

Search for exotic long-lived particles decaying into hadronic states in the calorimeter of the ATLAS detector at the Large Hadron Collider

Alice Polyxeni Morris

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
of
University College London.

Department of Physics and Astronomy
University College London

5th August 2020



I, Alice Polyxeni Morris, confirm that the work presented in this thesis is my own. Where information has been derived from other sources, I confirm that this has been indicated in the thesis.

Abstract

Exotic long-lived particles feature in many extensions to the Standard Model and, as a result of their unique decay signatures, may have been overlooked by previous searches for new physics, which have usually focused on promptly-decaying particles. This thesis presents a search for pairs of long-lived neutral particles decaying to hadronic final states in the calorimeters of the ATLAS detector, using data collected during Run 2 of the Large Hadron Collider. Long-lived neutral particles with masses between 5 and 400 GeV were considered, produced from decays of heavy bosons with masses between 125 and 1000 GeV, and decaying into Standard Model fermions. Machine learning techniques were employed to identify the displaced decays, and a data-driven estimate of the remaining background was performed. Since no events were observed in the search region, limits were set on the production cross section times branching ratio, and the result was extrapolated as a function of the decay length of the long-lived particles. The decay lengths probed range between a few centimetres and a few tens of metres. The combination of the results of this search with those from a search for displaced decays in the muon spectrometer is also presented, and provides a summary of the results of searches for the simplified hidden-sector model. Constraints on three new physics models not originally studied in the search are also shown, resulting from the reinterpretation of this search with the RECAST framework. The original data analysis workflow was completely captured using virtualisation techniques, allowing for an accurate and efficient reinterpretation of the search result in terms of the new signal models. For all three signal models considered, the results from the reinterpretation were complementary to the existing limits from previous dedicated searches, especially at shorter proper decay lengths.

Impact statement

Studies of particle physics have the potential to lead to ground-breaking discoveries, such as physics beyond the Standard Model and the development of new theories to better describe our Universe. The doctoral work presented in this thesis focused on a search for long-lived particles using the ATLAS detector at CERN's Large Hadron Collider (LHC). These searches are usually sensitive to a wide range of models of physics beyond the Standard Model and, as such, are useful in reducing the number of possible theories allowed by current observations. The results of the search were also combined with those of a similar search to provide a summary of the overall ATLAS sensitivity to the benchmark signal model they both considered.

In addition, the RECAST framework was used to preserve the analysis so that it could be easily and rapidly reinterpreted in terms of new signal models. Analysis preservation is of paramount importance in maximising the impact of the data recorded by the LHC experiments, so that the enormous person power (and money) invested in them will be useful for years to come. This search was the first long-lived particle search to be preserved in full-fidelity in the RECAST framework, and this will influence the future of long-lived particle searches at the LHC, setting a precedent for analyses to easily test models they may not have previously considered and may be sensitive to.

Particle physics collaborations such as the ATLAS experiment are international, uniting people from all over the world to push the frontiers of science and technology, for the benefit of all. New techniques and technologies with wide-ranging applications have been developed as a result of particle physics research, and particle physicists in turn benefit from the knowledge and tools used in industry. The RECAST framework illustrates the latter case well, in that it employs industry-standard tools such as Docker to the benefit of particle physics research.

In the wider community, CERN's LHC research programme, to which the research presented in this thesis belongs, has stimulated interest in physics on a remarkable scale—most notably with the discovery of the Higgs boson. In addition to enabling research at the forefront of human knowledge, CERN aims to inspire and nurture scientific awareness in the public, and it serves as a politically neutral voice for science, advocating investment in fundamental research and evidence-based policy.

Author contributions

This thesis describes the work carried out by the author as a member of the ATLAS collaboration. Due to the nature of particle physics research at the Large Hadron Collider and within the ATLAS collaboration, this work was carried out collaboratively by the author and several colleagues. The author's individual contributions are therefore outlined here.

As an active member of the ATLAS Exotics group, the author contributed to the search for neutral long-lived particles decaying to hadronic final states in the ATLAS calorimeter, described in **Chapter 4**. Within this analysis, the author's work included: studies of the cosmic-ray background; development of the analysis selection software; application of the analysis selections to the simulated signal samples; development of the lifetime extrapolation procedure; statistical interpretation of the results; and editing the internal documentation as well as the final paper. Further to the main search, the author was responsible for the statistical combination of the results with those from the search for displaced jets in the muon spectrometer, presented in **Chapter 5**. Additionally, the author led the preservation of the displaced jets search and its subsequent reinterpretation with respect to three further signal models, described in **Chapter 6**. This work included: containerisation of the analysis software; development and validation of the code defining the analysis workflow and job templates; and the application of the workflow to new signal models.

Acknowledgements

The completion of this PhD would not have been possible without the help and support of a great many people. First, I would like to thank my supervisor, Gavin Hesketh, for steering me towards long-lived particle searches in the first place, providing support and advice throughout my PhD, and in particular for reviewing my thesis. I would also like to thank Gabriel Facini who looked out for me while I was at CERN, advocated for my work in ATLAS meetings, and led me towards the RECAST project. It was a pleasure and an inspiration to be able to work with both of you.

I am profoundly grateful to Louie Corpe, who was such a font of knowledge and advice that I occasionally referred to him as ‘my’ post-doc. Thank you for the encouragement in every stage of my PhD, the invaluable feedback on all my presentations and written work, the hours spent debugging when things went badly, and the celebratory pints when things went well.

To the entire displaced jets analysis team, thank you for the many fruitful discussions during our weekly meetings—I could not have wished for a better group of colleagues. Without your hard work, expertise, and positivity, the work in this thesis would not have been possible. Special thanks also goes to Lukas Heinrich for the technical support while I was using the RECAST framework.

Over the last few years, the UCL High Energy Physics group was a fantastic place to work, and I extend my thanks to every member of the group for contributing to the welcoming and outgoing environment. In particular, I am grateful to the current and former students in the group for the windowsill conversations in D109, coffee and cake in D14, pub quizzes, and countless fun evenings. It has been a joy to get to know you all, and I look forward to continued shared experiences for years to come. I would also like to thank all the friends

I met while at CERN. Whether skiing in the most breathtaking scenery, rafting down the Rhône, or simply whiling away the evenings in R1, the time I shared with them made the 18 months I spent in Geneva some of the best in my life, and I doubt they will ever know the impact they had on me.

Of all the friends I found during my PhD, I especially want to thank Susy (who was there every step of the way with endless support and friendship, and frequently a sofa-bed), and Laurent (who was always up for a coffee break and a healthy dose of realism). I could not have had better companions to undertake this PhD journey with, and other journeys too, and I am so glad to have found such wonderful friends.

It is said that true friendships resist time, distance, and silence. I am incredibly fortunate to have received continual comfort from many friends I made before my PhD. I would like to thank Alexandra, Megan, Ellie, and Rhia for numerous hours filled with laughter, welcome distractions, and more often than not, wine. I am thankful that our friendships have thrived despite the distance between us—and to Alexandra, my lockdown-partner-in-crime, perhaps I should also say despite the *lack* of distance between us in the last few months! I would also like to thank Claire for her readiness to provide a sympathetic ear and hosting me in Amsterdam when I needed a relaxing getaway.

To my parents, Tony and Fotini, whose unwavering support and encouragement has made me the person I am today—I hope I've made you proud. To my not-so-little brother, Dimitris, who will probably never read this—thank you for making me laugh. Στη Παγιά, που μου δίδαξε το νόημα της σκληρής δουλειάς, και στον Παππού, που μου έδειξε το θαύμα των μικρών πραγμάτων.

Last, but by no means least, thank you to Simon, whose love and support kept me going even when we were separated by hundreds of miles.

Contents

1	Introduction	25
1.1	A particle physics primer	26
2	Theory	33
2.1	The Standard Model	34
2.1.1	Particles and forces	34
2.1.2	Quantum field theory	35
2.1.3	The electroweak interaction	36
2.1.4	QCD and hadronisation	36
2.1.5	The Higgs boson	37
2.2	Proton-proton collisions	38
2.3	Limitations of the Standard Model	40
2.3.1	Dark matter and dark energy	40
2.3.2	The hierarchy problem	40
2.3.3	Baryon asymmetry in the Universe	41
2.4	Physics beyond the Standard Model	41
2.4.1	Long-lived particles	42
2.4.2	Hidden sectors	43
2.4.3	Supersymmetry	45
2.4.4	Baryogenesis	48
2.5	Summary	49
3	The ATLAS experiment	51
3.1	The LHC	51
3.1.1	Overview	52
3.1.2	Bunch trains and filling schemes	53

3.1.3	Luminosity and pile-up	54
3.2	The ATLAS detector	56
3.2.1	Inner detector	57
3.2.2	Calorimeters	59
3.2.3	Muon spectrometer	62
3.2.4	Magnets	63
3.3	Triggering in ATLAS	63
3.3.1	CalRatio triggers	65
3.4	Event reconstruction	68
3.4.1	Track and vertex reconstruction	68
3.4.2	Electrons, photons and muons	70
3.4.3	Jets	72
3.4.4	Taus	76
3.4.5	Composite variables	77
3.5	Simulation	78
3.6	Long-lived particle signatures	80
4	Search for events with displaced jets in the ATLAS calorimeter	83
4.1	Analysis overview	83
4.1.1	Benchmark signal model	84
4.1.2	Previous searches for neutral LLPs	84
4.1.3	Background contributions	85
4.1.4	Search strategy	88
4.2	Data and simulated samples	90
4.2.1	Data samples	90
4.2.2	Signal and background simulation samples	91
4.2.3	Trigger efficiencies	92
4.3	Event selection	94
4.3.1	Definition of CalRatio jet candidates	94
4.3.2	Definition of CalRatio signal events	102
4.4	Background estimation	111
4.4.1	BIB background removal check	111

	15
4.4.2 Cosmic background removal check	113
4.4.3 Multijet background estimation: ABCD method	114
4.5 Systematic uncertainties	123
4.6 Statistical interpretation	126
4.6.1 Efficiency extrapolation	127
4.6.2 Limit setting procedure	128
4.7 Discussion of results	135
5 Combination with search for displaced vertices in the muon spectrometer	139
5.1 Search for displaced jets in the muon spectrometer	139
5.2 Orthogonality of searches	142
5.3 Statistical combination procedure	143
5.4 Results	143
6 Reinterpretation of search for displaced jets using RECAST	151
6.1 Reinterpretation	151
6.2 The RECAST framework	154
6.2.1 Analysis preservation in RECAST	157
6.3 Preservation of CalRatio displaced jets analysis	158
6.3.1 Validation with signal of original analysis	161
6.4 Constraints on new models	164
6.4.1 Stealth SUSY model results	164
6.4.2 Higgs-portal baryogenesis model results	168
6.4.3 Dark photon model results	174
6.5 Further work	177
7 Summary	181
Bibliography	183

List of figures

Figure 1.1:	Interactions of Standard Model particles with the ATLAS detector.	31
Figure 1.2:	Interactions of long-lived neutral particles with the ATLAS detector.	32
Figure 2.1:	The constituents of the Standard Model of particle physics: the 12 fundamental fermions and 5 fundamental bosons.	34
Figure 2.2:	The MMHT 2014 NNLO proton PDFs as a function of fractional energy of the proton that a parton claims.	39
Figure 2.3:	Particle decay length, $c\tau$, as a function of particle mass for a variety of Standard Model particles.	42
Figure 2.4:	Feynman diagram of the hidden sector scalar (HSS) model.	44
Figure 2.5:	Feynman diagrams of the FRVZ model with dark photon final states.	45
Figure 2.6:	Feynman diagram of the Stealth SUSY model.	47
Figure 2.7:	Feynman diagram of the Higgs-portal baryogenesis model.	48
Figure 3.1:	Overview of the CERN accelerator complex and the LHC.	52
Figure 3.2:	Cumulative pp luminosity by day delivered by the LHC and recor- ded by ATLAS for the 2016 data-taking period at $\sqrt{s} = 13$ TeV.	55
Figure 3.3:	The mean number of interactions per bunch crossing for the 2016 ATLAS pp dataset.	56
Figure 3.4:	Overview of the ATLAS detector with subdetectors labelled.	57
Figure 3.5:	Diagram of the ATLAS inner detector showing all its components, with the distances to the interaction point labelled.	58
Figure 3.6:	Diagram of the ATLAS calorimeters.	60
Figure 3.7:	The depth of each layer of the ATLAS calorimeters as measured by the nuclear interaction length.	61
Figure 3.8:	The ATLAS muon spectrometer.	62

Figure 3.9: The average number of reconstructed vertices as a function of the number of interactions per bunch crossing for the 2016 ATLAS pp dataset.	69
Figure 3.10: Representation of a $t\bar{t}h$ event as produced by an event generator.	79
Figure 4.1: A typical BIB event, showing a line of calorimeter energy deposits parallel to the beam.	86
Figure 4.2: The ATLAS detector in the experimental cavern.	87
Figure 4.3: A schematic overview of the analysis strategy.	89
Figure 4.4: Trigger efficiency for simulated signal events as a function of the LLP p_T and the p_T distribution of LLPs, for a selection of signal samples.	93
Figure 4.5: Trigger efficiency for simulated signal events as a function of the LLP decay position in the x - y plane for LLPs decaying in the barrel and in the z direction for LLPs decaying in the HCal endcaps, for three signal samples.	94
Figure 4.6: Probability density of MLP-predicted radial (L_{xy}) and longitudinal (L_z) LLP decay positions as a function of the truth LLP decay positions, for reconstructed jets matched to the LLP.	97
Figure 4.7: The distributions of the per-jet BDT weights for a multijet sample, a BIB sample and five signal samples.	100
Figure 4.8: Distribution of the jet leading cluster's z -position versus its timing for BIB data after the high- E_T preselection, for jets with low BIB-weight and jets with high BIB-weight.	102
Figure 4.9: Distribution of the low- E_T per-event BDT and high- E_T per-event BDT on main data, BIB data and five signal samples after preselection.	104
Figure 4.10: Distribution of the jet leading cluster's z -position versus its timing for BIB data and main data in three different ranges of high- E_T per-event BDT.	106
Figure 4.11: Distribution of the jet leading cluster's z -position versus its timing for BIB data and main data in three different ranges of low- E_T per-event BDT.	107
Figure 4.12: Diagram of the four regions in the ABCD plane.	115

Figure 4.13: The $\sum \Delta R_{\min}(\text{jet, tracks})$ and per-event BDT value distributions for data, BIB jets and signal after applying the event cleaning selection criteria, for the high- E_T analysis and low- E_T analysis.	116
Figure 4.14: The distributions of $\sum \Delta R_{\min}(\text{jet, tracks})$ vs. high- E_T per-event BDT and low- E_T per-event BDT for a signal sample, main data, and BIB events after event cleaning.	117
Figure 4.15: The $\sum \Delta R_{\min}(\text{jet, tracks})$ vs. per-event BDT distribution for several signal samples in the low- E_T selection and high- E_T selection.	119
Figure 4.16: The $\sum \Delta R_{\min}(\text{jet, tracks})$ vs. per-event BDT distribution for a signal sample, main data, and BIB data in $\text{VR}_{\text{high-}E_T}$ and $\text{VR}_{\text{low-}E_T}$	120
Figure 4.17: The extrapolated signal efficiencies as a function of proper decay length of the S for all simulated samples in the low- E_T and high- E_T selections.	128
Figure 4.18: The observed limits, expected limits and ± 1 and 2 sigma bands for the Φ masses of 125 GeV.	130
Figure 4.19: The observed limits, expected limits and ± 1 and 2 sigma bands for the Φ masses of 200 GeV.	131
Figure 4.20: The observed limits, expected limits and ± 1 and 2 sigma bands for the Φ masses of 400, 600 and 1000 GeV.	132
Figure 4.21: Summary of the observed limits for the Φ masses of 125, 200, 400, 600 and 1000 GeV.	133
Figure 4.22: Comparison of the observed and expected limits and $\pm 1\sigma$ bands obtained with the asymptotic method and the frequentist method for calculating the limit.	135
Figure 5.1: Combined limits for models with $m_\Phi = 125$ GeV from the CalRatio analysis and the MS analysis.	146
Figure 5.2: Combined limits for models with $m_\Phi = 200$ GeV from the CalRatio analysis and the MS analysis.	147
Figure 5.3: Combined limits for models with $m_\Phi = 400$ GeV from the CalRatio analysis and the MS analysis.	147

Figure 5.4: Combined limits for models with $m_\Phi = 600$ GeV from the CalRatio analysis and the MS analysis.	148
Figure 5.5: Combined limits for models with $m_\Phi = 1000$ GeV from the CalRatio analysis and the MS analysis.	148
Figure 5.6: Summary of combined limits from the CalRatio and MS analyses for models with $m_\Phi = 125$ GeV, shown alongside limits from another ATLAS search for exotic Higgs boson decays into a pair of scalar particles which decay into b-quarks promptly or with a short lifetime.	149
Figure 6.1: Schematic showing reconstruction and unfolding processes to convert between particle-, detector- and reconstruction-level information. . .	153
Figure 6.2: Diagram of a full analysis workflow.	156
Figure 6.3: Diagram of a reinterpretation workflow.	156
Figure 6.4: A schematic of the workflow used in the analysis preservation. . .	161
Figure 6.5: A more detailed view of the workflow used in the analysis preservation.	162
Figure 6.6: Comparison of the original published limits with those obtained from running the preserved analysis workflow with the low- E_T selection for $m_\Phi = 200$ GeV, $m_S = 8$ GeV and the high- E_T selection for $m_\Phi = 600$ GeV, $m_S = 150$ GeV.	163
Figure 6.7: A comparison of the constraints for the Stealth SUSY model with $m_{\tilde{g}} = 250, 500$ and 800 GeV from the high- E_T dataset of the newly-preserved CalRatio jet analysis and the results from the MS search. . . .	166
Figure 6.8: A comparison of the constraints for the Stealth SUSY model with $m_{\tilde{g}} = 1200, 1500$, and 2000 GeV from the high- E_T dataset of the newly-preserved CalRatio jet analysis and the results from the MS search. . . .	167
Figure 6.9: A comparison of the constraints on the $\chi \rightarrow \tau\tau\nu$ decay mode of the Higgs-portal baryogenesis model from the high- E_T dataset of the preserved CalRatio displaced jet analysis and the results from the MS search. . . .	169
Figure 6.10: A comparison of the constraints on the $\chi \rightarrow \nu b\bar{b}$ decay mode of the Higgs-portal baryogenesis model from the high- E_T dataset of the preserved CalRatio displaced jet analysis and the results from the MS search. . . .	170

- Figure 6.11:** A comparison of the constraints on the $\chi \rightarrow cbs$ decay mode of the Higgs-portal baryogenesis model from the high- E_T dataset of the preserved CalRatio displaced jet analysis and the results from the MS search. 171
- Figure 6.12:** A comparison of the constraints on the $\chi \rightarrow \ell cb$ decay mode of the Higgs-portal baryogenesis model from the high- E_T dataset of the preserved CalRatio displaced jet analysis and the results from the MS search. 172
- Figure 6.13:** A comparison of the extrapolated efficiencies from the preserved CalRatio displaced jet analysis and the hLJ-hLJ selection in the dedicated displaced lepton-jets search (where applicable) for the dark photon model. 175
- Figure 6.14:** The constraints on the dark photon model from the preserved CalRatio displaced jet analysis (compared to those from the hLJ-hLJ selection of the dedicated displaced lepton-jets search where applicable). . . 176

List of tables

Table 3.1: The ATLAS calorimetry energy and time resolutions.	60
Table 4.1: Summary of the data samples used for the analysis.	91
Table 4.2: Summary of the mass and decay length parameters simulated for the $\Phi \rightarrow SS$ hidden sector model.	92
Table 4.3: A summary of the requirements for event to enter the ABCD planes of the high- E_T and low- E_T selections.	110
Table 4.4: Sequential impact of each requirement on the number of events passing the selection for the high- E_T and low- E_T analyses.	112
Table 4.5: Application of the ABCD method to the high- E_T and low- E_T selec- tions.	118
Table 4.6: The definitions of the ABCD planes of the high- E_T and low- E_T se- lections and their respective validation regions.	122
Table 4.7: Application of the ABCD method to the high- E_T and low- E_T VRs.	122
Table 4.8: Summary of the systematic uncertainties.	123
Table 4.9: Application of the ABCD method to the high- E_T and low- E_T con- trol regions (CRs) used to determine the systematic uncertainty associated with the ABCD method.	124
Table 4.10: The simple ABCD estimate, the blinded estimate, and the unblinded estimate of the expected background in region A.	134
Table 4.11: Proper decay lengths excluded by the ATLAS searches that con- sidered the same HSS model used in the Run 2 CalRatio analysis, for two signal models common to them all.	137
Table 5.1: A summary of the selections applied in the MS vertex search strategies.	140
Table 5.2: Event counts in each of the four regions of the ABCD plane for the MS 1Vx search.	142

Table 6.1: Summary of the masses and proper decay lengths ($c\tau$) simulated for the Stealth SUSY model.	165
Table 6.2: The dominant systematic uncertainties on the Stealth SUSY model shown for all generated mass points.	165
Table 6.3: Ranges of singlino proper decay lengths excluded at 95% CL for the Stealth SUSY model, following the reinterpretation of the CalRatio displaced jets search.	168
Table 6.4: Summary of the masses and proper decay lengths ($c\tau$) simulated for the Higgs-portal baryogenesis model.	169
Table 6.5: The dominant systematic uncertainties on the Higgs-portal baryogenesis model shown for all generated mass points, averaged over all four decay channels.	169
Table 6.6: The χ proper decay lengths excluded at 95% CL for the Higgs-portal baryogenesis model, following the reinterpretation of the CalRatio displaced jets search.	173
Table 6.7: Summary of the masses and proper decay lengths ($c\tau$) simulated for the dark photon model.	174
Table 6.8: The dominant systematic uncertainties on the dark photon model shown for all generated mass points.	174
Table 6.9: The γ_d proper decay lengths excluded at 95% CL for the dark photon model with $m_H = 125$ GeV, following the reinterpretation of the CalRatio displaced jets search.	177
Table 6.10: The γ_d proper decay lengths excluded at 95% CL for the dark photon model with $m_H = 800$ GeV, following the reinterpretation of the CalRatio displaced jets search.	177

Chapter 1

Introduction

Over the last century, the collaborative scientific work of thousands of theoretical and experimental physicists and dozens of experiments has culminated in one of the triumphs of modern physics—the Standard Model of particle physics (SM). Although its success in describing most of the phenomena we observe has been remarkable, there are still many questions it remains unable to answer. Since the discovery of the Higgs boson in 2012—the last remaining puzzle piece of the SM—the field of high energy physics has been devoted to examining the SM and its possible extensions. There is not always a clear indication of which extensions to the SM are more likely to be correct than others, so search groups are sent out in all possible directions in the hopes of happening on a solution. One such direction is explored in this thesis; namely, that new physics may be uncovered by searching for exotic particles that decay some time after their production in high-energy proton collisions, resulting in distinctive displaced signatures.

The Large Hadron Collider (LHC) is one of a long line of collider experiments designed to probe the small distance scales of the Universe using accelerated protons in collisions at centre-of-mass energies ranging from 900 GeV to 14 TeV. This thesis describes a search for long-lived neutral particles decaying into hadronic states in the calorimeter of the ATLAS detector at the LHC, using data collected in 2016 at a centre-of-mass energy of 13 TeV.

An introduction to the theoretical framework of the SM is given in Chapter 2, alongside a discussion of its limitations and possible theoretical solutions, in particular for the case of theories involving long-lived particles. In Chapter 3, the LHC and the ATLAS detector are outlined, the reconstruction of physics objects from the recorded data is discussed, and the signatures expected from long-lived particle decays in the detector are summarised.

