs [7].

Next, we want to calculate the lower limit on the M_S parameter in each convention of summing KK states, GRW, Hewett, and HLZ, in both the pure graviton and interference cross-section schemes. The results of two methods for calculating these limits are presented in the following sections. First, a method converting the model-independent limit on n_{sig} into limits on M_S using an exponential fit to the SR yield of each generated ADD model signal sample. The second method is the CL_s method outlined in Sections 5.6 and 6.4.

The invariant mass distributions of the ADD model signal samples are harder than those of the CI model signals. For the ADD model samples, the signal-to-background separation

¹The uncertainties on the expected limits are not strictly uncertainties on the limit, but rather represent the distribution of possible limits under the background-only hypothesis

is better for a higher SR starting point. This is reflected in the limits on M_S , which are all found to be more stringent in the destructive SR than in the constructive SR, which has a lower starting point in invariant mass. Therefore, only the limits found in the destructive signal region are presented here.

6.6.1. Reinterpreting Model-Independent Limits

In addition to the frequentist CL_s method for calculating the lower limit on the ADD string scale parameter, we also calculate the limit from the model-independent upper limit on n_{sig} . This procedure, referred to from this point as the "conversion" method, was developed in the early stages of this analysis, before being abandoned in favour of the CL_s method. The limits on M_S obtained using the model-independent conversion method are given here as a cross-check of the final M_S limit results, presented in the next section.

To calculate the lower limits on M_S in the conversion method, we first calculate the SR event yield for each of the generated ADD signal histograms, between $M_S = 4$ TeV and $M_S = 8$ TeV. Next, we interpolate these signal yield values using an exponential fit. An exponential function is determined by eye to accurately describe the SR yield distribution as a function of string scale. Two such exponential fits are shown in Figure 6.14 for the signal samples following the GRW KK convention. The intersection of the signal yield exponential fit and the expected and observed n_{sig} upper limits define the corresponding lower limits on M_S . These intersection points can be seen for the GRW case in Figure 6.14. Similarly, the error bands on the expected M_S limit are defined by the intersection of the signal yield exponential fit and the expected n_{sig} limit uncertainty bands. The systematic uncertainty on the signal is added in quadrature to the exponential fit uncertainty. However, this uncertainty is not considered when converting the limits.





Figure 6.14.: Conversion method limit intersection for the GRW convention in the electron (a) and muon (b) channels. Corresponding plots are made for each choice of KK convention and in both the interference and pure graviton cross-section schemes. These figures are omitted from the text for brevity.


The above steps are carried out for each KK convention, and in both the interference and pure graviton cross-section schemes. The resulting expected and observed limits on M_S are listed in Table 6.8 and plotted in Figures 6.15 and 6.16, for the pure graviton and interference schemes, respectively.

Table 6.8.: Lower M_S limits, in TeV, calculated using the conversion method from the upper n_{sig} limits in the destructive SR.

Channel	Cross-section Term	GRW	Hewett		HLZ			
			$\lambda = +1$	$n = 3$	$n = 4$	$n = 5$	$n = 6$	$n = 7$
Exp. ee	Interference	6.6	6.3	7.2	6.6	6.3	6.1	5.9
Obs. ee		6.6	6.3	7.2	6.6	6.3	6.1	5.9
Exp. ee	Graviton	6.6	5.9	7.7	6.6	6.0	5.5	5.2
Obs. ee		6.6	5.9	7.8	6.6	6.0	5.5	5.2
Exp. $\mu\mu$	Interference	6.4	6.1	7.0	6.4	6.1	5.9	5.7
Obs. $\mu\mu$		6.5	6.1	7.0	6.5	6.2	6.0	5.8
Exp. $\mu\mu$	Graviton	6.4	5.8	7.5	6.4	5.8	5.4	5.0
Obs. $\mu\mu$		6.5	5.8	7.5	6.5	5.9	5.4	5.1



(a)



(b)

 Figure 6.15.: Electron- (a) and muon (b) channel lower M_S limits, set using the conversion method, in the pure graviton cross-section scheme. The error bands on the expected limits are only due to the uncertainty on the expected n_{sig} limit.




Figure 6.16.: Electron- (a) and muon (b) channel lower M_S limits, set using the conversion method, in the interference cross-section scheme. The error bands on the expected limits are only due to the uncertainty on the expected n_{sig} limit.

6.6.2. Direct M_S Limits

The main reason for moving away from the conversion method to the CL_s method is the ad hoc exclusion of signal systematic uncertainties. The CL_s method, using a profile likelihood ratio test statistic, is also more in line with recommended statistical practices in the ATLAS Collaboration and with the statistical analysis performed in the search for CIs.

Using the frequentist limit setting procedure outlined in Sections 6.4 and 5.6, we set expected and observed lower limits on M_S in both the interference and graviton cross-section schemes and for all conventions of summing KK states. The limits obtained using the asymptotic formulae are listed in Table 6.9. The asymptotic limits are used to validate the results obtained using pseudoexperiment ensembles and are presented here for completeness. The limits obtained using pseudoexperiments are listed in Table 6.10. In the electron channel, using the GRW KK convention, the expected and observed lower M_S limits are set at 6.5 and 6.6 TeV, respectively. In the muon channel, the corresponding expected and observed lower M_S limits are set at 6.3 and 6.4 TeV.

Table 6.9.: Lower M_S limits, in TeV, obtained using the asymptotic approximation.

<i>Channel</i>	<i>Cross-section Term</i>	GRW	Hewett		HLZ				
			$\lambda = +1$	$n = 3$	$n = 4$	$n = 5$	$n = 6$	$n = 7$	
Exp: ee	Interference	6.5	6.2	7.0	6.5	6.2	6.0	5.9	
Obs: ee		6.7	6.3	7.2	6.7	6.4	6.1	6.0	
Exp: ee	Graviton	6.5	5.9	7.5	6.5	5.9	5.5	5.1	
Obs: ee		6.7	6.0	7.8	6.7	6.0	5.6	5.2	
Exp: $\mu\mu$	Interference	6.3	6.0	6.8	6.3	6.0	5.8	5.6	
Obs: $\mu\mu$		6.4	6.1	7.0	6.4	6.1	5.9	5.7	
Exp: $\mu\mu$	Graviton	6.3	5.6	7.3	6.3	5.7	5.3	4.9	
Obs: $\mu\mu$		6.4	5.7	7.5	6.4	5.8	5.4	5.0	

Table 6.10.: Lower M_S limits, in TeV, obtained using pseudoexperiments.

Channel	Cross-section Term	GRW	Hewett		HLZ			
			$\lambda = +1$	$n = 3$	$n = 4$	$n = 5$	$n = 6$	$n = 7$
Exp: ee	Interference	6.5	6.2	7.0	6.5	6.2	6.0	5.8
Obs: ee		6.6	6.3	7.2	6.6	6.3	6.1	5.9
Exp: ee	Graviton	6.5	5.8	7.5	6.5	5.9	5.4	5.1
Obs: ee		6.6	5.9	7.7	6.6	6.0	5.5	5.2
Exp: $\mu\mu$	Interference	6.3	5.9	6.8	6.3	6.0	5.7	5.6
Obs: $\mu\mu$		6.4	6.0	6.9	6.4	6.1	5.8	5.7
Exp: $\mu\mu$	Graviton	6.3	5.6	7.2	6.3	5.7	5.2	4.9
Obs: $\mu\mu$		6.4	5.7	7.5	6.4	5.8	5.3	5.0

The expected and observed lower limits on M_S are plotted, for all KK conventions, in the pure graviton cross-section scheme in Figure 6.17 and in the interference cross-section scheme in Figure 6.18. Also included in the figures are the M_S limits obtained by the ATLAS exotic dilepton working group in 2014, using 20 fb^{-1} of data collected at $\sqrt{s} = 8 \text{ TeV}$ [114]. The 2014 limits are shown to illustrate how far the string scale limits have shifted since the previous time a dilepton ADD model analysis result was published by the ATLAS Collaboration.




Figure 6.17.: Expected and observed lower M_S limits for all KK conventions in the pure graviton cross-section scheme, in the electron (a) and muon (b) channels.

6.7. Signal Region Re-Optimisation

Although the limits on M_S have become more stringent since the corresponding 2014 ATLAS results, when compared to a recent paper put out by the CMS Collaboration, also using the full Run 2 LHC dataset, the limits presented here are less stringent [115]. The limit discrepancy between the ATLAS and CMS results are almost as large as 1 TeV for certain KK conventions. The difference is larger in the muon channel, which could be explained by the CMS muon acceptance times efficiency, which is quoted to be




Figure 6.18.: Expected and observed lower M_S limits for all KK conventions in the interference cross-section scheme, in the electron (a) and muon (b) channels.

considerably higher than that of this analysis, at around 80% for all dimuon invariant mass values. Another possible explanation is that this analysis does not optimise the SRs for an ADD signal, whereas the CMS analysis does. In the CMS analysis, the optimal search region, found to be above 1.8 TeV invariant mass, is divided into 400 GeV bins. We perform a simple study to quantify the possible gain in sensitivity from re-optimising the SRs for an ADD signal.

We want to determine if changing the lower edge of the SR bin, SR_{\min} increases the ability to discover an ADD signal. To achieve this we move SR_{\min} between 2 and 4 TeV in steps of 50 GeV. At each step, we calculate the expected number of background events in the SR by integrating the smooth background histogram template introduced in Chapter 3. The expected number of background events is used to calculate the expected upper limit on the number of signal events in the SR. Unlike the other statistical calculations performed in this chapter, this limit calculation uses a Bayesian interpretation of probability. The expected signal limit in both lepton channels is drawn as a function of SR_{\min} in Figure 6.19. Also included in this figure is the SR event yield of an $M_S = 7$ TeV ADD signal in the GRW convention. This event yield curve should be monotonically falling, but in both the electron and muon channels it is seen to increase slightly at low invariant mass. This behaviour is caused by negative bin counts in the signal histograms, caused by the DY subtraction procedure. To find the optimal SR starting point for this 7 TeV signal, we want to select the SR that puts the strongest expected limit on the number of signal events, while still accepting as many ADD events as possible. Such a SR would maximise the chances of discovering this 7 TeV ADD signal. Note that the number of integrated signal events is always smaller than the expected upper signal event limit. This means that we would not expect to be able to exclude a $M_S = 7$ TeV signal for any SR_{\min} .

Figure 6.20 shows the fraction of the expected limit to the ADD signal yield as a function of SR edge. The minimum value of the distribution of this fraction is taken as the optimal SR starting point. Note that this distribution is quite flat as a function of SR_{\min} , meaning that we expect the gain in sensitivity from moving SR_{\min} to be small. The optimal values of SR_{\min} are found to be 3 TeV in the electron channel, and 3.3 TeV in the muon channel. To estimate the expected M_S limit improvement gained when using these optimal SRs, the expected n_{sig} limit in these SRs are converted to expected limits in M_S using the limit




Figure 6.19.: Expected upper limit on n_{sig} as a function of SR_{min} , in the electron channel (a) and muon channel (b). Also included is the SR yield of a $M_S = 7$ TeV signal.




Figure 6.20.: The expected upper signal event limit, the integrated SR event yield, and the ratio between the two, for a $M_S = 7$ TeV signal sample, as a function of SR_{min} , in the electron channel (a) and muon channel (b). The minimum of this ratio curve defines the optimal choice of SR_{min} . The label on the y -axis applies only to the limit and event yield numbers, and not to the ratio.

conversion method introduced in Section 6.6.1. The expected limit on M_S becomes 180 GeV higher in the electron channel, and 120 GeV higher in the muon channel, than the expected M_S limit set in the destructive SR. Note that the conversion method uses the SR integrals of the ADD signal model histograms in the destructive SR, meaning that the conversion is only a rough estimate of what the expected limit on M_S would be in the new optimised SRs. This means that the expected limits calculated in this section are too optimistic, and hence we expect to slightly overestimate the potential gain from re-optimising the SRs.

7. Conclusions and Outlook

This thesis presents two analyses searching for increased production of dielectron and dimuon events in 139 fb^{-1} of pp collision data collected by the ATLAS detector. The analyses make use of all the data delivered to the ATLAS detector by the LHC, at a collision energy $\sqrt{s} = 13 \text{ TeV}$, during the Run 2 data-taking campaign that took place between 2015 and 2018.

Chapter 5 details a search for a resonant dilepton signature above 250 GeV. A data-driven method is used to estimate the SM background to the dilepton invariant mass distribution, by fitting a functional form to the data. Care is taken to avoid a spurious signal due to background mismodelling. A sliding window method for estimating the background is explored. A generic signal shape, of relative width between 0% and 10%, is used to model the new physics resonance hypotheses. No significant excess above the SM prediction is observed. Limits are set at 95% CL on the fiducial cross-section times branching ratio of a generic signal, as well as on the pole mass of three benchmark Z' boson models, two of which are based on the E_6 symmetry group. The analysis puts the most stringent lower mass limit, 5.1 TeV, on the SSM Z' boson to date. The Z' mass limits set in the analysis are more stringent than previous results put out by the ATLAS exotic dilepton working group by between 500 GeV and 800 GeV [10]. Figure 7.1 shows the cross-section limit of a generic dilepton resonance found in this and previous ATLAS dilepton searches. In the figure, the cross-section is expressed as a function of pole mass, in units of the SSM Z' boson cross-section.




Figure 7.1.: Limits on the generic cross-section of a dilepton resonance, as a function of resonance pole mass, in units of the SSM Z' boson cross-section, for six ATLAS dilepton analyses.

Chapter 6 presents the search for non-resonant high invariant dilepton mass excesses, consistent with the ADD model of large extra dimensions. The analysis uses a novel

method for estimating the SM background, in which a functional background form is fitted to the data in a low invariant mass control region, and extrapolated into the high invariant mass signal region. No excess is observed, in terms of events recorded in the signal region, above the SM prediction. Using the CL_s method we set lower limits at 95% CL on the string scale parameter of the ADD theory for three conventions for summing KK graviton modes, the GRW, Hewett, and HLZ conventions. In the GRW convention, the observed (expected) lower M_S limit is set to 6.6 (6.5) TeV in the electron channel, and 6.4 (6.3) TeV in the muon channel. The limits are the first ADD model limits set by the ATLAS Collaboration in dilepton final states using the Run 2 data set.

During the previous dilepton analysis undertaken by the ATLAS exotics dilepton working group, analysing 36 fb^{-1} of data [10], concerns were brought up by the analysis team, about the production of MC simulated data sets not being able to keep up with the data being delivered to the ATLAS detector by the LHC. In a physics analysis using MC samples to model the SM background, it is important for the statistical uncertainty on the MC background to be negligible compared to the statistical uncertainty in the data. For this reason, instead of purely relying on MC samples, the analyses presented in this thesis have used data-driven methods to estimate the SM background. In Run 3, and in the planned High-Luminosity Large Hadron Collider (HL-LHC), the issue of producing sufficiently large MC data sets is expected to get worse. There is reason to assume, therefore, that data-driven methods for background estimation in physics analysis will become even more commonplace in the LHC experiments in the future. This necessitates the development of efficient tools and methods for data-driven fits. However, as was the case in the searches presented in this work, MC data sets of high statistical precision are still needed in data-driven analyses.

A common problem facing data-driven analyses is the trade-off between bias and sensitivity. If the background function is too constrained, it may not model the background accurately, leading to a spurious signal. On the other hand, if the chosen parametric function contains too many free parameters, it may fit away a signal, leading to low sensitivity in the analysis. In this work, we address this issue by quantifying and including a spurious signal uncertainty in the statistical analyses. Other solutions are available, such as the *discrete profiling* method, which is favoured in the CMS Collaboration [116]. In the discrete profiling method, instead of using a spurious signal uncertainty, the uncertainty on the choice of function is considered a separate nuisance parameter in the statistical analysis.

Another issue facing data-driven physics analyses is that as the size of the data set being analysed increases, its statistical uncertainty decreases. As a result, the functional form used to fit the data must increase in complexity. When faced with a larger data set, as will be the case in ATLAS in Run 3 and beyond, it may be hard when using a global fit method, such as in our resonant analysis, to find a parametric function to accurately model the background across the whole range of the chosen discriminant variable. The sliding window fit method, described in Appendix A, may offer a solution to this issue in the future. Functional decomposition [117] and Gaussian process regression [118] are also promising methods, but are not considered in this thesis.

Appendices

A. Testing the Sliding Window Fit Method

As the size of the LHC data set, recorded by ATLAS, increases, the statistical uncertainty on said data decreases. For a data-driven search such as the one described in Chapter 5, this means that to accurately describe the shape of the SM background the chosen background fit function must increase in complexity. One might reach a point where it becomes hard to find a global fit function with acceptable performance across the full invariant search mass range. In this appendix, we explore an alternative method of data-driven background estimation, in which the dilepton invariant mass distribution can be fitted using a fit function with a small number of free parameters, by fitting the distribution more than once in shorter invariant mass subranges.

A.1. The Sliding Window Fit Method

In the global fit method, the dilepton invariant mass distribution is fitted in the whole invariant mass scan range, regardless of the pole mass hypothesis being tested. An alternative to this is to only fit a limited mass region, or "window", around each resonance pole mass. This approach is known as the sliding window fit (SWiFt) method.

First, some nomenclature. A *window* is a range in invariant mass, usually a subrange of the full global fit range. Since we are performing fits to binned data, the width of each window is some integer number of bins in the data histogram. The *central bin* of a window is the bin of the window centred at the pole mass hypothesis when searching for a signal. It is important to point out here that the central bin is not necessarily equivalent to the geometric centre of the window: a window may be asymmetric around the central bin. The low- and high-mass limits of each window are called the left edge m_L and right edge m_R of the window, respectively.

A.2. Optimising the Sliding Window Fit Method

The main obstacle in the global fit approach is finding a suitable fit function to describe the SM prediction. In the SWiFt approach, in addition to the choice of function, we must also determine a sensible shape of the sliding windows.

