

Searching for New Physics in Decays of the Higgs Boson with the ATLAS Detector

Elliot Thomas James Reynolds

*Thesis submitted for the degree of
Doctor of Philosophy*



Particle Physics Group,
School of Physics and Astronomy,
University of Birmingham

September 2019

Abstract

Three searches sensitive to physics beyond the Standard Model in Higgs boson decays are presented, using pp collision data collected at $\sqrt{s} = 13$ TeV by the ATLAS experiment during Run 2 of the LHC.

First, a direct search is performed for Higgs boson decays to charm quarks. An upper limit is set at $110\times$ the SM expectation on $\sigma(ZH)\times\text{BR}(H\rightarrow cc)$. The expected sensitivity is estimated for the ATLAS detector at the HL-LHC, assuming 3000 fb^{-1} of pp collision data at $\sqrt{s} = 14$ TeV. An upper limit of 6.3 times the Standard Model expectation for $\sigma(ZH)\times\text{BR}(H\rightarrow cc)$ is expected, in the absence of systematic uncertainties.

Second, a search is performed for Higgs boson decays to pairs of beyond the Standard Model resonances, in the four-muon final state. No events are observed, in agreement with the background-only expectation of 0.4 ± 0.1 events. Mass-dependent upper limits are set on two benchmark models, and model independent upper limits are set by defining a fiducial acceptance.

Third, a search is performed for Higgs boson decays to a Z boson and a light resonance in two lepton plus jet events. In the absence of a signal, upper limits are set on $\sigma(H)\times\text{BR}(H\rightarrow Za^0)/\sigma_{\text{SM}}(H)$, with values starting from 44.8%.

*This thesis is dedicated to my parents, Simon Michael Reynolds and Ella Jane Reynolds,
and to my grandad, Rex Turner.*

Acknowledgements

Over the last four years I have: been a member of a 3000 person scientific collaboration; been co-author to over 250 scientific papers; worked on 3 physics analyses and many more studies; given 14 talks at conferences, workshops and a seminar; lived abroad for a year and a half; and hopefully will also have achieved a PhD at the end of it! Needless to say, this would not have been possible without the many amazing people who were there for me before and throughout. If I were to write adequate thanks for everyone who has supported me over the course of this PhD, I would almost certainly miss its deadline! So I must apologise in advance for the brevity.

I am grateful to the University of Birmingham for hosting me during these investigations, and to STFC and the ERC for funding them. I would also like to give thanks to the ATLAS Collaboration, the LHC, and CERN for making all of this possible. As cliché as it sounds, to work on a CERN experiment was and is nothing less than a dream come true.

The first person I would like to thank is my supervisor, Kostas Nikolopoulos. I asked to work with Kostas soon after chatting with him at the PhD interview, and it was undoubtedly the best decision of my PhD. He has not just supported my research and offered thesis corrections, but has gone above and beyond as a supervisor, helped me develop as a scientist, and taught me many important lessons, both in physics, and in life (even if I did not always want to learn them). Many of these lessons will be remembered for years to come.

The next most obvious person to thank is Andy Chisholm, who is not just a brilliant scientist, but also one of the nicest people I have ever had the good fortune to work with. Andy invariably offered me wise advice when we were out at CERN, and continued doing so up until the end of my PhD, despite my scepticism on more than one occasion. For his collaboration on all three searches herein, and his seemingly everlasting patience, his contribution to this work can not be understated.

I would also like to mention Paul Thompson, Rhys, Phil Allport, and all the other Birmingham staff who have answered questions and offered suggestions throughout this work. I am also grateful to the members of the ATLAS Collaboration who I have worked with over the last four years. With special thanks to Christos, who is not just an e/γ expert, but also eagerly offers advice any time of the day or night! I would also like to thank Dan Tovey and Nigel Watson, for their careful reading of this thesis, their many helpful suggestions, and for their roles in its examination, for which I also thank Pete Watkins.

I have also been very lucky to be at Birmingham and at CERN with many friendly, helpful students, without whom these four years would not have been the same: Matt, Mark, James Broughton, Tim, Alasdair, Andy Foster, James Kendrick, Daniel Briglin, Jack, Russell, Rob Vallance, Daniel Lewis, Gov, Nandish, Patrick Freeman, Patrick Knights, Cameron, Alexis, Rob Ward, and Antonio. Of course I could not forget Kristian, whose company I was lucky to have on many great trips, especially the Fermilab school trip, which was not just one of the best trips of my life, but was undoubtedly the funniest. I was also incredibly lucky to spend one and a half years living with Georgios

out at CERN. While annoyingly, he did get better than me at bouldering, he more than made up for it with the many midnight physics conversations we shared! I wish everyone the best for the next stages of their journeys.

Of course, my time at Birmingham would not have been the same without all of the people and experiences who lead up to it. Of the many people I have to thank for this, a natural place to start is with my friends, especially Ryan and Aleks, two of the best friends I could hope for. As my undergraduate lab partner, and an amazing friend since, Ryan has offered reels of great advice before and throughout this PhD. I also have to thank Dave Edwards and Mike Braithewaite, for the formative roles they played in my education. They really were two of the best teachers a student could hope for. I would also like to thank two of my oldest friends, Dave and Nathan, for their friendship before and throughout this time.

I have also been extremely lucky to share this time, and many great adventures, with my wonderful partner Marbri. I can not imagine these last four years without her. I want her to know that even when I was rushing for a deadline, her patience, support and love were truly cherished.

I would like to thank Lisa and Serena. Without Lisa being there for my dad while he was ill, this PhD may never have got off the ground. I should also thank my brother and sister, Callum Reynolds and Jordan Reynolds, for checking that the Overview for Non-Physicists was appropriate for the intended audience. Last but not least, thank you so much to thank my mum and dad, brother and sister again, grandparents, and the rest of my family for their unwavering love and support, not just during this PhD, but long before it as well. This thesis would never have come to be, and I would not be who I am today, without it.

Since the start of my PhD, I have been saying that I want to continue on in particle physics research. I am looking forward to taking the next step in this journey, alongside the many brilliant people here at Birmingham.

Overview for Non-Physicists

About 100 meters under the border of France and Switzerland, and with a circumference of approximately 17 miles, the LHC is the largest particle collider ever built. It accelerates two counter-rotating beams of protons to about the momentum of a freight train, and then collides them with each other in order to study how they interact at these high energies. Because $E = mc^2$ allows energy and mass to be converted into each other, the large energy in these proton collisions can be converted into mass, producing heavy particles such as the Higgs boson. These particles make up everything in the universe, and everything that happens, through the four fundamental forces: electromagnetism, the weak nuclear force, the strong nuclear force, and gravity (which we are yet to establish is the result of a particle). The current best theory as to how the world works is the Standard Model of particle physics, which is illustrated in Figure 1. The study of these particles may reveal the nature of the universe on the most fundamental level.




Figure 1: Particles of the Standard Model of particle physics [1]. The quarks and leptons are the building blocks of the universe. The gauge boson carry the four fundamental forces: the gluon carries the strong nuclear force; the photon carries electromagnetism; and the W and Z bosons carry the weak nuclear force. The Higgs boson allows the other particles to have mass.

However, these particles are literally the smallest things in existence, and many of the most interesting vanish almost immediately after being created, leaving behind only a very subtle trace. This can mean that detecting them is a great challenge, and thus requires the world's most complex detectors. For this, there is the ATLAS detector, which surrounds the point at which the LHC collides the protons. The ATLAS detector is 46 m long and 25 m in diameter, and uses sophisticated technology to measure the particles which come away from the collision point. During this PhD, I have worked as part of the ATLAS Collaboration, contributing by collecting data, performing studies, developing software, and analysing the data collected with the detector.

One of these particles which I have analysed in detail is the Higgs boson. The Higgs boson is interesting because it is the only fundamental particle which does not have spin, it allows the other particles to have mass, and it helps give rise to the electromagnetic and weak nuclear forces that affect our universe today. It was also discovered only recently in 2012, and so there are still many things we do not know about it. More specifically, I have investigated whether the Higgs boson decays into anything which we do not expect, which may have major implications for our understanding of the universe. I searched for three different types of Higgs boson decay, which are described in the following three paragraphs.

First, I was part of a search for Higgs bosons decaying into pairs of new particles, where these new particles both decay into pairs of particles called muons, which can be thought of as heavy electrons. This decay is illustrated in Figure 2. These new particles could be many things, including lighter Higgs bosons or dark matter (an unknown substance which makes up about a quarter of the universe). This search benefits from the fact that the ATLAS detector is very good at detecting muons, and collisions which produce four of them are very rare according to the Standard Model of particle physics (our current best theory of how the world works), so seeing almost any of these decays at all would be a clear sign of new physics. We did not see any at all, which is what we were expecting. This result constrains new physics theories, guiding theoretical physicists in the theories they develop.




Figure 2: Feynman diagram showing decay of the Higgs boson (H) into a pair of lighter Higgs bosons (a^0), which then each decay to a pair of muons (μ^+ and μ^-).

Second, I was part of a search for Higgs bosons decaying into pairs of charm quarks, where the Higgs boson is produced from a Z boson, which is one of the particles which causes the weak nuclear force. This decay is illustrated in Figure 3. The quarks are a kind of particle which come in three generations, where the first generation is the lightest, and the third generation is the heaviest, as shown in Figure 1. The Higgs boson is expected to decay most often to heavy particles, so the decay of the Higgs boson to the third generation has been measured by other people. However, the charm quark is the heaviest second generation particle, so measuring this decay would give us the first information about how the Higgs boson decays to second generation quarks. This allows us to test a crucial prediction of the Standard Model: that the decays of the Higgs boson to the

second generation quarks continue to depend on their masses as we expect them to. We placed what, at the time they were published, were the world's tightest direct constraints on how often this occurs. These constraints were compatible with the predictions of the Standard Model.




Figure 3: Feynman diagram showing a Higgs boson (H) being produced from a Z boson (Z), and subsequently decaying into a charm quark (c) and an anti-charm quark ($\bar{c}$), while the Z boson decays into a pair of electrons (e^+ and e^-) or a pair of muons (μ^+ and μ^-).

Third, I was part of a search for Higgs bosons decaying into a Z boson and either a new particle or a particle made of two charm quarks, which decays via the strong nuclear force. This decay is illustrated in Figure 4. As in the first search, the new particle could be many things, but as a baseline we considered a lighter Higgs boson. We also looked for a particle made of two charm quarks, because it gave us another way to test how often the Higgs boson decays to second generation quarks. This analysis used a machine learning technique called an artificial neural network to search for the presence of these particles in the detector. We saw the amount of data predicted by the Standard Model, further constraining what new physics theories are possible.




Figure 4: Diagram showing a decay of the Higgs boson (H) into a Z boson (Z) and a lighter Higgs boson (a^0). The Z boson then decays into a pair of electrons (e^+ and e^-) or a pair of muons (μ^+ and μ^-), while the lighter Higgs boson decays via the strong nuclear force into a jet of lighter particles.

Lastly, by around 2026 the LHC and ATLAS detector will have received upgrades which allows ATLAS to take data 10 times faster. After this upgrade, the collider will be

known as the High-Luminosity LHC. Towards this effort, I performed a study into how well we will be able to detect electrons with the upgraded detector, which is important because the detector will have more activity from the increased rate of data taking, and will have lost one of its subdetectors which we currently use for detecting electrons to make room for a different subdetector. I also contributed to the testing of prototype sensors for one of the new subdetectors, which is designed to track the paths of charged particles. Finally, I predicted that the search for Higgs bosons decaying into pairs of charm quarks, which I mentioned above, will benefit significantly from the extra data we expect to collect.

Author's Contribution

As with any large collaborative endeavour, achievements of the ATLAS experiment are the result of the efforts of thousands of people. Most active members of the ATLAS Collaboration are on the author list of any ATLAS paper, for example see Ref. [2]. This list of people have qualified for authorship with ATLAS by performing a qualification task for the benefit of the experiment. Chapter 3 describes my qualification task. While throughout this thesis emphasis is placed on the work to which I contributed to directly, the entirety of each analysis is presented for context and clarity, including contributions from collaborators. As such, in the remaining paragraphs of this section, I will detail my contribution to each of the projects discussed. Also provided are links to any publications arising from this work, from which much of this material is drawn, including figures and tables from the paper.

Chapters 1 and 2 cover background material necessary to understand what follows, and is largely not the work of this author, with three main exceptions. First, as described in §2.2.7, this author made contributions to the ATLAS software framework, working as part of the ATLAS $e^\pm/\gamma$ reconstruction software team. Second, as described in the last paragraph of §2.2.8, this author contributed to recording data with the ATLAS detector, by performing Data Quality shifts in the ATLAS Control Room. Thirdly, the work described in §2.4.4 was conducted as part of the ATLAS HL-LHC Inner Tracker effort of the University of Birmingham. This work was essential to the upgrade of the inner tracker of the ATLAS detector, as it was one of the last chances to test the prototype sensors before the full sensor production begins. This author's contributions included irradiating and annealing sensors, data taking, the development of methods to measure the peak charge collection of the sensors, and the analysis of the collected data.

Chapter 3 is the work of this author, which benefited greatly from the support of Kostas Nikolopoulos and Christos Anastopoulos, the technical supervisor for this qualification task. This work resulted in contributions to multiple publications, these are Chapter 4.4.1 of Ref. [3] and Chapter 3 of Ref. [4].

Chapter 5 details the low and high mass parts of a search for Higgs boson decays into new light resonances, in the four muon final state. These were published, along with a search in which one of the intermediate resonances is replaced with the Standard Model Z boson, in Refs [5, 6]. This author made significant direct contributions to the low mass search alongside Kostas Nikolopoulos and Andrew Chisholm, who together performed most of the work necessary for this to be brought to publication. The major exception being the heavy flavour background estimate, for which this author offered cross checks, but the published works were conducted by Huacheng Cai, Roger Caminal Armadans and Verena Ingrid Martinez Outschoorn. The contributions of this author to the high mass component of this search are limited to what the analysts working directly on that search chose to adopt from the methods of the low mass part.

Chapter 4 discusses the first search for Higgs boson decays to charm quarks in ATLAS, which is published in Ref. [7]. This author contributed to several aspects of this analysis, including the generator-level simulation-based background modelling uncertainty estimation, the statistical analysis, and the ZV -based analysis validation. Section 4.11 describes

a study into the prospects for this analysis at the HL-LHC, which was conducted by this author, and published in Ref. [8]. This work was all performed with close collaboration with Andrew Chisholm and Kostas Nikolopoulos.

Finally, Chapter 6 results entirely from work performed by this author, in close collaboration from Andrew Chisholm and Kostas Nikolopoulos. At the time of writing, this analysis was unblinded, and in the final stages of approval by the ATLAS Collaboration.

Contents

Abbreviations

1	The Standard Model and Beyond	1
1.1	Principles of Quantum Field Theory	2
1.1.1	The Quantum Field Theory Framework	3
1.1.2	Feynman Diagrams	5
1.1.3	Particles	6
1.2	The Standard Model of Particle Physics	7
1.2.1	The Standard Model Fermions	7
1.2.2	The Electroweak Sector	9
1.2.3	Quantum Chromodynamics	10
1.2.4	The BEH Mechanism	10
1.3	Properties and Phenomenology of the SM Higgs Boson	15
1.3.1	Higgs Boson Production Mechanisms at the LHC	16
1.3.2	Higgs Boson Decay Channels	19
1.4	Extending the Higgs Sector	20
1.4.1	The Two Higgs Doublet Model	20
1.4.2	The Alignment and Decoupling Limits	23
1.4.3	The Two Higgs Doublet Model with an Additional Singlet	24
1.4.4	Previous Constraints on Light Higgs Bosons	25
2	The ATLAS Experiment at the Large Hadron Collider	27
2.1	The Large Hadron Collider	27
2.2	The ATLAS Detector	29
2.2.1	Conventions	30
2.2.2	Magnets	31
2.2.3	Inner Tracking Detector	32
2.2.4	Calorimeters	34
2.2.5	Muon Systems	36
2.2.6	Trigger	38
2.2.7	ATLAS Software and Physics Simulation	38
2.2.8	Data Taking	39
2.3	Physics Objects	41
2.3.1	Tracks	41
2.3.2	Interaction Vertices	43
2.3.3	Electrons	44
2.3.4	Electromagnetic Cluster Reconstruction Algorithm	47
2.3.5	Photons	47
2.3.6	Muons	49
2.3.7	Jets	50
2.3.8	Heavy Flavour Hadron Tagging	52

2.3.9	Missing Transverse Momentum	53
2.3.10	τ -Leptons	56
2.3.11	Overlap Removal	56
2.4	High Luminosity Upgrades	57
2.4.1	The High Luminosity LHC	57
2.4.2	ATLAS Upgrades For the HL-LHC	58
2.4.3	Potential HL-LHC Inner Tracker Layouts	60
2.4.4	Silicon Strip Sensor Testing for the Inner Tracker	61
3	Expected Electron and Photon Performance of ATLAS at the HL-LHC	65
3.1	Introduction	65
3.2	Backgrounds	66
3.2.1	Jet Fakes	66
3.2.2	Photon Conversions	67
3.2.3	Heavy Flavour Decays	68
3.3	Method	68
3.3.1	Simulated Samples	68
3.3.2	Discriminant Variables	68
3.3.3	Pre-Identification Requirements	74
3.3.4	Identification Requirements	74
3.3.5	Evaluation of the Menu	75
3.4	Results	77
3.5	Photon Performance at the HL-LHC	82
3.5.1	Introduction	82
3.5.2	Method	84
3.5.3	Results	87
3.6	Conclusion	87
4	Search for Higgs Boson Decays to Charm Quark Pairs	91
4.1	Introduction	91
4.2	Experimental and Simulated Data Samples	93
4.2.1	Generator-Level Jet Flavour Categorisation	94
4.3	Event Selection	95
4.4	Signal and Background Modelling	96
4.5	Systematic Uncertainties	97
4.6	Generator-Level Evaluation of Modelling Uncertainties	97
4.6.1	Background Samples Considered	98
4.6.2	Validation of Generator-Level Modelling	101
4.6.3	Yield Modelling Uncertainties	106
4.6.4	Shape Modelling Uncertainties	108
4.7	Statistical Interpretation	114
4.8	Diboson Validation	117
4.9	Grouped Uncertainty Breakdown	119
4.10	Results	122
4.11	HL-LHC Prospects Study	124
4.11.1	Methods	124
4.11.2	Results	126
4.11.3	Systematic Uncertainties	128
4.11.4	Simultaneous Evaluation of Systematic Uncertainties	129
4.11.5	Future Improvements	133
4.12	Conclusions	135

5	Search for Higgs Boson Decays to Pairs of Light Resonances in the Four-Muon Final State	137
5.1	Introduction	137
5.1.1	Pseudoscalar Higgs Bosons	139
5.1.2	Dark Vector Bosons	140
5.1.3	Kinematic Comparison of Models	140
5.2	Experimental and Simulated Data Samples	143
5.3	Event Selection	144
5.3.1	Triggers	144
5.3.2	Lepton Reconstruction	145
5.3.3	Complete Event Selection	146
5.4	Signal Modelling	148
5.4.1	Signal Shape in Mean Di-Muon Mass	149
5.4.2	Model-Dependent Efficiencies	151
5.4.3	Model-Independent Efficiency	151
5.5	Background Modelling	153
5.5.1	MC Estimate of Low Mass Electroweak Background	153
5.5.2	Data-Driven Estimate of Low Mass Heavy Flavour Background	153
5.5.3	MC Based Cross-Check of Heavy Flavour Background	156
5.5.4	Overall Background Model	161
5.6	Systematic Uncertainties	162
5.6.1	Theory Uncertainties	162
5.6.2	Experimental Uncertainties	162
5.7	Results and Interpretation	164
5.8	Conclusion	168
6	Search for Higgs Boson Decays to a Z Boson and a Light Hadronic Resonance	169
6.1	Introduction	169
6.1.1	Light Higgs Bosons	170
6.1.2	Charmonium States: J/ψ and η_c	171
6.2	Experimental and Simulated Data Samples	171
6.3	Event Selection	174
6.3.1	Event-Level Pre-Selection	175
6.3.2	Track Selection	184
6.3.3	Track-Based Multi-Layer-Perceptron	186
6.3.4	Complete Event Selection	196
6.4	Signal and Background Modelling	196
6.4.1	Signal Modelling	196
6.4.2	Background Modelling	197
6.5	Systematic Uncertainties	207
6.5.1	Modelling Uncertainties	208
6.5.2	Experimental Uncertainties	212
6.6	Statistical Interpretation	215
6.6.1	Statistical Model	215
6.6.2	Asimov Fits	216
6.6.3	Model-Independent Interpretation	220
6.7	Validation Strategy	221
6.8	Results	225
6.9	Conclusion	231
7	Conclusions	233
	Bibliography	237

Abbreviations

2HDM 2 Higgs Doublet Model.

2HDM+S 2 Higgs Doublet Model with an additional singlet.

ALICE A Large Ion Collider Experiment.

ATLAS A Toroidal LHC Apparatus.

BDT boosted decision tree.

BR branching ratio.

BSM beyond the Standard Model.

CL confidence level.

CMS Compact Muon Solenoid.

EM electromagnetic.

EW electroweak.

FCNC flavour changing neutral current.

HL-LHC High-Luminosity Large Hadron Collider.

HLT High-Level Trigger.

IBL Insertable B-Layer.

ITk Inner Tracker.

LAr liquid argon.

LHC Large Hadron Collider.

LHCb Large Hadron Collider Beauty.

MC Monte-Carlo simulated data.

MET missing momentum transverse to the beamline.

MLP multi-layer perceptron.

MS Muon Systems.

MSSM minimal Supersymmetric extension to the Standard Model.

MVA multivariate analysis.

NMSSM next-to-minimal Supersymmetric extension to the Standard Model.

QCD Quantum Chromodynamics.

QFT quantum field theory.

ROC Receiver Operating Characteristic.

SCT Semiconductor Tracker.

SM Standard Model.

SR signal region.

TMVA Toolkit for Multivariate Analysis.

TRT Transition Radiation Tracker.

TTVA track-to-vertex association.

VR validation region.

CHAPTER 1

THE STANDARD MODEL AND BEYOND

The Standard Model (SM) of particle physics is one of the greatest achievements of humanity, describing the fundamental constituents of nature and the interactions between them with astounding accuracy. However, there are major issues with the SM, which can be broadly categorised as unexplained phenomena, and issues of arbitrariness or *fine-tuning*. While a full discussion of these issues is beyond the scope of this thesis, a brief description of some of the most important will serve to illustrate the point.

Unexplained phenomena refers to any aspect of reality which is not explained by the SM. While any other theory can reduce its scope to a subset of reality, as a candidate theory of everything, the SM can not. Some of the most prominent examples are: gravity; baryon asymmetry, affecting $\approx 4.95\%$ of the universe; dark matter, $\approx 27.0\%$ of the universe; and dark energy, $(67.9 \pm 0.013)\%$ of the universe [9]. Another issue with the SM, is that some aspects of it appear arbitrary. For instance, there is no known reason why there are exactly three generation of fermions in the SM. Lastly, some variables in the SM appear fine-tuned, meaning they take values which would lead to a very different universe if they were even slightly different. An example being the mass of the Higgs boson, which as a scalar particle receives quadratic quantum corrections to its mass of

the order of the Plank scale [10]. However, the Higgs boson mass has been observed at the electroweak (EW) scale, requiring around sixteen orders of magnitude of cancellation between the bare (unrenormalised) Higgs boson mass and its quantum corrections. This is referred to as the *Hierarchy Problem*. The level of fine-tuning required is considered *unnatural*, so many new physics models have been proposed to avoid these divergences, once of the best known being *Supersymmetry* [11]. It should be noted that while issues of fine-tuning and arbitrariness are hard for people to accept, they are not a direct internal contradiction of the SM, or a contradiction with any experimental result. So while they motivate many new physics models, they do not guarantee that the SM is incorrect.

Section 1.1 provides a brief overview of quantum field theory (QFT), the language of particle physics. With this groundwork in place, Section 1.2 outlines the SM itself. While a full discussion of the various models which have been proposed to solve the issues outlined above is beyond the scope of this thesis, Section 1.4 covers two example models with extended scalar sectors, which will provide benchmarks for the two searches described in Chapters 5 and 6.

1.1 Principles of Quantum Field Theory

Two of the great scientific breakthroughs of the twentieth century are special relativity and quantum mechanics. Relativistic quantum mechanics combines the underlying principles of these, into a single, elegant framework. However, relativistic quantum mechanics can not describe situations in which the number of particles is changing, which is where QFT becomes necessary.

In the following subsection, the conceptual steps required to understand QFT will be outlined. Subsequently, interactions in QFT, and their connection to Feynman diagrams will be described. Natural units ($\hbar = c = 1$) will be used throughout.

1.1.1 The Quantum Field Theory Framework

Classical mechanics can be promoted to a quantum theory by applying *canonical quantisation* (also known as *first quantisation*), in which the generalised coordinates of the theory and their conjugate momenta are promoted to operators, and the energy is replaced with the Hamiltonian operator. These then satisfy the following definitions in the position representation

$$\mathbf{x} |\phi(\mathbf{x}, t)\rangle = \mathbf{x} |\phi(\mathbf{x}, t)\rangle$$

$$\mathbf{p} |\phi(\mathbf{x}, t)\rangle = -i \frac{\partial |\phi(\mathbf{x}, t)\rangle}{\partial \mathbf{x}},$$

$$\mathcal{H} |\phi(\mathbf{x}, t)\rangle = i \frac{\partial |\phi(\mathbf{x}, t)\rangle}{\partial t}.$$

where $|\phi(\mathbf{x}, t)\rangle$ is a quantum mechanical state.

Using these relations, the classical energy momentum relation translates directly into the time-dependent Schrödinger equation of quantum mechanics, in the position representation

$$E = \frac{\mathbf{p}^2}{2m} \rightarrow i \frac{\partial |\phi(\mathbf{x}, t)\rangle}{\partial t} = -\frac{1}{2m} \frac{\partial^2 |\phi(\mathbf{x}, t)\rangle}{\partial \mathbf{x}^2}.$$

Similarly, the Klein-Gordon equation of relativistic quantum mechanics can be derived, using these substitutions for the relativistic energy-momentum relation. It is then given by

$$E^2 = \mathbf{p}^2 + m^2 \rightarrow (\square + m^2) |\phi(\mathbf{x}, t)\rangle = 0.$$

The last conceptual leap in establishing QFT is to abandon a fixed number of particles represented by states ($|\phi(\mathbf{x}, t)\rangle$) for a *Fock space*, in which particles are represented by operators. These operators are the quantum analogue of the field, which represent particles

when applied to the vacuum state, $|0\rangle$. A multi-particle state is represented as:

$$|\mathbf{p}_1, \mathbf{p}_2, \dots, \mathbf{p}_n\rangle = \hat{a}^\dagger(\mathbf{p}_1)\hat{a}^\dagger(\mathbf{p}_2)\dots\hat{a}^\dagger(\mathbf{p}_n)|0\rangle,$$

where the $\hat{a}^\dagger(p_i)$ represent particles of momentum p_i . This is called *second quantisation*.

The Klein-Gordon Hamiltonian can now be expressed as

$$:\hat{\mathcal{H}}:=\hat{\mathcal{H}}-\langle 0|\hat{\mathcal{H}}|0\rangle=\int\frac{\mathbf{d}^3\mathbf{p}}{(2\pi)^32E_{\mathbf{p}}}E_{\mathbf{p}}\hat{a}^\dagger(\mathbf{p})\hat{a}(\mathbf{p}),$$

where the $:\hat{\mathcal{H}}:$ differs from the standard QFT Hamiltonian operator due to normal ordering, where the energy of the vacuum state (which is infinite) is subtracted.

A scalar field is introduced into a QFT model via the Lagrangian, and can be represented by an expression of the form

$$\hat{\mathcal{L}}=\frac{1}{2}(\partial_\mu\hat{\phi})^2-\frac{1}{2}m^2\hat{\phi}^2-\frac{\lambda}{4}\hat{\phi}^4,$$

where

$$\hat{\phi}=\int\frac{\mathbf{d}^3\mathbf{p}}{(2\pi)^32E_{\mathbf{p}}}(a(\mathbf{p})e^{-ipx}+\hat{a}^\dagger(\mathbf{p})e^{ipx}).$$

The first term in $\hat{\mathcal{L}}$ introduces the field to the model, the second gives it mass, and the last allows it to interact. Expanding the last term in terms of ladder operators results in terms with non-equal numbers of annihilation and creation operators, which represent the number of particles changing, the major benefit of QFT.

States in QFT then evolve according to *Dyson's formula*,

$$|\phi\rangle(t)=\hat{T}(e^{-i\int_{t_0}^t dt'\hat{\mathcal{H}}_{\mathcal{I}}(t')})|\phi\rangle(t_0),$$

where $\hat{T}$ is the time-ordering operator which commutes the operators with the lowest time value to the left-most position in the term, and $\hat{\mathcal{H}}_I$ is the Hamiltonian of only the interaction terms of the Lagrangian.

The *scattering* S -matrix elements for a given process can then be calculated as the limit of this time evolution matrix as $t \rightarrow \inf$,

$$S_{if} = \langle f | \hat{S} | i \rangle = \lim_{t \rightarrow \infty} \langle f | \hat{U}(t, -t) | i \rangle,$$

where

$$U(t, t_0) = T e^{-i \int_{t_0}^t dt' H_I(t')},$$

and T is the time ordering operator, which commutes operators with larger temporal arguments to the right of the expression. Decay rates and interaction cross sections are then obtained by calculating the value of the S -matrix and integrating over the available phase space using Fermi's golden rule,

$$P(i \rightarrow f) = 2\pi |S(i, f)|^2 \rho,$$

where ρ represents the properly normalised available phase space, and takes different values depending on whether this equation is being used to calculate an interaction cross section or a decay rate.

1.1.2 Feynman Diagrams

While this provides the starting point of all calculations in QFT, the calculation of the S -matrix element requires it to be expanded using *Wick's theorem*, which is often difficult to calculate, and beyond the scope of this thesis. However, this expansion is equivalent to a much simpler, more intuitive method: Feynman diagrams. The principle of Feynman

diagrams is very simple: you draw all diagrams connecting the initial and final states particles at a given order (number of vertices) of perturbation theory. You can then calculate the amplitude for each diagram as per the Feynman rules of the theory. The S -matrix is then simply the sum of all of these amplitudes. While the Feynman rules are different for different models, there are some similarities between the momentum-space Feynman rules for most models in QFT. These are:

1. Internal lines are represented by the momentum-space Feynman propagator, which varies between models, but usually has a denominator which scales with both the momentum and mass of the intermediate particle.
2. Each vertex introduces a constant (neglecting running from renormalisation) factor representing the strength of the interaction, and includes a δ -function which conserves momentum at the vertex.
3. All undefined momentum must be integrated over.

1.1.3 Particles

The entire known universe is made of particles, and all forces/interactions are thought to be mediated by particles. The only possible exception is gravity, for which we have not yet observed a mediating particle, but is thought to be mediated by a spin-2 massless particle called the graviton. More precisely, these particles are fluctuations of the quantum fields described previously. They are the initial and final states which Feynman diagrams connect together, and are produced (destroyed) by creation (annihilation) operators.

Particles at the edges of Feynman diagrams are *real*, which means that they are mass eigenstates of their respective field, and propagate through spacetime as quantum, relativistic objects. They correspond to the state of a particle when it is observed.

Conversely, the particles which connect the outer edges of Feynman diagrams are called *virtual* particles. These are off their mass-shell, which means that they do not necessarily have the mass of their respective field. They do not propagate large distances through

spacetime due to the uncertainty principle, and are used in the calculation of decay rates and scattering cross sections. They do not interact with anything but the particles at the edges of the Feynman diagram, and are therefore never observed themselves.

1.2 The Standard Model of Particle Physics

The Standard Model (SM) is a *gauge theory* respecting a $U(1)_Y \times SU(2)_L \times SU(3)_c$ local symmetry [12–15], in which the $U(1)_Y \times SU(2)_L$ symmetry is broken to a $U(1)_{em}$ symmetry by the BEH mechanism [16–18], named after Brout, Englert and Higgs. A gauge theory is a relativistic QFT, in which interactions are introduced to promote global gauge symmetries to local gauge symmetries. The twelve fermions of the SM make up all stable matter. The breaking of the $U(1) \times SU(2)$ gauge symmetry gives rise to the electromagnetic (EM) and weak interactions, which together form the EW sector. The $SU(3)$ gauge symmetry represents the quantum chromodynamic sector, which is responsible for the strong interaction. The following subsections describe each of these sectors, before being brought together to form the SM.

1.2.1 The Standard Model Fermions

The SM fermions consist of quarks and leptons, shown in Table 1.1 and Table 1.2, respectively. They are all spin-half particles.

All of these particles form weak isospin doublets, which couple to the gauge bosons of the weak interaction described in §1.2.2. Only the quarks are colour-charged and therefore couple to the gauge bosons of the strong interaction, described in §1.2.3. All of these particles except the neutrinos are electromagnetically charged, and therefore couple to the EM force.

Particle	Mass / GeV	Q	I_3
Up-type quarks			
Up-quark	$(2.16^{+0.49}_{-0.26}) \times 10^{-3}$	$+\frac{2}{3}$	$+\frac{1}{2}$
Charm-quark	1.27 ± 0.02	$+\frac{2}{3}$	$+\frac{1}{2}$
Top-quark	172.9 ± 0.4	$+\frac{2}{3}$	$+\frac{1}{2}$
Down-type quarks			
Down-quark	$(4.67^{+0.48}_{-0.17}) \times 10^{-3}$	$-\frac{1}{3}$	$-\frac{1}{2}$
Strange-quark	$(93^{+11}_{-5}) \times 10^{-3}$	$-\frac{1}{3}$	$-\frac{1}{2}$
Bottom-quark	$4.18^{+0.03}_{-0.02}$	$-\frac{1}{3}$	$-\frac{1}{2}$

Table 1.1: List of Standard Model quarks, including their mass, EM charge (Q), and the third component of their isospin (I_3). Values are taken from Ref. [19]. The value for the top mass is taken from a direct measurement.

Particle	Mass / GeV	Q	I_3
Leptons			
Electron	$(0.5109989461 \pm 0.0000000031) \times 10^{-3}$	-1	$-\frac{1}{2}$
Muon	$(105.6583745 \pm 0.0000024) \times 10^{-3}$	-1	$-\frac{1}{2}$
Tau-lepton	1.77686 ± 0.00012	-1	$-\frac{1}{2}$
Neutrinos			
Electron neutrino	$< 1.1 \times 10^{-9}$	0	$+\frac{1}{2}$
Muon neutrino	$< 1.1 \times 10^{-9}$	0	$+\frac{1}{2}$
Tau-lepton neutrino	$< 1.1 \times 10^{-9}$	0	$+\frac{1}{2}$

Table 1.2: List of Standard Model leptons and neutrinos, including their mass, EM charge (Q), and the third component of their isospin (I_3). Charged lepton mass values are taken from Ref. [19], while the neutrino mass upper limits are set at the 90% confidence interval by a procedure described in Ref. [20].

1.2.2 The Electroweak Sector

The SM respects a $U(1)_Y \times SU(2)_L$ global gauge symmetry. In order to promote this symmetry to a local gauge symmetry, one field is introduced for the $U(1)_Y$ symmetry (B^0), and three fields are introduced for the $SU(2)_L$ symmetry (W^0 and $W^\pm$). This requirement introduces the electroweak (EW) sector to the SM.

The three W fields couple to the third component of the weak isospin, I_3 (effectively all left-handed chiral states), with the same coupling constant, g , and do not interact with right-handed chiral states. This force is maximally parity violating due to the $V - A$ form of the interaction. The $SU(2)_L$ symmetry is non-Abelian, meaning that the fields associated with this symmetry interact with each other.

The B^0 couples to weak hypercharge $Y = 2(Q - I_3)$, where Q is the familiar electromagnetic (EM) charge and I_3 is the third component of the weak isospin, with a coupling constant: $g'/2$.

The neutral fields then rotate through the weak mixing angle, θ_W , to produce the familiar SM neutral vector-bosons: Z and A :

$$A_\mu = +B_\mu \cos \theta_W + W_\mu^3 \sin \theta_W$$

$$Z_\mu = -B_\mu \sin \theta_W + W_\mu^3 \cos \theta_W.$$

To successfully recover the familiar EM interaction of the Standard Model, the *Unification Condition* must be satisfied:

$$e = g \sin \theta_W = g' \cos \theta_W.$$

All of the above interactions conserve lepton number, except for in the neutrino sector, where the fields which couple weakly are obtained by a rotation of the mass eigenstates. This rotation is quantified by the PMNS matrix [21]. Interactions via the Z and γ

bosons also conserve quark number. However, the charged current weak interaction ($W^\pm$) interacts with weak eigenstates, which are obtained by a rotation of the quark mass eigenstates. This rotation is quantified by the CKM matrix [22, 23].

1.2.3 Quantum Chromodynamics

The final interaction of the Standard Model is Quantum Chromodynamics (QCD) [24], which is introduced by the promotion of the global $SU(3)_c$ gauge symmetry to a local gauge symmetry. This generates eight gluon fields, which couple to the colour charge (r , g and b) with a uniform strong coupling: g_s . QCD is a non-Abelian gauge theory, so the eight gluons carry colour charge, and are thus self-interacting.

QCD is also called the strong force, because it has the largest coupling at most energy scales. Due to the gluon loops in the gluon propagator, g_s decreases with increasing energy scale. At low energy scales, the strong coupling constant becomes so large that perturbation theory breaks down. Therefore, QCD calculations at low energy scale are typically performed using phenomenological models, based on quark and gluon splitting functions. However, at high energy scales, the strength of the strong force becomes negligible compared to the scales at hand. This property is called asymptotic freedom [25], and it ensures that quantities are perturbatively calculable at high energy scales. It also ensures that high-energy interacting partons can be considered as essentially free particles.

Phenomenological QCD models include a constant force term. This means that an infinite amount of energy is required to separate two colour charged objects, therefore no colour-charged object has ever been observed in isolation. This effect is called colour confinement.

1.2.4 The BEH Mechanism

The motivation behind the introduction of the Higgs sector into the SM [17] is threefold: first, it was introduced into the SM to facilitate EW symmetry breaking, by which the

EW gauge bosons acquire mass; second, the Higgs field is required to introduce fermionic mass terms into the SM Lagrangian without breaking its gauge symmetry and spoiling the renormalisability of the theory; third, without the Higgs boson the $W^+W^- \rightarrow W^+W^-$ longitudinal scattering cross section diverges.

Three of the four gauge bosons of the SM have non-zero masses, but writing a mass term into the Lagrangian as would naively be done in field theory breaks the local gauge symmetry, which is axiomatic to the SM. As such, the BEH symmetry breaking mechanism was introduced into the SM, facilitating the introduction of mass terms into the SM Lagrangian in a gauge invariant way.

The Higgs field is a complex scalar doublet, which is introduced into the SM via the Higgs sector Lagrangian:

$$L_h = D_\mu^\dagger \phi D^\mu \phi - V(\phi),$$

where the Higgs field potential takes the ‘Mexican-hat’ shape

$$V(\phi) = \lambda(\phi^\dagger \phi)^2 - \mu^2 \phi^\dagger \phi,$$

and the covariant derivative is given by

$$D_\mu \equiv \partial_\mu + ig\vec{\tau} \cdot \vec{w}_\mu + ig'Y B_\mu/2.$$

The covariant derivative has three terms: the first is a partial derivative, which is required to define the kinetic term; the second ensures that the Lagrangian respects a local $SU(2)_L$ gauge invariance; and the third ensures that the Lagrangian respects a local $U(1)_Y$ gauge invariance. The $\vec{W}_\mu$ and B_μ fields introduced here undergo EW mixing to become the EW gauge bosons of the SM (γ , Z and $W^\pm$), which therefore couple to the SM Higgs boson via this kinetic term.

The Higgs potential has a non-zero ground state, which can be calculated from the Higgs potential as

$$\phi_0^\dagger \phi_0 = \mu^2/2\lambda.$$

However, this does not fully constrain the full form of the Higgs field, so we choose the solution

$$\phi_0 = \frac{1}{\sqrt{2}} \begin{pmatrix} 0 \\ v \end{pmatrix}, \text{ where } v \equiv \frac{\mu}{\sqrt{\lambda}},$$

as by Goldstone's theorem [26], this choice results in the photon remaining massless.

This symmetry breaking procedure generates the mass terms for the EW gauge bosons. If we substitute this post-symmetry-breaking Higgs field into the Higgs kinetic term in the Lagrangian, after some algebra, we recover the mass terms of the EW gauge bosons

$$D_\mu^\dagger \phi_0 D^\mu \phi_0 = \frac{1}{8}v^2 g^2 W_\mu^+ W^{-\mu} + \frac{1}{8}v^2 \begin{pmatrix} W^{(3)}_\mu & B_\mu \end{pmatrix} \begin{pmatrix} g^2 & -gg' \\ -gg' & g'^2 \end{pmatrix} \begin{pmatrix} W^{(3)\mu} \\ B^\mu \end{pmatrix}.$$

It can immediately be read off from this that the W boson has a mass of

$$m_{W^\pm} = \frac{1}{2}vg.$$

However, the neutral bosons (the Z boson and photon) are mixed in this notation. To make their presence obvious, we must diagonalise the mass matrix to reveal the physical mass eigenstates. The resulting states are

$$A_\mu = B_\mu \cos \theta_W + W^{(3)}_\mu \sin \theta_W, \text{ and } Z_\mu = -B_\mu \sin \theta_W + W^{(3)}_\mu \cos \theta_W,$$

where $\tan \theta_W \equiv g'/g$.

Which have masses of

$$m_\gamma = 0 \text{ and } m_Z = \frac{1}{2}\sqrt{g^2 + g'^2},$$

respectively. θ_W is the weak mixing angle. This completes the derivation of the masses of the EW gauge bosons.

To make the mass terms of the Higgs boson and fermions explicit we must expand the Higgs field around its vacuum expectation value as shown:

$$\phi = \frac{1}{\sqrt{2}} \begin{pmatrix} 0 \\ v + h \end{pmatrix},$$

where h is a small perturbation. After some algebra, the Higgs potential now becomes

$$-V(\phi) = \frac{\mu^4}{4\lambda} - \mu^2 h^2 - \mu h^3 \sqrt{\lambda} - \frac{\lambda h^4}{4},$$

from which it can be clearly seen that the Higgs boson has cubic and quartic self coupling terms, and a mass of

$$m_h = v\sqrt{2\lambda}.$$

The fermion masses are not predicted by the SM, but are added in an ad-hoc fashion by the introduction of the Yukawa couplings. However, the SM predicts that the fermions couple to the Higgs boson with a strength proportional to their mass. The Lagrangian term for charged leptons can be written

$$L_l = -G_l \left[\begin{pmatrix} \bar{\nu}_l & \bar{l} \end{pmatrix}_L \begin{pmatrix} \phi^+ \\ \phi^0 \end{pmatrix} l_R + \bar{l}_R \begin{pmatrix} \phi^- & \bar{\phi}^0 \end{pmatrix} \begin{pmatrix} \nu_l \\ l \end{pmatrix}_L \right],$$

which after spontaneous symmetry breaking, becomes

$$\begin{aligned} L_l &= -\frac{G_l}{\sqrt{2}} [\bar{l}_L(v+h)l_R + \bar{l}_R(v+h)l_L] = -\frac{G_lv}{\sqrt{2}}(\bar{l}_L l_R + \bar{l}_R l_L) - \frac{G_l}{\sqrt{2}}(\bar{l}_L h l_R + \bar{l}_R h l_L) \\ &= -m_l \bar{l}l - \frac{m_l}{v} \bar{l}h l, \end{aligned}$$

where the mass of the charged lepton is given by

$$m_l = \frac{G_l v}{\sqrt{2}}.$$

To introduce the masses of the quarks, it becomes necessary to use the charge conjugated Higgs field

$$\phi_c \equiv -i\gamma^2 \phi^* = \begin{pmatrix} -\bar{\phi}_0 \\ \phi^- \end{pmatrix}.$$

Which, post symmetry breaking, takes the value

$$\phi_C = \begin{pmatrix} v+h \\ 0 \end{pmatrix}.$$

Now, analogously to the lepton mass term, a quark mass term can be introduced as

$$L_q = -G_d \begin{pmatrix} \bar{u} & \bar{d} \end{pmatrix}_L \phi d_R - G_u \begin{pmatrix} \bar{u} & \bar{d} \end{pmatrix}_L \phi_C u_R + \text{h.c.},$$

which after spontaneous symmetry breaking, becomes

$$L_q = -m_d \bar{d}d - m_u \bar{u}u - \frac{m_d}{v} \bar{d}h d - \frac{m_u}{v} \bar{u}h u,$$

where $m_q = \frac{G_q v}{\sqrt{2}}$.

This completes the introduction of locally gauge invariant mass terms into the SM Lagrangian.

1.3 Properties and Phenomenology of the SM Higgs Boson

In 2012, the ATLAS and CMS collaborations observed a particle compatible with the SM Higgs boson [27, 28]. The combined ATLAS and CMS Run 1 data has since shown this new particle to have a mass of $125.09 \pm 0.21(\text{stat}) \pm 0.11(\text{sys})$ GeV [29].

This particle has been observed to decay to a di-photon final state [30], meaning that it must be spin 0 or 2 by the Landau-Yang theorem, and that it must have $C = +1$. Analyses in its ZZ , $\gamma\gamma$ and W^+W^- decay channels have found its J^P properties to be consistent with 0^+ , and incompatible with 0^- , 2^+ and $1^\pm$ [31]. These observations make it probable that it has the J^{PC} properties of the SM Higgs boson: 0^{++} .

A SM Higgs boson of mass of 125 GeV will have a width of 4.2 MeV, though experimental constraints on the width are either loose or indirect. An attempt by CMS to directly probe the width using events in the $H \rightarrow ZZ^* \rightarrow 4\ell$ decay mode finds a 95% confidence level (CL) upper limit on the width at 1.1 GeV [32]. While a measurement from CMS of the off mass-shell production rates in the $H \rightarrow ZZ^*$ and $H \rightarrow W^+W^-$ decay modes places a 95% CL upper limit at 13 MeV [33], and a similar measurement by ATLAS in the $H \rightarrow ZZ^*$ channel places a 95% CL_s¹ upper limit at 14.4 MeV [34].

¹CL_s is used to avoid excluding hypotheses due to fluctuations of the data, in cases where you cannot differentiate between the null and alternative hypotheses. CL_s is defined as $p_1/(1 - p_0)$, where p_0 and p_1

These assume that the couplings to vector bosons are the same on mass-shell as they are off mass-shell.

An understanding of the production and decay modes of the Higgs boson is necessary to any Higgs physics programme. These are discussed in the following subsections, assuming SM production and decay modes and a Higgs boson mass of 125 GeV. While it should be noted that the Large Hadron Collider (LHC) can not independently measure the production or decay modes of the Higgs boson, model-dependent constraints can be obtained using combined fits to the search results for the various combined production and decay modes [35]. Figure 1.1 shows some of the measured combined production-decay modes of the Higgs boson, normalised to SM expectation. At the time of writing, no measured properties of the Higgs boson have any significant tensions with the SM expectations.

1.3.1 Higgs Boson Production Mechanisms at the LHC

The main Higgs boson production processes at the LHC are given in Table 1.3. All of these production modes have been observed except the associated production with a bottom quark pair, which has a small cross section, and a signature with large backgrounds from multijet production.

Production Mechanism	SM Expectation / pb
Gluon-fusion	$43.92^{+4.50}_{-4.36}$
Vector boson fusion	3.748 ± 0.123
WH Associated Production	$1.380^{+0.032}_{-0.037}$
ZH Associated Production	0.8696 ± 0.0382
$b\bar{b}H$ Associated Production	$0.5116^{+0.0781}_{-0.1267}$
$t\bar{t}H$ Associated Production	$0.5085^{+0.0533}_{-0.0651}$

Table 1.3: Standard model calculated cross sections for various Higgs boson production mechanisms at 13 TeV, calculated assuming a Higgs boson of mass 125 GeV. The total uncertainty is due to QCD scale, PDF and α_s uncertainties. Values are taken from the CERN Higgs Cross Section working group Handbook of LHC Higgs Cross Sections [37–39].

With a cross section of 43.9 pb at 13 TeV [37–39], by far the largest contribution to

are the p-values of the null and alternative hypotheses, respectively. It is then used in place of p_1 when setting exclusion intervals for a parameter of the alternative hypothesis.