The remaining chapters describe the analyses towards which the author has made significant contributions. Firstly, Chapter 4 details the so-called CalRatio displaced jets search, in which pairs of neutral long-lived particles are sought by exploiting the distinctive signature of jets with little energy deposited in the electromagnetic calorimeter and no associated tracks, and which was used to constrain a benchmark hidden sector model. The combination of the results of this search with an analogous search for displaced jets in the ATLAS muon spectrometer is presented in Chapter 5, where the statistical likelihoods are combined to improve and extend the limits set on the hidden sector model. Subsequently, the reinterpretation of the CalRatio displaced jets search in the context of three new signal models is presented in Chapter 6, which is the first reinterpretation of an ATLAS long-lived particle search using the RECAST framework. Finally, Chapter 7 summarises the main results and conclusions from the work outlined in previous chapters.

Considering that non-expert readers may find the majority of this thesis difficult to follow despite the effort made to explain detail while remaining as concise as possible, the remainder of this chapter contains an accessible introduction to the work presented herein.

1.1 A particle physics primer

All the matter we interact with—vegetable, animal and mineral—can be broken into smaller pieces called elements. A pure sample of any element, for example carbon, can be cut into smaller and smaller pieces and still retain the properties of that element. However, this is only true on the scale on which we see and interact with objects (the macroscopic scale). Eventually, the pieces will be small enough that one more cut will reach the point where the resulting pieces are no longer carbon.

The smallest possible amount of matter which still retains its identity as a particular element is called an *atom*, from the Greek word *ἄτομος*, meaning *indivisible*. Even the largest atoms are about 200,000 times smaller than the width of a human hair. The composition of an atom was the subject of much study for many years, and there were many proposed models before scientists converged on something that accurately described the interactions that happen on an atomic scale. The model of an atom that most people think of consists of a solid, positively-charged nucleus made of positively-charged protons and neutral neutrons,

orbited by negatively-charged electrons. The number of protons in the nucleus defines the element, and the number of neutrons defines the isotope of that element. The number of electrons is usually equal to the number of protons, unless the atom is *ionised*—either positively or negatively charged.

This could have been the end of the story. However, nuclear decay experiments showed that atoms could change into others via the decay of a neutron. In the following decades, a plethora of new particles and interactions were predicted or discovered. This eventually led to the development of the Standard Model of particle physics, which describes the 17 elementary particles and their interactions.

Matter particles occur in two basic types called quarks[†] and leptons. Each type consists of six particles, which are related in pairs, or *generations*. The lightest and most stable particles make up the first generation, while the heavier and less-stable particles belong to the second and third generations. All stable matter in the universe is made from particles that belong to the first generation; any heavier particles quickly decay to more stable ones. The six quarks are: the *up quark* and *down quark* (first generation); the *charm quark* and *strange quark* (second generation); and the *top quark* and *bottom (or beauty) quark* (third generation). Quarks also come in three different *colours* and only combine in ways as to form colourless objects, usually in triplets and doublets. The triplets are called baryons from the Greek βαρύς (heavy); the doublets are mesons from the Greek μέσος (medium). Collectively, baryons and mesons are known as hadrons, from the Greek αδρός (stout). The six leptons, named from the Greek word λεπτός (light), are: the *electron* and the *electron neutrino* (first generation); the *muon* and the *muon neutrino* (second generation); and the *tau* and the *tau neutrino* (third generation). The electron, the muon and the tau all have an electric charge and a sizeable mass, whereas the neutrinos are electrically neutral and have very little mass. At currently-accessible energies, these quarks and leptons are believed to be truly indivisible.

The interactions of the particles of matter are governed by four forces: the strong force, the weak force, the electromagnetic force, and the gravitational force. They work over different ranges and have different strengths. Gravity is the weakest but it has an infinite range. The

[†]The word *quark* comes from a phrase uttered by an inebriated seagull in James Joyce's novel *Finnegans Wake*: 'three quarks for Muster Mark'. At the time their name was chosen, the pre-SM theory only included three quarks, but the name stuck even though the current SM has six quarks.

electromagnetic force also has infinite range but it is many times stronger than gravity. The weak and strong forces are effective only over a very short range and dominate only at the scale of subatomic particles. Despite its name, the weak force is much stronger than gravity but it is the weakest of the other three. The strong force, as the name suggests, is the strongest of all four fundamental interactions.

Three of these forces result from the exchange of force-carrier particles, which belong to a broader group called *bosons*. Particles of matter transfer discrete amounts of energy by exchanging bosons with each other. Each fundamental force has its own corresponding boson – the strong force is carried by the *gluon*, the electromagnetic force is carried by the *photon*, and the *W* and *Z bosons* are responsible for the weak force. The Standard Model includes the electromagnetic, strong and weak forces and their four carrier particles, and explains how these forces act on all of the matter particles. However, the force that is most familiar to us in everyday life, gravity, is not part of the Standard Model. Gravity is described separately by Einstein’s theory of general relativity, which does not include the effects of quantum physics so can only be used to describe the macroscopic world. Luckily, at the subatomic scales studied by particle physics, the effect of gravity is so weak as to be negligible—only when matter is in bulk, for example at the scale of the human body or planets, does the effect of gravity dominate. So the Standard Model still works well despite its reluctant exclusion of one of the fundamental forces.

Returning for a moment to the model of electrons orbiting an atomic nucleus, we can now go even further in describing the constituents of the nucleus (which are examples of hadrons). The positively-charged protons are composed of two up-quarks and one down-quark, and the neutral neutrons are composed of two down-quarks and one up-quark[†]. The quarks in each are bound together by gluons (carriers of the strong force).

There is one more particle included in the Standard Model—the Higgs boson—which has not yet been discussed here. The Higgs boson is associated with the Higgs field, which pervades the universe and is responsible for giving mass to elementary particles. To understand this, consider the following famous analogy. Imagine a busy cocktail party: upon entering the room, an average person could wander through the crowd with ease, but a

[†]In fact, these quarks are the so-called *valence quarks* that determine the quantum properties of the hadrons, apart from which any hadron may contain an indefinite number of virtual *sea quarks*, but a discussion of this is somewhat beyond the scope of this introduction.

celebrity would be mobbed, making their passage more difficult. In this analogy, the partygoers make up the Higgs field, and people walking through the crowd represent particles to which the field gives mass. A person who is significantly impeded by guests is like a particle given a large mass by the Higgs field. Crucially, all attendees are equally unimportant (so have zero mass) before they enter the room—it is the interaction with the crowd that causes them to gain mass, and the strength of the interaction determines how much mass they gain. The Higgs boson is an excitation of the Higgs field, like a wave at the surface of the sea. Returning to the cocktail party, this kind of excitation might move through the crowd if a rumour spread from one end of the room to the other. People nearest the rumour-originator would lean in to hear it. They would then pass it along to their neighbours, drawing in a new group of people, and then return to their original positions to discuss it. The compression of the crowd would move from one end of the room to the other, analogous to a Higgs boson in the Higgs field.

The current cutting-edge technology for studying high-energy particles and their interactions is at the European Organisation for Nuclear Research (CERN), where the LHC accelerates and collides protons in a circular tunnel 27 km in circumference, 100 m below the ground, between Geneva and the foothills of the Jura mountains. Seven particle detectors are placed along the collider, and are designed to detect the particles that emerge from high-energy proton collisions.

When two protons collide at the energies generated by the LHC, it is their constituent quarks and gluons that interact. The energies of the colliding particles create new particles: leptons, quarks, gluons, or even Higgs bosons. The ATLAS detector is essentially a digital camera; it was designed to identify and study outgoing particles from collisions, in order to understand what happens in the collisions themselves, by taking a snapshot each time a collision occurs.

Proton-proton collisions may be used to search for many things. Well-known Standard Model processes can be precisely measured, so their parameters and rates can be compared with theoretical predictions and previous experiments. On the other hand, theoretical particles whose existence could explain unanswered questions, like the nature of dark matter, could be searched for in the vast quantities of data collected.

For the most part, these studies involve a careful analysis of the Standard Model particles that can be measured with the ATLAS detector, such as electrons, photons, hadrons and muons. The detector has four distinct parts that can distinguish these different particles and measure their energies. The inner detector is used to reconstruct tracks from charged particles passing through it, and sits in a magnetic field that causes the charged particles to follow a curved path proportional to their momentum and charge, thus allowing them to be distinguished. The calorimeters measure the energy deposited by both the charged and neutral particles they stop. Of these, the electromagnetic calorimeter absorbs and measures the energy of charged particles, such as electrons and photons, and may measure some of the energy of hadrons. The hadronic calorimeter, on the other hand, is designed to absorb and measure the energy of hadrons, and the energy from all particles except muons should be absorbed by the end of it. Finally, the muon spectrometer provides tracking for any particles, mainly muons, that are not absorbed by the first three detector layers.

Figure 1.1 shows how each type of particle interacts with each part of the detector. Clockwise from the left, the first are muons: they leave tracks in the inner detector and some energy deposits in the calorimeters, but also have tracks in the muon spectrometer because they are not fully absorbed by the calorimeters. Next is the result of a quark or gluon: when produced in collisions, they shower into many hadrons, producing a narrow *jet* of particles that leave tracks in the inner detector and deposit their energy in the calorimeters. Continuing clockwise, the next particle is an electron: it is charged so it is bent in the inner detector's magnetic field, leaves a track, and then deposits its energy in the electromagnetic calorimeter. The following particle is a photon, which also deposits energy in the electromagnetic calorimeter but is not charged and thus leaves no track in the inner detector. The dashed line is a neutrino, which does not interact with any of the detector, and whose energy escapes detection. The final signature is due to low-energy charged particles, which dissipate before reaching the calorimeters, and are the most common result of proton-proton collisions forming a *background* over which high-energy objects can be searched for.

Most searches for new physics involve looking for some combination of the above particles at specific energies, since new particles produced in collisions are mainly expected to decay to Standard Model particles very quickly, such that the decay products originate from




Figure 1.1: Interactions of Standard Model particles with the ATLAS detector. Clockwise from the left, the particles are: muons (pink), a jet (orange), an electron (green), a photon (yellow), and a neutrino (black). Dashed lines indicate that the trajectory is not directly measured. Half of the transverse cross-section of the detector is shown in the background. Each coloured layer is a different type of detector: starting from the inside, the inner tracking detectors (green), the magnet (red), the electromagnetic calorimeter (purple), the hadronic calorimeter (teal), and the muon spectrometer (grey).

the centre of the detector. However, there is also the possibility of new physics looking somewhat different. Instead of immediately decaying to Standard Model particles, what if the new particle travelled some distance first, and did not interact with the detector?

The signature left by such a particle would be striking: a burst of activity far from the collision point, with nothing in between. Figure 1.2 shows some examples of the signatures these *long-lived neutral particles* could leave, if they decayed in various sub-detectors: the sudden appearance of many tracks in the muon spectrometer; a jet with no tracks or energy in the electromagnetic calorimeter; or many tracks in the inner detector all coming from a location that is not the collision point. The second of these signatures is the focus of the search outlined in this thesis.




Figure 1.2: Interactions of long-lived natural particles with the ATLAS detector. Clockwise from the left are long-lived particle decays in the muon spectrometer, hadronic calorimeter, and inner detector. Dashed lines indicate that the trajectory is not directly measured. Half of the transverse cross-section of the detector is shown in the background. Each coloured layer is a different type of detector: starting from the inside, the inner tracking detectors (green), the magnet (red), the electromagnetic calorimeter (purple), the hadronic calorimeter (teal), and the muon spectrometer (grey).

Chapter 2

Theory

The motivation for long-lived particle searches begins with the incompleteness of the current model of particle physics, the Standard Model (SM). In the SM, three of the four known forces and all observed particles are represented, along with their interactions. It has been remarkably successful at predicting and describing the subatomic world; observed phenomena such as baryon asymmetry and dark matter, however, are not accounted for by the SM, and gravity is also not described by it. Evidence is necessarily sought of physics beyond the SM to explain the phenomena behind our current observations.

With numerous models produced by theorists, it is sensible to devise some general searches for new physics based upon the shared features of many models. The studies presented in this thesis focus on one such common feature, long-lived particles, with the optimistic intent to discover new physics, and in the absence of an observation to reduce the possible phase space for theories containing long-lived particles.

This chapter provides the theoretical framework for understanding the models of new physics considered in the studies presented in later chapters of this thesis, and the phenomenology of the particle collisions whose data is used to search for them. The SM is summarised in §2.1, and a brief description of proton-proton collisions is provided in §2.2. The limitations of the SM are discussed in §2.3. This is followed, in §2.4, by a discussion of the models of physics beyond the SM considered in the analyses presented in Chapters 4–6.




Figure 2.1: The constituents of the Standard Model of particle physics: the 12 fundamental fermions and 5 fundamental bosons. Each particle is labelled with its mass, charge, and spin.

2.1 The Standard Model

The theory that encapsulates our current understanding of the fundamental elements of matter and the forces which govern their interactions is referred to as the Standard Model of particle physics. It is the result of many years of collaboration between experiment and theory, and in addition to accurately describing many of the physical phenomena we have observed, it also has remarkable predictive power. Several particles, such as the top and charm quarks, have been discovered because they were predicted by the SM, validating the model as a theory of fundamental physics—the most recent of these, the Higgs boson, was discovered by the ATLAS and CMS experiments at the Large Hadron Collider in 2012.^[1,2]

2.1.1 Particles and forces

The SM treats all the fundamental particles as point-like entities with an internal angular momentum quantum number called spin. The spin value of fundamental particles classifies them into one of two categories; *fermions* have half-integer spin and *bosons* have integer spins. Fermions are often referred to as the constituents of matter, although it is only the first generation *u*-quark, *d*-quark and electron that make up the majority of matter. There

are four known fundamental forces in the Universe: the electromagnetic force, the weak nuclear force, the strong nuclear force, and the gravitational force. The SM describes only the first three of these as no adequate theory of quantum gravity currently exists. They are represented by the exchange of spin-1 *vector bosons*: photons, gluons, and the W and Z bosons. These vector bosons are absorbed and emitted by other particles, communicating the forces between them. The SM includes another boson, the Higgs boson, which has zero spin (scalar)—its role will be discussed in §2.1.5.

The known particle content of the SM is shown in Figure 2.1. For each fermion an anti-matter counterpart also exists with the same mass but opposite quantum numbers. Fermions are separated into two categories depending on whether or not they carry colour charge and therefore interact via the strong force. Those that do are referred to as *quarks* and those that do not are referred to as *leptons*. Neutrinos are nearly-massless neutral leptons that interact only via the weak force. Although the SM takes neutrinos as massless, their masses can be non-trivially added. The observation of neutrino flavour oscillations^[3,4] indicates that there is a non-zero difference between the neutrino masses, so at most one can be truly massless. The charged leptons interact via the electromagnetic and weak forces, while quarks interact via all three forces of the SM. There are three quark and lepton generations in Figure 2.1; each generation appears to be identical in every respect other than the particle masses.

2.1.2 Quantum field theory

The SM is constructed by a relativistic quantum field theory in which particle fields that permeate all space are quantised. Excitations in each field correspond to a particle of that field. The Lagrangian formulation is used to derive the dynamics of and interactions between different particle fields. The most general Lagrangian with massless neutrinos has 19 free parameters which have been experimentally determined, while straightforward extensions to the SM with massive neutrinos need a further 7 free parameters which are still uncertain. The application of the concept of symmetry to the Lagrangian by imposing local gauge invariance introduces new vector fields, and this gives rise to new interaction terms in the Lagrangian which couple the fermion fields to the vector fields.

Mathematically, the SM is described by a combination of the $SU(3)_C$ group symmetry, and the unified electroweak interaction corresponding to the $SU(2)_L \times U(1)_Y$ group symmetries. The $SU(3)_C$ term corresponds to the gauge group of QCD, and all particles that have colour charge are bound by the symmetries of $SU(3)$ as indicated by the subscript C.

2.1.3 The electroweak interaction

The weak charged-current, mediated by the $W^\pm$ bosons, couples to left-handed helicity states for particles and right-handed ones for anti-particles. Helicity is the sign of the dot product of a particle's spin with its direction of motion. More accurately, the weak force couples to left-handed *chiral* states for particles and right-handed chiral states for anti-particles. These correspond to left-handed and right-handed helicity states, respectively, in the limit of low mass with respect to particle energy. This weak charged-current interaction is associated with an $SU(2)_L$ local gauge symmetry, which gives rise to the W^+ and W^- bosons and a neutral gauge field. In the Glashow–Salam–Weinberg model of electroweak unification, this neutral field mixes with a photon-like field of the $U(1)_Y$ gauge symmetry to give the physical photon and Z boson fields. The field associated with this symmetry couples to a kind of charge termed weak hypercharge, Y .

Additionally, the weak interaction does not couple to the physical quarks but to weak eigenstates that correspond to a linear combination of the mass eigenstates with some mixing defined by the Cabbibo-Kobayashi-Maskawa (CKM) matrix. The result of this is that the weak interaction does not conserve quark flavour. There is an analogous mixing for neutrinos defined by the Pontecorvo-Maki-Nakagawa-Sakata (PMNS) matrix, where the weak flavour eigenstates— ν_e, ν_μ, ν_τ —are a linear combination of the neutrino mass eigenstates— ν_1, ν_2, ν_3 .

2.1.4 QCD and hadronisation

The structure of the $SU(3)_C$ group implies three colour charges: red (r), green (g), and blue (b). Quarks exist in the fundamental representation of $SU(3)$ and thus carry a single colour charge (r, g , or b), while anti-quarks carry an anti-colour charge ($\bar{r}, \bar{g}$, or $\bar{b}$). Gluons, however, exist in the adjoint representation of $SU(3)$, and thus have eight possible colour

charges. These can have many different representations but are always a combination of a colour and an anti-colour, such as $r\bar{b}$ or $(g\bar{b} + \bar{g}b)/\sqrt{2}$.

The coupling constant of QCD is large at low energy scales, meaning that low-energy QCD processes are not calculable using perturbation theory—this feature of QCD is known as *confinement*. Its consequence is that free quarks and gluons (partons) *hadronise*, to form colour singlet states of mesons ($r\bar{r}$, $g\bar{g}$, $b\bar{b}$) and baryons (rgb , $\bar{r}\bar{g}\bar{b}$), collectively termed *hadrons*. Quarks can also form higher multiplicity bound states—in recent years, both tetraquark and pentaquark states have been confirmed.^[5–9] However, at higher energies such as those in modern high-energy collider experiments, the QCD coupling constant is sufficiently small that perturbation theory can be used.

When a pair of quarks ($q\bar{q}$) or gluons ($g\bar{g}$) are produced in a high-energy collision or decay, if their momenta are sufficiently collimated they can bind together to produce a meson. If instead two quarks are produced with momenta in opposite directions, it will likely become energetically favourable for an additional $q\bar{q}$ pair to be produced from the vacuum to reduce the energy necessary to maintain the distance between the quarks. These new quarks will combine with the original pair to form two colourless mesons that proceed to travel apart. This is the simplest example of hadronisation. In general, especially at higher energies, there are many $q\bar{q}$ pairs produced and thus many hadrons are formed. These will generally be collimated around the direction of the original quark or gluon produced in the high-energy collision, and the sum of their momenta should be equal to that of the original particle. The set of collimated hadrons resulting from hadronisation is known as a particle *jet*.

2.1.5 The Higgs boson

Local gauge invariance has a solid experimental basis from the success of the SM in describing experimental data, including high-precision electroweak measurements. However, local gauge symmetry can only be satisfied if the gauge boson of an interaction is massless. While this is not a problem for QED and QCD where the gauge bosons are massless, it is in apparent contradiction with the observation of the large masses of the $W^\pm$ and Z bosons.^[10] This issue can be resolved through the introduction of spontaneous symmetry breaking and the Brout-Englert-Higgs mechanism. This mechanism is based on

introducing a doublet of complex scalar fields (the Higgs field) to the SM, as a result of which the vacuum state of the Universe is degenerate. The spontaneous breaking of this symmetry, when combined with the underlying $SU(2)_L \times U(1)_Y$ gauge symmetry of the electroweak model, provides masses to the $W^\pm$ and Z bosons while leaving the photon massless.

The Higgs field has a non-zero vacuum expectation value, unlike the fields associated with the other fundamental fermions and bosons, which sets the mass scale for the electroweak bosons. The interaction between the fermion fields and the non-zero expectation value of the Higgs field provides a gauge-invariant mechanism for generating the masses of the SM fermions, where the strength of their interactions with the Higgs field are directly proportional to their mass—known as Yukawa interactions.

There is a spin-0 particle associated with the Higgs field, the Higgs boson, whose mass is not directly predicted by theory. On 4th July 2012, the discovery of a new particle with a mass of 125 GeV was announced jointly by the ATLAS and CMS experiments.^[1,2] Since then, many properties of this particle such as its spin and charge have been found to be consistent with it being a SM Higgs boson.

2.2 Proton-proton collisions

At the high energies and correspondingly small distance scales probed at the LHC the colliding protons cannot be represented as point-like particles. Instead, collisions occur between quarks and gluons. The proton constituents, *valence* quarks (u -quarks and d -quarks) and gluons as well as *sea* quarks and anti-quarks that arise due to quantum fluctuations, can all take part in the interaction. These particles—collectively referred to as *partons*—carry some fraction, x , of the proton's energy. Parton distribution functions (PDFs) which describe the probability of a given parton having a particular x value are shown in Figure 2.2 for two different values of the square of the four-momentum transferred, Q^2 .

The LHC collides bunches of protons every 25 ns at a centre-of-mass energy, $\sqrt{s}$, of 13 TeV. Much of the time, protons in each bunch pass straight through the collision point without interacting. When two protons do have a high energy inelastic collision, this is referred to as a *hard scatter* event. Other than the hard scatter, additional interactions occur which




Figure 2.2: The MMHT 2014 NNLO proton PDFs at scale $Q^2 = 10 \text{ GeV}^2$ and 10^4 GeV^2 as a function of fractional energy of the proton that a parton claims. The gluon PDF is divided by 10 to improve the readability of the plot.^[11]

make up the *underlying event*. The two protons with partons involved in the hard scatter fragment—once the two partons have interacted, the protons are no longer stable hadrons. The remaining partons rearrange to form stable objects with different momenta, generally producing activity aligned with the direction of the incoming particles. Occasionally, another set of partons will interact in an elastic or inelastic collision, causing transverse activity. In addition to this, lower-energy parton collisions from the remaining protons in the bunch may occur. These additional collisions are referred to as *pileup*.

Charged particles (whether electromagnetically or colour charged) emit radiation when they are accelerated. Therefore, before partons interact, they may radiate low-momentum (*soft*) photons and gluons—this is called *initial state radiation*. Radiation can also occur from the outgoing partons and the decay products of particles produced in the interaction, in which case it is termed *final state radiation*. Usually, both initial and final state radiation are soft compared to the energy of the hard scatter, so the gluons or photons have little impact on the event. Particles can also be radiated with a much higher momentum, in which case they will impact the topology of the final event.

It is clear that the environment of a pp collision can be very complicated and messy. The simulation of the many contributions that constitute a pp collision event will be discussed in §3.5.

2.3 Limitations of the Standard Model

The Standard Model has evolved with our understanding of subatomic phenomena, and has been incredibly successful. However, there are still many phenomena that are outside the scope of the current SM. These include dark matter, dark energy, unification with gravity, and the matter–anti-matter asymmetry in the Universe today. There are myriad theories of the nature of physics beyond the Standard Model (BSM), and much work has been done to prove or disprove these. This thesis presents the results of a search for a few of these models, which could provide an explanation for the following three open questions of the SM.