The SWiFt windows may be constant in width, or have more complex shapes defined as functions of pole mass. Seeing as the dilepton mass resolution changes with increasing invariant mass, it is perhaps reasonable to assume that the optimal choice of window width is proportional to the dilepton invariant mass resolution. Based on this assumption we define the window parameterisation we call "resolution windows". In the resolution window scheme, the left and right window edges are defined using the dilepton invariant mass resolution, measured at the central bin. The window size is defined using the parameter

N . The right, high invariant mass, edge of each sliding window is defined as

$$m_R = m_{Z'} + N \times \text{Res}(m_{Z'}), \quad (\text{A.1})$$

where $\text{Res}(m)$ are the electron or muon mass resolution functions described in Section 3.6. The left window edge in the resolution scheme cannot be defined analogous to the right edge. Due to the large dimuon resolution at high mass, the term $N \times \text{Res}(m_{Z'})$ quickly becomes large at high mass in the muon channel. A left edge definition like that of Equation (A.1) would quickly lead to window widths of global fit size, even for modest values of N . Instead, the resolution window left edge is defined in an iterative way. Instead of finding the edge as some multiple of the resolution at the pole mass, the edge is found by "stepping down" from the pole mass in N iterations, updating the resolution at each step. Now, the resolution window left edge is found using the relations

$$\begin{aligned} m_L^0 &= m_{Z'} - \text{Res}(m_{Z'}), \\ m_L^i &= m_L^{i-1} - \text{Res}(m_L^{i-1}), \\ m_L &= m_L^N. \end{aligned} \quad (\text{A.2})$$

The shape of the windows are determined by the resolution window parameter N alone. Due to the difference in the left- and right edge definitions of Equations (A.1) and (A.2), resolution windows are wider on the right side of the central bin than they are on the left side. An alternate method of defining the shape of the sliding windows is to add a second window shape parameter, in addition to one controlling the window width, namely one controlling the *centrality* of the windows. We call this window parameterisation scheme Width and Centrality (WC) windows. WC windows may be set perfectly symmetrical, or they can be made longer on one side than the other. It is reasonable, perhaps, to assume that high pole-mass windows perform better with a longer side towards the high statistics region at low mass, than would symmetrical windows. Given the width and centrality parameters w and c , the left- and right edges of WC windows are given as

$$\begin{aligned} m_L &= m_{Z'} - w(1 - c)/2, \\ m_R &= m_{Z'} + w(1 + c)/2. \end{aligned} \quad (\text{A.3})$$

The centrality of a window controls the distance from the geometric centre of the window to the central bin of the window. A window of centrality $c = 0$ is perfectly symmetrical. In a window of centrality -1 or 1 , the central bin falls on the right or left edge of the window, respectively. In general, w and c may be functions of pole mass.

If the number of events in data within each sliding window becomes too low, background and signal fit stability may begin to suffer. There is a risk of this occurring, especially in the low statistics, high invariant mass, region if the width of the sliding windows become too narrow. Ensuring that windows are always wide enough to contain enough events can be achieved by defining sufficiently wide windows, or by employing a so-called "anchor point". An anchor point is an upper limit on the value of the left window edge. When employing the anchor point, if the window left edge, as calculated by Equations (A.2) or (A.3), is greater than the anchor point value, then the left edge is instead set to the value of the anchor point. An example of anchor point usage can be seen in Figure A.1. Here, an anchor point is set to 1 TeV for all windows.

A.2.1. Function Choice Optimisation for Sliding Window Fits

We want to determine a good choice of SWiFt fit function to describe the SM background in the resonant dilepton analysis. Candidate fit functions are evaluated using a spurious signal test. The ensemble of one- and two parameter¹ fit functions considered in the analysis are listed in Tables A.1 and A.2.

Table A.1.: One-parameter functions considered in the SWiFt function optimisation study. The function names are chosen to describe their mathematical form. The free parameter is a , and the dimensionless parameter x is defined as $x = m/\sqrt{s}$.

<i>Function Name</i>	<i>Expression</i>
EXP	e^{ax}
POW	x^a
LOG	$ \log(x) ^a$
POWPOLY1	x^{ax}
POWPOLYLOG1	$x^a \log(x)$

Resolution windows of $N = 10$ are used in the SWiFt function optimisation procedure, with an added anchor point set at 1 TeV, limiting the upper value of the window left edges. The signal strength parameter μ is restricted to positive values for pole masses above 1.5 TeV. In this study, windows are not allowed to extend outside of the scan range. A plot of the chosen resolution windows as a function of pole mass is shown in Figure A.1. The window shape parameterisation of the SWiFt method is further optimised in Section A.2.2.




Figure A.1.: Electron (a) and muon (b) channel window configurations used for the function choice optimisation procedure, as a function of pole mass. Each coloured line in the plot is a separate window. Black markers indicate the position of the central bins in each window. The effect on window size of the anchor point at 1 TeV is clearly seen, most notably in the electron channel.

¹One extra free parameter must be counted if one also considers the background normalisation parameter N_{bgr} .

Table A.2.: Two-parameter functions considered in SWiFt optimisation. Some function names are chosen based on their mathematical form, while other have names chosen for historical reasons. The free parameters in the functions are a and b , while the dimensionless parameter x is again defined as $x = m/\sqrt{s}$.

<i>Function Name</i>	<i>Expression</i>
EXPow	$e^{ax} x^b$
EXPOLYLOG2	$e^{a \log(x)^2 + b \log(x)}$
DIJET	$(1 - x)^a x^b$
CUTOFFEXP	$(1 - x)^a e^{bx}$
CUTOFFLOG	$(1 - x)^a \log(x) ^b$
CUTOFFEXPSQ	$(1 - x)^a e^{bx^2}$
CUTOFFPOWPOLY1	$(1 - x)^a x^{bx}$
CUTOFFPOWPOLYLOG1	$(1 - x)^a x^{b \log(x)}$
MULTIJET5	$(1 - x)^a (1 + x)^{bx}$
MULTIJET6	$(1 - x)^a (1 + x)^{b \log(x)}$
MULTIJET7	$(1 - x)^{a - b \log(x)} x^{-1}$
MULTIJET8	$(1 - x)^{a - b \log(x)} x^{-2}$
MULTIJET9	$(1 - x^{1/3})^a x^{bx}$
MULTIJET10	$(1 - x^{1/3})^a x^{b \log(x)}$
POWSPF	$(1 - \log(ex^a))^b$
EXPOLY2	$e^{ax^2 + bx}$
POWPOLYLOG2	$x^a \log(x)^2 + b \log(x)$
DIBoSON	$e^{ax} + e^{bx}$
LAURENT	$ax^{-4} + bx^{-5}$

Each candidate background function is used in a spurious signal test on the signal-free background MC template. The spurious signal significance is extracted at pole mass hypotheses spaced linearly at 5 GeV intervals between 250 GeV and 6 TeV. The maximum and mean value of the absolute spurious signal significance is shown for each candidate function in Table A.3. We want to choose a function which minimises these quantities. The mean- and maximum expected significance of each function is also tested. No function is seen to obtain low values of spurious signal at the expense of the power of the function. How we calculate expected significance is detailed in the next section.

The three best performing candidate functions are EXPOLYLOG2, EXPOW, and MULTIJET9. The spurious signal significance as a function of pole mass for these three functions is shown in Figure A.2. The three chosen function candidates are used in the window parameterisation optimisation study detailed in the next section.




Figure A.2.: Spurious signal significance in the electron channel (a) and muon channel (b), for the three best performing fit functions of the function optimisation study.

A.2.2. Window Shape Optimisation for Sliding Window Fits

The ideal choice of window shape may not be the same at every pole mass hypothesis. The decreasing number of events in the data in the high pole-mass region suggests that windows should widen as a function of pole mass. In this section, we will optimise the window shape parameters for discovering a new physics signal.

The figure of merit in the window optimisation test is the expected significance of a zero-width signal. The expected significance is calculated for a zero-width signal injected on an Asimov data set, where all observables are set to be identical to their expected values. The number of injected signal events is proportional to the square of the number of events in the background Asimov data in a resolution window of size $N = 4$ around the central bin of the window. The absolute size of the injected signal is not important, as long as it stays constant at each pole mass for different choices of window parameterisation. Similarly, the absolute value of the expected significance is not important. We are only interested in the relative significance between different candidate window configurations.

A spurious signal systematic uncertainty is added to the signal and background fits to the generated Asimov data set. If no such uncertainty is added, the "preferred" window size will always be as wide as possible, due to the additional constraining power of wider

Table A.3.: Function choice optimisation spurious signal test results. Smallest value in each column marked in red. We see that the EXPOLYLOG2 has the smallest maximum and mean SS/σ_{SS} in both the electron and muon channels. The EXPow and MULTIJET9 functions are also seen to perform well in both channels. Functions marked — are dropped due to poor performance or fit stability issues.

<i>Function Name</i>	<i>Electron Channel</i>		<i>Muon Channel</i>	
	<i>Max. SS/σ_{SS} [%]</i>	<i>Mean SS/σ_{SS} [%]</i>	<i>Max. SS/σ_{SS} [%]</i>	<i>Mean SS/σ_{SS} [%]</i>
EXP	830.0	26.2	855.8	168.6
POW	62.0	2.2	242.2	15.7
LOG	29.0	3.7	110.4	46.6
POWPOLY1	3,050.0	51.4	8,067.0	750.7
POWPOLYLOG1	—	—	—	—
EXPow	32.2	1.6	50.7	3.4
EXPOLYLOG2	25.8	1.6	34.3	3.0
DIJET	953.2	4.2	51.9	5.6
CUTOFFEXP	628.8	26.1	851.2	167.1
CUTOFFLOG	556.2	5.4	112.3	46.7
CUTOFFEXPSQ	—	—	—	—
CUTOFFPOWPOLY1	203.5	33.8	884.5	187.0
CUTOFFPOWPOLYLOG1	116.2	2.4	70.6	17.5
MULTIJET5	174.9	5.9	812.3	82.8
MULTIJET6	362.5	3.2	359.4	29.2
MULTIJET7	893.4	4.8	284.1	27.8
MULTIJET8	115.3	1.7	44.4	7.7
MULTIJET9	26.4	1.7	46.2	3.6
MULTIJET10	330,245.4	289.1	68.6	8.4
POWSPF	27,262.8	732.0	73,440.5	1,334.8
EXPOLY2	961.3	7.3	725.9	77.9
POWPOLYLOG2	—	—	—	—
DIBOSON	704.5	23.5	748.2	153.9
LAURENT	—	—	—	—

sideband regions. The spurious signal yield grows with increasing window width. At some point, the sensitivity gained by increasing the window width is offset by the sensitivity loss due to spurious signal. The aim in the window optimisation study is to find this point.

The test is performed using WC windows. The window width is varied in steps of 50 GeV, between 100 GeV and 2 TeV. The window centrality is varied in steps of 0.1 between -0.8 and 0.8. For each pole mass, in steps of 3 GeV, between 250 GeV and 1 TeV, each choice of window configuration is tested using the following three steps:

1. Perform a background plus signal fit on the smooth MC background template. Extract the size of the spurious signal constraint parameter.
2. Construct a smooth Asimov data set from the background PDF and inject a signal.
3. Calculate the significance of the signal injected in Step (2), with the added spurious signal constraint measured in Step (1).

The MC background template, described in Chapter 3, is smooth, but not entirely void of statistical fluctuations. Fluctuations in the template may lead to over- or underestimating the true value of the spurious signal for a chosen background model. We assume that the effects of such fluctuations will impact the different choices of window parameterisation more or less equally. Overall, over- and underestimation of the spurious signal at individual pole mass hypotheses should average out. This procedure bypasses the procedure outlined in Section 5.4, where the spurious signal is estimated using the envelope of local spurious signal maxima. This simplification of the treatment of the spurious signal uncertainty has room for improvement.

The expected significance distribution at each mass point is smoothed in the plane of window width and window centrality. This is done to reduce statistical fluctuations, allowing the broader trends we are interested in to become more visible. The smoothing is motivated by the assumption that neighbouring window configurations should have similar performance. The maximum expected significance at each pole mass, the point in the plane of width and centrality where the expected significance is at its greatest, is used to discriminate between the three candidate fit functions. The maximum expected significance, as a function of pole mass, for the three candidate functions, is shown in Figure A.3. The three fit functions perform similarly, but the ExPOW function has the highest mean maximum expected significance. This function is therefore kept as the final choice of background function for the SWiFt method.

The expected significance is drawn, as a function of window width and centrality for three example pole mass values, 250 GeV, 500 GeV, and 1 TeV, in Figure A.4. The points of greatest expected significance in the plots are marked with a red star.

In Figures A.5 and A.6 we show the maximum expected significance as a function of window centrality, optimised with respect to window width, and maximum expected significance as a function of window width, optimised with respect to window centrality, for the same three pole mass values shown in Figure A.4, 250 GeV, 500 GeV, and 1 TeV. Looking at these figures, we notice that the window parameterisation can be split into three regions of pole mass: a low pole mass region, a high pole-mass region, and a transition region between them. In the low pole mass region, there is a clear preference for narrow windows (see Figures A.5 (b) and A.6 (b)). In this region there is also a clear preference for positive centrality windows (see Figures A.5 (a) and A.6 (a)), meaning that windows skewed to the right, towards higher mass are preferred. With increasing pole mass, the maximum expected significance windows become wider, but there is no longer




Figure A.3.: Maximum expected significance, drawn as a function of pole mass, for the fit functions EXPow, EXPOLYLOG2, and MULTIJET9. The EXPow function has the highest mean maximum expected significance, and is therefore kept.

a sharp expected significance peak at some value of window width. Instead, the expected significance is seen to plateau. This means that there is more leeway when setting the window parameterisation in the high pole-mass region than there is in the low pole mass region. For a window of central bin 1 TeV (see Figures A.5 (f) and A.6 (f)), a window of width 1 TeV is expected to perform about as well as a window of width 2 TeV, given the right definition of window centrality.

Low pole mass windows have a tendency towards asymmetric windows. At higher pole mass values the window centrality distribution becomes flatter, and more centred around zero (see Figure A.5 (c) and (e); and Figure A.6 (c) and (e)). The window width and centrality values corresponding to maximum significance are plotted as functions of pole mass in Figure A.7.

Based on the observations above, we define the low pole mass region to run from 250 GeV to 300 GeV, and the high pole-mass region to be above 900 GeV. The so-called intermediate region runs between 300 GeV and 900 GeV. In the low pole mass region, the window width parameterisation is determined by fitting a constant value to the maximum expected significance value width distribution, shown in Figure A.7, in this region. In the intermediate region, the window width parameterisation is determined by a straight line fit to the maximum expected significance value width distribution in the intermediate pole mass region. The low edge of this fit is constrained by the value of the low pole mass region constant fit. The only free parameter in the fit is therefore the slope of the line. The high pole-mass region window width parameterisation is also a constant function. This value is not determined by a fit, but is rather set to the value of the intermediate fit value at the transition between the intermediate and high pole-mass region. The window centrality window parameterisations are determined similarly. The window centrality parameterisation in the low pole mass region is determined with a linear fit to the maximum expected significance value centrality distribution, in this region. For the centrality, the intermediate and high pole-mass regions are combined. The window centrality in this region is determined by fitting a constant value to the maximum expected significance value centrality distribution above 300 GeV. The window width and centrality







Figure A.4.: Expected significance as a function of sliding window width and centrality in the electron channel (left) and muon channel (right), for signal pole masses 250 GeV (a) (b), 500 GeV (c) (d), and 1 TeV (e) (f). The crescent shape in the lower right of the plots are due to the limited range of the background MC template; Certain window widths are, depending on their window centrality, too wide to fit inside the template range at some pole mass points.




(a)




(b)




(c)



(d)




(e)




(f)


Figure A.5.: Maximum expected significance, as a function of sliding window centrality, optimised with respect to window width (left), and maximum expected significance as a function of sliding window width, optimised with respect to window centrality (right), for signal hypothesis pole masses 250 GeV (a) (b), 500 GeV (c) (d), and 1 TeV (e) (f), in the electron channel.




(a)




(b)




(c)



(d)



(e)



(f)

Figure A.6.: Maximum expected significance, as a function of sliding window centrality, optimised with respect to window width (left), and maximum expected significance as a function of sliding window width, optimised with respect to window centrality (right), for signal hypothesis pole masses 250 GeV (a) (b), 500 GeV (c) (d), and 1 TeV (e) (f), in the muon channel.

parameterisation functions are drawn in the maximum expected significance plots in Figure A.7.




Figure A.7.: Maximum expected significance sliding window centrality (a) (b) and width (c) (d) in the electron channel (left) and muon channel (right). The final sliding window shape parameterisation function are also drawn in the figures.

We see in Figure A.7 that the chosen window parameterisation functions do not always perfectly fit the maximum expected significance distributions. The performance of the chosen window parameterisation scheme is tested by comparing the expected significance at the chosen window parameterisation to the maximum expected significance at each mass point. The ratio of expected significance at the chosen window to the maximum expected significance, as a function of pole mass, is shown in Figure A.8. Except for a region of inefficiency around 300 GeV pole mass in the muon channel, and a single spike in the electron channel around 600 GeV, the chosen window parameterisation functions show 90% efficiency or better.

In addition to the window parameterisation described above, an anchor point is set at 1 TeV. This is done to ensure adequate statistics for windows at very high pole-mass. The shape of the final window parameterisation in both the electron and muon channels are shown in Figure A.9. Here, the chosen window parameterisations are optimised for zero-width signals. For wide signal hypotheses, a new parameterisation of window size may be needed.




Figure A.8.: Ratio of expected significance at the chosen SWiFt window parameterisation to the maximum expected significance in the electron (a) and muon (b) channels, as a function of signal hypothesis pole mass.




Figure A.9.: The final SWiFt window parameterisation in the electron channel (left) and muon channel (right), as a function of signal hypothesis pole mass. Each line in the figure shows the range of the window. The black markers show where the central bin of each window falls.