Figure 1.1: Combined production-decay modes of the Higgs boson, normalised to SM expectation [36]. As measured by a simultaneous fit to all production and decay channels by the ATLAS Collaboration. The cross sections of the (ggF) $H \rightarrow b\bar{b}$, (VBF) $H \rightarrow WW^*$, (VH) $H \rightarrow \tau^+\tau^-$ processes are fixed to their SM expectations. The blue boxes show the systematic uncertainties, yellow boxes show the statistical uncertainties, and the black bars are the combined uncertainties on the measurements.

the total Higgs boson cross section at the LHC is through gluon fusion production, in which pairs of gluons fuse through a top-quark loop to form a Higgs boson, as shown in Figure 1.2. This production channel was the first to be observed, and was responsible for the discovery of the Higgs boson. It usually produces a single Higgs boson, with almost no transverse momentum, offering no signature with which to tag the production. However, jets can be emitted from the top-quark loop, which give the Higgs boson a momentum boost transverse to the beamline.




Figure 1.2: Feynman diagram for gluon-fusion Higgs boson production.

Vector boson fusion occurs when two Z or two W bosons are emitted from initial state quarks, and fuse to produce a Higgs boson. This is the second largest contribution to the total Higgs boson cross section with a cross section of 3.75 pb [37–39], and further benefits from the distinctive final state offered by the two high p_T , high invariant mass jets initiated by the two initial state quarks. The Higgs boson can also be produced with non-negligible transverse momentum, as it recoils against one of the initial state quarks. The production of Higgs bosons via this channel was only recently observed, through a combination of 6 different decay modes.

Higgs bosons can be produced after having been radiated from a vector boson, produced in the initial pp collision. The cross sections for WH and ZH production are 1.38 pb and 0.870 pb, respectively [37–39]. While this is a relatively small contribution to the total Higgs boson cross section, it offers a distinctive final state, with which to tag the production mode: the decay mode of the vector boson. For the 6.7% (21%) of events in which the Z (W) boson decays to a final state containing one or more electrons or muons,

or their anti-particles (throughout this thesis “electron” and “muon” will be used to refer to both the lepton and the anti-lepton). These production modes can be used to tag the presence of a Higgs boson. The Higgs boson and Z boson can also recoil against each other, and so be produced with significant transverse momentum. For this reason, this mode was first observed in the most common decay mode of the Higgs boson: $H \rightarrow b\bar{b}$.

Higgs bosons can be radiated from a bottom or top quark, which has been pair produced. This has a relatively small cross section of 0.512 pb and 0.509 pb [37–39], for Higgs bosons produced in association with a bottom and top quark pair, respectively. The $t\bar{t}H$ production mode was only recently observed [40], by combining many of the decay modes of the Higgs boson, making it the only coupling of the Higgs boson to an up-type fermion to be observed. However, the $b\bar{b}H$ production mode has not been observed due to the large backgrounds at the LHC.

1.3.2 Higgs Boson Decay Channels

The main decay modes of the Higgs boson, along with the associated branching ratio (BR) for each mode, are listed in the Table 1.4. They have all been observed except for the decays to pairs of gluons, pairs of charm-quarks, $Z\gamma$ and pairs of muons.

Final State	SM Expectation / %
$b\bar{b}$	57.7 ± 1.9
W^+W^-	21.5 ± 0.9
gg	$8.57^{+0.88}_{-0.86}$
$\tau^+\tau^-$	6.32 ± 0.36
$c\bar{c}$	$2.91^{+0.35}_{-0.36}$
ZZ	2.64 ± 0.11
$\gamma\gamma$	0.228 ± 0.011
$Z\gamma$	0.154 ± 0.014
$\mu^+\mu^-$	0.0219 ± 0.0013

Table 1.4: Standard model calculated BRs of the Higgs boson to various final states. Values are taken from the CERN Higgs Cross Section working group Handbook of LHC Higgs Cross Sections [37–39].

1.4 Extending the Higgs Sector

At this point, it should be noted that the SM BEH Mechanism is not the only mechanism which can facilitate EW symmetry breaking, and give mass to the gauge bosons and fermions; just the simplest. Being a complex doublet, the SM Higgs field possesses 4 degrees of freedom, 3 of which are required for the 3 massive gauge bosons to acquire mass, and the remaining degree of freedom gives rise to the Higgs boson.

In constructing beyond the Standard Model (BSM) Higgs sectors, $\rho \equiv m_W^2/(m_Z^2 \cos^2 \theta_W)$ must remain close to its SM value of 1, as this value has been confirmed experimentally. While this result could be achieved through a careful fine-tuning of the structure of the Higgs sector, a ‘natural’ solution arises from the realisation that any arrangement of Higgs singlets and doublets automatically satisfies this requirement. So the simplest natural extensions to the SM Higgs sector are: the Higgs Doublet Model with an additional singlet [41]; 2 Higgs Doublet Model (2HDM) [41]; and the 2 Higgs Doublet Model with an additional singlet (2HDM+S) [41, 42]. Many other combinations of singlets, doublets and even triplets are possible, with many leading to exotic signatures, such as doubly charged Higgs bosons [43]. These more exotic Higgs sector will not be discussed further here.

1.4.1 The Two Higgs Doublet Model

The 2 Higgs Doublet Model (2HDM) [19, 41, 41, 44, 45] is one of the simplest extensions to the SM Higgs sector, with a Higgs field content consisting of two complex Higgs doublet fields instead of one. The most general 2HDM potential, invariant under the SM gauge symmetry, is given by

$$\begin{aligned}
V(\phi_1, \phi_2) = & -\mu_1^2 \phi_1^\dagger \phi_1 - \mu_2^2 \phi_2^\dagger \phi_2 - [m_{12}^2 \phi_1^\dagger \phi_2 + \text{h.c.}] \\
& \lambda_1 (\phi_1^\dagger \phi_1)^2 \lambda_2 (\phi_2^\dagger \phi_2)^2 + \lambda_3 (\phi_1^\dagger \phi_1) (\phi_2^\dagger \phi_2) + \lambda_4 (\phi_1^\dagger \phi_2) (\phi_2^\dagger \phi_1) \\
& + [\frac{1}{2} \lambda_5 (\phi_1^\dagger \phi_2)^2 + \lambda_6 (\phi_1^\dagger \phi_1) (\phi_1^\dagger \phi_2) + \lambda_7 (\phi_1^\dagger \phi_2) (\phi_2^\dagger \phi_2) + \text{h.c.}],
\end{aligned}$$

which contains four real mass parameters $\mu_{1,2}^2$, $\text{Re}(m_{12}^1)$, $\text{Im}(m_{12}^1)$, and ten real quartic couplings $\lambda_{1,2,3,4}$, $\text{Re}(\lambda_{5,6,7})$ and $\text{Im}(\lambda_{5,6,7})$ [45]. This has vacuum solutions [45]

$$\langle \phi_1 \rangle = \frac{1}{2} \begin{pmatrix} 0 \\ v_1 \end{pmatrix} \text{ and } \langle \phi_2 \rangle = \frac{1}{2} \begin{pmatrix} 0 \\ v_2 \end{pmatrix},$$

where v_1 and v_2 are real. By expanding about these vacuum solutions and substituting them back into the potential, the presence of 5 Higgs bosons can be derived. This is because there are now 8 degrees of freedom, such that once 3 are dedicated to giving mass to the SM gauge bosons, 5 remain to produce scalar Higgs states. These states are: two CP-even neutral Higgs bosons, h^0 and H^0 , where h^0 is often associated with the observed Higgs boson; one CP-odd neutral Higgs boson, a^0 , often referred to as a ‘pseudoscalar’ as it behaves as a pseudoscalar in its interactions with fermions; and two charged Higgs bosons $H^\pm$. Their masses are given by

$$\begin{aligned}
m_{h^0, H^0}^2 = & \frac{1}{2} [(A + B) \pm \sqrt{(A - B)^2 + 4C^2}], \\
m_{a^0}^2 = & \frac{m_{12}^2}{\sin \beta \cos \beta} - v^2 \lambda_5 + \frac{v^2}{2 \sin \beta \cos \beta} (\lambda_6 \cos^2 \beta + \lambda_7 \sin^2 \beta), \\
\text{and } m_{H^\pm}^2 = & \frac{m_{12}^2}{\sin \beta \cos \beta} - \frac{v^2}{2} (\lambda_4 + \lambda_5) + \frac{v^2}{2 \sin \beta \cos \beta} (\lambda_6 \cos^2 \beta + \lambda_7 \sin^2 \beta),
\end{aligned}$$

where

$$\begin{aligned}
A = & m_{a^0}^2 \sin^2 \beta + v^2 (2\lambda_1 \cos^2 \beta + \lambda_5 \sin^2 \beta + 2\lambda_6 \sin \beta \cos \beta), \\
B = & m_{a^0}^2 \cos^2 \beta + v^2 (2\lambda_2 \sin^2 \beta + \lambda_5 \cos^2 \beta + 2\lambda_7 \sin \beta \cos \beta), \\
C = & -m_{a^0}^2 \sin \beta \cos \beta + v^2 (\lambda_{34} \sin \beta \cos \beta + \lambda_6 \cos^2 \beta + \lambda_7 \sin^2 \beta),
\end{aligned}$$

and $\tan \beta \equiv v_1/v_2$.

After constraints from symmetry requirements, measurements in the EW sector and CP conservation, there are 6 degrees of freedom left in the Higgs sector of the 2HDM Lagrangian, these are: $\tan \beta$, $\sin(\beta - \alpha)$, $m_{H^\pm}^2$, $m_{a^0}^2$, $m_{h^0}^2$ and $m_{H^0}^2$. Where α is the mixing angle in the neutral scalar Higgs sector.

A constraint on any extended Higgs sector is the need to avoid introducing tree-level flavour changing neutral currents (FCNCs), as large FCNCs are excluded by stringent experimental constraints [46–48]. This can be achieved in the 2HDM if all of the fermions of a given electric charge couple to just one Higgs doublet. Two natural ways of achieving this are by: coupling all of the fermions to just one of the doublets (Type-I); coupling all of the up-type fermions to one doublet, and down-type fermions to the other (Type-II). There is also a Type-III 2HDM, in which the quarks couple to one Higgs doublet, and leptons couple to the other, though this requires another mechanism to remove tree-level FCNCs. Finally, there is the Type-IV 2HDM, in which the up-type quarks and down-type leptons couple one Higgs doublet, and the down-type quarks couple to the other.

2HDM Higgs sectors have many strong motivations beyond the arbitrariness of the SM Higgs sector. First, the Type-II 2HDM is of particular interest, because it is required to generate all of the fermion masses in the minimal Supersymmetric extension to the Standard Model (MSSM). Supersymmetry is itself of interest because it automatically solves the Higgs boson mass hierarchy problem, in which loops in Feynman diagrams cause divergences in the mass of the Higgs boson. Supersymmetry also provides a dark matter candidate, which is usually the lightest stable supersymmetric particle. Second, CP violation can occur at tree-level in the 2HDM, which could explain the observed baryon asymmetry in the universe [49]. Third, 2HDMs can provide dark matter candidates [50–52]. One such example of this is the Inert Higgs Doublet Model, in which an additional Higgs doublet does not develop a vacuum expectation value meaning it does not couple to the gauge bosons, and it respects a Z_2 symmetry meaning it does not couple to the

leptons, making its Higgs bosons an ideal dark matter candidate.

1.4.2 The Alignment and Decoupling Limits

In the 2HDM the coupling of the lighter (heavier) scalar Higgs boson to pairs of W or Z bosons is equal to the SM value, multiplied by $\sin(\beta - \alpha)$ ($\cos(\beta - \alpha)$) [53]. As these couplings of the observed Higgs boson have been measured to be close to their SM expectations [54, 55], this constrains the value of $\sin(\beta - \alpha)$ ($\cos(\beta - \alpha)$) to be close to unity, assuming it is the lighter (heavier) scalar 2HDM Higgs boson. Furthermore, all of the other couplings of lighter (heavier) scalar Higgs bosons approach the SM values as $\sin(\beta - \alpha)$ ($\cos(\beta - \alpha)$) approaches unity [53]. Therefore, should the observed Higgs boson be a 2HDM scalar Higgs boson, either $\sin(\beta - \alpha)$ or $\cos(\beta - \alpha)$ must be close to unity. This is called the ‘alignment limit’ [41, 44, 45, 53, 56]. Figure 1.3 shows the constraints on $\cos(\beta - \alpha)$ and $\tan \beta$ for two 2HDM scenarios from fits to ATLAS data.




Figure 1.3: Regions of the $\cos(\beta - \alpha)$ and $\tan \beta$ plane excluded at the 95% confidence level, for the (a) Type-I and (b) Type-II 2HDM, from fits to various combined production-decay modes of the Higgs boson by the ATLAS Collaboration [36]. The cross marks the best fit value, and the dashed red line marks the alignment limit of $\cos(\beta - \alpha) = 0$.

There are three main ways of achieving the alignment limit. The first and simplest

possibility is that the observed Higgs boson is the lighter 2HDM scalar state, and all of the other 2HDM Higgs bosons either: have masses at a much higher scale [44], called large mass decoupling; or have very weak couplings to all other particles [56], called weak coupling decoupling. This is called the ‘decoupling limit’, and large mass decoupling can be formally expressed as $m_{a^0}^2 \gg \lambda_i v^2$ [44], where $\lambda_i = \mu_{1,2,12}$, or $\lambda_{1,2,3,4,5,6,7}$. Second, the alignment limit can be achieved without decoupling through the introduction of one of three symmetries, which are described in detail in Ref. [45]. Third, the alignment limit can be achieved through a careful fine-tuning of the parameters or the Higgs potential, though this last possibility is not natural [45].

The alignment limit has some noteworthy phenomenological features. First, there are no tree level $a^0 VV$ or $H^\pm VV$ couplings [53]. Second, if the observed Higgs state is the lighter (heavier) 2HDM Higgs boson then the XWW , XZZ , $Za^0 Y$, $WH^\pm Y$, $ZWH \pm Y$, $\gamma WH \pm Y$ couplings are suppressed by a factor of $\cos(\beta - \alpha)$ ($\sin(\beta - \alpha)$), where X (Y) is the lighter 2HDM Higgs boson and Y (X) is the heavier 2HDM Higgs boson [44]. Third, all vertices with at least one vector boson and exactly one non-minimal 2HDM Higgs boson are suppressed by a factor of $\cos(\beta - \alpha)$ in the decoupling limit [44]. Fourth, in the Type-I 2HDM the couplings of h^0 to fermions is proportional to $\cos \alpha / \sin \beta$, couplings of H^0 to fermions is proportional to $\sin \alpha / \sin \beta$, and the couplings of a^0 to fermions is proportional to $\cot \beta$ [53]. Fifth, in the Type-II 2HDM the couplings of h^0 to up-type (down-type) fermions is proportional to $\cos \alpha / \sin \beta$ ($\sin \alpha / \cos \beta$), the couplings of H^0 to up-type (down-type) fermions is proportional to $\sin \alpha / \sin \beta$ ($\cos \alpha / \cos \beta$), and the couplings of a^0 to up-type (down-type) fermions is proportional to $\cot \beta$ ($\tan \beta$) [53].

1.4.3 The Two Higgs Doublet Model with an Additional Singlet

The 2 Higgs Doublet Model with an additional singlet (2HDM+S) extends the 2HDM by one singlet superfield [42], predicting the existence of 7 Higgs bosons: the 5 Higgs bosons of the 2HDM; a neutral CP-even Higgs boson; and a neutral CP-odd Higgs boson. The Type-II version of the 2HDM+S is required to give masses to the fermions in the next-

to-minimal Supersymmetric extension to the Standard Model (NMSSM). The NMSSM is motivated by all the same considerations as the MSSM. Additionally, it greatly reduces the little hierarchy problem [57], and solves the μ -problem of the MSSM [58], in which the supersymmetric Higgs mass parameter, μ , is forced to be around the EW scale to avoid fine-tuning.

1.4.4 Previous Constraints on Light Higgs Bosons

In the 2HDM and 2HDM+S, the lightest CP-odd scalar boson can have a mass below half that of the observed Higgs boson, allowing the possibility of decays of the observed Higgs boson to pairs of these new light Higgs bosons. Alternatively, if the mass of the a^0 is less than the mass difference between the observed Higgs boson and the Z boson, then the Higgs boson can decay to a Za^0 final state. Over a large part of the 2HDM and 2HDM+S phase spaces, there is a significant BR for these decays [41]. The BRs of the light pseudoscalar in the 2HDM and 2HDM+S is dependent on $\tan\beta$, and is given for two $\tan\beta$ values in the 2HDM+S in Figure 1.4.




Figure 1.4: BRs of the light pseudoscalar in the Type-II 2HDM+S, for $\tan\beta$ values of (a) 0.5 and (b) 5 [41].

Chapters 5 and 6 describe searches for Higgs boson decays to light resonances, which are predicted in the 2HDM and 2HDM+S. This resonance can be interpreted as the a^0 , or as the h^0 if the observed Higgs boson is H^0 . In addition to indirect constraints from coupling measurements [36, 59], previous searches for such resonances have been performed. At the LHC, these include 95% CL_s upper limits from CMS a^0 production [60–64], and the ATLAS 95% CL_s upper limits in the $H \rightarrow a^0 a^0 \rightarrow 2\mu 2\tau$ [65], $H \rightarrow a^0 a^0 \rightarrow 4b$ [66], $H \rightarrow a^0 a^0 \rightarrow b\bar{b}\mu^+\mu^-$ [67], $H \rightarrow a^0 a^0 \rightarrow \gamma\gamma jj$ [68], and $H \rightarrow a^0 a^0 \rightarrow \gamma\gamma\gamma\gamma$ [69] channels. Previously, the DØ experiment at the Tevatron set 95% CL upper limits on $H \rightarrow a^0 a^0 \rightarrow 4\mu$ and $H \rightarrow a^0 a^0 \rightarrow \mu^+\mu^-\tau\tau$ in the low m_{a^0} range [70].

CHAPTER 2

THE ATLAS EXPERIMENT AT THE LARGE HADRON COLLIDER

2.1 The Large Hadron Collider

The CERN Large Hadron Collider (LHC) [71, 72] is a proton (p) synchrotron which collides 2808 bunches of up to 1.15×10^{11} protons at a frequency of approximately 4×10^7 Hz to produce a design peak instantaneous luminosity of $1.0 \times 10^{34} \text{ cm}^{-2}\text{s}^{-1}$. The protons are accelerated by superconducting electromagnets through two 27 km circular vacuum tubes, to collide at four interaction points around the collider, where the ATLAS, CMS, Large Hadron Collider Beauty (LHCb) and A Large Ion Collider Experiment (ALICE) experiments are based. Its design centre-of-mass energy is 14 TeV, though it currently operates at 13 TeV. The magnets are kept at a temperature of -269.15°C using a liquid helium based cooling system. Before the Phase 2 upgrades, described in §2.4.1, the LHC is expected to deliver 300 fb^{-1} of data to each of its general purpose particle detectors [73].

Run 2 of the LHC lasted between 2015 and 2018, and resulted in the delivery of 156 fb^{-1} of pp collision data, at a centre-of-mass energy of 13 TeV. Of this data, 147 fb^{-1} was recorded by the ATLAS detector, and 139 fb^{-1} was suitable for physics analysis. The

cumulative luminosity is illustrated in Figure 2.1. This luminosity is measured by measuring the calibrated and transferred signal [74] from the LUCID2 Cerenkov detector [75]. This was cross-checked using complementary methods, such as track counting [74].




Figure 2.1: Cumulative luminosity of data delivered by the LHC, recorded by the ATLAS detector and determined to be suitable for physics analysis [76].

The substantial luminosity of the LHC, combined with the 13 TeV proton-proton inelastic cross section of 78.1 ± 2.9 mb [77] leads to an inelastic collision rate of 1 GHz, and a mean number of inelastic pp interactions significantly affecting the detector at any given bunch crossing (pileup, μ) of around 33.7. The pileup distribution is shown for each year of data taking in Figure 2.2. This presents an experimental challenge, as every interaction of interest will be accompanied by around 33 other inelastic interactions. Furthermore, the high cross section for the production of jets of hadrons presents a large background to any hadronic signatures. Finally, the unknown fraction of the total proton momentum carried by the interacting partons mean that the initial state of each inelastic interaction is unknown. For this reason, some parameters of the events must be treated on a statistical basis, based on previously determined parton distribution functions, introducing associated uncertainties. To deal with these challenges, requirements were placed on the design of the ATLAS detector, which will be detailed in the next section.




Figure 2.2: Luminosity-weighted distribution of the mean number of interactions per bunch crossing, for each year of data taking, and for the combined 2015-2018 dataset [76].

2.2 The ATLAS Detector

The A Toroidal LHC Apparatus (ATLAS) detector [78] is a general purpose particle detector based at the CERN LHC. It is an approximately forward-backward symmetric detector, composed of a barrel region made of cylindrical sub-detectors, and end-caps formed of two series of disk shaped sub-detectors. It is characterised by its three large toroidal magnets, which allow ATLAS to be one of the largest particle detectors ever built: at 44 m in length and 25 m in diameter. ATLAS covers almost the entire solid angle around the interaction point, which is necessary for establishing the amount of missing momentum transverse to the beamline (MET) in any given event. Figure 2.3 shows a cut-away view of the ATLAS detector.

The ATLAS detector was designed with several purposes in mind: the discovery and study of the Higgs boson; high precision tests of the SM, especially the top quark which can now be produced in great abundance thanks to the high centre-of-mass energy collisions; and searches for new physics, including extended Higgs sectors, heavy gauge bosons, Supersymmetric particles such as squarks and gluinos, and exotic signatures such as those arising from models with black holes and extra dimensions. These physics goals




Figure 2.3: Cut-away view of the ATLAS detector [78].

place stringent requirements on the performance of the detector, which are summarised in Table 2.1.

Detector Component	Required Resolution	η Coverage	
		Measurement	Trigger
Tracking	$\sigma_{p_T}/p_T = 0.05\%p_T \oplus 1\%$	± 2.5	
EM Calorimetry	$\sigma_E/E = 10\%/\sqrt{E} \oplus 0.7\%$	± 3.2	± 2.5
Hadronic Calorimetry (Jets)			
Barrel and End-Caps	$\sigma_E/E = 50\%/\sqrt{E} \oplus 3\%$	± 3.2	± 3.2
Forward	$\sigma_E/E = 100\%/\sqrt{E} \oplus 10\%$	$3.1 < \eta < 4.9$	$3.1 < \eta < 4.9$
Muon Spectrometer	$\sigma_{p_T}/p_T = 10\%$ at $p_T = 1$ TeV	± 2.7	± 2.4

Table 2.1: Performance goals of the ATLAS detector [78]. p_T and E are measured in GeV unless stated otherwise.

2.2.1 Conventions

The origins of all coordinate systems at ATLAS are taken to be the interaction point of the pp collisions [78]. The right-handed Cartesian coordinate system about the interaction

point is then defined by an x-direction which points from the interaction point to the centre of the LHC ring, a y-direction which points upwards, and thus a z-direction parallel to the beamline. The polar (θ) and azimuthal (ϕ) angles then take their usual definitions from and around the z-axis, respectively.

The pseudorapidity is defined as $\eta \equiv -\ln \tan(\theta/2)$, which at energies where the particle mass is negligible is approximately equal to the rapidity, as defined in high energy physics: $y \equiv \frac{1}{2} \ln \left(\frac{E+p_L}{E-p_L} \right)$. This is useful because differences in rapidity are Lorentz invariant under boosts along the beam axis, particle production is approximately constant as a function of rapidity, and rapidities are additive, analogously to velocities in Galilean relativity.

Finally, angular separation is generally expressed as: $\Delta R \equiv \sqrt{(\Delta\eta)^2 + (\Delta\phi)^2}$.

2.2.2 Magnets

The ATLAS detector uses a unique magnet configuration, shown in Figure 2.4, to immerse the inner detector and muon spectrometer in a strong magnetic field, in order to measure the charge and momentum of charged particles [79]. A thin, superconducting solenoid of length 5.3 m and inner diameter of 2.44 m immerses the inner detector in magnetic field of up to 2.6 T. An 8-ring air-core superconducting toroid magnet, 25.3 m long and 9.4 m to 20.1 m in diameter, provides a magnetic field of up to 3.9 T to the muon detectors in the barrel region. Two smaller superconducting toroid magnets, each consisting of 8 rings of 5 m in length and 1.65 m to 10.7 m in diameter, supply magnetic fields to the end-caps. To reduce the passive material in the detector, the solenoid shares its cryostat with the liquid argon (LAr) calorimeters. This magnet configuration allows for the separation of the magnetic fields which cover the inner tracker and muon system, providing great freedom in the size of the detector. It is for this reason that ATLAS can have a 25 m diameter, while CMS, the other general purpose LHC experiment, has a diameter of just 15 m. This large diameter allows for large calorimeters, leading to the absorption of a larger fraction of the energy of electrons, photons, and hadrons.




Figure 2.4: Schematic of the ATLAS magnet systems [80].

2.2.3 Inner Tracking Detector

The inner tracking detector [81, 82], shown in Figure 2.5, measures the paths of passing charged particles; determining their charge and momenta by measuring the direction and radii of their curvature, under the applied magnetic field. With approximately 1000 particles being produced every 25 ns, low-latency electronics and sensor elements are required to minimise temporal particle-overlap in the detector components; while high granularity components near the interaction point are required to minimise spatial particle-overlap. Furthermore, high precision tracking close to the interaction point identifies the presence of tracks with a high impact-parameter from the interaction point and secondary vertices, which are used to establish the presence of charm-quarks, bottom-quarks and τ -leptons.

Precision tracking for the ATLAS detector is provided by four layers of silicon pixel detectors close to the beamline, and eight layers of less precise Semiconductor Tracker (SCT) surrounding them, together covering the region $|\eta| < 2.5$. The inner-most pixel layer is known as the Insertable B-Layer (IBL) [84], and is sensitive just 25.7 mm from the beamline, facilitating the efficient detection of high impact-parameter tracks and secondary vertices. These silicon-based detectors track the path of charged particles, by




Figure 2.5: Schematic of the ATLAS inner tracking detector [83].

monitoring the current induced on conducting plates by the movement of electron-hole pairs through an applied electric field. These electron-hole pairs are created in the silicon by the passage of the charged particle. The pixel detectors have approximately 80.4 million readout channels, with a typical resolution of $12\ \mu\text{m}$ in $R\text{-}\phi$ and $66\ (77)\ \mu\text{m}$ in $z\ (R)$ in the barrel (disks). The SCT has 6.3 million readout channels, with a typical resolution of $16\ \mu\text{m}$ in $R\text{-}\phi$ and $580\ \mu\text{m}$ in $z\ (R)$ in the barrel (disks).

Surrounding the SCT is a straw tube Transition Radiation Tracker (TRT) [85], which uses 4 mm diameter drift tubes to track charged particles in the region $|\eta| < 2$. The TRT typically receives around 36 hits per track, through the emission of transition-radiation (X-ray) photons in the xenon-based TRT gas. It does this by measuring the transition radiation emitted by the passing charged particle, as it transitions over the boundaries between the gas-filled straw tubes. It measures the transition radiation by monitoring the movement of electron-ion pairs, produced by pair-production from transition X-rays, in drift tubes. The TRT has about 351000 readout channels, and a typical resolution of $170\ \mu\text{m}$ per straw in $R\text{-}\phi$, but provides no z information. The total energy deposit in the

TRT provides robust separation between electrons and $\pi^\pm$ -mesons, as electrons deposit around 8-10 keV of transition radiation, while $\pi^\pm$ -mesons leave energy deposits of around just 2 keV; this has proven very useful for electron identification. Xenon is chosen as the main component of the gas mixture as it offers good X-ray absorption, while CO_2 and O_2 are present as a quencher [86] to reduce the latency of the detector, by increasing the electron drift velocity. However, due to leakages, the gas mixture now includes Ar.

2.2.4 Calorimeters

Fine-grain liquid argon (LAr) electromagnetic (EM) Calorimeters cover the $|\eta|$ region up to 3.2 to facilitate the identification of electrons and photons [87], as show in Figures 2.6 and 2.7. This calorimeter has a resolution of 0.025×0.0245 in η - ϕ space, with a strip-layer with a resolution of 0.0031×0.0982 in η - ϕ space. This corresponds to a resolution of 4.69 mm in the η direction, which is approximately one tenth of the Molière radius of the EM shower [88], allowing the fine structure of the shower to be probed in this direction. The EM calorimeter is a three-layer lead-liquid-argon detector with accordion-shaped electrodes and lead absorber plates. The accordion-shaped electrodes allow the signal to be read out of the calorimeter without the need for azimuthal cracks, which deteriorate the energy resolution of the calorimeter [87]. Electrons and photons multiply through conversions and bremsstrahlung producing an EM shower on the length scale of the radiation length of the detector, X_0 , until they reach their critical energy. They then deposit their remaining energy, mostly through absorption, leading to the ejection of electrons from the interacting atoms. Figure 2.7 shows that the EM calorimeters are at least 22 radiation lengths deep, allowing them to contain the entire EM shower in almost all cases. The fine granularity of the detector, particularly in the first layer, facilitates the efficient discrimination of electrons and photons from hadrons, which are required for electron and photon identification.

Coarser grain, higher volume hadronic calorimeters surround the EM in the region $1.4 < |\eta| < 4.8$ to reconstruct and measure the energy of jets of hadronic particles




Figure 2.6: Schematic of the ATLAS calorimeter systems [89].




Figure 2.7: Schematic of the ATLAS LAr calorimeter system [78].

and MET in events, and to help veto jets faking electrons or photons. The hadronic calorimeters are tile sampling calorimeters [90] which use steel as the absorber and plastic scintillator as the active material in the barrel, and copper as the absorber and liquid-argon gaps as the active material in the end-caps [87], as shown in Figure 2.6. Together, in the barrel region the calorimeters extend to an outer radius of 4.25 m. Hadrons multiply and deposit their energy through nuclear excitation and break-up caused by strong interactions with the nuclei, and ionisation processes, together known as the hadronic shower; they also tend to have a significant EM core, from π -mesons decaying into photons. These hadronic showers are much more complex than EM showers, and require a complex calibration procedure, as described in §2.3.7.

2.2.5 Muon Systems

The ATLAS Muon Systems (MSs) [78, 91], shown in Figure 2.8, use high precision tracking chambers, amongst the superconducting air-core toroid magnets described in §2.2.2, to track the passage and measure the momenta of charged particles passing through the calorimeters. At $\eta=0$, the calorimeters extend eleven interaction lengths from the interaction point, which is sufficient to prevent the vast majority of interacting particles besides muons from reaching the muon spectrometer, meaning that anything leaving a track in the muon detectors is probably a muon. These systems provide a robust momentum measurement over the pseudorapidity range $|\eta| < 2.7$ for muons with a $p_T < 3$ TeV. At the centre of the detector there is a gap in the muon spectrometers for services to reach the solenoidal magnet, leading to a loss of muon identification in an angular region of up to $|\eta| \leq 0.08$. Muon tracking is provided by Monitored Drift Tubes over the entire pseudorapidity range, which have a precision of 80 μm per tube. For $2 < |\eta| < 2.7$ higher granularity Cathode Strip Chambers are also used due to their superior time-resolution. They have a resolution of 40 μm in the bending plane.

Resistive Plate Chambers and Thin Gap Chambers provide triggering capabilities over the region $|\eta| < 2.4$, within a few tens of nanoseconds of the passage of the muon.




Figure 2.8: Schematic of the ATLAS Muon Systems (MS) subsystems [91].

The Resistive Plate Chambers are used in the barrel ($|\eta| < 1.05$), while the Thin Gap Chambers are used in the end-caps ($1.05 < |\eta| < 2.4$). Together they provide information about the multiplicity and approximate energy range of muons to the L1 trigger system (§2.2.6), and provide coarse tracking information to be used in the High-Level Trigger (HLT) (§2.2.6), for each bunch-crossing. These detectors have a granularity which varies with $|\eta|$ to compensate for the higher muon momentum and radiation levels at high $|\eta|$.

2.2.6 Trigger

The ATLAS trigger system has two independent levels: a hardware Level-1 Trigger; and a software-based High-Level Trigger (HLT) [92–95]. The hardware-based Level-1 Trigger uses a limited subset of coarsened detector information to reduce the event rate from the 40 MHz LHC event rate to 100 kHz. The Level-1 Trigger consists of separate Level-1 calorimeter and muon triggers, which provide information to the Level-1 Central Trigger Processor, which makes the Level-1 trigger decision and manages deadtime. It passes the surviving events, along with regions of interest in η - ϕ space, to the HLT based on a search for high transverse-momentum particles, MET, and displaced vertices. Once seeded by the Level-1 Trigger, the HLT uses full granularity information within the regions of interest to further reduce the rate of event recording down to 1 kHz. Triggers designed specifically for individual analyses are also applied at this stage. Selected events are then exported to the CERN Tier-0 computing facility for offline reconstruction.

2.2.7 ATLAS Software and Physics Simulation

In addition to a very complex detector, ATLAS uses sophisticated online and offline software which is necessary to turn the detector response into physics results [96]. The online software runs during data taking, in order to: define HLT objects for the trigger selection; control the detector; read data off the detector into storage; and display the data in real time to monitor the data as it comes off the detector. The offline software

processes the pre-stored data, in order to: reconstruct, identify and calibrate the various physics objects; and analyse the data.

Offline software also exists to simulate physics process using external Monte-Carlo simulated data (MC) generators, and simulate the response of the detector to these processes, so that it may be compared to data. To simulate physics processes, ATLAS uses a variety of MC generators, including but not limited to: SHERPA [97], PYTHIA [98], HERWIG [99, 100], MADGRAPH [101] and POWHEG [102–104]. The response of the ATLAS detector is then fully simulated [105] using a GEANT4 simulation [106] of the ATLAS detector, or a fast detector simulation [107], using parameterised detector responses to EM and hadronic calorimeter showers in place of fully simulated calorimetry, in conjunction with full simulations of other detector components. The fast simulation results have been validated against fully reconstructed samples. After the detector response has been simulated, but before any reconstruction algorithms are applied, pileup-induced detector activity is overlaid. The samples are then reweighted to reproduce the mean number of interactions per bunch crossing observed in the data. Finally, these simulated samples undergo the same reconstruction and identification procedures as applied to the data, and are processed using the same triggers and event selection.

This author developed an automated system to monitor electron and photon isolation variables. This was incorporated into the ATLAS ATHENA software framework, and is used to monitor the impact of software changes on the isolation variables. An example isolation variable plot is shown in Figure 2.9. This variable is the energy in the EM calorimeter, within ΔR of 0.2 around the photon, excluding the energy assigned to the photon itself.

2.2.8 Data Taking

The ATLAS detector is operated by a team of 8 shifters in the ATLAS Control Room. These include a shifter in charge of the triggers, inner detectors, calorimeters, and muons, in addition to a Run Control shifter, a Shift Leader in Matters of Safety, a Data Quality




Figure 2.9: Example isolation variable produced using automated electron and photon isolation monitoring tool described in the text. This variable is the energy in the EM calorimeter, within ΔR of 0.2 around the photon, excluding the energy assigned to the photon itself.

shifter, and the Shift Leader, respectively. There are also on-call expert shifters, who can be called if a more complex issue with the respective subsystem occurs.

The Shift Leader is the main responsible in the ATLAS Control Room, supervising all of the other activities, and remaining in contact with the CERN LHC control centre. The Run Control Shifter, ensures that the data taking runs smoothly, and starts and stops the physics runs. The Shift Leader in Matters of Safety is in charge of all safety-related aspects of the experiment, including granting cavern access. The various subdetector shifters must monitor the performance of their subdetectors, in order to ensure they are operational during data-taking. In addition to this, the trigger shifter must define the trigger prescale, to ensure the smooth operation of the readout system.

This author performed shifts as a Data Quality shifter in the ATLAS Control Room, checking the quality of the data for any issues during data-taking. The shifter was expected to monitor various properties of the incoming data using a variety of tools. Such properties include the luminosity, trigger rates, computing infrastructure, event displays, low-level and high-level reconstructed physics variables, and the data quality monitoring tools themselves. The shifter was expected to monitor plots of various reconstructed

physics variables update in real time, and look for any deviation from a set of reference histograms, such as the ones shown in Figure 2.10.




Figure 2.10: Example Data Quality Shifter reference histogram. Plot from ATLAS Data Quality Shifter Reference Histograms [108].

2.3 Physics Objects

The following subsections briefly describe the algorithms used to reconstruct and identify various physics objects relevant to the analyses described herein. ATLAS also reconstructs τ -leptons and MET, which are briefly described for completeness.

2.3.1 Tracks

Charged particles leave tracks in the inner detector, which are used as inputs to various higher-level physics objects, such as electrons, *converted* photons, muons and τ -leptons, and are sometimes used directly in physics analyses [83]. They are also used to form track-jets, analyse jet substructure, and perform flavour tagging.

The charged object leaves discrete hits in the silicon pixel and strip detectors, which are associated with each other to form clusters. These clusters are identified as *single* or

merged by an artificial neural network, depending on whether the reconstruction determines that they were likely produced from the induced charge of one or more charged particles. Tracks are then seeded from these clusters using three space points. This maximises the number of possible combinations, while allowing a first momentum measurement. Track seeds formed entirely from clusters in the SCT have the fewest fake signatures, followed by the pixels, and then clusters from both the SCT and the pixels. Requirements are placed on the momentum and impact parameters of the seeds, and on the compatibility with other clusters of the seeds.

A combinatorial Kalman filter [109] is then used to form track candidates by extrapolating from these seeds. Multiple track candidates can be formed from the same seed. Tracks are then ranked based on the χ^2 of the fit, the number of clusters, the momentum of the track candidates, and the absence of sensitive detector elements which intersect the track but do not have a cluster. The most highly ranked tracks are then prioritised in the reconstruction. Clusters can contribute to at most two tracks, and tracks can share at most two clusters, and clusters are assigned to tracks based on the aforementioned ranking. Track configurations in which *single* clusters are matched to multiple track candidates are penalised. Track candidates are then required to have: $p_T > 400$ MeV; $|\eta| < 2.5$; ≥ 7 silicon hits; maximum of one shared pixel cluster or two shared SCT clusters on the same layer; ≤ 2 silicon holes; ≤ 1 pixel holes; impact parameters compatible with the primary vertex.

The magnitudes of the momenta of tracks are taken from their curvature in the applied magnetic field, evaluated by a measurement of the sagitta of the tracks, and their directions of curvature are used to identify the sign of the charge of the particles. The track reconstruction efficiencies, as evaluated for 2015 conditions, as a function of η and p_T are shown in Figure 2.11.




Figure 2.11: Predicted track reconstruction efficiencies, as evaluated in MC for 2015 conditions, as a function of (a) η and (b) p_T [110].

2.3.2 Interaction Vertices

Interaction vertices are reconstructed using an iterative procedure, in which the z -positions of tracks at the beamline are used to seed vertices, and an iterative χ^2 fit is used to match nearby tracks to the seeds [111]. Tracks which are more than 7σ from any vertex are used to seed a new vertex, and the procedure is repeated until no additional vertices are found. Vertices are required to contain at least two tracks, and the position of the beam spot is applied as a 3-dimensional constraint on the vertex formation. The primary vertex is then chosen as the vertex with the largest sum of squared transverse momenta of the association tracks. The vertex reconstruction efficiency is shown in Figure 2.12 for low pileup data and MC.

Tracks are often required to be compatible with the primary vertex, to ensure that they originate from the pp collision of interest. For this, requirements are placed on the transverse and longitudinal impact parameters, and their significances. z_0 is the longitudinal impact parameter of a track from the primary vertex, though $z_0 \sin \theta$ is often used for this purpose. d_0 is the transverse impact parameter of a track from the primary vertex, though d_0/σ_{d_0} is often used for this purpose, where σ_{d_0} is the uncertainty on d_0 .




Figure 2.12: Vertex reconstruction efficiency for low pileup data and MC [112].

2.3.3 Electrons

In all of this thesis except Chapter 6, electron candidates are reconstructed by a four part process [113–116] as described in this subsection. The differences relevant to Chapter 6 are described in the next subsection. First, a sliding-window of size 3×5 , in units of the granularity of the EM calorimeter, is translated around the EM calorimeter cells searching for energy seed-clusters of $E_T > 2.5$ GeV, removing duplicates which have a large window overlap with already reconstructed electron candidates. Second, tracks are reconstructed under the assumption that the particle leaving the track is a $\pi^\pm$ -meson. If a track is not found with at least 3 hits in the pixel detectors and $E_T > 1$ GeV that points to a seed-cluster in the calorimeter, an electron hypothesis which allows for energy loss of up to a 30% at each material surface is employed instead. Third, extrapolated tracks are matched to the EM calorimeter seed-clusters, and basic selection requirements are applied to the potential candidates to ensure the track came from the primary vertex, and to veto photon conversions. If multiple tracks are matched to a cluster, one of them is designated as the primary track, based on its quality and distance from the seed-cluster. Finally, the momentum of the electron is calculated from the energy deposited in

the EM calorimeter and the direction of the track, then the reconstruction procedure is completed by calibrating the candidate energy to the energy of the generator-level electron as determined by multivariate techniques applied to MC simulated data. Figure 2.13 shows a reconstructed electron candidate, and the various detector components relevant to its reconstruction and identification.




Figure 2.13: Electron candidate reconstruction at the ATLAS detector [113].

Variables are defined based on the shape of the EM shower, the leakage into the hadronic calorimeter, and the response of the TRT, which is designed with electron identification in mind. These variables are input to a likelihood-based multivariate discriminant, which is used to identify electrons against backgrounds from hadrons and photon conversions. The three identification levels are then defined to be subsets of each other, *Loose*, *Medium* and *Tight*, and the identification menu is optimised in bins of E_T and η . Electrons which are reconstructed in the transition region between the barrel and end-caps, $1.37 < |\eta| < 1.52$, have a much higher fake rate due to the large amount of material from the carbon-fibre support structures in this region. The (predicted) electron reconstruction and identification efficiencies are shown in Figures 2.14 and 2.15, respectively. The reconstruction, identification and isolation efficiencies, as well as the momentum, of these electrons are calibrated using data, from $Z \rightarrow e^+e^-$ events.




Figure 2.14: Predicted cluster, track, cluster and track, and reconstruction efficiencies for electrons as a function of the generator-level E_T , as evaluated for 2015-2016 conditions [113].




Figure 2.15: Electron identification efficiencies, as evaluated using 2015-2016 data, as a function of (a) η and (b) p_T [113]. The rise of the *Loose* and *Medium* identification efficiencies towards low E_T is due to a mismodelling of the variables used in the identification in the MC samples used to optimise the discriminant cuts.

Electrons are also reconstructed in the forward region, outside the acceptance of the inner tracker, based only on the calorimeter information. Due to the greater material in the forward region, and the absence of a track, these have much larger backgrounds, and are almost indistinguishable from photons.

2.3.4 Electromagnetic Cluster Reconstruction Algorithm

For the search in Chapter 6, the algorithm used to reconstruct EM clusters in ATLAS was changed relative to the algorithm used elsewhere in this thesis. This new algorithm uses dynamic, variable-size superclusters [117]. This allows the cluster size to adapt to recover energy from bremsstrahlung or photon conversions. First, topological clusters in the EM calorimeter are seeded by cells with readout energy magnitude at least four times higher than the expected noise in the cell from electronic noise and pileup. The presampler and first LAr EM layer are excluded from this first step, while in the region $1.37 < |\eta| < 1.63$ cells in the presampler and scintillator between the calorimeter cryostats are also used. Second, these seeds collect neighbouring cells with readout energy magnitude at least two times higher than the expected noise, and then any cells next to these collected cells are considered in the same way. This process continues, with any clusters sharing a cell being merged, until all the neighbouring cells with sufficient readout energy magnitude have been collected, after which one final layer of cells is added without a requirement on the energy. Third, clusters with two local maxima with cell energy above 500 MeV are split into separate clusters. Finally, any clusters with total energy below 400 MeV, or EM cluster energy less than half that of the corresponding cluster formed from both the EM and hadronic calorimeters, are discarded.

2.3.5 Photons

Photons are reconstructed in much the same way as electrons, using clusters of energy in the EM calorimeter, but with different track requirements. As with electrons, the recon-

struction algorithm was changed as per the last subsection, though this new algorithm is not used anywhere in this thesis. Photons can be broadly categorised as *unconverted* photons, which reach the calorimeter as photons, or as *converted* photons, which convert to electron-positron pairs before reaching the calorimeter. *Unconverted* photons are required to have no track matched to them, while *converted* photons are required to have a pair of tracks which form a vertex compatible with the decay of a massless particle. As with electrons, likelihood-based identification criteria are then applied, but without any information from the TRT. These criteria differ for *converted* and *unconverted* photons due to the effect of the magnetic field on the direction of the electrons, and thereby the shape of the EM showers. For *unconverted* photons the momentum is taken from the EM calorimeter measurement, while for *converted* photons the magnitude of the momentum is taken from the calorimeter measurement, and the directionality is taken from the conversion tracks. The predicted *converted* photon reconstruction efficiency is shown in Figure 2.16, and the predicted *unconverted* and *converted* photon identification efficiencies are shown in Figure 2.17.




Figure 2.16: Predicted reconstruction efficiency for *converted* photons as a function of the generator-level E_T , as evaluated for 2015-2017 conditions [117]. The open (full) markers represent the previous (current) reconstruction methods.




Figure 2.17: Photon identification efficiencies, as evaluated using 2015-2016 data and MC, for (a) *unconverted* and (b) *converted* photons [118] with $|\eta| < 0.6$. These results use the EM cluster reconstruction method described in subsection 2.3.3.

2.3.6 Muons

In ATLAS, muons are reconstructed in five different categories, depending on the location in the detector where they are detected. These are listed below:

- **Combined muons:** muons reconstructed when a track in the ID is matched to a track in the MS [119]. The momentum of these muons is established by combining the measurements from the ID and the MS. The vast majority of signal muons are of this type.
- **Segment tagged muons:** muons reconstructed when a track in the ID is matched to a partially reconstructed track in the MS. The momentum of these muons is taken from the ID measurement.
- **Stand alone muons:** muons reconstructed as a track in the MS, not associated with the ID. This occurs for muons with $|\eta|$ in the region 2.5 to 2.7, beyond the ID acceptance but within the MS acceptance.
- **Silicon associated forward muons:** muons reconstructed as a track in the MS,

associated with a partial track in the ID. This occurs for muons with $|\eta|$ around 2.5, only partially in contact with the ID.

- **Calorimeter muons:** muons identified when a track in the ID with $p_T > 15$ GeV is matched to a calorimeter deposit, which is compatible with the energy deposit of a minimal ionising particle. This occurs in the central barrel region, $|\eta| \leq 0.1$, where there is a gap in the active region of the MS.

The muon reconstruction efficiencies are shown in Figure 2.18. *Loose* muons are defined as any of the above muon types, with calorimeter and segment tagged muons required to have $|\eta| < 0.1$ [120]. *Medium* muons only include combined, stand alone, and silicon associated forward muons, to minimise the systematic uncertainties associated with muon reconstruction and calibration [120]. *Tight* muons are combined muons which have additional reconstruction requirements applied to maximise their purity, as described in Ref. [120]. Further requirements are typically applied on the transverse and longitudinal impact parameters of muons to reject muons from fake inner detector tracks, and from cosmic rays. The reconstruction, isolation and track-to-vertex association efficiencies, as well as the momentum, of these muons are calibrated using data from $Z \rightarrow \mu^+ \mu^-$ events.

2.3.7 Jets

Hadronisation of final state quarks and gluons from hard interactions, or from final state partons from the underlying event, lead to jets of highly collimated hadrons. There are reconstructed as topological energy clusters in the calorimeters [88, 121], using the anti- k_t jet algorithm [122], with a typical distance parameter of 0.4. Other jet algorithms and distance parameters are used, though not in any of the analyses presented herein. The resulting jet is used as a proxy for the initial state quark or gluon.

The jet energies are calibrated using p_T and η dependent scale factors, followed by corrections [123] based on the jet-area. Further corrections from internal jet properties, in situ measurements, and muons found within jets are also applied. All jets must have




Figure 2.18: Muon reconstruction efficiencies, as evaluated using 2015 data, for (a) *Loose* and *Medium* and (b) *Tight* muons [120].