2.3.1 Dark matter and dark energy

Dark matter’s existence was first surmised from the rotation curves of stars and galactic clusters, and is now supported by many other measurements and observations including the gravitational lensing of background objects in the presence of stars, cosmic microwave background radiation properties, and large-scale structure formation.^[12–15] By fitting a theoretical model of the composition of the Universe to all cosmological observations, the matter described by the SM is estimated to account for only about 5% of the energy in the observable Universe. Around 27% is in the form of dark matter, and the remaining 68% is dark energy.^[16] Yet still the particle nature of dark matter is unknown and the SM provides no plausible candidates. Neutrinos are too relativistic (hot) to be adequate candidates for the majority of dark matter, while neutral hadrons are found in an abundance too small to account for the amount of dark matter observed in other measurements. Dark energy is even more mysterious—it was hypothesised to account for observations of distant type Ia supernovae, which show the Universe going through an accelerated period of expansion. It is thought to be a form of energy that exerts a negative, repulsive pressure, behaving like the opposite of gravity.

2.3.2 The hierarchy problem

Generally, a hierarchy problem is a large and seemingly unexplained difference in the magnitudes of related quantities. In the calculation of the Higgs boson mass (m_H), the higher order corrections to m_H from top quark loops result in a predicted mass of the order of

the Planck mass, $m_{\text{Planck}} = 10^{19}$ GeV. However, the Higgs boson has been observed with a mass of just 125 GeV.

The simplest solution to the hierarchy problem involves so-called *fine tuning*, in which parameters in the SM have to match to an extremely large number of significant figures to accurately describe the observed Higgs boson mass. This is not impossible, but the unlikelihood of it occurring by chance alone is certainly motivation to probe for other underlying mechanisms.

A more elegant solution to this problem could be introducing supersymmetry into the SM, where each SM particle has a superpartner, and each loop contribution from a quark is exactly cancelled out by the contribution from its superpartner. One of the models considered in this thesis is a supersymmetry model, and will be discussed in more detail in §2.4.3.

2.3.3 Baryon asymmetry in the Universe

The asymmetry in the presence of matter versus anti-matter in the observable Universe, evidenced by our very existence, is also unexplained by the SM. The amount of asymmetry that could result from currently-known interactions allowed by the SM that violate charge–parity invariance is insufficient to account for the abundance of matter in the current Universe. A model which could generate the observed baryon asymmetry via a Higgs-portal mechanism is considered in this thesis.

2.4 Physics beyond the Standard Model

Making a discovery is possible in two ways—trying to discover physics that has been predicted in a well-motivated way, and searching for new phenomena that are not necessarily predicted by any theoretical models but are not ruled out by current constraints. Given the wide range of outstanding issues in the SM, it is not surprising that there are a plethora of distinct BSM models proposed by theorists. However, it is inefficient and often impossible for experimentalists to perform dedicated searches for each of these. Instead, analyses are increasingly designed to be sensitive to features common to many models. One such fea-




Figure 2.3: Particle decay length, $c\tau$, in metres, as a function of particle mass, in GeV, for a variety of Standard Model baryons (blue), mesons (green), leptons (red), bosons (orange), and quarks (purple).

ture, appearing in many models intended to explain phenomena which the SM does not address, is long-lived particles (LLPs).

This thesis presents the results of a search for neutral LLPs decaying hadronically within the ATLAS detector (corresponding to proper lifetimes up to ~ 100 ns). The analysis described in Chapter 4 was initially designed to be sensitive to a hidden sector benchmark model and was subsequently reinterpreted in the context of a variety of other benchmark models, all of which feature LLPs, as described in Chapter 6. This section presents the general motivation for LLP searches and a description of the models considered in the analysis.

2.4.1 Long-lived particles

Particles in the SM have lifetimes spanning many orders of magnitude, from the Z boson ($\tau \sim 2 \times 10^{-25}$ s) to the electron ($\tau \gtrsim 10^{28}$ yr) and proton ($\tau \gtrsim 10^{34}$ yr). The variety of SM particles' masses and proper decay lengths is illustrated in Figure 2.3. Similarly, BSM theories often predict new particles with a variety of masses and lifetimes. In particular, weak-scale particles can have long lifetimes for several reasons, including small couplings between the LLP and lighter states, and suppressed phase space available for decays.

Particles travelling near the speed of light can have macroscopic, detectable displacements between the production and decay points.

The possible characteristics of LLPs and their decays are varied, depending on the model in which they are produced: they may be neutral, electrically- or colour-charged; they may decay to photons, a single photon and invisibles, hadronically, leptonically, or semi-leptonically. Combined with the displacement of the decay, this results in possible experimental signatures that are very different from those expected in SM processes, discussed further in §3.6.

2.4.2 Hidden sectors

In general, hidden sector models are a class of BSM theories in which there is another sector of new particles invisible to the SM except through a massive mediator particle.^[17–20] If the mediator has a high mass and the mixing is sufficiently weak, then such a sector could remain hidden until experiments generate collisions at high enough energies to access it regularly. Hidden sector particles that decay back to the SM may do so with a long lifetime, such that in a collider experiment a substantial fraction of their decays would occur at macroscopic distances from the point at which they were produced. This also provides an easy means of accommodating dark matter—if the lightest hidden sector particle is stable, meaning it cannot decay to the SM via the communicator particle, then it is a candidate for dark matter.

2.4.2.1 Hidden sector scalar model

In the model first considered by the analysis described in Chapter 4, the exact nature of the decays and particles in the hidden sector is not important. A generic benchmark model was used where the hidden sector is weakly coupled to the SM sector by a mediator particle which mixes with its hidden sector counterpart. In this model, a neutral scalar boson, Φ , decays to a pair of scalars or pseudoscalars S , as shown in Figure 2.4, with the details at the vertex taken to be unimportant. The Φ boson could be the SM Higgs boson,^[21] or some other BSM boson that has yet to be discovered, and this analysis considered m_Φ ranging from 125 GeV to 1000 GeV. The scalars, S , decay back to SM sector particles via the mediator (off-shell), and thus possess long lifetimes. They are invisible to the SM,




Figure 2.4: Diagram of the hidden sector scalar (HSS) model. A heavy neutral scalar boson Φ decays to two neutral long-lived scalars S which in turn decay to pairs of SM fermions.

and decay to SM fermions. Their coupling mixes with the SM Higgs Yukawa coupling which increases with fermion mass, so the decays are predominantly to kinematically-accessible[†] heavy quarks. The branching ratio among the different decays depends on the mass of the scalar, but for $m_S \geq 25$ GeV it is almost constant: 85% to $b\bar{b}$, 5% to $c\bar{c}$, and 8% to $\tau^+\tau^-$. The lifetime of the S is tunable in the model, allowing for macroscopic decay lengths, and the analysis considered proper decay lengths from a few tens of millimetres to 50 m.

2.4.2.2 Dark photon model

Another possible hidden sector model, the Falkowski–Ruderman–Volansky–Zupan (FRVZ) model,^[22,23] was considered in the study presented in Chapter 6. In this model a pair of dark fermions, f_{d2} , are produced via a heavy boson decay, shown in Figure 2.5. Either the dark fermion decays into a dark photon, γ_d , and a lighter dark fermion assumed to be the hidden lightest stable particle (HLSP), or the dark fermion decays into the HLSP and a dark scalar, s_d , that in turn decays into a pair of dark photons.

The hidden sector is connected with the SM through kinetic mixing of the dark photon and the SM photon. With this mixing, a dark photon with a mass, m_{γ_d} , up to a few GeV will decay into light SM fermions, with branching fractions that depend on m_{γ_d} .^[24–26] In the case where the dark photon mass is small, they are expected to be produced with large

[†]The range of scalar masses considered includes one case ($m_S = 400$ GeV) for which decays to top quarks ($m_t = 173$ GeV) would be kinematically accessible, but would not give a significantly different signature since all hadronic jets look alike in the calorimeter, so for consistency these decays were not included in the model.




Figure 2.5: Diagrams of the FRVZ model with dark photon final states. The dark fermions f_{d_2} each decay into an HLSP and a dark photon γ_d in the diagram on the left, and an HLSP and a dark scalar s_d that in turn decays into a pair of γ_d in the diagram on the right. The γ_d decay into SM fermions, denoted by ℓ^+ and ℓ^- .

boosts, so their decays result in collimated groups of leptons and light hadrons in a jet-like structure, known as lepton-jets.

In the model considered in Chapter 6, the heavy boson could be the SM Higgs boson,^[21] or an as-yet undiscovered heavier boson with mass 800 GeV. Additionally, the model is simplified by imposing that only electron, muon or pion pairs are produced by the dark photons, even though other light hadrons or τ leptons could be produced for higher dark photon masses. More dark photons could be produced by hidden sector radiation but the number of additional dark photons introduced in this way is highly model-dependent and scales with the size of the dark gauge coupling. This radiation is not included in the FRVZ models, corresponding to a dark gauge coupling of $\lesssim 0.01$.

2.4.3 Supersymmetry

As mentioned previously, a natural extension to the SM which would solve the hierarchy problem is supersymmetry (SUSY). As the name suggests, this adds an additional symmetry to the SM between particles of different spin. In SUSY, every SM fermion (spin- $\frac{1}{2}$) gains a bosonic (spin-1) superpartner, denoted by adding an “s” to the start of the particle name (e.g. the electron has a superpartner called the selectron). Every SM boson (spin-1) gains a fermionic (spin- $\frac{1}{2}$) superpartner, denoted by adding “ino” to the end of the particle

name (e.g. the gluon's superpartner is the gluino).[†] In the simplest SUSY theories, where SUSY is an unbroken symmetry, each pair of superpartners would share the same mass and internal quantum numbers besides spin. However, these theories predict that the superpartners would have been discovered at the same mass as the SM particles. As no evidence of this has been seen, if SUSY exists it must be a spontaneously broken symmetry allowing superpartners to differ in mass. There are many constraints from squark and gluino searches on the masses allowed for the superpartners—recent results push this scale very high, to the point where fine-tuning is required to explain the mass discrepancies, or something additional must be necessary.

In SUSY, each particle has an additional property called R -parity, which, in most theories, must be conserved for all interactions and decays. There is also a special class of SUSY theories that seek to avoid this restriction by allowing R -parity violating interactions. The R -parity of a particle is determined by its baryon number, B , lepton number, L and spin, s : $R = (-1)^{3B+L+2s}$. The product of the R -parities of outgoing particles must be equal to the product of the R -parities of all incoming particles. SM particles have an R -parity of +1 and SUSY particles have an R -parity of -1 . Consequently, any SUSY particle that decays to an SM particle *must additionally produce a SUSY particle*. Therefore, the lightest supersymmetric particle (LSP) cannot decay—there are no lighter SUSY particles for it to decay to while conserving R -parity.

Many SUSY models have the neutralino as the LSP, meaning it is stable, heavy and non-interacting. This has two advantages: the LSP is a dark matter candidate; and it leaves a very distinct signature when it is produced in a high-energy collision, namely a large amount of energy escaping the detector. Most SUSY searches setting strong limits on the masses of gluinos or squarks rely on the presence of heavy, stable neutralinos, and so it is interesting to construct a model that circumvents these limits without this LSP structure.

2.4.3.1 Stealth SUSY

Stealth SUSY models^[27,28] are a subclass of theories with R -parity conserving SUSY^[29–34] whose signatures do not include large amounts of undetected energy. In Stealth SUSY, particles are nearly mass-degenerate with their superpartners, such that in the decay

[†]The symbols for most superpartners are the same as their SM partners but with an added tilde over the top.




Figure 2.6: Diagram of the Stealth SUSY model. The LLP $\tilde{S}$ is represented by double lines and final-state SM gluons are labeled as g .

$\tilde{X} \rightarrow X + \tilde{N}$, $m_{\tilde{N}} \ll m_{\tilde{X}} \approx m_X$. To ensure that there is no large missing-energy signature, two requirements must be true: $\tilde{N}$ must not carry away a large amount of energy even though it is neutral and escapes the detector (guaranteed due to the decay kinematics); and $\tilde{X}$ must decay within the detector.

This can be accomplished in many ways, but the study described in Chapter 6 explores a model in which an approximately supersymmetric hidden sector is added at the electroweak scale as a sector that the Minimal Supersymmetric Standard Model (MSSM) can decay to. The model introduces a singlet superfield, S , with a superpartner singlino, $\tilde{S}$. The hidden sector is weakly coupled to the MSSM, and if SUSY-breaking occurs at a low scale, the mass difference between these states, $\delta M(\tilde{S}, S)$ will be small.[†]

In the process shown in Figure 2.6, a gluino $\tilde{g}$ decays to a gluon g and a singlino $\tilde{S}$. The singlino decays to a singlet plus a gravitino ($\tilde{S} \rightarrow S\tilde{G}$). The decay width of the singlino, $\Gamma_{\tilde{S} \rightarrow S\tilde{G}}$, and therefore its lifetime, $\tau_{\tilde{S}}$, is determined by the size of the mass splitting and the SUSY-breaking scale F : $\Gamma_{\tilde{S} \rightarrow S\tilde{G}} \propto m_{\tilde{S}}(\delta M)^4/F^2$.^[27] Since F is not fixed, the singlino could travel appreciably through a particle detector before decaying, resulting in a displaced signature. The gravitino carries off very little energy and the singlet promptly decays to two gluons, which are collimated. Since R -parity is conserved, two gluinos are produced in each event, and the final state is two prompt jets and two displaced jets.

[†]This is also possible, albeit more complex, in the case of higher-scale SUSY-breaking.




Figure 2.7: Diagram of the Higgs-portal baryogenesis model. The LLP χ is represented by double lines and final-state SM fermions are labeled as f .

2.4.4 Baryogenesis

The final class of BSM theories explored in this thesis seeks to explain the baryon-antibaryon asymmetry in the observed Universe. The generation of this asymmetry, a process termed baryogenesis, requires a departure from thermal equilibrium in the early Universe.^[35] Many baryogenesis models feature the out-of-equilibrium decay of a new massive particle. When the cosmic temperature drops below the mass of this particle, the out-of-equilibrium requirement is optimally satisfied. In the case where the mass is electroweak-scale, the out-of-equilibrium condition requires a proper decay length for the new particle of $c\tau \gtrsim \mathcal{O}(1 \text{ mm})$.^[36] In a proton-proton collider, the new particle would be pair-produced, and leave a distinctive signature of two displaced jets or decay vertices. A set of low-scale ($< \text{TeV}$) baryogenesis models that generate the observed baryon asymmetry via decays of weak-scale states can be directly tested at colliders.^[36–38]

2.4.4.1 Higgs-portal baryogenesis

In the specific baryogenesis model considered in the study in Chapter 6, the baryon asymmetry is produced through the freeze-out and decay of a meta-stable weakly-interacting massive particle (WIMP). Baryogenesis occurs if the WIMP is sufficiently long-lived and decays through baryon- or lepton-number violating processes.^[39] In this case, the LLP is chosen to be a Majorana particle (meaning it is its own antiparticle) and provides a simple realisation of baryon-number, lepton-number and CP-symmetry violation. A Higgs-portal model^[39] is considered where a scalar boson mixes with the SM Higgs boson.^[21] This boson is assumed to be heavy and decouples, leaving the production of a

long-lived singlet, χ , via the exchange of a single SM-like Higgs boson after mixing, as shown in Figure 2.7.

For the production of χ through the Higgs-portal, two different regimes are considered:

- $m_\chi < m_h/2$: here the dominant production mechanism is through an *on-shell* Higgs boson. The χ production at $\sqrt{s} = 13$ TeV is expected to be copious, and despite existing strong constraints, this regime remains potentially very interesting.
- $m_\chi > m_h/2$: in this regime the Higgs is *off-shell* and the signal rate falls rapidly for increasing m_χ , even for large mixing. The cross section expected for a χ mass of 100 GeV is approximately 7 fb at leading order, so the sensitivity is not expected to be very high.

The decay modes for χ must violate baryon- and/or lepton-number conservation to generate baryonic asymmetry; the corresponding lowest-dimensional interactions allow decays to three SM fermions. The following four decay modes and their charge conjugates are considered: $\chi \rightarrow \tau^+ \tau^- \nu_\ell$, $\chi \rightarrow cbs$, $\chi \rightarrow \ell^\pm cb$, $\chi \rightarrow \nu b\bar{b}$.

2.5 Summary

The incompleteness of the SM necessitates searching for BSM physics processes, but it is unclear which new models are more likely than others among the large number that have been proposed. Therefore, an effective strategy for searching for such processes is to focus on common features of many BSM models. The existence of long-lived particles is predicted by all the models described here, whose signatures are distinctive in particle detectors and may previously have been overlooked. The following chapters present a search for long-lived particles using the ATLAS detector at the LHC.

Chapter 3

The ATLAS experiment

The work presented in this thesis relies on data recorded during high energy proton-proton (pp) collisions by the ATLAS detector. This chapter provides an overview of the experimental apparatus and reconstructed objects used throughout.

The Large Hadron Collider (LHC), which delivers pp collisions to the ATLAS experiment among others, is described in §3.1. An overview of the ATLAS detector and its subsystems is provided in §3.2, with the trigger system described in §3.3. The reconstruction of physics objects produced in collisions is outlined in §3.4. Next, §3.5 describes the production of simulated events for use in physics analyses. Finally, §3.6 details the variety of experimental signatures of long-lived particles (LLPs) decaying in the ATLAS detector, one of which is the focus of the remainder of this thesis.

3.1 The LHC

The LHC is the world's largest particle accelerator, at 27 km in circumference, operating at a centre-of-mass energy, $\sqrt{s}$, of 13 TeV. It lies approximately 100 m underground, below the European Center for Nuclear Research (CERN) close to the border between France and Switzerland, near Geneva. Occupying the tunnel that previously housed the Large Electron-Positron Collider (LEP), it was designed to collide hadrons at a maximum centre-of-mass energy of 14 TeV, and with an instantaneous luminosity of $10^{34} \text{ cm}^{-2}\text{s}^{-1}$ (defined in §3.1.3). Although the LHC is capable of accelerating heavy ions such as lead, for most of its operation it accelerates protons, and all studies presented here will focus on pp collisions. To date, there have been two operational runs of the LHC: Run 1 began in 2009 and ended




Figure 3.1: Overview of the CERN accelerator complex and the LHC.^[44]

in 2013, and Run 2 occurred between 2015 and 2018. Run 3 is expected to begin in 2021, and the high-luminosity LHC (HL-LHC), which will increase the instantaneous luminosity by a factor of 10, is scheduled to begin taking data in 2027.

The entire CERN accelerator complex is shown in Figure 3.1, including the series of smaller accelerators that protons are sent through to incrementally increase their energy before they are injected into the LHC in two beams travelling in opposite directions. Situated around the LHC ring are the four main physics experiments of the LHC—ALICE,^[40] CMS,^[41] LHCb,^[42] and ATLAS^[43]—located at crossing points where the two beams are collided.

3.1.1 Overview

Every proton to be collided in the LHC starts in a small bottle of hydrogen gas. Molecular hydrogen from this bottle is passed through an electric field and the hydrogen nuclei (protons) are extracted. From there they are sent into the LINAC2, a linear accelerator,

which accelerates the protons to an energy of 50 MeV using electric fields which switch from positive to negative at a given frequency such that the particles accelerate in closely-spaced bunches. The proton bunches then pass to the Proton Synchrotron Booster, which accelerates them to 1.4 GeV before sending them into the Proton Synchrotron. The latter accelerator further boosts the energy of the protons to 25 GeV before they are directed into the Super Proton Synchrotron (SPS), where they are accelerated to 450 GeV. Following this, they are split into two beams and delivered into the LHC.

The 27 km LHC ring is made up of a series of radio frequency (RF) cavities for acceleration, as well as magnets for focusing and bending the beams. These cavities are shaped to resonate at specific frequencies, allowing radio waves to interact with passing bunches. Each time a beam passes the electric field in an RF cavity, some energy is transferred from the radio waves to the protons. There are eight points on the ring at which beams may cross—four house the detectors, and the others are occupied by kickers and collimators. In Run 2, the LHC accelerated proton beams to 6.5 TeV each, only 0.5 TeV below the design energy.

3.1.2 Bunch trains and filling schemes

During the acceleration process, and before they reach the LHC, the original proton bunches from the LINAC2 are split into more bunches using a multiple splitting process.^[45] Upon reaching the LHC, each of these bunches contains up to 1.2×10^{11} protons. As a result of the parameters of the LHC machinery, the number of potential wells (known as *buckets*) in each beam in which the proton bunches may sit exceeds the number of proton bunches. There are a total of 35640 possible buckets that can contain a proton bunch in each beam, but only every tenth bucket is ever filled, so there are a maximum of 3564 proton bunches. However, not every one of these ends up filled with protons either—for example, a set of empty bunches is left after each bunch train to allow enough time for the magnets to be turned on to direct the beams to the beam dump.

The 2016 dataset used in this thesis was collected with a bunch spacing[†] of 25 ns. This means that in principle the ATLAS detector has the possibility of recording a collision every 25 ns. A *bunch crossing* is recorded each time two bunches pass through the collision

[†]It should be noted that the *bucket spacing* is in fact 2.5 ns, but since only every tenth bucket is ever filled by a proton bunch, the *bunch spacing* is recorded as 25 ns.

point from each side. When the bunches are both filled with protons, the bunch crossing is considered *paired*. For the study of some types of non-collision backgrounds, as was required in the analysis described in Chapter 4, it is useful to have data collected during bunch crossings when one or both of the bunches are empty. When two empty bunches cross, the bunch crossing is labelled *empty*. If the empty bunch crossing is within five bunch crossings of a paired bunch crossing, it falls into the category *empty before/after paired*. If one bunch is filled and the other is not, the crossing is *unpaired*, and if the unpaired bunch crossing is at least seven bunch crossings away from a paired crossing, it is *isolated*. In practice, the distance from a paired crossing is important because it means any residual activity from a collision should have dissipated. However, some protons in filled bunches may be knocked into neighbouring empty buckets, and when this happens in both beams, there can be *ghost* collisions out-of-time with the bunch crossing by multiples of 2.5 ns (the bucket spacing).

3.1.3 Luminosity and pile-up

To calculate the number of physics events expected during any data-taking period, one must know the rate at which protons collide. This is called instantaneous luminosity, $\mathcal{L}$, and is determined from parameters of the beam according to,

$$\mathcal{L} = f_r n_b \frac{N_b^2}{2\pi \Sigma_x \Sigma_y}, \quad (3.1)$$

where f_r is the LHC revolution frequency, n_b is the number of bunches in the beam, N_b is the number of particles in each bunch, and $\Sigma_{x(y)}$ is the mean beam width in the x (y) direction.

The integrated luminosity is related to the instantaneous luminosity by,

$$L = \int \mathcal{L} dt. \quad (3.2)$$

This is measured by detectors at the experiments along the ring, not simply calculated from the delivered beam parameters. ATLAS, for example, has two dedicated detectors to measure the luminosity—LUCID and the Beam Conditions Monitor (BCM).^[46,47] Having two separate detectors decreases the uncertainty on the luminosity measurement, and they are complemented by luminosity measurement algorithms in the standard subdetectors.




Figure 3.2: Cumulative pp luminosity by day delivered by the LHC (green) and recorded by ATLAS (yellow) for the 2016 data-taking period at $\sqrt{s} = 13$ TeV.^[48] The recorded luminosity is lower than the delivered luminosity due to inefficiencies in data acquisition.

Figure 3.2 shows the cumulative luminosity delivered by the LHC and recorded by the ATLAS detector for the 2016 data-taking period at $\sqrt{s} = 13$ TeV.^[48]

Finally, the number of times a given process occurs at the LHC, N , depends on the integrated luminosity, L , as

$$N = \sigma L, \quad (3.3)$$

where σ , the process cross-section, is the underlying quantum mechanical probability that the process interaction will occur.