A.3. Comparing The Global Fit and Sliding Window Fit Methods

Here, we compare the global fit method used in Chapter 5 to the optimised SWiFt method from Section A.2.2. In Section 5.4, the spurious signal is parameterised as a function of pole mass by fitting a parametric function shape to the maxima of the signal yields extracted in a spurious signal test. For the global fit and SWiFt comparison study performed in the next section, a simpler approach is taken to obtain a conservative spurious signal estimate. Here, we draw linear interpolations between neighbouring points of spurious signal maxima, instead of fitting a parameterised function to these points. The spurious signal was extracted for a zero-width signal at 500 logarithmically spaced pole mass points between 250 GeV and 6 TeV. The spurious signal yields, with their linear interpolation envelopes, for the SWiFt and global fit methods can be seen in Figures A.10 and A.11, respectively. The spurious signal significances of the SWiFt and global methods can be seen in Figure A.12. We see that although the two methods for background estimation have similar spurious signal yields at low pole masses, the SWiFt method seems to be more susceptible to spurious signal at higher pole masses. It is worth pointing out here that the chosen sliding window optimisation procedure based on expected significance naturally leads to a non-negligible spurious signal.




Figure A.10.: Spurious signal yields extracted in a zero-width spurious signal test using the SWiFt method in the electron (a) and muon (b) channels. Also shown are the linear interpolations between neighbouring points of maximum spurious signal yield, used to give a conservative estimate of the spurious signal systematic uncertainty.

The global fit and SWiFt methods for background estimation are compared by calculating the significance of a zero-width signal, tested at one thousand logarithmically spaced pole mass points between 250 GeV and 6 TeV. The tests are performed on the full 139 fb^{-1} Run 2 data set. The local significances calculated with the two methods are shown together as a function of pole mass in Figure A.13. These significance calculations only take into account statistical uncertainty on the background and signal fits. Some noise spikes can be seen in the significance curve calculated using the SWiFt method, where the sliding window fits fail or give wrong results. Next, the linearly interpolated spurious signal estimates from the previous section are included as signal normalisation nuisance parameters in the fits. The resulting local significance values are shown in Figure A.14.




Figure A.11.: Spurious signal yields extracted in a zero-width spurious signal test using the global fit method in the electron (a) and muon (b) channels. Also shown are the linear interpolations between neighbouring points of maximum spurious signal yield, used to give a conservative estimate of the spurious signal systematic uncertainty.




Figure A.12.: Spurious signal significance in a zero-width spurious signal test using the SWiFt and global fit methods, in the electron (a) and muon (b) channels.




Figure A.13.: Comparing significance as a function of pole mass, calculated, only taking into account statistical uncertainty, using the SWiFt and global fit methods, in the electron (a) and muon (b) channels.




Figure A.14.: Comparing significance as a function of pole mass, calculated, with spurious signal, using the SWiFt and global fit methods, in the electron (a) and muon (b) channels.

The global fit and SWiFt methods show similar results in the electron channel, with or without including the constraint of spurious signal uncertainty. In the muon channel, the SWiFt method yields a higher significance than the global fit method in the very low pole mass range, while the opposite is the case at intermediate pole masses. The discrepancy in the low pole mass region can perhaps be explained by the difference in fit range between the two methods; the global fit window always starts at 250 GeV, at the search range lower limit, while the SWiFt low mass window edges can extend below this (see Figure A.9). Some muon channel windows start at around 190 GeV. A study is performed to investigate the effect of this difference in low mass fit window edge between the competing fit methods. In this test, the low pole mass edge of the SWiFt method is forced to be 250 GeV in a significance scan of the low pole mass region, only taking into account statistical uncertainty. The upper edges of the sliding windows are not altered from the optimal choices found in Section A.2. The resulting significance values are plotted in Figure A.15. The significance values obtained using the SWiFt method are closer to those obtained using the global fit method when the low mass fit window edges are set to 250 GeV. However, some discrepancy remains between the two methods. The muon channel discrepancy in the intermediate pole mass range cannot be explained by the difference in spurious signal significance between the global fit and SWiFt methods, as the discrepancy is present with and without the inclusion of the spurious signal nuisance parameter.




Figure A.15.: Low pole mass significance values in the muon channel, taking into account only statistical uncertainty on the background, for three background model methods: the global fit method (green), the SWiFt method with windows optimised per Section A.2 (black), and the SWiFt method with low pole mass window edges set to 250 GeV (red). The upper edges of the sliding windows are the same in both cases.

A.4. Conclusion

The SWiFt method is a promising method for doing data-driven background estimation on large data sets. However, in the comparison carried out here, it is outperformed by the global fit method. Due to the involved optimisation procedure of the SWiFt method, it would need to outperform the global fit method to be practical. The global fit method

is therefore chosen for the analysis performed in Chapter 5. There is perhaps room for improvement in the window shape and function choice optimisation procedures here. It is also possible that choosing a three-parameter fit function could perform better than the one and two-parameter functions considered here.

B. Auxiliary ADD Model Material

In the ATLAS Collaboration, a dedicated team is in charge of producing MC simulation samples. Whenever an analysis team needs a new MC sample, simulating a SM background- or new physics process, they put in a request for the samples to be produced to the ATLAS central production group. Having a dedicated MC production team ensures consistent and correct use of simulation tools within the Collaboration, minimises the risks of biased or erroneous samples, and avoids redundant production between analysis teams.

Analysers requesting a new set of MC samples must first produce a set of so-called validation plots. These are distributions of key kinematic variables of a test sample, often containing fewer events than the requested final sample size, visualised at the event generator level. These validation plots are checked to verify the production setup. In our case, the plots must pass scrutiny by the ATLAS exotics group before being handed over to ATLAS central production. For the non-resonant limit reinterpretation analysis performed in Chapter 6, I was responsible for the ADD model signal MC request. Presented here are the validation plots I prepared for the ADD signal sample request, and technical details on how the samples are generated. An introduction to the ADD model of large extra dimensions is given in Section 1.2.2 of Chapter 1.

B.1. Generating ADD Model Signal Samples

Seven ADD model samples are generated for different hypotheses of string scale M_S : 3 TeV, 4 TeV, 5 TeV, 6 TeV, 7 TeV, 8 TeV, and 50 GeV. The choice of string scales is motivated by the string scale limit set, between 3.2 TeV and 5.0 TeV, in the previous ATLAS search for large extra dimensions [114]. The hard scattering of the ADD model events is generated, at NLO, using the SHERPA event generator [83]. The cross-sections of the ADD process are calculated in the GRW convention. The samples are produced in two separate invariant mass intervals, one ranging from 500 GeV to 2 TeV, and one ranging from 2 TeV to 6 TeV. This is done to enhance statistical precision in the high invariant mass region. The CT10NLO PDF distributions are used [75].

The electron channel generator level validation samples each contain ten thousand events in the low mass interval, and between one thousand and five thousand events in the high-mass interval, depending on the string scale hypothesis. The muon channel validation samples each contain five thousand events in the low mass interval and one thousand events in the high-mass interval. Figure B.1 shows the invariant mass distributions of the validation samples. Similarly, Figures B.2 and B.3 show the distributions of the transverse momenta, azimuthal angles, and pseudorapidities of both final state leptons in the electron and muon channel, respectively. The invariant mass distribution of the ADD samples follows a DY exponentially falling distribution at low invariant mass. Then, at some point determined by M_S , the KK production enhances the dilepton event production. A similar effect is seen in the lepton transverse momentum distributions. In Figures B.2 (e) and (f), and B.3 (e) and (f), the presence of an ADD signal can be seen to enhance lepton production at central values of pseudorapidity.




Figure B.1.: Generator level invariant mass of the ADD signal samples in the electron (a) and muon (b) channel. The lower panel shows the ratio of each distribution to that of the $M_S = 50$ TeV sample.

What we refer to in this text as "ADD" MC signal samples, are in reality "SM + ADD" MC signal samples; in the SHERPA event generator, the matrix element of the ADD model process contains a SM spin-1 DY component. This DY contribution to the production cross-section cannot be disentangled at the generator level. In the non-resonant analysis, the SM DY contribution is already included in the data-driven background estimation procedure. Therefore, to avoid double-counting in the analysis, the DY component of the ADD samples is removed. By setting the string scale parameter in the SHERPA data card sufficiently high during event generation, far beyond the energy reach of the LHC, the resulting data sample will for all practical purposes be made up of pure DY events. Such a sample can then be used to remove the DY component from the other signal samples before analysis. The string scale of the "pure DY" sample is set to 50 TeV. To ensure that any ADD component is sufficiently suppressed at this string scale, two test samples are constructed at generator level: one with $M_S = 50$ TeV, and one with $M_S = 100$ TeV. Because any residual ADD signal is expected to reside at high-mass, both test samples are only generated in the high invariant mass interval, between 2 TeV and 6 TeV. The two test samples are generated using the same random number generator seed. This means that any noticeable difference between the two samples will be due to differences in their underlying physical properties and not due to statistical fluctuations. The invariant mass distributions of the two signal-free test samples are shown, along with their ratio, in Figure B.4. Aside from some minor differences assumed to stem from floating-point rounding, the two distributions are identical. We conclude that the 50 TeV string scale sample only contains DY events. This test is only performed in the electron channel. At the generator level, before detector simulation, and at the invariant masses being generated in the ADD model signal samples, the kinematic differences between the electron and muon channels are negligible.

Particle interaction with the ATLAS detector material is modelled for the ADD samples using AtIfast-II. In the ATLAS Collaboration, MC production occurs in so-called MC campaigns. In each of these campaigns, MC production and simulation is set up to mimic the real running conditions of a specific ATLAS data-taking period. The LHC Run 2




Figure B.2.: Generator level ADD model kinematic distributions for the leading and sub-leading electrons in the electron channel. Figures (a) and (b) show electron transverse momentum p_T , Figures (c) and (d) show electron azimuthal angle ϕ , and Figures (e) and (f) show electron pseudorapidity η . The left plots, (a), (c), and (e), show the kinematic distributions of the leading electron, while the right plots, (b), (d), and (f), show the kinematic distributions of the sub-leading electrons. The lower panels show the ratio of each distribution to that of the $M_S = 50$ TeV sample. Note that in (e), the leading electron has been mislabelled as the sub-leading electron.




Figure B.3.: Generator level ADD model kinematic distributions for the leading and sub-leading muons in the muon channel. Figures (a) and (b) show muon transverse momentum p_T , Figures (c) and (d) show muon azimuthal angle ϕ , and Figures (e) and (f) show muon pseudorapidity η . The left plots, (a), (c), and (e), show the kinematic distributions of the leading muon, while the right plots, (b), (d), and (f), show the kinematic distributions of the sub-leading muons. The lower panels show the ratio of each distribution to that of the $M_S = 50$ TeV sample. Note that in (e), the leading muon has been mislabelled as the sub-leading muon.




Figure B.4.: Generator level dilepton invariant mass of the $M_S = 50$ TeV test sample and the $M_S = 100$ TeV test sample, in the high invariant mass interval.

data-taking period is divided into three MC production campaigns: *mc16a*, *mc16d*, and *mc16e*, corresponding to the ATLAS data-taking conditions of the years 2015 and 2016, 2017, and 2018, respectively. Eighty thousand events are generated in each ADD signal sample, forty thousand each in the low and high-mass interval. The samples are roughly divided between the three MC production campaigns according to the luminosity of their corresponding data-taking periods: ten thousand events in *mc16a*, ten thousand events in *mc16d*, and twenty thousand events in *mc16e*. A list of all the ADD signal samples, with their cross-sections in the GRW convention and effective luminosities, is given in Table B.1.

Table B.1.: Table summarising the ADD model signal samples produced for the non-resonant analysis of Chapter 6. The cross-sections listed are calculated in the GRW convention for summing over KK states. The effective luminosity is the integrated luminosity required, at $\sqrt{s} = 13$ TeV, to produce forty thousand ADD model dilepton events, which is the number of events produced at each row in the table.

Channel	String Scale [TeV]	Invariant Mass [TeV]	cross-section (GRW) [fb]	Effective Luminosity [fb^{-1}]
ee	3	$0.5 < m_{ee} < 2$	1.0964E+02	3.6483E+02
	3	$2 < m_{ee} < 6$	1.2046E+01	3.3206E+03
	4	$0.5 < m_{ee} < 2$	9.7574E+01	4.0995E+02
	4	$2 < m_{ee} < 6$	2.8617E+00	1.3978E+04
	5	$0.5 < m_{ee} < 2$	9.7181E+01	4.1160E+02
	5	$2 < m_{ee} < 6$	7.9658E-01	5.0215E+04
	6	$0.5 < m_{ee} < 2$	9.5537E+01	4.1869E+02
	6	$2 < m_{ee} < 6$	3.3016E-01	1.2115E+05
	7	$0.5 < m_{ee} < 2$	9.5866E+01	4.1725E+02
	7	$2 < m_{ee} < 6$	2.1308E-01	1.8772E+05
	8	$0.5 < m_{ee} < 2$	9.5794E+01	4.1756E+02
	8	$2 < m_{ee} < 6$	1.8166E-01	2.2019E+05
	50	$0.5 < m_{ee} < 2$	9.5109E+01	4.2057E+02
	50	$2 < m_{ee} < 6$	1.6531E-01	2.4197E+05
	$\mu\mu$	3	$0.5 < m_{\mu\mu} < 2$	1.0846E+02
3		$2 < m_{\mu\mu} < 6$	1.2043E+01	3.3214E+03
4		$0.5 < m_{\mu\mu} < 2$	9.7306E+01	4.1107E+02
4		$2 < m_{\mu\mu} < 6$	2.8744E+00	1.3916E+04
5		$0.5 < m_{\mu\mu} < 2$	9.5708E+01	4.1794E+02
5		$2 < m_{\mu\mu} < 6$	7.9500E-01	5.0314E+04
6		$0.5 < m_{\mu\mu} < 2$	9.6120E+01	4.1615E+02
6		$2 < m_{\mu\mu} < 6$	3.3112E-01	1.2080E+05
7		$0.5 < m_{\mu\mu} < 2$	9.5711E+01	4.1792E+02
7		$2 < m_{\mu\mu} < 6$	2.1368E-01	1.8720E+05
8		$0.5 < m_{\mu\mu} < 2$	9.6311E+01	4.1532E+02
8		$2 < m_{\mu\mu} < 6$	1.8064E-01	2.2143E+05
50		$0.5 < m_{\mu\mu} < 2$	9.5787E+01	4.1759E+02
50		$2 < m_{\mu\mu} < 6$	1.6514E-01	2.4222E+05

B.2. Kaluza-Klein Graviton Angular Distribution

The hypothetical force-mediating boson of gravity, the graviton, is assumed to be a spin-2 particle. Most of the SM background events to the dilepton analyses presented in this work stem from spin-1 particles. There are measurable angular differences between spin-0, spin-1, and spin-2 dilepton decay events. This can be seen in Figure B.5. In the figure, the cosine of the polar angle of each fermion in difermion decays of spin-0, spin-1, and spin-2 bosons, is measured in the COM frame of the difermion system. This variable is also known as $\cos\theta^*$.




Figure B.5.: Polar angle of difermion decays from bosons, in the COM frame of the fermions. The black, orange, and yellow lines show the angular distributions of spin-0, spin-1, and spin-2 bosons, respectively [119].

It should be possible then, in principle, to use the $\cos\theta^*$ variable as a discriminant in a graviton search for spin-2 processes at the LHC. In Figure B.6, the 3 TeV and 50 TeV string scale test samples are compared in both the high and low invariant mass intervals in the electron channel. Figure B.7 shows the corresponding distributions in the muon channel.¹ The 50 TeV sample consists purely of spin-1 DY events. In the 3 TeV samples, on the other hand, the spin-2 process contributes significantly to their cross-sections, at least in the high invariant mass interval. Here, we see the angular distribution predicted from Figure B.5 repeated in the dilepton $\cos\theta^*$ -distribution, albeit with some statistical fluctuations.

We can distinguish spin-2 boson from spin-1 bosons in dilepton decays using angular information at the event generator level. However, an angular analysis may not be so straightforward in reconstructed data. Detector effects will significantly affect the resolution of any kinematic discriminant. See for example the dilepton invariant mass resolution in Figure 3.6. Moreover, finding the boost to the COM frame of the reconstructed leptons in the event is non-trivial in a hadron collider such as the LHC, due to the unknown initial momenta of the colliding partons. The non-resonant search for ADD model graviton excitations carried out in Chapter 6 does not use angular information in the analysis. Future

¹An error when scaling the samples to cross-section means that the distributions differ between the electron and muon channels. Emphasis should be placed on the kinematic shape of Figures B.6 and B.7, and not on their scale.




Figure B.6.: Comparing the electron channel $\cos\theta^*$ distribution between the $M_S = 3$ TeV test sample and the $M_S = 50$ TeV test sample, in the low (a) and high (b) invariant mass intervals.




Figure B.7.: Comparing the muon channel $\cos\theta^*$ distribution between the $M_S = 3$ TeV test sample and the $M_S = 50$ TeV test sample, in the low (a) and high (b) invariant mass intervals.

dilepton searches could perhaps benefit from including an angular discriminant along with invariant mass in their searches.

C. ATLAS Authorship Qualification Report

C.1. Introduction

To become an author in the ATLAS Collaboration one has to fulfil certain criteria. One of these criteria is to spend the equivalent of at least 80 working days on a technical task for the experiment. This is a report of the technical task I performed to qualify as an ATLAS author. The project started on 10 October 2016 and ended on 10 October 2017.

I wanted to contribute to the computing operations of ATLAS and was given a project in the ATLAS Distributed Computing (ADC) analytics group. The project description as it was recorded into the internal ATLAS authorship database was:

Automating distributed computing operations with machine learning algorithms

The project delivers a method to automate workflows usually done by human operators on shift. The goal is to develop a system to classify operational system metrics into steady-state and anomalous behaviour, across a wide range of infrastructures and services. This includes data management, workload management, databases and networks. With the ATLAS Open Analytics Platform, we have a data store that collects system metrics, which can be used for correlation and classification. The first part involves manual and algorithm-assisted inspection of the historic data to obtain insights into the anomalies. The second part provides a framework with an integrated feedback loop that learns the decisions of the shifter, yielding a better classification of events and proposals for notification or automated resolution. A final milestone would be to automate the recommendations and decisions where statistically a significant decision can be made.