$p_T > 20$ GeV, and be reconstructed inside the EM calorimeter, which requires that $|\eta| < 2.5$. The total jet energy scale uncertainty is shown in Figure 2.19.




Figure 2.19: Fractional jet energy scale uncertainty and its components, evaluated using 2015-2017 data, as a function of (a) η (b) p_T [124].

2.3.8 Heavy Flavour Hadron Tagging

Hadronic jets which are initiated by the hadronisation of the c -quark or b -quarks have three main properties which allow them to be distinguished from jets initiated by light flavour quarks, or gluons. First, the large lifetime of the heavy-flavour hadron causes it to decay away from the primary vertex, producing displaced tracks. These displaced tracks can be extrapolated backwards, and their impact parameter with respect to the primary vertex can be used as a measure of how displaced the track is. The transverse and longitudinal impact parameters are considered separately, as pileup and the magnetic field affect them differently. Second, in the event that multiple displaced tracks are found, they can often be used to reconstruct a displaced vertex. This then provides a wealth of information, such as the mass of the vertex, the decay length of the heavy flavour hadron, the number of decay products, and the energy carried by the decaying hadron. Third, as b -jets primarily decay through c -jets, tertiary vertices can often be reconstructed, sometimes by matching single tracks to the axis of the calorimeter jet. This full decay chain is difficult to reconstruct without the presence of a heavy-flavour hadron. Finally, these variables are all given to at least one boosted decision tree (BDT), which is designed to separate b -jets, c -jets and light-flavour jets from each other. Nearby low- p_T muons are sometimes also used to tag the presence of a heavy flavour hadrons. The transverse impact parameter significance, and the fraction of energy of all tracks reconstructed in the jet represented by the tracks from a displaced vertex, are shown in Figure 2.20.

A c -jet identification algorithm is used for the search in Chapter 4, which consists of two BDTs: one to select c -jets against a background of b -jets, where b -hadrons have longer lifetimes; and another to select c -jets against a background of light-jets, where light-flavour hadrons usually have shorter lifetimes. The effect of the intermediate lifetime of c -hadrons makes c -jet identification a difficult task, as most of the discriminant variables have distributions which are between those of b -jet and light-jets, as shown in Figure 2.20. The presence of tertiary vertices can be used to veto b -jets during c -jet tagging. The distributions of c -jets, b -jets and light-jets in the space of these two BDTs is shown in





Figure 2.20: Distributions of the (a) transverse impact parameter significance and (b) fraction of energy of all tracks reconstructed in the jet represented by the tracks from a displaced vertex, as determined in $t\bar{t}$ simulation [125].


Figure 2.21, along with the efficiency point adopted for the search in Chapter 4. The predicted b -jet and light-jet rejection of this are shown for three efficiency points of this algorithm in Figure 2.22.

2.3.9 Missing Transverse Momentum


The missing momentum transverse to the beamline (MET) can be used to reconstruct particles which escape the detector, such as neutrinos [126]. Several such algorithms are used in ATLAS, some using the information from the inner tracker, others using information from the calorimeters, and others using both. When both are used, procedures are applied to avoid objects being double counted. The missing transverse momentum is calculated as the negative vector sum of the three-momenta transverse to the beamline of the objects in the detector. Two sources of three-momentum which contribute to the calculation of the missing transverse momentum are distinguished: those from hard objects, defined as electrons, photons, muons, τ -leptons and jets; and those from soft



(a)



(b)



(c)

Figure 2.21: Distributions of (a) c -jets, (b) b -jets and (c) light-jets in the space of the two BDTs used to tag c -jets against backgrounds of b -jets and light-jets [7], as determined in simulation. The efficiency point adopted for the search in Chapter 4 is shown by the solid black line.




Figure 2.22: Predicted c -tagging efficiency as a function of b -jet and light-jet rejection, obtained from $t\bar{t}$ simulation. The cross designates the efficiency point chosen for the analysis described in Chapter 4 [7].

objects, defined as any signal in the detector not associated with those high- p_T objects. A dedicated calibration is performed for each of the two types of missing transverse momentum.

2.3.10 τ -Leptons

τ -leptons which decay leptonically provide a signature of a charged lepton and missing transverse momentum. Conversely, τ -leptons which decay hadronically require a dedicated reconstruction procedure [127]. First, they are seeded using an anti- k_t jet, with a radius parameter of 0.4, and required to have $p_T > 10$ GeV and $|\eta| < 2.5$, but not $1.37 < |\eta| < 1.52$. Track-based reconstruction and isolation requirements are applied, followed by identification requirements. The identification requirement uses two Boosted Decision Trees, one trained for one-track τ -lepton decays, and the other for three-track τ -lepton decays. *Loose*, *Medium* and *Tight* identification efficiency points are defined. A dedicated calibration is applied to hadronic decays of τ -leptons.

2.3.11 Overlap Removal

To avoid double counting physics objects in the event, overlapping electrons, muons and jets are removed [128]. The procedure is outlined as follows:

1. Any electron sharing an inner detector track with a muon is removed from the event.
2. Any jets within $\Delta R < 0.2$ of electrons are removed.
3. Electrons within a ΔR of $\min(0.4, 0.04 + 10 \text{ GeV}/p_T^{\text{electron}})$ of jets which satisfy $JVT > 0.59$ [129], $p_T < 60$ GeV and $|\eta| < 2.4$ are removed. The jet vertex tagger variable, JVT , is described in Ref. [129].
4. Remove jets with $\Delta R < 0.2$ of muons if the jets have less than three associated tracks, or if the muon has 70% of the associated track momentum and less than 50% of the jet momentum

5. Muons with $\Delta R < \min(0.4, 0.04 + 10 \text{ GeV}/p_T^{\text{muon}})$ of jets passing $JVT > 0.59$ [129], $p_T < 60 \text{ GeV}$ and $|\eta| < 2.4$ are removed.

2.4 High Luminosity Upgrades

2.4.1 The High Luminosity LHC

By 2025 the LHC is expected to have undergone a major upgrade, increasing its instantaneous luminosity from $2 \times 10^{34} \text{ cm}^{-2}\text{s}^{-1}$ to $5\text{--}7.5 \times 10^{34} \text{ cm}^{-2}\text{s}^{-1}$, at which point it will be known as the High-Luminosity Large Hadron Collider (HL-LHC) [73, 130]. The HL-LHC aims to deliver $3000\text{--}4000 \text{ fb}^{-1}$ of 14 TeV pp collision data to each of its general purpose detectors by 2038, while the LHC is expected to deliver 300 fb^{-1} operating at 13 TeV or 14 TeV. However, this will result in the average number of interactions per bunch crossing increasing from 33.7 to 140–200. Figure 2.23 shows a $t\bar{t}$ event at a pileup of 200, to illustrate the pileup conditions. The proposed programme for the upgrade is shown in Figure 2.24.




Figure 2.23: $t\bar{t}$ event in the $\mu = 200$ pileup conditions, expected to the HL-LHC [131].

This increase in data will not just result in a major reduction to the statistical uncertainties in physics analyses, but also in the statistical component of systematic un-




Figure 2.24: HL-LHC upgrade programme [132].

certainties, due to the greater availability of calibration data. Many analyses will also have access to high-statistics control regions which can be used to constrain some of their backgrounds. Lastly, with this additional data, parton distribution function uncertainties are expected to be reduced by a factor of 2-4 for invariant masses above 100 GeV [133].

2.4.2 ATLAS Upgrades For the HL-LHC

As described in the previous subsection, the pileup rate is expected to increase by a factor of 4-6. This will place stringent demands on the radiation hardness and pileup robustness of the detector, requiring major upgrades [130] to the inner tracking detector and the trigger. The aim of the upgrades to the ATLAS detector will be to maintain the current performance in these conditions. The greater luminosity also increases the radiation damage by up to a factor of 10, meaning about 2×10^{16} n-eq/cm² at the innermost pixel layer, over the lifetime of the detector. Where 1 n-eq is equivalent to the damage caused by a single 1 MeV neutron. For this reason, stringent requirements must be met regarding the radiation hardness of the detector, especially near the interaction point.

Due to the expected radiation damage it will have received, the entire inner detector will be replaced with an all-silicon Inner Tracker (ITk). The increased resolution of the

ITk will also be required to efficiently reconstruct tracks in the high-pileup environment of the HL-LHC. The ITk will fill all the available space within the calorimeters, with a length of 6 m and an active radius of 1 m. This allows the pixel layers of the ITk to cover the region $|\eta| \leq 4.0$. In the central barrel region, the ITk will have 5 silicon pixel layers and 4 silicon strip layers, 2 of which will have longer strips than the other 2. In the end-caps, it will have 5 silicon pixel rings and 6 silicon strip rings. Figure 2.25 shows the intended layout of the ITk.




Figure 2.25: Nominal ITk layout [134], at the time of writing, with *Inclined* pixel configuration.

The trigger system will be redesigned, with the introduction of a new 2-stage hardware trigger being considered at present. The L0 and L1 triggers are being designed to operate at rates of up to 1 MHz and 400 kHz respectively. The new HLT will use multithreading to be able to operate at rates of up to 10 kHz. New Feature Extractors [135] will be installed to allow off-line style event reconstruction at the trigger level. Lastly, a new Fast Tracker [136] system will be used to deliver full event track reconstruction at an earlier stage in the trigger selection than ever before, providing more efficient event rejection.

The forward calorimeters will have significant upgrades to their readout electronics, giving access to the full granularity of the calorimeters [137, 138]. The muon systems will also be upgraded with the introduction of the New Small Wheel [139], which will increase

the muon fake rate at L0 of the trigger, and improve muon reconstruction and momentum resolution. Lastly, a thin silicon pixel layer with a 30 ps time resolution, called the High Granularity Timing Detector [140], will be added in the region $2.4 < |\eta| < 4$ to provide time-based separation of pileup vertices.

2.4.3 Potential HL-LHC Inner Tracker Layouts

During the first years of this PhD programme, multiple Inner Tracker (ITk) layouts were under consideration. These were studied as part of the HL-LHC electron performance work of Chapter 3, and are briefly described below. These layouts have since been superseded by the layout in Figure 2.25, but are nonetheless given below as they are used for the studies in Chapter 3.

Extended Layout

The *Extended* ITk layout is the more standard of the two ITk layouts which were under consideration at the time of the studies for Chapter 3, with all of the pixels parallel to the beamline as shown in Figure 2.26. Besides its simplicity, the advantage of this layout is that charged particles which leave deposits at high $|\eta|$ transverse more active material, leaving a short track stub in each layer. These stubs provide more information than the simple hits left at low $|\eta|$, which can improve the reconstruction of the track direction.




Figure 2.26: ITk pixel layout [130], with *Extended* pixel configuration.

Inclined Layout

The layout which is taken as the nominal layout at the time of writing is called the *Inclined* layout, shown in Figure 2.27. In this layout, some of the pixel layers are *inclined* to face the beamspot, reducing the material which particles must transverse and the total silicon required to cover the full solid angle about the beamspot. The sensors in the *inclined* layout also provide two or more hits in the first layer, improving track reconstruction. The *Inclined* layout is preferred because the reduced material transversed leads to better impact parameter resolution in the forward region, and reduced showering of EM and hadronic particles before they hit the calorimeters.




Figure 2.27: ITk pixel layout [130], with *Inclined* pixel configuration.

Previously Considered Inner Tracker Layouts

A previous ITk layout which was also studied had 4 pixel and 5 strip layers, as shown Figure 2.28. This was changed due to general tracking performance improvements, but in particular due to improvements in the ability to resolve close-by tracks. This has direct applications for τ -lepton reconstruction and track-based jet substructure measurements.

2.4.4 Silicon Strip Sensor Testing for the Inner Tracker

This author made direct contributions to the development of the ITk, by testing prototype silicon strip sensors. These were irradiated using the University of Birmingham 27 MeV






Figure 2.28: Old ITk layout [130], with one fewer silicon pixel layer and one more silicon strip layer.

proton cyclotron [141], and then annealed for the standard time of 80 mins at 60°C. The collected charge and cluster width were plotted against the bias voltage, before the irradiation, after the irradiation, and after the annealing, at different temperatures. The charge collected was found to decrease, and the cluster widths were found to increase, after irradiation to 6×10^{14} n-eq/cm², as shown in Figures 2.29 and 2.30, respectively. The charge collected was found to increase and the leakage current was found to increase after annealing, which were not expected, as the annealing process is expected to repair some of the defects introduced by the irradiation process. These results are believed to be due to an issue with the cyclotron.




(a) Unirradiated




(b) Irradiated to 6×10^{14} n-eq/cm²

Figure 2.29: Charge collected against bias voltage, for ATLAS 12 silicon strip sensor prototypes (a) unirradiated and (b) irradiated to 6×10^{14} n-eq/cm².



(a) Unirradiated



(b) Irradiated to 6×10^{14} n-eq/cm²

Figure 2.30: Cluster width against bias voltage, for ATLAS 12 silicon strip sensor prototypes (a) unirradiated and (b) irradiated to 6×10^{14} n-eq/cm².

CHAPTER 3

EXPECTED ELECTRON AND PHOTON PERFORMANCE OF ATLAS AT THE HL-LHC

3.1 Introduction

The upgrade of the LHC to the HL-LHC will lead to an increased pileup of up to $\mu=200$, for which the ATLAS detector will receive substantial upgrades, including the replacement of the current inner detector with the new all-silicon ITk, as described in Section 2.4. This means that the detector will lose its TRT, which is currently used to identify electrons through the detection of transition radiation X-ray photons. The higher pileup, combined with the loss of the TRT raises concerns over how effectively the post-upgrade ATLAS detector will be able to identify electrons. Therefore, a study was conducted to identify pileup-robust variables, and use them to design and assess a menu of three electron identification efficiency points for the HL-LHC. This study was performed for multiple candidate ITk layouts to inform the design of the ITk. This study forms the basis of the following chapter, and the results for the electron identification were published in Refs [3, 4].

MC samples have been generated for this study, which simulate these conditions for

various pileup scenarios ($\mu=70-90$, $\mu=130-150$ and $\mu=190-210$), and for multiple ITk layouts. Electron candidates in the central region, $|\eta| < 2.47$, are reconstructed from these samples by matching energy deposits in the EM calorimeter to tracks in the ITk [114, 115], as described in §2.3.3. Prior to designing the identification menu, a set of pre-identification requirements was identified and applied to the reconstructed electron candidates, to remove any candidates for which the full identification procedure could not be meaningfully applied. Electrons are then selected in 3 identification efficiency points: *Loose*, *Medium* and *Tight*. These have signal identification efficiencies of about 95%, 90% and 70% respectively, against a background of hadrons faking electrons (fakes) and photons converting to electrons via pair-production (conversions). These categories are defined using single sided requirements on individual variables. The variables use three different physical principles: tracking, shower shapes, and the level of matching between tracks and showers. These requirements are varied in transverse energy, E_T , and pseudorapidity, $|\eta|$, to keep the efficiency fixed at the desired level.

The pre-identification and identification efficiencies and fake probabilities are evaluated for each efficiency point of the menu. The charge misidentification probabilities and energy resolution are then evaluated. Short studies are performed into the effect of different pileup scenarios, and the different ITk layouts being considered, on the performance of the menu. Finally, 1 photon identification efficiency point is optimised and evaluated using methods similar to those for electrons.

3.2 Backgrounds

3.2.1 Jet Fakes

The main background to electrons at ATLAS are jets of hadrons faking electrons. Given that a reconstructed electron candidate is a track matched to a calorimeter energy cluster, practically any jet containing a charged particle will be a candidate electron. Most of these

jet fakes are easily rejected by the identification algorithm, because hadrons generally leave the majority of their energy in the hadronic calorimeter, because they tend to have a longer shower which starts later. However, some jets have one or more neutral pions which decay to photons near the EM calorimeter, which then shower electromagnetically, leaving a signature almost identical to an electron. Others simply leave very little energy in the hadronic calorimeter, and thus have to be rejected purely on the basis of the shower shapes. Shower shape based rejection is possible because hadronic showers consist of a wider range of interactions, some of which are high energy and inelastic due to nuclear break-up, causing wider opening angles between the outgoing particles, meaning broader showers. The larger penetration depth of hadrons means that the hadronic showers are usually longer as well, which can offer discrimination through the longitudinal shower profile, and the leakage of the shower into the hadronic calorimeters.

3.2.2 Photon Conversions

A second background to electron identification is conversions of photons to an electron-positron pair by pair-production, in the EM field of the nucleus. The nucleus is required to conserve momentum, which means that conversions are most likely to occur in regions of high matter density. This is one reason why it is sensible to minimise the amount of matter in the innermost layers of the detector. It is also why conversions are most problematic in the high material regions of the detector, which for ATLAS are the transition region between the barrel and the end-caps, and to a lesser extent the end-caps themselves. If a conversion happens at small radius, the electron positron pair produced will leave a complete track, mimicking the signature of an electron which originated from the primary vertex. Furthermore, if the momentum of the converted e^+e^- system is high, then the magnetic field in the ITk will not be able to separate them sufficiently for their showers to be distinguished from that of a single electron. However, frequently the conversion happens later in the detector, and therefore these backgrounds can be rejected by requiring that the electron candidate has left a track segment in the innermost part of the ITk.

Only photons which convert in the beam pipe or the first silicon layer can produce this signature.

3.2.3 Heavy Flavour Decays

Physics analyses which use electrons, are usually interested in electrons which are promptly produced from the hard pp interaction. However, electrons are often produced in secondary decays of c - and b -hadrons, which have identical EM showers to prompt electrons, except that they often overlap with the showers of the other decay products of the heavy flavour hadron. They can therefore be rejected on the basis of their shower shapes, or by using isolation or transverse energy requirements.

3.3 Method

3.3.1 Simulated Samples

The MC signal samples used for this study were $Z \rightarrow e^+e^-$ samples generated using POWHEG [102–104] with the CT10 PDF set [142, 143], interfaced to PYTHIA 8 [98] with the AU2 [144] set of tuning parameters for the modelling of the parton shower. The MC background samples used are dijet samples generated using PYTHIA 8. The samples are normalised to the respective theoretical cross section, with any available higher order corrections applied. The upgraded ATLAS detector is fully simulated [105] using GEANT4 [106].

3.3.2 Discriminant Variables

The twelve variables used in the identification menu, are summarised by efficiency point in Table 3.1. The *Loose* and *Medium* efficiency points use the same variables. The purely tracking variables are able to veto conversion backgrounds, while the calorimetric

and track-calorimeter matching variables primarily veto fakes, but can also discriminate against conversions. The variables are then described in order of the distance of the relevant sub-detector from the beamline in the barrel region of the detector.

Starting with the variables defined solely in the ITk, n_{BL} , which is shown in Figure 3.1, is highly discriminant against conversions because the only opportunity a photon has to convert before it reaches the first pixel layer is in the beam-pipe. n_{2nd} , shown in Figure 3.2, is similar to n_{BL} except that the photon can also convert in the first pixel layer; this variable is only used for the *Tight* efficiency point. Electrons will sometimes traverse a silicon layer without leaving a hit, thus these variables are not perfectly efficient, which is why the more flexible n_{Pix} , shown in Figure 3.3, is also used for all efficiency points. A requirement of at least 1 hit in the first pixel layer, and at least 3 in the pixel detector is required for all efficiency points; while the *Tight* efficiency point also requires at least 1 hit in the second inner-most pixel layer.

The only matching variable used is $\Delta\eta_1$, which is shown in Figure 3.4. This variable is highly discriminant against fakes, because only the charged component of the jet leaves a track in the ITk, but this component will often not be the main cause of the energy deposit in the EM calorimeter. This means that there will be significant displacement between the dominant component of the energy deposit and the projected track. Then if one hadron leaves most of the hits in the ITk, but the other leaves the bulk of the energy deposit in the EM calorimeter, $\Delta\eta_1$ should be able to effectively veto them. This variable uses the full extent of the new ITk, and so is expected to improve with the upgrade.

E_{Ratio} and $w_{s,Tot1}$, shown in Figures 3.5 and 3.6 respectively, are both shower shape variables defined in the strip layer, the first layer, of the EM calorimeter. They are both measures of the spread of the energy deposit, and provide effective discrimination against fakes because typically EM showers are narrower than hadronic showers. They are also expected to suffer reduced effectiveness at high pileup, due to the increased probability of overlapping with particles from pileup interactions.

Similarly, w_{η_2} , R_η and R_ϕ , shown in Figures 3.7, 3.8 and 3.9 respectively, are shower

Type	Description	Name
Pre-Identification Requirements		
Tracking	Number of hits in the pixel detector	n_{Pix}
	Number of hits in the silicon detector	n_{Si}
	Impact parameter of the track to the beam-spot	$z_0 \sin \theta$
	Transverse impact parameter of the track to the beam-spot	d_0/σ_{d_0}
Strip layer of EM calorimeter	Ratio of energy in the strip layer to energy in the whole EM calorimeter	f_1
<i>Loose</i> (contains parameters of Pre-identification requirements)		
Tracking	Number of hits in the innermost pixel layer	n_{BL}
Track-cluster matching	$ \Delta\eta $ between the cluster position in the strip layer and the extrapolated track	$ \Delta\eta_1 $
Strip layer of EM calorimeter	Shower width $\sqrt{(\sum E_i(i - i_{max})^2)/(\sum E_i)}$, where i runs over all strips in a window of $\Delta\eta \times \Delta\phi \approx 0.0625 \times 0.2$, where i_{max} is the index of the highest-energy strip	$w_{s,Tot1}$
	Ratio of the energy difference between the largest and second largest energy deposits in the cluster, divided by their sum	E_{Ratio}
Middle layer of EM calorimeter	Lateral shower width, $\sqrt{(\sum E_i \eta_i^2)/(\sum E_i) - ((\sum E_i \eta_i)/(\sum E_i))^2}$, where the sum is calculated within a window of 3×5 cells	$w_{\eta 2}$
	Ratio of the energy in $\eta \times \phi$ of 3×7 cells and 7×7 cells, centred on the energy cluster	R_η
	Ratio of the energy in $\eta \times \phi$ of 3×3 cells and 3×7 cells, centred on the energy cluster	R_ϕ
Back layer of EM calorimeter	Ratio of the energy in the back layer to the total energy in the EM calorimeter	f_3
Hadronic leakage	Ratio of E_T in the hadronic calorimeter to E_T of the EM cluster (used over the range $0.8 < \eta < 1.37$)	R_{had}
	Ratio of E_T in the first layer of the hadronic calorimeter to E_T of the EM cluster (used over the range $ \eta < 0.8$ and $ \eta > 1.37$)	$R_{had,1}$
<i>Medium</i> (contains parameters of <i>Loose</i> efficiency point)		
<i>Tight</i> (contains parameters of <i>Medium</i> efficiency point)		
Tracking	Number of hits in the second innermost pixel layer	n_{2nd}

Table 3.1: Electron identification variables. Adapted from Ref. [115].

shape variables, which veto fakes by the same principle as E_{Ratio} and $w_{\text{s,Tot1}}$; but these are instead defined in the middle layer of the EM calorimeter. These variables are also expected to be sensitive to pileup.

The EM calorimeter is designed to fully absorb and measure the energy of the vast majority of photons and electrons passing through it, and therefore, electrons passing through the EM calorimeter are expected to lose a significant fraction of their energy by the end of the calorimeter. f_3 , shown in Figure 3.10, is a variable defined in the final layer of the EM calorimeter, which capitalises on the fact that electrons rarely have much energy left to deposit by the end of the EM calorimeter to distinguish them from fakes. This is possible because hadronic showers start later and leave a more uniform energy distribution in the EM calorimeter. These variables are also expected to lose some discrimination power due to the higher pileup.

Finally, the hadronic calorimeter surrounds the EM calorimeter and is designed to measure the energy of hadrons, which deposit energy through different processes. Electrons usually leave their energy in the EM calorimeter, and very little in the hadronic calorimeter behind it. R_{had} and $R_{\text{had},1}$, shown in Figures 3.11 and 3.12 respectively, are hadronic leakage variables, which veto fakes that leave more energy in the hadronic calorimeter than could reasonably be expected of an electron of that energy. The discrimination power of these variables is also expected to deteriorate with the increased pileup at the HL-LHC.




Figure 3.1: n_{BL} .




Figure 3.2: n_{2nd} .




Figure 3.3: n_{Pix} .




Figure 3.4: $\Delta\eta_1$.




Figure 3.5: E_{Ratio} .




Figure 3.6: $w_{s,Tot1}$.




Figure 3.7: $w_{\eta 2}$.




Figure 3.8: R_{η} .




Figure 3.9: R_ϕ .




Figure 3.10: f_3 .




Figure 3.11: R_{had} .




Figure 3.12: $R_{\text{had},1}$.

3.3.3 Pre-Identification Requirements

Before producing the electron identification menu, requirements are placed on the electron candidates to ensure that the discriminating variables are meaningful for the candidate in question. First, at least 1 pixel hit and at least 9 silicon hits are required to ensure that the track was reconstructed sufficiently well, such that the matching variable ($\Delta\eta_1$) is well defined. Second, we required that $d_0/\sigma_{d_0} \leq 10$ to ensure that the projected origin of the electron candidate is broadly consistent with the interaction point. Finally, the electron E_T was required to be greater than 5 GeV, and the overlap removal procedure outlined in §2.3.11 is applied. These requirements are then placed on all electron candidates before the identification stage.

3.3.4 Identification Requirements

In producing the electron identification menu, requirements are first applied to the tracking variables to veto conversions. As these distributions consist of integer numbers of hits, they were evaluated without the use of sophisticated optimisation procedures. Second, the fraction of energy in the strip layer of the EM calorimeter was required to be at least 0.5% of the total in the EM calorimeter, to ensure that the variables defined using the strip layer, E_{Ratio} and $w_{\text{s,Tot1}}$, are well defined. Third, a set of 8 continuous calorimeter and matching variables are chosen based on their approximate discriminating power, as determined through the application of the Toolkit for Multivariate Analysis (TMVA) [145] software package. The 8 variables are selected to be those described in §3.3.2. R_{had} is used in the region $0.8 < |\eta| \leq 1.37$ due to the presence of a gap in the first layer of the hadronic calorimeter, else $R_{\text{had},1}$ is used. Fourth, the data was split into 5×5 bins in $(E_T, |\eta|)$ before further requirements were optimised, to ensure that the signal efficiency was flat in this space; the E_T bins are 5-10, 10-15, 15-20, 20-30 and >30 in GeV; and the $|\eta|$ bins are <0.8 (inner barrel), 0.8-1.37 (outer barrel), 1.37-1.52 (transition region), 1.52-2.37 (end-caps) and 2.37-2.47 (the edge of the end-caps). Fifth, requirements on

$\Delta\eta_1$ and R_ϕ were chosen by inspecting both plots of the variables, and the preliminary TMVA output. This was necessary because TMVA failed to find the global minimum of the solution when trying to optimise in the 8-dimensional space of the full variable set, as determined by a 2-fold cross-validation applied to an earlier data set. Finally, TMVA driven by a Genetic Algorithm was used to simultaneously optimise the requirements on the remaining 6 continuous variables, setting each requirement as a maximum or minimum by hand. The Genetic Algorithm randomly samples potential sets of requirements with a given signal efficiency, allowing the requirements to change by small amounts called *mutations*, it then averages the properties of the solutions with the lowest background efficiency in an iterative manner analogous to biological evolution. The Genetic Algorithm used here sampled from a population of 1000 sets of requirements, and the nominal solution after each iteration is taken as the one with the lowest background efficiency. The algorithm is defined to have converged if the background efficiency does not improve by more than 0.00001 after 60 iterations of the algorithm. This entire process is performed 10 times to ensure that the algorithm has converged correctly. The *Loose* efficiency point was designed to contain all the events contained by the *Medium* efficiency point, which in turn was required to contain all the events contained by the *Tight* efficiency point. Electrons in the transition region and at the edges of the end-caps were not used for the optimisation, due to insufficient MC statistics in these regions. Instead, the requirements for the corresponding E_T bins, in the region $1.52 < |\eta| \leq 2.37$, are applied to electron candidates in the transition region and the edges of the end-caps.

3.3.5 Evaluation of the Menu

The efficiencies are evaluated for each stage of the reconstruction process. The cluster efficiency is calculated as the fraction of generator-level electrons which leave a cluster in the EM calorimeter. The reconstruction efficiency is calculated as the fraction of generator-level electrons that left an energy cluster in the EM calorimeter, which have a track successfully matched to them, completing the reconstruction procedure. Lastly,

the efficiency of the track quality requirements is calculated as the fraction of fully reconstructed electrons matched to generator-level electrons which pass the track-quality requirements. The generator-level electron or jet relevant to each electron candidate was selected as the one with the lowest ΔR from the candidate. If the reconstructed electron candidate does not have a generator-level electron with $\Delta R < 0.1$, and no generator-level jet with $\Delta R < 0.3$, then it is discarded.

The signal and background electron identification efficiencies are calculated for each efficiency point, as the number of identified electron candidates divided by the number of electron candidates passing the track-quality requirements. These were evaluated for backgrounds of jets reconstructed using the anti- k_t algorithm with a radius parameter of 0.4, hadrons, conversions and heavy flavour fakes. The total efficiencies were also calculated for true electrons (jet fakes) as the fraction of generator-level electrons (jets) which passed all stages of the reconstruction and identification, both inclusively and in bins of E_T and $|\eta|$.

The charge misidentification probability and electron energy resolution were calculated on the set of generator-level electrons surviving the *Medium* identification efficiency point. This was calculated as the fraction of electrons for which the reconstructed charge does not match that of the generator-level electron.

To avoid the bremsstrahlung tail, the resolution was taken as the standard deviation of a Gaussian distribution, iteratively fitted to the asymmetric peak $[\mu - 2.5\sigma, \mu + 1.5\sigma]$ of $\frac{E_T^{\text{Reconstructed}} - E_T^{\text{Generator}}}{E_T^{\text{Generator}}}$. The fit was optimised by minimising the negative-log-likelihood of the fit. The fit result to the 20-30 GeV E_T bin is illustrated in Figure 3.13. The resolution was fitted with the standard energy resolution formula to parameterise the efficiency as a function of E_T .

The effect of different pileup scenarios on the efficiency and fake rate of the *Medium* efficiency point was evaluated. This comparison was performed on the Step 1.5 ITk layout, which was an old ITk layout that is now superseded by the Step 1.6 layouts. This serves to demonstrate that the identification algorithms are robust against pileup.




Figure 3.13: Asymmetric fit to $\frac{E_T^{\text{Reconstructed}} - E_T^{\text{Generator}}}{E_T^{\text{Generator}}}$ distribution for events in the barrel with $20 \text{ GeV} < E_T < 30 \text{ GeV}$, evaluated using MC samples with $\mu=190\text{-}210$ and the Step 1.6 ITk layout.

Finally, the performance of the various efficiency points using the Step 1.6 Inclined and Extended ITk layouts are compared. This provides input to the layout of ITk.

For the entire evaluation, unless specified otherwise, further restrictions were applied. First, for signal the generator-level electron E_T was required to be greater than 7 GeV, and for backgrounds the reconstructed electron E_T was required to be greater than 7 GeV. This differs from the 5 GeV minimum E_T requirement which was applied before the optimisation of the identification requirements. Second, overlapping electrons are removed. Finally, the transition region ($1.37 < |\eta| \leq 1.52$) and high $|\eta|$ region ($2.37 < |\eta| \leq 2.47$) are excluded from all efficiency evaluations except for those which are given as a function of $|\eta|$.

3.4 Results

The efficiencies for each part of the reconstruction, and the track-quality requirements, are shown in Table 3.2. The quantities are defined in §3.3.5.

The inclusive identification efficiencies and fake probabilities are shown in Table 3.3.

	Efficiency
Cluster Reconstruction / Generator-Level	(99.8 $\pm$ 0.1)%
Reconstruction / Cluster Reconstruction	(97.7 $\pm$ 0.1)%
Track-Quality / Reconstruction	(98.6 $\pm$ 0.1)%

Table 3.2: Reconstruction and track-quality efficiencies, evaluated using MC samples with $\mu=190$ -210 and the Step 1.6 ITk layout.

The large uncertainty on conversions is due to low MC statistics. While the total efficiencies and fake probabilities are shown, binned in $|\eta|$ and E_T , in Figures 3.14 and 3.15 respectively.

Category	<i>Loose</i>	<i>Medium</i>	<i>Tight</i>
Identification Efficiency (%)			
Electrons	92.4 $\pm$ 0.1	85.2 $\pm$ 0.1	65.3 $\pm$ 0.1
Jet Fakes	6.2 $\pm$ 0.2	2.7 $\pm$ 0.1	0.90 $\pm$ 0.08
Hadrons	5.0 $\pm$ 0.1	2.01 $\pm$ 0.07	0.72 $\pm$ 0.04
Conversions	10 $\pm$ 2	4 $\pm$ 2	0.6 $\pm$ 0.5
Heavy Flavour	42 $\pm$ 6	23 $\pm$ 5	11 $\pm$ 3
Total Efficiency (%)			
Electrons	88.9 $\pm$ 0.1	82.0 $\pm$ 0.1	62.8 $\pm$ 0.1
Jet Fakes	0.150 $\pm$ 0.005	0.065 $\pm$ 0.003	0.022 $\pm$ 0.002

Table 3.3: Identification efficiencies and fake probabilities, evaluated using MC samples with $\mu=190$ -210 and the Step 1.6 ITk layout. The identification efficiencies are the probability that the relevant generator-level particle/jet will be identified as an electron, given that it has already satisfied the reconstruction and track-quality requirements. The total efficiency is the probability that the relevant generator-level particle/jet will be identified as an electron.

The fraction of generator-level electrons surviving the identification menu, which have an incorrectly identified charge, are shown as a function of $|\eta|$ and E_T , in Figure 3.16.

The transverse energy resolution for generator-level electrons at pileup 190-210, which pass the *Medium* identification efficiency point, is shown in Figure 3.17 as a function of the transverse energy of the generator-level electron, for electrons in both the central barrel and the end-caps. As expected, the resolution improves with higher transverse energy.

The standard energy resolution formula is




Figure 3.14: Total electron efficiency against (a) $|\eta|$ and (b) E_T , evaluated using MC samples with $\mu=190-210$ and the Step 1.6 ITk layout [3].




Figure 3.15: Total jet fake probability against (a) $|\eta|$ and (b) E_T , evaluated using MC samples with $\mu=190-210$ and the Step 1.6 ITk layout [3].




Figure 3.16: Charge misidentification probability against (a) $|\eta|$ and (b) E_T , evaluated using MC samples with $\mu=190$ -210 and the Step 1.6 ITk layout [3].




Figure 3.17: The transverse energy resolution for different E_T -values in the barrel and end-caps, evaluated using MC samples with $\mu=190$ -210 and the Step 1.6 ITk layout. The black line is a fit of the standard energy resolution formula to the barrel resolution distribution, as described in the main text.

$$\frac{\sigma}{E_T} = \frac{a}{E_T} \oplus \frac{b}{\sqrt{E_T}} \oplus c \text{ [146]}.$$

The first term is the ‘noise term’, and it is due to electrical noise in the detector. The second term is the ‘stochastic term’, and it is due to natural fluctuations in the shower development. The third term is the ‘constant term’, it covers instrumental effects which are not dependent on the particle energy, such as non uniformities in the detector response due to structural imperfections or radiation damage. Fitting this to the generator-level electrons in the central region passing the *Medium* selection criterion for pileup 190-210, we find: a noise term of $a = 1.06 \pm 0.03$; a stochastic term of $b = 0.201 \pm 0.017$; and a constant term of $c = 0 \pm 1$. Meaning the resolution can be approximately parameterised by:

$$\frac{\sigma}{E_T} = \frac{1.06}{E_T/\text{GeV}} \oplus \frac{0.201}{\sqrt{E_T/\text{GeV}}},$$

The dependence of the *Medium* identification efficiency point efficiency and fake probability on μ is shown for the LoI samples in Figure 3.18(a). The efficiencies for electrons and jet fakes to be identified as electrons are lower for higher pileup scenarios, while no statistically significant trend can be observed for conversion fakes. These trends are due to the fact that any calorimeter deposit from an electron or jet which overlaps with a pileup jet, will be made broader and deeper due to the additional energy deposited by the pileup jet, and is thus less likely to be identified as an electron. The fact that the electron efficiency changes by less than 5% between $\mu = 80$ and $\mu = 200$ demonstrates that the ITk is successfully mitigating the effect of the high pileup. This can be contrasted to Figure 3.18(b), which shows that even at the relatively modest pileup of Run 2, which can be seen in Figure 2.2, the electron identification efficiency of the current inner tracker is effected on the level of a few percent.

A comparison of the Inclined and Extended Step 1.6 ITk layouts reveals that the Inclined layout shows superior electron performance. The *Tight* efficiency point of the Inclined layout providing a total inclusive efficiency of $(62.8 \pm 0.1)\%$ and a jet fake prob-




Figure 3.18: (a) Comparison of jet and conversion fake probabilities for three different μ -values, for the *Medium* efficiency point and Step 1.6 Inclined sample [4]. (b) Run 2 electron identification efficiency for μ above and below 40 [147].

ability of $(0.022 \pm 0.002)\%$, while the Extended layout which provides a total inclusive efficiency of $(61.4 \pm 0.1)\%$ and a jet fake probability of $(0.028 \pm 0.002)\%$, an almost 30% higher fake rate. The difference in electron total efficiency can be seen to be concentrated in the high $|\eta|$ region where the two layouts most differ. This can be seen by comparing parts (a) and (b) of Figure 3.19, which shows the total electron efficiency against $|\eta|$, for the Step 1.6 Inclined and Extended ITk layouts, respectively. This effect is possibly due to the reduced material in the forward region for the Inclined layout, as compared to the Extended layout.

3.5 Photon Performance at the HL-LHC

3.5.1 Introduction

Further to the electron performance investigation, a study was conducted into the expected photon performance of the HL-LHC. This study lent heavily on the lessons learnt from the electron performance investigation, and was structured in very much the same way. This optimisation was performed in the context of the $HH \rightarrow bb\gamma\gamma$ channel, which




Figure 3.19: Total electron efficiency against $|\eta|$ for the Step 1.6 (a) Inclined and (b) Extended layouts [3], for $\mu=190\text{--}210$.

required an efficiency of around 70% over a range of p_T values.

Unlike electrons, unconverted photons lack tracking information. However, they can convert into an electron-positron pair before they reach the EM detector. These photons are referred to as “converted photons”, as opposed to “unconverted photons”. Ideally, these two types of photon would be treated separately, however this is beyond the scope of this study, for which they were treated inclusively. The photons considered in this study are reconstructed as per Section 2.3.5. As the photons do not have a track to rely on, they can only be reliably reconstructed up to $|\eta|$ of 2.37, beyond which some of the EM shower escapes from the detector.

The scope of this analysis is reduced in several ways with respect to the electron performance study. First, there is only one efficiency point, as opposed to the full 3 efficiency point menu developed for electrons. Second, due to the lack of tracking information with photons, there are no track-based pre-identification requirements. Third, the only background considered to this analysis are hadronic fakes, which are described as they relate to electron and photons in §3.2.1. Finally, only the efficiencies and fake rates are estimated for the single efficiency point.

3.5.2 Method

The MC background samples used in this analysis were the same as was used for the electron study, which are described in §3.3.1. The signal samples used were produced using the same MC generators as with the electron study, and also had a full GEANT4 simulation of the ATLAS detector, but the $H \rightarrow \gamma\gamma$ process was simulated in place of $Z \rightarrow e^+e^-$.

Ten variables are used in the identification of photons, these are: f_1 , $w_{s,\text{Tot1}}$, E_{Ratio} , ΔE , $w_{\eta 2}$, R_η , R_ϕ , R_{had} or $R_{\text{had},1}$, and $E_{\text{T}}^{\text{cone20}}$. All of these variables except ΔE and $E_{\text{T}}^{\text{cone20}}$ are common to the electron identification menu, and are described in §3.3.2.

ΔE is defined as the difference between the energy associated with the second maximum in the strip layer of the EM calorimeter, and the energy reconstructed in the strip with the minimal value found between the first and second maxima; i.e. it is the height of the second maxima above the minimum between the first and second maxima. This variable is able to detect showers induced by multiple particles, which is the case in jets, especially in the case of an EM-like $\pi^0 \rightarrow \gamma\gamma$ shower, but often not with isolated photons.

$E_{\text{T}}^{\text{cone20}}$ is the only isolation variable used in the menu. It is defined as the sum of transverse energy calorimeter cells within ΔR of 0.2 around the barycentre of the EM cluster, excluding a 5×7 grid at the centre, which is assumed to contain the photon.

The distributions of these variables are shown for photons and jet fakes in Figures 3.20 to 3.28, except for f_1 , which is used to ensure that the strip-layer variables are well defined.

The identification methods used for photons closely follow those for electrons, described in §3.3.4, and so this subsection will focus on the differences with that procedure. The R_ϕ requirement was chosen by eye with some guidance from TMVA, which resulted in a requirement of $R_\phi < 0.8$ for $p_{\text{T}} < 40$ GeV, or $R_\phi < 0.9$ otherwise. All the other variables were optimised entirely in TMVA. Different efficiency points were derived for each bin, in order to reflect the different efficiencies of the R_ϕ requirement in each bin, producing an overall efficiency distribution which is flatter in p_{T} . Lastly, due to low MC




Figure 3.20: E_{Ratio} .




Figure 3.21: $w_{s,\text{Tot1}}$.




Figure 3.22: $w_{\eta 2}$.




Figure 3.23: R_{η} .




Figure 3.24: R_{ϕ} .




Figure 3.25: R_{had} .




Figure 3.26: $R_{\text{had},1}$.




Figure 3.27: ΔE .




Figure 3.28: $E_T^{\text{cone}20}$.

statistics, the highest p_T regions of the efficiency point were optimised using events of all p_T values. This was shown to increase the performance at high p_T , as the low statistics was causing sub-optimal performance.

First, the reconstruction efficiency is calculated. Second, the signal and jet fake total efficiencies, defined as the number of identified candidates divided by the number of generator-level candidates, were calculated both inclusively and in bins of p_T and $|\eta|$.

The generator-level photon/jet relevant to each candidate was selected as the one with the lowest ΔR from the candidate. If the reconstructed photon candidate does not have a generator-level photon with $\Delta R < 0.1$, and no generator-level jet with $\Delta R < 0.3$, then it is discarded from further analysis.

For the entire evaluation, unless specified otherwise, further restrictions were applied. First, for signal the generator-level photon p_T was required to be greater than 20 GeV, and for backgrounds the reconstructed photon p_T was required to be greater than 20 GeV. Second, overlapping photons are removed. Finally, the transition region ($1.37 < |\eta| \leq 1.52$) is excluded from all efficiency calculations except the efficiency binned in $|\eta|$.


3.5.3 Results

The reconstruction efficiency is the probability of a generator-level photon being reconstructed. The reconstruction efficiency is $(96.2 \pm 0.4)\%$.


The total efficiencies and fake probabilities are shown, binned in $|\eta|$ and p_T , in Figures 3.29 and 3.30 respectively.

3.6 Conclusion

In preparation for the conditions expected at the HL-LHC, a complete identification menu for electrons has been optimised, with three efficiency points: *Loose*, *Medium* and *Tight*. These were evaluated on MC samples of various ITk layouts and pileup scenarios, and lower pileup events were found to have consistently higher efficiencies. For 190-210




(a)




(b)

Figure 3.29: Total photon efficiency against (a) $|\eta|$ and (b) E_T , evaluated using MC samples with $\mu=190$ -210 and the Step 1.6 ITk layout.



(a)



(b)

Figure 3.30: Total jet fake probability against (a) $|\eta|$ and (b) E_T , evaluated using MC samples with $\mu=190$ -210 and the Step 1.6 ITk layout.

interactions per bunch crossing, and an electron efficiency of $(62.8 \pm 0.1)\%$, the expected jet fake efficiency is $(0.022 \pm 0.002)\%$. This corresponds to an efficiency for identifying electron candidates from hadrons as electrons of $(0.72 \pm 0.04)\%$. During Run 2, the electron identification algorithm described in §2.3.3 results in an electron identification efficiency of 78% for an efficiency for identifying electrons candidates from hadrons as electrons of 0.3%, for electron candidates with $E_T = 25$ GeV. For a slightly looser efficiency point, the HL-LHC results described in this chapter predict about a factor of 2 lower background rejection. However, for the Run 2 results a more powerful likelihood-based identification algorithm is applied, which is in contrast to the cut-based identification procedure used herein. In addition to the electron performance study, a single identification efficiency point for photons has been optimised. This was evaluated on MC samples with the Step 1.6 Inclined ITk layout. For 190-210 interactions per bunch crossing, and a photon efficiency of over 70%, less than 1 fake for every 1000 jets is expected. These results demonstrate that, despite the increased pileup and loss of the TRT at the HL-LHC, ATLAS will still have considerable electron and photon identification performance.

CHAPTER 4

SEARCH FOR HIGGS BOSON DECAYS TO CHARM QUARK PAIRS

4.1 Introduction

Unlike the vector bosons, whose couplings to the Higgs field are integral to the gauge theory, the fermion couplings to the Higgs field are generated separately, in an ad-hoc procedure called the Yukawa mechanism. Of the mechanisms which can give rise to fermion masses, the Yukawa mechanism is motivated purely by its simplicity. For this reason, the couplings of the Higgs boson to fermions are an ideal place to search for BSM physics.

With a SM BR of 2.9% [148], Higgs boson decays to charm (c) quarks represent the fermionic decay mode with the largest BR for which no experimental evidence exists to date. This high BR makes this decay mode the most promising window through which to probe the Yukawa couplings of the second-generation quarks. Despite this, the Yukawa coupling of the charm quark is still small in the SM, meaning that new physics affecting this sector could lead to notable modifications [149, 150], with some models [151, 152] predicting values within the reach of the ATLAS. For instance, Ref [149] shows that the

2HDM can facilitate around a four-fold increase in the coupling of the Higgs boson to the charm quark. For these reasons, the first direct search for Higgs boson decays into charm quarks was made at the ATLAS experiment, and the results are published in Ref. [7].

Despite direct evidence for couplings of the Higgs boson to the top [153] and bottom [154] quarks, as of yet, no successful measurement has been made of second generation quarks. Previous searches for Higgs boson decays into a light vector meson and a photon have been conducted [155–158], with the tightest constraint on the charm Yukawa coupling coming from the J/ψ channel [156, 159], which resulted in an indirect 95% CL_s upper limit of about 220 times the SM expectation on the production cross section times BR, with mild theoretical assumptions. Lastly, global fits to observed SM channels, combined with some theoretical assumptions, impose an indirect 95% CL upper limit of 22% on the decay of the Higgs boson to unobserved particles [160]. Within these constraints, large modifications to the decay rate of the Higgs boson to charm quarks from new physics are still possible.

This chapter describes a direct search for Higgs boson decays into charm quarks [7], where the Higgs boson is produced in association with a Z boson which decays to electrons or muons, which will be collectively called leptons in this chapter. The Feynman diagram for this process is given in Figure 4.1. The search used the ATLAS 2015 and 2016 dataset, collected at 13 TeV, with a total integrated luminosity of $(36.1 \pm 0.8) \text{ fb}^{-1}$ [161]. However, this search is particularly challenging at hadron colliders like the LHC, due to the small BR, large hadronic backgrounds, and requirement to tag charm-flavour jets against backgrounds of light- and bottom-flavour jets. For these reasons, the Higgs boson is tagged by its associated production with a Z boson, providing a powerful trigger signature, high signal to background ratio, simple SM background composition, and low exposure to experimental uncertainties due to the leptonic final state. Four signal categories are used in this analysis, each defined as having either: *low* ($75 \text{ GeV} < p_{\text{T}}^Z < 150 \text{ GeV}$) or *high* ($p_{\text{T}}^Z > 150 \text{ GeV}$) dilepton p_{T} ; and one or two of the highest p_{T} jets, which are used to reconstruct the Higgs boson candidate, being c -tagged. The analysis methods are

validated by searching for the ZV process, where $V = Z$ or W , and confirming that the observed yield is compatible with the SM expectation.




Figure 4.1: Feynman diagram showing a Higgs boson produced in association with a Z boson, decaying into a pair of charm quarks, while the Z boson decays dileptonically.

The search also pioneered the use of the new ATLAS c -tagging algorithms, with an efficiency point defined and calibrated specifically for the analysis.