Each bunch entering the LHC contains $\sim 10^{11}$ protons, and is about 1 mm wide in the plane perpendicular to the beam direction (the *transverse* plane) and several centimetres long. While generally there is one hard scattering process of interest in any paired bunch crossing, there are also a large number of softer collisions taking place, referred to as pile-up[†]. The mean number of interactions per bunch crossing, μ , is calculated as,

$$\mu = \frac{\mathcal{L} \sigma}{n_b f_r} \quad (3.4)$$

where $\mathcal{L}$ is the instantaneous luminosity per bunch, σ is the inelastic cross section, n_b is the number of bunches, and f_r is the LHC revolution frequency. Figure 3.3 shows the

[†]As well as this *in-time* pile-up, detector elements may also record collisions from bunch crossings preceding or following the current bunch crossing if their response time is greater than the bunch separation; in this case the pile-up is described as *out-of-time*.




Figure 3.3: The mean number of interactions per bunch crossing for the 2016 ATLAS pp dataset.^[48]

mean number of interactions per bunch crossing for the 2016 dataset used in this thesis, with $\langle\mu\rangle = 24.9$.^[48]

3.2 The ATLAS detector

The ATLAS experiment^[43] is a multi-purpose detector at the LHC, designed to give as complete a picture as possible of the production and decay of particles created in proton-proton collisions, thus allowing a wide physics program to be delivered, from high-precision measurements to searches for new physics. The detector (illustrated in Figure 3.4) was designed with multiple subsystems, in order to make precise measurements of the properties of a variety of particles within a large volume.

ATLAS is barrel-shaped, and oriented with its axis of radial symmetry along the beam pipe. The long cylindrical middle portion is called the *barrel* and the end portions on either side (wheel-shaped to cover as much area as possible in the transverse plane) are termed the *endcaps*. From the beam pipe outwards, the ATLAS subdetectors are the inner detector, the calorimeters and the muon spectrometer. Each of these subdetectors is composed of several subsystems, which will be described in more detail later in this section. In addition to these subdetectors, ATLAS has a magnet system which provides a 2 T solenoidal magnetic field within the inner detector and a 0.5 T toroidal magnetic field within the muon spectrometer. These magnetic fields bend the trajectories of charged particles passing through them, allowing measurements of charge and momentum to be made.




Figure 3.4: Overview of the ATLAS detector with subdetectors labelled.^[43]

A right-handed coordinate system is used to describe the ATLAS detector, with the z -axis following the beam line, the x -axis pointing inwards to the centre of the LHC ring, and the y -axis pointing vertically upwards. The azimuthal angle is represented by ϕ , spans $(-\pi, \pi]$ and starts from the x -axis. The polar angle θ spans the range $[0, \pi]$, but is usually replaced by the pseudorapidity, η , defined as,

$$\eta = -\ln(\tan(\theta/2)). \quad (3.5)$$

A measurement often used in analyses as a measure of the closeness of two particles or objects is ΔR , defined as,

$$\Delta R = \sqrt{(\Delta\eta)^2 + (\Delta\phi)^2}. \quad (3.6)$$

3.2.1 Inner detector

The inner detector (ID) is designed to measure tracks from charged particles with energy as low as 100 MeV, and is formed of four main subsystems, shown in Figure 3.5. It has a total radius, $r = \sqrt{x^2 + y^2}$, of 1.2 m and length in z of 6.2 m. Each subsystem is located concentrically around the beam pipe, inside a 2 T axial magnetic field created by the ATLAS solenoid magnet. Closest to the beam line is the insertable b -layer (IBL), followed by the




Figure 3.5: Diagram of the ATLAS inner detector showing all its components, with the distances to the interaction point labelled.^[49]

pixel detector, then the semiconductor tracker (SCT) and finally the transition radiation tracker (TRT). As particles traverse the layers of the ID they interact with the active material in each, resulting in local ionisation. By combining the locations at which the interactions occurred, ATLAS software is able to reconstruct the paths along which particles passed (the *tracks*). The magnetic field within the ID causes the trajectory of charged particles to bend, allowing both their charge and momentum to be determined from the track curvature. The target track momentum resolution, σ_{p_T} , for the ATLAS ID as a function of the track transverse momentum, p_T , is given by,

$$\frac{\sigma_{p_T}}{p_T} = 0.05\% p_T[\text{GeV}] \oplus 1\%. \quad (3.7)$$

The tracks are also extrapolated back to the beam pipe in order to determine the interaction point (IP)—the location of the highest energy collision within the detector.

The IBL was installed during the first long shut-down of the LHC, between 2013 and 2015, and is the innermost layer of the ATLAS detector, situated at a radius of approximately 33 mm.^[50] It covers $|\eta| < 2.5$ and is formed of silicon pixels which are smaller than those in the other layers. Increasing the number of pixel layers from three to four increases the

tracking efficiency for short-lived b -quark decays, and also protects the other pixel layers from radiation effects in higher luminosity conditions.

The pixel detector is the next subsystem, and consists of three cylindrical pixel layers in the central barrel region, and two sets of three disks (one at either end). In the barrel region, the layers cover $50.5 < r < 122.5$ mm (again up to $|\eta| < 2.5$). Each silicon pixel sensor is $50 \times 400 \mu\text{m}^2$, providing a hit accuracy of $10 \mu\text{m}$ in $(r - \phi)$ and $115 \mu\text{m}$ in z .

The next layer is the SCT, which consists of four layers of silicon microstrip pairs in the barrel region, covering $255 < r < 549$ mm, and two sets of nine disk layers in the endcap regions, covering $275 < r < 560$ mm for $839 < |z| < 2735$ mm. The uncertainties in the position measurements from the SCT are $17 \mu\text{m}$ in the transverse direction, and $580 \mu\text{m}$ in the longitudinal direction.

The outermost layer of the ID is the transition radiation tracker. The TRT is constructed from 2 mm radius straw trackers, filled with xenon gas. In the barrel, the straws are parallel to the beam pipe and cover $560 < r < 1080$ mm for $|z| < 720$ mm. In the endcap, straws are perpendicular to the beam pipe in a fan-like configuration, covering $617 < r < 1106$ mm and $827 < |z| < 2774$ mm. The overall coverage is $|\eta| < 2.0$, with an accuracy of $130 \mu\text{m}$ per straw and ~ 35 hits per track. The TRT additionally offers some particle identification capabilities. The space between tracker straws is filled with materials of various refractive indices, which causes passing particles to radiate photons. Since particles with the same momentum but different masses have different Lorentz factors, this aids in their identification; lighter particles such as electrons have a larger Lorentz factor which means they radiate more photons than heavier hadrons.

3.2.2 Calorimeters

The calorimeter system is split into two parts: the electromagnetic calorimeter (ECal), and the hadronic calorimeter (HCal). The layout of the calorimeters is shown in Figure 3.6. The ECal generally absorbs all electromagnetic energy, since the final decay products in electromagnetic showers will always be electrons and photons, which the calorimeter measures. However, hadronic energy is not fully measured in the HCal because hadrons can decay into undetectable neutrinos or minimally-interacting muons, and because hadrons can lose some of their energy during nuclear interactions with the material of the detector.




Figure 3.6: Diagram of the ATLAS calorimeters.^[43]

Table 3.1: The ATLAS calorimetry energy and time resolutions.

	Component	Energy resolution	Time resolution
ECal	Barrel and endcap	$\frac{\sigma_E}{E} = \frac{10\%}{\sqrt{E[\text{GeV}]}} \oplus 0.7\%$	$\sim 0.1 \text{ ns}$
HCal	Barrel and endcap	$\frac{\sigma_E}{E} = \frac{50\%}{\sqrt{E[\text{GeV}]}} \oplus 3\%$	$\sim 1 \text{ ns}$
	Forward	$\frac{\sigma_E}{E} = \frac{100\%}{\sqrt{E[\text{GeV}]}} \oplus 10\%$	

These unmeasurable components of hadronic energy are accounted for with a scale factor to convert the measured energy to the estimated true energy. The ATLAS calorimeter energy and time resolutions are summarised in Table 3.1.

Combined, the two calorimeters provide coverage in the region $|\eta| < 4.9$ with between 10 and 20 interaction lengths of material[†]. Both calorimeters are constructed of alternating layers of absorbers and active sampling material. The absorbers cause particles travelling through the detector to shower, after which the energy of the shower can be measured by the sampling material. This allows the shower to be measured by the calorimeter within a smaller volume, but the exact energy of the shower must be estimated, as some of the

[†]The interaction length is the mean distance a particle will travel before undergoing an inelastic interaction with a nucleus in the surrounding medium.




Figure 3.7: The depth of each layer of the ATLAS calorimeters as measured by the nuclear interaction length.^[43]

original particle's energy is deposited in the absorbers. The depth of the calorimeters in interaction lengths as a function of $|\eta|$ is shown in Figure 3.7.

The ECal is divided into one barrel and two endcap calorimeters, each with three layers and a presampler. The sampling material in the ECal is liquid argon (LAr) and lead plates are used as absorbers. The lead plates are folded into an accordion shape, which allows for full coverage and symmetry in ϕ . Copper electrodes are located inside the folds to detect electrons. The second layer of the ECal collects most of the energy from an electromagnetic shower, since it comprises the most radiation lengths[†]. The thin presampler corrects for any energy lost by the particle as it travelled from the interaction point to the ECal. The first layer is finely grained in order to provide an accurate η measurement, and the third (outermost) layer measures the tail of the shower not absorbed by the other layers. The barrel of the ECal covers $1.42 < r < 1.99$ m, and the endcaps cover $3.68 < |z| < 4.25$ m.

The HCal is made up of a number of distinct regions and materials. The barrel ($|\eta| < 0.8$) and extended barrel ($0.8 < |\eta| < 1.7$) regions are constructed from plastic scintillator tiles and steel absorbers. Signal in the scintillator tiles is detected with photomultiplier tubes. The HCal endcaps (HEC) cover $1.5 < |\eta| < 3.2$ and use LAr as a sampling material like the ECal, though the absorber is copper, and there is no accordion geometry. Because of the LAr usage, the ECal and HEC are often referred to collectively as the LAr calorimeters. The HCal has poorer resolution than the ECal—typical of hadronic calorimeters intended to be used alongside electromagnetic calorimeters. Hadrons generally require a greater

[†]The radiation length is the distance over which a particle loses all but $1/e$ of its energy.




Figure 3.8: The ATLAS muon spectrometer.^[51]

depth of material to be fully stopped than charged leptons such as electrons, so the HCal is thicker than the ECal. Each barrel section has an inner radius of 2.28 m and an outer radius of 4.25 m, while the LAr endcap sections cover $4.35 < |z| < 6.05$ m. The HCal is intended to measure the energy of all jets (groups of collimated particles from the hadronisation of quarks and gluons), though a small fraction of the time they may *punch through* the calorimeters and end up interacting with the muon system.

3.2.3 Muon spectrometer

The largest part of the ATLAS detector is the muon spectrometer, which is the outermost subdetector system. Muons are much more penetrating than other charged particles generated by collisions and decays. They do not interact via the strong force, radiate bremsstrahlung far less than the lighter electron, yet are light enough that they can travel some distance before decaying. This means they are able to pass through the ATLAS calorimeters with minimal interactions, and so a dedicated detector is required to measure them.

Four different subsystems make up the full muon spectrometer (MS), shown in Figure 3.8. Tracking information is provided by monitored drift tubes (MDTs) covering $|\eta| < 2.7$ and cathode strip chambers (CSCs) covering $2.0 < |\eta| < 2.7$. Resistive plate chambers (RPCs) covering $|\eta| < 1.05$, and thin gap chambers (TGCs) covering $1.05 < |\eta| < 2.4$ in the endcap region, provide triggering and additional hit information for offline reconstruction. Overall, the muon spectrometer occupies $4.3 < r < 10.7$ m in the barrel ($|\eta| < 1.05$), and

$7.0 < |z| < 22.0$ m in the endcaps ($1.0 < |\eta| < 2.7$). It is designed to provide momentum measurements with a relative resolution better than 3% over a wide p_T range and up to 10% at $p_T \approx 1$ TeV.^[52]

3.2.4 Magnets

The ATLAS magnet system comprises two different subsystems. A superconducting solenoid magnet generates a magnetic field for the ID in the r - ϕ plane, while air-core toroid magnets generate a magnetic field for the MS in the r - z plane. The solenoid magnet is aligned along the beam axis and rests just inside the electromagnetic calorimeter, which shares its cryostat; it produces a 2 T magnetic field throughout the ID. A large toroid magnet (formed of eight coils arranged radially about the beam axis) covers the region $|\eta| < 1.4$, while there is a smaller toroid magnet in each endcap covering the region $1.6 < |\eta| < 2.7$. The region between the toroids is covered by the overlapping edges of the fields from the barrel and endcap toroids. The magnetic field produced by the toroid magnets is 0.5 T.

3.3 Triggering in ATLAS

The frequency of bunch crossings in the ATLAS detector can be as high as 40 MHz. With current technology and financial limitations, it is impossible to store and process the data from all these events, so a trigger system is employed to decide which events are interesting enough to keep, and which can be discarded. A rate of 40 MHz is also too high to fully process events before deciding to store them or not, so there are two distinct steps in the trigger system.

A hardware-based trigger is first used to make a fast, coarse selection of candidate events to record. This is referred to as the *Level-1* (L1) trigger, and outputs events at a rate of 100 kHz. Due to limited latencies, the L1 trigger only has information from the calorimeters and muon spectrometers. Regions of interest (RoIs) are identified by the L1 trigger, selecting regions which could contain high- p_T jets, leptons or photons.

The analysis described in Chapter 4 uses L1 calorimeter triggers, which comprise three types of L1 items: jet (L1JET), electromagnetic (L1EM), and τ (L1TAU). All use trigger towers of granularity $(\Delta\eta \times \Delta\phi) = (0.1 \times 0.1)$ made of calorimeter cell energy depos-

its along lines of constant η . The analysis uses L1 τ RoIs, which are formed as follows. For each possible 4×4 cluster of trigger towers, the electromagnetic energy in every possible 2×1 group from the centre 4 towers in the cluster is calculated and added to the hadronic energy in the central 2×2 region of the cluster. If this sum passes one of the τ energy thresholds, and is isolated from other deposits (checked by summing the EM energy deposited in the 12 outer trigger towers of the cluster, and requiring this to be less than an allowable energy threshold), an RoI is created.

Between Run 1 and Run 2 of the LHC, the L1 trigger system was upgraded to allow for topological decisions, providing scope for lowered thresholds given a preferred event topology. For example, another L1 trigger item used in the analysis described in Chapter 4 selects events where no more than 3 GeV is deposited in the ECal within a $\Delta R = 0.2$ cone of the most energetic L1 tau RoI and the second-most energetic, if there is more than one (see §3.3.1.1).

Anything selected by the L1 trigger is then passed onto the *High Level Trigger* (HLT), which reconstructs events in a much more detailed way, using a computing farm. Offline techniques allow for more accurate selections compared to the low granularity of the L1 trigger, which improves the triggering efficiency. Events from the HLT are selected at a maximum rate of 1.5 kHz, although this decreases over the course of a run as the luminosity decreases.

At each level of the trigger system, events are selected according to a list of trigger items and chains. These are defined separately at L1 and HLT, where each HLT chain is generally seeded by a specific L1 item (though a single L1 item can be shared between many HLT chains). Triggers may also be assigned a prescale, defining the fraction of events passing the criteria that are actually kept, to prevent individual triggers overloading or using too large a share of the processing resources. This would happen if low energy-thresholds are required to select low-energy processes which occur at a high rate. A prescaled trigger only runs on a pre-defined fraction of possible events—at L1 a prescale is applied after the trigger decision, while at the HLT only the fraction of events equal to the prescale are reconstructed, saving CPU time. However, applying a prescale can be detrimental to an analysis' sensitivity for very rare processes, so triggers for these processes are kept unprescaled where possible and the selection criteria are modified to keep the rate low enough to record. A set of L1 items,

HLT chains and their associated prescales comprise a trigger menu, and trigger menus are modified to take into account changes in the instantaneous luminosity.

ATLAS employs a number of triggers designed to select events with standard physics objects, such as leptons and jets. These triggers, although efficient for selecting the topologies they are designed for, are not always well-suited to selecting events with LLPs, which often lack the signatures of traditional decays. For example, muons from LLP decays do not have associated tracks in the ID so will not be identified by standard muon triggers, while jets from LLP decays may have relatively low energy without the usual pattern of energy deposition and arrive later than expected so would be missed by standard jet triggers. As a result, using standard triggers in an LLP analysis often restricts the search to topologies that have a prompt object in addition to the displaced object if the displaced decay is far from the interaction point, and it also limits the search to decays involving high-energy objects since the threshold for an unprescaled trigger must be sufficiently high to have an acceptable rate and using a prescaled trigger would lower the sensitivity of the search by the prescale factor. Therefore, dedicated triggers for LLP searches were developed to ensure that as many potential signal events as possible are recorded. In 2016, two sets of these triggers were running: so-called *CalRatio* triggers for selecting decays in the hadronic calorimeter, and the muon RoI cluster trigger for selecting decays in the muon spectrometer.^[53] The CalRatio triggers used in the analysis described in Chapter 4 are outlined in more detail below.

3.3.1 CalRatio triggers

Jets that result from neutral LLPs decaying near the outer radius of the ECal or within the HCal, like those sought in the analysis described in Chapter 4, are identified by two dedicated triggers. These make use of three main characteristics of these jets: they are narrow with a high ratio of energy deposited in the HCal compared to the ECal (the CalRatio) and typically have no tracks pointing towards the jet. These two triggers differ only in the L1 item used as a seed for the HLT selection.

The *high- E_T* trigger was designed for Run 1, but in Run 2 due to higher energy and pile-up the E_T threshold in the L1 seed was raised in order to keep the rate low enough, resulting in a decreased efficiency for some LLP models with lower mediator masses. An additional

trigger was therefore introduced part-way through the 2016 data-taking period, the *low- E_T trigger*, to provide an efficient selection for these models with a lower threshold.

3.3.1.1 L1 CalRatio triggers

The L1 seeds for both versions of the CalRatio triggers are based on L1TAU items, rather than L1JET items which are the usual seed type for triggering on standard jets. An L1TAU item sums energy in a narrower region than an L1JET item: 0.2×0.2 in $\eta \times \phi$ versus the 0.8×0.8 square the L1JET triggers use. Since displaced jets are expected to be reconstructed as narrow, this choice of L1 seed provides a higher efficiency. In the case of the L1 seed used in the high- E_T trigger (L1TAU60), the E_T threshold is 60 GeV, above which the event progresses to the HLT. This high threshold is necessary to keep the trigger rate low enough so that it can remain unprescaled. The trigger seeded by L1TAU60 was active for the whole 2016 data-taking period, collecting 33.0 fb^{-1} of data.

In contrast, LLP-NOMATCH is an L1 topological trigger item which can make use of combined information from different regions of both calorimeters, and is used as the seed for the low- E_T trigger. The basis is an L1TAU30 item (with an E_T threshold of 30 GeV), followed by a requirement that there are no energy deposits in the ECal with $E_T > 3 \text{ GeV}$ within a 0.2 cone around the most energetic high- E_T energy deposit (and the second-most energetic deposit if there is more than one). Using the information from the HCal and the ECal separately at L1 allows a low E_T threshold to be kept while still maintaining a low enough rate to avoid any prescaling. The low- E_T trigger, seeded by LLP-NOMATCH, collected only 10.8 fb^{-1} of data since it was introduced later in the 2016 data-taking period, but its advantage for signal models with a lower mediator mass remains despite this.

3.3.1.2 HLT CalRatio triggers

The HLT chain for the CalRatio triggers is the same for both L1 seeds. Partial-scan (rather than full-scan) jet reconstruction is used, considering only sections of the detector centered around the regions of interest (RoIs) identified by the L1 trigger. These individual RoIs are combined into a super-RoI in which jet reconstruction is performed with the anti- k_t algorithm, described in more detail in §3.4.3.

Due to the unusual nature of the jets studied, the standard jet cleaning requirements used in most HLT jet triggers are not applicable here. The standard cleaning is designed to reject

signal from noisy cells and poorly reconstructed jets, so it removes jets with a large part of their energy in a single calorimeter cell or sampling layer, which are also characteristics of jets from LLP decays. Applying the standard cleaning would result in a significant decrease in the signal efficiency. Instead, an LLP-specific jet cleaning algorithm is used, which is described in more detail in §3.4.3.3. At least one jet passing this dedicated cleaning algorithm is required to satisfy $E_T > 30$ GeV and $|\eta| < 2.5$. To identify jets with a high fraction of their energy in the hadronic calorimeter, triggering jets are also required to pass $\log_{10}(E_H/E_{EM}) > 1.2$. Jets satisfying these requirements are used to determine 0.8×0.8 square regions in $\eta \times \phi$ centred on the jet axis in which to perform track reconstruction. After this, triggering jets are required to have no tracks with $p_T > 2$ GeV within $\Delta R < 0.2$ of the jet axis.

Finally, jets satisfying all of the above criteria are required to pass a beam-induced background (BIB)[†] removal algorithm. This algorithm relies on the fact that calorimeter hits resulting from BIB would be aligned in a relatively horizontal line parallel to the beam pipe and would have hit times inconsistent with the time taken for a particle to travel at the speed of light from the IP to each hit location. It rejects the event if at least four cells at the same ϕ and in the same calorimeter layer (but separated in ΔR) from the triggering jet have timing consistent with that of a BIB muon from the current bunch crossing.

Additional CalRatio triggers are used for the study of non-collision background in the analysis described in Chapter 4. The first of these is only active during empty bunch crossings and is seeded by an L1TAU30 item with the same HLT requirements as the main triggers; this is used to study cosmic-ray background events. Another two background triggers run alongside the standard CalRatio triggers and are identical to them except that the BIB removal algorithm is not applied. Selecting collision events which pass one of these background triggers but not the corresponding standard CalRatio trigger provides a dataset which can be used to study BIB in collisions.

[†]BIB muons may be produced when a proton bunch hits the collimators on the $\pm z$ side of the IP. As the muons travel parallel to the beam-line through the calorimeter they undergo bremsstrahlung, and may leave energy deposits in the calorimeter. Particles originating from the collision at the IP and subsequently depositing energy in the calorimeter take a different, longer, path. As the distances traveled by each are different, the time at which the BIB energy deposit is made in the calorimeter relative to the collision time ($t = 0$) will be negative.

3.4 Event reconstruction

As mentioned in Chapter 1 and shown in Figure 1.1, particles passing through the ATLAS detector interact with its subsystems in several ways. In order to identify and describe the particles producing signals in the detector, details of the various interactions with the detector need to be combined. There are many algorithms and combinations of detector elements utilised to achieve this, since the high multiplicity of particles in a collision event complicates matters.

3.4.1 Track and vertex reconstruction

Individual interactions of charged particles with the layers of the ID are used to reconstruct their tracks, which are also extrapolated beyond the material layers to the beam pipe and combined to locate the hard scatter interaction (primary vertex). Energy deposits recorded in each detector element are reconstructed as hits, corresponding to a 3D space-point where the charged particle traversed the sensor.^[54] Track seeds are formed from three hits in the IBL or pixel layers or the SCT which are compatible with a track. By adding hits consistent with the trajectory of the initial track seed, the track is propagated out from the seed towards the TRT.

3.4.1.1 Primary vertex reconstruction

Due to the high instantaneous luminosity of the LHC, multiple pp collisions can occur in any filled bunch crossing. In 2016, the mean number of interactions per bunch crossing was 24.9, as was shown in Figure 3.3, and the number of reconstructed vertices as a function of the number of interactions per bunch crossing from the same period is shown in Figure 3.9. Interaction vertices are reconstructed in two stages: vertex finding and vertex fitting.^[55] The former stage involves the association of reconstructed tracks to vertex candidates, and the fitting stage deals with reconstruction of the actual vertex position.