Large computing infrastructure is required to fulfil the needs of the LHC experiments. It would be difficult, both financially and in terms of infrastructure, to keep all these resources in one place. The Worldwide LHC Computing Grid (Grid) is a global network of computer centres, called *sites*. Computing tasks, so-called *jobs*, such as the reconstruction of physics events or detector modelling for Monte Carlo simulations, are created by CERN users and submitted to the Grid. These jobs are then distributed to computing sites where computing resources are available at any given moment. Computing jobs may also be distributed according to where the data to be processed are located.

The aim of my authorship qualification project was to develop a system for automating anomaly detection on ATLAS distributed computing operations. The Grid is monitored around the clock by humans, so-called *shifters*. Shifters observe the status of the Grid and take action when anomalies appear. It could free up a not insignificant amount of person-power if some or all of the tasks currently performed by shifters could be automated. The

first phase of the project would be to make some algorithm for detecting anomalous states on the Grid. The final goal of the project would then be to make some kind of automated system taking over shifter responsibilities.

C.2. File Transfer Anomalies

As a starting point, we chose to focus on file transfers. The ATLAS Collaboration has a large need for digital storage space, to store measurements from the detector as well as simulated particle collisions. Digital files are stored at CERN or at one of the other computing centres around the world that make up the Grid. This method of scattered data storage is known as distributed data management (DDM). The files stored on the Grid must often be transferred between storage locations when they are needed by the ATLAS users. Rucio [120] is the data management system developed and employed by the ATLAS Collaboration. It tracks the storage location of all ATLAS data sets and files and provides an interface for moving files between grid sites and ATLAS users.

Every so often, file transfers between storage locations fail. There are a lot of reasons a transfer may fail. The error may be on the user end, such as attempting to move a non-existent file, or it may be due to technical errors on the Grid, such as grid sites being offline. Shifters are monitoring the Grid at all times. If the rate of failed transfers to or from a Grid site becomes too large over time, the shifter will take action to see to it that the problem is fixed.

Metadata about each transfer made on the Grid is stored in a database. This data consists of information such as the source and destination of the transfer, when it was made, and whether or not it was successful. Another database contains information about the state of the Grid at any given moment. This includes information about how many files have been transferred in the last hour and the throughput and latency for transfers between any two Grid sites. These databases are updated in real-time. Our goal was to develop some automated system that, based on these transfer metadata and Grid metrics, could recognise an anomalous state of the Grid or predict file transfer failures.

C.3. Correlating Anomalies and Shifter Tickets

The instructions given to shifters monitoring the DDM system state that shifters should look for problematic clouds or sites, meaning locations where a large fraction of the transfers made in or out of the system is failing. Shifters are then instructed to monitor these problematic sites. If the problems persist over time, the shifter is to file a ticket to the Global Grid User Support (GGUS) ticketing system. GGUS is the primary channel for users to request support while using the Grid. Tickets entered into GGUS are stored, and can be viewed and searched for by keyword, even after being resolved. The instructions given to shifters are vague, not stating a specific rate of file transfer failure or period of inactivity at a site before action is required. This allows shifters to exercise their expertise and judgement on a case by case basis, whether or not to file a support ticket.

It sounds feasible to replace the DDM shifters with an automated system. A monitoring system could be put in place that autonomously files a support ticket if the rate of file transfer failures to or from a Grid location stays above some predetermined threshold for some predetermined period of time. It does seem unlikely that one threshold or cutoff time would suffice for all Grid locations. Some sites may experience a higher failure rate

than others, without this being symptomatic of an anomalous state. The thresholds may also vary with time. Certain times of the day or days of the week may be more prone to errors occurring, for example, due to increased network traffic. A more sophisticated system should also take into account such factors.

We had the idea to make a monitoring system that would learn how to correctly identify an anomalous state in the DDM system by observing shifter decisions. The system would need to learn where to put the thresholds for acceptable file transfer failure rates and inefficiency periods before filing an automatic support ticket. This would be done by studying old tickets filed to GGUS in the past, and examining the state of the system at the Grid sites being mentioned in the ticket, at the time of, and the period leading up to, submission of the ticket.

As a starting point, to test the feasibility of such a system, we carried out a study to see if there really was a correlation between extended periods with a high rate of failure for file transfers between Grid locations, and the timestamps of historical DDM support tickets. Such a correlation is expected given how the shifter instructions are formulated. If the fraction of file transfer failures between two grid sites is elevated for an extended period of time, shifters are directed to file a ticket. For the kind of system outlined here to be possible, such a correlation would need to be observed. We arbitrarily chose a week of file transfer metadata from the first week of February 2017 as the basis of the study. Collecting metadata from longer periods in time would have been possible, but would have increased the computation time of the scripts used to perform the study. We split the chosen time period into five-minute intervals. In each interval, we calculated the rate of failed transfers between each unique pair of source and destination of file transfers. From this point, such pairs will be referred to as links. When the failure rate of a link stayed above a given threshold for more than a given number of consecutive five-minute intervals, the testing script triggered an alarm. The length of time required before triggering an alarm is called the failure time. For an alarm to be triggered we also required at least ten attempted transfers across the link in each of the five-minute intervals that triggered the alarm. This prevented sites with an insignificant number of transfers, or no transfers, showing up as faulty.

With this testing script in place, we could now vary the thresholds for alarms to be triggered, and see how that influenced the number of alarms being put out. This is shown in Figure C.1. For values of failure rate between 10% and 90%, and failure time between 15 minutes and 45 minutes, the number of alarms triggered numbered a few hundred. We were hoping that by selecting specific values for these variables, the number of alarms triggered in our tests would be of the same order as the number of tickets filed to GGUS in the same testing period. Unfortunately, when looking up these tickets, we were only able to find three. The number of tickets found for any period, not only the week selected for testing, was of the order 10. This low number of tickets is problematic, not only for the fact that this is two orders of magnitude lower than the number of alarms triggered in our study, but also because such a low number of tickets is far fewer than what would be needed to train an automated system properly. We also looked at which Grid sites showed up in more alarms. One of the sites in the top five was one mentioned in one of the three support tickets found. We realised in hindsight that we should not have organised the transfers by links for this study, but rather on a site by site basis. The failure rate of file transfers going in and out of each site individually is what shifters are studying. This would very probably have reduced the number of alarms triggered by my scripts significantly, as one faulty site may trigger alarms over a large number of links.




Figure C.1.: The number of alarms triggered in the testing script as a function of file transfer failure rate and failure time.

Making a system that automatically submits a support ticket when the file transfer failure rate into or out of a Grid site stays above some critical value over some period of time would certainly be possible. One would only have to decide on which threshold values to use for acceptable failure rates and failure times. What we showed was that, due to the low number of historical shifter tickets, such a system could not be trained in any meaningful way by observing past shifter decisions. This result led us to abandon the approach of using support tickets in the project.

C.4. Seasonal Trend Decomposition

As the method of using shifter tickets to train an anomaly system did not pan out, we shifted our attention towards another method of anomaly detection, namely one based on *seasonal trend decomposition*. Seasonal trend decomposition is a tool used in time series analysis. A time series can be defined as a set of data points sorted by time. The number of failed or successful transfers to or from a Grid site in some time interval is a good example of a time series. When using time series analysis for anomaly detection, the aim is usually to find outliers in the data relative to some expected or regular signal.

Seasonal trend decomposition means splitting the time series into three distinct parts: The overall trend of the data, the seasonal component, and the randomly fluctuating residue. This method is useful when the time series displays periodicity or cyclic behaviour. There are several tools and methods available for performing this decomposition. One is outlined in Ref. [121]. An example of such a decomposition is shown in Figure C.2. Any large fluctuation in the residue is labelled as an anomaly.

Seasonal trend decomposition could automatically take into account the fact that what




Figure C.2.: Seasonal trend decomposition. From top to bottom: Original time series data, seasonal component, overall trend, and residual. A large deviation from zero in the residual is classed as an anomaly [122].

constitutes an anomaly may change over time. High network traffic at certain times of day or on certain days of the week could lead to higher failure rates, without this necessarily indicating an anomaly in the system. However, when investigating the time series data of failed and successful file transfers going in and out of all the different Grid storage locations, it quickly became apparent that this data is too chaotic for such a method to be practical.

We selected two weeks of file transfer metadata from July 2017 and separated it into one-hour intervals. For each storage location on the Grid, we plotted the number of failed and successful transfers, as well as the fraction of failed transfers, being sent to and from it. Two such plots can be seen in Figure C.3. The plots shown were two of the more densely populated plots that were made. Although there are thousands of file transfers taking place on the Grid every hour, on a site by site basis, the data is too sparse for any trends to appear over the random fluctuations. Plotting the transfers in wider time intervals could lead to less noisy plots in which cyclical trends could appear. However, anomalies would need to be detected and resolved as quickly as possible. Any system that would need more than an hour to detect an error would not be practical. On account of these results, the approach using seasonal trend decomposition as a tool in anomaly detection was abandoned.

C.5. Machine Learning Approach

Instead of trying to make a system to identify anomalous states of the DDM system itself, we had the idea of making a system that could predict whether or not a transfer was likely to fail before being transferred. To achieve this goal we wanted to employ machine learning (ML), a branch of computer science and statistics in rapid development. Such a system could be used to stop transfers before being submitted if the probability of failure,




Figure C.3.: All file transfers being transferred *from* (a) Grid site *FZK-LCG2* and *to* (b) Grid site *CA-MCGILL* in the chosen time period. An extended period of high file transfer failure rate can be seen in (b), starting at around 200 hours.

calculated by the system, became too large. This could help bring down the rate of failed transfers and thus increase efficiency. The system could also be used as an anomaly detection system if a large number of file transfers classified as likely to succeed suddenly started failing.

C.5.1. A Brief Introduction to Machine Learning

Machine Learning is a branch of applied statistics. ML methods are usually employed using computers, due to the size of the calculations involved. The goal of ML is to build computer systems that can make decisions or predictions based on experience, rather than having to be explicitly programmed. A commonly used definition of ML, which encapsulates this idea in a general form, was given by Tom Mitchell in 1997:

A computer program is said to learn from experience E with respect to some class of tasks T and performance measure P , if its performance at tasks in T , as measured by P , improves with experience E [123].

It should be pointed out that ML is a deep and complex field of study. ML algorithms are used to solve a host of problems in society: automatic fraud detection in credit card transactions, facial recognition software, and self-driving cars are just a few examples. We will barely be scratching its surface in this text. For a more thorough introduction to the subject we recommend the book *Deep Learning* by Goodfellow, Bengio, and Courville [124].

Several tools and methods exist for performing machine learning, but they all contain the same four basic components: A data set, a model or models, a cost function, and an optimisation algorithm. The cost function is sometimes called the error function. Using Mitchell's definition of ML, the model is the task T that needs to be solved, the data set provides the experience E , while the loss function gives the performance measure P . The optimisation procedure is what the model uses to improve performance with experience. Linear regression is a simple example that falls under this definition of ML. In this case, the model is simply a linear equation $y = \theta_1 x + \theta_2$. If some data set containing values \mathbf{x}, \mathbf{y} is expected to follow a linear distribution, the regression procedure aims to determine the values of the parameters θ_i that best fit the data. This is done by minimising the error function, which is usually the sum of the squared residuals between the model prediction

and the data set values in each data point. Minimising this loss function in this way is known as least squares regression. Adding additional points to the data set used in the regression would lead to a better linear fit with less uncertainty on the fit parameters. This embodies the ML concept of improving, or learning, with experience.

In ML literature, events or entries in the input data set are often referred to as *examples*. The variables, measurements or characteristics belonging to each of these examples are referred to as *features*. There are two main types of ML tasks: *supervised learning* tasks and *unsupervised learning* tasks. Supervised learning algorithms aim to map the features of the input data set \mathbf{x} to some output function $f(\mathbf{x})$. This is achieved by providing the chosen ML algorithm with information about the true or desired value of the output function for each example in the input data set. The algorithm tries to minimise the residual between the goal value and chosen value of the output function for each example in the input data set. This is done during the optimisation or training step. Supervised learning tasks usually falls into one of two categories: *classification* tasks or *regression* tasks. Classification tasks involve putting one or more labels on the data being analysed. Image recognition is an example of a classification problem. Regression tasks involve determining some continuous output function. The example of linear regression given previously is a simple case of such a task. Unsupervised learning algorithms are not given any information about the true or desired value of the input. Unsupervised algorithms aim to deduce the underlying structure of the input data without any human steering. So-called *clustering algorithms* are the most common type of unsupervised ML tasks. Clustering algorithms try to group examples in the input data based on their closeness in the space of their features. The field of ML is further divided into specific algorithms and tools. In this text, we will only focus on supervised learning techniques, specifically the method based on so-called boosted decision trees (BDTs). This method is covered in more detail in the next section.

The concepts of *underfitting* and *overfitting* are important in ML development. Underfitting occurs when the chosen ML model fails to conform to the input data. One reason for underfitting may be the lack of complexity in the chosen model. Take the linear regression example used above. This straight-line model will never be able to accurately agree to the input data if the data points are drawn from a non-linear distribution. Underfitting may also arise if the optimisation or training step is prematurely ended, or if the input data set is too sparse. Overfitting, on the other hand, occurs when the model starts conforming to the random fluctuations of the input data, rather than the underlying statistical distributions of the data. This means that the model will generalise poorly to new data. Using a model of too high complexity, one having too many tunable parameters, may lead to overfitting. It is hard to determine whether a model is suffering from over- or underfitting without testing the model on data not included in the optimisation of the model. For this reason, it is common to use two or more input data sets instead of just one in the development of ML models. One of these sets is the so-called *training set*, used to optimise the chosen ML model. Another is the so-called *testing set*. The testing set is not used to tune the parameters of the model, but rather to check that the model generalises well to unseen data. A clear sign of overfitting is that the model performance on the training set is much better than on the testing set. It is crucial that the training and testing data sets follow the same probability distributions. One way of ensuring that they do is to use the so-called *holdout method*. Instead of using two separately collected data sets for optimising and cross-validation, a random subset of one single input data set is set aside to become the testing set.

In addition to BDTs, two other common ML methods were considered for this project: artificial neural networks (ANNs) and support vector machines (SVMs). ANNs are inspired by the brain, made up of interconnected artificial "neurons". The neurons are connected in a graph, typically organised in layers; consisting of an input layer, one or more so-called hidden layers, and an output layer. Each layer consists of several neurons, and each neuron takes the information from the neurons in the previous layer, calculates an output, and transmits it to the next layer of neurons. SVMs are most often used for binary classification problems. This method aims to find the hyperplane in the feature space of the examples in the training set that best separates the two classes.

Any parameter of an ML method that is not determined through the training step of the algorithm is known as a *hyperparameter*. The number of hidden layers in an ANN or the number of decision trees used in a BDT algorithm are examples of hyperparameters. The choice of certain hyperparameters may significantly impact the performance of the ML model. A common way of determining hyperparameters is to tune them, automatically or by hand, to maximise the performance of the algorithm on the testing data set. However, tuning the hyperparameters in this way is sometimes problematic as it may reintroduce a kind of overfitting. The model may perform well on the testing set, but may fail to generalise to unseen data. The solution to this problem is often to use a third data set when building an ML model, along with the training and testing set, the so-called *validation set*. The validation set is analogous to the testing set and is used to check that the model generalises well to new data with the chosen hyperparameters.

Machine Learning algorithms can be very powerful tools, able to find patterns and higher-order correlations in the input data that simpler methods fail to capture. However, ML algorithms are often criticised for being "black boxes", meaning that it can be hard to know or interpret their internal state. ML methods can also be very computationally expensive.

C.5.2. Boosted Decision Trees

A decision tree is a graph-like structure, often used to solve classification problems. Decision trees are so named for their loose resemblance to trees, with a single trunk furcating into branches and leaves. A decision tree is, in essence, a nested list of "if-then-else" statements. Each of the cases being classified starts at the so-called *root node* of the decision tree. A decision or cut is made on one of the attributes or features of the case. The root node is linked to two or more nodes in the tree and the outcome of the decision determines which of these nodes the case is sent to next. At each node in the tree, a new decision is made on one of the features of the case, and so on, until the case ends up in a so-called *leaf node*. Each leaf node is labelled as belonging to one of the output classes of the classifier. Every case ending up in a specific leaf node is assigned to the class associated with said node. A schematic of a simple decision tree is shown in Figure C.4.

Decision trees can be used as-is in ML applications. A decision tree is then constructed, and the threshold values at each node are optimised such as to maximise performance on the input training data set. However, such trees tend to generalise poorly to unseen data. Single decision trees also tend to grow very large, even for relatively simple tasks. A more robust use of decision trees in ML applications is to pair them with a so-called *boosting* procedure. Boosting is a technique that combines several weak models or *learners*, with poor discriminatory power, into one strong learner. The power of the individual weak models can be only slightly better than random guessing, and the ensemble learner can still




Figure C.4.: A short decision tree. The input is allocated into one of two classes, S_1 or S_2 , based on decisions made at each of the five nodes. The cuts are made on three features, a_1 , a_2 and a_3 , of the input data.

achieve good precision. ML methods combining decision trees with boosting are simply called Boosted Decision Trees BDTs. The weak learners in question may be as simple as "tree stumps", decision trees consisting of just a single node.