4.2 Experimental and Simulated Data Samples

The data used in this search correspond to the portion of the Run 2 dataset collected by the ATLAS detector between 2015-2016. This represents a total integrated luminosity of 36.1 fb^{-1} . This corresponds to: 3.2 fb^{-1} of 2015 data; and 33.0 fb^{-1} of 2016 data.

The MC samples used in this search are described below:

- The $q\bar{q} \rightarrow ZH$ signal sample was produced at NLO using POWHEG-BOX v2 [104], GoSAM [162] and MINLO [163, 164] interfaced to the NNPDF3.0NLO and PDF4LHC15NLO [162] PDF sets for modelling the hard scatter, with the Higgs boson mass set to 125 GeV. PYTHIA 8 was used to model the parton shower and underlying event, with the AZNLO [165] set of tuning parameters.
- The $gg \rightarrow ZH$ signal sample was produced at LO using POWHEG-BOX v2 interfaced to the NNPDF3.0NLO and PDF4LHC15NLO PDF sets for modelling the hard scatter, with the Higgs boson mass set to 125 GeV. PYTHIA 8 was used

to model the parton shower and underlying event, with the AZNLO set of tuning parameters.

- The nominal $Z + \text{jets}$ and ZV background samples were generated by interfacing SHERPA 2.2.1 [97] to COMIX [166] and OPENLOOPS [167] for the calculation of the matrix element, using the NNPDF3.0NNLO PDF set [162]. The parton showering and underlying events were also modelled using SHERPA 2.2.1, using the SHERPA tuning parameters [168]. Events with <2 jets are generated at NLO, and events with >2 jets generated at LO. The samples represent slices in $\max(H_T, p_T^Z)$, and have various filters and vetoes placed on the jet flavours.
- The $t\bar{t}$ background was produced at NNLO and NNLL using POWHEG-BOX v2 for modelling the hard scatter, interfaced to the NNPDF3.0NLO PDF set, with the top quark mass set to 172.5 GeV. PYTHIA 8 is used to model the parton shower, interfaced to EVTGEN to improve the modelling of heavy flavour decays, using the A14 set of tuning parameters, and HDAMP set to 1.5 times to top quark mass. The sample has a dilepton filter applied to increase statistics.

All samples are normalised to their theoretical cross-sections, with any higher order corrections applied. The MC samples described here use the full GEANT4-based [106] ATLAS detector simulation [105], as described in §2.2.7.

4.2.1 Generator-Level Jet Flavour Categorisation

Jets which are simulated in the MC samples, are categorised based on the presence of a nearby heavy flavour hadron. The hadron must have $p_T > 5$ GeV, and be within $\Delta R < 0.3$ of the jet cone axis. If a b -hadron is found, the jet is labelled a b -jet. If a c -hadron but no b -hadron is found, then the jet is labelled a c -jet. If no b - or c -hadrons are found, then the jet is labelled as a light flavour jet.

4.3 Event Selection

Events of interest were selected using the unprescaled single-electron or single-muon trigger with the lowest p_T threshold. For 2015 data taking, these thresholds corresponded to 24 (20) GeV for electrons (muons). This was raised to 26 GeV for all leptons for 2016 data-taking, due to the increased instantaneous luminosity of the LHC.

Anti- k_t jets [122] with a distance parameter of 0.4 are reconstructed, and their energy calibrated, as described in §2.3.7. All jets must have $p_T > 20$ GeV, and be reconstructed inside the EM calorimeter, which requires that $|\eta| < 2.5$.

At least one of these jets in each event is then required to be c -tagged, as described in §2.3.8. An efficiency point of the c -tagging algorithm is optimised for this analysis, which results in an efficiency of 41%, for a b -jet fake probability of 25% and a light jet fake probability of 5%. These efficiencies are calibrated in data, using b -jets from $t \rightarrow Wb$ and c -jets from $W \rightarrow cs$. To reduce the effect of limited MC statistics, these efficiencies are accounted for in simulation by weighting events by the relevant tagging efficiency. These weights are parameterised as a function of the jets p_T and $|\eta|$, and the angular distance of the closest jet.

Electrons and muons are reconstructed and identified for use in this search, as per the descriptions in 2.3.3 and 2.3.6, respectively. All leptons used in this analysis must have a p_T of at least 7 GeV, and must pass a *Loose* (efficiency >99%) track-isolation criteria. The overlap procedure given in §2.3.11 is applied to electrons, muons and jets, to ensure that no object is double counted.

Once the various physics objects are defined, the full event selection for the analysis is summarised as follows:

1. The event must pass at least one of the triggers used.
2. At least one vertex must be reconstructed, as per §2.3.2.
3. Two same-flavour leptons must be identified.

4. If the leptons are muons they must have opposite charge. This requirement is not applied to electrons due to the higher charge misidentification rate.
5. At least one lepton which was matched to an object which caused the event to pass the trigger must have a $p_T > 27$ GeV.
6. The invariant mass of the dilepton system must be compatible with the mass of the Z boson: $81 \text{ GeV} < m_{\ell^+\ell^-} < 101 \text{ GeV}$.
7. At least two jets must be identified, as per the requirements of Section 4.3.
8. At least one of the two highest p_T jets must be c -tagged, as per the requirements of Section 4.3.
9. The ΔR between the two highest p_T jets must be: < 2.2 for $p_T^Z < 150 \text{ GeV}$; < 1.5 for $150 \text{ GeV} < p_T^Z < 200 \text{ GeV}$; or < 1.3 for $200 \text{ GeV} < p_T^Z$. This cut is primarily to reject $Z + \text{jets}$ and $t\bar{t}$ backgrounds.
10. The invariant mass of the dijet system (formed from the two highest p_T jets) must be compatible with the mass of the Higgs boson: $50 \text{ GeV} < m_{c\bar{c}} < 200 \text{ GeV}$.

This selection was largely optimised in the context of the $H \rightarrow b\bar{b}$ analysis, which has a very similar final state. The two main exceptions were the flavour tagging, and the ΔR requirements, for which a dedicated optimisation was performed.

4.4 Signal and Background Modelling

The shapes of the signal and all of the backgrounds in this analysis are modelled using MC simulation. The normalisation of the signal is extracted from the fit to data. The background for this analysis is heavily dominated by $Z + \text{jets}$, with smaller contributions from ZV , $t\bar{t}$ and $ZH(b\bar{b})$. Higgs boson decays to $b\bar{b}$ are treated in the same way as the other backgrounds. The normalisation of the $Z + \text{jets}$ background is determined entirely

by the fit to data, allowing the normalisation to vary independently between analysis categories. In contrast, the normalisation of the other backgrounds are constrained to the MC prediction, and assumed to be fully correlated across the analysis categories, but with independent uncorrelated uncertainties of around 10% on the acceptance in each category. Due to the di-lepton selection, the WW , $W + \text{jets}$ and single top processes are found to be negligible, and are modelled by proxy from the $t\bar{t}$ background.

4.5 Systematic Uncertainties

Systematic uncertainties are incorporated into the fit using nuisance parameters which modify the shape or normalisation of the various distributions from which the model is constructed. Some of these uncertainties are constrained using auxiliary measurements, which are then included in the fit.

The systematic uncertainties from the modelling of the shape and normalisation, including the normalisation of the free $Z + \text{jets}$ background, of both the signal and background processes are considered. Various uncertainties caused by the flavour tagging calibration are also included. Experimental uncertainties on the calibration of the leptons and jets used in the analysis are included, along with uncertainties on the luminosity and pileup reweighting. Lastly, the effects of limited statistics in the simulated datasets are also included in the model.

4.6 Generator-Level Evaluation of Modelling Uncertainties

Some of the largest systematic uncertainties for this analysis are the modelling uncertainties on the various background processes considered. The hadronisation modelling uncertainties are estimated for the signal and background processes by comparing the predictions of different MC generators. The PDF, factorisation scale and renormalisation

scale, which accounts for missing higher order corrections in the calculation, uncertainties are estimated by varying the PDFs used, or by varying the factorisation and renormalisation scales used by the generators. Other external generator parameters are varied, as described below. The uncertainty is then taken as the largest difference found with the chosen generator, which is itself selected for its good agreement with data by the ATLAS Physics Modelling Group. Taking the largest difference found with a large set of generators is a conservative approach, which is standard in ATLAS analyses. Many more alternative samples are available at generator-level, than are fully reconstructed. As such, generator-level studies are conducted to find the largest variation from the nominal sample for each background, after which, fully reconstructed samples are used to evaluate the variation where they are available. Both the shape and acceptance uncertainties are evaluated at generator-level for three of the largest backgrounds: $Z + \text{jets}$, $t\bar{t}$ and diboson. Comparisons of the generator-level and reconstruction-level backgrounds are used to validate the use of the generator-level samples to evaluate background uncertainties. This modelling uses generator-level MC samples, with a generator-level analysis selection, and the same parameterised c -tagging efficiencies as used in the rest of the analysis. The distributions of relevant variables are compared to fully reconstructed distributions for the nominal samples to validate the methods described here.

4.6.1 Background Samples Considered

In this subsection the various background samples are described, with focus on how the alternative samples differ from the nominal sample.

$Z + \text{jets}$

The nominal $Z + \text{jets}$ sample is described in Section 4.2. MC files in which the Z boson decays into muons, tau-leptons or neutrinos are not included in the comparison with the

alternative samples, for which these files are not available.

One alternative $Z + \text{jets}$ sample has a hard-scatter event generated at NNLO with MADGRAPH, and is interfaced to EVTGEN for improved modelling of the heavy flavour decays, with the NNPDF2.3LO PDF. PYTHIA 8 is used to model the underlying event and parton showering, with the A14 set of tuning parameters. The samples are sliced in the total transverse hadronic energy in the event.

The remaining four alternative $Z + \text{jets}$ samples have an underlying event generated as per the nominal sample, but with some of the external parameters of the generator being varied. The CKKW parameter [169] is varied between 15 and 30, while the resummation scale is varied between 0.25 and 4.

$t\bar{t}$

All of the $t\bar{t}$ samples are filtered for dilepton events to increase their contribution to the signal region of the analysis. The nominal sample is described in Section 4.2.

The first alternative sample considered has a hard-scatter event generated by MADGRAPH 5, with NLO corrections from AMC@NLO, and the NNPDF2.0NLO PDF set. It is interfaced to EVTGEN for improved heavy flavour modelling, using the scale $\sqrt{\Sigma_i (m_{\text{top}}^{\text{T}})^2}/2$. The parton shower and underlying event are modelled by PYTHIA 8, with the A14 set of tuning parameters.

The second alternative sample considered is generated using POWHEG, EVTGEN and HERWIG 7.0. The H7UE set of tuning parameters is used, with the MMHT2014LO68CL PDF set used for the shower modelling and multi-parton interactions, while the matrix element is calculated using the NNPDF3.0NLO PDF set and HDAMP set equal to 1.5 times the top mass. Only leptonically decaying W bosons are allowed.

Diboson

The diboson samples which represent large backgrounds in this analysis can be broadly categorised as ZZ and ZW , where one Z boson decays dileptonically, and the other vector boson decays hadronically. The nominal samples are described in Section 4.2.

The first alternative diboson sample has a hard-scatter event generated at NLO using POWHEG, with the CT10 PDF set used for the ME calculation, and the CTEQ6L1 PDF set used for the shower and multi-parton interactions. EVTGEN is interfaced to improve the modelling of heavy flavour decays. HERWIG is used to model the parton showering and underlying event, with the CTEQ6L1-UE-EE-5 set of tuning parameters.

The second alternative sample has a hard-scatter event generated at NLO using POWHEG, with the CT10 PDF set. EVTGEN is interfaced to improve the modelling of heavy flavour decays. PYTHIA 8 is used to model the parton shower and underlying event, with the AZNLO CTEQ6L1 generator set of tuning parameters. Events with a dijet or dilepton mass below 20 GeV are filtered out.

The third alternative sample is generated using SHERPA 2.1, and the CT10 generator PDF set and set of tuning parameters. Events with <2 jets are generated at NLO, and events with >2 jets generated at LO.

Half of the remaining alternative samples are generated in the same way as the POWHEG+PYTHIA 8 alternative sample, but with variations in various external parameters. These variations consist of an upward and downward variation in the renormalisation scale factor variables, and a parameter controlling multi-parton interactions.

The other half of the remaining alternative samples are all generated in the same way as the SHERPA 2.1 alternative sample, also with variations in the external parameters. These variations consist of an upward and downward variation in the factorisation, renormalisation and resummation scales. To most directly analyse the effect of these modified parameters, these variations were compared to the SHERPA 2.1 alternative sample in order to isolate the effect of the different parameters from the effect of the different generator. The factorisation, renormalisation and resummation scales are each varied between 0.25

and 4 for this sample. While variations of 0.5 and 2 are usually standard, only variations of 0.25 and 4 were present in this MC sample.

4.6.2 Validation of Generator-Level Modelling

The variations are evaluated using generator-level information, with a generator-level version of the selection, and the same parameterised c -tagging efficiencies used in the main analysis. However, this generator-level analysis does not account for the smearing of the jet energies due to the detector, and thus is validated using two methods. The first validation method compares the distributions produced in the generator-level study to the fully reconstructed distributions for the nominal signal and $t\bar{t}$ samples (for which they are available). The second method compares the final shape systematics derived from this method, to the shape systematics derived from the reconstruction-level evaluation method for the largest background ($Z + \text{jets}$).

All variables show good agreement except $m_{c\bar{c}}$, shown in Figure 4.2, for which the effect of the jet energy resolution of the detector is clearly visible. However, it is expected that any discrepancies in $m_{c\bar{c}}$ will cancel to first order when we take the ratio of the nominal sample results with the variations. This assumption is validated by the comparison to the reconstruction-level shape systematics.

Generator- vs Reconstruction-Level Variable Comparison

Figures 4.2, 4.3 and 4.4 show comparisons of the generator-level and reconstruction-level distributions for signal samples for evaluating the shape variations, while Figures 4.5, 4.6 and 4.7 show the same distributions for the $t\bar{t}$ background.




Figure 4.2: Comparison of $m_{c\bar{c}}$ generator-level and reconstruction-level distributions, for (a) the $75 \text{ GeV} < p_T^Z < 150 \text{ GeV}$ category and (b) the $p_T^Z > 150 \text{ GeV}$ category. 2 c -tags are required for the nominal signal samples.






Figure 4.3: Comparison of p_T^Z generator-level and reconstruction-level distributions, for (a) the $75 \text{ GeV} < p_T^Z < 150 \text{ GeV}$ category and (b) the $p_T^Z > 150 \text{ GeV}$ category. 2 c -tags are required for the nominal signal samples.




(a)




(b)

Figure 4.4: Comparison of $\Delta R(j1, j2)$ generator-level and reconstruction-level distributions, for (a) the $75 \text{ GeV} < p_T^Z < 150 \text{ GeV}$ category and (b) the $p_T^Z > 150 \text{ GeV}$ category. 2 c -tags are required for the nominal signal samples.



(a)



(b)

Figure 4.5: Comparison of $m_{c\bar{c}}$ generator-level and reconstruction-level distributions, for (a) the $75 \text{ GeV} < p_T^Z < 150 \text{ GeV}$ category and (b) the $p_T^Z > 150 \text{ GeV}$ category. 2 c -tags are required for the nominal $t\bar{t}$ samples.




Figure 4.6: Comparison of p_T^Z generator-level and reconstruction-level distributions, for (a) the $75 \text{ GeV} < p_T^Z < 150 \text{ GeV}$ category and (b) the $p_T^Z > 150 \text{ GeV}$ category. 2 c -tags are required for the nominal $t\bar{t}$ samples.




Figure 4.7: Comparison of $\Delta R(j1, j2)$ generator-level and reconstruction-level distributions, for (a) the $75 \text{ GeV} < p_T^Z < 150 \text{ GeV}$ category and (b) the $p_T^Z > 150 \text{ GeV}$ category. 2 c -tags are required for the nominal $t\bar{t}$ samples.

Generator- vs Reconstruction-Level Shape Systematic Comparison

To further validate the generator-level method of evaluating the modelling systematics, the final shape systematics, as derived from the generator-level reconstruction-level methods, are compared, as shown in Figure 4.8. These comparisons demonstrate the validity of the assumption that the effects of the jet resolution will cancel upon taking the ratio between nominal and alternative samples.




Figure 4.8: Comparison of $m_{c\bar{c}}$ shape systematics derived using the generator-level and the full reconstruction method for the Z +jets background, for (a) the $75 \text{ GeV} < p_T^Z < 150 \text{ GeV}$ 2 c -tags category, (b) the $p_T^Z > 150 \text{ GeV}$ 2 c -tags category, (c) the $75 \text{ GeV} < p_T^Z < 150 \text{ GeV}$ 1 c -tag category, and (d) the $p_T^Z > 150 \text{ GeV}$ 1 c -tag category.

4.6.3 Yield Modelling Uncertainties

Uncertainties on the yield, acceptance, and relative yields in the different p_{T}^Z categories are studied for different MC samples using the generator-level analysis.

$Z + \text{jets}$

As the largest background, the normalisation of the $Z + \text{jets}$ background is taken directly from the fit to data. This choice is justified by the large normalisation modelling uncertainties shown in Tables 4.1 and 4.2. The large magnitude of these uncertainties arise from the sensitivity of the analysis to the flavour of the final state, which is modelled using low-scale, non-perturbative QCD.

Variation	$75 \leq p_{\text{T}}^Z \leq 150 \text{ GeV}$	$p_{\text{T}}^Z > 150 \text{ GeV}$
<i>Sherpa2.2.1</i>	26600 ± 70 (1)	6400 ± 20 (1)
MADGRAPH+PYTHIA8	27900 ± 60 (1.05)	6750 ± 20 (1.05)
SHERPA CT10 CKKW=15	12800 ± 80 (0.48)	3060 ± 20 (0.478)
SHERPA CT10 CKKW=30	16500 ± 100 (0.62)	6230 ± 30 (0.974)
SHERPA CT10 QSF=0.25	14900 ± 90 (0.56)	6320 ± 30 (0.987)
SHERPA CT10 QSF=4	12900 ± 100 (0.487)	5170 ± 30 (0.809)

Table 4.1: Yields of $Z(\rightarrow ee) + \text{jets}$ variations in the m_{jj} variable, for the $75 \leq p_{\text{T}}^Z \leq 150 \text{ GeV}$ and $p_{\text{T}}^Z > 150 \text{ GeV}$ categories, requiring jet flavour inc, 1 c -tags. The ratio with respect to nominal is shown in parentheses. The nominal sample is shown in italics, while the most discrepant variation is shown in bold for each category.

$t\bar{t}$

As it is by far the dominant non- $Z(\ell^+\ell^-)$ background to this analysis, $t\bar{t}$ is used as a proxy to model all non- $Z(\ell^+\ell^-)$ backgrounds. As such, the agreement between data and MC in a $\mu - e$ control region without any c -tagging applied is used to estimate the normalisation uncertainty on this background. This uncertainty is estimated to be 14% and 38% for

Variation	$75 \leq p_T^Z \leq 150$ GeV	$p_T^Z > 150$ GeV
<i>Sherpa2.2.1</i>	2040 $\pm$ 10 (1)	517 $\pm$ 2 (1)
MADGRAPH+PYTHIA8	2300 $\pm$ 10 (1.13)	558 $\pm$ 3 (1.08)
SHERPA CT10 CKKW=15	1290 $\pm$ 10 (0.633)	310$\pm$4 (0.599)
SHERPA CT10 CKKW=30	1590 $\pm$ 20 (0.778)	485 $\pm$ 3 (0.938)
SHERPA CT10 QSF=0.25	1450 $\pm$ 10 (0.71)	506 $\pm$ 4 (0.98)
SHERPA CT10 QSF=4	1260$\pm$10 (0.619)	414 $\pm$ 3 (0.802)

Table 4.2: Yields of $Z(\rightarrow ee) + \text{jets}$ variations in the m_{jj} variable, for the $75 \leq p_T^Z \leq 150$ GeV and $p_T^Z > 150$ GeV categories, requiring jet flavour inc, 2 c -tags. The ratio with respect to nominal is shown in parentheses. The nominal sample is shown in italics, while the most discrepant variation is shown in bold for each category.

the *low* and *high* p_T^Z regions, respectively. The order of magnitude of these numbers are validated by comparison to the relative yields in Tables 4.3 and 4.4.

Variation	$75 \leq p_T^Z \leq 150$ GeV	$p_T^Z > 150$ GeV
<i>Powheg+Pythia8</i>	1970 $\pm$ 10 (1)	69.9 $\pm$ 3.2 (1)
AMC@NLO+PYTHIA8	1700$\pm$20 (0.863)	55.9$\pm$3.9 (0.8)
POWHEG+HERWIG7	1960 $\pm$ 10 (0.997)	68.7 $\pm$ 3.4 (0.983)

Table 4.3: Yields of $t\bar{t}$ variations in the m_{jj} variable, for the $75 \leq p_T^Z \leq 150$ GeV and $p_T^Z > 150$ GeV categories, requiring jet flavour inc, 1 c -tags. The ratio with respect to nominal is shown in parentheses. The nominal sample is shown in italics, while the most discrepant variation is shown in bold for each category.

Variation	$75 \leq p_T^Z \leq 150$ GeV	$p_T^Z > 150$ GeV
<i>Powheg+Pythia8</i>	256 $\pm$ 4 (1)	7.44 $\pm$ 1.66 (1)
AMC@NLO+PYTHIA8	213$\pm$4 (0.831)	6.14$\pm$1.61 (0.824)
POWHEG+HERWIG7	252 $\pm$ 4 (0.982)	7.23 $\pm$ 1.66 (0.971)

Table 4.4: Yields of $t\bar{t}$ variations in the m_{jj} variable, for the $75 \leq p_T^Z \leq 150$ GeV and $p_T^Z > 150$ GeV categories, requiring jet flavour inc, 2 c -tags. The ratio with respect to nominal is shown in parentheses. The nominal sample is shown in italics, while the most discrepant variation is shown in bold for each category.

The uncertainty on the relative yield of $t\bar{t}$ events in the different p_T^Z categories is estimated by comparing the relative yields of various generators, as shown in Table 4.5. A conservative estimate of 7% is placed on the relative yields in both c -tag categories.

Variation	$\frac{p_T^Z > 150 \text{ GeV}}{75 \leq p_T^Z \leq 150 \text{ GeV}}$ 1-tag	$\frac{p_T^Z > 150 \text{ GeV}}{75 \leq p_T^Z \leq 150 \text{ GeV}}$ 2-tag
<i>Powheg+Pythia8</i>	0.0355 ± 0.0016 (1)	0.029 ± 0.007 (1)
AMCATNLO+PYTHIA8	0.0329 ± 0.0023 (0.927)	0.0288 ± 0.0076 (0.993)
POWHEG+HERWIG7	0.035 ± 0.002 (0.986)	0.0287 ± 0.0066 (0.989)

Table 4.5: Ratio of yields in the $p_T^Z > 150 \text{ GeV}$ to $75 \text{ GeV} \leq p_T^Z \leq 150 \text{ GeV}$ categories for $t\bar{t}$ variations in the m_{jj} variable, for the 1 and 2 c -tag categories, requiring jet flavour inc. The ratio with respect to nominal is shown in parentheses. The nominal sample is shown in italics, while the most discrepant variation is shown in bold for each category.

Diboson

Comparing the cross-section predictions for the ZZ and ZW processes of the nominal sample (SHERPA 2.2.1) and an alternative (POWHEG+PYTHIA 8), allows us to place an estimate on the uncertainty on the total production rates of these processes. This uncertainty is found to be +4.5% for ZZ , and -4.6% for ZW . We therefore assign a conservative uncertainty of 5% on the production rates on both processes, which is uncorrelated in the fit as the production mechanisms of these processes are different.

Uncertainties are also placed on the acceptances of our analysis selection for these processes. They are evaluated as 13% for ZZ and 12% for ZW , by comparison to the POWHEG+PYTHIA 8 sample at reconstruction-level. These uncertainties are shown to be approximately the same when evaluated by comparing generator-level acceptances, as shown in Tables 4.6 to 4.9. It is clear that the POWHEG+PYTHIA 8 variations are not the largest variation (which is the POWHEG+HERWIG), but estimates derived from these samples are used for consistency with the shape variations, and because the reconstruction-level samples include detector effects.

4.6.4 Shape Modelling Uncertainties

To model the uncertainty on the shape of the background model in the fit, we take a bin-by-bin ratio of the $m_{c\bar{c}}$ distributions for the variations with the nominal background sample, both normalised to unity. We then fit this ratio with a first order polynomial in

Variation	$75 \leq p_T^Z \leq 150$ GeV	$p_T^Z > 150$ GeV
<i>Sherpa2.2.1</i>	780 ± 7 (1)	264 ± 5 (1)
POWHEG+HERWIG	617 ± 6 (0.792)	216 ± 5 (0.817)
POWHEG+PYTHIA8	677 ± 6 (0.868)	230 ± 5 (0.872)
SHERPA2.1	786 ± 8 (1.01)	268 ± 6 (1.02)

Table 4.6: Acceptances ($\times 10^5$) of ZZ variations in the m_{jj} variable, for the $75 \leq p_T^Z \leq 150$ GeV and $p_T^Z > 150$ GeV categories, requiring jet flavour inc, 1 c -tags. The ratio with respect to nominal is shown in parentheses. The nominal sample is shown in italics, while the most discrepant variation is shown in bold for each category.

Variation	$75 \leq p_T^Z \leq 150$ GeV	$p_T^Z > 150$ GeV
<i>Sherpa2.2.1</i>	111 ± 4 (1)	41 ± 3 (1)
POWHEG+HERWIG	86.7 ± 3.6 (0.782)	32.2 ± 2.8 (0.784)
POWHEG+PYTHIA8	96.5 ± 3.7 (0.871)	35.3 ± 2.8 (0.860)
SHERPA2.1	114 ± 4 (1.03)	40.4 ± 3.1 (0.985)

Table 4.7: Acceptances ($\times 10^5$) of ZZ variations in the m_{jj} variable, for the $75 \leq p_T^Z \leq 150$ GeV and $p_T^Z > 150$ GeV categories, requiring jet flavour inc, 2 c -tags. The ratio with respect to nominal is shown in parentheses. The nominal sample is shown in italics, while the most discrepant variation is shown in bold for each category.

Variation	$75 \leq p_T^Z \leq 150$ GeV	$p_T^Z > 150$ GeV
<i>Sherpa2.2.1</i>	775 ± 5 (1)	278 ± 4 (1)
POWHEG+HERWIG	634 ± 6 (0.819)	227 ± 4 (0.814)
POWHEG+PYTHIA8	688 ± 5 (0.889)	247 ± 4 (0.887)
SHERPA2.1	760 ± 5 (0.981)	275 ± 4 (0.988)

Table 4.8: Acceptances ($\times 10^5$) of ZW variations in the m_{jj} variable, for the $75 \leq p_T^Z \leq 150$ GeV and $p_T^Z > 150$ GeV categories, requiring jet flavour inc, 1 c -tags. The ratio with respect to nominal is shown in parentheses. The nominal sample is shown in italics, while the most discrepant variation is shown in bold for each category.

Variation	$75 \leq p_T^Z \leq 150$ GeV	$p_T^Z > 150$ GeV
<i>Sherpa2.2.1</i>	61.4 ± 2.4 (1)	24.4 ± 1.9 (1)
POWHEG+HERWIG	50.1 ± 2.4 (0.816)	19.1 ± 1.9 (0.785)
POWHEG+PYTHIA8	55.3 ± 2.4 (0.901)	21.9 ± 1.9 (0.896)
SHERPA2.1	61.2 ± 2.6 (0.997)	24.6 ± 2.1 (1.01)

Table 4.9: Acceptances ($\times 10^5$) of ZW variations in the m_{jj} variable, for the $75 \leq p_T^Z \leq 150$ GeV and $p_T^Z > 150$ GeV categories, requiring jet flavour inc, 2 c -tags. The ratio with respect to nominal is shown in parentheses. The nominal sample is shown in italics, while the most discrepant variation is shown in bold for each category.


the $m_{c\bar{c}}$ range 30 GeV to 230 GeV, which is just greater than the fit range (50 GeV to 200 GeV). The variation which diverges furthest from unity over the fit range is selected as the largest variation. Each background modelled in the likelihood fit is then allowed to vary according to the derived uncertainty, which is parameterised by a nuisance parameter.

$Z + \text{jets}$


In all but the *high* p_{T}^Z 1 c -tag category, the MADGRAPH sample is the largest shape variation from the nominal. Figure 4.9 shows the $m_{c\bar{c}}$ distribution of the nominal and the largest MADGRAPH variation in the $Z + \text{jets}$ background, for the different p_{T}^Z and c -tag categories. For the remaining category the CKKW=15 variation is the largest, but the difference is small, as shown by comparison with Figure 4.10. As such, and as the MADGRAPH sample is available at reconstruction-level, the fully reconstruction-level MADGRAPH sample was used to estimate the shape uncertainty on all analysis categories.

$t\bar{t}$


The variations are very small for the $t\bar{t}$ background, with both of the *high* p_{T}^Z categories showing no significant deviation. The largest variation for the *low* p_{T}^Z 2 c -tag category is the HERWIG alternative sample, while the largest variation for the *low* p_{T}^Z 1 c -tag category is the AMCATNLO based sample. It was judged that, while both small, the variations were slightly larger for the HERWIG alternative sample, and so for consistency it was decided that these were chosen to represent the shape uncertainty in all categories. Figure 4.11 shows the $m_{c\bar{c}}$ distribution of the nominal and the HERWIG variation in the $t\bar{t}$ background, for the different p_{T}^Z and c -tag categories. For comparison, Figure 4.12 shows the significant deviations for the AMCATNLO variation.




(a)



(b)



(c)



(d)

Figure 4.9: Comparison of $m_{c\bar{c}}$ generator-level distributions of the nominal sample and MADGRAPH variation of the $Z(\rightarrow e^+e^-) + \text{jets}$ background, for (a) the $75 \text{ GeV} < p_T^Z < 150 \text{ GeV}$ 2 c -tags category, (b) the $p_T^Z > 150 \text{ GeV}$ 2 c -tags category, (c) the $75 \text{ GeV} < p_T^Z < 150 \text{ GeV}$ 1 c -tag category, and (d) the $p_T^Z > 150 \text{ GeV}$ 1 c -tag category.







Figure 4.10: Comparison of $m_{c\bar{c}}$ generator-level distributions of the nominal sample and CKKW=15 variation of the $Z(\rightarrow e^+e^-) + \text{jets}$ background, for the *high* p_T^Z and 1 *c*-tag category.




(a)



(b)



(c)



(d)

Figure 4.11: A comparison of the $m_{c\bar{c}}$ distribution for the $t\bar{t}$ process for the nominal POWHEG+PYTHIA 8 sample and the POWHEG+HERWIG 7 samples, for (a) the $75 \text{ GeV} < p_T^Z < 150 \text{ GeV}$ 2 c -tags category, (b) the $p_T^Z > 150 \text{ GeV}$ 2 c -tags category, (c) the $75 \text{ GeV} < p_T^Z < 150 \text{ GeV}$ 1 c -tag category, and (d) the $p_T^Z > 150 \text{ GeV}$ 1 c -tag category.





Figure 4.12: A comparison of the $m_{c\bar{c}}$ distribution for the $t\bar{t}$ process for the nominal POWHEG+PYTHIA 8 sample and the AMCATNLO samples, shown in the categories for which a significant deviation is seen from the nominal. These are the $75 \text{ GeV} < p_T^Z < 150 \text{ GeV}$ categories with (a) 1 c -tag and (b) 2 c -tags.

Diboson


The POWHEG-BOX+HERWIG variations were chosen to model the background shape uncertainties as these were available and reconstruction-level. Figures 4.13 and 4.14 show the $m_{c\bar{c}}$ distributions of the nominal and these variations in the ZZ and ZW backgrounds, respectively, for the different p_T^Z and c -tag categories.

4.7 Statistical Interpretation


The signal yield is estimated using a profile likelihood fit, referred to as the fit hereafter, to the invariant mass distribution of the two highest p_T jets, simultaneously across four event categories. The four event categories are each defined having either: *low* ($75 \text{ GeV} < p_T^Z < 150 \text{ GeV}$) or *high* ($p_T^Z > 150 \text{ GeV}$) dilepton p_T ; and one or two of the highest p_T jets being c -tagged. The fit is performed in 15 bins of equal width per category,




(a)



(b)




(c)




(d)


Figure 4.13: Comparison of $m_{c\bar{c}}$ generator-level distributions of the nominal sample and POWHEG-BOX+HERWIG variation of the ZZ background, for (a) the $75 \text{ GeV} < p_T^Z < 150 \text{ GeV}$ 2 c -tags category, (b) the $p_T^Z > 150 \text{ GeV}$ 2 c -tags category, (c) the $75 \text{ GeV} < p_T^Z < 150 \text{ GeV}$ 1 c -tag category, and (d) the $p_T^Z > 150 \text{ GeV}$ 1 c -tag category.




(a)



(b)



(c)



(d)

Figure 4.14: Comparison of $m_{c\bar{c}}$ generator-level distributions of the nominal sample and POWHEG-BOX+HERWIG variation of the ZW background, for (a) the $75 \text{ GeV} < p_T^Z < 150 \text{ GeV}$ 2 c -tags category, (b) the $p_T^Z > 150 \text{ GeV}$ 2 c -tags category, (c) the $75 \text{ GeV} < p_T^Z < 150 \text{ GeV}$ 1 c -tag category, and (d) the $p_T^Z > 150 \text{ GeV}$ 1 c -tag category.

over the $m_{\ell+\ell^-}$ range 50 GeV to 200 GeV. In the absence of an observed signal, a 95% upper CL will be set on the signal yield using a modified frequentist CL_s method based on the profile likelihood ratio already described.

4.8 Diboson Validation

The methods used in this analysis are validated by measuring the flavour inclusive diboson yield, henceforth referred to as diboson signal. While either individual ZZ or ZW processes or specifically c -jet final states could be used, the expected significance in any of these scenarios is too small to provide a meaningful cross-check.

To understand the sensitivity of this validation channel, five thousand toy datasets were produced based on the post-fit Asimov dataset¹. The observation significance for the diboson signal was calculated for each dataset using the asymptotic formula for a one-sided profile likelihood ratio test statistic [170]. The resulting distribution of significances are then fit with a Gaussian distribution, as shown in Figure 4.15, from which the significance is estimated to be 2.2 ± 0.9 standard deviations.

The fraction of the ZZ yield with $Z \rightarrow c\bar{c}$ is 55% in the 2 c -tag category, and 20% in the 1 c -tag category. While the fraction of the ZW yield from $W \rightarrow cl$ is 65% in both the 1 and 2 c -tag categories. The diboson signal strength (defined as the ratio of observed to expected number of signal events extracted from the fit) is measured to be $\mu_{ZV} = 0.6^{+0.5}_{-0.4}$, consistent with unity within the uncertainty of the validation. The observed significance of the diboson signal is 1.4σ , consistent with the SM expectation of $(2.2 \pm 0.9)\sigma$. These results are summarised in Table 4.10.

¹The Asimov data is an artificial dataset in which every observable takes its median value. Fitting the Asimov dataset allows the median expected sensitivity and the fluctuation about this to be derived [170].




Figure 4.15: Distribution of the expected diboson signal observation significance, evaluated using toy datasets sampled from the post-fit Asimov dataset.

	Observed	Expected
Signal Strength	$0.6^{+0.5}_{-0.4}$	$1.0^{+0.5}_{-0.4}$
Observation Significance	1.4	2.2 ± 0.9

Table 4.10: The observed and expected signal strengths and significances of the diboson validation.

4.9 Grouped Uncertainty Breakdown

The breakdown of the various sources of uncertainty on the signal strength, both in terms of the absolute yield and relative to the SM expectation, $\hat{\mu}$, is evaluated for both the ZH and the ZV fits. First, the uncertainty is split into statistical and systematic components, then the systematic uncertainty is further broken down into uncertainties due to distinct groups of nuisance parameters. The groups of nuisance parameters are:

- Flavour tagging
- Signal and background modelling
- Limited MC statistics
- Lepton, jet, luminosity and pileup

All uncertainties used in this breakdown are calculated using the following procedure. First, the asymmetric impact of uncertainties on the signal strength are estimated by fitting the model to the data and then finding the point where the profiled negative twice log-likelihood increases by one unit, using the MINOS algorithm [171]. Next these uncertainties are symmetrised by taking their mean, producing a quantity which represents the overall uncertainty on the signal yield.

Using this method of evaluating uncertainties, the impact of a given group of NPs is then calculated as follows:

- The total uncertainty on the signal yield is calculated including all available NPs.
- The partial uncertainty on the signal yield is then calculated having fixed all NPs in the group to be evaluated to their post-fit values, as determined from the full fit.
- The partial uncertainty on the signal yield is then subtracted in quadrature from the total uncertainty, to give a measure of the impact of the given group of NPs.

For the statistical uncertainty, the fit is performed with all NPs fixed to their post-fit values, and the uncertainty is evaluated. The total systematic uncertainty is then given by subtracting this result in quadrature from the total uncertainty.

The breakdown of the uncertainty on the signal strength is computed on the observed dataset for both the nominal analysis and the diboson validation, as shown in Tables 4.11 and 4.12, respectively. The observed breakdown on the extracted signal yield is also given in Table 4.13. The expected breakdown is also calculated for the nominal analysis using the post-fit Asmiov dataset, with an injected signal of 100 times the SM expected signal yeild, as shown in Table 4.14. This is to represent the uncertainty on the signal yield in the presence of a large signal.

NP Group	Uncertainty on $\hat{\mu}$
Total	100.9
Statistical	38.5
Systematic	93.3
Background Modelling	85.5
Flavour Tagging	73.6
Signal Modelling	28.7
Lepton, Jet and Miscellaneous	28.5
MC Statistics	6.1

Table 4.11: Breakdown of the uncertainty on $\hat{\mu}$ for the ZH analysis, evaluated on the observed dataset.

NP Group	Uncertainty on $\hat{\mu}$
Total	0.449
Statistical	0.138
Systematic	0.427
Background Modelling	0.414
Lepton, Jet and Miscellaneous	0.188
Flavour Tagging	0.142
Signal Modelling	0.082
MC Statistics	0.008

Table 4.12: Breakdown of the uncertainty on $\hat{\mu}$ for the ZV analysis, evaluated on the observed dataset.

The resulting uncertainties do not sum in quadrature to equal the total uncertainty,

NP Group	Uncertainty on fitted $\sigma(pp \rightarrow ZH) \times \mathcal{B}(H \rightarrow c\bar{c})$ [pb]
Total	2.6
Statistical	1.0
Systematic	2.4
Background Modelling	2.2
Flavour Tagging	1.9
Signal Modelling	0.7
Lepton, Jet and Miscellaneous	0.7
MC statistical	0.2

Table 4.13: Breakdown of the uncertainty the fitted value of $\sigma(pp \rightarrow ZH) \times \mathcal{B}(H \rightarrow c\bar{c})$ for the ZH analysis, evaluated on the observed dataset. The SM value of $\sigma(pp \rightarrow ZH) \times \mathcal{B}(H \rightarrow c\bar{c})$ is taken to be 0.025 pb.

NP Group	Uncertainty on $\hat{\mu}$
Total	59.2
Statistical	29.3
Systematic	51.5
Background Modelling	32.5
Flavour Tagging	32.3
Lepton, Jet and Miscellaneous	13.8
Signal Modelling	6.66
MC Statistics	2.86

Table 4.14: Breakdown of the uncertainty on $\hat{\mu}$ for the ZH analysis, evaluated on the post-fit Asimov dataset with a $ZH(c\bar{c})$ signal of $\mu = 100$ injected.

because the different groups of nuisance parameters have non-zero correlations between them.

4.10 Results

The post-fit yields for the signal and each background process, in each analysis category is shown in Table 4.15. Each signal category is shown in Figure 4.16.

Sample	Yield, $50 \text{ GeV} < m_{c\bar{c}} < 200 \text{ GeV}$			
	$75 \leq p_{\text{T}}^Z < 150 \text{ GeV}$		$p_{\text{T}}^Z > 150 \text{ GeV}$	
	1 c -tag	2 c -tags	1 c -tag	2 c -tags
$Z + \text{jets}$	69400 ± 500	5320 ± 100	15650 ± 180	1280 ± 40
ZZ	490 ± 70	55 ± 18	180 ± 28	26 ± 8
ZW	750 ± 130	53 ± 13	290 ± 50	20 ± 5
$t\bar{t}$	2020 ± 280	240 ± 40	130 ± 50	13 ± 6
$ZH(b\bar{b})$	32 ± 2	4.1 ± 0.4	19.5 ± 1.5	2.7 ± 0.2
$ZH(c\bar{c})$	-143 ± 170	-30 ± 40	-84 ± 100	-20 ± 29
$ZH(c\bar{c}) \text{ (SM)}$	2.36 ± 0.01	0.73 ± 0.01	1.401 ± 0.004	0.490 ± 0.003
Total	72500 ± 320	5650 ± 80	16180 ± 140	1320 ± 40
Data	72504	5648	16181	1320

Table 4.15: Post-fit yields for signal and all background processes, as well as observed data yields, in each analysis category. The post-fit signal yields normalised to their SM expectation are also shown. The uncertainties represent both statistical and systematic effects.

These results show no clear signal, and an observed (expected) 95% CL_s upper limit is set on $\sigma(pp \rightarrow ZH) \times \text{BR}(H \rightarrow c\bar{c})$ of 2.7 pb ($3.9_{-1.1}^{+2.1}$ pb), which translates to an upper limit on the signal strength of 110 (150_{-40}^{+80}) times the SM expectation on the signal yield. The best fit value for the $ZH(\rightarrow c\bar{c})$ signal strength is $\mu_{ZH} = -69 \pm 101$.

Due to the larger SM BR and the fact it has very similar kinematic features, the $ZH(b\bar{b})$ background is particularly troublesome. If this analysis were more sensitive, this would be a significant issue, and for this reason this background is treated differently in the next section. However, due to the dominance of the $Z + \text{jets}$ background, this limit is not sensitive to modifications in the $ZH(b\bar{b})$ background yield. This was demonstrated




Figure 4.16: Observed and post-fit simulated $m_{c\bar{c}}$ distributions, for (a) the $75 \text{ GeV} < p_T^Z < 150 \text{ GeV}$ 2 c -tags category, (b) the $p_T^Z > 150 \text{ GeV}$ 2 c -tags category, (c) the $75 \text{ GeV} < p_T^Z < 150 \text{ GeV}$ 1 c -tag category, and (d) the $p_T^Z > 150 \text{ GeV}$ 1 c -tag category [7]. The expected signal is scaled up by a factor of 100 so that it is visible. The total pre-fit background yield is also shown. The error bands show the quadrature sum of statistical and systematical uncertainties on the background prediction.

by varying this background between 0 and 2 times the SM expectation, which resulted in a limit within $\sim 5\%$ of the nominal limit.

4.11 HL-LHC Prospects Study

The end of this chapter describes the prospects for this search at the upgraded HL-LHC ATLAS detector [130, 172], by extrapolating the simulated datasets used in the Run 2 analysis [7], to the expected integrated luminosity of 3000 fb^{-1} . The results of this extrapolation are published in Ref. [8]. The analysis strategy is almost identical to the Run 2 analysis [7], with any differences stated explicitly. However, at the HL-LHC, both the ATLAS detector and analysis strategy are expected to change significantly because the amount of data can be used to constrain the backgrounds and uncertainties, so it is difficult to make reliable predictions for the systematic uncertainties which will affect the HL-LHC analysis. Due to these limitations, the main result of this study is an evaluation of the statistical power of the search. The impact of the sources of the largest systematic uncertainties on the search are estimated, based on their Run 2 values. Finally, the effect of expected flavour tagging performance improvements from upgrades to the ATLAS detector on the sensitivity of this search is estimated.

4.11.1 Methods

The expected HL-LHC result is obtained by scaling the signal and background expectations from the Run 2 search [7] by process- and category-dependent scale factors (SF). These SFs account for the higher integrated luminosity and various production cross sections, and the modified c -jet tagging efficiencies. The performance for all physics objects relevant to this analysis is assumed to be unchanged from Run 2 for the nominal result of this study, due to a cancellation between the effect of pileup and the upgrades to the ATLAS detector, which is a design goal of many aspects of upgrade of the ATLAS detector [173]. However, the effect of c -jet tagging performance improvements expected with

the HL-LHC ITk are estimated.

The higher luminosity is taken into account by scaling the expected yields of all physics processes by the ratio of the Run 2 to expected HL-LHC integrated luminosities: $3000 \text{ fb}^{-1}/36.1 \text{ fb}^{-1}$. The increased centre-of-mass energy, $\sqrt{s}$ from 13 TeV to 14 TeV, of the LHC is taken into account by scaling the expected yields for the $ZH(c\bar{c})$ and $ZH(b\bar{b})$ processes, in a p_{T}^Z -category dependent way. The overall scaling is taken as the ratio of $pp \rightarrow ZH$ cross sections from [174], accounting for differences between the p_{T}^Z -categories using a PYTHIA 8-based generator-level MC simulation [175]. The Z + jets and diboson backgrounds are scaled using the predicted ratio of $q\bar{q}$ parton luminosities [174] for 14 TeV and 13 TeV, while the $t\bar{t}$ background is scaled using the ratio of gg parton luminosities [174].

In this study, the higher signal yield expected at the HL-LHC is exploited by using a tighter c -jet tagging efficiency point, with greater b - and light-flavour jet rejection than used in the Run 2 analysis. This tighter efficiency point has a c -jet efficiency of 18%, for background efficiencies of 5% and 0.5% for b -jet and light-flavour jets, respectively. This reduces the total background by a factor of 5.3, and lowers relative contribution of the kinematically irreducible $ZH(b\bar{b})$ background, for a 54% loss in signal. These effects improve the upper limit by 7%.

The normalisation of the Z + jets background is free to vary independently in the four fit categories, given its importance and the limited precision with which it is modelled by MC, with the expectation that it will be constrained by the large number of expected Z + jets background events in the HL-LHC scenario. The diboson and $t\bar{t}$ background normalisations are fixed to the SM expectation. It is estimated that the normalisation of the Z + jets background will be determined from the data with an uncertainty of less than 2% in all categories. As the Z + jets background normalisation is determined in data, its uncertainty is considered as statistical.

In the Run 2 analysis, the $ZH(b\bar{b})$ background normalisation was constrained using the uncertainty on its SM expectation. However, for the HL-LHC study, the normalisation

of the $ZH(b\bar{b})$ background is constrained to the expected uncertainty of the $VH(b\bar{b})$ measurement for an integrated luminosity of 3000 fb^{-1} at the HL-LHC [176]: 14%. This represents the probable scenario in which the $VH(b\bar{b})$ normalisation will be determined in data, possibly in a simultaneous measurement of the $VH(b\bar{b})$ and $VH(c\bar{c})$ processes. The uncertainty on the $VH(b\bar{b})$ normalisation is considered as a statistical uncertainty for this reason.

4.11.2 Results

In the absence of systematic uncertainties, the 95% CL_s [177] upper limit on the signal strength is expected to be $\mu_{ZH(c\bar{c})} < 6.3^{+2.5}_{-1.8}$, where the uncertainty is calculated as the $\pm 1\sigma$ interval of background-only pseudo-experiments. The expected best fit value for the $ZH(c\bar{c})$ signal strength is $\mu_{ZH(c\bar{c})} = 1.0 \pm 3.2$. Figure 4.17 shows the $m_{c\bar{c}}$ distributions in the four analysis categories, and the expected yields for the signal and background processes in each category are shown in Table 4.16.

Sample	Yield			
	$75 \leq p_{\text{T}}^Z < 150 \text{ GeV}$		$p_{\text{T}}^Z > 150 \text{ GeV}$	
	1 c -tag	2 c -tags	1 c -tag	2 c -tags
$Z + \text{jets}$	$271\,000 \pm 14\,000$	4350 ± 220	$59\,300 \pm 3000$	890 ± 40
WZ	4080 ± 200	48.5 ± 2.4	1700 ± 90	29.6 ± 1.5
ZZ	2570 ± 130	96 ± 5	1020 ± 50	49.7 ± 2.5
$t\bar{t}$	$16\,000 \pm 800$	241 ± 12	860 ± 40	26.3 ± 1.4
$ZH(b\bar{b})$	441 ± 17	10.7 ± 0.4	327 ± 12	9.4 ± 0.4
$ZH(c\bar{c})$	74.4 ± 2.8	8.54 ± 0.32	52.6 ± 2.0	6.89 ± 0.26
Total	$294\,000 \pm 14\,000$	4750 ± 220	$63\,300 \pm 3000$	1010 ± 40
$S/\sqrt{S+B}$	0.137 ± 0.008	0.124 ± 0.007	0.209 ± 0.013	0.216 ± 0.013

Table 4.16: The expected yields for the signal and each background process, in each analysis category for $100 \text{ GeV} < m_{c\bar{c}} < 150 \text{ GeV}$. The yields are taken from MC simulation. The uncertainties on the cross section, which are not included in the fit, and shown for each sample. The final row shows $S/\sqrt{S+B}$ for the $ZH(c\bar{c})$ signal (S), considered within the context of the sum of all background contributions (B).