The input to the vertex reconstruction algorithm consists of tracks satisfying the following criteria:

- $p_T > 400 \text{ MeV}$; $|d_0| < 4 \text{ mm}$; $\sigma(d_0) < 5 \text{ mm}$; $\sigma(z_0) < 10 \text{ mm}$;
- at least four hits in the SCT detector and nine SCT/pixel hits;
- no expected but missing pixel hits.




Figure 3.9: Distribution of the average number of reconstructed vertices as a function of the number of interactions per bunch crossing, μ , for the 2016 ATLAS pp dataset.^[56]

The symbols d_0 and z_0 denote the transverse and longitudinal impact parameters of tracks with respect to the centre of the luminous region (*beam spot*), and $\sigma(d_0)$ and $\sigma(z_0)$ are their uncertainties. The convention is to calculate these geometry-dependent parameters of tracks at the track *perigee*—the point of closest approach to the beam line. The impact parameter requirements reduce contamination from tracks originating from secondary interactions, and the above requirements are tighter than the standard ATLAS track selection criteria in order to maintain a low rate of tracks mistakenly reconstructed from a random combination of hits (*fakes*). With this set of tracks defined, the seed position of the first vertex is chosen. The x - and y -coordinates of this point are taken from the geometric centre of the beam spot, and the z -coordinate is calculated as the mode of the z -coordinates of tracks at their respective points of closest approach to the centre of the beam spot. After the seed position has been determined, an iterative primary vertex finding procedure begins. An adaptive vertex fitting algorithm with an annealing procedure is used. Using the seed position as a starting point and the parameters of reconstructed tracks as input, the algorithm performs an iterative χ^2 minimisation, finding the optimal vertex position. Each input track is assigned a weight representing its compatibility with the vertex, the vertex position is recalculated using the weighted tracks, and then the procedure is repeated with the track weights recalculated with respect to the new vertex position. Following the last iteration, the final weight of each track used in the vertex fit is calculated. Tracks found to be incompatible with the vertex by more than $\sim 7\sigma$ are removed from the vertex candidate and returned to the pool of unused tracks, and all the rejected tracks are used as input for

finding the next vertex. This procedure is repeated until no unassociated tracks are left in the event or no additional vertex can be formed from the remaining set of tracks. All vertices with at least two associated tracks are retained as valid primary vertex candidates, described as a set of three-dimensional vertex positions and their covariance matrices.

For most hard-scatter physics processes, it is effective to identify the hard-scatter primary vertex as the primary vertex with the highest sum of the squared transverse momenta of contributing tracks, $\sum p_T^2$. This is based on the assumption that charged particles produced in a hard-scatter interaction have on average a harder transverse momentum spectrum than those produced in a pile-up collision. While this choice is valid for most processes, there are a few for which different primary vertex selection criteria must be used—for example, when studying $H \rightarrow \gamma\gamma$ or processes with τ leptons in the final state.

3.4.1.2 Large radius tracking

For physics analyses involving displaced objects, the reconstruction of displaced tracks with a large impact parameter may be required, and this follows a two-step procedure.^[57] In the first iteration, the default track identification algorithm is applied, which uses hits in the pixel system, SCT, and TRT to reconstruct tracks with a small impact parameter. The hits not associated to a track during the first pass are used in a second run of the track finder, with loose requirements on the transverse and longitudinal impact parameters (d_0 and z_0) and the number of silicon hits that are shared (or not shared) with another track. This two-step procedure is referred to as the large radius tracking (LRT) algorithm. Applying the LRT procedure is CPU-intensive, and thus it is only run once per data-processing campaign, on a subset of specially-requested events.

3.4.2 Electrons, photons and muons

Reconstruction of electrons and photons is performed using a combination of the ID and ECal, since they are both usually stopped by the ECal. Muons, on the other hand, deposit only a small fraction of their energy in the calorimeters and are reconstructed primarily using the MS.

The reconstruction of electrons begins from seed clusters in the middle ECal layer. For each seed, if a suitable track is found near the cluster, a sliding window algorithm is used

to find the $\eta \times \phi$ region that encloses the most energy. Each electron's ID track is chosen as the track closest to the cluster's centre. All candidate electrons are subject to further cuts on the shower shape and leakage into the HCal to reject backgrounds from misidentified hadrons, electrons from photon conversions and non-isolated electrons from heavy flavour decays.^[58]

Photon reconstruction is intended to identify both converted and unconverted photons. Converted photons undergo pair production within the detector material, while unconverted photons undergo no material interactions prior to reaching the ECal. The reconstruction is seeded in the same way as for electrons, but seeds are selected if they lack associated tracks or possess a track consistent with a photon conversion.^[59] Requirements are placed on the shower shape and on leakage into the HCal to reject backgrounds—photons reaching the calorimeter are not bent by the magnetic field so will usually create very narrow showers, and these showers do not reach the HCal. Converted photons are reconstructed from the electron showers and tracks left by their decay products. Additional selections are applied to unconverted photons to discriminate between single photon showers and pairs of photons that result from neutral meson decays.

Muons can be reconstructed using only the MS (*standalone muons*) or by combining MS and ID tracks (*combined muons*).^[52] ID tracks with corresponding MS segments, but not complete tracks, are a third type of reconstructed muon called *tagged muons*. In all cases, MDT hits are combined within single MDT chambers into segments, approximated as a straight line due to the short depth of the chambers. These segments may incorporate trigger coordinate information from the RPCs or TGCs, if available, to improve their position resolution. As with ID tracks, the track parameters of muon tracks are calculated at the perigee of the extrapolated track.

The MS track reconstruction requires at least two muon segments pointing towards the interaction point which result in a good fit to a possible track. These tracks are extrapolated towards the beam pipe. If an extrapolated standalone muon track lies near an ID track, they are combined. If the χ^2 for the combined fit is low enough, the muon is a combined muon, and its track parameters are a weighted average of the MS and ID tracks. Low- p_T muons may not produce enough hits in the MS to directly reconstruct a MS track for the muon, but will still leave ID hits. In this case, if an ID track extrapolated to the MS is associated

with one or more muon segments, the track is denoted a tagged muon, with parameters identical to those of the ID track.

3.4.3 Jets

Due to the short range of the strong force and colour confinement, it is not possible to observe quarks and gluons directly. Partons produced in collisions hadronise, creating a collimated shower of particles with a net momentum equal to that of the initiating quark or gluon. Some of these particles (like the π^0) decay more or less instantaneously while others (like the K^0 and $\pi^\pm$) are more likely to reach the calorimeter without decaying. This shower of particles is called a *jet* and is an extremely common object in LHC collisions. The majority of the energy in the average jet is carried by photons and charged pions, while neutral hadrons carry less than 20%. Jets are not only produced in the initial pp collision; they may also result from hadronic decays of particles such as the Z or τ (see §3.4.4).

A jet may have one or more charged tracks in the ID corresponding to the passage of particles forming the jet. It may also have muon tracks or segments associated with it if a muon was produced as part of the jet or if the hadronic calorimeter fails to fully absorb the energy of the shower and some particles punch through. Ideally, the reconstructed energy of a jet will be well-calibrated and include the full shower.

Jets resulting from the decay of a neutral LLP will have properties that depend on how displaced from the primary vertex their decay is, such as the spatial extent and distribution of their energy deposits in the calorimeter, or the presence of charged tracks in the inner detector pointing towards those deposits. The accuracy and precision of the jet reconstruction determines how well these jets can be separated from showers of particles originating at the primary vertex.

3.4.3.1 Clustering and reconstruction

Any jet reconstruction algorithm should be both infrared and collinear safe. An infrared-safe algorithm will reconstruct the same jet in the presence of soft particles, which are numerous from ISR and the underlying event and should not affect the jet shape or have a significant impact on its final energy. On the other hand, two collinear, hard particles

are likely to be products of the same shower and should not be reconstructed as separate jets—collinear-safe algorithms correctly group hard objects into the same jet.

In ATLAS, the energy collected by the calorimeters is used to reconstruct the jet and determine its energy and position measurements. Incoming particles interact with the calorimetry and deposit their energy across many cells both parallel to and perpendicular to their momentum. Instead of tracks, roughly cone-shaped collections of energy deposits in the calorimeter cells are reconstructed and given as input to clustering algorithms. Three-dimensional topological cell clusters (*topoclusters*) are built using a nearest-neighbour algorithm that groups together calorimeter cells with energy significance above specific thresholds.

Many algorithms exist for combining clusters into jets. The principal strategy of any clustering algorithm is to determine if the separation of two clusters is large enough, with respect to their energy, that they are likely to be the result of different showers rather than part of the same jet. The algorithm loops over clusters i and j and compares the distances between each pair of clusters, d_{ij} , and the distance between each cluster i and the beamline, d_{iB} to identify the parameter with the minimum value. If the smallest is d_{ij} clusters i and j are merged, while if it is d_{iB} the i is identified as belonging to a separate jet and removed from the set of clusters.

Individual jet algorithms differ in how d_{ij} and d_{iB} are defined. In general,

$$d_{ij} = \min(k_{T,i}^{2p}, k_{T,j}^{2p}) \frac{(y_i - y_j)^2 + (\phi_i - \phi_j)^2}{R^2}, \quad (3.8)$$

and,

$$d_{iB} = k_{T,i}^{2p}, \quad (3.9)$$

where $k_{T,i}$ is the transverse momentum of i , y_i is its rapidity, and ϕ_i is its azimuthal angle. The distance parameter, R , has several common values. Larger R values will make the algorithm less likely to underestimate the energy of wider jets, but also more prone to overestimating energies due to incorporating soft jets not originating from the hard scatter. The case $p = 1$ corresponds to the k_t jet reconstruction algorithm while $p = -1$ corresponds to the anti- k_t algorithm.^[60] The Cambridge/Aachen algorithm uses $p = 0$.^[61,62] Each final jet's four-vector is then the sum of the four-vectors of its constituent topoclusters.

The analysis described in Chapter 4 uses anti- k_t ($p = -1$) jet reconstruction with $R = 0.4$. In this case, d_{ij} is large either when the clusters are both soft, or they are highly spatially separated. This ensures that very low- p_T topoclusters are combined with high- p_T topoclusters before low- p_T clusters are combined with each other, and therefore this algorithm produces very circular jets compared to other combination algorithms.

3.4.3.2 Jet energy calibration

Once the jets have been reconstructed, their energy is calibrated to accurately match the energy of the original decaying particle.^[63] In general, a shower of pure hadrons will be reconstructed with a lower energy compared to a pure electromagnetic shower with the same true energy. This difference in response is in part due to energy lost via escaping muons and material interactions. A single jet energy scale (JES) is used to compensate for energy not measured by the calorimeters due to these effects and others, and ensures that jets have a uniform response across all regions of the calorimeter.

3.4.3.3 Jet quality selections

In addition to the jet energy scale calibration, a set of cleaning criteria are applied to make sure that jets originating from true hadronic activity from the hard scatter are well-measured. Most analyses use a definition of good jets defined by a dedicated data quality group,^[64] but some searches for exotic events may require modification of the standard criteria if they involve jets that would otherwise be removed.

The common jet parameters for determining the quality of a jet include:

- EMF: the fraction of the total jet energy deposited in the ECal;
- FMax: the largest fraction of energy deposited in a single layer;
- HECf: the fraction of energy deposited in the hadronic endcap;
- LArQ: the fraction of LAr cells with a Q-factor greater than 4000 (the Q-factor is a measure of the difference in shape between the measured pulse and the predicted pulse used to reconstruct the cell energy, where values near zero indicate similarity of the shapes);
- HECQ: similar to LArQ, but calculated only in the hadronic endcap;
- negE: the sum of the energy of cells with negative energy (this may arise from noise and noise suppression);

- timing: the energy-weighted mean of the cell timing, where each cell's timing is calculated with respect to the time of flight of a particle travelling at the speed of light between the IP and the cell;
- η : the reconstructed jet's pseudorapidity.

LAr noise can result in a noisy cell being reconstructed as a jet or poor reconstruction of a jet from the collision. Jets associated with noise spikes generally have large negative energy or poor signal quality in the readout from the calorimeter. Jets from non-collision background (such as cosmic rays or BIB muons) more frequently have larger and more negative timing than a typical collision jet and often contain fewer high energy cells than a collision jet.

For SM jets, showers of purely neutral hadrons are unusual enough that removing low-EMF jets will not remove too many collision jets. The search described in Chapter 4, however, explicitly sought low-EMF jets as a signal, and risked losing a substantial number of desired events if these criteria remained. As a result, a dedicated cleaning for events with these jets was developed without the low-EMF criteria, and other selections were implemented to remove events resulting from non-collision backgrounds.

In this dedicated cleaning algorithm, reconstructed jets are required to pass the `cleanLLP` flag, which makes use of a number of the parameters described above. The first term of the `cleanLLP` flag contains the HEC spike-cleaning part of the standard jet cleaning requirements, vetoing on

$$(\text{HECf} > 0.5) \cap (|\text{HECQ}| > 0.5) \cap (|\text{LArQ}|/65535 > 0.8), \quad (3.10)$$

which is very efficient for noise burst rejection in the LAr calorimeter. A second term vetoing on

$$(\text{negE} < 10 \text{ GeV}) \cap (\text{FMax} > 0.85) \quad (3.11)$$

is applied to further eliminate LAr noise, and is also useful for non-collision background removal. In the search for displaced jets described in Chapter 4, if the jet is vetoed by either of these two terms, it is rejected and not considered later.

3.4.4 Taus

With a mass of 1.78 GeV and proper decay length ($c\tau$) of 87 μm , τ leptons decay either leptonically ($\tau \rightarrow \ell \nu_\ell \nu_\tau$, $\ell = e, \mu$) or hadronically ($\tau \rightarrow \text{hadrons } \nu_\tau$). They typically decay inside the beam pipe before reaching the detector, and can thus only be identified by their decay products. Since the leptonic decay products of the τ cannot be distinguished from prompt electrons or muons, only the reconstruction of hadronic decays (which represent 65% of all possible decay modes) is discussed here.

In hadronic τ decays, one or three charged pions are produced in 77% and 18% of all cases, respectively. Charged kaons are present in the majority of the remaining hadronic decays. In 61% of all hadronic decays, one or more associated neutral pions are also produced.^[10] The neutral and charged hadrons stemming from the τ decay make up the visible decay products, which appear in the ATLAS detector as a narrow jet in the calorimeters, associated to one or three tracks in the ID.^[65,66] The main background to hadronic τ decays is from jets produced via the fragmentation of quarks and gluons, and electrons and muons which can mimic the signature of τ decays with one charged hadron.

The τ reconstruction starts with a jet obtained using the anti- k_t algorithm with distance parameter $R = 0.4$. These jets are required to have $p_T > 10 \text{ GeV}$ and $|\eta| < 2.5$. The vertex normally chosen as the primary vertex (the one with the highest $\sum p_T^2$, see §3.4.1.1) does not always correspond to the one at which the τ was produced. Instead, the vertex associated with the τ is identified among the previously reconstructed primary vertex candidates in the event by summing the p_T of the tracks in a cone with $\Delta R < 0.2$ around the jet direction, and choosing the vertex to which the largest fraction of the p_T sum is matched.

Since many jets formed by quarks or gluons can fake the signature of τ leptons, a multivariate analysis algorithm based on Boosted Decision Trees (BDTs) is used to distinguish between τ -jets and jets from other particles. It is trained separately for 1-track and 3-track decays, and the input variables are chosen mostly among shower shape variables that use the narrowness of the τ -jet to distinguish it from other jets. Three working points corresponding to different τ identification efficiencies are defined, with values between 45% and 60%.

3.4.5 Composite variables

In addition to describing individual objects, it is often useful to be able to describe the overall properties of an event. For example, events with many jets could be described using the total transverse energy deposited in the calorimeters, but this would not give any indication of the amount of energy escaping the detector or the imbalance of the jet energy, so additional variables are computed.

3.4.5.1 Missing energy

From conservation of momentum, the vector sum of the momentum in the transverse plane of the detector should sum to zero. The imbalance of momentum in this plane is called missing transverse momentum, and can be the result of neutrinos (which escape the detector carrying away momentum), detector acceptance effects, mis-measured objects, or unreconstructed objects. It can also include the transverse momentum of reconstructed muons due to their tendency to escape the calorimeters. Another non-negligible source of missing transverse momentum are non-collision backgrounds. For example, BIB muons can leave energy deposits in the calorimeters, and cosmic ray muons or showers may also do the same. These energy deposits will rarely be symmetric around the collision. Other possible sources arise in BSM physics models, such as those involving escaping long-lived particles.

The missing transverse momentum is first calculated in the x (y) direction as,

$$E_{x(y)}^{\text{miss}} = E_{x(y)}^{\text{miss}, e} + E_{x(y)}^{\text{miss}, \gamma} + E_{x(y)}^{\text{miss}, \tau} + E_{x(y)}^{\text{miss}, \text{jets}} + E_{x(y)}^{\text{miss}, \mu} + E_{x(y)}^{\text{miss}, \text{soft}}, \quad (3.12)$$

where $E_{x(y)}^{\text{miss}, \text{object}}$ is the x (y) component of the negative vector sum of the momentum of all reconstructed objects, and $E_{x(y)}^{\text{miss}, \text{soft}}$ is the vector sum of all remaining detector objects not passing the selection of the main physics objects. These objects could be low- p_T tracks in the inner detector, or calorimeter deposits not associated to hard objects. The total missing transverse momentum is then given by,

$$E_T^{\text{miss}} = \sqrt{(E_x^{\text{miss}})^2 + (E_y^{\text{miss}})^2}. \quad (3.13)$$

3.4.5.2 Total jet energy

The total jet energy in the event is usually described with H_T , the scalar sum of all transverse jet energies, defined as,

$$H_T = \sum_{\text{all jets}} |\vec{E}_T|. \quad (3.14)$$

The vector sum can also be determined, and should equal zero unless there is a substantial component of missing jet energy, denoted H_T^{miss} . Then,

$$\vec{H}_T^{\text{miss}} = - \sum_{\text{all jets}} \vec{E}_T. \quad (3.15)$$

Combining the total jet energy and the missing jet energy defines the effective mass of an event,

$$m_{\text{eff}} = H_T + |\vec{H}_T^{\text{miss}}|. \quad (3.16)$$

3.5 Simulation

To understand the ways in which particles will interact with the detector and the signatures they will leave, it is necessary to generate simulated data. Monte Carlo (MC) simulation is an essential part of most analyses in ATLAS. It is used for comparison of the data to predictions from simulation and to estimate SM backgrounds when data-driven techniques are not available. In searches for new physics, it is required to simulate the signal process and informs the analysis selections to improve the sensitivity. Additionally, generator-level particle information (*truth information*) can be key to understanding the composition and behaviour of signal and background in an analysis.

Event simulation can be rather complicated at hadron colliders due to the large number of processes occurring in not only the pp collision but also the underlying event, as discussed in §2.2. This is shown schematically in Figure 3.10. Additionally, the running of the strong coupling constant, described in §2.1.4, leads to markedly different behaviour of QCD at different distance or energy scales. At short distances (i.e. high energies), asymptotic freedom means QCD is weakly interacting. Therefore, the calculations of the hard scatter can be performed using perturbation theory. However, at larger distances (lower energies) soft processes like hadronisation and the formation of the underlying event are




Figure 3.10: Representation of a $t\bar{t}h$ event as produced by an event generator. Image adapted from Ref. [67].

non-perturbative and must be computed from QCD-inspired models rather than directly from QCD, and this is referred to as the *parton shower*. There are many MC tools available for the generation of a wide range of processes, each using different methods for the modelling. The calculation of the hard scatter process and the parton shower are usually performed separately, followed by the detector simulation.

Stable particles that travel more than a few mm after hadronisation and decays of unstable particles are input to the ATLAS detector simulation.^[68] This simulation is performed using either *full* GEANT4^[69] simulation or a *fast* simulation in which calorimeter showers are simulated with a parametrised description^[70] whilst all other interactions are simulated with GEANT4.

Following the simulation of the hard scattering process, additional softer collisions are added to simulate pile-up. Reconstruction of the events after digitisation of the hits in the detector uses the same software as that applied to data. The MC simulation is then weighted to match the μ distribution in the data.

3.6 Long-lived particle signatures

The experimental signatures of SM processes, or BSM processes where decays occur promptly back to SM particles, tend to be composed of combinations of the objects described in §3.4, for which the ATLAS detector and standard reconstruction software were designed. However, the experimental signatures of BSM LLPs produced in collisions, such as the signature considered in this thesis, can be varied and unusual. They may include:

- tracks with unusual ionisation properties;
- localised deposits of energy inside calorimeters without associated tracks;
- decays which are out-of-time with the collision;
- displaced vertices;
- disappearing, appearing, or kinked tracks.

Many such LLP signatures would be vetoed during standard data cleaning or completely missed by searches for prompt activity, so dedicated LLP searches are necessary to ensure that possible avenues for the discovery of new physics are not overlooked. These searches can be challenging and time-consuming, but are also often not dominated by irreducible SM background processes because SM LLPs (such as B-mesons and K-short and K-long mesons) have masses $\lesssim 5$ GeV and have well-understood experimental signatures. Unusual BSM LLP signatures therefore offer excellent prospects for the discovery of new physics at particle colliders.

Although small compared to the large number of searches for prompt decays of new particles, many searches for LLPs at the ATLAS, CMS, and LHCb experiments have already been performed, including the ATLAS search described in Chapter 4. However, in cases of ultra-low-mass particles, ultra-long lifetimes, or unusual LLP charges, it is hard or impossible to trigger on and/or reconstruct such events in the ATLAS, CMS, and LHCb detectors. This has led to new proposals for dedicated experiments to look for LLPs in regimes that are otherwise inaccessible at the LHC. Together, these experiments provide the best sensitivity to new fractionally-charged LLPs, magnetic monopoles, and other LLPs arising from models such as those containing Higgs-portal hidden sectors, dark photons, and Majorana neutrinos.

These dedicated experiments[†] include the milliQan milli-charged particle experiment,^[71] the MoEDAL monopole search,^[72] the MATHUSLA surface detector for ultra-LLPs,^[73] the SHiP hidden particle search,^[74] the CODEX-b proposal for a new detector near LHCb,^[75] and the FASER proposal^[76] for a long, narrow detector located in the forward direction 480 m downstream of the ATLAS collision point. Each of these dedicated experiments is sensitive to different LLP lifetimes, masses, and production modes depending on their position and orientation.

[†]Of these, only the FASER and MoEDAL experiments have been approved at the time of writing.

Chapter 4

Search for events with displaced jets in the ATLAS calorimeter

This chapter presents a search for neutral long-lived particles (LLPs) in the calorimeters of the ATLAS detector, using data collected during Run 2 of the LHC.

An overview of the search strategy and how it fits in with the landscape of previous neutral LLP searches is provided in §4.1. The datasets and simulated samples used in the analysis are detailed in §4.2. In §4.3, the full series of analysis selections is presented, including descriptions of the machine-learning algorithms used to separate signal events from background. A description of the data-driven background estimation performed in this analysis is contained in §4.4, in addition to the results obtained. The statistical interpretation of the results and the extrapolation as a function of LLP proper decay length are presented in §4.6, followed by a discussion of the results and the future of the search in §4.7.