Several algorithms exist for performing some form of boosting. In this text, we will focus on adaptive boosting (AdaBoost) [125]. As mentioned previously, the basic concept of boosting is to combine several weak or simple ML models into one powerful model. AdaBoost achieves this goal in an iterative, sequential manner, where the outcome of the training or optimisation of one weak model influences the training of the next. First, the initial model is trained on the training data set. For this first optimisation procedure, each example in the input data is given equal importance. For the next iteration of the boosting procedure, examples in the training set that were misclassified by the first model are weighted up, while correctly classified examples are weighted down. This re-weighting of the input data is performed after each new model is trained. This procedure of training and re-weighting is then repeated a predetermined number of times. How many weak learners to include in the strong model is an important hyperparameter of BDTs. When all the weak models have been trained, they are combined with a weighted sum into the strong model. The weight of each weak model in this sum is inversely proportional to the rate of misclassification of that model. A schematic of an example boosting procedure is shown in Figure C.5. In this figure the aim is to build a classifier that can separate the input into two classes, "+" and "-", based on two features. The training set contains ten examples that have been correctly labelled, making this a supervised learning problem. The example uses three weak learners, which are simple tree stumps, meaning single splits in one of the two features. The first tree is trained on the initial distribution, splitting it in two such as to best separate the input into the correct classes. Some of the examples

are seen to be misclassified. Before training the next tree, these misclassified examples are given higher importance or weight than the rest. These examples are circled in the figure. The next tree is then trained using the same procedure as the first tree. Due to the re-weighting procedure, the optimal split found by this tree is not the same as that of the first tree. The re-weighting and retraining is then performed one more time for the third tree. Finally, the BDT classifier is built from a linear combination of the three tree stumps. We see that even though each simple tree has some misclassification, the misclassification rate of the final classifier is zero.




Figure C.5.: Schematic of the boosting procedure. The two dimensions of the boxes represent the two features of the data. The first two weak classifiers discriminate in the first feature, while the third classifier discriminates on the second feature. Since the rate of misclassification (3 examples) is the same for each weak learner, the weights α_i will be equal in this example. Each example in the input data is correctly classified by two out of three weak classifiers, which means that the final classifier has a 100% success rate.

BDTs have become ubiquitous in the high energy particle physics (HEPP) community. They are powerful and relatively simple to set up. The AdaBoost algorithm has also been shown to be resilient against overfitting [126]. However, like other ML algorithms, BDTs are criticised for their obscurity. It is hard to determine their internal state. This problem is alleviated somewhat because BDTs allow for the easy extraction of metadata such as *feature importance* from the model. Feature importance is a tally of how many times each feature is used to split the input data at a node in a decision tree. This means that features with high feature importance scores have more influence on the BDTs decision-making process.

C.5.3. Collecting File Transfer Metadata

Before the ML algorithm could be constructed, we had to obtain a data set for it to train on. The goal of the project was to train the algorithm to detect whether a file transfer on the Grid would fail or be successful. This meant that we needed historical metadata

of previous file transfers, containing features suspected of being correlated with the file transfer failure rate. Table C.1 lists the features selected for training. This information was gathered off the two Grid metadata databases mentioned previously. One database contains every transfer previously made on the Grid. Whether or not a file transfer was successful, the size of the file being transferred, and the file transfer protocol used for the transfer is stored in this database. This database became known in the project as The CERN Database due to its physical location at CERN. The CERN Database also contains information about the state of the Grid. The number of file transfers queued, network throughput, and how many successful file transfers already occurred in the last one or six hours is stored for every link in the network. This information is updated every five minutes in the database. The second database tracks the latency and packet loss of each link in the network. This information is updated every 10 minutes. This database was referred to as the Chicago Database in the project, due to its physical location at the University of Chicago. All of the features concerning the state of the Grid are variable in time except for the so-called network closeness. This feature is static in time, and its value for each link on the Grid was stored locally in JSON format for easy access.

Table C.1.: A list of the file transfer and Grid state features used for ML training in this project.

Whether or not the file transfer was successful. This is the target feature of the classifier.
The size of the file being transferred.
The file transfer protocol used to perform the file transfer.
The number of successful file transfers performed on the same link in the last hour.
The number of successful file transfers performed on the same link in the last six hours.
The mean throughput, in MB/s, across the link in the last hour.
The number of files queued for transfer across the link at the time of the file transfer.
The latency across the link.
The packet loss across the link.
The network closeness of the link.

Collecting the file transfer and Grid status data proved non-trivial due to the information being scattered over several sources. Even the file transfer and Grid status data stored in the CERN Database were not automatically linked. The data used in this project were collected in September 2017 and correspond to several days worth of file transfers on the Grid. First, the information on each file transfer performed in the selected period was collected by querying The CERN Database. Grid state data for the same period were collected from the CERN and Chicago databases. For each file transfer, the correct Grid state features were added to its data set entry by looking up the information in the five- or ten-minute interval of the Grid state data corresponding to its timestamp. The example was dropped from the data set in the cases where these look-ups failed. Some complications arose because of differences in naming conventions of Grid sites between the CERN and Chicago databases. This problem was solved by making a script mapping the site names on one database to the corresponding site names on the other. The file transfer protocol feature is not a numeric value. Each possible file transfer protocol was therefore assigned an integer number, as most ML tools are only able to read numeric input.

Two features of the transfers were used in testing, but later dropped: whether or not the file had been retried, meaning that the file attempted to transfer but failed, and the time elapsed between the file transfer was initiated until it was executed. A file having

already failed to transfer was shown to be a strong indicator that the transfer would fail again. These features were not used to train the ML algorithm, as they can only be known for past file transfers and not when initiating a new transfer. These features may be used if instead of making some system predicting whether or not a file transfer is going to fail, we wanted to make some anomaly detection system that monitored file transfers being made. In this case, the aim of the system would not be to stop transfers with high failure probability, but rather to throw an alarm if the rate of transfer failure grows larger than what is being predicted by the system.

After collecting the file transfer examples for the training set, with all the relevant features, the data were shaped into an N by 10 matrix, where N is the number of file transfers included in the data set. For the selected period, N was of the order 10^5 . This matrix is often referred to in statistics and ML as a design matrix. Each row in a design matrix correspond to a single data point or example, a transfer in this case, while each column correspond to a feature of the transfers. The so-called *target feature*, whether or not the file transfer is successful, is what the ML algorithm will use to optimise its behaviour.

Some number of duplicate entries, identical rows in the design matrix, emerged when compiling the data set. These may stem from quickly retried file transfers. Identical entries in the training set of an ML algorithm may negatively influence its performance, as it essentially doubles the weight of certain examples in the data set. Therefore, we removed these entries before training the BDT classifier.

C.5.4. Building the Classifier and Analysing Hyperparameters

As mentioned previously, three ML algorithms were considered for this project: BDTs, SVMs, and ANNs. The BDT, ANN, and SVM classifiers were all built using scikit-learn, an open source ML library for the Python programming language [127]. Scikit-learn is well supported and documented, making it well suited for this project. BDTs performed better out-of-the-box than both SVM and ANNs in our initial tests, and were therefore selected for further development. ANNs are powerful tools, but they have a large space of hyperparameters. Achieving high performance using an ANN often requires a lot of effort finding and tuning these parameters.

After selecting a BDT method for the project, we researched how the choice of hyperparameters influenced the performance of the classifier. We focused on how many trees to include in the boosting procedure, and the maximum depth of these trees, meaning the number of allowed layers of nodes in the decision trees. The decision tree in Figure C.4 has a depth of 3. To determine good choices for these hyperparameters, an AdaBoost BDT classifier was trained for maximum depths between 1 and 8 and between 1 and 2,000 trees. Figure C.6 shows the rate of misclassification on the training and testing data sets as a function of the number of trees for every choice of maximum depth. The important graph to observe in the figure is the error or misclassification rate on the testing set. Whereas the error on the training set falls to zero, the error on the testing set approaches some non-zero value in all the plots. The mode of the misclassification rate on the testing set, meaning the most frequently occurring value of the graph, is plotted as a red line in the figure. This value is a makeshift indicator of the smallest possible misclassification rate for unseen data for the chosen tree depth. The aim should be to minimise this value. We observe that this value is more or less constant for values of maximum tree depth of three or more. This means that the classifier does not perform noticeably better by increasing the complexity


of the individual trees making up the BDT above this number. Based on this finding we chose to use a maximum tree depth of four in our BDT classifier. In Figure C.6 (d) we see that for this choice of maximum tree depth, the misclassification rate on the testing set is flat when the BDT is made up of more than around 500 trees. This means that using more than 500 trees in the classifier will not lead to any improvement in performance. In the computationally costly field of ML, such inefficiencies are best avoided. Therefore, we chose 500 as the number of trees to make up the BDT.

Figure C.6 also illustrates AdaBoost's tendency to avoid overfitting. In a model displaying overfitting behaviour, the misclassification rate on the testing set would not only decrease and flatten out towards some minimal value, but would at some point of the optimisation start to grow. This behaviour is not observed here; Even long after the training error has reached zero, the testing error does not seem to rise.


The number of successful transfers in the input data set outnumbered the failed transfers approximately ten to one, meaning that the global file transfer failure rate on the Grid is around 10%. We noticed during testing that the BDT algorithm was better able to correctly classify successful transfers than failed transfers. By balancing the number of failed and successful examples in the training set, the precision of the classifier on failed transfers improved to the same level as the precision on successful transfers. We balanced the training by splitting the original training set into two parts, one containing only failed transfers, and the other containing only successful transfers. The new training set was then constructed by drawing half of the desired number of training examples from one sample at random, and the other half from the second.

A common measure of classifier performance is the area under curve (AUC) of the so-called receiver operating characteristic (ROC) curve. The ROC curve is a graph of the true positive rate of the classifier, plotted against the false positive rate, for all possible *class thresholds*. The output of a binary classifier, like the one used in this project, is a number referred to as the score of the classifier. An example receiving a high score means that the classifier deems the example more like the first class while receiving a low score means that the classifier deems the example more like the second class. In our case, the two classes are whether or not the file transfer is going to fail. The class threshold is the score separating the two classes. Every case receiving a score higher than the class threshold will be classified as belonging to the positive class, while cases receiving a score lower than this threshold are classified as negative. The integral of the ROC curve is called the AUC. Figure C.7 illustrates how the ROC curve emerges when varying the class threshold. The worst possible binary classifier, one making random guesses will have an AUC of 0.5. A perfect classifier, on the other hand, one that never misclassifies a single case, will have an AUC of one.

We wanted to investigate how the size of the training data set influenced the performance of the classifier. Collecting the training data from the CERN and Chicago databases was quite time consuming. Therefore, we investigated the amount of data needed to attain a good performance of the classifier. This test was performed by training the BDT algorithm several times for different amounts of training data, between 10 file transfers and 5×10^5 file transfers. The performance of the algorithm was measured for each of these cases using the AUC on a testing set containing 5,000 entries. AUC as a function of training set size is plotted in Figure C.8. The classifier performance can be seen to increase as a function of training sample size, as expected. However, there are diminishing returns on the performance gained by adding more training data. Figure C.8 shows the AUC growing by around 0.2 when increasing the training data sample size from 10 entries to 10^3 entries,




(a) Area under ROC curve: 0.89




(b) Area under ROC curve: 0.93




(c) Area under ROC curve: 0.94




(d) Area under ROC curve: 0.94




(e) Area under ROC curve: 0.93



(f) Area under ROC curve: 0.93



(g) Area under ROC curve: 0.94



(h) Area under ROC curve: 0.95

Figure C.6.: Training and testing misclassification error as a function of trees used in the BDT estimators, for eight values of maximum depth of the tree graphs.




Figure C.7.: Schematic of the construction of the ROC curve [128]. The green arrows indicate the movement of the curve as the class threshold is varied.

while the gain is much smaller when increasing the training sample size from 10^3 entries to 10^5 entries.




Figure C.8.: AUC as a function of training data set size.

We neglected to set aside a portion of the input data set for validation after tuning the hyperparameters of the BDT. There is a slight risk that this may have introduced a bias during training; rather than the model generalising well to unseen data, the hyperparameters chosen for the BDT may be optimised for the testing set. However, the probability of this causing issues was deemed low enough to not warrant any action.

C.5.5. Results and Performance

Figure C.9 shows the BDT classifier score distributions of the entries in the training and testing data sets. The two classes appear well separated in both cases. The ROC curve of

the testing data classifier scores is shown in Figure C.10. The curve has an AUC of 0.96.




Figure C.9.: BDT classifier score distributions for the training (top) and testing (bottom) data sets.




Figure C.10.: ROC curve of the BDT classifier scores on the testing data set.

We extracted the feature importance distribution of the BDT to study which features had the largest impact on the classification. This distribution can be seen in Figure C.11. From this figure, we see that the size of the file being transferred and the network latency of the link are the two features with the largest discriminatory properties. Packet loss, the file transfer protocol used for the file transfer, and network closeness were the least important features.




Figure C.11.: Feature importance distribution of the BDT classifier.

C.6. Conclusion

The goal of this project was to make a system for automating Grid monitoring jobs currently being done by human shifters. We focused on the ATLAS DDM system. It seems feasible to make a system that automatically files a support ticket when the number or rate of failed file transfers becomes too large. We fear that such a system would be too simplistic to be practical; what constitutes an anomalous state of the DDM system is not static in time, but changes with the situation. A hard-coded automated system would not have the same capability of exercising judgement based on previous experiences. We had the idea of using past shifter decisions in the form of historical support tickets to construct a more sophisticated system. The lack of past support tickets made this impractical.

We also investigated building an anomaly detection system for the DDM system using seasonal trend decomposition. This method was abandoned because the available data was too noisy for any seasonal component to emerge from the time series of file transfers being performed over each link in the network.

We demonstrated an ML method, using BDTs, for predicting whether or not a file transfer on the DDM system was likely to fail, based on the features of the file being sent and the state of the Grid at the time of transmission. This system could be used to lower the rate of failed transfers being performed on the grid by stopping transfers with large probabilities of failing from taking place. It could also be used to construct an anomaly detection system by throwing an alarm if the rate of failed transfers grows larger than that predicted by the classifier. The largest obstacle for such a system to be practical, at the time of the project, was the cumbersome process of collecting file transfer and Grid status metadata to train the ML algorithm. Due to this issue, the system has so far not been put to use monitoring the ATLAS DDM system.

D. Physics Education Projects

D.1. Introduction

During my PhD, my contract at the University of Oslo required me to participate in teaching in addition to my research. I found that I enjoyed teaching, and I became interested in the didactics of physics. While working as a teaching assistant in a laboratory course at the department of physics, I started thinking about educational physics experiments that can be performed with limited laboratory equipment.

Modern smartphones contain a lot of sensors: microphones, cameras, accelerometers, gyroscopes, etc. Having become so commonplace in society, smartphones can provide economical alternatives to costly laboratory equipment. I wanted to test out a few different ways of using a smartphone as a physics laboratory, with the aim of using these methods when teaching physics laboratory courses.

I performed three different physics measurements using a smartphone. This chapter contains a summary of these three experiments, as well as copies of the three papers written by me presenting their results.

D.2. Measuring the Acceleration of Gravity

If the observer of an emitted sound signal is moving with respect to the source of the sound, the frequency f of the observed signal will be different from the emitted frequency f_0 . This phenomenon is known as the Doppler effect. The magnitude of this shift in frequency can be calculated using Doppler's equation

$$f = \frac{c \pm v_O}{c \pm v_K} f_0, \quad (\text{D.1})$$

Where c is the speed of sound in air, while v_O and v_K are the speeds of the observer and source respectively, with respect to the air.

If the source is stationary, and the emitted frequency is known, Equation (D.1) can be used to measure the velocity of the observer. Or, with repeated frequency measurements, the *acceleration* of the observer. Such a measurement is what I perform in my article *Measuring the acceleration of gravity g using Doppler Shift*, published in Norwegian in Ref. [129]. The experiment is a recreation of the one described in Ref. [130].

A smartphone is set to record audio. The smartphone is held underneath a speaker and dropped to the floor. The speaker is sending out a pure 2 kHz sine wave. In the resulting audio recording, one can clearly hear the recorded frequency being shifted to lower frequencies as the phone falls.

The audio recording is split into 15 equal intervals. The recorded frequency is assumed to be constant in each of these intervals. The observed frequency in each interval is determined using the *zero crossing method* [131]. The zero crossing method works by finding the length of time between successive crossings from negative to positive values in the audio waveform. This procedure is illustrated in Figure 3 in the paper.

The acceleration of the smartphone is proportional to the slope of the shift in the observed frequency. The smartphone is free-falling, so its acceleration is equal to the gravitational acceleration g . The slope of observed frequency values is found by linear regression. The gravitational constant is measured to be $g = 10.1 \text{ m/s}^2$, which is within 3% of the current best-fit value of 9.81 m/s^2 . No attempt was made to calculate the uncertainty of the measurement.

To calculate the value of g using the method outlined here, the speed of sound in air must be known. A procedure for measuring the speed of sound in air using a smartphone is outlined in the next section.

D.3. Measuring the Speed of Sound in Air

In my paper *Measuring the Speed of Sound in Air Using a Smartphone and a Cardboard Tube* I measure, as the title suggests, the speed of sound in air [132]. The only piece of equipment needed for this experiment, in addition to a smartphone, is a tube closed in one end.

For a pure sine wave of frequency f and wavelength λ , propagating in a medium, the propagating speed in the medium c is given by

$$c = f\lambda. \tag{D.2}$$

When an audio sine wave travels down the half-closed tube and meets the end, part of the signal will be reflected back up the tube. For certain frequencies of the incoming wave, the reflected sine wave will interfere constructively with the incoming sine wave, setting up a resonance. A microphone placed outside the tube will pick up an increase in recorded amplitude for these frequencies compared to signals of other arbitrary frequencies. These standing wave resonances occur when the incoming sine waves have wavelengths equal to an odd number of acoustic lengths of the tube. Therefore, if the resonance frequency and the length of the tube are known, the propagating speed c of the wave can be found. This quantity is also known as the speed of sound in the medium.

I place my smartphone outside the tube. A slow audio sine wave sweep from low to high frequency is emitted from a signal generator app installed on the phone. The phone is set to record audio during the sweep. At the resonance frequencies of the tube, the waveform of the recorded audio spikes. The frequency of the recorded audio is determined using the zero crossing method described in the previous section. The resonance frequencies are linearly spaced. I perform a linear regression through the resonance frequencies. The speed of sound is given by $c = 4aL$, where a is the slope of the regression, and L is the acoustic length of the tube.

The speed of sound in air is measured to be $c = 335 \text{ m/s}$. This result is within 3% of the theoretical speed of sound in an ideal gas at the same temperature as air when performing the experiment, which is $c_T = 345 \text{ m/s}$.