Allowing the $Z + \text{jets}$ normalisation to float in the fit has an impact of +21% on the expected upper limit for an integrated luminosity of 3000 fb^{-1} , relative to the case




Figure 4.17: Post-fit $m_{c\bar{c}}$ distributions for the HL-LHC projection, in (a) the $75 \text{ GeV} < p_T^Z < 150 \text{ GeV}$ 2 c -tags category, (b) the $p_T^Z > 150 \text{ GeV}$ 2 c -tags category, (c) the $75 \text{ GeV} < p_T^Z < 150 \text{ GeV}$ 1 c -tag category, and (d) the $p_T^Z > 150 \text{ GeV}$ 1 c -tag category [8]. The expected signal is scaled by a factor of 100 to be visible on the plots. The Asimov Dataset is constructed from the sum of expected signal and background events, while the stacked histogram corresponds only to the background expectations. The error bars represent the statistical uncertainty of the expected number of data events.

where it is fixed to the nominal prediction in each category. The upper limit improves by only $\sim 1\%$ when the $ZH(b\bar{b})$ background normalisation is fixed to its SM expectation, showing that the expected $ZH(b\bar{b})$ signal strength measurement is sufficient to constrain this background.

In the context of this study, the sensitivity at the end of Run 3 was also evaluated under the assumption of 300 fb^{-1} of 13 TeV pp collision data. The analysis method and systematic uncertainties assumed are identical to those of the Run 2 analysis. At the end of Run 3 the 95% CL_s expected upper limit on the $ZH(c\bar{c})$ signal strength is $\mu_{ZH(c\bar{c})} < 38^{+18}_{-10}$, estimated with a fit close to that of the Run 2 search [7].

4.11.3 Systematic Uncertainties

Evaluation of Systematic Uncertainties by Category

Systematic uncertainties affecting the Run 2 analysis are modelled as nuisance parameters in the fit. Due to the changes to the detector and the analysis strategy expected for the HL-LHC analysis, it is difficult to estimate the precise sources and effects of systematic uncertainties. However, the impact of some of the dominant sources of systematic uncertainty in the Run 2 analysis, estimated based on their impact in the Run 2 analysis, are studied to estimate the susceptibility of the HL-LHC sensitivity to pertinent systematic uncertainties. Individual nuisance parameters from the Run 2 analysis are assigned to broad groups (e.g. c -jet tagging and background shape) based on the nature of the systematic uncertainties they correspond to. The effect of each group of nuisance parameters on the 95% CL_s expected upper limit on the $ZH(c\bar{c})$ signal strength is evaluated by repeating the fit with all of the nuisance parameters in a given group introduced to the fit. The impact of these groups of uncertainties on the sensitivity is summarised in Table 4.17. The uncertainties associated with the nuisance parameters can be constrained in the fit. The largest constraints occur for the nuisance parameters associated with the $Z + \text{jets}$ background shape, c -jet tagging efficiency and jet energy scale and resolution un-

certainties, representing the ability to more precisely control these sources of uncertainty using the large amount of data at the HL-LHC. Compared to the Run 2 analysis, the HL-LHC analysis experiences a reduced exposure to the uncertainties associated with the tagging efficiency measurements, due to the reduced light-flavour jet component in the background as a result of the tighter operating point, and due to the uncertainty on the c -jet tagging efficiency being constrained in the fit.

The uncertainty of the shape of the $Z + \text{jets}$ background, due to the modelling of the underlying event and the parton shower, is likely to represent the dominant limitation to the sensitivity of the analysis, and will therefore require careful consideration in a HL-LHC analysis. However, the impact of the experimental systematics uncertainties (e.g. the c -jet tagging efficiency uncertainty) in a HL-LHC scenario will likely reduce relative to their effect on the Run 2 analysis given the large datasets available, allowing precise performance studies to be conducted. This effect is estimated in this study through the constraints on the associated nuisance parameters.

Source of uncertainty	Change in limit
Background shape	+36%
Jet energy scale and resolution	+17%
Lepton reconstruction and identification	+12%
c -jet tagging efficiency	+11%

Table 4.17: The increase in the nominal 95% CL_s upper limit on the $ZH(c\bar{c})$ signal strength due to the introduction of typical systematic uncertainties, based on their effect on the Run 2 analysis. The “Background shape” uncertainties refer to the shape uncertainty of the $Z + \text{jets}$, diboson, $t\bar{t}$ and $ZH(b\bar{b})$ backgrounds as estimated from MC generator comparisons in the Run 2 analysis. The “ c -jet tagging efficiency” uncertainties refer to the uncertainty of the efficiencies of c -, b - and light-flavour jets in data, determined within the context of the Run 2 analysis.

4.11.4 Simultaneous Evaluation of Systematic Uncertainties

In an alternative approach, systematic uncertainties affecting this analysis are all modelled in the likelihood nuisance parameters, based on those of the Run 2 analysis. Given the large number of expected background events, the uncertainties associated with NPs


affecting the shape and normalisation of the background are expected to be significantly reduced with respect to their prior values during the fit. For this reason, the uncertainties affecting the background only are applied only to the kinematically irreducible $H \rightarrow b\bar{b}$ background. While those affecting both signal and background are conservatively decorrelated between signal and background, and applied only to the signal and the $H \rightarrow b\bar{b}$ background. Uncertainties affecting only the signal are retained unchanged.

Due to the kinematic differences in the production of di-jets from diboson decays and those produced in association with a Z -boson, c -tagging systematic uncertainties are considered separately between these backgrounds. To account for this, a set of c -tagging NPs are assigned to the diboson background, uncorrelated with the c -tagging NPs on the Higgs processes. c -tagging NPs are neglected for the dominant $Z + \text{jets}$ background, due to its large statistical power to reduce the fitted uncertainty. It is further assumed that the c -tagging uncertainties on the $t\bar{t}$ background are fully correlated with those of the $Z + \text{jets}$ background, and therefore are also assumed to be negligible.

As with the nominal result of the study, the $ZH(b\bar{b})$ background normalisation is constrained to the uncertainty with which ATLAS expects to measure the $\mu_{VHb\bar{b}}$ signal strength using a dataset with an integrated luminosity of 3000 fb^{-1} at the HL-LHC. Furthermore, it is assumed that current limitations on the production of MC will not be an issue at the HL-LHC. Consequently, all uncertainties due to limited MC statistics are ignored in this study.

The correlation matrix and pull plot for this fit are shown in Figures 4.18 and 4.19, respectively. The likelihood scan for this fit, in addition to a fit using the *Loose* Flavour tagging efficiency point, and a fit without any nuisance parameters are shown in Figure 4.20.

Upon introducing these uncertainties, the expected 95% CL_s upper limit on μ_{ZHcc} is found to be $7.8^{+4.2}_{-2.2} \times$ the SM expectation. However, due to the aforementioned difficulties in estimating the sources and effects of the systematic uncertainties affecting this search, an accurate projection of the search results including all systematics is not possible at






Figure 4.19: Plot showing the pulls on the nuisance parameters and the uncertainties on the normalisation of the $Z + \text{jets}$ event categories for the fit including Run 2 based estimates of the systematics, as described in §4.11.4.




Figure 4.20: Likelihood scan for the parameter of interest in the fit including Run 2 based estimates of the systematics, as described in §4.11.4.

this stage.

Despite the higher number of expected background events in the HL-LHC scenario, a non-negligible uncertainty in the shape of the $m_{c\bar{c}}$ distribution for the background processes will likely remain. The effect of such a shape uncertainty is estimated by varying the MC event generator used to model the shape of the dominant $Z + \text{jets}$ background, and evaluating the change in the upper limit. The alternative generator used is MADGRAPH5_AMC@NLO [101], as opposed to the nominal generator: SHERPA 2.2.1 [97]. This results in a 7% uncertainty on the upper limit. The $m_{c\bar{c}}$ distributions in the 2 c -tag, $p_T^Z > 150$ GeV category are shown for the nominal and alternative shapes in Figure 4.21.

4.11.5 Future Improvements

This upper limit represents an overestimation of the sensitivity of the analysis, due to the absence of systematic uncertainties. However, various improvements to the analysis strategy could increase the sensitivity significantly.

In particular, other production channels, such as $Z(\nu\bar{\nu})H$ and $W(\ell\nu)H$, are known




Figure 4.21: Comparison of (a) the nominal SHERPA 2.2.1 and (b) alternative MADGRAPH5_AMC@NLO generators in the $m_{c\bar{c}}$ distributions of the 2 c -tag, $p_T^Z > 150$ GeV event category.

to exhibit comparable sensitivity to the $Z(\ell^+\ell^-)H$ channel in the analogous analysis for $H \rightarrow b\bar{b}$ decays [154]. Estimated based on the gain in sensitivity for the $H \rightarrow b\bar{b}$ analysis, we can expect a $\sim 40\%$ improvement in the upper limit.

Furthermore, the use of multivariate analysis (MVA) techniques was also shown to provide a sensitivity improvement in the $H \rightarrow b\bar{b}$ analysis [154], where a BDT improved the expected sensitivity of the search by 7%. While this seems modest, the ‘cut-based’ selection for the $H \rightarrow b\bar{b}$ analysis is significantly more mature than that of the $H \rightarrow c\bar{c}$ analysis. Therefore, this corresponds to an underestimate of the improvement that can be expected from an MVA-based analysis strategy.

The further tightening of the flavour tagging efficiency point used, the introduction of pseudo-continuous flavour tagging, and the fitting of additional b -tagged control regions to further constrain the $VHb\bar{b}$ background all represent promising improvements to the analysis. However, estimates of the effect of these improvements can not be easily quantified at this stage.

Finally, preliminary studies into the b -jet tagging performance of ATLAS at the HL-

LHC suggest an improvement of around a factor of 2.5 [178] in the light-flavour jet rejection. Assuming a factor of 2.5 improvement for the light-flavour jet rejection, for a fixed b -jet rejection and c -jet efficiency, an 8% improvement in the upper limit can be expected. Furthermore, c -jet tagging in a hadron collider environment is a very active area of research which is currently less mature than b -jet tagging. Significant improvements in the performance of c -jet tagging algorithms can be expected in coming years.

4.12 Conclusions

A search for Higgs boson decays into charm quarks, produced in association with a leptonically decaying Z boson, has been performed over 36.1 fb^{-1} of $\sqrt{s} = 13 \text{ TeV}$ pp collision data. No significant excess is observed, and so a 95% upper CL is set at 110 times the SM expectation on the signal yield. This represented the first search for inclusive Higgs boson decays to charm quarks at ATLAS, and at the time of publishing set world-leading direct constraints on the decay of the Higgs to charm quarks.

Furthermore, the expected sensitivity of this search has been evaluated by extrapolating the results of the Run 2 search. Assuming an integrated luminosity of 3000 fb^{-1} of $\sqrt{s} = 14 \text{ TeV}$ of pp collision data at the HL-LHC, a 95% CL_s upper limit on the $ZH(c\bar{c})$ signal strength of $\mu_{ZH(c\bar{c})} < 6.3$ can be expected, in the absence of systematic uncertainties. Based on this projection, if the coupling of the Higgs boson to the c -quark takes a value near to the SM expectation, an observation is not expected. However, this upper limit would provide strong constraints on new physics models, and provide competitive direct constraints on the Yukawa coupling of the Higgs boson to charm quarks. An observation may require a collider beyond the HL-LHC, and the potential of e^+e^- colliders has been highlighted for this search [179].

CHAPTER 5

SEARCH FOR HIGGS BOSON DECAYS TO PAIRS OF LIGHT RESONANCES IN THE FOUR-MUON FINAL STATE

5.1 Introduction

The SM of particle physics does a remarkably good job of describing the fundamental constituents of the universe, and their interactions. However, as was described in Chapter 1, there are various issues with the SM which suggest that physics beyond the SM should be a reality. In particular, Chapter 1 described the reasons to believe that new physics should be present in the Higgs sector. For these reasons, searches for Higgs boson decays into new light resonances are an essential part of the LHC Higgs boson physics programme. Due to its small natural width, even small couplings to BSM particles could result in substantial BRs to BSM particles. Such Higgs boson decays to BSM final states are only constrained at the 95% CL for BRs above 22% from global fits to the couplings of the Higgs boson, provided that the couplings of the Higgs boson to SM particles take the SM values [160]. This upper limit loosens to exclude Higgs boson invisible/undetected

final states with BRs above 50% if the couplings of the Higgs boson to SM particles are not assumed to take SM values [160].

This analysis [5, 6] searches for decays of the SM Higgs boson to a pair of light on-shell bosons in the 4μ final state, using 36.1 fb^{-1} of pp collision data [2], collected at the LHC centre-of-mass energy of 13 TeV. The clean experimental signature and low SM backgrounds to this search, provided by the 4μ final state, make it a key channel in which to search for Higgs boson decays to beyond the SM resonances. The mass range of the search is between 1 and 15 GeV, excluding mass windows around the charm and bottom quarkonia. The results are interpreted in terms of two benchmark models: a light pseudoscalar Higgs boson from the 2HDM+S; and a dark vector boson from a dark sector extension to the SM.

The analysis strategy is to define a set of requirements for selecting signal events, while rejecting SM backgrounds which are $H \rightarrow ZZ^* \rightarrow 4\mu$, $ZZ^* \rightarrow 4\mu$, higher order EW processes, and 4μ production from the decay of heavy flavour hadrons. The $H \rightarrow ZZ^* \rightarrow 4\mu$ and $ZZ^* \rightarrow 4\mu$ processes are estimated using the results of MC simulations, while the heavy flavour background is estimated using a data-driven technique. 95% CL_s upper limits are then established using a maximum likelihood fit of the Gaussian-modelled signal distribution and histograms describing the background distribution in the discriminating observable: $\langle m_{\mu^+\mu^-} \rangle = \frac{1}{2}(m_{12} + m_{34})$, the mean invariant mass of the di-muon systems.

Leptonic final states other than 4μ are ignored in this search, because boosted di-electron or di- τ systems would overlap in the detector, and currently no dedicated reconstruction is available for such final states. Furthermore, in the 2HDM+S interpretation the a^0 couplings are Yukawa ordered, and thus the $H \rightarrow a^0 a^0 \rightarrow 4\mu$ channel is expected to be the most important leptonic channel in the mass range $2m_\mu \lesssim m_{a^0} \lesssim 2m_\tau$.

This search was designed as an extension to a previous search for Higgs boson decays into four leptons (e or μ), through two intermediate dark vector bosons over the Z_d mass range 15 GeV to 60 GeV [5, 6]. The boosted final states and very different background composition are among the reasons why the low mass search was not included in the

original 8 TeV search. This high mass search was also updated with new data recorded at 13 TeV, and while the high mass search was not primarily the work of this author, some details will be provided to ensure complete context for the low mass search. The analysis resulted in the publication in Ref. [180].

Two benchmark models are used to design the analysis selection, and to provide specific interpretations of the search results. These models are described in §5.1.1 and §5.1.2.

5.1.1 Pseudoscalar Higgs Bosons

The 2HDM+S [41, 42], described in §1.4.3, predicts the existence of a CP-odd scalar resonance a^0 . The BR for $a^0 \rightarrow \mu^+ \mu^-$ can be significant in the range $2m_\mu \lesssim m_{a^0} \lesssim 2m_\tau$ where di-muon decays are the heaviest kinematically accessible leptonic decay modes, especially for high $\tan\beta$ in the Type-II 2HDM(+S) where the couplings of the a^0 to down-type fermions are enhanced, as shown in Figure 1.4(b). This motivates the search for Higgs boson decays to resonances with masses below 15 GeV, as performed in this analysis. The Feynman diagram for this process is given in Figure 5.1.




Figure 5.1: Feynman diagram for exotic Higgs boson decays to four muons induced by intermediate pseudoscalar resonances.

5.1.2 Dark Vector Bosons

Many BSM theories [41, 181–190] attempt to explain the existence of dark matter in the universe [191], by introducing a hidden, or dark sector. One such model considered herein, the Hidden Abelian Higgs Model [188, 189], introduces an additional $U(1)_d$ dark gauge symmetry to the SM Lagrangian [41, 185–189], manifesting in the form of a new dark vector boson (Z_d). This model could also explain astrophysical observations of positron excesses [192–194].

The dark sector could be accessed through a coupling to the SM in the form of kinematic mixing with the hypercharge gauge boson [195–197]. In this hypercharge portal scenario the strength of the coupling would be parameterised by the kinetic mixing parameter (ϵ), although stringent constraints already exist on this final state from Drell-Yan processes [185, 198–200] and beam dump experiments [201–203]. However, if the $U(1)_d$ symmetry is broken by a dark Higgs boson, in analogy to the spontaneous symmetry breaking in the SM, then there could also be mixing between the SM and dark sector Higgs bosons [41, 185–189], which would be parameterised by a Higgs portal coupling parameter: κ . This production mode can be probed for the first time at the LHC due to the production of Higgs bosons, and is searched for in this analysis. Figure 5.2 shows this production mode, which allows the (κ, m_{Z_d}) parameter space to be probed. For Z_d masses in the range searched for here, the Z_d is expected to have BRs to muons of order 10%, and to decay promptly. Previous constraints on the production of pairs of Z_d resonances exist from a similar search by CMS [204].

5.1.3 Kinematic Comparison of Models

The different spins of the pseudoscalar and vector bosons, collectively referred to as X in this chapter, lead to different kinematic distributions of the final state muons. The primary production mechanism of Higgs bosons at the LHC is gluon fusion, for which the Higgs boson is produced with low transverse momentum. When decaying to vector




Figure 5.2: Feynman diagram for exotic Higgs boson decays to four leptons induced by intermediate dark vector bosons via Higgs portal.

bosons, the Higgs boson decays preferentially to bosons in the longitudinal polarisation state, carrying an amplitude of $1 - \cos^2 \theta$, where θ is the angle between the direction of travel of the vector boson and one of its decay products measured in the rest frame of the boson. Therefore, the muons are produced preferentially perpendicular to the vector boson, while in the rest frame of the pseudoscalar the muons are produced isotropically. This difference can be seen clearly in generator-level distributions of the signal samples in Figure 5.3(a). However, the highly boosted final states ensure good agreement between the two models for most kinematic variables, as illustrated in Figure 5.3(b), motivating the dual-interpretation of the analysis.

A result of this angular difference is that a greater fraction of the muon pairs from a^0 decays will be produced co-linearly to the decaying boson, causing a greater spread in the transverse momenta of the final-state muons. This means that the lowest p_T muon in the quadruplets is softer in the case of the a^0 decays than the Z_d decays, resulting in a lower acceptance. This effect is demonstrated in Figure 5.4.

However, for X mass (m_X) below 1 GeV, the mass of the intermediate boson ceases to produce a significant effect of the kinematics of the distributions, and thus both models tend to the same efficiency. This results in an upward slope in the a^0 efficiency curve as m_X approaches zero.




Figure 5.3: Generator-level distributions of (a) the cosine of the angle between the p_T of the lead di-muon and the p_T of the anti-muon from that di-muon in the centre-of-mass frame of the di-muon, for 2 GeV Z_d and a resonances, and (b) the η of the leading muon.




Figure 5.4: Generator-level distributions of the p_T of (a) the softest and (b) the hardest muons in the quadruplet, for the 1 GeV a^0 and Z_d signal samples.

5.2 Experimental and Simulated Data Samples

The data used in this search correspond to the portion of the Run 2 dataset, collected by the ATLAS detector between 2015-2016. This represents a total integrated luminosity of 36.1 fb^{-1} . This corresponds to: 3.2 fb^{-1} of 2015 data; and 33.0 fb^{-1} of 2016 data.

The hard-scatter process for the a^0 signal samples are produced using three-steps: first, a gluon fusion SM Higgs Boson is generated using POWHEG [102–104]; second, this SM Higgs boson is replaced with a neutral scalar Higgs boson of the same mass from the 2HDM+S model; third, this BSM Higgs boson is forced to decay to a 4μ final state through two light pseudoscalar Higgs bosons using PYTHIA 8. The parton shower and underlying event is simulated by PYTHIA 8, using the AZNLO set of tuning parameters and the CTEQ6L1 PDF set [205]. The MC simulated a^0 signal samples with an a^0 mass of 0.5, 2, 5 and 15 GeV used a full simulation of the ATLAS detector, while the 1, 2.5, 4, 7 and 9 GeV samples use fast simulation [107], as described in §2.2.7. Further fast simulation samples were generated for use in the high mass search, at a^0 masses of 20, 30, 40 and 60 GeV, and the full simulation 15 GeV a^0 mass signal sample was used for both.

The Z_d signal samples are produced using the Hidden Abelian Higgs Model [188, 189] of MADGRAPH 5 [206], with the CTEQ6L1 PDF set. PYTHIA 8 [98] is used for the modelling of the parton shower, hadronisation, and the underlying event, with the A14 NNPDF23LO set of tuning parameters. These samples are produced using the fast detector simulation for Z_d mass points of 0.5, 1, 2, 5 and 15 GeV for the low mass search, and 15 (used for both), 20, 25, 30, 35, 40, 45, 50, 55 and 60 GeV for the high mass search.

The gluon fusion and vector boson fusion SM Higgs processes are considered as a background to this search. They are simulated using POWHEG [102–104] interfaced to the CT10 PDF set [143] for the hard-scatter event. PYTHIA 8 interfaced to the CTEQ6L1 PDF set is used for the parton shower, with the AZNLO set of tuning parameter [207]. The SM background where a Higgs boson is produced in association with a vector boson (Z or W) is modelled using PYTHIA 8, interfaced to the CTEQ6L1 PDF set. The SM diboson (ZZ and WZ) backgrounds are modelled using SHERPA interfaced to PYTHIA 8.

Triboson EW processes are simulated with SHERPA, using the CT10 PDF set.

Some MC samples are used in just the high mass search. The SM background where a Higgs boson is produced in association with a pair of top quarks is modelled using MADGRAPH 5, and JIMMY [208] for the underlying event. The parton shower is modelled using HERWIG++, interfaced to the CTEQ6L1 PDF set is used. The SM ZJ/ψ and $Z\Upsilon$ backgrounds are modelled using PYTHIA 8, interfaced to the CTEQ6L1 PDF sets, and a fast detector simulation is used. The $Z + \text{jets}$ backgrounds are simulated using POWHEG and PYTHIA 8, interfaced to the CT10 PDF set. The $t\bar{t}$ process is simulated with POWHEG and PYTHIA 6, interfaced to the CT10 PDF set. Lastly, the $t\bar{t}Z$ processes is simulated with MADGRAPH 5, interfaced to the CT10 PDF set.

All samples are normalised to their theoretical cross-sections, with any higher order corrections applied. The detector is fully simulated [105] using GEANT4 [106], unless stated otherwise. Some samples are produced using the fast detector simulation [107], as described in §2.2.7. For all-muon final states, these two simulation methods are largely equivalent.

5.3 Event Selection

The low mass search only involves 4μ final states, which the high mass search also includes $2e2\mu$ and $4e$ final states. This selection described the requirements placed on an event for it to enter our signal regions, which only vary where necessary between the low- and high mass analyses.

5.3.1 Triggers

A combination of dilepton triggers are used to select events of interest. An event will be considered for further analysis if it passes any one of these triggers, providing that the detector was functioning within acceptable limits while it was being recorded, and that it contains a reconstructed primary vertex formed from at least 2 tracks.

Two different types of trigger are used in the low mass search, both of which have transverse momentum requirements which varied with the instantaneous luminosity of the data-taking period. The first trigger required 2 muons with a p_T of either 10 GeV each, or 14 GeV each, depending on the data-taking period. The second trigger required a muon with p_T of at least 8 GeV which does not need to have been reconstructed by the L1 trigger, and another muon with a p_T of 18 GeV each, 20 GeV each, or 22 GeV each, depending on the data-taking period. For the high mass search, all of these triggers in addition to all other available mono-lepton, di-lepton and tri-lepton triggers were used, with a combined trigger efficiency of close to 100%, with respect to signal events passing the rest of the event selection.

5.3.2 Lepton Reconstruction

Muons are used in both the low mass and the high mass search, and are reconstructed as described in §2.3.6. All muons in the event must have $p_T > 5$ GeV, and be in the active region of the MS ($|\eta| < 2.7$). Any muons associated with a track in the ID must also have a longitudinal impact parameter, $z_0 \sin \theta$, of less than 0.5 mm with respect to the primary vertex; which is defined as the vertex with at least three associated tracks, for which the sum of the squared transverse momenta of the tracks is the highest. To reject muons from cosmic rays, the impact parameter of the muon in the bending plane, d_0 , is also required to be within 1 mm of the primary vertex. At most a single muon per quadruplet is allowed to be reconstructed from a stand alone track in the muon spectrometer, or from a calorimeter energy deposit.

Muons in this analysis must pass loose isolation requirements, with close-by corrections applied. This requirement specifies that: the transverse energy of topo clusters in a cone of $\Delta R < 0.2$ around the muon to be less than 30% of the p_T of the muon; and that the sum of the p_T of tracks within $\Delta R < \min(0.3, 10 \text{ GeV}/p_T)$ of the muon be less than 15% of the p_T of the muon. The close-by correction accounts for the presence of the other close-by muons in the quadruplet by: subtracting the transverse energy of any topo

cluster that has $0.1 < \Delta R < 0.2$ from the muon being corrected, and $\Delta R < 0.1$ from the muon being subtracted; subtracting the inner detector track p_T of any quadruplet muons in the isolation cone of the muon.

Electrons are only used in the high mass search, and are reconstructed and identified as described in §2.3.3. Electrons are required to be in the EM calorimeter ($|\eta| < 2.47$), and have $p_T > 7$ GeV for this analysis. The electron candidate is then identified as an electron if it passes the *Loose* efficiency point of the likelihood-based discriminant described in §2.3.3. As with muons, $z_0 \sin \theta$ is required to be < 0.5 mm for electrons, but the requirement on d_0 is replaced with a requirement of $|d_0^{BL}/\sigma_{d_0}| < 5$.

5.3.3 Complete Event Selection

The event selection for the low mass search is organised into the following stages:

1. Event Preselection: Events failing these requirements are discarded.
2. Baseline Muon Selection: Muons failing these requirements are discarded.
3. Quadruplet Formation: All possible muon quadruplets are formed from baseline muons - no charge selection). If no quadruplet can be formed, the event is discarded.
4. Quadruplet Selection: Quadruplets failing these requirements are discarded. If no quadruplets pass this stage, the event is discarded.
5. Quadruplet Ranking: The quadruplet with the lowest value of the ranking metric is selected for the event, and all other quadruplets are discarded. The ranking metric is: $\Delta m_{\mu^+\mu^-} = |m_{12} - m_{34}|$, where m_{12} is the invariant di-muon mass closest to the Z-boson mass, and m_{34} is the other di-muon mass.
6. Event selection: These requirements are applied to the chosen quadruplet. If any requirement is failed, the event is discarded.

Category	Requirement Name	Requirement
Event Preselection	Trigger	See §5.3.1
	Vertex	At least one vertex reconstructed with 2 or more associated tracks
Baseline Muons	p_T η ID $z_0 \sin \theta$ $ d_0 $	>5 GeV (15 GeV if calo-tagged) $ \eta < 2.7$ Loose <0.5 mm if muon is not StandAlone <1 mm if muon is not StandAlone
Quadruplet Formation	Quadruplet	At least one quadruplet
Quadruplet Selection	Opposite sign Overlap Removal p_T Trigger Matched Muon Quality	Di-Muons are $\mu^+\mu^-$ pairs No overlap-removed muons as per §2.3.11 $p_T^{\mu_1} > 20$ GeV, $p_T^{\mu_2} > 15$ GeV, $p_T^{\mu_3} > 10$ GeV Muons in quadruplet responsible for firing at least one trigger At most 1 StandAlone or CaloTagged μ
Quadruplet Ranking	Minimal Δm	Select quadruplet with smallest $\Delta m_{\mu^+\mu^-} = m_{12} - m_{34} $
Event Selection	Isolation Impact Parameter Quarkonia Veto Higgs Window Low Mass Dilepton Compatibility	All leptons in quadruplet pass isolation requirements (described in main text), with close-by corrections $ d_0^{BL}\sigma < 3$ for all muons in quadruplet Reject event if either of: $m_{J/\psi} - 0.25$ GeV $< m_{12,34,14,23} < m_{\psi(2S)} + 0.3$ GeV, $m_{\Upsilon(1S)} - 0.7$ GeV $< m_{12,34,14,23} < m_{\Upsilon(3S)} + 0.75$ GeV 120 GeV $< m_{4\mu} < 130$ GeV 0.88 GeV $< m_{12,34} < 20$ GeV $ m_{12} - m_{34} /m_{12} < 0.15$

Table 5.1: Summary of the selection used in the low mass search. The quarkonia veto uses $m_{J/\psi} = 3.096$ GeV, $m_{\psi(2S)} = 3.686$ GeV, $m_{\Upsilon(1S)} = 9.461$ GeV and $m_{\Upsilon(3S)} = 10.355$ GeV, with windows corresponding to approximately triple the mass resolution for these resonances, as taken from Ref. [209] and Ref. [210]).

This selection is summarised in Table 5.1. The high mass search also includes a baseline electron selection.

The high mass search has a similar event selection to the low mass, only differing where necessary. The differences of the high mass search with respect to the low mass search are listed below:

1. More triggers are used, as described in §5.3.1
2. A Baseline Electron Selection stage is required
3. ΔR must be greater than 0.1 between same flavour leptons, and greater than 0.2 for different flavour leptons
4. A low mass veto of $m_{12,34,14,23}$ greater than 5 GeV is applied
5. The Higgs boson mass window is widened to $115 \text{ GeV} < m_{4\ell} < 130 \text{ GeV}$ to account for the bremsstrahlung losses of the electrons
6. A Z boson veto of $m_{12,34} < 64 \text{ GeV}$ and $m_{14,32} < 75 \text{ GeV}$ is applied
7. Instead of the low mass region of $0.88 \text{ GeV} < m_{12,34} < 20 \text{ GeV}$, a high mass region of $20 \text{ GeV} < m_{12,34}$ is applied

5.4 Signal Modelling

In order to interpret the results of the search, model dependent efficiencies and signal shapes in the $\langle m_{\mu^+\mu^-} \rangle$ variable are modelled using dedicated analytical functions and Gaussian distributions, respectively. Additionally, to provide the model independent interpretation, a fiducial volume is defined such that the efficiencies of events contained within this volume are expected to be independent of the dynamics of the model in question, and these efficiencies are found to be approximately equal for the two benchmark interpretations considered in this search.

The rapidly changing signal efficiency near the quarkonia veto regions necessitates establishing regions in which we do not try to interpret the search results which cover the $\langle m_{\mu^+\mu^-} \rangle$ space in and just either side of the quarkonia veto regions. Furthermore, our data-driven background estimate could not be validated for $\langle m_{\mu^+\mu^-} \rangle < 0.088$ GeV, requiring that we veto a very low $\langle m_{\mu^+\mu^-} \rangle$ region. These regions were chosen to be taken as $\langle m_{\mu^+\mu^-} \rangle < 1$ GeV, $2.6 \text{ GeV} < \langle m_{\mu^+\mu^-} \rangle < 4.3 \text{ GeV}$, and $8.1 \text{ GeV} < \langle m_{\mu^+\mu^-} \rangle < 11.8 \text{ GeV}$.

5.4.1 Signal Shape in Mean Di-Muon Mass

The signal shapes for both interpretations are modelled as Gaussian distributions in the $\langle m_{\mu^+\mu^-} \rangle$ variable, as exemplified in Figure 5.5. The mean of the Gaussian is assumed to be equal to the m_X value in the model; this assumption is justified in Figure 5.6, which shows the strong agreement between the m_X in the model and the reconstructed $\langle m_{\mu^+\mu^-} \rangle$. The standard deviation of the Gaussian is assumed to be entirely due to the resolution of the detector, resulting from the small-width approximation in the models used, which is parameterised as the sum of an exponential function and a constant term as shown in Figure 5.7. The signal shape is thus assumed to be model independent.




Figure 5.5: Gaussian fits to $\langle m_{\mu^+\mu^-} \rangle$ distributions for the 2 GeV (a) a^0 and (b) Z_d signal samples. The signal files are normalised to unity, this makes the Z_d normalisation lower because the Z_d also decays to electrons.




Figure 5.6: Mean of a Gaussian fit to $\langle m_{\mu^+\mu^-} \rangle$ normalised to m_X over the mass range considered, for (a) a^0 and (b) Z_d signal samples. The red line denotes the function used to describe the mean of the Gaussian distribution which models the signal.




Figure 5.7: Standard deviation of a Gaussian fit to $\langle m_{\mu^+\mu^-} \rangle$ normalised to m_X over the mass range considered, for (a) a^0 and (b) Z_d signal samples. The red line denotes the function used to describe the standard deviation of the Gaussian distribution which models the signal.

5.4.2 Model-Dependent Efficiencies

The signal efficiency is found to differ between the two models considered for the reasons explained in §5.1.3. As such, for the purposes of the model-dependent interpretation, model-dependent efficiencies are calculated using all of the events in the two benchmark MC samples.

The a^0 signal efficiency can be modelled as: $\frac{c_0 + c_1 \times e^{-\langle m_{\mu^+\mu^-} \rangle / c_2}}{1 + c_3 \times e^{-\langle m_{\mu^+\mu^-} \rangle / c_4}}$, and the Z_d signal efficiency as: $c_0 + c_1 \times e^{-\langle m_{\mu^+\mu^-} \rangle / c_2}$. This efficiency, and the regions of applicability are justified in Figure 5.8.




Figure 5.8: Selection efficiencies for various mass points, and interpolated efficiency models, for the (a) a^0 and (b) Z_d . These efficiency models have been corrected for the effect of the quarkonia veto regions by modelling the invariant mass of each di-muon pair as independent Gaussian distributions.

5.4.3 Model-Independent Efficiency

A fiducial volume is defined to mirror the analyses selection in such a way as to absorb all of the model dependent kinematics, leading to the possibility of a model-independent

interpretation. This fiducial volume differs between the low- and high mass searches. Events inside the low mass search fiducial volume must have four muons, leptons for the high mass search, and satisfy the fiducial definitions shown in Table 5.2. A model-independent efficiency is then defined as the ratio of the expected yield of reconstructed events in the signal region, to the expected yield of generator-level events in the fiducial volume. This model-independent efficiency is then evaluated using both the a^0 and the Z_d MC, and the results are found to agree to within statistical uncertainties, demonstrating that most of the model dependence of the efficiency has been absorbed into the definition of the fiducial phase space. The Z_d MC is then used to evaluate the model-independent efficiencies because there are no electrons in the final states of the a^0 MC sample.

All of the definitions of the fiducial volume are applied to generator-level muons. In order to emulate the effects of quasi-collinear EM radiation from the charged leptons on the resolution of the detector, all prompt photons within ΔR of 0.1 of a lepton are added to the four-momentum of the closest lepton. This process is called dressing the leptons.

Object	Low Mass Fiducial Definition
Muons	Dressed with prompt photons within ΔR of 0.1 $p_T > 5 \text{ GeV}$ $ \eta < 2.7$
Quadruplet	Three leading p_T muons satisfy $p_T > 20 \text{ GeV}$, 15 GeV , and 10 GeV Reject event if either of: $m_{J/\psi} - 0.25 \text{ GeV} < m_{12,34,14,23}$ $< m_{\psi(2S)} + 0.3 \text{ GeV}$, $m_{\Upsilon(1S)} - 0.7 \text{ GeV} < m_{12,34,14,23}$ $< m_{\Upsilon(3S)} + 0.75 \text{ GeV}$ $0.88 \text{ GeV} < m_{12,34} < 20 \text{ GeV}$ $ m_{12} - m_{34} /m_{12} < 0.15$

Table 5.2: Definitions of fiducial volume used in model independent interpretation of low mass search.

The high mass fiducial volume differs from that of the low mass search as follows:

- Electrons are included, must be dressed, and have $p_T > 7 \text{ GeV}$ and $|\eta| < 2.5$
- The low mass requirement ($0.88 \text{ GeV} < m_{12,34} < 20 \text{ GeV}$) is replaced by a high mass requirements ($10 \text{ GeV} < m_{12,34} < 64 \text{ GeV}$)

- ΔR must be >0.1 between same flavour leptons, and >0.2 for different flavour leptons
- A requirement of $5 \text{ GeV} < m_{14,32} < 75 \text{ GeV}$ is applied if the quadruplet is $4e$ or 4μ

5.5 Background Modelling

The main backgrounds to the low mass search are heavy flavour for very low $\langle m_{\mu^+\mu^-} \rangle$, and EW otherwise. The heavy flavour background is caused in large part by a b -hadron from a $b\bar{b}$ pair decaying semi-leptonically into a muon and a charm quark, which then decays producing a second muon, resulting in two collimated di-muon systems. Various resonances produced in the decay of the b -quark (ρ , ω , ϕ , J/ψ) which result in a pair of muons are also a relevant part of this background. This estimate of this background used in the fit is data-driven. At higher $\langle m_{\mu^+\mu^-} \rangle$, the main backgrounds come from $ZZ^* \rightarrow 4\mu$ decays, both from resonant Higgs boson production, and non-resonant SM EW processes. The high mass backgrounds are also discussed briefly.

5.5.1 MC Estimate of Low Mass Electroweak Background

The yield of the Higgs boson and non-resonant EW backgrounds was estimated using MC, and distributed assuming a uniform distribution in $\langle m_{\mu^+\mu^-} \rangle$, as there was insufficient MC statistics to determine the shape. This yielded a Higgs boson background estimate of 0.100 ± 0.013 events, a ZZ background estimate of 0.104 ± 0.057 events, and a tri-boson EW background estimate of 0.074 ± 0.023 events.

5.5.2 Data-Driven Estimate of Low Mass Heavy Flavour Background

The heavy flavour background is modelled using a data-driven method, based on a similar approach developed by CMS [64]. The method models the 4μ background as a function of

the invariant masses of the di-muon pairs, which it factorises into the product of the two di-muon spectra, multiplied by an efficiency function which accounts for the Higgs boson mass compatibility requirement. The final quadruplet selection requirements are then applied to this template, and the final background estimate is given as a function of the $\langle m_{\mu^+\mu^-} \rangle$ variable. The modelling of the two di-muon spectra as independent distributions is made possible because the heavy flavour background is dominated by pairs of decays of b -quarks via two semi-leptonic decays, or decays of resonances (ρ , ω , ϕ , J/ψ) to di-muons; both of which have two independent legs of the quadruplet. The baseline muon selection is applied to the muons used in this method. A small contribution to the heavy flavour background shape in which the di-muon pairs originate from separate b -quark decays is included.

The 4-muon template is produced from the product of a high and a low p_T di-muon template, which are taken from data. These di-muon templates have been shown to represent 97% of signal events using MC. The high p_T di-muon template is taken from a fit to data events in which the trigger was fired by a pair of muons with p_T of at least 20 GeV and 10 GeV, and a third muon with a p_T of at least 5 GeV. The low p_T di-muon template is taken from a fit to data events in which the trigger was fired by a pair of muons with p_T of at least 5 GeV each, and a third muon must have a p_T of at least 25 GeV. The individual di-muon templates and the 4-muon template are shown in Figure 5.9.

The application of the Higgs boson mass window requirement introduces correlations in data, which can not be factorised into the di-muon templates. These correlations in the event kinematics are modelled using an efficiency function derived from a $b\bar{b}$ enriched data sample, in which the isolation and impact parameter requirements are inverted, which increases the $b\bar{b}$ event statistics without effecting the kinematics of the events. This efficiency function is a double 2D Gaussian function in the space of the two di-muon invariant masses. The heavy flavour shape in the signal region can then be estimated by multiplying the previously obtained 2D template by with efficiency function, resulting in the shape shown in Figure 5.10. The shape of the $b\bar{b}$ background is then obtained by




Figure 5.9: (a) di-muon and (b) 4-muon kinematic templates, extracted from fits to data [211].

applying the di-muon mass compatibility requirement.

Finally, the background estimate is normalised using a fully data-driven ABCD-style method, in which the normalisation in the signal region is estimated by extrapolating from a region in which the Higgs boson mass and di-muon mass compatibility requirements are inverted, by applying two transfer factors, under the assumption that these variables are uncorrelated. The transfer factor from the di-muon incompatibility region into the signal region is estimated in the region of inverted Higgs boson mass requirement. However, the transfer factor from the inverted Higgs boson mass window region is obtained in a region of both inverted di-muon mass compatibility, and inverted isolation and impact parameter requirements on the muons, to maximise data statistics. This process predicts a total of 0.08 ± 0.05 heavy flavour background events in the signal region.

In addition to the normalisation uncertainty on the yield, which is estimated by error propagation, a subdominant shape uncertainty is derived by varying each parameter in the fit independently by $\pm 1\sigma$. This shape uncertainty is then varied to estimate the effect on the di-muon mass compatibility transfer factor, and thus the total yield. The final




Figure 5.10: 4-muon kinematic template, extracted from fits to data and corrected for the efficiency of the Higgs boson mass requirement [211].

data-driven heavy flavour estimate and its uncertainty is shown in Figure 5.11.

5.5.3 MC Based Cross-Check of Heavy Flavour Background

The shape of the data-driven heavy flavour MC estimate was validated using a MC-based method. Ideally, we would generate the shape using a full MC simulation. However, the rare 4μ final state meant that the production was far too slow using filtered MC samples. As such, quadruplets of muon momenta were drawn from generator-level MC templates, which were then smeared to account for the detector resolution. Two processes were modelled in this way: $b\bar{b} \rightarrow 4\mu$ and $b\bar{b}b\bar{b} \rightarrow 4\mu$, as described in the following paragraphs.

Simulating $b\bar{b} \rightarrow 4\mu$

To simulate the decay of the b -hadrons, first a sample of $b\bar{b}$ events was generated using PYTHIA 8. The generator-level record of these events was scanned for b -hadrons which




Figure 5.11: Data-driven heavy flavour background estimate [211]. The yellow band represents the uncertainty on the estimate. This figure is normalised as described in §5.5.2.

decay to a pair of muons. The invariant masses of the di-muon system, and of the rest of the products of the decay of the b -hadron are used to construct a 2-dimensional PDF describing the kinematics of the b -hadron decay.

Another sample of $b\bar{b}$ events was generated using SHERPA to model the kinematics of the initial b -hadrons. The decay of these b -hadrons is then modelled in a three-step process. First, masses of the di-muon system and the remaining decay products are randomly sampled from the PDF constructed in the previous step, requiring that the decay is kinematically possible. Second, the momentum of the di-muon system and the sum of the remaining decay products are assigned random values from an isotropic distribution, representing the decay of the initial state b -hadron, into systems with the masses drawn from the PDF. Third, the individual muon momenta are assigned random values drawn from an isotropic decay of the di-muon system. The momenta of the final state muons are then smeared to account for the effect of the detector.

The PYTHIA 8 sample was chosen to model the decay of the b -hadrons because it had a di-muon filter which provides a large number of $b \rightarrow \mu^+\mu^-$ decays. The SHERPA sample was chosen to model the kinematics of the di-muon system because it populates the kinematic region of interest ($m_{b\bar{b}}$ near the Higgs boson mass).

Simulating $b\bar{b}b\bar{b} \rightarrow 4\mu$

The $b\bar{b}b\bar{b} \rightarrow 4\mu$ process is modelled in a very similar way to $b\bar{b} \rightarrow 4\mu$. A PYTHIA 8 sample is used to construct a PDF describing the decay of b -hadrons into individual muons. The PDF is a function of the energy of the final state muon, and cosine of the angle between the initial b -hadron and the final state muon in the rest-frame of the b -hadron.

A multijet sample generated using PYTHIA 8 is then searched for events with 4 b -hadrons in the generator-level record. Quadruplets of muons are then generated by randomly assigning muons to each b -hadron according to the kinematics of the single muon PDF.

Modelling of Detector Resolution

The resolution of the generator-level MC muon momenta were smeared to account for the detector resolution. The resolution on both the direction and the magnitude of the momentum of the muons were modelled by comparing the kinematic properties of reconstructed muons (from a Z_d MC signal sample) to generator-level muons, matched to them with a $\Delta R < 0.1$ requirement. The resolutions on p_T , η and ϕ were parameterised as a function of the p_T of the generator-level muon. The resolution on p_T was taken as the σ -parameter of a Gaussian distribution fitted to the $|p_T^{reco} - p_T^{generator}|/p_T^{generator}$ distribution of the matched muons, and then fitted with a first order polynomial. The angular resolutions were taken from the σ -parameter of Gaussian distributions fitted to the $|\Phi^{reco} - \Phi^{generator}|$ distributions of the muons, where $\Phi = \phi$ or η , and these were then fitted with a function of the form: $a + b/\sqrt{p_T}$. The effect of this smearing was very small.

Comparison to Data-Driven Shape

This MC-based heavy flavour background shape is derived without accounting for the effect of the isolation and impact parameter requirements, this makes it incomplete as a background model. However, by relaxing the isolation and impact parameter requirements on the data-driven model, the MC-based shape can provide a cross-check of the shape of the data-driven method, before the application of the isolation and impact parameter requirements. Figure 5.12 shows the MC-based and data-driven heavy flavour shape background estimates, without the effect of the isolation and impact parameter requirements. The shapes can be seen to be qualitatively similar, except that the toy MC based shape estimate lacks an estimate beyond about 3 GeV, implying that some processes could be unaccounted for.




Figure 5.12: Comparison of the (a) toy MC based and (b) data-driven [211] heavy flavour background estimates. The effect of the isolation and impact parameter requirements on the shape is not considered in either shape estimate. The toy-based estimate does not extend above about 3 GeV due to the limited masses of the decaying b -hadrons considered in the estimate, while other processes are clearly present in the data which were not considered in the toy-based estimate.

5.5.4 Overall Background Model

The total background for the low mass search is shown as a function of $\langle m_{\mu^+\mu^-} \rangle$ in Figure 5.13. It is also summarised inclusive in Table 5.3.




Figure 5.13: Total low mass background estimate [180].

Process	Expected Yield
$H \rightarrow ZZ^* \rightarrow 4\mu$	0.1 ± 0.1
ZZ^*	0.10 ± 0.01
Higher Order EW	0.06 ± 0.03
Heavy Flavour	0.07 ± 0.04
Total	0.4 ± 0.1

Table 5.3: Expected background yields for low mass search [211]. Uncertainties include both MC statistical and systematic components.

The backgrounds to the high mass search are determined entirely from MC simulation. The dominant irreducible backgrounds to the analysis come from Higgs boson decays to four leptons, and non-resonant EW processes with four leptons in the final state. There are also small $t\bar{t}$ and Z + jets contributions from jets being incorrectly identified as electrons. Finally, there are small contributions from quarkonia states produced in association with EW bosons.

5.6 Systematic Uncertainties

With such a low background estimate, the uncertainty for this analysis is almost entirely statistical. However, theoretical and experimental systematic uncertainties are incorporated into this analysis as nuisance parameters.

5.6.1 Theory Uncertainties

The theory uncertainties considered are summarised in Table 5.4. All uncertainties assume a Higgs boson mass of 125 GeV.

Processes	Uncertainty	Value
Signal and $H \rightarrow ZZ^* \rightarrow 4\ell$	Renormalisation and factorisation scales [212, 213]	+7.6% -8.1%
Signal and $H \rightarrow ZZ^* \rightarrow 4\ell$	PDF+ α_s	$\pm 3.1\%$
VBF production $H \rightarrow ZZ^* \rightarrow 4\ell$	QCD scale	+0.4% -0.3%
VBF production $H \rightarrow ZZ^* \rightarrow 4\ell$	PDF+ α_s	$\pm 2.1\%$
WH production $H \rightarrow ZZ^* \rightarrow 4\ell$	QCD scale	+0.5% -0.7%
WH production $H \rightarrow ZZ^* \rightarrow 4\ell$	PDF+ α_s	$\pm 1.9\%$
ZH production $H \rightarrow ZZ^* \rightarrow 4\ell$	QCD scale	+3.9% -3.1%
ZH production $H \rightarrow ZZ^* \rightarrow 4\ell$	PDF+ α_s	$\pm 1.6\%$
ttH production $H \rightarrow ZZ^* \rightarrow 4\ell$	QCD scale	+5.8% -9.2%
ttH production $H \rightarrow ZZ^* \rightarrow 4\ell$	PDF+ α_s	$\pm 3.6\%$
$qq/gg \rightarrow ZZ$	QCD scale and PDF	$\pm 5\%$

Table 5.4: Summary of the theory uncertainties in the analysis. Here VBF stands for vector boson fusion.