4.1 Analysis overview

The search described in this chapter targets pair-produced neutral LLPs decaying to hadronic final states within the ATLAS calorimeters, resulting in displaced jets with unusual features compared to jets from SM processes. A jet from an SM process is a tightly collimated group of charged and neutral particles originating from the IP. The charged particles leave tracks in the ID, and both the charged and neutral particles deposit energy in the

ECal and HCal. A jet from a sufficiently displaced LLP decay will contain charged and neutral particles originating from a point late in the ECal or in the HCal rather than the IP, and therefore leave no tracks in the ID. This also results in a high ratio of energy deposited in the HCal to energy deposited in the ECal (E_H/E_{EM} , called the *CalRatio*, so the jets are called *CalRatio jets*). Further, CalRatio jets are narrower on average than their SM counterparts since the particles have less time to become separated in the detector volume before depositing energy in the calorimeters. The decay products from each LLP decay are collimated enough that they are reconstructed as a single jet, so the resulting event signature is two narrow, trackless jets with little or no energy deposited in the ECal. Requiring two displaced jets also significantly reduces the expected background.

4.1.1 Benchmark signal model

In the simplified hidden sector (HS) model considered in this search, detailed further in §2.4.2.1, a heavy mediator boson, Φ , with mass, m_Φ , between 125 and 1000 GeV is produced in a pp collision and decays to two long-lived scalar bosons, S , with mass, m_S , between 5 and 400 GeV (depending on the mediator mass)[†]. The simulated signal samples used in the analysis are described in §4.2.2.

4.1.2 Previous searches for neutral LLPs

Several previous searches have considered the same benchmark HS model as the analysis described in this chapter. In Run 1 (at $\sqrt{s} = 8$ TeV), searches were performed for displaced jets in the ATLAS ID, HCal and MS in 20.3 fb^{-1} of data.^[77,78] The ID and MS searches were combined in one analysis, where either a displaced muon or single-jet plus E_T^{miss} trigger was used to search for two displaced vertices in the ID, two displaced vertices in the MS or one displaced vertex in each of the ID and MS subdetectors.^[77] Models with $m_\Phi = 100\text{--}900$ GeV and $m_S = 10\text{--}150$ GeV were explored, for which there was good sensitivity for proper decay length, $c\tau = 0.01\text{--}100$ m. The search in the HCal was somewhat similar to the Run 2 analysis described in this chapter, but was a simpler cut-based analysis.^[78]

[†]The range of masses chosen was partly motivated by the coverage resulting from other searches, such as dedicated SUSY searches with the ATLAS experiment, for LLP masses above $\mathcal{O}(100 \text{ GeV})$. The excellent sensitivity of these searches, together with the lack of a definitive signal in any prompt channels at the LHC, have focused attention on low-mass LLP signatures as an interesting avenue for further study.

It used a CalRatio trigger to select events with displaced jets originating in the HCal, and explored the same models as the above analysis with good sensitivity for $c\tau = 0.2 - 20$ m. The higher level of pile-up in Run 2 meant that a cut-based approach would have resulted in a much lower sensitivity, since the displaced jets would have a lower CalRatio due to there being more pileup deposits in the ECal, and so machine-learning algorithms were used in the Run 2 version of the analysis.

In Run 2, another search for displaced vertices in the MS was performed in 36.0 fb^{-1} of data, considering the same HS model as the analysis in this chapter (in addition to a Stealth SUSY model and a baryogenesis model).^[79] Two topologies were considered: two displaced vertices in the MS, or one displaced vertex in the MS and E_T^{miss} . The addition of this latter topology allowed the analysis to probe longer decay lengths for cases where one LLP escaped the detector. For models with $m_\Phi = 100 - 1000 \text{ GeV}$ and $m_S = 5 - 400 \text{ GeV}$, there was a good sensitivity in the range $c\tau = 0.04 - 100$ m. Following the completion of the CalRatio analysis described in this chapter, a statistical combination of its results with those from the MS search was performed, and is the subject of Chapter 5. Other ATLAS searches have considered models different to the benchmark HS model used here, but still resulting in neutral LLPs decaying hadronically in the ATLAS detector. These include: a search for displaced collimated leptons or light hadrons, with sensitivity in the range $c\tau = 0.01 - 3$ m for $m_{\text{LLP}} = 400 \text{ MeV}$;^[80] a search for a long-lived dark photon produced in association with a Z boson using lepton triggers, sensitive to $c\tau = 0.1 - 100$ m for $m_{\text{LLP}} = 5 - 200 \text{ GeV}$;^[81] and a search for decays of a Higgs boson via two long-lived bosons which produce displaced b -jets ($H \rightarrow aa \rightarrow 4b$) using b -triggers to target small displacements, with sensitivity to $c\tau = 0.1 - 3$ mm for $m_{\text{LLP}} = 20 - 60 \text{ GeV}$.^[82]

4.1.3 Background contributions

Three contributions to the background were considered: SM multijet events, and two sources of non-collision background. Of these, the SM multijet background had the largest contribution to the final background estimate. These background contributions were estimated with the methods described in §4.4, and their characteristics are described in further detail below.




Figure 4.1: A typical BIB event, showing a line of calorimeter energy deposits (yellow) parallel to the beam. The purple segments are muon segments and the red lines represent tracks from the collision.^[53]

SM multijet background. It is possible for an SM jet to contain only neutral hadrons, thus mimicking a signal jet. Even though the probability that this will occur during a collision is low, the multijet cross section is sufficiently large that the rate is high enough for this to be the dominant background in this search. However, also as a result of this low probability, it is difficult to simulate enough SM multijet events to allow the background to be properly studied, so a data-driven method was used to estimate the contribution from this background in the search region.

Beam-induced background. As mentioned previously in §3.3.1.2, another important source of background was from so-called beam-induced background (BIB) muons. Protons in the LHC may scatter inelastically with residual gases in the vacuum chambers or with components such as the collimators, resulting in the production of pions and thus muons. The circulating proton bunches are therefore accompanied by a halo of muons, which can have energies above a few GeV and even up to 1 TeV. These BIB muons travel parallel to the beam pipe and roughly in time with proton bunches. At the calorimeter level, the BIB muons are preferentially, although not entirely, at $\phi \approx 0$ and $\phi \approx \pi$. When one of these muons passes through the calorimeter, if it has a high enough energy (above a few tens of GeV) it can emit bremsstrahlung radiation and leave behind energy deposits which may be reconstructed as jets. If this occurs in the HCal, then the reconstructed jet will naturally lack ECal energy deposits and ID tracks, so would be reconstructed as track-less jets with a high CalRatio. However, a BIB muon at a calorimeter radius and in time with a bunch crossing will leave energy deposits in cells out of time with what is expected




Figure 4.2: The ATLAS detector in the experimental cavern. Above the cavern are the two access shafts used for the detector installation.^[83]

from a particle traveling from the IP to the cell. A row of out-of-time cells at the same ϕ is therefore suggestive of a BIB event, an example of which is shown in Figure 4.1, and this was exploited in the trigger-level BIB removal algorithm used in this analysis described in §3.3.1.2.

Cosmic ray muon background. Cosmic ray muons reaching ATLAS tend to travel in the downward direction, primarily arriving via the two access shafts (see Figure 4.2). The muons are minimally-ionizing particles which traverse the detector, and may bypass the ECal and ID while passing through the HCal volume and undergoing bremsstrahlung emission there. The reconstruction software will reconstruct jets from the energy deposits they leave as if they were due to particles originating from the IP. These jets are therefore a background to the CalRatio jet search. They are created independently of any proton-proton collision in the detector so they arrive at a constant rate. This background was not explicitly removed during the analysis selections, since it was not expected to dominate. Rather, an estimate of the cosmic contribution to the final search region was determined by applying the analysis selections to a cosmic-specific dataset and then scaling to the main dataset, and was found to be negligible.

4.1.4 Search strategy

The search strategy for the detection of neutral LLPs is very dependent on where in ATLAS the LLP decays. If the LLP decays in the ID or MS, then reconstruction of a displaced vertex is possible. On the other hand, LLPs decaying within the calorimeters will produce displaced jets. The search described here focused on events in which both LLPs decayed either in the HCal or near the outer edge of the ECal. The analysis strategy is outlined schematically in Figure 4.3.

The data for this analysis were collected using signature-driven triggers. Although the jets from neutral LLP decays would pass various standard jet triggers used by ATLAS, the high prescales of those triggers motivated the use of dedicated triggers. Two triggers were developed: the high- E_T CalRatio trigger and the low- E_T CalRatio trigger. These both exploited the jets' expected lack of ID tracks, high CalRatio, and narrowness. The triggers differed in their L1 seed, and therefore E_T thresholds, so they were efficient for different signal models. The high- E_T trigger had a threshold of 60 GeV so was efficient for models with $m_\Phi > 200$ GeV (*high- m_Φ models*), up to 90%, while the low- E_T trigger had an E_T threshold of 30 GeV and was more efficient than the high- E_T trigger for models with $m_\Phi \leq 200$ GeV (referred to as *low- m_Φ models*), up to 50%. Further information on the specific requirements in these triggers was provided in §3.3.1, and their differences in efficiency are explored further in §4.2.3. The subsequent analysis stages were split into two paths, each beginning with a different trigger.

Offline, the selections were designed to first identify signal-like *jets* and subsequently identify signal-like *events*. Signal jets were distinguished from those likely to have been produced by a background process by making use of the differences in the properties of the jets. Since there were many properties which could be used to distinguish signal from background, machine learning techniques were used to classify jets in two steps, and then to classify events in a third step.

First, a multi-layer perceptron (MLP) was used to predict the decay position of the particle that generated each reconstructed jet, given only the calorimeter layer energy distributions as input. This is presented in §4.3.1.2. The second step, a per-jet boosted decision tree (BDT), used the result of the MLP alongside other jet shape variables to score each jet as




Figure 4.3: A schematic overview of the analysis strategy. Further detail on each stage of the analysis is provided in the sections indicated.

signal-like, multijet-like, or BIB-like, as described in §4.3.1.3. Finally, a per-event BDT was used to classify events as likely to have been produced by a signal or background process. Two separate versions of this BDT were trained: one optimised for high- m_Φ models and the other for low- m_Φ models. More detail on the per-event BDT can be found in §4.3.2.1.

The final sample was constructed by making selections that removed almost all the non-collision background, leaving only multijet background, and maximised the signal-to-background ratio in the final search region. A data-driven estimate of the remaining background was performed using a modified version of the ABCD method, described in §4.4.3. Since no events were observed in the search region, limits were set on the production cross section times branching ratio, and the result was extrapolated as a function of the LLP decay length. The results are presented in §4.6, where the decay lengths probed range between a few centimetres and a few tens of metres.

4.2 Data and simulated samples

4.2.1 Data samples

All datasets used in the analysis are summarised in Table 4.1. The main dataset used corresponded to a total integrated luminosity of $\int \mathcal{L} dt = 33.0 \text{ fb}^{-1}$ recorded by the ATLAS experiment in 2016 with a centre-of-mass energy $\sqrt{s} = 13 \text{ TeV}$. This was collected with the high- E_T CalRatio trigger, described previously in §3.3.1. A second dataset, containing $\int \mathcal{L} dt = 10.8 \text{ fb}^{-1}$, was collected using the low- E_T CalRatio trigger. The low- E_T trigger was introduced to provide an efficient selection for low- m_Φ models part-way through the 2016 data-taking period and so has a lower integrated luminosity. The difference in efficiency between these two triggers can be seen in Figure 4.4 and is discussed in more detail in §4.2.3.

A further two datasets were collected to study non-collision backgrounds. The first of these used a dedicated trigger running in empty bunch crossings to give a sample of cosmic ray events, which was active in approximately half as many empty bunch crossings as the main triggers. The second used a modified HLT requirement for BIB selection to provide a sample of BIB events in collisions, and was active in the same bunch crossings as the main triggers. These were used when training machine-learning algorithms to reject

Table 4.1: Summary of the data samples used for the analysis.

Dataset name	L1 seed (bunch crossing type)	HLT	$\int \mathcal{L}$	Usage
High- E_T dataset	L1TAU60 (paired)	CalRatio trigger	33.0 fb^{-1}	$m_\Phi \in [400, 1000] \text{ GeV}$ limits
Low- E_T dataset	LLP-NOMATCH (paired)	CalRatio trigger	10.8 fb^{-1}	$m_\Phi \in [125, 200] \text{ GeV}$ limits
Cosmics dataset	L1TAU30 (empty)	CalRatio trigger	-	background estimation
BIB dataset	L1TAU60 (paired)	$(\text{NoIso}) \cap \neg(\text{CalRatio})$	33.0 fb^{-1}	BDT training, bkg estim.

background, and for the optimisation of the final selections. More detail on these triggers is provided in §3.3.1.

4.2.2 Signal and background simulation samples

The $\Phi \rightarrow SS$ signal samples used in this analysis were generated using MadGraph^[84] at leading order (LO) with the NNPDF2.3LO parton distribution function (PDF) set,^[85] with the shower process carried out using PYTHIA 8.210^[86] and the A14 set of tuned parameters (A14 tune).^[87] Several samples were generated, each modelling different combinations of m_Φ and m_S , with $m_\Phi \in [125, 1000] \text{ GeV}$ and $m_S \in [5, 400] \text{ GeV}$, as summarised in Table 4.2. For each combination of m_Φ and m_S , two samples were generated with different LLP decay lengths. One sample with a lab-frame decay length (LF) for the LLPs of approximately 5 m was used to study the signal throughout the analysis. The other sample (with $\text{LF} = 9 \text{ m}$) was used in training the machine learning algorithms as well as in validating the procedure for extrapolating the final limits as a function of the proper decay length of the long-lived scalar S .

The main SM background in this analysis was multijet production, and this was estimated using a data-driven method. However, some simulated multijet events were needed in order to train the machine learning algorithms to reject multijet-like jets and events, and to evaluate some of the systematic uncertainties. These samples were generated using the NNPDF2.3LO PDF set, and PYTHIA 8.186^[88] with the A14 tune for the parton showering and hadronisation.

During simulation, simulated hard scatter events were initially generated in the absence of pile-up. Therefore to correctly model the effect of multiple pp interactions in the same or neighbouring bunches, simulated inclusive pp events were overlaid on each generated event. The multiple interactions were simulated with PYTHIA 8.186 using the A2 tune^[89] and the MSTW2008LO PDF set.^[90] Also, the detector response to the simulated events

Table 4.2: Summary of the mass and decay length parameters simulated for the $\Phi \rightarrow SS$ hidden sector model. Both the lab-frame decay lengths (LF) and proper decay lengths ($c\tau$) are listed for each sample.

m [GeV]	m_S [GeV]	LF=5 m $c\tau$ [m]	LF=9 m $c\tau$ [m]
125	5	0.127	0.229
	8	0.200	0.375
	15	0.580	0.715
	25	0.760	1.210
	40	1.180	1.900
	55	1.540	2.730
200	8	0.170	0.290
	25	0.540	0.950
	50	1.070	1.900
400	50	0.700	1.260
	100	1.460	2.640
600	50	0.520	0.960
	150	1.720	3.140
1000	50	0.380	0.670
	150	1.170	2.110
	400	3.960	7.200

must be evaluated to provide an accurate representation of how a signal or background event would appear in the detector. The response was simulated with the GEANT4-based detector simulation.^[68,69] A full simulation of the detector was used for all the samples, and the standard ATLAS reconstruction software was used for both simulation and pp data.

4.2.3 Trigger efficiencies

The trigger efficiency as a function of the LLP p_T is calculated from simulated signal as the number of LLPs firing the CalRatio trigger for a given p_T bin divided by the number of LLPs generated at that p_T , and is shown in Figure 4.4 for LLPs decaying in the HCal. The figure shows the efficiency for the low- E_T CalRatio trigger in open markers as well as the high- E_T CalRatio trigger in filled markers for three different signal samples. The high- E_T trigger starts to be efficient for LLPs with $p_T > 50$ GeV and reaches its plateau at 200 GeV. On the other hand, the low- E_T trigger recovers a fraction of the LLPs with $p_T < 100$ GeV (especially visible in the low- m_Φ models), while it gives a poorer efficiency for high- p_T jets.




Figure 4.4: Trigger efficiency for simulated signal events as a function of the LLP p_T (left) and the p_T distribution of LLPs (right) for a selection of signal samples.

This decrease in efficiency for high- p_T jets was due to the L1 requirement that there are no EM deposits matched to the most energetic HCal deposit (and the second-most energetic deposit if there is more than one). Low- p_T LLPs deposit little energy in the ECal, unlike higher- p_T LLPs which undergo more bremsstrahlung radiation. Since the requirement was on both HCal deposits, events where one of the LLPs decays in the HCal but the other decays in the ECal or before were much more likely to pass for low- p_T jets than for high- p_T ones. The LLP p_T distribution is also shown in Figure 4.4 for a range of signal samples considered in the analysis. The combination of these two shows that the low- E_T trigger gave a better selection for low- m_ϕ signal models and the high- E_T trigger performed better for the high- m_ϕ signal models.

The trigger efficiency was also calculated as a function of the LLP decay position, by computing the number of LLPs decaying at a given position and firing the CalRatio trigger divided by the number of LLPs generated at that position. This efficiency was strongly dependent on the LLP decay position, as shown in Figure 4.5. Here, the efficiency as a function of the LLP decay length in the x - y plane is shown for LLPs decaying in the barrel ($|\eta| < 1.4$) and in the z direction for LLPs decaying in the endcaps ($|\eta| \geq 1.4$). The efficiency is shown for the low- E_T trigger in open markers and the high- E_T trigger in filled markers for three signal samples. Again, the low- E_T trigger gave a higher efficiency for low- m_ϕ models, while the high- E_T trigger gave a better efficiency for high- m_ϕ mod-




Figure 4.5: Trigger efficiency for simulated signal events as a function of the LLP decay position in the x - y plane for LLPs decaying in the barrel (left, $|\eta| < 1.4$) and in the z direction for LLPs decaying in the HCal endcaps (right, $1.4 \leq |\eta| < 2.5$) for three signal samples.

els. The selection was most efficient in the HCal ($2.25 \text{ m} < L_{xy} < 4.25 \text{ m}$ in the barrel, $4.3 \text{ m} < L_z < 6.05 \text{ m}$ in the endcaps) for both triggers. The efficiency decreases before the end of the HCal both in the barrel and the endcaps because not enough of the jet's energy is deposited in the HCal to pass the trigger threshold, and the majority of the jet's energy punches through to the MS.

4.3 Event selection

4.3.1 Definition of CalRatio jet candidates

The jets used in this analysis were first selected by applying quality requirements to candidate jets. If they passed this selection, these jets were evaluated by a neural-network-based regression algorithm to predict their decay position in L_{xy} and L_z , and then by a per-jet BDT which classified them as signal-like, multijet-like or BIB-like.

4.3.1.1 Preselection

Following selection by either of the CalRatio triggers, events were required to satisfy the *preselection* before being considered in further analysis stages, as shown in Figure 4.3. Two requirements were made: that the events contained so-called *clean* jets, defined below, and that these jets were not close to tracks.

The standard cleaning requirements used within the ATLAS jet cleaning tool exclude jets with low ECal activity. Since using these requirements would eliminate most of the CalRatio jets, the analysis instead used an alternative, LLP-specific, set of cleaning requirements. These did not include the explicit selection on electromagnetic fraction, and instead included a requirement of $(\text{negE} < 4 \text{ GeV}) \cap (\text{FracSamplingMax} > 0.85)$, similar to that in the `cleanLLP` trigger flag (see §3.4.3.3). To be labelled as clean, jets were required to pass the LLP-specific jet cleaning selections, and have $p_T > 40 \text{ GeV}$ and $|\eta| < 2.5$.

In order to select events with trackless jets, an additional event-level variable was defined: $\sum \Delta R_{\min}(\text{jet}, \text{tracks})$, calculated as follows. For each jet, $\Delta R_{\min}(\text{jet}, \text{tracks})$ is the distance from the jet's axis to the closest track with $p_T > 2 \text{ GeV}$ that passes some loose track selections. Summing this value for each clean jet in the event with $p_T > 50 \text{ GeV}$ and $|\eta| < 2.5$ gives $\sum \Delta R_{\min}(\text{jet}, \text{tracks})$. This variable has good separation power between signal-like events and SM-like ones because the displaced jets are trackless, so the distance between the jet axis and the closest track is large. In events with two displaced jets, the sum of minimum distances will also be large, and if there is an additional prompt jet in the event, the distance to its closest track will add only a small contribution to the total sum. On the other hand, regular SM jets have nearby tracks and therefore $\Delta R_{\min}(\text{jet}, \text{tracks})$ for most of these jets will be very small. Events used by the analysis were required to have passed the trigger, to contain at least two clean jets and to have $\sum \Delta R_{\min}(\text{jet}, \text{tracks}) > 0.5$.

4.3.1.2 Decay position prediction for signal jets

In the first of the series of machine learning steps used in the analysis (see Figure 4.3), each clean jet was evaluated by an MLP to predict the decay position of the particle that produced the jet. An MLP is a specific implementation of an artificial neural network (ANN), in which the neurons making up the network are organised in layers, with an input layer using the jet's fraction of energy deposited in each of the ECal and HCal layers as input variables, and an output layer that produces the regression or the classification result, and some number of intermediate fully connected layers. Within the network, each layer is only allowed to directly communicate with the next layer. The MLP used in this analysis was implemented with TMVA^[91] using a regression algorithm to predict the radial (x - y) and longitudinal (z) decay position of the LLP which created each jet, using calorimeter-based variables as input. The algorithm was trained to determine the character and strength of

the relationship between the input variables and the decay position, using the true decay positions of LLPs from simulated signal events during the training.

As described in Chapter 3, both the ECal and HCal are separated into central and forward regions and those regions are further divided into layers along the path from the interaction point. A late-developing jet, such as one from an LLP, will leave energy only in the outer layers of the calorimeter. The pattern of energy left in the layers depends on where the LLP decays in the calorimeter and thus a combination of the following variables is useful in predicting the LLP's decay position and were used as input to the MLP:

- hadronic energy fractions in the HCal central barrel, divided into three layers in the x - y direction;
- hadronic energy fractions in the HCal extended barrel, divided into three layers in the x - y direction;
- hadronic energy fractions in the HCal endcap, divided into four layers in the z direction;
- electromagnetic energy fractions in the ECal barrel, divided into four sections in the x - y direction;
- electromagnetic energy fractions in the ECal endcap, divided into four sections in the z direction;
- jet η , included in the training so the MLP learned which group of layers it should look for when predicting decays.

The separate training and testing samples contained individual jets taken from simulated signal events for a range of models with m_Φ in the range [200, 1000] GeV and lab-frame decay length of 9 m. The samples with a 9 m decay length were chosen as they are statistically independent from those with a decay length of 5 m, which were used in the signal efficiency calculation and statistical analysis. Each reconstructed signal jet was required to be within $\Delta R < 0.2$ of a true LLP, to pass the jet cleaning cuts described in §4.3.1.1, and to have a $p_T > 50$ GeV. Approximately 1 million simulated signal jets were used for training and testing.

Figure 4.6 shows a two-dimensional plot comparing the true radial decay position of an LLP leading to a jet to the MLP-predicted radial decay position, L_{xy} , while the right plot shows the same comparison for the longitudinal decay position, L_z . The first clearly shows




Figure 4.6: Probability density of MLP-predicted radial (L_{xy} , left) and longitudinal (L_z , right) LLP decay positions as a function of the truth LLP decay positions, for reconstructed jets matched to the LLP. Dotted lines show where the MLP value equals the truth value.

the different layers of both the ECal and HCal in areas where decays in the same layer lead to constant MLP radial decay position even as the true decay position changes. However, the overall prediction aligns closely with the true decay position, making the MLP-predicted radial decay position an extremely useful variable. The second figure shows less obvious layering, but a very good correlation between prediction and truth for the whole range of the forward calorimeters.