D.4. Measuring the Strength of Earth's Magnetic Field

Most modern smartphones have compasses. The compass in a smartphone operates by measuring the direction of the Earth's magnetic field using a Hall sensor. Or, more accurately, three Hall sensors, each measuring the magnetic field strength in one of the three spatial coordinates.

I was supervising an introductory laboratory course in electromagnetism in the autumn of 2018. One of the laboratory exercises saw the students measuring the field strength of the Earth's magnetic field using a rotating induction coil. To check their results, the students compared their results to that given by a magnetometer app on their smartphones. I noticed that the absolute value of the magnetic field strength as measured by the smartphones often varied with the orientation of the phone itself. I suspected that there were either a difference in calibration between the three independent Hall sensors in the magnetometer, or that the measured values were being offset by internal magnetic disturbances inside the smartphone itself.

In my paper *Measuring Earth's Magnetic Field Using a Smartphone Magnetometer* [133], I devise a method for eliminating any offsets due to internal magnetic fields in the smartphone when measuring magnetic field strength. I then use this method to measure the magnetic field strength in Oslo.

I perform two magnetic field measurements. For one of these measurements, I rotate the smartphone 180 degrees. This rotation flips the sign of the measured magnetic field strength, while the sign of a potential internal offset remains the same. When the two measurements are subtracted and divided by two the internal offset is cancelled.

This subtraction procedure is performed three times for each individual spatial coordinate. This means that one measurement is performed for each individual Hall sensor in the smartphone. The average magnetic field strength of the three axes is taken as the final magnetic field strength measurement. The earth's magnetic field strength in Oslo is measured to be $B_{\text{earth}} = 50.5 \pm 0.4 \mu\text{T}$.

Måling av tyngdeakselerasjonen g ved hjelp av Dopplerskift

Målet med dette forsøket er å beskrive eit eksperiment der ein kan måle tyngdeakselerasjonen g ved hjelp av ein mikrofon i fritt fall vekk frå ei lydkjelde. Forsøket skal kunne utførast med enkle verktøy og forkunnskapar, slik at det skal kunne vere mogleg å gjenskape for studentar.

Simen Hellesund Fysisk institutt, UiO

Når den relative avstanden mellom ei kjelde og ein observatør av ein lyd endrar seg med tid, vil frekvensen f som observatøren måler, vere ein annan enn den frekvensen f_0 kjelda sender ut. Denne effekten kallar me dopplereffekten. Endringa i frekvens er avhengig av den relative farten mellom kjelda og observatøren og kan bli rekna ut med formelen


$$f = \frac{c \pm v_0}{c \pm v_k} f_0, \quad (1)$$

der c er lydfarten i mediet som lyden forplantar seg i, mens v_0 og v_k er farten til observatøren og kjelda i forhold til mediet. Forsteikna i likning 1 er bestemt av om kjelda og observatøren beveger seg mot (+) eller frå (-) einannan. Likninga antek òg at rørsla til kjelde og observatør er lineær i forhold til einannan. I dette forsøket vil mediet lyden forplantar seg i, vere luft. c er difor lydfarten i luft.


Om ein observatør er i fritt fall vekk frå ei lydkjelde i ro, vil den observerte frekvensen grunna dopplerskift vere

$$f(t) = \frac{c - gt}{c} f_0 = f_0 - \frac{g f_0}{c} t, \quad (2)$$

der g er tyngdeakselerasjonen, og t er tida sidan observatøren starta å dette. Dette er likninga for ei rett linje. Oppførselen til den målte frekvensen frå observatøren startar å dette ved tida t_0 til hen



Figur 1. Målt frekvens som funksjon av tid av observatør som er i fritt fall vekk frå ei lydkjelde. Observatør startar å falle ved tida t_0 og stoppar opp ved tida t_1 .



Figur 2. Telefonen blir halden like under høgtalaren. Han blir deretter sleppt og dett fritt ned til golvet. Tidene t_0 og t_1 er markert for samanlikning med figur 1.

sluttar å dette ved tida t_1 er avbilda skjematisk i figur 1.

Ein kan sjå av likning 2 at stigningstalet a til den rette linja som oppstår i frekvensutviklinga til den dettande observatøren, er proporsjonal med tyngdeakselerasjonen som

$$a = -\frac{g f_0}{c}. \quad (3)$$


Eksperiment

Forsøket som er utført her er ei gjenskaping av det som står beskrive i [1].

Forsøket vart utført i eit lydtempa rom på Fysisk institutt ved Universitetet i Oslo. Det er nok ikkje naudsynt for forsøket å ha eit slikt rom tilgjengeleg, men ein bør prøve å eliminere bakgrunnsstøy så mykje som råd.

Ein høgtalar kopla til ein signalgenerator vart plassert i taket på rommet, om lag 2 meter over golvet. Høgtalaren vart sett til å sende ut ei konstant sinusbølge på 2000 Hz.

Ein smarttelefon med ein applikasjon for innspeling vart brukt som observatør. Telefonen vart sett til å ta opp lyd med ei samplingsfrekvens på 44100 Hz. Telefonen vart halden like under høgtalaren før han fekk dette i fritt fall ned til golvet. Dette er skjematisk framstilt i figur 2. Lydfilene vart lagra i .wav-format. Dette formatet gjer det enkelt å analysere innsamla data. I dette forsøket



Figur 3. Nullkryssingsmetoden. Den svarte kurva er lyd-signalet som vert utsendt av kjelda. Dei raude punkta viser samplingspunkta som vert gjort av telefonen. Dei gule punkta er nullkryssingspunkta som vert brukt til å fastslå frekvensen til signalet.

vart analysen utført i programmeringsspråket Python. På lydopptaket kan ein klart høyre at frekvensen endrar seg på grunn av dopplerskift.


Lydfarten i lufta vart estimert ved å ta temperaturen i rommet for så å slå opp verdien i ein tabell. I dette forsøket vart $c = 345,8 \text{ m/s}$ brukt.

Måling av frekvens

Sidan me vil gjere ei tilpassing til frekvensen på lydopptaket, treng me ein metode for å måle frekvens. Dette er ikkje så uproblematisk som det kanskje synest ved første augekast. Lydfila som vart spelt inn på mobilen, inneheld i første rekke berre informasjon om utslaget til mikrofonmembranen i kvart samplingspunkt. Dette må konverterast til frekvensinformasjon. Det viser seg at det er umogleg å måle den momentane frekvensen i kvart samplingspunkt. Ein er nøydd til å sjå på utviklinga til lyden i eit tidsintervall. Dess finare oppløysing ein vil ha på frekvensmålingane, dess lengre tidsintervall treng ein.

I tilfellet der ein har høg signalstyrke over bakgrunnen kan ein bruke den såkalla nullkryssingsmetoden (*zero crossing method*). Denne metoden går ut på å finne punkta i bølgeforma i lydfila der utslaga kryssar frå negative til positive (eller omvendt) verdier. Når ein har funne to punkt i bølgeforma på kvar side av ei slik kryssing, kan ein trekke ei rett linje mellom dei. Punktet der denne linja kryssar tidsaksen er ei tilnærming til nullpunktet til lydbølga. Avstanden i tid mellom eit slikt punkt og neste blir då ei tilnærming til perioden til bølga. Dette er framstilt i figur 3.

I dette forsøket vart lydopptaket av telefonen som dett, først skore til slik at kun tida i fritt fall er med på opptaket. Deretter vart fila delt inn i 15 delar. I kvar av desse subseksjonane vart null-



Figur 4. Dei målte frekvensane, bestemt med nullkryssingsmetoden, med den rette linja som gir best tilpassing til data.

kryssingsmetoden brukt for å finne frekvensen ved kvar nullkryssing. Frekvensen i intervallet vart tatt som snittet av dei individuelle frekvensmålingane.

Resultat

Frekvensane i dei 15 tidsintervalla er tilpassa med ei rett linje. Dette kan ein sjå i figur 4. Stigningstalet til denne linja sett inn i likning 3 gir $g = 10,1 \text{ m/s}^2$.

Forsøket vart gjentatt fleire gonger. Berre eit av resultatata (tilfeldig valt) er gjengitt her. Det vart ikkje gjort noko estimat av uvisse på den målte verdien for g . Både frekvensmålingane og verdien for lydfarten vil vere kjelder til uvisse. I samanlikning med desse kjeldene til uvisse kan luftmotstand reknast å vere neglisjerbar.

Konklusjon

I dette forsøket har tyngdeakselerasjonen g vorte målt ved hjelp av dopplerskift. Dette kan gjerast utan avansert utstyr og med beskjedne programmeringskunnskapar. Forsøket burde kunne gjennomførast av studentar på eit innføringskurs i eksperimentalfysikk. Forsøket gir læringsmoglegheiter både om dopplerskift, digitalisering av lyd, og om programmering.

Takk

Ein stor takk til professor Alex Read ved Fysisk institutt ved Universitetet i Oslo for utlån av utstyr for å gjennomføre forsøket.

Takk òg til Andreas Heggelund, Eli Bæverfjord Rye og Knut Oddvar Høie Vadla, for god hjelp som handlangarar og laboratorieassistentar.

Referansar

1. S.M. Torres & W.J. González-Espada. «Calculating g from acoustic doppler data». *The Physics Teacher* 44(8): s. 536–539 (2006).

Measuring the speed of sound in air using a smartphone and a cardboard tube

Simen Hellesund 

University of Oslo, Oslo, Norway

E-mail: simen.hellesund@fys.uio.no



CrossMark

Abstract

This paper demonstrates a variation on the classic Kundt's tube experiment for measuring the speed of sound. The speed of sound in air is measured using a smartphone and a cardboard tube, making the experiment very economical in terms of equipment. The speed of sound in air is measured to within 3% of the theoretical prediction.

Introduction

In recent years, smartphones have become packed with sensors; microphones, cameras, accelerometers, magnetometers, gyroscopes, thermometers, proximity sensors etc. Having become ubiquitous in society, smartphones can provide economical alternatives to expensive laboratory equipment in physics education.

Several papers have examined the use of smartphones in acoustic experiments for educational purposes [1–5]. A common experiment in physics education is to measure the speed of sound c in air, or other gasses, by observing standing acoustic waves in a tube. August Kundt first described this experiment in 1866 [6]. Such an experiment is therefore often referred to as *Kundt's tube*.

Parolin and Pezzi have shown how the experiment can be performed using two smartphones [7]. Yavuz has shown how it can be done using a single smartphone by partially submerging the tube in water [8]. The aim of this paper is to attempt to perform the experiment using only a smartphone and a cardboard tube, offering an alternative method to the one outlined in [8].

Theory

For a sinusoidal wave with constant frequency f and wavelength λ , propagating in a medium, the speed of sound in said medium is given by:

$$c = \lambda f. \quad (1)$$

This means that if we can determine both the frequency and wavelength of the wave, we can measure the speed of sound in the medium. For this experiment, the medium in question is air at room temperature and atmospheric pressure.

When an acoustic wave enters through the open end of a half-closed tube and hits the closed end, part of the wave is reflected back down the tube towards the open end. At specific wavelengths, the incident and the reflected wave form a standing wave. In the *antinodes* of the standing wave, the points on the standing wave where the amplitude is maximal, the amplitude of the standing wave is greater than the amplitude of the incident wave alone. The opening of the tube will always be a displacement antinode of the standing waves. The wavelengths at which the standing waves occur are called the *resonance wavelengths* of the tube. For the half-closed tube, the resonances occur when the length of the tube



equals an odd number of quarter wavelengths of the incident wave:

$$\lambda_n = \frac{4L}{n}, \quad n = 1, 3, 5, \dots \quad (2)$$

The number n is often referred to as the n th harmonic of the tube. L is the length of the tube.

The resonance frequencies f_n of the tube, the frequencies at which standing waves occur in the tube, can be found by combining equations (1) and (2):

$$f_n = \frac{cn}{4L}, \quad n = 1, 3, 5, \dots \quad (3)$$

We see that the resonance frequencies as a function of n is the equation of a straight line. The slope a of this line is given by

$$a = \frac{c}{4L}. \quad (4)$$

This means that if we can identify the resonance frequencies of the tube and fit a straight line to them, we can calculate the speed of sound as

$$c = 4aL. \quad (5)$$

It turns out that the acoustic length of the tube is slightly longer than its physical length. The position of the antinode at the tube's open end will be a small distance outside of the tube. A more accurate measurement can therefore be performed by adding a correction term δL to the length of the tube. Levine and Schwinger found this correction term to be $\delta L = 0.61r$, where r is the radius of the tube [9]. Using this correction to the length of the tube, the speed of sound is given by

$$c = 4a(L + 0.61r). \quad (6)$$

The theoretical speed of sound in air c_T can be calculated as

$$c_T = \sqrt{\frac{\gamma RT}{M}}, \quad (7)$$

where γ is the adiabatic index of air; R is the molar gas constant; T is the temperature of the air in Kelvin; and M is its molar mass. Equation (7) is only valid for an ideal gas. At room temperature and normal atmospheric pressure, air behaves close enough to an ideal gas for our purpose.

Setup

The following equipment is used for this experiment: a smartphone with a signal generator app

and a recording app installed; a cardboard tube, closed in one end; a thermometer; a tape measure; a computer.

The smartphone used is a Motorola Moto g6 Plus. There are several apps available both for generating audio sine waves and for recording. The function generator app must be able to perform a sweep over frequencies. The generator app used in this experiment is called *Function Generator* and is available in the Google Play Store [10]. The recording app must record audio with a sampling frequency known to the user. The recording app used for this experiment is called *Smart Recorder*, also available in the Google Play Store [11]. For users of Apple products, the recording app *Voice Recorder Lite: HD Audio Recording & Playback* [12] and the signal generator *Audio Function Generator* [13], both available from the Apple App Store, may be used.

The cardboard tube used is one intended for storing or shipping posters. One end is stopped by a plastic plug. The length of the tube is measured, using the tape measure, to be 47.6 cm. The inner diameter of the tube is measured to be 7.5 cm. The tube and the smartphone used in the experiment can be seen in figure 1.

The experiment is performed in an anechoic chamber at the physics department at the University of Oslo. Having access to such a room is not critical to the experiment, although one should aim to limit background noise as much as possible.

The thermometer is not used in the measurement directly but is used to measure the temperature in the room during the experiment. This is then used to calculate c_T . The temperature in the anechoic chamber is measured to be 24 °C.

Procedure

We place the smartphone such that the microphone is located in the opening of the tube. This is shown in figure 2. The phone is set to record audio with a sampling frequency of 44.1 kHz. While the phone is recording, the function generator app emits a pure sine wave. The sine wave sweeps from 50 Hz to 3000 Hz at a rate of 1 Hz s⁻¹. The audio recording is stored in .wav format. This format makes for easy data analysis later.

Measuring frequency

We need to know the frequencies at which the resonances occur. The function generator does not store any information on which frequency is emitted at which time. Thus, the frequency has to be inferred from the recording.

Frequency cannot be instantaneously measured. The signal must be monitored over some period of time to count how many times it oscillates per second.

In this case, where we are trying to identify a pure sine wave in a high signal to noise ratio environment, we can use the method known as the *zero crossing method* (ZCM) [14]. The ZCM works by determining the points in time where the waveform of the recording crosses from negative to positive values (or vice versa). These points are approximated by identifying the points on either side of such a crossing and drawing a straight line between them. The zero crossing point of the signal is approximated as the point at which this straight line becomes zero. The distance from one such zero crossing point to the next is an approximation of the period of oscillation of the signal. The inverse of this period is the frequency of the signal. The ZCM procedure is illustrated in figure 3.

The samples of the audio recording are split into one-second intervals. In each of these intervals, the ZCM is used to find all the zero crossings. The measured frequency in each interval is taken to be the mean of all the frequencies measured from these zero crossings. These frequencies are plotted as a function of time in figure 4.

Looking at figure 4, we notice that there are a few points where the ZCM clearly fails to identify the frequency emitted by the function generator. In our case, these points can be ignored, as they are far away from the points where the sound intensity of the recording peaks.

The results obtained using the ZCM are cross-checked by measuring the frequency in the recording using a *fast fourier transfer* (FFT) in each second interval. The most prominent frequency in the FFT spectrogram of each interval is taken to be the frequency of the signal. This method yields identical results to those obtained by the ZCM. However, the FFT method is significantly slower and is therefore not used in the final analysis.



Figure 1. The smartphone and the cardboard tube used in the experiment.




Figure 2. Schematic illustrating the placement of the tube and smartphone during the experiment.

Data analysis

The start of the audio recording proves too noisy to be of use. The low frequencies emitted by the function generator are perhaps causing the smartphone to vibrate. This region of the recording is therefore cut before the analysis.

The data analysis is performed in Python. The SciPy python library contains the packages needed for reading the .wav file, as well as for signal processing and curve fitting [15].

We need to identify the points where the sound intensity of the recording is maximal, as these maxima will occur at the resonance frequencies of the tube. First, the recording is split into the same one-second increments used to measure the frequency. The absolute value of the digital samples is taken in each of these intervals. The peaks in the resulting signal are identified using the `find_peaks` function of the SciPy signal processing package. The mean height of the peaks found in each second increment is taken to be the sound amplitude. The resulting distribution is shown in figure 5.

Next, the `find_peaks` function is used once more to identify the peaks in the sound amplitude distribution. These peaks are also shown in figure 5. The frequencies measured at the times of




Figure 3. The ZCM illustrated. The black line is the pure sine wave signal. The red points are the points sampled by the smartphone. The yellow points are the zero crossing points used to approximate the frequency of the signal.




Figure 4. Frequencies measured using the ZCM.




Figure 5. Sound amplitude of the audio recording.

Table 1. Frequencies at peak amplitudes.

n	Frequency (Hz)
1	544
3	867
5	1130
7	Unknown
9	1865
11	2247
13	2594
15	2906




Figure 6. Resonance frequencies of the cardboard tube with line of best fit.

these peaks are the resonance frequencies of the cardboard tube. These frequencies are recorded in table 1 in the results section. It seems one resonance frequency fails to give a peak in the sound amplitude distribution. The reason for this is unknown. The frequencies at $n = 7$ and above are therefore shifted up by one.