5.6.2 Experimental Uncertainties

The ATLAS detector has been observed to respond differently to real proton-proton collisions, than its GEANT4 simulation responds to MC-simulated collisions. This difference is corrected for by applying $|\eta|$ and p_T dependent scale factors, equal to the binned ratio of data to MC. These scale factors are used to correct the muon track-to-vertex-association, reconstruction, identification and isolation efficiencies.

Similar corrections are also applied to correct for mismodelling of the resolution and scale of the muon p_T . The smearing correction is applied to the inner detector and MS p_T measurements independently, while the calibration correction is applied to the final muon p_T .

Other detector-related systematic uncertainties include: the uncertainty on the total integrated luminosity delivered to ATLAS by the LHC, which is estimated to be 3.2% [214]; and the pileup-reweighting scale factor, which corrects for the mismodelling of the inelastic activity in the detector that arises due to the MC set of tuning parameters.

The various uncertainties are described below, along with labels for the different uncertainties, which will be used to refer to them later:

- STAT: Statistical uncertainty due to the limited number of MC sample events
- MUON_EFF_STAT: Statistical error on the muon identification efficiency SF, derived for muons of $p_T > 15$ GeV using a $Z \rightarrow \mu^+\mu^-$ tag-and-probe sample
- MUON_EFF_STAT_LOWPT: Statistical error on the muon identification efficiency SF, derived for muons of $p_T < 15$ GeV using a $J/\psi \rightarrow \mu^+\mu^-$ tag-and-probe sample
- MUON_EFF_SYS: Systematic error on the muon identification efficiency SF, derived for muons of $p_T > 15$ GeV using a $Z \rightarrow \mu^+\mu^-$ tag-and-probe sample
- MUON_EFF_SYS_LOWPT: Systematic error on the muon identification efficiency SF, derived for muons of $p_T < 15$ GeV using a $J/\psi \rightarrow \mu^+\mu^-$ tag-and-probe sample
- MUON_ISO_STAT: Statistical error on the muon isolation efficiency SF derived using a tag-and-probe sample
- MUON_ISO_SYS: Systematic error on the muon isolation efficiency SF derived using a tag-and-probe sample
- MUON_TTVA_STAT: Statistical error on the muon track-to-vertex association efficiency SF derived using a tag-and-probe sample

- MUON_TTVA_SYS: Systematic error on the muon track-to-vertex association efficiency SF derived using a tag-and-probe sample
- MUONS_ID: Systematic error on the smearing of the muons ID track
- MUONS_MS: Systematic error on the smearing of the muons ID track
- MUONS_SCALE: Systematic error on the scale of the muons momentum

Table 5.5 shows the breakdown of systematic uncertainties for various signal models.

With final states containing electrons, the high mass search has additional systematic uncertainties on the reconstruction, isolation and identification efficiencies of electrons. There are also uncertainties on the electron momentum calibration scale, and associated with the smearing of the electron p_T .


Systematic	2 GeV a^0	2 GeV Z_d	15 GeV a^0	15 GeV Z_d
STAT	± 1.4	± 1.9	± 1.4	± 1.9
MUON_EFF_STAT	± 0.6	± 0.7	± 0.7	± 0.7
MUON_EFF_STAT_LOWPT	± 0.2	± 0.2	± 0.3	± 0.2
MUON_EFF_SYS	$^{+2.2}_{-2.1}$	± 2.3	$^{+2.3}_{-2.2}$	± 2.4
MUON_EFF_SYS_LOWPT	± 0.4	± 0.3	± 0.4	± 0.3
MUON_ISO_STAT	± 0.6	± 0.5	± 0.6	± 0.5
MUON_ISO_SYS	± 1.4	± 1.2	± 1.4	± 1.2
MUON_TTVA_STAT	± 1.0	± 0.9	± 1.0	± 0.9
MUON_TTVA_SYS	± 2.0	± 1.5	$^{+2.1}_{-2.0}$	± 1.5
PRW_DATASF	$^{+1.6}_{-1.8}$	$^{+2.5}_{-3.5}$	$^{+1.8}_{-1.6}$	$^{+0.8}_{-0.3}$
MUONS_ID	$^{+0.3}_{-0.2}$	$^{+0.6}_{-0.0}$	± 0.1	$^{+0.3}_{-0.1}$
MUONS_MS	$^{+0.2}_{-0.1}$	$^{+0.6}_{-0.1}$	$^{+0.4}_{-0.1}$	$^{+0.0}_{-0.3}$
MUONS_SCALE	± 0.1	$^{+0.3}_{-0.0}$	± 0.3	$^{+0.2}_{-0.3}$

Table 5.5: Breakdown of the various systematic uncertainties (in %) for the a^0 and Z_d 2 GeV and 15 GeV mass points for the low mass event selection.


5.7 Results and Interpretation

No events were found in the signal region of the low mass search, compatible with the expectation. The Higgs boson mass compatibility requirement was relaxed as a check for

any resonances outside the Higgs boson mass window. No resonances were observed, as the 16 events observed showed no local deviations from the SM estimates, as shown in Figure 5.14.



(a)



(b)

Figure 5.14: (a) 1D and (b) 2D m_X distributions of the observed data in the low mass search after relaxing the Higgs boson mass compatibility requirement [180]. The SM background and three different Z_d mass signal hypotheses are shown in (a), with the signal yields normalised to 1/10 the SM $H \rightarrow ZZ^* \rightarrow 4\ell$ expectation. The green shaded region in (b) shows the events passing the mass compatibility requirement.


For the high mass search, 6 events were observed in total: 0 in the $4e$ channel; 3

in the $2e2\mu$ channel; and 3 in the 4μ channel. This is compatible with the background expectation of 3.9 ± 0.3 (MC stat.) ± 2.0 (data stat.). The largest deviation from the SM is from a single event at $\langle m_{\ell+\ell-} \rangle \approx 20$ GeV, with a local significance of 3.2σ and a global significance [215] of 1.9σ .


Given the absence of evidence for the $H \rightarrow XX \rightarrow 4\ell$ signal process, these results are used to set 95% CL_s upper limits as described in Section 5.4. First, the upper limits on the BRs to the benchmark processes are described, then the upper limits on the fiducial cross-section. The upper limits calculated all use the CL_s frequentist formalism [177] and the profile-likelihood test statistic [170].

95% CL_s upper limits are also placed on the cross-section, normalised to the SM Higgs boson production cross-section, times BR for the two benchmark processes. These are calculated using the full model-dependent efficiencies, Figure 5.8, for each of the benchmark models, described in Section 5.4. These cross-sections are converted to BRs using the theoretical BRs of $X \rightarrow \ell^+\ell^-$ from each benchmark model [41, 185], assuming Type-II fermion couplings with $\tan\beta = 5$ for the 2HDM+S model. Furthermore, the Higgs boson production cross-section is assumed to not deviate significantly from the SM expectation [213]. Figure 5.15 shows the 95% CL_s upper limits on the BR for Higgs boson decays to the two benchmark models.

The model-independent efficiencies described in Section 5.4 are used to calculate model-independent upper limits at the 95% confidence level on the cross-section of any new physics process entering the fiducial region. Figure 5.16 shows the 95% CL_s upper limits on the fiducial cross-section. As no events were observed, the upper limit can be seen to be roughly constant, with a value of about 0.14 fb. This can be compared to a similar result from CMS [204], which used 20.7 fb⁻¹ of pp collision data to set an upper limit of around 0.24 fb over the same mass range. Scaling the CMS upper limit down by the square-root of the luminosity ratio gives an upper limit of around 0.18 fb, which is similar but slightly looser than this upper limit.



(a)



(b)

Figure 5.15: 95% CL_s upper limits on the BR for Higgs boson decays to the two benchmark models: (a) Z_d and (b) a^0 [180]. The shaded areas are the quarkonia veto regions. The step change at 15 GeV in (a) is due to the addition of the $4e$ and $2e2\mu$ final states. The upper limit in (b) does not appear on the plot above the Υ veto as it is greater than 1, i.e. the analysis has no sensitivity in this region. The disjoint limit either side of the quarkonia veto regions is primarily due to changing efficiency in (a) and changing $\text{BR}(a^0 \rightarrow \mu^+ \mu^-)$ in (b).




Figure 5.16: Model-independent 95% CL_s upper limit on cross-section within fiducial volume [180]. The discrete change in the 4μ upper limit is due to the different fiducial regions and analysis selections used in the low and high mass searches.

5.8 Conclusion

Searches were performed for decays of the observed Higgs boson into light bosons from two benchmark models, a^0 and Z_d , using 36.1 fb^{-1} of proton-proton collision data produced at $\sqrt{s} = 13 \text{ TeV}$, recorded using the ATLAS detector and the LHC in 2015 and 2016. Four lepton final states were used, with only muons being used in the low mass search.

No globally significant deviation from the SM was observed. 95% CL_s upper limits were set on a model-independent fiducial cross-section, and on the BR of Higgs boson decays to new resonances predicted in two benchmark models, as a function of the mass of the intermediate exotic boson. 95% CL_s upper limits of around 0.01% were set over the entire Z_d mass range, and upper limits of $\text{BR} < 0.1\%$ on a^0 were set below the J/ψ mass window, as the a^0 decays to τ leptons dominate above this mass range.

CHAPTER 6

SEARCH FOR HIGGS BOSON DECAYS TO A Z BOSON AND A LIGHT HADRONIC RESONANCE

6.1 Introduction

This chapter describes a search for Higgs boson decays to a Z boson, and a light (≤ 4 GeV) resonance, using the full ATLAS Run 2 dataset of 139 fb^{-1} . The Z boson is required to decay to leptons (ℓ), specifically electrons (e) or muons (μ), although the selection also has some acceptance for decays to pairs of τ leptons if they both decay leptonically. Hadronic decays of the light resonance are targeted, and due to its low mass and large boost, they are reconstructed as a single jet of hadrons. This process is illustrated in Figure 6.1. This is a dual interpretation analysis: first, a light Higgs boson, nominally from the 2HDM, or the 2HDM+S; second, a SM charmonium resonance: either η_c or J/ψ .

The analysis strategy for this search starts by designing an event selection using MC, which selects signal events while rejecting background events. This consists of an event-level pre-selection, followed by matching tracks to the jet, and then applying quality and vertex requirements to the tracks. Machine-learning techniques are used to discriminate jets that possess the features of jets initiated by our target signal resonances from the




Figure 6.1: Diagram for Higgs boson decays to a dileptonically decaying Z boson, and a light resonance which decays hadronically. The light resonance, X , has multiple interpretations, as described in §6.1.1 and §6.1.2.

large QCD background at the LHC, using variables defined based on the selected tracks. The signal yield passing this selection is modelled using MC, while the background yield is modelled using a data-driven ABCD estimate, corrected using reweighted MC. A single-bin profile likelihood fit to the number of data events passing this selection is used to interpret the results of the search, in terms of the BSM and SM charmonium interpretations. Various sources of systematic uncertainty are evaluated, and implemented in the likelihood fit as nuisance parameters.

6.1.1 Light Higgs Bosons

In many scenarios for BSM Higgs sectors, there exists a light pseudoscalar or scalar Higgs boson, with large BR to hadronic final states [216]. Two such models are the 2HDM and 2HDM+S, described in §1.4.1 and §1.4.3, respectively. As shown in Figure 1.4, in the mass range being considered, the a^0 BR to light hadronic final states is dominant across most of the low mass phase space. This is because the Yukawa-like coupling of the light Higgs boson, the limited final states available due to kinematic considerations, and the large indirect coupling to gluons entail a dominantly hadronic final state. Furthermore, for low $\tan\beta$ the BR of the a^0 to leptons and down-type quarks is suppressed, potentially

making this decay mode the only reliable probe of these parts of the 2HDM+S phase space, motivating this search.

6.1.2 Charmonium States: J/ψ and η_c

Higgs boson decays to a Z boson and light SM resonances are still unconstrained. However, the potential of searches for Higgs boson decays to bosons and light resonances has been demonstrated by the ATLAS experiment [217, 218], though only ever in exclusive decay modes of the resonance. This is the first search at the LHC for Higgs boson decays to a final state containing a meson which decays to an inclusive hadronic final state. It can provide a test of the SM, and a low Q^2 probe of $H \rightarrow ZZ^*$. This interpretation can also potentially probe the Yukawa coupling of charm quark. This channel is also sensitive to modifications from new physics [219, 220].

6.2 Experimental and Simulated Data Samples

The data used in this search correspond to the full Run 2 dataset, collected by the ATLAS detector between 2015-2018. This represents a total integrated luminosity of 139 fb^{-1} . This corresponds to: 3.2 fb^{-1} of 2015 data; 33.0 fb^{-1} of 2016 data; 44.3 fb^{-1} of 2017 data; and 58.5 fb^{-1} of 2018 data. This dataset is then further processed in order to reduce the size of the dataset, which results in the loss of any events with less than 2 leptons with $p_T > 18 \text{ GeV}$. Finally, it was noticed that some data events were duplicated in the final data files. These events are removed in the analysis-level software.

MC samples were used to guide the development of this analysis, and are used directly in the signal and background modelling. These MC samples are summarised in Table 6.1. All MC samples used in this analysis have been processed in the same way as the data, only keeping events with at least 2 leptons with $p_T > 18 \text{ GeV}$.

Having the largest cross section, MC samples in which Higgs bosons are produced via gluon-gluon Fusion (ggF) are generated for all signal hypotheses in this analysis. The

	Process	Generator	Shower	Notes
Signal	$H \rightarrow Z(\ell^+\ell^-) + \eta_c$	POWHEG+PYTHIA 8	PYTHIA 8+EVTGEN	
	$H \rightarrow Z(\ell^+\ell^-) + J/\psi$	POWHEG+PYTHIA 8	PYTHIA 8+EVTGEN	
	$H \rightarrow Z(\ell^+\ell^-) + a^0$	POWHEG+PYTHIA 8	PYTHIA 8+EVTGEN	$m_{a^0} = 0.5 \text{ GeV}$ $m_{a^0} = 0.75 \text{ GeV}$ $m_{a^0} = 1 \text{ GeV}$ $m_{a^0} = 1.5 \text{ GeV}$ $m_{a^0} = 2 \text{ GeV}$ $m_{a^0} = 2.5 \text{ GeV}$ $m_{a^0} = 3 \text{ GeV}$ $m_{a^0} = 3.5 \text{ GeV}$ $m_{a^0} = 4 \text{ GeV}$ $m_{a^0} = 8 \text{ GeV}$
Background	$Z(\ell^+\ell^-) + jets$	SHERPA 2.2.1		$\max(H_T, p_T^V)$ slices and flavour filtered
	$Z(\ell^+\ell^-) + Z(q\bar{q})$	SHERPA 2.2.2		Gluon induced
	$Z(\ell^+\ell^-) + Z(q\bar{q})$	SHERPA 2.2.1		
	$Z(\ell^+\ell^-) + W(q\bar{q})$	SHERPA 2.2.1		
	$t\bar{t}$	POWHEG	PYTHIA 8+EVTGEN	$\geq 2\ell$

Table 6.1: Signal and background processes simulated with MC for this analysis.

Higgs boson is produced in POWHEG [102–104], using the AZNLO tune [221]. The decay, hadronisation, parton shower and underlying event are modelled using PYTHIA 8 [175] (v8.212) and EVTGEN [222], interfaced to the CT10 [223] and CTEQ6L1 PDF sets [205]. For the BSM signal hypotheses, the SM Higgs boson (pdgId=25) is replaced by the heavy neutral scalar Higgs from the 2HDM (pdgId=35), which is then decayed in PYTHIA 8 to a Z boson and a neutral pseudo-scalar a^0 (pdgId=36). The a^0 is allowed to decay to any final state, to ensure that any decay mode with a significant selection efficiency is considered, and the default PYTHIA 8 2HDM $\tan\beta$ value of 1 is used to generate the decays of the BSM Higgs bosons. These resulting BRs are shown in Table 6.2.

The background for this analysis is dominated by $Z + jets$ events, which is modelled using the ATLAS recommended MC sample [224]. This choice is further motivated by this sample having a NLO ME calculation, and showing the best data to MC agreement of the samples we have studied. The calculation of the hard scatter and parton shower tuning is done with SHERPA 2.2.1 [225] interfaced to the NNPDF [226] PDF sets. The

a^0 Mass	BR
0.5 GeV	gg (92%), $\mu^+\mu^-$ (8%)
0.75 GeV	gg (88%), $\mu^+\mu^-$ (12%)
1 GeV	gg (88%), $\mu^+\mu^-$ (12%)
1.5 GeV	gg (76%), $s\bar{s}$ (16%), $\mu^+\mu^-$ (8%)
2 GeV	gg (82%), $s\bar{s}$ (13%), $\mu^+\mu^-$ (5%)
2.5 GeV	gg (88%), $s\bar{s}$ (8%), $\mu^+\mu^-$ (4%)
3 GeV	gg (86%), $s\bar{s}$ (9%), $\mu^+\mu^-$ (4%)
3.5 GeV	$c\bar{c}$ (88%), gg (10%), $s\bar{s}$ (1%)
4 GeV	$c\bar{c}$ (57%), $\tau^+\tau^-$ (37%), gg (5%)
4.5 GeV	$c\bar{c}$ (52%), $\tau^+\tau^-$ (43%), gg (4%)
5 GeV	$c\bar{c}$ (50%), $\tau^+\tau^-$ (45%), gg (4%)
8 GeV	$\tau^+\tau^-$ (45%), $c\bar{c}$ (40%), gg (14%)
12 GeV	$b\bar{b}$ (81%), $\tau^+\tau^-$ (10%), $c\bar{c}$ (7%), gg (2%)

Table 6.2: BRs of the main decay modes (BR > 1%), for various a^0 mass points. Values are determined in PYTHIA 8 using the default BSMHIGGS $\tan\beta$ value of 1.

samples are sliced according to the maximum of the scalar sum of the p_T of all jets and leptons in the event (H_T), and the p_T of the Z boson, at generator-level. They are further split by the presence of heavy flavour quark induced hadronic showers at generator-level. The inclusive production cross sections are known to NNLO in QCD [227].

The ZZ and ZW processes constitute small (< 1%) backgrounds to this analysis. The diboson backgrounds are also modelled according to the ATLAS recommendations [228]. SHERPA 2.2.1 is interfaced to the NNPDF 3.0 PDF set for the modelling of the hard interaction and parton shower.

The $t\bar{t}$ process constitutes a further, small (< 1%) background to this analysis. The hard interaction for the $t\bar{t}$ background is modelled using POWHEG, while the decay, hadronisation, parton shower and underlying event are modelled using PYTHIA 8 and EVTGEN.

In addition to the nominal background MC samples, the dominant Z +jets background is modelled using an alternative MC generator, in order to have a second estimate of the main background with which to cross check the first. This sample uses events generated from MADGRAPH_AMC@NLO, PYTHIA 8 and EVTGEN. The generator tune is A14, and the PDF set is NNPDF23LO. The $Z \rightarrow e^+e^-$ and $Z \rightarrow \mu^+\mu^-$ samples are sliced based on H_T and heavy flavour filters, while the $Z \rightarrow \tau^+\tau^-$ sample is sliced based on the

number of additional final state particles.

A full simulation of the ATLAS detector [229] in GEANT4 [230] is used to estimate the response of the ATLAS detector in all of the above samples. Data-driven corrections are applied to the event-level trigger efficiency, the jet vertex tagging efficiency, the electron reconstruction, identification and isolation efficiencies, and the muon reconstruction, isolation and track-to-vertex association efficiencies. Pathological events with a weight of magnitude greater than 100 are produced in the SHERPA MC samples that are used in this study. Based on the ATLAS recommendations, the weights of these events are set to 1 in the background distribution plots in this Chapter, and in the background estimate.

Lastly, the design of the substructure-based selection was assisted using additional signal samples, with finely sampled a^0 mass values of: 0.4, 0.5, 0.75, 0.8, 1, 1.2, 1.5, 1.6, 2, 2.4, 2.5, 2.8, 3, 3.2, 3.5, 3.6 and 4 GeV, and a high-statistics Z +jets background sample of 20 million events. These additional samples were generated using PYTHIA 8, and the response of the ATLAS detector was simulated using DELPHES with the ATLAS card. These samples were used to explore different strategies to maximise the substructure-based selection efficiency over the entire a^0 mass range, which was not possible with the limited number of full-simulation samples we had at hand, to ensure that the selection efficiency did not dip significantly for a^0 mass points not used in the optimisation, and as a second validation of the final selection.

6.3 Event Selection

The full event selection can be broken down into three main stages. First, an event-level pre-selection is applied, which targets a di-lepton plus jet (used to reconstruct the target BSM or SM charmonium resonance) final state. This final state is contaminated by a large Z +jets background, and so track-based substructure techniques are used to discriminate between the signal resonances and the QCD background jets. The high resolution of the inner-tracker is required to do this, and so the second stage of the selection involved

ghost-associating [231] tracks to the reconstructed jet, and applying a loose track-selection requirement in order to reject backgrounds from pileup, the underlying event, multiple parton interactions and fake tracks, as described in §6.3.2. Finally, the tracks surviving this track-selection are used to form substructure-based discriminants, which are given to a classification machine learning algorithm, to which we apply a requirement that discriminates signal from background.

6.3.1 Event-Level Pre-Selection

Events are triggered for offline storage using the lowest unscaled single lepton trigger for each period. These are listed in Table 6.3. The use of dilepton triggers was studied, but found to offer less than a 3% gain in signal efficiency, while increasing the background contamination and the uncertainties on the trigger efficiency, and so these triggers were not used for this search. The trigger objects which fired the triggers are required to be matched to the offline leptons, based on a requirement of $\Delta R < 0.1$ between the trigger object and the corresponding offline object.

Electron candidates are reconstructed and identified offline as described in §2.3.3, with the differences pointed out in §2.3.4, using the *Medium* identification efficiency point. An isolation selection is also applied, requiring the electrons to have a transverse energy sum in a cone of $\Delta R < 0.2$ around the electron of less than 20% of its p_T , and the p_T of tracks within a variable-width cone of $\Delta R < \min(0.2, 10 \text{ GeV}/p_T)$ of the electron must be less than 15% of its p_T . Close-by leptons are removed from these isolation cones. The electrons must also have a longitudinal impact parameter ($z_0 \sin \theta$) of less than 0.5 mm with respect to the reconstructed primary vertex, defined as the vertex with the highest sum of square track p_T . They are required to have a p_T of at least 18 GeV due to the way the data is processed, and they must be found in the central body of the detector ($|\eta| < 2.47$), but not the transition region ($1.37 < |\eta| < 1.52$). The leading lepton is required to have a p_T of at least 27 GeV, due to the trigger requirement.

Muon candidates reconstructed offline as described in §2.3.6, using the *Medium* quality

	Triggers
2015	Loosely isolated $p_T > 20$ GeV muon at HLT, matched to $p_T > 15$ GeV muon at L1 $p_T > 40$ GeV muon at HLT <i>Medium</i> identified $p_T > 24$ GeV electron at HLT, matched to $p_T > 20$ GeV electron at L1 <i>Medium</i> identified $p_T > 60$ GeV electron at HLT <i>Loose</i> identified $p_T > 120$ GeV electron at HLT
2016 A	Medium isolated $p_T > 26$ GeV muon at HLT $p_T > 40$ GeV muon at HLT <i>Medium</i> identified $p_T > 26$ GeV electron at HLT, loose isolation requirement <i>Medium</i> identified $p_T > 60$ GeV electron at HLT <i>Loose</i> identified $p_T > 140$ GeV electron at HLT
2016 B-E	Medium isolated $p_T > 26$ GeV muon at HLT $p_T > 50$ GeV muon at HLT <i>Tight</i> identified $p_T > 26$ GeV electron at HLT, loose isolation requirement <i>Medium</i> identified $p_T > 60$ GeV electron at HLT <i>Loose</i> identified $p_T > 140$ GeV electron at HLT
2016 F-L	Medium isolated $p_T > 26$ GeV muon at HLT $p_T > 50$ GeV muon at HLT <i>Tight</i> identified $p_T > 26$ GeV electron at HLT, loose isolation requirement <i>Medium</i> identified $p_T > 60$ GeV electron at HLT <i>Loose</i> identified $p_T > 140$ GeV electron at HLT
2017	Medium isolated $p_T > 26$ GeV muon at HLT $p_T > 50$ GeV muon at HLT <i>Tight</i> identified $p_T > 26$ GeV electron at HLT, loose isolation requirement <i>Medium</i> identified $p_T > 60$ GeV electron at HLT <i>Loose</i> identified $p_T > 140$ GeV electron at HLT
2018	Medium isolated $p_T > 26$ GeV muon at HLT $p_T > 50$ GeV muon at HLT <i>Tight</i> identified $p_T > 26$ GeV electron at HLT, loose isolation requirement <i>Medium</i> identified $p_T > 60$ GeV electron at HLT <i>Loose</i> identified $p_T > 140$ GeV electron at HLT

Table 6.3: Triggers used to select events in data for the full Run 2 dataset. The letters correspond to the different run periods.

efficiency point. An isolation selection is applied, requiring the transverse energy sum in a cone of $\Delta R < 0.2$ around the muon of less than 30% of its p_T , and the p_T of tracks within a variable-width cone of $\Delta R < \min(0.15, 10 \text{ GeV}/p_T)$ of the muon must be less than 15% of its p_T . If an inner detector track is present, the muons must also have a longitudinal impact parameter ($z_0 \sin \theta$) of less than 0.5 mm, and a transverse impact parameter (d_0) of less than 1 mm, with respect to the primary vertex. They are required to have a p_T of at least 18 GeV due to the way the data is processed, and be found within the acceptance of the muon spectrometers ($|\eta| < 2.7$). The leading lepton is required to have a p_T of at least 27 GeV, due to the trigger requirement.

Due to the low mass of the resonance, and the relatively large kinetic energy imparted to it from the mass difference between the Higgs and Z bosons, the resonance is highly boosted. The spread of the a^0 decay products is contained in a cone of width $\Delta R \approx 2m_{a^0}/p_T^{a^0}$, which for our largest mass point (4 GeV) and lowest p_T jet (20 GeV), gives a cone size of 0.4. For this reason, it is reconstructed as a single anti- k_t jet with a radius parameter of 0.4, formed of topological calorimeter clusters and calibrated to the EM energy scale, as described in §2.3.7. Jet energies are corrected for contributions from pile-up interactions using a jet-area based technique, and calibrated using p_T - and η -dependent correction factors determined from simulation, with residual corrections from in situ measurements applied to data and internal jet properties, referred to as the Global Sequential Calibration [123, 232]. Jets are required to satisfy $p_T > 20 \text{ GeV}$, $|\eta| < 2.5$ and the “jet cleaning” selection of Ref. [233]. To reject jets from pile-up interactions, jets with $p_T < 60 \text{ GeV}$ and $|\eta| < 2.4$ are required to have a “Jet Vertex Tagger” [129] value in excess of 0.59.

To avoid double counting, overlapping electrons, muons and jets are then removed, as described in §2.3.11. At least two leptons are required to survive the overlap removal procedure, two of which must be same-flavour opposite-sign leptons. These are required to have an invariant mass compatible with the Z boson, $81 \text{ GeV} < m_{\ell^+\ell^-} < 101 \text{ GeV}$. If multiple same flavour opposite sign lepton pairs mass these requirements, then the pairing

with an invariant mass closest to that of the Z boson is chosen.

The three body system is then required to have an invariant mass passing a loose pre-selection requirement: $m_{j\ell\ell} < 250$ GeV. This variable is used to form the regions from which the background is estimated, thus the large high side window is to allow sufficient background to determine the background normalisation from data. Many events contain multiple jets which could be taken as the resonance candidate, in which case the jet with the highest p_T is chosen. This criteria correctly selects the resonance in 81.8% to 88.3% of cases for the < 4 GeV a^0 signal hypotheses, while not biasing the jet selection in such a way as to create a fake peak in the invariant mass of the three body system, as would be the case for a three body mass based selection. The fraction of events in which the correct jet was selected was also higher than with selecting the jet which gives a three body mass closest to 125 GeV. The correct reconstruction efficiency was also higher than in the three body mass case. Finally, the three body mass is required to be between 120 GeV and 135 GeV. This requirement was chosen to maximise the $S/\sqrt{B}$, as motivated by Figure 6.5(b).

Kinematic distributions for the selected calorimeter jet are shown in Figure 6.2. Kinematic distributions for the leptons chosen to reconstruct the Z boson are shown Figure 6.3. Kinematic distributions for the reconstructed Z boson (the sum of the 2 lepton 4-momenta) are shown in Figure 6.4. Kinematic distributions for the reconstructed Higgs boson (the sum of the 2 lepton and the jet 4-momenta) are shown in Figure 6.5. Various event level angular distributions are shown in Figure 6.6.

The signal peak in Figure 6.5(b) is not centred on 125 GeV due to the calibration of the calorimeter jets used in the construction of this variable being designed for QCD jets, rather than the jets produced from the decay of one of these light resonances. This is demonstrated by Figure 6.7. The central value shifts and resolution of the three body mass peak degrades for high a^0 masses, due to the presence of neutrinos in the decay of the resonance.




Figure 6.2: (a) p_T , (b) mass, (c) η and (d) ϕ of the leading jet in the event, after the full event-level pre-selection. Data, signal and background distributions are shown. The background in these distributions has been reweighted as per §6.4.2.




Figure 6.3: (a) p_T of the leading lepton, (b) p_T of the subleading lepton, (c) η of the leading lepton and (d) η of the subleading lepton used to reconstruct the Z boson, after the full event-level pre-selection. Data, signal and background distributions are shown. The background in these distributions has been reweighted as per §6.4.2.




Figure 6.4: (a) p_T , (b) mass, (c) η and (d) ϕ of the Z boson (the sum of the 2 lepton 4-momenta) in the event, after the full event-level pre-selection. Data, signal and background distributions are shown. The background in these distributions has been reweighted as per §6.4.2.




Figure 6.5: (a) p_T , (b) mass, (c) η and (d) ϕ of the reconstructed Higgs boson (the sum of the 2 lepton and calo jet 4-momenta) in the event, after the full event-level pre-selection. Data, signal and background distributions are shown. The background in these distributions has been reweighted as per §6.4.2.




Figure 6.6: ΔR between (a) the jet and the furthest lepton, (b) the jet and the closest lepton, (c) the leptons and (d) the jet and the Z boson system, after the full event-level pre-selection. Data, signal and background distributions are shown. The background in these distributions has been reweighted as per §6.4.2.





Figure 6.7: $(p_{T, \text{calo jet}} - p_{T, a^0})/p_{T, a^0}$ distribution.

6.3.2 Track Selection


Ultimately, a track-based discriminant will be used to separate the signal resonance from the background QCD jets. These tracks are selected by a method known as ghost-association [231], in which the tracks in the event are assigned a negligible p_T , and then the jet reconstruction algorithm is re-run, on both the calorimeter clusters (as per the original algorithm) and these new ghost-tracks; any tracks clustered in the jet are considered ghost-associated to it.

The majority of tracks ghost-associated to the jet come from pileup, the underlying event, multiple parton interactions and fake tracks. This is demonstrated by matching the tracks to the generator-level particles which caused them, and then following the generator-level record upwards to see if the particle originated from a decay of the signal resonance. The p_T distributions of the tracks ghost-associated to the calorimeter jet, with and without generator-level matching, are shown in Figure 6.8. The p_T distributions of the tracks which are not matched to particles from the signal resonance follow that of pileup, while the signal events have a significantly higher p_T spectrum. This lower p_T spectrum in fakes means that most of the variables constructed from these tracks, which

are dependent on the p_T of the jet, are somewhat resilient to change from the addition of fake tracks.



(a) Generator-Level Matched



(b) Generator-Level Anti-Matched

Figure 6.8: p_T distributions of the tracks ghost-associated to the calorimeter jet, where (a) were and (b) were not matched to the generator-level a^0 or η_c . Only signal distributions are shown here.

Requirements are applied to the tracks to reject the large contamination from fake tracks, loose track quality and track-to-vertex association (TTVA). The track quality


selection requires that the track has $p_T > 500$ MeV, $|\eta| < 2.5$, at least 7 silicon hits, at most one shared module, at most 1 hole in the pixel, and at most 2 holes in the pixels or strips. The TTVA selection requires that $|d_0| < 2$ mm and $|\Delta z_0 \sin \theta| < 3$ mm, where d_0 and z_0 are the transverse and longitudinal impact parameters described in §2.3.2. The effect of this requirement on the track assisted mass [234] is shown in Figure 6.9. The track assisted mass is defined as the mass of the track system multiplied by the ratio of the p_T of the jet as measured in the calorimeter and tracker, to use the higher angular resolution of the inner detector to estimate the mass, and then corrects for the missing contribution from neutral particles using the calorimeter. This is not used later in the analysis, but provides a reasonable proxy for the mass, which illustrates the effectiveness of the track-selection.

The resulting discriminant (§6.3.3) relies on at least 2 tracks being ghost-associated to the jet. As such, any events in which the selected jet has less than two ghost-associated tracks which passing the requirements detailed in this section are discarded.


6.3.3 Track-Based Multi-Layer-Perceptron

A multi-layer perceptron (MLP) is used to discriminate between jets from hadronic X decays, and background multijet production. TMVA is used to train the MLP. As these are thin ($R = 0.4$) jets, the information from the inner tracker is expected to be more useful due to the higher angular resolution. The reconstruction and modelling of the tracks are also better understood than that of the calorimeter objects. We therefore use the set of ghost-associated tracks, selected as per §6.3.2, as the basis for the inputs variables to this MLP.

Variables defined on these tracks are selected based on their ability to separate the various signals from the total background. These variables are summarised in Table 6.4, briefly described in the next paragraph, and displayed in Figures 6.10, 6.11, and 6.12. These variables are chosen to be dimensionless in order to reduce their correlation with event-level kinematic quantities; also minimising the correlation between the MVA output



(a) Without Track Selection



(b) With Track Selection

Figure 6.9: Track assisted mass distributions of the calorimeter jet, (a) without and (b) with the TTVA and track selection requirements. Data, signal and background distributions are shown.

and the event-level kinematic quantities. The correlation coefficient between each of these variables and the three body mass is also shown in Table 6.4, with none having a correlation greater than 20%. What correlation remains between the three body mass and these variables occurs primarily through the transverse momentum of the calorimeter jet.

Variable	Description	$m_{\ell+\ell-j}$ Corr.
$p_T^{\text{lead track}}/p_T^{\text{tracks}}$	Ratio of transverse momentum of the leading track to total	7.5%
$\Delta R^{\text{lead track, calo jet}}$	ΔR between the leading track and the calorimeter jet axis	19.6%
τ_2	NSubJettiness 2 [235]	1.0%
$U_1(0.7)$	Modified energy correlation function (2, 1, 0.7) [236]	9.4%
$M_2(0.3)$	Ratio of modified energy correlation functions (3, 1, 0.3) and (2, 1, 0.3) [236]	12.5%
angularity(2)	Angularity 2 [237]	13.4%

Table 6.4: Variables chosen to discriminate the resonance signals from the combined background. All tracks have been ghost-associated to the calorimeter jet as per §6.3.2. The Pearson correlation coefficient between each of these variables and the three body mass is shown in the last column for the background.




Figure 6.10: (a) ratio of p_T of leading track to the vector sum of the total p_T of the tracks ghost-associated to the jet, and (b) ΔR between the leading track ghost-associated to the jet and the calorimeter jet axis, after the full event-level pre-selection. Data, signal and background distributions are shown. The background in these distributions has been reweighted as per §6.4.2.




Figure 6.11: Modified correlation functions (a) $U_1(0.7)$ and (b) $M_2(0.3)$, after the full event-level pre-selection. Data, signal and background distributions are shown. The background in these distributions has been reweighted as per §6.4.2.




Figure 6.12: (a) NSubJettiness 2 and (b) angularity with an argument of 2, after the full event-level pre-selection. Data, signal and background distributions are shown. τ_2 is defined to be 0 for less than three tracks. The background in these distributions has been reweighted as per §6.4.2.

$p_{\text{T}}^{\text{lead track}}/p_{\text{T}}^{\text{tracks}}$ is the ratio of transverse momentum of the highest p_{T} track to transverse momentum of the ghost-associated track system. $\Delta R^{\text{lead track, calo jet}}$ is the angular separation (ΔR) between the highest p_{T} track and the calorimeter jet axis. τ_2 is the NSubJettiness 2 variable [235], where NSubJettiness is a measure of how well the jet can be described in terms of two sub-jets. It is used due to the fact that most of the decays of the resonances of interest are to pairs of final state partons. It is defined as $\tau_2 = \sum_t p_{\text{T},t} \min(\Delta R_{1,t}, \Delta R_{2,t}) / \sum_t p_{\text{T},t} R_0$, where the sums run over the ghost-associated tracks, the angles $\Delta R_{i,t}$ are between the track and one of two subjets, and R_0 is the radius parameter of the larger jet (0.4). The two subjets are chosen using the exclusive k_t sub-jet algorithm [235]. $U_1(0.7)$ and $M_2(0.3)$ are both modified energy correlation functions [236], designed for quark-gluon discrimination and to target 2-pronged substructure, respectively. $U_1(0.7) = {}_1e_2^{(0.7)}$ and $M_2(0.3) = {}_1e_3^{(0.3)} / {}_1e_2^{(0.3)}$, where ${}_1e_2^{(\beta)} = \sum_{1 \leq i < j \leq n} z_i z_j \theta_{ij}^\beta$ and ${}_1e_3^{(\beta)} = \sum_{1 \leq i < j < k \leq n} z_i z_j z_k \min(\theta_{ij}^\beta, \theta_{ik}^\beta, \theta_{jk}^\beta)$, and z_i is the ratio of the p_{T} of the track, to that of the track system. Lastly, *angularity*(2) is an angularity variable, defined analogously to Ref. [237] as *angularity*(2) = $\sum_t p_{\text{T},t} \sin^2(\pi \theta_t / 2R_0) (1 - \cos(\pi \theta_t / 2R_0))^{-1}$. The form of this variable is motivated by the different matrix elements for resonance- and QCD-induced jets, as detailed in Ref. [237].

This is not a typical classification task, as there is a continuous spectrum of signals being classified against a large background distribution. As such, before being given to a classification algorithm, the input variables are first given to a regression multi-layer-perceptron, as implemented in the TMVA software package [145]. This regression MLP is trained on the 0.5, 0.75, 1, 1.5, 2, 2.5, 3, 3.5 and 4 GeV a^0 signal samples, and is designed to estimate the mass of the a^0 based on the input variables. Only events with $100 \text{ GeV} < m_{\ell^+\ell^-j} < 170 \text{ GeV}$ were used. It uses an architecture of 4 layers of 12 neurons, though various other architectures and hyper-parameters were tried, and these were found to be optimal. Negative weight events are ignored in the training of the MLPs. This regression is then given, along with the original input variables, to a classification MLP. This way, the classifier is indirectly informed as to which signal mass it should

be targetting, so that it can use the other feature variables to target this signal more effectively. The regression output variable is shown in Figure 6.13.

Once trained, a MLP is simply a non-linear function of multiple input variables. Prior to being input to the MLP, the input variables are transformed to the range $[-1, 1]$, which helps convergence during training. The response of each neuron in the first layer is calculated as a sigmoid function of the sum of a constant bias and all of the input variables, where each input variable is multiplied by a neuron-specific weight that is calculated during training. The outputs of each neuron in the second layer is then calculated in the same way from the response of the neurons in the first layer, and another constant bias and set of weights. This continues until the last layer is reached, at which point the output of the MLP is calculated in the same way from the responses of the neurons in the last layer, except that no sigmoid is used this time. For a regression MLP this output variable estimates the target variable, while for a classification MLP this output variable is designed for maximal separation between signal and background.

The constant weights and biases used in the calculation of the MLP response are derived from a training procedure using data of known target: a boolean which defines whether the event is signal or background for classification; or the target variable for regression. Prior to training, as with the MLP application, the input variables are transformed to the range $[-1, 1]$ to help with convergence. The weights and biases are then optimised to minimise the mean squared deviation of the network response from the known target in training data, using the back-propagation procedure. This is done by adjusting the weights by a fixed amount in the direction in the weight-bias space which maximises the gradient of the mean squared deviation: $\mathbf{w}^{(\rho+1)} = \mathbf{w}^{(\rho)} - \eta \nabla_{\mathbf{w}} E$, where $w^{(\rho)}$ are the weights and bias for iteration ρ , E is the mean squared deviation of the network response, and η is the learning rate. The learning rate starts at 0.02, and is decreased by 1% (10%) after each of the first (last) 475 (25) cycles, to assist convergence. The weights and biases go through 500 of these cycles in total. At every tenth cycle, the network is tested for overtraining and lack of convergence, if either is found, the training is terminated early.




Figure 6.13: Output distribution for the regression MLP, after the full event-level pre-selection. Signal and background distributions are shown. The background in these distributions has been reweighted as per §6.4.2.

This output variable of the regression MLP is then given, along with the original input variables, to a classification MLP. This MLP uses the same hyper-parameters as the regression MLP, and an architecture of 2 layers of 6 and 5 neurons, to separate the background from the sum of the following signals: 0.5, 1, 1.5, 2, 2.5, 3, 3.5 and 4 GeV a^0 signal samples. As with the regression MLP, various other architectures and hyper-parameters were tried, and these were found to be optimal. As with the regression MLP, only events with $100 \text{ GeV} < m_{\ell+\ell-j} < 170 \text{ GeV}$ were used. The 0.75 GeV a^0 signal sample was removed from the training because it was found to bias the classifier towards lower mass signal samples, resulting in a deterioration of the performance towards higher masses. The classification output variable is shown in Figure 6.14, and the Receiver Operating Characteristic (ROC) curves for the testing and training events are given in Figure 6.15. The ROC curves show good signal to background discrimination (average $\sim 50\%$ signal efficiency for $\sim 90\%$ background rejection), as well as little-to-no overtraining.

The design of this MVA was assisted, and concepts validated, by the use of PYTHIA 8 MC samples, with a DELPHES [238] detector simulation using the ATLAS card. Signal samples were produced with a^0 masses of: 0.4, 0.5, 0.75, 0.8, 1, 1.2, 1.5, 1.6, 2, 2.4, 2.5, 2.8, 3, 3.2, 3.5, 3.6 and 4 GeV. Each signal sample had 100k signal events, and 20M Z +jets background events were simulated. These samples were not used in any of the results in this note directly, but were used to test the concepts which were then implemented using the nominal MC samples.

MLPs were used as the MVA in this analysis because they showed a good ability to interpolate between the masses of the MC samples provided during the training. This was tested using both DELPHES- and GEANT4-based MC by removing a signal sample from the training of the MVA, and comparing the performance of the MVAs, with and without the inclusion of this signal sample in the training, on this signal sample. Using the 2.2 GeV DELPHES sample, the MLP was found to have an $S/\sqrt{B}$ 11% lower with the sample excluded from the training, as opposed to a 30% loss in $S/\sqrt{B}$ improvement for a BDT. The MLP also showed greater overall performance and less overtraining as




Figure 6.14: Output distribution for the classification MLP, after the full event-level pre-selection. Signal and background distributions are shown. The background in these distributions has been reweighted as per §6.4.2.




Figure 6.15: ROC curves for the testing events and the training events.

compared to the BDT.

A single requirement on the resulting MLP is then used to reject background events. The requirement is chosen to maximise $S/\sqrt{B}$, as this minimises the uncertainty on the signal strength, and the 95% CL_s upper limit, in the high stats limit. However, as there are multiple signal samples, the requirement is chosen to maximise the average $S/\sqrt{B}$, with each signal samples weighted by how much of the a^0 mass phase space is closer to that signal sample than to any other. This also maximises the expected $S/\sqrt{B}$, assuming a flat Bayesian prior on the signal mass. This results in a requirement of $MLP > 0.0524$, which results in a background efficiency of 1.01% for $110 \text{ GeV} < m_{\ell+\ell-j} < 170 \text{ GeV}$, and signal efficiencies as given in Table 6.5. The inclusion of the regression output variable in the classification MVA was found to result in an improvement in $S/\sqrt{B}$ of 13%, when averaged across the signal samples with an $S/\sqrt{B}$ gain due to the MLP greater than unity.

a^0 mass / GeV	MLP Eff (%)	MLP $S/\sqrt{B}$ Gain
0.5	45.8 ± 0.8	5.25
0.75	42.0 ± 0.8	4.82
1	38.3 ± 0.7	4.39
1.5	31.9 ± 0.6	3.66
2	25.5 ± 0.5	2.92
2.5	15.7 ± 0.4	1.80
3	8.32 ± 0.30	0.953
3.5	5.98 ± 0.25	0.685
4	1.96 ± 0.16	0.225
η_c	6.09 ± 0.25	0.698
J/ψ	6.93 ± 0.27	0.794

Table 6.5: Efficiencies of the $MLP > 0.0524$ requirement on each signal sample. This requirement results in a background efficiency of 0.761% for $120 \text{ GeV} < m_{\ell+\ell-j} < 135 \text{ GeV}$. The $S/\sqrt{B}$ gains due to the application of just the MLP requirement are also shown. PYTHIA 8 a^0 BRs are assumed, using the default BSMHIGGS $\tan\beta$ value of 1, as given in Table 6.2.

The performance of the MLP is then tested on both the testing and training samples to check for overtraining. No statistically significant overtraining is seen in any of the signal or background samples. Considering background events with $100 \text{ GeV} < m_{\ell+\ell-j} < 175 \text{ GeV}$, 3100 more events from the training sample than the validation sample pass the MLP re-

quirement. This is the opposite of what would be expected if overtraining were present, though it is consistent with equality within the statistical uncertainty. Considering events with $100 \text{ GeV} < m_{\ell+\ell-j} < 175 \text{ GeV}$, taking the 1.5 GeV signal sample as an example, 43 more events from the training sample than the validation sample pass the MLP requirement. This is also consistent with equality within the statistical uncertainty.

6.3.4 Complete Event Selection

The full selection is summarised in Table 6.6.

Requirement	Details
Triggers	Single lepton triggers requiring $p_T > 27 \text{ GeV}$
Leptons	e or $\mu \geq 2$ with $p_T > 18 \text{ GeV}$
Z boson	2 same-flavour opposite-sign leptons, with $ m_{\ell+\ell-} - m_Z < 10 \text{ GeV}$
Select X -candidate as anti- k_t 4 jet ($p_T^{jet} > 20 \text{ GeV}$), with highest p_T , for which $m_{\ell+\ell-j} < 250 \text{ GeV}$	
> 2 tracks	≥ 2 tracks ghost associated to the calo jet, surviving track selection
Higgs boson	$120 \text{ GeV} < m_{\ell+\ell-j} < 135 \text{ GeV}$
MLP	$MLP > 0.0524$

Table 6.6: Summary of full event selection.

6.4 Signal and Background Modelling

6.4.1 Signal Modelling

The signal efficiency for the selection is taken directly from MC. This scales the expected Higgs production yield, taken as the product of the luminosity of 139 fb^{-1} and the total SM Higgs production cross section of 55.7 pb . The contributions to the total cross section are taken from the LHC Higgs Cross Section Working Group [239]. The total Higgs production cross section is taken as the sum of gluon fusion, vector boson fusion, ZH , WH , $b\bar{b}H$, $t\bar{t}H$ and tH associated production. This is scaled by the branching fraction of the Z boson to electrons, muons or tau-leptons, which is taken from the

Particle Data Group to be 10.1% [240]. Finally, this is scaled by the signal strength: $\mu = \sigma(H) \times \text{BR}(H \rightarrow Za^0)/\sigma_{\text{SM}}(H)$, to give the total number of expected signal events, assuming the default PYTHIA BRs given in Table 6.2. μ is left free in the likelihood fit, as described in the §6.6.1. Table 6.7 shows the expected signal yields for each of the signal hypotheses considered, assuming $\mu = 1$.

a^0 mass / GeV	Total Efficiency (%)	Total Yield (1000×)
0.5	3.27 ± 0.05	25.5 ± 0.4
0.75	2.77 ± 0.05	21.7 ± 0.4
1	2.88 ± 0.05	22.5 ± 0.4
1.5	2.52 ± 0.05	19.7 ± 0.4
2	2.03 ± 0.04	15.86 ± 0.34
2.5	1.332 ± 0.035	10.41 ± 0.27
3	0.712 ± 0.025	5.57 ± 0.20
3.5	0.529 ± 0.022	4.14 ± 0.17
4	0.145 ± 0.012	1.13 ± 0.09
η_c	0.560 ± 0.023	4.38 ± 0.18
J/ψ	0.581 ± 0.023	4.54 ± 0.18

Table 6.7: Efficiencies of the full selection (pre-selection and MLP requirement) and total expected signal yields (assuming $\mu = 1$) for each signal sample. PYTHIA 8 a^0 BRs are assumed, using the default BSMHIGGS $\tan \beta$ value of 1, as given in Table 6.2.