4.3.1.3 Per-jet signal, BIB, and multijet separation

A per-jet BDT was used to select jets produced by signal and reject those likely to have been produced by BIB or SM multijet processes by classifying them.

The BDT used the result of the MLP and other jet- and cluster-related variables as input. Variables were chosen for their potential to separate signal and background, and their effectiveness within the BDT was determined in order to choose the set of variables that would be used in the BDT. This was done by dropping each variable in turn from the BDT training, and evaluating the signal to background ratio obtained. Variables whose removal had little detrimental effect on the BDT were discarded, following which the remaining variables used to perform the final classification were:

- jet p_T ;
- $\Delta t(\text{BIB}, \text{jet})$: the difference between the timing of the highest- p_T (leading) cluster in the jet and the timing expected from BIB muons originating from the $\pm z$ direction. This variable was designed to be zero if the jet timing was consistent with originating

from a BIB muon, so the distribution had a strong peak at zero for BIB jets while for QCD and signal jets the distribution was broader;

- $\sum p_T(\text{tracks})$ and $\max(p_T(\text{tracks}))$: the sum of and maximum of the p_T of all tracks within $\Delta R < 0.2$ of the jet axis passing the track selection;
- leading cluster lateral and longitudinal width: the width of the leading cluster in the lateral and longitudinal direction respectively;
- leading cluster radius: the distance from the centre of the beam spot to the leading cluster;
- leading cluster energy density: defined as $\sum(E_i^2/V_i)/\sum(E_i)$ where E_i is the energy of cell i in the leading cluster, and V_i is the volume of that cell. The sums are over each cell with positive energy in the cluster;
- leading cluster shower center: calculated as the distance from the inner face of the ECal to the leading cluster's center along the jet axis;
- hadronic layer 1 fraction: fraction of the jet's hadronic energy in the first layer of the hadronic calorimeter;
- predicted L_{xy} and L_z : the output from the MLP predicting the position of the LLP decay that formed the jet.

Of the above variables, those related to jet shape—such as the cluster width, radial position, and energy density—differ between signal and background jets and therefore offer discriminating power; the energy deposits from signal jets are narrower than QCD jets, yet wider than those from BIB jets. Other variables related to tracks associated with jets differentiate strongly between QCD jets and signal and BIB jets, as the former have associated tracks while both of the latter do not. The decay positions predicted by the MLP offer particularly good discrimination between signal and QCD jets, since QCD jets are produced at the interaction point. Finally, the energy deposited in the first HCal layer differentiates well between signal jets, which deposit energy in several layers, and BIB jets, which usually deposit energy in a single layer.

The per-jet BDT was trained using the `scikit-learn`^[92] package, using a 3-class BDT with gradient boosting. The first of three training samples contained simulated signal events for models with m_Φ in the range [125, 1000] GeV and lab-frame lifetime of 9 m. Each reconstructed signal jet was required to be within $\Delta R < 0.2$ of a truth-level LLP,

while LLPs associated with reconstructed signal jets with $|\eta| < 1.4$ had to have a true $L_{xy} > 1250$ mm, and those with $|\eta| \geq 1.4$ had to have a true $L_z > 3500$ mm. This decay-position-based truth matching eliminated signal jets that resembled SM multijet background from the signal sample used in the training. The second training sample contained jets from simulated multijet events, and the third contained jets from the BIB dataset. In the latter sample, to obtain a clean BIB sample with low contamination from multijet events, only the triggering jet in each event was considered by the training, which corresponded to the BIB jet in most cases. These two background samples contained 800,000 events each, while the signal sample had 450,000 events. In all cases, jets with $p_T > 40$ GeV and $|\eta| < 2.5$ were used.

The jet p_T spectrum was very different in each of the three training samples; events in each sample were therefore weighted such that the p_T distribution was flat, ensuring that the per-jet BDT remained as p_T -agnostic as possible. Since the jet p_T was highly correlated with a number of BDT input variables, the jet p_T was included as a variable in the BDT.

The output of the per-jet BDT was a set of three weights that summed to one: signal-weight, BIB-weight and multijet-weight. The distribution of each of these weights is shown in Figure 4.7 for the three types of samples used in the training: a range of signal samples, the BIB sample and the multijets sample. Only events passing the preselection before the trigger are shown. The left column shows the distributions for all clean jets in the event, so signal jets which were not produced by an LLP are included such as jets from recoil or pileup as well as jets produced by a promptly decaying scalar. The right column shows the same distributions but for the signal jets which are matched to a truth LLP as described previously. These were the type of jets used in the per-jet BDT training and the ones that were targeted by this analysis. The distributions for multijets and BIB data are the same between the left and right plots.

The first row in Figure 4.7 represents the distribution of the signal-weight of the jets. For signal events this distribution peaked between 0.34 and 0.35 when truth-matched jets were considered, while for background jets the distribution peaked around 0.325. By comparing the left and right plots, it can be seen that jets in signal events which did not come from a truth-level LLP were typically also assigned low signal-weights, as desired. The separation between background and signal jets was most pronounced for benchmark sig-




Figure 4.7: The distributions of the per-jet BDT weights for a multijet sample, a BIB sample and five signal samples. For the signal samples, the left column shows all clean jets in the events, and the right column shows clean jets matched to an LLP decaying in the calorimeter. The multijet and BIB distributions are the same in the left and right plots, and show all clean jets in the event.

nal models with m_Φ in the range [400, 1000] GeV. A small peak at 0.34 can be seen for the BIB sample. Since this sample contained events passing the CalRatio trigger before the BIB rejection algorithm was applied, every event contained at least one trackless jet with high CalRatio that fired the trigger. These characteristics, similar to those of signal jets, caused the triggering jets to be assigned a high signal-weight.

The second row in Figure 4.7 represents the distribution of the multijet-weight. The jets from multijet events were assigned high values of this weight as expected. The distribution for signal events peaks at low values when truth-matched jets are considered, while the BIB sample is known to be contaminated with standard jets from the collision, and therefore it also contains some multijet-like jets. The multijet-weight for this sample can be separated into three regions: one centered around 0.35, which corresponds to standard multijets; an intermediate region around 0.33 for trackless jets with low CalRatio; and a second peak at 0.32 corresponding to trackless jets with higher CalRatio, which are more likely to be BIB. Once again, the separation between background and signal jets was most pronounced for benchmark signal models with m_Φ in the range [400, 1000] GeV.

The third row of Figure 4.7 shows the distribution of the BIB-weight. The multijet and signal events peaked at intermediate values of this weight, as did the multijet-like events contained within the BIB sample, while higher values were assigned to BIB jets. Figure 4.8 shows the z position of the leading cluster in each jet versus its timing for the BIB sample in two ranges of the BIB-weights. The typical shape for BIB in this plane corresponds to the ‘banana’-shaped bands indicated by the black curves. The bands starting near $t = -2$ ns at each edge of the calorimeter, marked by the dotted lines, correspond to the exit point of the BIB muon associated to the current bunch crossing. When BIB muons enter the calorimeters they have very negative reconstructed timing, which becomes increasingly less negative as they approach the opposite end of the calorimeter where the paths followed by particles produced at the IP and BIB muons become more similar. The plot on the left of Figure 4.8 corresponds to jets in the BIB sample with a BIB-weight < 0.34 —these BIB-weight values identify standard multijet-like jets. The plot on the right presents the same distribution for jets with BIB-weight > 0.34 , showing that BIB jets (both from the current BC and from the next BC) are accurately identified by the per-jet BDT since most jets fall inside the typical BIB ‘banana’ regions.




Figure 4.8: Distribution of the jet leading cluster’s z -position versus its timing for BIB data after high- E_T preselection, for jets with BIB-weight < 0.34 (left) and jets with BIB-weight > 0.34 (right). The regions marked by the dashed lines correspond to BIB from the next bunch crossing, separated by 25ns; and those marked by the solid lines correspond to BIB from the next-to-next bunch crossing. The flat region of the distribution with near-zero but positive timing corresponds to jets associated with the collision.

The per-jet BDT had better signal-to background discrimination for high- m_Φ models than for low- m_Φ models, in which case it was more difficult to distinguish the signal from BIB and jets from pile-up[†]. The main reason for this lies in the p_T distributions—both BIB and pile-up jets have relatively soft p_T spectra, and even though these backgrounds were mitigated by the jet-cleaning requirements, their remaining contributions were harder to distinguish from signal at low p_T .

4.3.2 Definition of CalRatio signal events

4.3.2.1 Per-event signal and background separation

Following the classification of jets, another BDT was used to separate signal-like events from background-like events—the last of the three machine-learning steps in the analysis (see Figure 4.3). This was designed with the main objective of separating BIB events from signal events where at least two jets have a considerable displacement. Although a large

[†]The presence of pile-up jets has two effects: they can leave energy deposits in the ECal, changing the fraction of energy per calorimeter layer and worsening the signal-to-background discrimination. In addition, pile-up jets’ tracks do not point back to the IP in many cases and hence are not removed by track isolation requirements. These jets can therefore be reconstructed as nearly trackless, making them more similar to signal.

portion of the BIB background was removed by the BIB rejection algorithm within the trigger, some BIB still remained. Furthermore, the trigger algorithm was effective at removing BIB associated with the current bunch crossing but not for BIB in the next or next-to-next bunch crossing.

A combination of signal samples and the BIB dataset were used to train the per-event BDT. The BIB dataset could contain SM multijet events in addition to the BIB event that caused them to be selected by the trigger. Consequently, the per-event BDT was able to discriminate signal from multijet background as well as BIB even though no multijet sample was used in the training.

In signal, the two jets with the highest per-jet BDT signal-weight in the event were the most likely to come from an LLP decay in the calorimeter. These are referred to as *CalRatio jet candidates*. In BIB data, the two jets with the highest per-jet BDT BIB-weight in the event (*BIB jet candidates*) were the most likely to be BIB-jets capable of faking signal. Hence, these jets were selected and their per-jet weights together with other characteristics of events were used as input variables to the per-event BDT. The variables chosen for the per-event BDT were selected using the same elimination procedure as in the per-jet BDT. The final set of input variables which gave the highest signal and background separation, and therefore used in the training were:

- per-jet BDT signal-weights of the CalRatio jet candidates and BIB jet candidates;
- per-jet BDT BIB-weights of the CalRatio jet candidates and BIB jet candidates;
- p_T of the CalRatio jet candidates;
- H_T^{miss}/H_T , defined in §3.4.5.2, for jets with $p_T > 30$ GeV and $|\eta| < 3.2$;
- m_{eff} , defined in §3.4.5.2;
- $\Delta\phi(j_1, j_2)$ and $\Delta R(j_1, j_2)$, the opening angle and distance between the two most signal-like jets in the event;
- the mean signal-weight of all the clean jets in the event;
- the mean BIB-weight of all the clean jets in the event.

As in the per-jet-BDT, the separation of signal from BIB was much more difficult for the low- m_Φ models, since BIB-jets have a similar p_T distribution to the LLPs in these models. Event-level variables like H_T and m_{eff} are also similar in BIB events and events from low- m_Φ signal models. Therefore, to optimise the selection, two versions of the BDT were




Figure 4.9: Distribution of the low- E_T per-event BDT (left) and high- E_T per-event BDT (right) on main data, BIB data and five signal samples after preselection.

trained for the analysis of the high- E_T dataset and the low- E_T dataset separately, where both BDTs used the input variables listed above. The only difference between the two BDTs was the combination of signal samples used in the training and the trigger selection. The high- E_T per-event BDT training used a combination of low-, intermediate- and high- m_Φ signal samples with no trigger selection applied. The low- E_T per-event BDT training used a combination of low- m_Φ signal samples and additionally used only events passing the low- E_T trigger.[†]

Figure 4.9 shows the performance of both per-event BDTs for five signal models as well as for the main dataset and the BIB sample. The output of both BDTs has a common structure with three different regions: very low values of the BDT score are assigned to events with a high level of BIB-contamination (BDT < -0.2). An intermediate region (-0.2 < BDT < 0) is dominated by multijet events with low BIB contamination, and higher BDT values are assigned to signal-like events. This can be seen by comparing the distributions for the BIB sample, which is dominated by BIB with some multijets contamination, to the main dataset, which is dominated by multijets with some BIB contamination. Taking this separation between BIB and multijets into account, the per-event BDT served two purposes in the analysis: first, it was used as part of the event cleaning described in §4.3.2.2 to reject BIB

[†] All the machine-learning steps were trained on a mixture of HS signal samples (with different mediator masses) to maintain model-dependence, but it should be noted that in later studies where the same selection was applied to different signal models (described in Chapter 6), it became apparent that some amount of model-dependence remained—especially in the low- E_T case, which was trained only on the low- E_T signal samples.

so that the only remaining background was from multijets; second, it was used as one of the two variables which defined the ABCD plane for background estimation in the search region (see §4.4), the other being $\sum \Delta R_{\min}(\text{jet}, \text{tracks})$.

4.3.2.2 Event cleaning

The ABCD method used for the background estimation in this analysis, described in §4.4, is designed for the case of a single background population in the final selection. Contamination from BIB risks impeding the proper functioning of the method because the BIB and SM multijet backgrounds are two populations with quite different distributions in the ABCD plane. Hence, an additional set of selections to eliminate BIB prior to the application of the ABCD method was applied, described below.

Figures 4.10 and 4.11 show that the per-event BDT was able to identify BIB events at low BDT values. Figure 4.10 includes jets from events passing the high- E_T preselection and Figure 4.11 includes jets from events passing the low- E_T preselection, with trigger matching ($\Delta R(\text{offline jet}, \text{HLT jet}) < 0.2$) required in both. These show the jet's leading cluster z -position versus its timing, where BIB is expected to fall into the bands defined by the black curves. The first column in each figure corresponds to the BIB sample while the second column shows the main dataset.

The plots at the top of Figures 4.10 and 4.11 show events with BDT score < -0.2 , with many of the events falling into the BIB bands. The plots in the middle rows are for events with $-0.2 < \text{BDT} < 0.1$, which are dominated by multijet events but still include some BIB. The plots at the bottom are for events with BDT score > 0.1 , where there is no evidence of BIB events remaining. Events which entered the high- E_T data analysis were therefore required to satisfy high- E_T per-event BDT > 0.1 , and those in the low- E_T analysis required low- E_T per-event BDT > 0.1 , to ensure that most BIB was removed at this level.

After enforcing the preselection and the per-event BDT cuts, some events still remained with characteristics atypical of multijet background or signal. Further selections were therefore applied to remove these events. The full list of *event cleaning* requirements were therefore as follows:

- high- E_T per-event BDT > 0.1 (low- E_T per-event BDT > 0.1);




Figure 4.10: Distribution of the jet leading cluster's z -position versus its timing. Left plots are for BIB data; right plots are for main data. Events with high- E_T per-event $\text{BDT} < -0.2$ are shown at the top; with $-0.2 < \text{high-}E_T \text{ per-event BDT} < 0.1$ in the middle; and with high- E_T per-event $\text{BDT} > 0.1$ at the bottom.




Figure 4.11: Distribution of the jet leading cluster's z -position versus its timing. Left plots are for BIB data; right plots are for main data. Events with low- E_T per-event $\text{BDT} < -0.2$ are shown at the top; with $-0.2 < \text{low-}E_T \text{ per-event BDT} < 0.1$ in the middle; and with low- E_T per-event $\text{BDT} > 0.1$ at the bottom.

- trigger matching: at least one of the CalRatio jet candidates had to be matched to the HLT jet that fired the trigger, defined by $\Delta R(\text{offline jet}, \text{HLT jet}) < 0.2$;
- a timing window of $-3 < t < 15$ ns for the CalRatio jet candidates and the BIB jet candidates. This was designed to reject out-of-time pileup from the previous and next bunch crossings which peaked close to $t = \pm 25$ ns, and ultimately removed a significant part of the remaining background while retaining a large fraction of the signal events.

4.3.2.3 Final selections

An optimisation procedure was used to determine the final selections defining the signal region. Several variables were used in this process, and different combinations of selections were tested. For each selection and each signal model, the $S/\sqrt{B}$ ratio was calculated, where S is the number of signal events entering the search region and B is the background estimate obtained by applying the ABCD method as described in §4.4. For each analysis path, the selection giving the best $S/\sqrt{B}$ ratio was chosen among those where the signal contamination in regions B, C and D was lower than 30%. The variables considered in this process were:

- $\sum_{j_1, j_2} \log_{10}(E_H/E_{EM})$: in every event passing a CalRatio trigger there is at least one jet with high $\log_{10}(E_H/E_{EM})$: the triggering jet. If that is the only high CalRatio jet in the event, this variable shows a peak close to one for data. In signal this peak is shifted towards higher values. In events with a second high CalRatio jet, a second peak appears at values higher than two.
- $p_T(j_1)$ and $p_T(j_2)$: signal jets tend to have high p_T , with a distribution that peaks close to $m_\Phi/2$ while data has a softer distribution. Cutting on $p_T(j_2)$ was found to increase the $S/\sqrt{B}$ ratio in the low- E_T selection only.
- H_T^{miss}/H_T : this variable was originally designed to reject BIB and hence it was used as input to the per-event BDTs. In events composed of only one jet from BIB, the calorimeter energy is completely uncompensated with a value for H_T^{miss} very close to that of H_T . In events with very low contributions to H_T^{miss} the variable peaks at low values; this is the case for signal. In events with some H_T^{miss} there is a tail towards higher values. After the selection this tail was more visible in data than in signal in the high- E_T selection.

The optimisation was performed separately for the high- E_T and low- E_T analyses. As a result, two final selections were defined, which are referred to as the *high- E_T selection* and *low- E_T selection*. The maximum sensitivity for low masses ($m_\Phi = 125$ and 200 GeV) was obtained from the low- E_T selection, while the optimal selection for higher masses ($m_\Phi = [400, 1000]$ GeV) was obtained from the high- E_T selection. The definitions of these two selections were as follows:

- low- E_T selection:
 - low- E_T per-event BDT > 0.1 ,
 - $(\sum_{j_1, j_2} \log_{10}(E_H/E_{EM})) > 2.5$,
 - $p_T(j_1) > 80$ GeV and $p_T(j_2) > 60$ GeV;
- high- E_T selection:
 - high- E_T per-event BDT > 0.1 ,
 - $H_T^{\text{miss}}/H_T < 0.6$,
 - $(\sum_{j_1, j_2} \log_{10}(E_H/E_{EM})) > 1$,
 - $p_T(j_1) > 160$ GeV and $p_T(j_2) > 100$ GeV.

The full list of all selections applied in each of the two analysis paths is shown in Table 4.3.

Table 4.3: A summary of the requirements for event to enter the ABCD planes of the high- E_T and low- E_T selections.

Analysis Selection	High- E_T	Low- E_T
L1 Trigger	L1TAU60 [Main, BIB datasets] L1TAU30_EMPTY, L1TAU $E_T \geq 60$ GeV [Cosmics dataset]	L1LLP-NOMATCH [Main, BIB datasets] L1TAU30_EMPTY [Cosmics dataset]
HLT Trigger	CalRatio Trigger [Main, Cosmics datasets] (Trigger with no BIB removal) $\cap$ ($\neg$ CalRatio Trigger) [BIB dataset]	
Preselection	$\sum \Delta R_{\min}$ (tracks, clean jets) > 0.5 (sum over clean jets with $p_T > 50$ GeV; tracks with $p_T > 2$ GeV, Loose vertex association, and ChargedLoosePion selection)	
Event cleaning	high- E_T per-event BDT > 0.1 Trigger matching: at least one candidate CalRatio jet matched to HLT triggering jet within $\Delta R < 0.2$ Signal-like jet timing: both candidate CalRatio jets have timing t between $-3 < t < 15$ ns BIB-like jet timing: the two jets with highest per-jet BDT BIB-weight have timing t between $-3 < t < 15$ ns	
Final Selection	$H_T^{\text{miss}}/H_T < 0.6$ $\sum_{j_1, j_2} \log_{10}(E_H/E_{\text{EM}}) > 1.0$ $p_T(j_1) > 160$ GeV, $p_T(j_2) > 100$ GeV $\sum_{j_1, j_2} \log_{10}(E_H/E_{\text{EM}}) > 2.5$ $p_T(j_1) > 80$ GeV, $p_T(j_2) > 60$ GeV	

4.4 Background estimation

A data-driven method was used to estimate the contribution from the dominant background (SM multijet events) to the final search regions. This method, known as the ABCD method, relies on the assumption that the distribution of background events can be factorised in the plane of two relatively uncorrelated variables. In this plane, the method uses three control regions (B, C and D) to estimate the contribution of background events in the search region (A), described in more detail in §4.4.3. It is important that there is only one source of background in the ABCD planes, so the number of BIB and cosmic events entering the planes was first verified to be negligible.

4.4.1 BIB background removal check

The per-event BDTs defined in §4.3.2.1 were designed to reject BIB events. Their good performance and the proper BIB rejection shown in Figures 4.10 and 4.11 was cross-checked using the BIB data sample. The number of BIB events passing each stage of the analysis selections is shown in Table 4.4.

For the high- E_T selection, fifteen BIB events passed all selection criteria and entered the ABCD plane. This was within the statistical uncertainty for the number of events passing all selections in the main dataset (459 ± 21), and furthermore only one of these fifteen events entered the search region (region A, defined in Table 4.6). Similarly, in the low- E_T selection, only four events passed all the selections (well within the statistical uncertainty for the number of events passing in the main dataset, 136 ± 12), and of these none entered region A.

A significant proportion of events in the BIB sample are SM multijet events, which would form part of the same population as the main background so would not invalidate the ABCD method. Therefore, it was also checked that the BIB events which did enter the ABCD plane did not have the typical characteristics of BIB jets, by checking that they fell outside the typical BIB ‘banana’ shapes in time versus z -position. This confirmed that all events containing real BIB jets in the BIB data sample were removed after the selections.

Table 4.4: Sequential impact of each requirement on the number of events passing the selection for the high- E_T (top) and low- E_T (bottom) analyses. The quoted number of events for the cosmic rays was estimated from the cosmic-ray dataset by applying corrections for the proportion of empty crossings relative to filled crossings, and a factor to account for the change in analysis efficiency in a zero-pile-up collision. Signal yields are quoted as a percentage of the total number of generated events.

High- E_T selection:		Main data	BIB	Cosmic rays	Signal (m_Φ, m_S) = (1000, 150) GeV $c\tau = 1.17$ m	Signal (m_Φ, m_S) = (600, 150) GeV $c\tau = 1.72$ m	Signal (m_Φ, m_S) = (400, 100) GeV $c\tau = 1.46$ m
Preselection:	Pass trigger, 2 clean jets & $\sum \Delta R_{\min} > 0.5$	1375483	183015	526.0	26.2%	22.4%	17.5%
Event cleaning:	High- E_T per-event BDT > 0.1	4515	192	7.6	25.4%	21.2%	15.3%
	Trigger matching	3627	119	3.8	24.5%	20.4%	15.0%
	$-3 < t < 15$ ns	3388	110	3.2	24.0%	20.0%	14.8%
High-E_T selection:	$\sum_{j_1, j_2} \log_{10}(E_H/E_{EM}) > 1$	1815	61	2.7	21.7%	16.8%	11.5%
	$H_T^{\text{miss}}/H_T < 0.6$	1421	41	2.1	18.1%	15.2%	10.9%
	$p_T(j_1) > 160$ GeV	774	26	0	17.5%	13.6%	7.50%
	$p_T(j_2) > 100$ GeV	459	15	0	16.5%	11.8%	5.56%
Region A :		10	1	0	10.7%	7.74%	3.10%

Low- E_T selection		Main data	BIB	Cosmic rays	Signal (m_Φ, m_S) = (200, 50) GeV $c\tau = 1.07$ m	Signal (m_Φ, m_S) = (125, 25) GeV $c\tau = 0.76$ m
Preselection:	Pass trigger, 2 clean jets & $\sum \Delta R_{\min} > 0.5$	2180349	95247	319.1	7.58%	4.33%
Event cleaning:	Low- E_T per-event BDT > 0.1	40474	678	65.1	6.26%	2.73%
	Trigger matching	34567	538	42.1	5.97%	2.51%
	$-3 < t < 15$ ns	33680	519	23.4	5.86%	2.46%
Low-E_T selection:	$\sum_{j_1, j_2} \log_{10}(E_H/E_{EM}) > 2.5$	722	13	18.3	0.92%	0.39%
	$p_T(j_1) > 80$ GeV	304	6	7.3	0.69%	0.16%
	$p_T(j_2) > 60$ GeV	136	4	3.5	0.60%	0.10%
Region A:		7	0	0.4	0.43%	0.07%

4.4.2 Cosmic background removal check

The analysis selections detailed in §4.3 also removed almost all the cosmic ray background—this was checked using the cosmics dataset, which was collected in empty bunch crossings. To obtain an estimate of the number of cosmic events which might contaminate the ABCD plane, two factors were considered.