Results

The resonance frequencies identified in the previous section are listed in table 1. A straight line is fitted to these frequencies as a function of n using the `polyfit` function of the Python module NumPy [16]. The frequencies as well as the line of best fit is shown in figure 6. The slope of the straight line is used to calculate c using equation (6). The speed of sound in air at 24 °C is measured to be $c = 335 \text{ m s}^{-1}$.

Using equation (7), the theoretical speed of sound in an ideal gas at 24 °C is calculated to be $c_T = 345 \text{ m s}^{-1}$. The measured value of c is thus within 3% of the theoretical prediction.

Conclusion

In this paper, a method for measuring the speed of sound in air has been demonstrated using a smartphone and a cardboard tube. The experiment is very economical, and can thus be performed in places with limited access to laboratory equipment.

The experiment requires some amount of programming. It can provide learning opportunities for students both in experimental methods as well as in data analysis.

ORCID iDs

Simen Hellesund  <https://orcid.org/0000-0002-4424-4643>

Received 22 January 2019, in final form 14 February 2019
Accepted for publication 8 March 2019
<https://doi.org/10.1088/1361-6552/ab0e21>

References

- [1] Kuhn J, Vogt P and Hirth M 2014 Analyzing the acoustic beat with mobile devices *Phys. Teach.* **52** 248–9
- [2] Kuhn J and Vogt P 2013 Analyzing acoustic phenomena with a smartphone microphone *Phys. Teach.* **51** 118–9
- [3] Parolin S O and Pezzi G 2013 Smartphone-aided measurements of the speed of sound in different gaseous mixtures *Phys. Teach.* **51** 508–9
- [4] Klein P *et al* 2014 Classical experiments revisited: smartphones and tablet PCs as experimental tools in acoustics and optics *Phys. Educ.* **49** 412
- [5] Gómez-Tejedor J A *et al* 2014 The acoustic doppler effect applied to the study of linear motions *Eur. J. Phys.* **35** 025006
- [6] Kundt A 1866 Ueber eine neue art akustischer staubfiguren und über die anwendung derselben zur bestimmung der schallgeschwindigkeit in festen körpern und gasen *Ann. Phys.* **203** 497–523
- [7] Parolin S and Pezzi G 2015 Kundt's tube experiment using smartphones *Phys. Educ.* **50** 443–7
- [8] Yavuz A 2015 Measuring the speed of sound in air using smartphone applications *Phys. Educ.* **50** 281
- [9] Levine H and Schwinger J 1948 On the radiation of sound from an unflanged circular pipe *Phys. Rev.* **73** 383–406
- [10] Function generator <https://play.google.com> (Accessed: 27 November 2018)
- [11] Smart recorder <https://play.google.com> (Accessed: 27 November 2018)
- [12] Voice recorder lite: HD audio recording & playback <https://itunes.apple.com> (Accessed: 05 February 2019)
- [13] Audio function generator <https://itunes.apple.com> (Accessed: 05 February 2019)
- [14] Boashash B 1992 Estimating and interpreting the instantaneous frequency of a signal—part 2: algorithms and applications *Proc. IEEE* **80** 540–68
- [15] Jones E *et al* 2001 SciPy: open source scientific tools for Python (Accessed: 28 November 2018)
- [16] Oliphant T 2006 *NumPy: a Guide to NumPy* (Spanish Fork, UT: Trelgol Publishing) (Accessed: 29 November 2018)



Simen Hellesund is a PhD student at the University of Oslo. His main research interest is experimental high energy physics, and he is a member of the ATLAS collaboration at CERN. He also enjoys teaching physics and is researching ways to improve physics education. In particular, he

is interested in how to perform physics experiments or measurements for educational purposes, with limited access to laboratory equipment.

References

- [1] J. J. Aubert et al. *Experimental Observation of a Heavy Particle J.* Phys. Rev. Lett., 33:1404–1406, Dec 1974. doi: 10.1103/PhysRevLett.33.1404.
- [2] J. E. Augustin et al. *Discovery of a Narrow Resonance in e^+e^- Annihilation.* Phys. Rev. Lett., 33:1406–1408, Dec 1974. doi: 10.1103/PhysRevLett.33.1406.
- [3] S. W. Herb et al. *Observation of a Dimuon Resonance at 9.5 GeV in 400-GeV Proton-Nucleus Collisions.* Phys. Rev. Lett., 39:252–255, Aug 1977. doi: 10.1103/PhysRevLett.39.252.
- [4] G. Arnison et al. *Experimental Observation of Lepton Pairs of Invariant Mass Around 95 GeV/c² at the CERN SPS Collider.* Phys. Lett. B, 126(5):398–410, 1983. doi: [https://doi.org/10.1016/0370-2693\(83\)90188-0](https://doi.org/10.1016/0370-2693(83)90188-0).
- [5] P. Bagnaia et al. *Evidence for $Z^0 \rightarrow e^+e^-$ at the CERN pp Collider.* Phys. Lett. B, 129(1):130–140, 1983. doi: [https://doi.org/10.1016/0370-2693\(83\)90744-X](https://doi.org/10.1016/0370-2693(83)90744-X).
- [6] The ATLAS Collaboration. *Search for High-Mass Dilepton Resonances using 139 fb⁻¹ of pp Collision Data Collected at $\sqrt{s} = 13$ TeV With the ATLAS Detector.* Phys. Lett. B, 796:68–87, 2019. doi: 10.1016/j.physletb.2019.07.016.
- [7] The ATLAS Collaboration. *Search for New Non-Resonant Phenomena in High-Mass Dilepton Final States With the ATLAS Detector.* JHEP, 2020(11):5, Nov 2020. doi: 10.1007/JHEP11(2020)005.
- [8] N. Arkani-Hamed, S. Dimopoulos and G. Dvali. *The Hierarchy Problem and New Dimensions at a Millimeter.* Phys. Lett. B, 429(3-4):263–272, Jun 1998. doi: 10.1016/s0370-2693(98)00466-3.
- [9] The ATLAS Collaboration. *Search for Non-Resonant Extra Dimensions in High-Mass Dilepton Final States at $\sqrt{s} = 13$ TeV Using 139 fb⁻¹.* Technical Report ATL-PHYS-PUB-2021-021, CERN, Geneva, Jun 2021. URL <https://cds.cern.ch/record/2772628>.
- [10] The ATLAS Collaboration. *Search for New High-Mass Phenomena in the Dilepton Final State Using 36 fb⁻¹ of Proton-Proton Collision Data at $\sqrt{s} = 13$ TeV With the ATLAS Detector.* JHEP, 2017(10):182, Oct 2017. doi: 10.1007/JHEP10(2017)182.
- [11] P. Zyla et al. *Review of Particle Physics.* PTEP, 2020(8), 2020. doi: 10.1093/ptep/ptaa104.
- [12] V. E. Barnes et al. *Observation of a Hyperon With Strangeness Minus Three.* Phys. Rev. Lett., 12:204–206, Feb 1964. doi: 10.1103/PhysRevLett.12.204.
- [13] L. Di Lella and C. Rubbia. *The Discovery of the W and Z Particles.* Adv. Ser. Dir. High Energy Phys., 23:137–163, 2015. doi: 10.1142/9789814644150_0006.

- [14] F. Mandl and G. Shaw. *Quantum Field Theory*. Wiley, 2010.
- [15] E. Fermi. *Tentativo di una Teoria Dei Raggi*. Il Nuovo Cimento (1924-1942), 11, 1934. doi: 10.1007/BF02959820.
- [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*. Phys. Lett. B, 716:1–29, 2012. doi: 10.1016/j.physletb.2012.08.020.
- [17] The CMS Collaboration. *Observation of a New Boson at a Mass of 125 GeV With the CMS Experiment at the LHC*. Phys. Lett. B, 716:30–61, 2012. doi: 10.1016/j.physletb.2012.08.021.
- [18] M. E. Peskin and D. V. Schroeder. *An Introduction to Quantum Field Theory*. Addison-Wesley, 1995.
- [19] D. Hanneke, S. Fogwell and G. Gabrielse. *New Measurement of the Electron Magnetic Moment and the Fine Structure Constant*. Phys. Rev. Lett., 100(12), Mar 2008. doi: 10.1103/physrevlett.100.120801.
- [20] S. L. Glashow. *Partial-Symmetries of Weak Interactions*. Nucl. Phys., 22(4):579–588, 1961. doi: [https://doi.org/10.1016/0029-5582\(61\)90469-2](https://doi.org/10.1016/0029-5582(61)90469-2).
- [21] S. Weinberg. *A Model of Leptons*. Phys. Rev. Lett., 19:1264–1266, Nov 1967. doi: 10.1103/PhysRevLett.19.1264.
- [22] A. Salam. *Weak and Electromagnetic Interactions*. Conf. Proc. C, 680519:367–377, 1968. doi: 10.1142/9789812795915_0034.
- [23] L. Álvarez Gaumé and J. Ellis. *Eyes On a Prize Particle*. Nat. Phys., 7(1):2–3, 2011. doi: 10.1038/nphys1874.
- [24] J. Goldstone, A. Salam and S. Weinberg. *Broken Symmetries*. Phys. Rev., 127: 965–970, Aug 1962. doi: 10.1103/PhysRev.127.965.
- [25] The LHCb Collaboration. *Test of Lepton Universality in Beauty-Quark Decays*, 03 2021. arXiv: 2103.11769.
- [26] E. M. Henley and A. Garcia. *Subatomic Physics*. World Scientific Publishing Company, 3rd edition, Jul 2007.
- [27] A. Riess. *Dark Matter*, Dec 2017. URL <https://www.britannica.com/science/dark-matter>. Accessed: 26.10.2021.
- [28] V. C. Rubin and W. K. Ford Jr. *Rotation of the Andromeda Nebula from a Spectroscopic Survey of Emission Regions*. Astrophys. J., 159, Feb 1970. doi: 10.1086/150317.
- [29] I. Gogoladze, B. He and Q. Shafi. *New fermions at the LHC and Mass of the Higgs Boson*. Phys. Lett. B, 690(5):495–500, 2010. doi: <https://doi.org/10.1016/j.physletb.2010.05.076>.
- [30] P. Langacker. *The Physics of Heavy Z' Gauge Bosons*. Rev. Mod. Phys., 81(3): 11991228, Aug 2009. doi: 10.1103/revmodphys.81.1199.

-
- [31] D. London and J. L. Rosner. *Extra Gauge Bosons in E_6* . Phys. Rev. D, 34:1530, 1986. doi: 10.1103/PhysRevD.34.1530.
- [32] D. Pappadopulo, A. Thamm, R. Torre and A. Wulzer. *Heavy Vector Triplets: Bridging Theory and Data*. JHEP, 2014(9), Sep 2014. doi: 10.1007/jhep09(2014)060.
- [33] A. Djouadi. *The Anatomy of Electro-Weak Symmetry Breaking. II: The Higgs bosons in the Minimal Supersymmetric Model*. Phys. Rept., 459:1–241, 2008. doi: 10.1016/j.physrep.2007.10.005.
- [34] C. Hays, A. Kotwal and O. Stelzer-Chilton. *New Techniques in the Search for Z Bosons and Other Neutral Resonances*. Mod. Phys. Lett. A, 24(30):23872403, Sep 2009. doi: 10.1142/s021773230903179x.
- [35] L. Randall and R. Sundrum. *A Large Mass Hierarchy from a Small Extra Dimension*. Phys. Rev. Lett., 83:3370–3373, 1999. doi: 10.1103/PhysRevLett.83.3370.
- [36] A. Shivaji, V. Ravindran and P. Agrawal. *Production of a KK -Graviton and a Vector Boson in ADD Model via Gluon Fusion*. JHEP, 2012(2), Feb 2012. doi: 10.1007/jhep02(2012)057.
- [37] T. Gleisberg et al. *Helicity Formalism for Spin-2 Particles*. JHEP, 2003(09):1, Sep 2003. doi: 10.1088/1126-6708/2003/09/001.
- [38] G. F. Giudice, R. Rattazzi and J. D. Wells. *Quantum Gravity and Extra Dimensions at High-Energy Colliders*. Nucl. Phys. B, 544(1-2):3–38, Apr 1999. doi: 10.1016/s0550-3213(99)00044-9.
- [39] J. L. Hewett. *Indirect Collider Signals for Extra Dimensions*. Phys. Rev. Lett., 82(24):4765–4768, Jun 1999. doi: 10.1103/physrevlett.82.4765.
- [40] T. Han, J. D. Lykken and R.-J. Zhang. *Kaluza-Klein States from Large Extra Dimensions*. Phys. Rev. D, 59(10), Mar 1999. doi: 10.1103/physrevd.59.105006.
- [41] A. V. Manohar. *Introduction to Effective Field Theories*. In Proceedings of Les Houches Summer School on Effective Field Theories, 2017. arXiv: 1804.05863.
- [42] S. Willenbrock and C. Zhang. *Effective Field Theory Beyond the Standard Model*. Annu. Rev. Nucl. Part. Sci., 64(1):83100, Oct 2014. doi: 10.1146/annurev-nucl-102313-025623.
- [43] E. J. Eichten, K. D. Lane and M. E. Peskin. *New Tests for Quark and Lepton Substructure*. Phys. Rev. Lett., 50:811–814, Mar 1983. doi: 10.1103/PhysRevLett.50.811.
- [44] B. Grzadkowski, M. Iskrzyski, M. Misiak and J. Rosiek. *Dimension-Six Terms in the Standard Model Lagrangian*. JHEP, 2010(10), Oct 2010. doi: 10.1007/jhep10(2010)085.
- [45] E. Eichten, I. Hinchliffe, K. Lane and C. Quigg. *Supercollider Physics*. Rev. Mod. Phys., 56:579–707, Oct 1984. doi: 10.1103/RevModPhys.56.579.
- [46] D. J. Griffiths. *Introduction to Electrodynamics*. Pearson, 3rd edition, 1999.

- [47] K. Schindl. *The Injector Chain for the LHC*. Technical Report CERN-PS-99-018-DI, CERN, Mar 1999. URL <https://cds.cern.ch/record/384396>.
- [48] O. S. Brüning et al. *LHC Design Report*. CERN Yellow Reports: Monographs. CERN, 2004. doi: 10.5170/CERN-2004-003-V-1.
- [49] J. Wenninger. *Operation and Configuration of the LHC in Run 2*. Technical Report CERN-ACC-NOTE-2019-0007, CERN, Mar 2019. URL <https://cds.cern.ch/record/2668326>.
- [50] E. Mobs. *The CERN Accelerator Complex*, Aug 2018. URL <http://cds.cern.ch/record/2636343>.
- [51] The ATLAS Collaboration. *ATLAS Luminosity Public Results Run 2*. URL <https://twiki.cern.ch/twiki/bin/view/AtlasPublic/LuminosityPublicResultsRun2>. Accessed: 03.02.2020.
- [52] The ATLAS Collaboration. *ATLAS Detector and Physics Performance: Technical Design Report, Volume 1*. Technical Design Report ATLAS. CERN, 1999.
- [53] The ATLAS Collaboration. *The ATLAS Experiment at the CERN Large Hadron Collider*. JINST, 3(08), Aug 2008. doi: 10.1088/1748-0221/3/08/s08003.
- [54] J. Pequeno. *Computer Generated Image of the Whole ATLAS Detector*, Mar 2008. URL <https://cds.cern.ch/record/1095924>.
- [55] The ATLAS Collaboration. *ATLAS Pixel Detector Electronics and Sensors*. JINST, 3(07), Jul 2008. doi: 10.1088/1748-0221/3/07/p07007.
- [56] A. La Rosa. *The ATLAS Insertable B-Layer: from Construction to Operation*. JINST, 11(12), Dec 2016. doi: 10.1088/1748-0221/11/12/c12036.
- [57] J. R. Pater. *The ATLAS Semiconductor Tracker Operation and Performance*. JINST, 7(04), Apr 2012. doi: 10.1088/1748-0221/7/04/c04001.
- [58] H. Kolanoski and N. Wermes. *Particle Detectors: Fundamentals and Applications*. Oxford University Press, 2020.
- [59] The ATLAS Collaboration. *Performance of the ATLAS Transition Radiation Tracker in Run 1 of the LHC: Tracker Properties*. JINST, 12, 05 2017. doi: 10.1088/1748-0221/12/05/P05002.
- [60] The ATLAS Collaboration. *Track Reconstruction Performance of the ATLAS Inner Detector at $\sqrt{s} = 13$ TeV*. Technical Report ATL-PHYS-PUB-2015-018, CERN, Geneva, Jul 2015. URL <http://cds.cern.ch/record/2037683>.
- [61] A. Yamamoto et al. *The ATLAS Central Solenoid*. Nucl. Instrum. Methods Phys. Res. A, 584:53–74, 01 2008. doi: 10.1016/j.nima.2007.09.047.
- [62] The ATLAS Collaboration. *Operation of the ATLAS Trigger System in Run 2*. JINST, 15:P10004. 60 p, Jul 2020. doi: 10.1088/1748-0221/15/10/P10004.
- [63] T. Cornelissen et al. *The New ATLAS Track Reconstruction (NEWT)*. J. Phys. Conf. Ser., 119(3), Jul 2008. doi: 10.1088/1742-6596/119/3/032014.