6.4.2 Background Modelling

A semi data-driven background model is used to estimate the SM background content in the signal region (SR), using three steps. First, the simulated background is reweighted to match the data to improve the modelling of key variables. Second, a fully data-driven ABCD estimate of the background in the SR is produced, which assumes no correlation between the three body mass and the MLP output variable. Third, the reweighted MC is used to correct the data-driven ABCD estimate for the correlation between the three body mass and the MLP output variable. Finally, this background estimation method is compared to data in 13 validation regions.

Simulated Background Reweighting

Before being used to construct the background estimate, the simulated data is reweighted to improve the modelling in the key variables used in the ABCD estimate correction. These variables are the three body mass, and the MLP output variable. The modelling of the MLP output variable is improved by improving the modelling of the input variables. All of the variables are reweighted against data in a blinded data region, consisting of the events passing the full selection except either the Higgs boson mass or the MLP requirements, but not passing both $110 \text{ GeV} < m_{\ell+\ell-j} < 155 \text{ GeV}$ and the MLP requirement.

It was observed that the three body mass is well modelled for each given number of ghost-associated tracks. Therefore, the mismodelling in the three body mass distribution is entirely due to the mismodelling in the ghost-associated track multiplicity. Hence, the ghost-associated track multiplicity is reweighted against data to improve the three body mass distribution.

Reweighting the simulated data based on the $U_1(0.7)$ variable was observed to improve the modelling of the other track-based substructure variables input to the MLP. In doing so, this improves the modelling of the regression and classification MLP output variables. However, this introduces a mismodelling in the p_T distribution of the calorimeter jet and the ghost-associated track multiplicity, and thus in the three body mass.

To simultaneously improve the modelling in both the three body mass, and the MLP output variable, a fully-correlated 3D reweighting is applied based on n_{tracks} , $U_1(0.7)$, and p_T^{jet} . The reweighting is performed by applying corrections derived from the ratio of 3D histograms in data and background MC. Each value of n_{tracks} between 2 and 6 has a dedicated bin in the reweighting, events with n_{tracks} of 7 or 8 share a bin, and events with $n_{\text{tracks}} \geq 9$ share a bin. The $U_1(0.7)$ range 0 to 0.25 is split into 25 equal bins 0.01 wide, and one overflow bin is used for events with $U_1(0.7) > 0.25$. The p_T^{jet} region between 20 GeV and 50 GeV is split into 6 bins 5 GeV wide, the region 50 GeV to 60 GeV represents another bin, and the region above 60 GeV represents a final overflow bin. These three distributions, before and after the reweighting procedure is applied, are


shown in Figure 6.16. This results in significant improvements to the modelling of the three body mass, the MLP input variables and the MLP output variables, as shown in Figure 6.17, Figures 6.18 and 6.19, and Figure 6.20, respectively.

ABCD-Based Background Estimation


A semi-data-driven estimate is used to estimate the background contribution to the SR. The first step towards this estimate is to calculate a fully data-driven ABCD estimate of the background contribution in the signal region. To do this, 4 regions are defined in the space of the three body mass and MLP classifier output variables, which are shown in Figures 6.5(b) and 6.14, respectively. Region *A* is the SR, Region *B* shares the same three body mass requirement as the signal region but also requires that $0.0108 < MLP < 0.0524$, Region *C* shares the MLP requirement of the SR but has $155 < m_{\ell+\ell-j} < 175$ GeV, and Region *D* is defined by $0.0108 < MLP < 0.0524$ and $155 < m_{\ell+\ell-j} < 175$ GeV. The region $0.0108 < MLP < 0.0524$ is chosen to contain approximately 10% of the background. An estimate of the background in the SR is then given by $A = BC/D$. This estimate is accurate if the MLP and three body mass variables are uncorrelated, and there is negligible signal contamination in regions *B*, *C* and *D*. The signal contamination in these regions is negligible in this case, as described in §6.4.2.

While the MLP input variables were selected to minimise the correlation with the three body mass, a non-negligible correlation remains. A correction factor is derived to account for this correlation, using the half of the MC events which were not used in the training of the MLP. This correction factor is defined as the ratio of the MC background events in the signal region, to the MC-based ABCD estimate in the SR: $A/(BC/D)$. This correction factor multiplies the data-driven ABCD estimate to produce the final background estimate:


$$A_{\text{SR}}^{\text{ABCD Est.}} = \underbrace{\frac{B_{\text{data}} C_{\text{data}}}{D_{\text{data}}}}_{\text{Data-driven ABCD Estimate}} \times \underbrace{\frac{A_{\text{MC}}}{\frac{B_{\text{MC}} C_{\text{MC}}}{D_{\text{MC}}}}}_{\text{MC-based ABCD Correction Factor}}.$$




(a) Pre-Rewighting




(b) Post-Rewighting




(c) Pre-Rewighting



(d) Post-Rewighting



(e) Pre-Rewighting



(f) Post-Rewighting

Figure 6.16: Distributions of the three variables used to reweight the background simulation, after the full event-level pre-selection, in data and background MC. These variables are the ghost-associated track multiplicity (a) before and (b) after reweighting, the modified correlation variable $U_1(0.7)$ (c) before and (d) after reweighting, and the transverse momentum of the calorimeter jet (e) before and (f) after reweighting. The remaining mismodelling in (b) is due to multiple bins being treated together in the reweighting.





Figure 6.17: Distributions of the three body mass distribution, after the full event-level pre-selection, in data and background MC (a) before and (b) after reweighting.

This results in an expectation of 82400 ± 2900 background events in the SR, where the uncertainty is derived from the statistical uncertainties in the MC and data inputs to the estimate.


This method allows an estimate of the background in the SR, in which only a double ratio of numbers of events are taken from MC. Only taking ratios of events from MC causes background normalisation uncertainties to fully cancel. While the double ratio ensures that any residual mismodelling in the shape of either of the ABCD variables will cancel insofar as the variables can be considered uncorrelated. The Pearson (Spearman) correlation coefficient between the three body mass and the MLP output variable plane, for post-reweighting background MC, is 6.48% (13.0%).

Validation of Background Modelling


The background model is compared to data in 13 validation regions. 15 regions, including the signal region, are defined by values of $m_{\ell+\ell-j}$ of 100-110 GeV, 110-120 GeV, 120-135 GeV, 135-150 GeV, 150-155 GeV, and the *MLP* ranges of > 0.052 , 0.037-0.052 and 0.026-0.037. The two MLP validation regions are defined by the ranges in the MLP




(a) Pre-Rewighting




(b) Post-Rewighting




(c) Pre-Rewighting



(d) Post-Rewighting



(e) Pre-Rewighting



(f) Post-Rewighting

Figure 6.18: Distributions of the variables input to the MLP, after the full event-level pre-selection, in data and background MC. These variables are the ratio of p_T of leading track to the vector sum of the total p_T of the tracks ghost-associated to the jet (a) before and (b) after reweighting, the ΔR between the leading track ghost-associated to the jet and the calorimeter jet axis (c) before and (d) after reweighting, and the NSubJettiness 2 variable (e) before and (f) after reweighting.







Figure 6.19: Distributions of the variables input to the MLP, after the full event-level pre-selection, in data and background MC. These variables are $M_2(0.3)$ (a) before and (b) after reweighting, and $angularity(2)$ (c) before and (d) after reweighting. The modified correlation variable $U_1(0.7)$ is also input to the MLP, but it is used in the reweighting, and shown in Figure 6.16.




(a) Pre-Rewighting



(b) Post-Rewighting



(c) Pre-Rewighting



(d) Post-Rewighting

Figure 6.20: Distributions of the output of the regression MLP (a) before and (b) after the reweighting, and the the output of the classification MLP (c) before and (d) after the reweighting, after the full event-level pre-selection, in data and background MC.

output variable nearest the signal region, which contain equal amounts of background to the signal region. The MC-based correction factors are presented in Table 6.8, which are then used in the calculation of the full background estimates, which are calculated following the procedure described in §6.4.2, and shown in Table 6.9. The data in each of these regions is shown in Table 6.10, along with a comparison with the full background estimates in Figure 6.21. Good agreement is seen between the data and the background estimates in these regions.

MLP Range	$m_{\ell+\ell-j}/\text{GeV}$ Range				
	100-110	110-120	120-135	135-150	150-155
> 0.052	0.350 ± 0.0571	0.443 ± 0.0240	<i>0.702 ± 0.0242</i>	0.866 ± 0.0284	1.03 ± 0.0472
$0.037 - 0.052$	0.797 ± 0.0728	0.899 ± 0.0382	1.01 ± 0.0320	0.989 ± 0.0322	1.05 ± 0.0525
$0.026 - 0.037$	0.966 ± 0.0832	0.989 ± 0.0420	1.03 ± 0.0346	1.00 ± 0.0325	0.966 ± 0.0473

Table 6.8: MC-based correction factors used in the calculation of the background estimates, calculated following the procedure described in §6.4.2, in background estimate validation regions. As the number of events in the 155 – 175 GeV bins are used to calculate the background estimates, the background estimate method can not provide a prediction in these regions. The quoted uncertainties are due to limited MC statistics, and the estimate in the SR is written in italics.

MLP Range	$m_{\ell+\ell-j}/\text{GeV}$ Range				
	100-110	110-120	120-135	135-150	150-155
> 0.052	2190 ± 357	$20\,900 \pm 1140$	<i>82400 ± 2860</i>	$91\,300 \pm 3010$	$31\,200 \pm 1430$
$0.037 - 0.052$	3800 ± 348	$32\,100 \pm 1370$	$89\,200 \pm 2840$	$78\,200 \pm 2570$	$24\,000 \pm 1200$
$0.026 - 0.037$	4430 ± 383	$34\,000 \pm 1450$	$87\,000 \pm 2940$	$76\,100 \pm 2480$	$21\,200 \pm 1040$

Table 6.9: Background estimates, calculated following the procedure described in §6.4.2, in background estimate validation regions. As the number of events in the 155 – 175 GeV bins are used to calculate the background estimates, the background estimate method can not provide a prediction in these regions. The quoted uncertainties are due to limited data and MC statistics, and the estimate in the SR is written in italics.

As a further test of the background modelling strategy, these estimates are evaluated using a alternative Z +jets MC generator: MADGRAPH. The results of this test are given in Table 6.11. This sample has been reweighted using a procedure designed to mitigate the observed mismodelling in the MADGRAPH Z +jets sample, using the p_T of the calorimeter

	$m_{\ell+\ell-j}/\text{GeV}$ Range					
MLP Range	100-110	110-120	120-135	135-150	150-155	155-175
> 0.052	2479	23049	<i>83106</i>	95040	29408	100926
$0.037 - 0.052$	3919	32512	90091	82703	23508	76404
$0.026 - 0.037$	4358	33992	88695	79289	22220	73980

Table 6.10: Data in background estimate validation regions, with the estimate in the SR is written in *italics*.




Figure 6.21: Data and background estimates, calculated following the procedure described in §6.4.2, in background estimate validation regions. As the number of events in the 155 – 175 GeV bins are used to calculate the background estimates, the background estimate method can not provide a prediction in these regions. The uncertainties are due to limited data and MC statistics, and the SR is denoted by red dashed lines. These numbers correspond to those presented in Tables 6.9 and 6.10.

jet, the p_T of the three body system, and the multiplicity of tracks ghost-associated to the calorimeter jet.

Finally, if there is signal contamination in the ABCD regions from which data is taken, this would lead to the background in the SR being overestimated. However, the background estimate is highly robust against such contamination. For example, if the signal from a 1.5 GeV a^0 produced with a $\text{BR}(H \rightarrow Za^0) = 20\%$ is injected into the data, this would lead to a 0.53% increase in the background estimate, demonstrating that this background estimation method is robust against signal contamination.

MLP Range	$m_{\ell+\ell-j}/\text{GeV}$ Range				
	100-110	110-120	120-135	135-155	155-175
> 0.052	2780 ± 282	$24\,600 \pm 872$	<i>$85\,400 \pm 1990$</i>	$96\,400 \pm 2170$	$31\,200 \pm 1010$
$0.037 - 0.052$	3350 ± 301	$30\,400 \pm 1050$	$88\,300 \pm 2190$	$79\,400 \pm 2030$	$21\,700 \pm 823$
$0.026 - 0.037$	4550 ± 388	$35\,200 \pm 1190$	$86\,100 \pm 2190$	$79\,400 \pm 2060$	$20\,900 \pm 795$

Table 6.11: Background estimates, calculated following the procedure described in §6.4.2, in background estimate validation regions, using the alternative MADGRAPH generator for the $Z + \text{jets}$ background. As the number of events in the 155 – 175 GeV bins are used to calculate the background estimates, the background estimate method can not provide a prediction in these regions. The quoted uncertainties are due to limited data and MC statistics, and the estimate in the SR is written in italics.

6.5 Systematic Uncertainties

Systematic uncertainties are the dominant sources of uncertainty for this analysis. The systematic uncertainties relevant to this analysis have been implemented in the statistical model as nuisance parameters (NP). The systematic uncertainties are of two types: modelling uncertainties; and experimental uncertainties, here meaning detector and reconstruction uncertainties. The following subsections describe the systematic uncertainties relevant to this analysis.

6.5.1 Modelling Uncertainties

The following subsections describe the seven modelling uncertainties relevant to this analysis in order of magnitude.

Scale and PDF Uncertainties

The renormalisation and factorisation scale, and PDF, uncertainties are investigated for both the signal and $Z + \text{jets}$ background samples. The half renormalisation scale uncertainty ($\mu_R = 0.5$) is found to be the largest of these by far, and so is implemented in the fit for both the signal and $Z + \text{jets}$ background processes. This is illustrated in Table 6.12. It is implemented using internal weights in the relevant MC samples. The large asymmetry in the renormalisation scale variation is not understood.

Variation	Uncertainty
$\mu_R = 0.5$	5.7%
$\mu_R = 2$	0.67%
$\mu_F = 0.5$	0.53%
$\mu_F = 2$	3.8%
$\mu_R = 0.5 \ \& \ \mu_F = 0.5$	2.5%
$\mu_R = 2 \ \& \ \mu_F = 2$	1.2%
MMHT2014nnlo68cl	1.2%
CT14nnlo	1.1%

Table 6.12: Scale and PDF uncertainties on the total background, evaluated by scaling the $Z + \text{jets}$ background and evaluating the change in the background estimate without a dedicated reweighting.

For the signal samples, the systematic is taken from the change in the signal efficiency under the scale variation. The internal weights used in the derivation of these uncertainties are not present in the first generation of MC signal samples: 0.5 GeV, 2.5 GeV and 8 GeV a^0 , and the η_c samples. As such, the renormalisation scale systematic uncertainty on the missing a^0 signal samples are interpolated if possible, else they are taken from the nearest signal sample. The renormalisation scale systematic uncertainty on the η_c signal sample is taken from the J/ψ signal sample. The renormalisation scale systematic uncertainty

on the signal estimate varies between 0.195% and 2.00% of the total signal normalisation, for the 0-4 GeV a^0 signal samples, and is 1.07% for the η_c and J/ψ samples.

For the $Z + \text{jets}$ background sample, the systematic uncertainty is taken as the change in the background estimate after a dedicated reweighting is applied to the $Z + \text{jets}$ background sample. This dedicated reweighting is calculated analogously to the nominal reweighting, but the $Z + \text{jets}$ MC sample used in the calculation of the reweighting factors has the dominant renormalisation scale variation applied. The renormalisation scale uncertainty on the background normalisation was found to be 4.64%; this was reduced from 5.7% due to the dedicated reweighting, showing that the reweighting is successfully reducing the reliance of the background estimate on the chosen MC samples.

Background Modelling Uncertainty

The hadronisation and ME uncertainties are evaluated for the dominant $Z + \text{jets}$ background, by comparing the background estimate derived with the nominal SHERPA MC sample to the background as estimated using an alternative MADGRAPH MC sample. The only difference in the estimation method is the reweighting, which uses different variables due to the different nature of the mismodelling in MADGRAPH. These three variables are the p_T of the calorimeter jet, the p_T of the three body system, and the multiplicity of tracks ghost-associated to the calorimeter jet. This results in an uncertainty of 3.61% on the background normalisation. The key distributions in the derivation of this estimate are compared for MADGRAPH and SHERPA in Figure 6.22.

Statistical Uncertainty of Background Estimate

One of the largest uncertainties for this analysis is the statistical uncertainty on the background estimate. This is due primarily to MC statistical uncertainty in the ABCD correction described in §6.4.2. However, there is also a smaller contribution from the statistical uncertainty in data on the pre-correction ABCD estimate described in §6.4.2. This results in a total uncertainty of 3.47% on the background estimate.




Figure 6.22: Distributions of (a) the three body mass and (b) the output of the classification MLP, for background MC, where the $Z + \text{jets}$ process is being modelled by MADGRAPH and SHERPA. No reweighting is applied for either background.

Signal Hadronisation Uncertainty

The effect of the signal hadronisation modelling uncertainty on the MLP output has been evaluated by calculating the change in acceptance after reweighting events based on generator-level track multiplicity. These reweightings will be derived from an alternative signal sample produced using HERWIG 7. This is based on the assumption that the largest impact of the modelling uncertainty is on the MLP via the track multiplicity. Due to technical limitations to do with the constituent masses in HERWIG 7, only quark decays are used in the calculation of the scale factors, and the 1.5 GeV a^0 scale factors are used as a proxy for all lower masses.

The hadronisation uncertainty on the signal estimate varies between 4.13% and 17.6% of the total signal normalisation, for the 0-4 GeV a^0 signal samples, 0.575% for the η_c sample, and 27.9% for the J/ψ sample. For the 3.5 GeV a^0 signal, this uncertainty was estimated to be much smaller than for the mass points surrounding it. In order to ensure that the estimate has not fluctuated to a small value for this mass point, the value of the systematic is conservatively interpolated from the surrounding mass points.

Higgs Cross Section Uncertainty

Theory uncertainties from truncation, unknown N³LO PDFs, unknown finite-mass effects, renormalisation and factorisation scales, combined PDF, α_s , and flavour scheme uncertainties (only for tH associated production) are applied to the various Higgs production mode cross sections. These are taken from the official recommendations of the CERN Higgs Cross Section Working Group [37, 38]. The different types of uncertainty are summed in quadrature, then the uncertainties on the various Higgs production processes are summed (weighted by the relevant cross sections) to calculate the total uncertainty. For all uncertainties, a Higgs boson mass of 125 GeV is assumed. The uncertainty on the ggF process is by far the largest contribution to the total uncertainty, with a combined theory, renormalisation and factorisation scale uncertainty of $^{+4.6\%}_{-6.7\%}$, a PDF uncertainty of 1.9%, and an $+\alpha_s$ uncertainty of 2.6%. The uncertainty on the inclusive cross section is $^{+9.4\%}_{-9.3\%}$.

Statistical Uncertainty of Signal Estimate

The MC statistical uncertainty on the signal estimate varies between 1.7% and 7.9% of the total signal normalisation, for the 0-4 GeV a^0 signal samples, 4.0% for the η_c sample, and 3.9% for the J/ψ sample.

Signal Production Modelling Uncertainty

The full inclusive Higgs production cross section is used to normalise the signal yields, while only the ggF production mode is used to model the signal samples. We account for this by applying a systematic uncertainty on the signal acceptance, derived from a comparison of generator-level MC. 100k ggF and 100k VBF events are generated using PYTHIA, the acceptance of these are compared in a generator-level fiducial acceptance. The Higgs boson is more highly boosted in the case of VBF production, which means the individual objects are more likely to pass the minimum p_T requirements. However, the

boost means that the angular separation between the objects is smaller, so they are more likely to fail the overlap removal for VBF production. This leads to the generator-level acceptances differing by just 0.85%. This uncertainty is scaled down to 14% to account for the fraction of Higgs boson events in the SM which are not produced by ggF, leading to an overall systematic uncertainty of just 0.12%.

6.5.2 Experimental Uncertainties

The following subsections describe the seven experimental systematic uncertainties relevant to this analysis in order of magnitude.

Jet Energy Scale Uncertainties

The uncertainties on the jet energy scale are expected to be among the largest experimental uncertainties for this analysis. This uncertainty has many sub-components, including those derived from: in-situ analysis, η calibration, high- p_T jets, pileup, flavour composition, flavour response, b -jets and punch-through jets. A principle component analysis is used to combine the different components of the jet uncertainty where possible, in such a way as to preserve correlations in certain jet-kinematic regions, resulting in 30 nuisance parameters. This results in an uncertainty that varies between $^{+0.41}_{-12}\%$ and $^{+3.9}_{-21}\%$ of the total signal normalisation for the 0-4 GeV a^0 signal hypotheses, an uncertainty of $^{+0.62}_{-19}\%$ for the η_c signal sample normalisation, and an uncertainty of $^{+1.9}_{-12}\%$ for the J/ψ signal sample normalisation.

The impact of the jet energy scale on the signal selection efficiencies is asymmetric due to the requirement applied to the three body mass. If the jet energy scale is increased (decreased), the signal peak in the three body mass distribution is shifted up (down), and more events are lost through the upper (lower) side of the window that are gained through the lower (upper). So a shift up or down in the jet energy scale causes a lowering of the signal efficiency. This causes in an asymmetric likelihood, as shown in Figure 6.23.

This asymmetry causes a deterioration in the expected 95% CL_s upper limit in the event of a null observation, but has almost no effect on the discovery potential.

Pileup Uncertainty

The uncertainty on the pileup distribution affects many aspects of the event, especially the MLP input variables, which in turn affects the efficiency of the MLP requirement. This is modelled by reweighting the pileup distributions in MC to match data. This uncertainty covers the discrepancy seen between predicted and measured inelastic cross-section in the fiducial volume defined by $M_X > 13$ GeV, where M_X is the mass of the non-diffractive hadronic system. This discrepancy arises because of mismodelling of the central activity by the MC tune, so can be incorporated into analysis as an uncertainty on the mean number of hard interactions per bunch crossing that a given MC event corresponds to. This results in a total uncertainty that varies between ${}_{+1.4}^{-1.3}\%$ and ${}_{+1.9}^{-2.0}\%$ of the total signal normalisation for the different 0-4 GeV a^0 signal hypotheses, an uncertainty of ${}_{+0.14}^{-0.59}\%$ for the η_c signal sample normalisation, and an uncertainty of ${}_{+1.5}^{-1.1}\%$ for the J/ψ signal sample normalisation. The anti-correlation between the pileup NP and the signal normalisations is due to the fact that signal jets have less ghost-associated tracks, and so events with higher pileup are less likely to pass the MLP requirement.

Luminosity Uncertainty

The uncertainty on the total integrated luminosity is 1.7%. It is derived from the calibration of the luminosity scale using x - y beam-separation scans, following a methodology similar to that detailed in Ref. [241], and using the LUCID-2 detector for the baseline luminosity measurements [75].

Lepton Uncertainties

Data-driven uncertainties on the reconstruction, identification and isolation of electrons, as well as the reconstruction, isolation and track-to-vertex association of muons are consid-

ered in this analysis. These uncertainties are derived using $Z \rightarrow \ell^+ \ell^-$ events, in addition to $J/\psi \rightarrow \mu^+ \mu^-$ events for muons. As all of these lepton uncertainties small, for convenience they are added in quadrature as though they were uncorrelated, and evaluated as one total lepton uncertainty. This results in a total uncertainty that varies between 0.98% and 1.1% of the signal normalisation for the various 0-4 GeV a^0 signal hypotheses, an uncertainty of 1.1% for the η_c and J/ψ signal sample normalisations.

Jet Vertex Tagging Uncertainties

Requirements are placed on the impact parameters of the jets from the primary vertex. These impact parameters have associated experimental uncertainties, which influence the efficiency of the jet vertex tagging requirement. These uncertainties are between about 0.2% and 0.6% for $|\eta| < 2.4$, and between about 0.1% and 0.2% for $2.4 < |\eta| < 2.5$, depending on the p_T of the jet. This results in an uncertainty that varies between 0.59% and $^{+0.65}_{-0.64}\%$ of the total signal normalisation for the 0-4 GeV a^0 signal hypotheses, an uncertainty of $^{+0.62}_{-0.61}\%$ for the η_c signal sample normalisation, and an uncertainty of 0.63% for the J/ψ signal sample normalisation.

Trigger Efficiency Uncertainty

Data-driven uncertainties on the trigger efficiencies are also considered in this analysis. Due to the multiple triggers used for this analysis, the calculation of these uncertainties is non-trivial, and are calculated using dedicated software. As these are all fairly small uncertainties, for convenience they are added in quadrature as though they were uncorrelated, and evaluated as one total trigger uncertainty. This results in an uncertainty that varies between $^{+0.27}_{-0.11}\%$ and $^{+0.32}_{-0.14}\%$ of the total signal normalisation for the 0-4 GeV a^0 signal hypotheses, an uncertainty of $^{+0.35}_{-0.15}\%$ for the η_c signal sample normalisation, and an uncertainty of $^{+0.33}_{-0.12}\%$ for the J/ψ signal sample normalisation.

6.6 Statistical Interpretation

6.6.1 Statistical Model

The statistical interpretation of the result is performed using a single-bin profile likelihood fit to the signal region. This uses a binned profile likelihood fit to extract the final results from the observed number of events, and the signal and background estimates described in the previous two sections. This likelihood is given by the product of the Poisson probability term for the signal region, and the Gaussian constraints on the various nuisance parameters which represent the systematic uncertainties described in Section 6.5:

$$\mathcal{L} = \text{Pois}\left(N_{\text{SR}}^{\text{D}}; \boldsymbol{\mu} \times S_{\text{SR}}^{\text{MC}} \times \prod_{i \in S} \boldsymbol{\alpha}_i + A_{\text{SR}}^{\text{ABCD Est.}} \times \prod_{i \in B} \boldsymbol{\alpha}_i\right) \times \prod_{i \in S+B} \text{Gaus}\left(1; \boldsymbol{\alpha}_i, \sigma_i\right),$$

where variables in bold are free in the fit. N_{SR}^{D} is the observed number of data events in the signal region, and $S_{\text{SR}}^{\text{MC}}$ is the signal estimate in the signal region as evaluated in §6.4.1. The parameter of interest μ scales the signal in the fit, and is left free. The α parameters represent the nuisance parameters, which model the effect of the systematic uncertainties, and are described in §6.5.1 and §6.5.2. A likelihood ratio test statistic is then defined as

$$q_{\mu} = -2 \ln\left(\mathcal{L}(\mu, \hat{\hat{\alpha}})/\mathcal{L}(\hat{\mu}, \hat{\alpha})\right),$$

where $\hat{\mu}$ and $\hat{\alpha}$ are the values of the parameters which maximise the likelihood, and $\hat{\hat{\alpha}}$ are the values which maximise $\mathcal{L}$ given a certain value of μ . This test statistic is used to measure the compatibility between the background-only model and the data. We then define the local p_0 value as the probability, assuming the background-only model, that we would have observed a test statistic at least as incompatible with the background-only model than the observed test statistic. This is then used to form exclusion intervals using the CL_s method [170, 177].

6.6.2 Asimov Fits

Asimov datasets are defined based on the background hypothesis, and varying levels of signal. These are then fit with the signal plus background model, and the asymptotic approximation is used to produce uncertainties and 95% CL_s upper limits. The uncertainties on the signal normalisation, along with 95% CL_s upper limits on μ , expected in the absence of a signal are summarised in Table 6.13.

a^0 mass / GeV	$\Delta\mu$ (%)	95% CL_s Limit (%)
0.5	32.9	$47.3^{+18.5}_{-13.2}$
0.75	32.8	$52.9^{+20.7}_{-14.8}$
1	34.5	$52.9^{+20.7}_{-14.8}$
1.5	36.7	$59.7^{+23.4}_{-16.7}$
2	41.5	$73.3^{+28.7}_{-20.5}$
2.5	56.5	106^{+42}_{-30}
3	105	209^{+82}_{-59}
3.5	141	281^{+110}_{-79}
4	510	1040^{+410}_{-290}
η_c	132	266^{+104}_{-74}
J/ψ	135	287^{+113}_{-80}

Table 6.13: Uncertainties on the signal strength parameters for $\mu = 1$, and 95% CL_s upper limits on $\sigma(H) \times \text{BR}(H \rightarrow Za^0)/\sigma_{\text{SM}}(H)$ expected in the absence of a signal. $\Delta\mu$ is the post-fit uncertainty on the parameter of interest: $\sigma(H) \times \text{BR}(H \rightarrow Za^0)/\sigma_{\text{SM}}(H)$. All uncertainties are the mean of the upward and downward MINOS uncertainties. PYTHIA 8 a^0 BRs are assumed, using the default BSMHIGGS $\tan\beta$ value of 1, as given in Table 6.2. The η_c and J/ψ BRs are also taken from PYTHIA 8.

The profile likelihood curve, pull plot and correlation plot for the fit to the 1.5 GeV a^0 signal hypothesis, are shown in Figures 6.23, 6.24 and 6.25, respectively. Signal injection tests are performed, in which various amounts of 1.5 GeV a^0 signal are injected into the Asimov dataset, and then fitted to recover the fitted level of signal. The fitted signal consistently recovers the injected signal, as shown in Figure 6.26.

The breakdown of the expected uncertainties on μ are given for three example signal hypotheses in Table 6.14. The total uncertainty is 99.9% systematic, the vast majority of which is due to the background modelling uncertainty.




Figure 6.23: The profile likelihood curve for the fit to the Asimov dataset using the 1.5 GeV a^0 signal hypothesis, with (a) no systematics except the background MC statistical uncertainty and (b) with all systematics, for a dataset with $\mu = 1$. PYTHIA 8 a^0 BRs are assumed, using the default BSMHIGGS $\tan\beta$ value of 1, as given in Table 6.2.

a^0 mass	0.5 GeV	1.5 GeV	2.5 GeV
Total Uncertainty	0.23 (100%)	0.29 (100%)	0.55 (100%)
Total Statistical Uncertainty	0.011 (5.0%)	0.015 (5.0%)	0.028 (5.1%)
Total Systematic Uncertainty	0.23 (99.9%)	0.29 (99.9%)	0.55 (99.9%)
Signal			
Jet Energy Scale	0.036 (16.1%)	0.040 (13.8%)	0.050 (9.1%)
Parton Shower	0.025 (10.9%)	0.037 (12.8%)	0.050 (9.1%)
Higgs Cross Section and Acceptance	0.014 (6.3%)	0.023 (7.8%)	0.043 (7.9%)
Luminosity, Pileup, Trigger, Leptons, & JVT [129]	0.0061 (2.7%)	0.015 (5.2%)	0.032 (5.9%)
MC Statistics	0.0047 (2.1%)	0.015 (5.1%)	0.031 (5.7%)
Renormalisation Scale	0.0022 (1.0%)	0.012 (4.1%)	0.027 (4.9%)
Background			
Renormalisation Scale	0.15 (67.9%)	0.20 (68.0%)	0.37 (68.0%)
Parton Shower and ME	0.12 (52.8%)	0.15 (53.0%)	0.29 (53.0%)
MC Statistics	0.11 (50.8%)	0.15 (50.9%)	0.28 (50.9%)

Table 6.14: Breakdown of the expected uncertainties on μ for 3 a^0 mass hypotheses, derived from fits to the background-only Asimov dataset. The fraction of the total uncertainty is given in parentheses. The uncertainties are evaluated by removing them from the fit, and subtracting the overall uncertainty on μ without that parameter from with it in quadrature. PYTHIA 8 a^0 BRs are assumed, using the default BSMHIGGS $\tan\beta$ value of 1, as given in Table 6.2.




Figure 6.24: The pull plot for the fit to the Asimov dataset using the 1.5 GeV a^0 signal hypothesis, with systematics. PYTHIA 8 a^0 BRs are assumed, using the default BSMHIGGS $\tan \beta$ value of 1, as given in Table 6.2.




Figure 6.25: The correlation matrix for the fit to the Asimov dataset using the 1.5 GeV a^0 signal hypothesis, with systematics. Only entries with values greater than 0.5% are plotted. PYTHIA 8 a^0 BRs are assumed, using the default BSMHIGGS $\tan \beta$ value of 1, as given in Table 6.2.




Figure 6.26: Plot of injected signal $\sigma(H) \times \text{BR}(H \rightarrow ZX)/\sigma_{SM}(H)$ against fitted signal $\sigma(H) \times \text{BR}(H \rightarrow ZX)/\sigma_{SM}(H)$ for the 1.5 GeV a^0 signal hypothesis. PYTHIA 8 a^0 BRs are assumed, using the default BSMHIGGS $\tan \beta$ value of 1, as given in Table 6.2.

6.6.3 Model-Independent Interpretation

The nominal interpretation strategy assumes either a 2HDM+S or SM charmonium signal hypothesis, which although well motivated do not describe all possible processes which can produce this final state. The assumed BRs for these final states in each model leads to model-dependent 95% CL_s upper limits which are only valid for the model under consideration, and not valid for other models.

To generalise the results of this search, model-independent results will be provided under the following assumptions. First, as the focus of this search, only hadronic decays are considered. Second, due to the Yukawa-ordering of the decays of Higgs bosons, only decays to gluons and the heaviest kinematically accessible quark will be considered. Third, due to the low masses of the first generation quarks, decays to these final states will not be considered. Fourth, the systematic uncertainties on each exclusive decay for any given sample, are the same as those for the inclusive decay for that sample. The motivation and justification for the fourth assumption is given in the next paragraph.

The efficiencies are then reevaluated for exclusive decays of the a^0 to gluons, and either s - or c -quarks, as shown in Table 6.15. It can be seen that the efficiencies for quarks are

consistently slightly higher than that for gluons. The 95% CL_s upper limits from the inclusive search are then multiplied by the inclusive signal efficiency, and divided by the signal efficiencies to exclusive gluon or quark final states. These results then represent the 95% CL_s upper limits on $\sigma(H) \times \text{BR}(H \rightarrow Za^0) \times \text{BR}(a^0 \rightarrow q/g)/\sigma_{\text{SM}}(H)$, under the assumption that the systematic uncertainties on the signal acceptance are the same for quark and gluon final states, and are given in Table 6.16. This assumption allows the upper limit on any specific decay of the a^0 to be calculated, by a linear superposition of the two exclusive upper limits, which would not be possible if the systematics were considered exclusively. The impact of the systematic uncertainties on the signal, which may vary between gluon or quark final states, has an expected impact on $\Delta\mu$ of just $\sim 1.94\%$ for the 1.5 GeV a^0 signal hypothesis, justifying the assumption that these uncertainties can be approximated by their inclusive values. The only exception to this is for very low BR decays, such as the 8 GeV a^0 to $c\bar{c}$, in which the signal MC statistical uncertainty could be much larger than in the inclusive case.

a^0 mass / GeV	$a^0 \rightarrow gg$	$a^0 \rightarrow s\bar{s}$	$a^0 \rightarrow c\bar{c}$
0.5	3.61 ± 0.08	-	-
0.75	2.95 ± 0.06	-	-
1	3.31 ± 0.06	-	-
1.5	2.75 ± 0.06	2.96 ± 0.13	-
2	2.14 ± 0.05	2.39 ± 0.13	-
2.5	1.38 ± 0.04	1.66 ± 0.14	-
3	0.736 ± 0.028	0.895 ± 0.094	-
3.5	0.496 ± 0.068	-	0.543 ± 0.024
4	0.170 ± 0.057	-	0.131 ± 0.015

Table 6.15: Efficiencies of the full selection (pre-selection and MLP requirement) for exclusive gluon or quark decays of each signal sample.

6.7 Validation Strategy

A validation region (VR) is defined in the sideband of the MLP variable. It is as close as possible to the signal requirement, containing the same amount of background as passes

a^0 mass / GeV	$a^0 \rightarrow gg$	$a^0 \rightarrow s\bar{s}$	$a^0 \rightarrow c\bar{c}$
0.5	$42.7^{+16.7}_{-11.9}$	-	-
0.75	$49.7^{+19.5}_{-13.9}$	-	-
1	$46.1^{+18.0}_{-12.9}$	-	-
1.5	$54.8^{+21.5}_{-15.3}$	$50.9^{+19.9}_{-14.2}$	-
2	$69.4^{+27.2}_{-19.4}$	$62.1^{+24.3}_{-17.4}$	-
2.5	103^{+40}_{-29}	$85.2^{+33.4}_{-23.8}$	-
3	202^{+79}_{-57}	167^{+65}_{-47}	-
3.5	300^{+117}_{-84}	-	274^{+107}_{-77}
4	891^{+349}_{-249}	-	1150^{+450}_{-322}

Table 6.16: 95% CL_s upper limits on $\sigma(H) \times \text{BR}(H \rightarrow Za^0) \times \text{BR}(a^0 \rightarrow q/g)/\sigma_{\text{SM}}(H)$ expected in the absence of a signal. These results are derived for exclusive gluon or quark decays for each signal sample, using the Asimov dataset.

the signal requirement. This includes all events with $0.0341 < MLP < 0.0524$. This region is used as a first validation of the analysis methods. This VR has an $S/\sqrt{B}$ improvement of less than 0.800 for all signal hypotheses relative to no jet substructure or three body mass requirements, which means that this region can be analysed in data without unblinding the analysis. The full analysis can therefore be performed in the VR, as though it was the SR.

90091 events were observed in this VR, to be compared with the background-only expectation of 89200 ± 2800 events. This region was fit with the full signal plus background hypothesis, including all the systematic uncertainties mentioned in Section 6.5. The priors for the systematics were evaluated separately for this region. When fit with the 1.5 GeV a^0 signal and background hypotheses, the best fit number of signal events was -870, with a best fit number of background events of 89200. The corresponding best fit values for the signal strength parameter is: $\hat{\mu} = 0.14^{+0.55}_{-0.53}$. To test the relevant machinery, the 95% CL_s upper limit on μ was found to be 133%. Figures 6.27 and 6.28 show the resulting pull plot and correlation matrix, respectively.




Figure 6.27: The pull plot for the fit to the data in the validation region using the 1.5 GeV a^0 signal hypothesis, with systematics. PYTHIA 8 a^0 BRs are assumed, using the default BSMHIGGS $\tan \beta$ value of 1, as given in Table 6.2.




Figure 6.28: The correlation matrix for the fit to the data in the validation region using the 1.5 GeV a^0 signal hypothesis, with systematics. Only entries with values greater than 0.5% are plotted. PYTHIA 8 a^0 BRs are assumed, using the default BSMHIGGS $\tan \beta$ value of 1, as given in Table 6.2.

6.8 Results

83106 events were observed in the signal region. This result is compatible with the SM background only expectation of 82400 ± 5600 ($2900 \oplus 4800$) events, where the total uncertainty is followed by the uncertainty due to limited data and MC statistics, and then the systematic uncertainty. In the absence of a significant excess, 95% CL_s upper limits are set on $\sigma(H) \times \text{BR}(H \rightarrow Za^0)/\sigma_{\text{SM}}(H)$, for both the nominal PYTHIA 8 BRs, and the model independent interpretation described in Section 6.6.3. These are given in Tables 6.17 and 6.18, respectively. Due to the lower efficiency of the MLP for the charmonium states, BR limits below 100% can not be set without assuming an enhanced Higgs boson production cross section.

The pull and correlation plots for the fit to the observed data are given for the 1.5 GeV a^0 signal hypothesis in Figures 6.29 and 6.30, respectively. The regression and classification MLP output variables are given in Figures 6.32 and 6.33 for the events in the signal region. Finally, Figure 6.31 shows the three body mass distribution in the MLP signal region, without the three body mass cut applied.

a^0 mass / GeV	$\hat{\mu}$	95% CL _s Limit (%)
0.5	$0.022^{+0.228}_{-0.224}$	49.6
0.75	$0.022^{+0.271}_{-0.260}$	55.6
1	$0.030^{+0.253}_{-0.258}$	55.6
1.5	$0.029^{+0.293}_{-0.290}$	62.6
2	$0.042^{+0.356}_{-0.364}$	76.8
2.5	$0.064^{+0.539}_{-0.553}$	111
3	$0.11^{+1.02}_{-1.04}$	220
3.5	$0.16^{+1.38}_{-1.40}$	295
4	$0.56^{+5.07}_{-5.12}$	1090
η_c	$0.17^{+1.27}_{-1.35}$	279
J/ψ	$0.13^{+1.31}_{-1.30}$	302

Table 6.17: The observed model-dependent signal strength parameters and 95% CL_s upper limits on $\sigma(H) \times \text{BR}(H \rightarrow Za^0)/\sigma_{\text{SM}}(H)$ for the observed dataset. PYTHIA 8 a^0 BRs are assumed, using the default BSMHIGGS $\tan\beta$ value of 1, as given in Table 6.2. The η_c and J/ψ BRs are also taken from PYTHIA 8.




Figure 6.29: The pull plot for the fit to the data in the signal region using the 1.5 GeV a^0 signal hypothesis, with systematics. PYTHIA 8 a^0 BRs are assumed, using the default BSMHIGGS $\tan\beta$ value of 1, as given in Table 6.2.




Figure 6.30: The correlation matrix for the fit to the data in the signal region using the 1.5 GeV a^0 signal hypothesis, with systematics. Only entries with values greater than 0.5% are plotted. PYTHIA 8 a^0 BRs are assumed, using the default BSMHIGGS $\tan \beta$ value of 1, as given in Table 6.2.

a^0 mass / GeV	$a^0 \rightarrow gg$	$a^0 \rightarrow s\bar{s}$	$a^0 \rightarrow c\bar{c}$
0.5	44.8	-	-
0.75	52.2	-	-
1	48.4	-	-
1.5	57.5	53.4	-
2	72.8	65.1	-
2.5	108	89.4	-
3	213	175	-
3.5	315	-	287
4	934	-	1210

Table 6.18: 95% CL_s observed upper limits on $\sigma(H) \times \text{BR}(H \rightarrow Za^0) \times \text{BR}(a^0 \rightarrow q/g)/\sigma_{\text{SM}}(H)$. These results are derived for exclusive gluon or quark decays for each signal sample.






Figure 6.31: Distribution of the three body mass distribution for the unblinded data, reweighted background and various signal hypotheses, after the signal region classification MLP output cut ($MLP > 0.052$), but no three body mass cut.



(a)

Figure 6.32: Distributions of the classification MLP output for the unblinded data, reweighted background and various signal hypotheses, in the signal region.



(a)

Figure 6.33: Distribution of the regression MLP output for the unblinded data, reweighted background and various signal hypotheses, in the signal region.

6.9 Conclusion

A search has been described which is a promising new way to probe the Higgs boson decays to a Z boson and charmonium or 2HDM a^0 final states. It is well-motivated both in the SM, and from its sensitivity to BSM final states. No excess is observed, and so 95% CL_s upper limits are set on $\sigma(H) \times \text{BR}(H \rightarrow Za^0)/\sigma_{\text{SM}}(H)$, with values starting from 44.8%, for the signal hypothesis of a 0.5 GeV a^0 decaying to gluons, and with a value of 279% (302%) for the η_c ($J\psi$) state.

CHAPTER 7

CONCLUSIONS

The discovery of the Higgs boson in 2012, at the end of Run 1 of the LHC, represents both the completion of the experimental validation of the particle content of the SM, and a historic opportunity to probe reality on the most fundamental level. As the only seemingly non-composite scalar in the Standard Model, associated with an all-permeating, ever-present field thought to generate the mass of all other particles, the study of the Higgs boson is a highlight of the LHC physics programme. To date, measurements of the properties of the Higgs boson have been found to be consistent with the Standard Model predictions. However, having been discovered just 3 years before the start of Run 2 of the LHC, ample room remains for new physics to be discovered in the Higgs sector. After a 2 year long shut down, 2015 saw the start of Run 2 of the LHC, which resulted in the collection of 139 fb^{-1} of pp collision data by the ATLAS detector. A significant role was played in both the collection of this dataset, and in the development of the ATLAS software framework. This data was used to perform three searches for new physics in Higgs boson decays, which have been described herein.

The recent observations of $t\bar{t}H$ production [40] and $H \rightarrow b\bar{b}$ decays [242] establish the couplings of the Higgs boson to the third generation quarks. However, despite strong efforts, the couplings of the Higgs boson to the first and second generation quarks still lack experimental evidence. Having the largest Yukawa coupling, the charm quark provides the

best probe of this part of the Higgs sector. Furthermore, the ad-hoc nature of the Yukawa mechanism makes the coupling of the Higgs boson to charm quarks an ideal place to search for new physics, with several well-motivated models predicting values within the reach of the ATLAS detector with the Run 2 dataset. 36.1 fb⁻¹ of data are used to search for Higgs boson decays to pairs of charm quarks, and as no significant excess is observed, a 95% CL_s upper limit is placed at 110× the SM expected rate of $\sigma(pp \rightarrow ZH) \times \text{BR}(H \rightarrow c\bar{c})$ [7].

Various scenarios for physics beyond the SM predict Higgs boson decays into pairs of light resonances, which are well motivated by the considerations of Chapter 1. Two such models include the 2HDM+S, and the Hidden Abelian Higgs models, described in Chapter 5. In the case that these resonances have low masses, they can have significant BRs to di-muons, resulting in a four-muon final state. This final state is searched for using 36.1 fb⁻¹ of data, resulting in no events passing the selection, compatible with the SM expectation. As such, 95% CL_s upper limits are set on a fiducial cross section of around 0.14 fb, in addition to model dependent upper limits [5, 6].

With a dominant BR across most of the 2HDM phase space, and the only significant BR for low values of $\tan\beta$, the hadronic decays of a light pseudoscalar resonance, a^0 , represent a powerful probe of this part of the phase space. The decays of Higgs bosons to a Za^0 final state, where the Z boson decays leptonically, provides a powerful signature to trigger events and reject the large QCD backgrounds at the LHC. By using track-based substructure variables as the input to a machine-learning based classifier, the hadronic decays of the a^0 can be searched for directly. No excess was observed, so 95% CL_s upper limit were set, depending on the mass and decays of the signal hypothesis. The upper limit starts at a $\text{BR}(H \rightarrow Za^0)$ of 44.8% for a 0.5 GeV a^0 decaying to gluons, assuming a SM Higgs production cross section.

Both the LHC and the ATLAS detector are currently preparing for the HL-LHC upgrade, after which the ATLAS detector is expecting to collect 3000 fb⁻¹ of 14 TeV pp collision data. This will lead to a regime of precision Higgs physics, in which many searches and measurements are expecting to make significant gains in sensitivity. Contributions

described herein include testing the radiation hardness of prototype silicon strip sensors for the new ITk, and studies into the prospects for electron and photon identification [3, 4]. Lastly, using this dataset, ATLAS is expecting to set a 95% CL_s upper limit at $6.3\times$ the SM expected rate of $\sigma(pp \rightarrow ZH) \times \text{BR}(H \rightarrow c\bar{c})$, assuming the SM and in the absence of systematic uncertainties [8].

Bibliography

- [1] Cush, “Wikimedia: Standard Model of Elementary Particles.” https://commons.wikimedia.org/wiki/File:Standard_Model_of_Elementary_Particles.svg. Accessed: 25/09/19.
- [2] ATLAS Collaboration, “Improved luminosity determination in pp collisions at $\sqrt{s} = 7$ TeV using the ATLAS detector at the LHC,” *Eur. Phys. J. C*, vol. 73, p. 2518, 2013.
- [3] ATLAS Collaboration, “Technical Design Report for the ATLAS Inner Tracker Strip Detector,” Tech. Rep. CERN-LHCC-2017-005. ATLAS-TDR-025, CERN, Geneva, Apr 2017.
- [4] ATLAS Collaboration, “Expected performance for an upgraded ATLAS detector at High-Luminosity LHC,” Tech. Rep. ATL-PHYS-PUB-2016-026, CERN, Geneva, Oct 2016.
- [5] ATLAS Collaboration, “Search for new light gauge bosons in Higgs boson decays to four-lepton final states in pp collisions at $\sqrt{s} = 8$ TeV with the ATLAS detector at the LHC,” *Phys. Rev. D*, vol. 92, p. 092001, 2015.
- [6] ATLAS Collaboration, “Search for Higgs boson decays to Beyond-the-Standard-Model light bosons in four-lepton events with the ATLAS detector at $\sqrt{s} = 13$ TeV,” Tech. Rep. ATLAS-CONF-2017-042, CERN, Geneva, Jul 2017.
- [7] ATLAS Collaboration, “Search for the Decay of the Higgs Boson to Charm Quarks with the ATLAS Experiment,” *Phys. Rev. Lett.*, vol. 120, no. 21, p. 211802, 2018.
- [8] ATLAS Collaboration, “Prospects for $H \rightarrow c\bar{c}$ using Charm Tagging with the ATLAS Experiment at the HL-LHC,” Tech. Rep. ATL-PHYS-PUB-2018-016, CERN, Geneva, Aug 2018.
- [9] N. Aghanim et al., “Planck 2018 results. VI. Cosmological parameters,” 2018.
- [10] G. Bhattacharyya, “Electroweak Symmetry Breaking and BSM Physics (A Review),” *Pramana*, vol. 72, pp. 37–54, 2009.
- [11] S. P. Martin, “A Supersymmetry primer,” pp. 1–98, 1997. [Adv. Ser. Direct. High Energy Phys.18,1(1998)].
- [12] S. L. Glashow, “Partial Symmetries of Weak Interactions,” *Nucl. Phys.*, vol. 22, pp. 579–588, 1961.
- [13] S. Weinberg, “A Model of Leptons,” *Phys. Rev. Lett.*, vol. 19, pp. 1264–1266, 1967.
- [14] A. Salam, “Weak and electromagnetic interactions,” *Proc. of the 8th Nobel Symposium*, p. 367, 1969.

- [15] A. Salam, in *Elementary Particle Theory*. Stockholm: Almqvist and Wiksell, 1968.
- [16] F. Englert and R. Brout, “Broken Symmetry and the Mass of Gauge Vector Mesons,” *PRL*, vol. 13, p. 321, 1964.
- [17] P. W. Higgs, “Broken Symmetries and the Masses of Gauge Bosons,” *PRL*, vol. 13, p. 508, 1964.
- [18] P. W. Higgs, “Broken Symmetries, Massless Particles and Gauge Fields,” *PL*, vol. 12, p. 132, 1964.
- [19] M. Tanabashi et al., “Review of Particle Physics,” *Phys. Rev. D*, vol. 98, p. 030001, Aug 2018.
- [20] M. Aker et al., “An improved upper limit on the neutrino mass from a direct kinematic method by KATRIN,” 2019.
- [21] Z. Maki, M. Nakagawa, and S. Sakata, “Remarks on the Unified Model of Elementary Particles,” *Progress of Theoretical Physics*, vol. 28, pp. 870–880, 11 1962.
- [22] N. Cabibbo, “Unitary Symmetry and Leptonic Decays,” *Phys. Rev. Lett.*, vol. 10, pp. 531–533, 1963. [648(1963)].
- [23] M. Kobayashi and T. Maskawa, “CP Violation in the Renormalizable Theory of Weak Interaction,” *Prog. Theor. Phys.*, vol. 49, pp. 652–657, 1973.
- [24] M. Gell-Mann, “A schematic model of baryons and mesons,” *Physics Letters*, vol. 8, no. 3, pp. 214 – 215, 1964.
- [25] D. J. Gross and F. Wilczek, “Ultraviolet behavior of non-abelian gauge theories,” *Phys. Rev. Lett.*, vol. 30, pp. 1343–1346, Jun 1973.
- [26] J. Goldstone, A. Salam, and S. Weinberg, “Broken Symmetries,” *Phys. Rev.*, vol. 127, pp. 965–970, 1962.
- [27] 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.*, vol. B716, pp. 1–29, 2012.
- [28] CMS Collaboration, “Observation of a new boson at a mass of 125 GeV with the CMS experiment at the LHC,” *Phys. Lett.*, vol. B716, pp. 30–61, 2012.
- [29] ATLAS Collaboration, “Combined Measurement of the Higgs Boson Mass in pp Collisions at $\sqrt{s} = 7$ and 8 TeV with the ATLAS and CMS Experiments,” *Phys. Rev. Lett.*, vol. 114, p. 191803, 2015.
- [30] ATLAS Collaboration, “Measurements of Higgs boson properties in the diphoton decay channel with 36 fb⁻¹ of pp collision data at $\sqrt{s} = 13$ TeV with the ATLAS detector,” *Phys. Rev.*, vol. D98, p. 052005, 2018.
- [31] ATLAS Collaboration, “Study of the spin and parity of the Higgs boson in diboson decays with the ATLAS detector,” *Eur. Phys. J.*, vol. C75, no. 10, p. 476, 2015. [Erratum: *Eur. Phys. J.*C76,no.3,152(2016)].
- [32] A. M. Sirunyan et al., “Measurements of properties of the Higgs boson decaying into the four-lepton final state in pp collisions at $\sqrt{s} = 13$ TeV,” *JHEP*, vol. 11, p. 047, 2017.

- [33] V. Khachatryan et al., “Search for Higgs boson off-shell production in proton-proton collisions at 7 and 8 TeV and derivation of constraints on its total decay width,” *JHEP*, vol. 09, p. 051, 2016.
- [34] ATLAS Collaboration, “Constraints on off-shell Higgs boson production and the Higgs boson total width in $ZZ \rightarrow 4\ell$ and $ZZ \rightarrow 2\ell 2\nu$ final states with the ATLAS detector,” *Phys. Lett.*, vol. B786, pp. 223–244, 2018.
- [35] ATLAS and CMS Collaborations, “Measurements of the Higgs boson production and decay rates and constraints on its couplings from a combined ATLAS and CMS analysis of the LHC pp collision data at $\sqrt{s} = 7$ and 8 TeV,” *JHEP*, vol. 08, p. 045, 2016.
- [36] T. A. collaboration, “Combined measurements of Higgs boson production and decay using up to 80 fb⁻¹ of proton–proton collision data at $\sqrt{s} = 13$ TeV collected with the ATLAS experiment,” 2019.
- [37] LHC Higgs Cross Section Working Group, S. Dittmaier, C. Mariotti, G. Passarino, and R. Tanaka (Eds.), “Handbook of LHC Higgs Cross Sections: 1. Inclusive Observables,” *CERN-2011-002*, CERN, Geneva, 2011.
- [38] LHC Higgs Cross Section Working Group, “Handbook of LHC Higgs Cross Sections: 2. Differential Distributions,” *CERN-2012-002*, 2012.
- [39] LHC Higgs Cross Section Working Group, S. Heinemeyer, C. Mariotti, G. Passarino, and R. Tanaka (Eds.), “Handbook of LHC Higgs Cross Sections: 3. Higgs Properties,” *CERN-2013-004*, CERN, Geneva, 2013.
- [40] ATLAS Collaboration, “Observation of Higgs boson production in association with a top quark pair at the LHC with the ATLAS detector,” *Phys. Lett.*, vol. B784, pp. 173–191, 2018.
- [41] D. Curtin et al., “Exotic decays of the 125 GeV Higgs boson,” *Phys. Rev. D*, vol. 90, p. 075004, Oct 2014.
- [42] A. Belyaev et al., “LHC discovery potential of the lightest NMSSM Higgs in the $h_1 \rightarrow a_1 a_1 \rightarrow 4\mu$ channel,” *Phys. Rev.*, vol. D81, p. 075021, 2010.
- [43] S. Blunier, G. Cottin, M. A. Díaz, and B. Koch, “Phenomenology of a Higgs triplet model at future e^+e^- colliders,” *Phys. Rev.*, vol. D95, no. 7, p. 075038, 2017.
- [44] J. F. Gunion and H. E. Haber, “The CP conserving two Higgs doublet model: The Approach to the decoupling limit,” *Phys. Rev.*, vol. D67, p. 075019, 2003.
- [45] P. S. B. Dev and A. Pilaftsis, “Natural standard model alignment in the two higgs doublet model,” *Journal of Physics: Conference Series*, vol. 631, p. 012030, jul 2015.
- [46] M. Ilyushin, P. Mandrik, and S. Slabospitsky, “Constraints on the Higgs boson anomalous FCNC interactions with light quarks,” 2019.
- [47] F. J. Botella, G. C. Branco, A. Carmona, M. Nebot, L. Pedro, and M. N. Rebelo, “Physical Constraints on a Class of Two-Higgs Doublet Models with FCNC at tree level,” *JHEP*, vol. 07, p. 078, 2014.
- [48] R. Aaij et al., “Measurement of the $B_s^0 \rightarrow \mu^+\mu^-$ branching fraction and effective lifetime and search for $B^0 \rightarrow \mu^+\mu^-$ decays,” *Phys. Rev. Lett.*, vol. 118, no. 19, p. 191801, 2017.

- [49] S. Ipek, “Perturbative analysis of the electron electric dipole moment and CP violation in two-Higgs-doublet models,” *Phys. Rev.*, vol. D89, no. 7, p. 073012, 2014.
- [50] D. A. Camargo, M. D. Campos, T. B. de Melo, and F. S. Queiroz, “A Two Higgs Doublet Model for Dark Matter and Neutrino Masses,” *Phys. Lett.*, vol. B795, pp. 319–326, 2019.
- [51] L. Wang and X.-F. Han, “A simplified 2HDM with a scalar dark matter and the galactic center gamma-ray excess,” *Phys. Lett.*, vol. B739, pp. 416–420, 2014.
- [52] G. Arcadi, “2HDM portal for Singlet-Doublet Dark Matter,” *Eur. Phys. J.*, vol. C78, no. 10, p. 864, 2018.
- [53] J. Bernon et al., “Scrutinizing the alignment limit in two-Higgs-doublet models: $m_h=125$ GeV,” *Phys. Rev.*, vol. D92, no. 7, p. 075004, 2015.
- [54] ATLAS Collaboration, “Measurements of gluon-gluon fusion and vector-boson fusion Higgs boson production cross-sections in the $H \rightarrow WW^* \rightarrow e\nu\mu\nu$ decay channel in pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector,” *Phys. Lett.*, vol. B789, pp. 508–529, 2019.
- [55] ATLAS Collaboration, “Combined measurement of differential and total cross sections in the $H \rightarrow \gamma\gamma$ and the $H \rightarrow ZZ^* \rightarrow 4\ell$ decay channels at $\sqrt{s} = 13$ TeV with the ATLAS detector,” *Phys. Lett.*, vol. B786, pp. 114–133, 2018.
- [56] H. E. Haber, “The Higgs data and the Decoupling Limit,” in *Proceedings, 1st Toyama International Workshop on Higgs as a Probe of New Physics 2013 (HPNP2013): Toyama, Japan, February 13-16, 2013*, 2013.
- [57] R. Barbieri and A. Strumia, “The ‘LEP paradox’,” in *4th Rencontres du Vietnam: Physics at Extreme Energies (Particle Physics and Astrophysics) Hanoi, Vietnam, July 19-25, 2000*, 2000.
- [58] K. S. Babu and Y. Mimura, “Solving the muon problem in gauge mediated supersymmetry breaking models with flavor symmetry,” 2001.
- [59] The Gfitter Group, “Update of the global electroweak fit and constraints on two-Higgs-doublet models,” *Eur. Phys. J.*, vol. C78, no. 8, p. 675, 2018.
- [60] CMS Collaboration, “Search for the exotic decay of the Higgs boson to two light pseudoscalar bosons with two taus and two muons in the final state at $\sqrt{s} = 8$ TeV,” Tech. Rep. CMS-PAS-HIG-15-011, CERN, Geneva, 2016.
- [61] CMS Collaboration, “Search for exotic decays of the Higgs boson to a pair of new light bosons with two muon and two b jets in final states,” Tech. Rep. CMS-PAS-HIG-14-041, CERN, Geneva, 2016.
- [62] CMS Collaboration, “Search for a very light NMSSM Higgs boson produced in decays of the 125 GeV scalar boson and decaying into τ leptons in pp collisions at $\sqrt{s} = 8$ TeV,” *JHEP*, vol. 01, p. 079, 2016.
- [63] CMS Collaboration, “Search for Higgs Decays to New Light Bosons in Boosted Tau Final States,” Tech. Rep. CMS-PAS-HIG-14-022, CERN, Geneva, 2015.
- [64] CMS Collaboration, “A search for pair production of new light bosons decaying into muons,” *Phys. Lett.*, vol. B752, pp. 146–168, 2016.

- [65] ATLAS Collaboration, “Search for Higgs bosons decaying to aa in the $\mu\mu\tau\tau$ final state in pp collisions at $\sqrt{s}=8$ TeV with the ATLAS experiment,” *Phys. Rev.*, vol. D92, 2015.
- [66] ATLAS Collaboration, “Search for the Higgs boson produced in association with a W boson and decaying to four b -quarks via two spin-zero particles in pp collisions at 13 TeV with the ATLAS detector,” *Eur. Phys. J.*, vol. C76, no. 11, p. 605, 2016.
- [67] ATLAS Collaboration, “Search for Higgs boson decays into a pair of light bosons in the $b\bar{b}\mu\mu$ final state in pp collision at $\sqrt{s}=13$ TeV with the ATLAS detector,” *Phys. Lett.*, vol. B790, pp. 1–21, 2019.
- [68] ATLAS Collaboration, “Search for Higgs boson decays into pairs of light (pseudo)scalar particles in the $\gamma\gamma jj$ final state in pp collisions at $\sqrt{s}=13$ TeV with the ATLAS detector,” *Phys. Lett.*, vol. B782, pp. 750–767, 2018.
- [69] ATLAS Collaboration, “Search for new phenomena in events with at least three photons collected in pp collisions at $\sqrt{s}=8$ TeV with the ATLAS detector,” *Eur. Phys. J.*, vol. C76, no. 4, p. 210, 2016.
- [70] V. M. Abazov et al., “Search for NMSSM Higgs bosons in the $h \rightarrow aa \rightarrow \mu\mu \mu\mu, \mu\mu \tau\tau$ channels using $p\bar{p}$ collisions at $\sqrt{s}=1.96$ TeV,” *Phys. Rev. Lett.*, vol. 103, p. 061801, 2009.
- [71] O. S. Brüning et al., *LHC Design Report*. CERN Yellow Reports: Monographs, Geneva: CERN, 2004.
- [72] M. Lamont, “LHC Performance in Run 2 and Beyond,” *PoS*, vol. LeptonPhoton2015, p. 001. 8 p, 2016.
- [73] I. Bejar Alonso and L. Rossi, “HiLumi LHC Technical Design Report: Deliverable: D1.10,” Tech. Rep. CERN-ACC-2015-0140, Nov 2015.
- [74] ATLAS Collaboration, “Luminosity determination in pp collisions at $\sqrt{s}=13$ TeV using the ATLAS detector at the LHC,” Tech. Rep. ATLAS-CONF-2019-021, CERN, Geneva, Jun 2019.
- [75] G. Avoni et al., “The new LUCID-2 detector for luminosity measurement and monitoring in ATLAS,” *JINST*, vol. 13, no. 07, p. P07017, 2018.
- [76] ATLAS Collaboration, “ATLAS Luminosity Public Results Run 2.” <https://twiki.cern.ch/twiki/bin/view/AtlasPublic/LuminosityPublicResultsRun2>. Accessed: 20/08/19.
- [77] ATLAS Collaboration, “Measurement of the Inelastic Proton-Proton Cross Section at $\sqrt{s}=13$ TeV with the ATLAS Detector at the LHC,” *Phys. Rev. Lett.*, vol. 117, no. 18, p. 182002, 2016.
- [78] ATLAS Collaboration, “The ATLAS Experiment at the CERN Large Hadron Collider,” *JINST*, vol. 3, p. S08003. 437 p, 2008. Also published by CERN Geneva in 2010.
- [79] ATLAS Collaboration, *ATLAS magnet system: Technical Design Report, 1*. Technical Design Report ATLAS, Geneva: CERN, 1997.
- [80] ATLAS Collaboration, “ATLAS Magnet System.” <http://atlas-ma.web.cern.ch/atlas-ma/>. Accessed: 20/08/19.

- [81] ATLAS Collaboration, “ATLAS inner detector: Technical design report. Vol. 1,” 1997.
- [82] ATLAS Collaboration, “ATLAS inner detector: Technical design report. Vol. 2,” 1997.
- [83] ATLAS Collaboration, “Performance of the ATLAS Track Reconstruction Algorithms in Dense Environments in LHC Run 2,” *Eur. Phys. J.*, vol. C77, no. 10, p. 673, 2017.
- [84] M Capeans et al., “ATLAS Insertable B-Layer Technical Design Report,” Tech. Rep. CERN-LHCC-2010-013. ATLAS-TDR-19, Sep 2010.
- [85] A. Vogel, “ATLAS Transition Radiation Tracker (TRT): Straw Tube Gaseous Detectors at High Rates,” Tech. Rep. ATL-INDET-PROC-2013-005, CERN, Geneva, Apr 2013.
- [86] ATLAS Collaboration, “Performance of the ATLAS Transition Radiation Tracker in Run 1 of the LHC: tracker properties,” *JINST*, vol. 12, no. 05, p. P05002, 2017.
- [87] ATLAS Collaboration, *ATLAS liquid-argon calorimeter: Technical Design Report*. Technical Design Report ATLAS, Geneva: CERN, 1996.
- [88] ATLAS Collaboration, “Topological cell clustering in the ATLAS calorimeters and its performance in LHC Run 1,” *Eur. Phys. J.*, vol. C77, p. 490, 2017.
- [89] J. Pequeno, “Computer Generated image of the ATLAS calorimeter.” Mar 2008.
- [90] ATLAS Collaboration, *ATLAS tile calorimeter: Technical Design Report*. Technical Design Report ATLAS, Geneva: CERN, 1996.
- [91] ATLAS Collaboration, “ATLAS muon spectrometer: Technical design report,” 1997.
- [92] ATLAS Collaboration, “Performance of the ATLAS Trigger System in 2015,” *Eur. Phys. J.*, vol. C77, no. 5, p. 317, 2017.
- [93] ATLAS Collaboration, “Trigger Menu in 2017,” Tech. Rep. ATL-DAQ-PUB-2018-002, CERN, Geneva, Jun 2018.
- [94] P. Jenni, M. Nelli, M. Nordberg, and K. Smith, *ATLAS high-level trigger, data-acquisition and controls: Technical Design Report*. Technical Design Report ATLAS, Geneva: CERN, 2003.
- [95] ATLAS Collaboration, *ATLAS level-1 trigger: Technical Design Report*. Technical Design Report ATLAS, Geneva: CERN, 1998.
- [96] ATLAS Collaboration, *ATLAS Computing: technical design report*. Technical Design Report ATLAS, Geneva: CERN, 2005.
- [97] T. Gleisberg et al., “Event generation with SHERPA 1.1,” *JHEP*, vol. 02, p. 007, 2009.
- [98] T. Sjostrand, S. Mrenna, and P. Z. Skands, “A Brief Introduction to PYTHIA 8.1,” *Comput. Phys. Commun.*, vol. 178, pp. 852–867, 2008.
- [99] M. B. et al., “Herwig++ physics and manual,” *The European Physical Journal C*, vol. 58, pp. 639–707, Dec 2008.

- [100] J. Bellm et al., “Herwig 7.0/Herwig++ 3.0 release note,” *Eur. Phys. J.*, vol. C76, no. 4, p. 196, 2016.
- [101] J. Alwall et al., “MadGraph 5 : Going Beyond,” *JHEP*, vol. 06, p. 128, 2011.
- [102] P. Nason, “A New method for combining NLO QCD with shower Monte Carlo algorithms,” *JHEP*, vol. 11, p. 040, 2004.
- [103] S. Frixione, P. Nason, and C. Oleari, “Matching NLO QCD computations with Parton Shower simulations: the POWHEG method,” *JHEP*, vol. 11, p. 070, 2007.
- [104] 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*, vol. 06, p. 043, 2010.
- [105] ATLAS Collaboration, “The ATLAS Simulation Infrastructure,” *Eur. Phys. J.*, vol. C70, pp. 823–874, 2010.
- [106] S. Agostinelli et al. (GEANT4 Collaboration), “GEANT4: A simulation toolkit,” *Nucl. Instrum. Meth.*, vol. A506, pp. 250–303, 2003.
- [107] ATLAS Collaboration, “The simulation principle and performance of the ATLAS fast calorimeter simulation FastCaloSim,” 2010.
- [108] ATLAS Collaboration, “ATLAS Data Quality Shifter Manual Combined Performance Reference Histograms Web Page.” Accessed: 22/09/19 (Unpublished).
- [109] R. Frühwirth, “Application of kalman filtering to track and vertex fitting,” *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 262, no. 2, pp. 444 – 450, 1987.
- [110] ATLAS Collaboration, “Early Inner Detector Tracking Performance in the 2015 data at $\sqrt{s} = 13$ TeV,” Tech. Rep. ATL-PHYS-PUB-2015-051, CERN, Geneva, Dec 2015.
- [111] ATLAS Collaboration, “Performance of the ATLAS Inner Detector Track and Vertex Reconstruction in the High Pile-Up LHC Environment,” Tech. Rep. ATLAS-CONF-2012-042, CERN, Geneva, Mar 2012.
- [112] ATLAS Collaboration, “Vertex Reconstruction Performance of the ATLAS Detector at $\sqrt{s} = 13$ TeV,” Tech. Rep. ATL-PHYS-PUB-2015-026, CERN, Geneva, Jul 2015.
- [113] 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.*, vol. C79, no. 8, p. 639, 2019.
- [114] ATLAS Collaboration, “Electron reconstruction and identification efficiency measurements with the ATLAS detector using the 2011 LHC proton-proton collision data,” *Eur. Phys. J.*, vol. C74, no. 7, p. 2941, 2014.
- [115] ATLAS Collaboration, “Electron efficiency measurements with the ATLAS detector using the 2015 LHC proton-proton collision data,” Tech. Rep. ATLAS-CONF-2016-024, CERN, Geneva, Jun 2016.
- [116] ATLAS Collaboration, “Electron efficiency measurements with the ATLAS detector using 2012 LHC proton–proton collision data,” *Eur. Phys. J. C*, vol. 77, p. 195, 2017.
- [117] ATLAS Collaboration, “Electron and photon performance measurements with the ATLAS detector using the 2015-2017 LHC proton-proton collision data,” 2019.

- [118] M. Aaboud *et al.*, “Measurement of the photon identification efficiencies with the ATLAS detector using LHC Run 2 data collected in 2015 and 2016,” *Eur. Phys. J.*, vol. C79, no. 3, p. 205, 2019.
- [119] ATLAS Collaboration, “Measurement of the muon reconstruction performance of the ATLAS detector using 2011 and 2012 LHC proton–proton collision data,” *Eur. Phys. J.*, vol. C74, no. 11, p. 3130, 2014.
- [120] ATLAS Collaboration, “Muon reconstruction performance of the ATLAS detector in proton–proton collision data at $\sqrt{s} = 13$ TeV,” *Eur. Phys. J. C*, vol. 76, no. 5, p. 292, 2016.
- [121] ATLAS Collaboration, “Properties of jets and inputs to jet reconstruction and calibration with the ATLAS detector using proton–proton collisions at $\sqrt{s} = 13$ TeV.” ATL-PHYS-PUB-2015-036, 2015.
- [122] M. Cacciari, G. P. Salam, and G. Soyez, “The anti- k_t jet clustering algorithm,” *JHEP*, vol. 04, p. 063, 2008.
- [123] ATLAS Collaboration, “Jet energy scale measurements and their systematic uncertainties in proton-proton collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector,” *Phys. Rev.*, vol. D96, no. 7, p. 072002, 2017.
- [124] ATLAS Collaboration, “Jet energy scale and uncertainties in 2015-2017 data and simulation.” <https://atlas.web.cern.ch/Atlas/GROUPS/PHYSICS/PLOTS/JETM-2018-006/>. Accessed: 20/08/19.
- [125] “Expected performance of the ATLAS b -tagging algorithms in Run-2,” Tech. Rep. ATL-PHYS-PUB-2015-022, CERN, Geneva, Jul 2015.
- [126] ATLAS Collaboration, “Performance of algorithms that reconstruct missing transverse momentum in $\sqrt{s} = 8$ TeV proton-proton collisions in the ATLAS detector,” *Eur. Phys. J.*, vol. C77, no. 4, p. 241, 2017.
- [127] ATLAS Collaboration, “Measurement of the tau lepton reconstruction and identification performance in the ATLAS experiment using pp collisions at $\sqrt{s} = 13$ TeV,” Tech. Rep. ATLAS-CONF-2017-029, CERN, Geneva, May 2017.
- [128] T. Masubuchi, J. Benitez, and A. S. Bell, “Search for a Standard Model Higgs boson produced in association with a vector boson and decaying to a pair of b -quarks,” Tech. Rep. ATL-COM-PHYS-2016-1724, CERN, Geneva, Nov 2016.
- [129] ATLAS Collaboration, “Tagging and suppression of pileup jets with the ATLAS detector,” Tech. Rep. ATLAS-CONF-2014-018, CERN, Geneva, May 2014.
- [130] ATLAS Collaboration, “ATLAS Phase-II Upgrade Scoping Document,” Tech. Rep. CERN-LHCC-2015-020. LHCC-G-166, CERN, Geneva, Sep 2015.
- [131] ATLAS Collaboration, “ATLAS Event Displays from Upgrade Physics Simulated Data.” <https://twiki.cern.ch/twiki/bin/view/AtlasPublic/UpgradeEventDisplays>. Accessed: 20/08/19.
- [132] ATLAS Collaboration, “The HL-LHC Project.” <http://hilumilhc.web.cern.ch/about/hl-lhc-project>. Accessed: 20/08/19.
- [133] P. Azzi *et al.*, “Standard Model Physics at the HL-LHC and HE-LHC,” Tech. Rep. arXiv:1902.04070, CERN, Geneva, 2018.

- [134] ATLAS Collaboration, “Expected Tracking Performance of the ATLAS Inner Tracker at the HL-LHC,” Tech. Rep. ATL-PHYS-PUB-2019-014, CERN, Geneva, Mar 2019.
- [135] S. Tang et al., “The development of Global Feature eXtractor (gFEX) - the ATLAS calorimeter Level 1 trigger for ATLAS at High Luminosity LHC,” Tech. Rep. ATL-DAQ-PROC-2017-035, CERN, Geneva, Oct 2017.
- [136] T. Iizawa, “The ATLAS Fast Tracker System,” *PoS*, vol. TWEPP-17, p. 139, 2017.
- [137] ATLAS Collaboration, “Technical Design Report for the Phase-II Upgrade of the ATLAS LAr Calorimeter,” Tech. Rep. CERN-LHCC-2017-018. ATLAS-TDR-027, CERN, Geneva, Sep 2017.
- [138] ATLAS Collaboration, “Technical Design Report for the Phase-II Upgrade of the ATLAS Tile Calorimeter,” Tech. Rep. CERN-LHCC-2017-019. ATLAS-TDR-028, CERN, Geneva, Sep 2017.
- [139] B. Stelzer, “The new small wheel upgrade project of the atlas experiment,” *Nuclear and Particle Physics Proceedings*, vol. 273-275, pp. 1160 – 1165, 2016. 37th International Conference on High Energy Physics (ICHEP).
- [140] C. Allaire, “A High-Granularity Timing Detector in ATLAS: Performance at the HL-LHC,” *Nucl. Instrum. Meth.*, vol. A924, pp. 355–359, 2019.
- [141] P. Allport et al., “Recent results and experience with the birmingham MC40 irradiation facility,” *Journal of Instrumentation*, vol. 12, pp. C03075–C03075, mar 2017.
- [142] J. Gao et al., “CT10 next-to-next-to-leading order global analysis of QCD,” *Phys. Rev.*, vol. D89, no. 3, p. 033009, 2014.
- [143] H.-L. Lai et al., “New parton distributions for collider physics,” *Phys. Rev.*, vol. D82, p. 074024, 2010.
- [144] ATLAS Collaboration, “Summary of ATLAS Pythia 8 tunes,” Tech. Rep. ATL-PHYS-PUB-2012-003, CERN, Geneva, Aug 2012.
- [145] A. Hocker et al., “TMVA - Toolkit for Multivariate Data Analysis with ROOT: Users guide. TMVA - Toolkit for Multivariate Data Analysis,” Tech. Rep. physics/0703039, CERN, Geneva, Mar 2007. TMVA-v4 Users Guide: 135 pages, 19 figures, numerous code examples and references.
- [146] N. Ilic, “Performance of the ATLAS Liquid Argon Calorimeter after three years of LHC operation and plans for a future upgrade,” *Journal of Instrumentation*, vol. 9, pp. C03049–C03049, mar 2014.
- [147] ATLAS Collaboration, “Electron identification discriminating variables and efficiency measurements in 2017 data : evolution versus pileup.” <https://atlas.web.cern.ch/Atlas/GROUPS/PHYSICS/PLOTS/EGAM-2018-005/index.html>. Accessed: 12/11/19.
- [148] A. Djouadi, J. Kalinowski, and M. Spira, “HDECAY: A program for Higgs boson decays in the Standard Model and its supersymmetric extension,” *Comput. Phys. Commun.*, vol. 108, pp. 56–74, 1998.
- [149] C. Delaunay, T. Golling, G. Perez, and Y. Soreq, “Enhanced Higgs boson coupling to charm pairs,” *Phys. Rev.*, vol. D89, no. 3, p. 033014, 2014.

- [150] A. Lenz, M. Spannowsky, and G. Tetlalmatzi-Xolocotzi, “Double-charming Higgs boson identification using machine-learning assisted jet shapes,” *Phys. Rev.*, vol. D97, no. 1, p. 016001, 2018.
- [151] S. Bar-Shalom and A. Soni, “Universally enhanced light-quarks Yukawa couplings paradigm,” *Phys. Rev.*, vol. D98, no. 5, p. 055001, 2018.
- [152] G. Abbas, “Solving the fermionic mass hierarchy of the standard model,” *Int. J. Mod. Phys.*, vol. A34, no. 20, p. 1950104, 2019.
- [153] ATLAS and CMS Collaborations, “Measurements of the Higgs boson production and decay rates and constraints on its couplings from a combined ATLAS and CMS analysis of the LHC pp collision data at $\sqrt{s} = 7$ and 8 TeV,” *JHEP*, vol. 08, p. 045, 2016.
- [154] ATLAS Collaboration, “Evidence for the $H \rightarrow b\bar{b}$ decay with the ATLAS detector,” *JHEP*, vol. 12, p. 024, 2017.
- [155] G. T. Bodwin, F. Petriello, S. Stoynev, and M. Velasco, “Higgs boson decays to quarkonia and the $H\bar{c}c$ coupling,” *Phys. Rev.*, vol. D88, no. 5, p. 053003, 2013.
- [156] A. Collaboration, “Search for Higgs and Z Boson Decays to $J/\psi\gamma$ and $\Upsilon(nS)\gamma$ with the ATLAS Detector,” *Phys. Rev. Lett.*, vol. 114, no. 12, p. 121801, 2015.
- [157] ATLAS Collaboration, “Search for Higgs and Z Boson Decays to $\phi\gamma$ with the ATLAS Detector,” *Phys. Rev. Lett.*, vol. 117, no. 11, p. 111802, 2016.
- [158] ATLAS Collaboration, “Search for exclusive Higgs and Z boson decays to $\phi\gamma$ and $\rho\gamma$ with the ATLAS detector,” *JHEP*, vol. 07, p. 127, 2018.
- [159] G. Perez, Y. Soreq, E. Stamou, and K. Tobioka, “Constraining the charm Yukawa and Higgs-quark coupling universality,” *Phys. Rev.*, vol. D92, no. 3, p. 033016, 2015.
- [160] A. Falkowski, F. Riva, and A. Urbano, “Higgs at last,” *JHEP*, vol. 11, p. 111, 2013.
- [161] ATLAS Collaboration, “Luminosity determination in pp collisions at $\sqrt{s} = 7$ TeV using the ATLAS detector at the LHC,” *Eur. Phys. J. C*, vol. 71, p. 1630, 2011.
- [162] R. D. Ball et al., “Parton distributions for the LHC Run II,” *JHEP*, vol. 04, p. 040, 2015.
- [163] K. Hamilton, P. Nason, and G. Zanderighi, “MINLO: Multi-Scale Improved NLO,” *JHEP*, vol. 10, p. 155, 2012.
- [164] G. Luisoni, P. Nason, C. Oleari, and F. Tramontano, “ $HW^\pm/HZ + 0$ and 1 jet at NLO with the POWHEG BOX interfaced to GoSam and their merging within MiNLO,” *JHEP*, vol. 10, p. 083, 2013.
- [165] ATLAS Collaboration, “Measurement of the Z/γ^* boson transverse momentum distribution in pp collisions at $\sqrt{s} = 7$ TeV with the ATLAS detector,” *JHEP*, vol. 2014, p. 55, 2014.
- [166] T. Gleisberg and S. Höche, “Comix, a new matrix element generator,” *JHEP*, vol. 0812, p. 039, 2008.
- [167] F. Cascioli, P. Maierhofer, and S. Pozzorini, “Scattering Amplitudes with Open Loops,” *Phys. Rev. Lett.*, vol. 108, p. 111601, 2012.

- [168] S. Schumann and F. Krauss, “A Parton shower algorithm based on Catani-Seymour dipole factorisation,” *JHEP*, vol. 0803, p. 038, 2008.
- [169] K. Hamilton, P. Richardson, and J. Tully, “A Modified CKKW matrix element merging approach to angular-ordered parton showers,” *JHEP*, vol. 11, p. 038, 2009.
- [170] G. Cowan, K. Cranmer, E. Gross, and O. Vitells, “Asymptotic formulae for likelihood-based tests of new physics,” *Eur. Phys. J. C*, vol. 71, p. 1554, 2011. [Erratum: *Eur. Phys. J. C* **73** (2013) 2501].
- [171] M. R. F. James, “Minuit: A System for Function Minimization and Analysis of the Parameter Errors and Correlations,” *Comput. Phys. Commun.* *10*, pp. 343–367, 1975.
- [172] ATLAS Collaboration, “Letter of Intent for the Phase-II Upgrade of the ATLAS Experiment,” Tech. Rep. CERN-LHCC-2012-022. LHCC-I-023, CERN, Geneva, Dec 2012.
- [173] A. Collaboration, “Technical Design Report for the ATLAS Inner Tracker Pixel Detector,” Tech. Rep. CERN-LHCC-2017-021. ATLAS-TDR-030, CERN, Geneva, Sep 2017.
- [174] D. de Florian et al., “Handbook of LHC Higgs Cross Sections: 4. Deciphering the Nature of the Higgs Sector,” 2016.
- [175] T. Sjöstrand, S. Mrenna, and P. Z. Skands, “A Brief Introduction to PYTHIA 8.1,” *Comput. Phys. Commun.*, vol. 178, pp. 852–867, 2008.
- [176] ATLAS Collaboration, “Projections for measurements of Higgs boson signal strengths and coupling parameters with the ATLAS detector at a HL-LHC,” Tech. Rep. ATL-PHYS-PUB-2014-016, CERN, Geneva, Oct 2014.
- [177] A. L. Read, “Presentation of search results: The CL(s) technique,” *J. Phys.*, vol. G28, pp. 2693–2704, 2002. [11(2002)].
- [178] ATLAS Collaboration, “ATLAS page for the Pixel TDR plots approved by LHCC in April 2018,” tech. rep., CERN, Geneva, Apr 2018.
- [179] H. Abramowicz *et al.*, “Higgs physics at the CLIC electron–positron linear collider,” *Eur. Phys. J.*, vol. C77, no. 7, p. 475, 2017.
- [180] ATLAS Collaboration, “Search for Higgs boson decays to beyond-the-Standard-Model light bosons in four-lepton events with the ATLAS detector at $\sqrt{s} = 13$ TeV,” *JHEP*, vol. 06, p. 166, 2018.
- [181] P. Fayet, “Light spin 1/2 or spin 0 dark matter particles,” *Phys.Rev.*, vol. D 70, p. 023514, 2004.
- [182] D. P. Finkbeiner and N. Weiner, “Exciting Dark Matter and the INTEGRAL/SPI 511 keV signal,” *Phys.Rev.*, vol. D 76, p. 083519, 2007.
- [183] N. Arkani-Hamed, D. P. Finkbeiner, T. R. Slatyer, and N. Weiner, “A Theory of Dark Matter,” *Phys.Rev.*, vol. D 79, p. 015014, 2009.
- [184] E. Dudas, Y. Mambrini, S. Pokorski, and A. Romagnoni, “Extra U(1) as natural source of a monochromatic gamma ray line,” *J. High Energy Phys.*, vol. 1210, p. 123, 2012.

- [185] D. Curtin, R. Essig, S. Gori, and J. Shelton, “Illuminating Dark Photons with High-Energy Colliders,” *JHEP*, vol. 02, p. 157, 2015.
- [186] H. Davoudiasl, H.-S. Lee, I. Lewis, and W. J. Marciano, “Higgs Decays as a Window into the Dark Sector,” *Phys.Rev.*, vol. D 88, no. 1, p. 015022, 2013.
- [187] H. Davoudiasl, H.-S. Lee, and W. J. Marciano, “‘Dark’ Z implications for Parity Violation, Rare Meson Decays, and Higgs Physics,” *Phys. Rev.*, vol. D85, p. 115019, 2012.
- [188] J. D. Wells, “How to Find a Hidden World at the Large Hadron Collider,” 2008.
- [189] S. Gopalakrishna, S. Jung, and J. D. Wells, “Higgs boson decays to four fermions through an abelian hidden sector,” *Phys. Rev.*, vol. D78, p. 055002, 2008.
- [190] J. Alexander et al., “Dark Sectors 2016 Workshop: Community Report,” 2016.
- [191] D. Clowe et al., “A direct empirical proof of the existence of dark matter,” *Astrophys.J.*, vol. 648, pp. L109–L113, 2006.
- [192] O. Adriani et al. (PAMELA Collaboration), “An anomalous positron abundance in cosmic rays with energies 1.5–100 GeV,” *Nature*, vol. 458, pp. 607–609, 2009.
- [193] J. Chang et al. (ATIC Collaboration), “An excess of cosmic ray electrons at energies of 300–800 gev,” *Nature News*, vol. 456, p. 362, 2008.
- [194] M. Aguilar et al. (AMS Collaboration), “First Result from the Alpha Magnetic Spectrometer on the International Space Station: Precision Measurement of the Positron Fraction in Primary Cosmic Rays of 0.5–350 GeV,” *Phys. Rev. Lett.*, vol. 110, p. 141102, 2013.
- [195] P. Galison and A. Manohar, “TWO Z’s OR NOT TWO Z’s?,” *Phys.Lett.*, vol. B136, p. 279, 1984.
- [196] B. Holdom, “Two U(1)’s and Epsilon Charge Shifts,” *Phys.Lett.*, vol. B166, p. 196, 1986.
- [197] K. R. Dienes, C. F. Kolda, and J. March-Russell, “Kinetic mixing and the supersymmetric gauge hierarchy,” *Nucl.Phys.*, vol. B492, pp. 104–118, 1997.
- [198] A. Hook, E. Izaguirre, and J. G. Wacker, “Model Independent Bounds on Kinetic Mixing,” *Adv. High Energy Phys.*, vol. 2011, p. 859762, 2011.
- [199] M. Pospelov, “Secluded U(1) below the weak scale,” *Phys. Rev.*, vol. D80, p. 095002, 2009.
- [200] I. Hoenig, G. Samach, and D. Tucker-Smith, “Searching for dilepton resonances below the Z mass at the LHC,” *Phys. Rev.*, vol. D90, no. 7, p. 075016, 2014.
- [201] E. M. Riordan et al., “A Search for Short Lived Axions in an Electron Beam Dump Experiment,” *Phys. Rev. Lett.*, vol. 59, p. 755, 1987.
- [202] J. D. Bjorken et al., “Search for Neutral Metastable Penetrating Particles Produced in the SLAC Beam Dump,” *Phys. Rev.*, vol. D38, p. 3375, 1988.
- [203] A. Bross et al., “A Search for Shortlived Particles Produced in an Electron Beam Dump,” *Phys. Rev. Lett.*, vol. 67, pp. 2942–2945, 1991.
- [204] CMS Collaboration, “A search for pair production of new light bosons decaying into muons,” *Phys. Lett.*, vol. B752, pp. 146–168, 2016.

- [205] J. Pumplin et al., “New generation of parton distributions with uncertainties from global QCD analysis,” *JHEP*, vol. 07, p. 012, 2002.
- [206] J. Alwall et al., “MadGraph 5 : Going Beyond,” *JHEP*, vol. 1106, p. 128, 2011.
- [207] ATLAS Collaboration, “Measurement of the Z/γ^* boson transverse momentum distribution in pp collisions at $\sqrt{s} = 7$ TeV with the ATLAS detector,” *JHEP*, vol. 09, p. 145, 2014.
- [208] J. M. Butterworth, J. R. Forshaw, and M. H. Seymour, “Multiparton interactions in photoproduction at HERA,” *Z. Phys.*, vol. C72, pp. 637–646, 1996.
- [209] ATLAS Collaboration, “Search for Higgs and Z Boson Decays to $J/\psi \gamma$ and $\Upsilon(nS) \gamma$ with the ATLAS Detector,” *Phys. Rev. Lett.*, vol. 114, no. 12, p. 121801, 2015.
- [210] ATLAS Collaboration, “Measurement of the differential cross-sections of prompt and non-prompt production of J/ψ and $\psi(2S)$ in pp collisions at $\sqrt{s} = 7$ and 8 TeV with the ATLAS detector,” *Eur. Phys. J.*, vol. C76, no. 5, p. 283, 2016.
- [211] ATLAS Collaboration, “Search for Higgs decays to beyond the standard model light gauge bosons in four-lepton events with the ATLAS detector $\sqrt{s} = 13$ TeV,” Tech. Rep. ATL-COM-PHYS-2016-1500, CERN, Geneva, Oct 2016. (Unpublished).
- [212] LHC Higgs Cross Section Working Group, S. Dittmaier, C. Mariotti, G. Passarino, and R. Tanaka (Eds.), “Handbook of LHC Higgs Cross Sections: 1. Inclusive Observables,” *CERN-2011-002*, CERN, Geneva, 2011.
- [213] LHC Higgs cross section working group, S. Dittmaier, C. Mariotti, G. Passarino, and R. Tanaka (Eds.), “Handbook of LHC Higgs Cross Sections: 2. Differential distributions,” 2012. CERN-2012-002.
- [214] ATLAS Collaboration, “Luminosity determination in pp collisions at $\sqrt{s} = 8$ TeV using the ATLAS detector at the LHC,” *Eur. Phys. J.*, vol. C76, no. 12, p. 653, 2016.
- [215] E. Gross and O. Vitells, “Trial factors or the look elsewhere effect in high energy physics,” *Eur. Phys. J.*, vol. C70, pp. 525–530, 2010.
- [216] A. S. Chisholm, S. Kuttimalai, K. Nikolopoulos, and M. Spannowsky, “Measuring rare and exclusive Higgs boson decays into light resonances,” *Eur. Phys. J.*, vol. C76, no. 9, p. 501, 2016.
- [217] The ATLAS collaboration, “Search for exclusive higgs and z boson decays to $\phi\gamma$ and $\rho\gamma$ with the atlas detector,” *Journal of High Energy Physics*, vol. 2018, p. 127, Jul 2018.
- [218] The ATLAS collaboration, “Searches for exclusive higgs and z boson decays into $j/\psi\gamma$, $\phi(2s)\gamma$, and $v(ns)\gamma$ at $\sqrt{s}=13$ tev with the atlas detector,” *Physics Letters B*, vol. 786, pp. 134 – 155, 2018.
- [219] S. Alte, M. König, and M. Neubert, “Exclusive Weak Radiative Higgs Decays in the Standard Model and Beyond,” *JHEP*, vol. 12, p. 037, 2016.
- [220] G. Isidori, A. V. Manohar, and M. Trott, “Probing the nature of the Higgs-like Boson via $h \rightarrow V\mathcal{F}$ decays,” *Phys. Lett.*, vol. B728, pp. 131–135, 2014.
- [221] ATLAS Collaboration, “Measurement of the Z/γ^* boson transverse momentum distribution in pp collisions at $\sqrt{s} = 7$ TeV with the ATLAS detector,” *JHEP*, vol. 09, p. 145, 2014.

- [222] A. Ryd et al., “EvtGen: A Monte Carlo Generator for B-Physics,” 2005.
- [223] H.-L. Lai et al., “New parton distributions for collider physics,” *Phys. Rev.*, vol. D82, p. 074024, 2010.
- [224] ATLAS Collaboration, “Monte Carlo Generators for the Production of a W or Z/γ^* Boson in Association with Jets at ATLAS in Run 2,” Tech. Rep. ATL-PHYS-PUB-2016-003, CERN, Geneva, Jan 2016.
- [225] T. Gleisberg et al., “Event generation with SHERPA 1.1,” *JHEP*, vol. 02, p. 007, 2009.
- [226] R. D. Ball et al., “Parton distributions with LHC data,” *Nucl. Phys.*, vol. B 867, pp. 244–289, 2013.
- [227] J Butterworth et al., “Single Boson and Diboson Production Cross Sections in pp Collisions at $\sqrt{s}=7$ TeV,” Tech. Rep. ATL-COM-PHYS-2010-695, CERN, Geneva, Aug 2010.
- [228] ATLAS Collaboration, “Multi-Boson Simulation for 13 TeV ATLAS Analyses,” Tech. Rep. ATL-PHYS-PUB-2017-005, CERN, Geneva, May 2017.
- [229] ATLAS Collaboration, “The atlas simulation infrastructure,” *The European Physical Journal C*, vol. 70, pp. 823–874, Dec 2010.
- [230] S. Agostinelli et al., “Geant4—a simulation toolkit,” *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 506, no. 3, pp. 250 – 303, 2003.
- [231] M. Cacciari and G. P. Salam, “Pileup subtraction using jet areas,” *Phys. Lett.*, vol. B 659, pp. 119–126, 2008.
- [232] ATLAS Collaboration, “Jet global sequential corrections with the ATLAS detector in proton-proton collisions at $\sqrt{s} = 8$ TeV,” Tech. Rep. ATLAS-CONF-2015-002, CERN, Geneva, Mar 2015.
- [233] ATLAS Collaboration, “Selection of jets produced in 13TeV proton-proton collisions with the ATLAS detector,” Tech. Rep. ATLAS-CONF-2015-029, CERN, Geneva, Jul 2015.
- [234] T. A. collaboration, “Jet mass reconstruction with the ATLAS Detector in early Run 2 data,” 2016.
- [235] J. Thaler and K. Van Tilburg, “Identifying Boosted Objects with N-subjettiness,” *JHEP*, vol. 03, p. 015, 2011.
- [236] I. Moutl, L. Necib, and J. Thaler, “New Angles on Energy Correlation Functions,” *JHEP*, vol. 12, p. 153, 2016.
- [237] L. G. Almeida et al., “Substructure of high- p_T Jets at the LHC,” *Phys. Rev.*, vol. D79, p. 074017, 2009.
- [238] J. de Favereau et al., “DELPHES 3, A modular framework for fast simulation of a generic collider experiment,” *JHEP*, vol. 02, p. 057, 2014.
- [239] CERN Higgs Cross Section Working Group, “<https://twiki.cern.ch/twiki/bin/view/lhcphysics/cernyellowreportpageat1314tev2014>,” 2016.
- [240] Particle Data Group, “<http://pdg.lbl.gov/2018/tables/rpp2018-sum-gauge-higgs-bosons.pdf>,” 2018.

- [241] ATLAS Collaboration, “Luminosity determination in pp collisions at $\sqrt{s} = 8$ TeV using the ATLAS detector at the LHC,” *Eur. Phys. J. C*, vol. 76, p. 653, 2016.
- [242] ATLAS Collaboration, “Observation of $H \rightarrow b\bar{b}$ decays and VH production with the ATLAS detector,” *Phys. Lett.*, vol. B786, pp. 59–86, 2018.