-
- [64] B. Resende. *Muon Identification Algorithms in ATLAS*. Technical Report ATL-PHYS-PROC-2009-113, CERN, Geneva, Sep 2009. URL <https://cds.cern.ch/record/1209632>.
- [65] The ATLAS Collaboration. *Reconstruction of Primary Vertices at the ATLAS Experiment in Run 1 Proton-Proton Collisions at the LHC*. Eur. Phys. J. C, 77(5), May 2017. doi: 10.1140/epjc/s10052-017-4887-5.
- [66] M. Cacciari, G. P. Salam and G. Soyez. *The Anti- k_t Jet Clustering Algorithm*. JHEP, 2008(04):063063, Apr 2008. doi: 10.1088/1126-6708/2008/04/063.
- [67] The ATLAS Collaboration. *Performance of Missing Transverse Momentum Reconstruction With the ATLAS Detector Using Proton-Proton Collisions at $\sqrt{s}=13$ TeV*. Eur. Phys. J. C, 78(11), Nov 2018. doi: 10.1140/epjc/s10052-018-6288-9.
- [68] C. Kourkoumelis and S. Vourakis. *HYPATIA — An Online Tool for ATLAS Event Visualization*. Phys. Educ., 49:21, 01 2014. doi: 10.1088/0031-9120/49/1/21.
- [69] 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*. Eur. Phys. J. C, 79(8):639, 2019. doi: 10.1140/epjc/s10052-019-7140-6.
- [70] B. Mindur. *ATLAS Transition Radiation Tracker (TRT): Straw Tubes for Tracking and Particle Identification at the Large Hadron Collider*. Technical Report ATL-INDET-PROC-2016-001, CERN, Geneva, Mar 2016. URL <http://cds.cern.ch/record/2139567>.
- [71] The ATLAS Collaboration. *Muon Reconstruction and Identification Efficiency in ATLAS Using the Full Run 2 pp Collision Data Set at $\sqrt{s} = 13$ TeV*. Eur. Phys. J. C, 81(578), 2021.
- [72] S. Agostinelli et al. *Geant4 — A Simulation Toolkit*. Nucl. Instrum. Methods Phys. Res. A, 506(3):250–303, 2003. doi: [https://doi.org/10.1016/S0168-9002\(03\)01368-8](https://doi.org/10.1016/S0168-9002(03)01368-8).
- [73] W. Lukas. *Fast Simulation for ATLAS: Atlfast-II and ISF*. Technical Report ATL-SOFT-PROC-2012-065, CERN, Geneva, Jun 2012. URL <https://cds.cern.ch/record/1458503>.
- [74] S. Alioli, P. Nason, C. Oleari and E. Re. *A General Framework for Implementing NLO Calculations in Shower Monte Carlo Programs: the POWHEG BOX*. JHEP, 06:43, 2010. doi: 10.1007/JHEP06(2010)043.
- [75] H.-L. Lai et al. *New Parton Distributions for Collider Physics*. Phys. Rev. D, 82, Oct 2010. doi: 10.1103/PhysRevD.82.074024.
- [76] C. Anastasiou, L. Dixon, K. Melnikov and F. Petriello. *High-precision QCD at hadron colliders: Electroweak gauge boson rapidity distributions at next-to-next-to leading order*. Phys. Rev. D, 69, May 2004. doi: 10.1103/PhysRevD.69.094008.
- [77] A. Arbuzov et al. *Update of the MCSANC Monte Carlo Integrator, v. 1.20*. JETP Letters, 103(2):131136, Jan 2016. doi: 10.1134/s0021364016020041.

- [78] S. Dulat et al. *New Parton Distribution Functions from a Global Analysis of Quantum Chromodynamics*. Phys. Rev. D, 93, Feb 2016. doi: 10.1103/PhysRevD.93.033006.
- [79] P. Golonka and Z. Was. *PHOTOS Monte Carlo: A Precision Tool for QED Corrections in Z and W Decays*. Eur. Phys. J. C, 45:97–107, 2006. doi: 10.1140/epjc/s2005-02396-4.
- [80] R. D. Ball et al. *Parton Distributions for the LHC Run II*. JHEP, 04:40, 2015. doi: 10.1007/JHEP04(2015)040.
- [81] P. Artoisenet, R. Frederix, O. Mattelaer and R. Rietkerk. *Automatic Spin-Entangled Decays of Heavy Resonances in Monte Carlo Simulations*. JHEP, 2013, 12 2012. doi: 10.1007/JHEP03(2013)015.
- [82] M. Czakon and A. Mitov. *Top++: A Program for the Calculation of the Top-Pair Cross-Section at Hadron Colliders*. CPC, 185(11):2930–2938, 2014. doi: <https://doi.org/10.1016/j.cpc.2014.06.021>.
- [83] T. Gleisberg et al. *Event Generation With SHERPA 1.1*. JHEP, 2009(02):7, Feb 2009. doi: 10.1088/1126-6708/2009/02/007.
- [84] E. W. Varnes. *A Poisson Likelihood Approach to Fake Lepton Estimation With the Matrix Method*, Jun 2016. arXiv: 1606.06817.
- [85] V. Morisbak. *Searching for New Forces of Nature: New Neutral Gauge Bosons in Dimuon Final States With the ATLAS Detector at the LHC*. PhD thesis, University of Oslo, Jan 2021.
- [86] T. Sjöstrand et al. *An Introduction to PYTHIA 8.2*. CPC, 191:159–177, 2015. doi: <https://doi.org/10.1016/j.cpc.2015.01.024>.
- [87] R. D. Ball et al. *Parton Distributions With LHC Data*. Nucl. Phys. B, 867:244–289, 2013. doi: 10.1016/j.nuclphysb.2012.10.003.
- [88] D. Lange. *The EvtGen Particle Decay Simulation Package*. Nucl. Instrum. Methods Phys. Res. A, 462:152–155, 04 2001. doi: 10.1016/S0168-9002(01)00089-4.
- [89] S. Frixione, P. Nason and C. Oleari. *Matching NLO QCD Computations With Parton Shower Simulations: the POWHEG Method*. Journal of High Energy Physics, 2007 (11):70, Nov 2007. doi: 10.1088/1126-6708/2007/11/070.
- [90] T. Sjöstrand, S. Mrenna and P. Skands. *A Brief Introduction to PYTHIA 8.1*. CPC, 178(11):852–867, 2008. doi: <https://doi.org/10.1016/j.cpc.2008.01.036>.
- [91] The ATLAS Collaboration. *ATLAS Pythia 8 Tunes to 7 TeV data*. Technical Report ATL-PHYS-PUB-2014-021, CERN, Geneva, Nov 2014. URL <https://cds.cern.ch/record/1966419>.
- [92] D. Stump et al. *Inclusive Jet Production, Parton Distributions, and the Search for New Physics*. JHEP, 2003(10):46, Oct 2003. doi: 10.1088/1126-6708/2003/10/046.

-
- [93] P. J. Falke. *Search for New Heavy Resonances in the Dilepton Final State and Calibration of the LAr Phase-1 Upgrade Demonstrator at the ATLAS Experiment*. PhD thesis, Université Grenoble Alpes, Oct 2019. URL <https://cds.cern.ch/record/2693468>.
- [94] The ATLAS Collaboration. *Luminosity Determination in pp Collisions at $\sqrt{s} = 8$ TeV Using the ATLAS Detector at the LHC*. Eur. Phys. J. C, 76(12):653, 2016. doi: 10.1140/epjc/s10052-016-4466-1.
- [95] G. Avoni et al. *The New LUCID-2 Detector for Luminosity Measurement and Monitoring in ATLAS*. JINST, 13(07), Jul 2018. doi: 10.1088/1748-0221/13/07/p07017.
- [96] The ATLAS Collaboration. *Muon Reconstruction Performance of the ATLAS Detector in Proton-Proton Collision Data at $\sqrt{s} = 13$ TeV*. Eur. Phys. J. C, 76(5):292, 2016. doi: 10.1140/epjc/s10052-016-4120-y.
- [97] The ATLAS Collaboration. *Electron and Photon Energy Calibration With the ATLAS Detector Using 2015-2016 LHC Proton-Proton Collision Data*. JINST, 14(03), 2019. doi: 10.1088/1748-0221/14/03/P03017.
- [98] The ATLAS Collaboration. *Search for New Phenomena in Dijet Events Using 37 fb $^{-1}$ of pp Collision Data Collected at $\sqrt{s} = 13$ TeV With the ATLAS Detector*. Phys. Rev. D, 96, Sep 2017. doi: 10.1103/PhysRevD.96.052004.
- [99] The CMS Collaboration. *Search for Diphoton Resonances in the Mass Range from 150 to 850 GeV in pp Collisions at $\sqrt{s} = 8$ TeV*. Phys. Lett. B, 750:494–519, 2015. doi: <https://doi.org/10.1016/j.physletb.2015.09.062>.
- [100] The ATLAS Collaboration. *Search for Resonances in Diphoton Events at $\sqrt{s} = 13$ TeV With the ATLAS Detector*. JHEP, 09:1, 2016. doi: 10.1007/JHEP09(2016)001.
- [101] The ATLAS Collaboration. *Search for New Phenomena in High-Mass Diphoton Final States Using 37 fb $^{-1}$ of Proton-Proton Collisions Collected at $\sqrt{s} = 13$ TeV With the ATLAS Detector*. Phys. Lett. B, 775:105–125, 2017. doi: 10.1016/j.physletb.2017.10.039.
- [102] W. Verkerke and D. P. Kirkby. *The RooFit Toolkit for Data Modeling*. In Proceedings of CHEP, 2003. arXiv: physics/0306116.
- [103] L. Moneta et al. *The RooStats Project*. In Proceedings of PoS(ACAT), 2010. doi: 10.22323/1.093.0057. arXiv: 1009.1003.
- [104] O. Behnke, K. Kroninger, G. Schott and T. Schorner-Sadenius, editors. *Data Analysis in High Energy Physics: A Practical Guide to Statistical Methods*. Wiley-VCH, 1st edition, 2013.
- [105] The ATLAS Collaboration, The CMS Collaboration, The LHC Higgs Combination Group. *Procedure for the LHC Higgs boson search combination in Summer 2011*. Technical Report CMS-NOTE-2011-005, ATL-PHYS-PUB-2011-11, CERN, Geneva, Aug 2011. URL <https://cds.cern.ch/record/1379837>.

- [106] S. S. Wilks. *The Large-Sample Distribution of the Likelihood Ratio for Testing Composite Hypotheses*. Ann. Math. Stat., 9(1):60–62, 1938. doi: 10.1214/aoms/1177732360.
- [107] G. Cowan, K. Cranmer, E. Gross and O. Vitells. *Asymptotic formulae for likelihood-based tests of new physics*. Eur. Phys. J. C, 71(2), Feb 2011. doi: 10.1140/epjc/s10052-011-1554-0.
- [108] E. Gross and O. Vitells. *Trial factors for the look elsewhere effect in high energy physics*. Eur. Phys. J. C, 70:525–530, 2010. doi: 10.1140/epjc/s10052-010-1470-8.
- [109] A. L. Read. *Presentation of search results: the CL_s technique*. J. Phys. G: Nucl. Part. Phys., 28(10):2693–2704, Sep 2002. doi: 10.1088/0954-3899/28/10/313.
- [110] ATLAS Statistics Forum. *Frequentist Limit Recommendation*, 2011. URL https://indico.cern.ch/event/126652/contributions/1343592/attachments/80222/115004/Frequentist_Limit_Recommendation.pdf.
- [111] C. Anastasiou, L. Dixon, K. Melnikov and F. Petriello. *High-Precision QCD at Hadron Colliders: Electroweak Gauge Boson Rapidity Distributions at NNLO*. Phys. Rev. D, 69, May 2004. doi: 10.1103/PhysRevD.69.094008.
- [112] D. Abhayasinghe. *Searches for New Non-Resonant Phenomena in High-Mass Dilepton Final States With the ATLAS Detector at the Large Hadron Collider*. PhD thesis, Royal Holloway, University of London, Jul 2020.
- [113] A. White. *Searches for the Rare Decay of the Higgs Boson to Two Muons and for New Physics in Dilepton Final States With the ATLAS Experiment*. PhD thesis, Horace H. Rackham School of Graduate Studies, University of Michigan, Jun 2021.
- [114] The ATLAS Collaboration. *Search for Contact Interactions and Large Extra Dimensions in the Dilepton Channel Using Proton-Proton Collisions at $\sqrt{s} = 8$ TeV With the ATLAS Detector*. Eur. Phys. J. C, 74(12), Dec 2014. doi: 10.1140/epjc/s10052-014-3134-6.
- [115] The CMS Collaboration. *Search for Resonant and Nonresonant New Phenomena in High-Mass Dilepton Final States at $\sqrt{s} = 13$ TeV*. JHEP, 2021(7), Jul 2021. doi: 10.1007/jhep07(2021)208.
- [116] P. D. Dauncey, M. Kenzie, N. Wardle and G. J. Davies. *Handling Uncertainties in Background Shapes: the Discrete Profiling Method*. JINST, 10(04):P04015P04015, Apr 2015. doi: 10.1088/1748-0221/10/04/p04015.
- [117] R. Edgar, D. Amidei, C. Grud and K. Sekhon. *Functional Decomposition: A New Method for Search and Limit Setting*, 2018. arXiv: 1805.04536.
- [118] M. Frate, K. Cranmer, S. Kalia, A. Vandenberg-Rodes and D. Whiteson. *Modeling Smooth Backgrounds and Generic Localized Signals With Gaussian Processes*, 2017. arXiv: 1709.05681.
- [119] R. Diener, S. Godfrey and T. Martin. *Using Final State Pseudorapidities to Improve s -Channel Resonance Observables at the LHC*. Phys. Rev. D, 80, 09 2009. doi: 10.1103/PhysRevD.80.075014.

-
- [120] V. Garonne et al. on behalf of the ATLAS Collaboration. *Rucio — The Next Generation of Large Scale Distributed System for ATLAS Data Management*. J. Phys. Conf. Ser., 513(4), 2014.
- [121] R. B. Cleveland, W. S. Cleveland, J. E. McRae and I. Terpenning. *STL: A Seasonal-Trend Decomposition Procedure Based on Loess*. J. Off. Stat., 6:3–33, 01 1990.
- [122] P. Tiunov. *Time Series Anomaly Detection Algorithms*, 08.06.2017. URL <https://medium.com/cube-dev/time-series-anomaly-detection-algorithms-1cef5519aef2>. Accessed: 18.11.2021.
- [123] T. M. Mitchell. *Machine Learning*. McGraw-Hill Inc., Mar 1997.
- [124] I. Goodfellow, Y. Bengio and A. Courville. *Deep Learning*. MIT Press, 2016.
- [125] Y. Freund and R. E. Schapire. *A Decision-Theoretic Generalization of On-Line Learning and an Application to Boosting*. JCSS, 55(1):119–139, 1997. doi: <https://doi.org/10.1006/jcss.1997.1504>.
- [126] R. E. Schapire. *Explaining AdaBoost*. In B. Schölkopf, Z. Luo and V. Vovk, editors, Empirical inference. Festschrift in honor of Vladimir N. Vapnik, chapter 5, pages 37–52. Springer, 2013. doi: 10.1007/978-3-642-41136-6_5.
- [127] F. Pedregosa et al. *Scikit-Learn: Machine Learning in Python*. JMLR, 12:2825–2830, 2011.
- [128] Wikipedia contributors. *Receiver Operating Characteristic — Wikipedia, The Free Encyclopedia*, 2019. URL https://en.wikipedia.org/wiki/Receiver_operating_characteristic. [Accessed 05.09.2019].
- [129] S. Hellesund. *Måling av tyngdeakselerasjonen g ved hjelp av Dopplerskift*. Fra fysikkens verden, 4:108–109, 2018.
- [130] S. M. Torres and W. J. González-Espada. *Calculating g from Acoustic Doppler Data*. Phys. Teach., 44(8):536–539, 2006. doi: 10.1119/1.2362950.
- [131] B. Boashash. *Estimating and Interpreting the Instantaneous Frequency of a Signal — Part 2: Algorithms and Applications*. Proc. IEEE, 80(4):540–568, 1992.
- [132] S. Hellesund. *Measuring the Speed of Sound in Air Using a Smartphone and a Cardboard Tube*. Phys. Educ., 54(3), Apr 2019. doi: 10.1088/1361-6552/ab0e21.
- [133] S. Hellesund. *Measuring Earth’s Magnetic Field Using a Smartphone Magnetometer*, 2018. arXiv: 1901.00857.

Acronyms and Abbreviations

- AdaBoost** Adaptive boosting
- ADC** ATLAS Distributed Computing
- ADD** Arkani-Hamed, Dimopoulos, and Dvali
- ANN** Artificial neural network
- ATLAS** A Toroidal LHC ApparatuS
- AUC** Area under curve
- BDT** Boosted decision tree
- BEH** Brout-Englert-Higgs
- BSM** Beyond standard model
- CERN** The European Organization for Nuclear Research
- CI** Contact interaction
- CMS** Compact Muon Solenoid
- COM** Centre of mass
- CR** Control region
- CSC** Cathode strip chamber
- DDM** Distributed data management
- DY** Drell-Yan
- EFT** Effective field theory
- EWSB** Electroweak symmetry breaking
- GGUS** Global Grid User Support
- Grid** Worldwide LHC Computing Grid
- GRW** Giudice, Rattazzi, and Wells
- GUT** Grand unification theory
- HEPP** High energy particle physics

HL-HLC High-Luminosity Large Hadron Collider

HLZ Han, Lykken, and Zhang

HVT Heavy vector triplet

IBL Insertable B-Layer

KK Kaluza-Klein

LEAR Low Energy Antiproton Ring

LEP The Large Electron Positron Collider

LHC Large Hadron Collider

LO Leading order

MC Monte Carlo

MDT Monitored drift tube

ML Machine learning

MS Muon Spectrometer

MSSM Minimal supersymmetric standard model

NLO Next-to-leading order

NNLO Next-to-next-to-leading order

PDF Parton density function

PDF Probability density function

pp Proton-proton

PS Proton Synchrotron

PSB Proton Synchrotron Booster

QCD Quantum chromodynamics

QED Quantum electrodynamics

QFT Quantum field theory

ROC Receiver operating characteristic

RPC Resistive plate chamber

RS1 Randall-Sundrum 1

- SCT** SemiConductor Tracker
- SM** Standard Model
- Sp \bar{p} S** Super Proton-Antiproton Synchrotron
- SPS** Super Proton Synchrotron
- SR** Signal region
- SSB** Spontaneous symmetry breaking
- SSM** Sequential Standard Model
- SVM** Support vector machine
- SWiFt** Sliding window fit
- TF** Transfer function
- TGC** Thin gap chamber
- TRT** Transition Radiation Tracker