First, the fact that the number of empty bunch crossings was not the same as the number of filled bunch crossings was quantified by a filled-to-empty live-time factor R_{live} . This was obtained from the trigger record, which stores the triggers that were running for a particular run, whether they were prescaled and the prescale factor, and the number of bunch crossings per run. Both the high- E_T and cosmics triggers were running unprescaled for the whole 2016 data-taking period, during 6.31×10^7 filled bunch crossings and 3.31×10^7 empty bunch crossings. For the high- E_T selection, R_{live} was therefore equal to 1.9. The number of filled bunch crossings where the low- E_T trigger was active was 7.1×10^6 . Therefore, for the low- E_T selection, $R_{\text{live}} = 0.21$.

Second, events selected during empty bunch crossings are much cleaner than events collected during regular collisions, so they are much more likely to pass the analysis selections (in particular, the requirement at HLT that the jets have no tracks with $p_T > 2$ GeV within $\Delta R < 0.2$ around their axis). An additional factor, R_{clean} , was applied to convert the number of events observed in an empty bunch crossing (largely trackless) to the number expected in the more noisy environment of a filled bunch crossing with potentially many tracks. Assuming that R_{clean} does not depend strongly on the trigger used to collect the data, this factor was conservatively estimated by dividing the number of events entering the ABCD plane in the main dataset by the number passing the selection if all tracks in the event were ignored. This was done by reprocessing the main dataset but manually setting the number of tracks in each event to be exactly zero. All analysis steps, including the evaluation of all BDTs, were then repeated using the modified track information. This approach is conservative because the low- E_T trigger may be more efficient in empty bunch-crossings due to the extra topological information it has available, and the CalRatio candidates from cosmic rays would likely also be narrower. These effects were not estimated here, but would only reduce the estimated number of cosmic events contaminating the ABCD plane in the main dataset, and therefore these estimates should be seen as upper

limits. For the high- E_T analysis, 459 events entered the ABCD plane in the main dataset, versus 8640 when tracks were ignored, giving $R_{\text{clean}} = 0.05$. For the low- E_T analysis, 136 events entered the ABCD plane in the main dataset, versus 523 when tracks were ignored, giving $R_{\text{clean}} = 0.26$.

In summary, the upper limit on the expected number of cosmic ray events which entered the ABCD plane ($N_{\text{cosmics}}^{\text{main dataset}}$) was estimated by:

$$N_{\text{cosmics}}^{\text{main dataset}} = N_{\text{cosmics}}^{\text{cosmics dataset}} \cdot R_{\text{live}} \cdot R_{\text{clean}},$$

where $N_{\text{cosmics}}^{\text{cosmics dataset}}$ is the number of events passing the selection in the cosmics dataset. After application of the high- E_T analysis selection (and additionally requiring the L1TAU item to have $E_T > 60$ GeV to match the high- E_T trigger threshold), zero events from the cosmics dataset remained. Applying the above formula therefore gave an estimate of $N_{\text{cosmics}}^{\text{main dataset}} = 0 \pm 0.04$, which was entirely negligible. When applying the low- E_T selection, a more significant number of cosmic events, 64, passed the selection and entered the ABCD plane, of which seven were in region A. The upper bound for $N_{\text{cosmics}}^{\text{main dataset}}$ was then 3.5 ± 0.4 . Although this value was not negligible, it was well within the statistical uncertainty for events entering the ABCD plane (126 ± 12) and was therefore deemed to be unlikely to interfere with the functioning of the ABCD method. In region A in particular, the same formula suggested that fewer than 0.4 cosmics events would be present, which once again was well within the statistical uncertainty on the estimated number of events in that region. The estimated number of cosmic events passing each stage of the analysis selections is shown in Table 4.4.

4.4.3 Multijet background estimation: ABCD method

Although the probability of producing a jet via SM processes that satisfies the CalRatio jet candidate selection criteria is low, the SM multijet cross-section is large enough that this background is not negligible. A modified version of the ABCD method was used to estimate the contribution from this background in the search region.

Four regions are defined based on a two-dimensional ABCD plane, as shown in Figure 4.12. The choice of the two variables defining the ABCD plane has to be such that they provide a good separation of signal from background, and that signal contamination




Figure 4.12: Diagram of the four regions in the ABCD plane, defined by two uncorrelated variables.

into regions B, C and D is small. Additionally, there should be no correlation between the two variables in the background distribution so that the relation $A/B = C/D$, and by extension $A/C = B/D$, are valid. In the background-only hypothesis, or if all the signal events are concentrated in region A, the number of background events in region A can be predicted from the population of the other three regions using the following formula:

$$N_A = \frac{N_B \times N_C}{N_D}, \quad (4.1)$$

where N_X is the number of background events in region X .

In reality, some signal events may lie outside of region A, and this can skew the result of the standard ABCD method. To account for this, a likelihood-based ABCD method was used, which involved fitting to background and signal simultaneously during the calculation of the limits. This is discussed in detail in §4.6.2.1. The two variables found to fulfil the requirements of the ABCD method were: $\sum \Delta R_{\min}(\text{jet, tracks})$, as defined in §4.3.1.1; and high- E_T per-event BDT or low- E_T per-event BDT depending on the signal model and dataset being analysed.

Figure 4.13 shows the variables defining the ABCD plane after event cleaning, for the high- E_T and low- E_T analyses respectively. These two variables have a good separation between signal and multijet background, as can be seen in Figure 4.14 which shows the distribution of events in the final ABCD planes in signal, main data and BIB data.




Figure 4.13: The $\sum \Delta R_{\min}(\text{jet, tracks})$ (left) and per-event BDT value (right) distributions for data, BIB jets and signal after applying the event cleaning selection criteria, for the high- E_T analysis (top row) and low- E_T analysis (bottom row). The error bars account for statistical uncertainties only.

Region A is defined by the last selection criterion shown in Table 4.4 in each case:

- high- E_T per-event BDT ≥ 0.22 and $\sum \Delta R_{\min} \geq 1.5$ for the high- E_T analysis;
- low- E_T per-event BDT ≥ 0.22 and $\sum \Delta R_{\min} \geq 1.5$ for the low- E_T analysis.

Regions B, C, and D are obtained by reversing one or both of the cuts. The definitions are listed explicitly in Table 4.6.

These requirements gave a signal contamination of $< 25\%$ in regions B, C and D with respect to the total ABCD plane content. The correlations between the ABCD plane variables were checked by dividing regions C and D into several sub-regions and fitting the ratios of




Figure 4.14: The distributions of $\sum \Delta R_{\min}(\text{jet, tracks})$ vs. high- E_T per-event BDT (left column) or low- E_T per-event BDT (right column) for a signal sample (top), main data (middle), and BIB events (bottom) after event cleaning. The signal sample with ($m_\phi = 600$ GeV, $m_S = 150$ GeV) is shown for the high- E_T selection, while ($m_\phi = 125$ GeV, $m_S = 25$ GeV) is shown for the low- E_T selection.

Table 4.5: Application of the ABCD method to the high- E_T and low- E_T selections. Only statistical uncertainties are considered in this table. The estimated number of events in A is obtained from the ABCD relation $N_A = N_B \times N_C / N_D$, and propagating the statistical uncertainties.

Selection	Estim. A	A	B	C	D
High- E_T selection	6.7 ± 2.3	10	9	187	253
Low- E_T selection	2.5 ± 1.8	7	2	70	57

the number of events in each sub-region C to the number of events in each sub-region D to a straight line. The slope in these fits was found to be consistent with zero.

The estimated and observed number of events in region A and the observed numbers in regions B, C, and D for the high- E_T and low- E_T selections are shown in Table 4.5, where the estimated contribution (Estim. A) is obtained with the simple ABCD relation assuming no signal. The final ABCD planes after the high- E_T and low- E_T selections are shown in Figure 4.15 for relevant signal samples in each of the selections.

4.4.3.1 Validation of the ABCD method

The ABCD method was verified by defining validation regions (VRs) orthogonal to the nominal signal regions and checking that the background estimation performed as expected. The exact definitions of all VRs are listed in Table 4.6.

For the high- E_T selection, $VR_{\text{high-}E_T}$ was defined in the same way as the nominal selection but in a low p_T range and removing the cut on H_T^{miss}/H_T to retain enough events. In order to have enough statistics in all ABCD regions, the boundary for high- E_T per-event BDT was lowered to 0.12 and an upper bound of high- E_T per-event BDT < 0.22 was applied to reduce the amount of signal contamination, as seen in the left plots of Figure 4.16. In this figure the ABCD plane for $VR_{\text{high-}E_T}$ is shown for one of the high mass signal samples, main data and BIB data.

For the low- E_T selection it was difficult to find an alternative plane orthogonal to the nominal selection with high enough statistics and low enough signal contamination. Reversing one of the cuts as done with $VR_{\text{high-}E_T}$ did not give enough statistics for a realistic check, and reversing all the cuts took the VR too far away from the SR. The best region found to validate the method was $VR_{\text{low-}E_T}$, defined with the same selection as the nominal, but modifying the boundaries on the ABCD plane variables, such that all events entering the




Figure 4.15: The $\sum \Delta R_{\min}(\text{jet, tracks})$ vs. per-event BDT distribution for several signal samples in the low- E_T selection (top row) and high- E_T selection (middle and bottom row).




Figure 4.16: The $\sum \Delta R_{\min}(\text{jet, tracks})$ vs. per-event BDT distribution for a signal sample (top), main data (middle) and BIB data (bottom) in $VR_{\text{high-}E_T}$ (left column) and $VR_{\text{low-}E_T}$ (right column).

plane were excluded from the nominal region A. A lower range in low- E_T per-event BDT was selected as the VR, where the signal contamination was relatively small. This range could not be extended to lower values of low- E_T per-event BDT since it was necessary to make sure that BIB was still completely eliminated. The right plots of Figure 4.16 show the ABCD plane for $VR_{\text{low-}E_T}$ for one of the low mass signal samples, main data and BIB data. Signal contamination in all VRs was small, confirmed using $S/\sqrt{B}$, which was 7 times smaller in $VR_{\text{high-}E_T}$ than in the nominal high- E_T selection and 3.2 times smaller in $VR_{\text{low-}E_T}$ than in the nominal low- E_T selection. The ABCD plane variables were also found to have no correlation in the VRs. In all VRs, the estimated number of background events was in good agreement with the number of data events observed in region A, confirming the validity of the ABCD method; this is summarised in Table 4.7.

Table 4.6: The definitions of the ABCD planes of the high- E_T and low- E_T selections and their respective validation regions.

Analysis Selection	High- E_T		Low- E_T	
	High- E_T Search Region	High- E_T Validation Region	Low- E_T Search Region	Low- E_T Validation Region
Region	$\frac{H_T^{\text{miss}}}{H_T} < 0.6$ $\sum_{j_1, j_2} \log_{10}(E_H/E_{\text{EM}}) > 1.0$ $p_T(j_1) > 160 \text{ GeV}, p_T(j_2) > 100 \text{ GeV}$	$\frac{H_T^{\text{miss}}}{H_T} < 1.0$ $\sum_{j_1, j_2} \log_{10}(E_H/E_{\text{EM}}) > 1.0$ $100 < p_T(j_1) < 160 \text{ GeV}, p_T(j_2) > 100 \text{ GeV}$	- - $\sum_{j_1, j_2} \log_{10}(E_H/E_{\text{EM}}) > 2.5$ $p_T(j_1) > 80 \text{ GeV}, p_T(j_2) > 60 \text{ GeV}$	$\frac{\log_{10}(E_H/E_{\text{EM}})}{\text{low-}E_T \text{ per-event BDT}} < 0.22$ - $\sum_{j_1, j_2} \log_{10}(E_H/E_{\text{EM}}) > 2.5$ $p_T(j_1) > 80 \text{ GeV}, p_T(j_2) > 60 \text{ GeV}$
Region A	$\sum \Delta R_{\text{min}} \geq 1.5$ high- E_T per-event BDT ≥ 0.22	$\sum \Delta R_{\text{min}} \geq 1.5$ high- E_T per-event BDT ≥ 0.12	$\sum \Delta R_{\text{min}} \geq 1.5$ low- E_T per-event BDT ≥ 0.22	$\sum \Delta R_{\text{min}} \geq 1.5$ low- E_T per-event BDT ≥ 0.13
Region B	$\sum \Delta R_{\text{min}} < 1.5$ high- E_T per-event BDT ≥ 0.22	$\sum \Delta R_{\text{min}} < 1.5$ high- E_T per-event BDT ≥ 0.12	$\sum \Delta R_{\text{min}} < 1.5$ low- E_T per-event BDT ≥ 0.22	$\sum \Delta R_{\text{min}} < 1.5$ low- E_T per-event BDT ≥ 0.13
Region C	$\sum \Delta R_{\text{min}} \geq 1.5$ high- E_T per-event BDT < 0.22	$\sum \Delta R_{\text{min}} \geq 1.5$ high- E_T per-event BDT < 0.12	$\sum \Delta R_{\text{min}} \geq 1.5$ low- E_T per-event BDT < 0.22	$\sum \Delta R_{\text{min}} \geq 1.5$ low- E_T per-event BDT < 0.13
Region D	$\sum \Delta R_{\text{min}} < 1.5$ high- E_T per-event BDT < 0.22	$\sum \Delta R_{\text{min}} < 1.5$ high- E_T per-event BDT < 0.12	$\sum \Delta R_{\text{min}} < 1.5$ low- E_T per-event BDT < 0.22	$\sum \Delta R_{\text{min}} < 1.5$ low- E_T per-event BDT < 0.13

Table 4.7: Application of the ABCD method to the high- E_T and low- E_T VRs. Only statistical uncertainties are considered in this table. The estimated number of events in A is obtained from the ABCD relation $N_A = N_B \times N_C / N_D$, and propagating the statistical uncertainties.

Selection	Estim. A	A	B	C	D
VR _{high-E_T}	66.3 ± 15.0	70	64	57	55
VR _{low-E_T}	54.1 ± 17.4	36	35	34	22

Table 4.8: Summary of the signal systematic uncertainties for models of a given m_Φ are shown in the upper table. Signal model systematics were found to be consistent between signal models of the same m_Φ but different lifetime or m_S . The quoted uncertainties above were taken from the model with the lowest statistical uncertainty. The systematic errors in the lower table were the same for all samples.

m_Φ Point (GeV)	JES	JES _{EMF}	JER	Trigger	Pileup	PDFs	BDT	Stat
125	9%	17%	5%	1%	12%	3%	2%	6%
200	3%	5%	2%	1%	10%	3%	1%	2%
400	3%	8%	1%	2%	4%	4%	1%	< 1%
600	< 1%	2%	1%	1%	3%	5%	< 1%	< 1%
1000	< 1%	< 1%	1%	0.5%	1%	8%	< 1%	< 1%

Systematic Error	Value
Luminosity	$\pm 2\%$
ABCD uncertainty for high- E_T Plane	$\pm 22\%$
ABCD uncertainty for low- E_T Plane	$\pm 24\%$

4.5 Systematic uncertainties

This section describes the systematic uncertainties affecting the predicted background rate and signal yields. Table 4.8 summarises all the systematic uncertainties, which are described individually below.

Experimental uncertainties

Uncertainty on the ABCD method. A dijet-enriched sample was selected using a single-jet-based trigger, vetoing on the CalRatio triggers to ensure orthogonality, and applying the following requirements: the leading and subleading jets had to pass standard quality selections and have $p_T > 320$ GeV and $p_T > 60$ GeV respectively, where the first threshold was from the single-jet trigger and the second threshold removed any remaining BIB; they had to be roughly back-to-back ($\Delta\phi > 3$); their p_T had to be approximately balanced $((p_{j0}^T - p_{j1}^T)/(p_{j0}^T + p_{j1}^T)) < 0.3$; and their H_T^{miss} was required to be below 120 GeV. These selections were designed to ensure the sample contained only QCD events with no signal contamination, while remaining as close as possible to the nominal selection.

The ABCD plane was then defined in the same way as for the main analysis but adjusting the boundaries so that each region contained enough events to ensure that any potential discrepancy was not statistically dominated. The statistical uncertainty on the estimate was

Table 4.9: Application of the ABCD method to the high- E_T and low- E_T control regions (CRs) used to determine the systematic uncertainty associated with the ABCD method. Only statistical uncertainties on the number of events are considered in this table.

Control Region	Estim. A	A	B	C	D
High- E_T CR	60 ± 10	76	57	202	193
Low- E_T CR	306 ± 61	247	231	65	49

below 20%. The discrepancy between the number of events in region A predicted by the ABCD method and the observed number of events in region A was taken as the systematic uncertainty associated with the method. Given that the ABCD plane was defined using the per-event BDT value and this was different in the high- E_T and low- E_T selections, this systematic uncertainty was calculated for each of the planes separately. The number of events in each of the control regions is shown in Table 4.9. The final uncertainty was 21.5% in the high- E_T ABCD plane and 23.9% in the low- E_T plane.

Luminosity. The uncertainty on the integrated luminosity was 2.2%. It was derived from a calibration of the luminosity scale using x - y beam-separation scans, following the standard ATLAS methodology.^[93,94]

Jet energy scale. The uncertainty associated with the Jet Energy Scale (JES) was provided by a common ATLAS package, where the effect of the JES variation was compared to the nominal scale for each of four uncertainty sets, as a function of the jet pseudorapidity. The four scenarios all gave changes in the final signal yields of fairly similar size, typically between 0.5% and 9% depending on the model (the low-mass models were more affected by the JES than the high-mass models). The largest of the four variations was taken as the uncertainty, which was very small compared to the size of the uncertainty on the ABCD method.

The JES uncertainty was also re-evaluated as a function of $E_{EM}/(E_{EM} + E_H)$ (known as EMF) to understand the effect on low-EMF (i.e. high CalRatio) jets. The jet calibration itself was not changed, only the uncertainty was re-evaluated. The method was the same as that used in the standard ATLAS in-situ jet pseudorapidity intercalibration,^[95] but the uncertainty was determined as a function of EMF (as well as η) because the BDTs most strongly depended on it and other parameters were correlated with it.

The two uncertainties (standard JES and JES_{EMF}) were calculated separately for each reference sample and the final results were treated as independent uncertainties. The final uncertainties are listed in Table 4.8. This approach overestimated the uncertainties somewhat, because while the regular JES systematic was applied to every jet in the event, the JES_{EMF} uncertainty was also applied to low-EMF jets, so these jets received a systematic contribution from both methods.

Jet energy resolution. The uncertainty associated with the jet energy resolution (JER) was determined by the application of a standard ATLAS software package. A low-EMF version of this systematic was also checked. The width of the distributions of asymmetry and response were compared between data and simulation in different bins of EMF, and the result was extrapolated to low-EMF. However, the size of this uncertainty was found to be negligible compared to the regular JER, and therefore was omitted.

Trigger efficiency. To pass the CalRatio trigger at the HLT, events were required to satisfy selections on three main variables: E_T of the leading jet; $\log_{10}(E_H/E_{\text{EM}})$ of the leading jet; and the number of tracks with $p_T > 2$ GeV within $\Delta R < 0.2$ of the jet axis. If there was mismodelling of any of these trigger variables between the HLT- and the offline-reconstructed quantities, then this could affect the trigger efficiency for signal events entering the analysis.

To evaluate the size of this uncertainty, a tag-and-probe technique using b -jet triggers was used to obtain a pure sample of multijet events in both data and MC. Using these events, a parametric scale factor was derived which represented the degree of mismodelling in each variable. The scale factor was applied to each variable in an emulation of the CalRatio trigger. The change in yield with respect to the nominal (unscaled) trigger after the full analysis selection was applied was taken as the size of the systematic uncertainty.

BDT modelling systematic. An additional systematic was included to account for potential mismodelling of the 19 per-jet and per-event BDT input variables excluding the per-jet BDT scores. This was calculated using the same control sample of dijet events used in the calculation of the ABCD systematic uncertainty. In this dijet control sample, the distributions of the variables were compared between data and MC. For each variable, the typical distance of the data/MC ratio from one as a function of the value of the variable, v , was

taken as an approximation of the variation for that variable, $\sigma(v)$. The discrepancies were translated into a systematic uncertainty using a sampling procedure. For a benchmark signal model from each m_Φ scenario, a large number of pseudo-experiments were generated. During each pseudo-experiment, a set of 19 random numbers (one for each input variable) was drawn from a Gaussian distribution of mean 0 and width $1/\sqrt{19}^\dagger$. The signal efficiency was recalculated by processing the signal sample in question and re-evaluating the per-jet BDT (for every clean jet) and the per-event BDT, but with each of the variables shifted by its individual random number multiplied by its $\sigma(v)$. The relative difference of this efficiency with respect to the nominal signal efficiency was then calculated. The distribution of these relative differences for all pseudo-experiments was determined, and the standard deviation of this distribution was taken as the BDT modelling uncertainty for that m_Φ scenario. The final uncertainties obtained with this method are shown in Table 4.8.

Pile-up uncertainty. The MC is re-weighted to describe the pile-up in data using a centrally-provided ATLAS tool, by scaling the pile-up in MC by the standard recommendation of $1.09^{+0.09}_{-0.09}$. The systematic uncertainty associated with this re-weighting was evaluated with the standard ATLAS procedure of varying this scale factor within its errors. The uncertainties in the final yields given by this tool are quoted in Table 4.8.

Theory uncertainties

Parton distribution functions. The NNPDF2.3LO PDF set was used to generate using the signal samples. The corresponding uncertainty was obtained by taking the standard deviation of the 100 set of weights, propagated through the analysis.

4.6 Statistical interpretation

The data-driven background estimation and signal hypothesis test was performed simultaneously in all four regions of the ABCD plane. As can be seen in Table 4.7, no excess of

[†]The width of the Gaussians was chosen because this procedure corresponds to sampling a 19-dimensional multivariate normal distribution. Drawing from Gaussians of width 1 would result in over-estimation of the uncertainty, since the radius of the contour which contains 68% of the distribution (corresponding to 1σ in the 1-dimensional case) scales approximately with the square root of the number of dimensions of the distribution. Sampling from Gaussian distributions of mean 0 and width $1/\sqrt{19}$ instead leads to a 68% containment contour at approximately 1σ , as desired.

events was observed in region A for either of the analysis selections. The CL_s method^[96] was therefore used to set upper limits on $\sigma(\Phi) \times B_{\Phi \rightarrow ss}$ in the benchmark model. The expected limits were obtained by setting the observed number of events in region A to the number expected using the ABCD relation (effectively blinding region A), while the observed limits used the unblinded number.

Since each signal sample was generated for a particular LLP proper decay length[†] ($c\tau_{\text{gen}}$), it was necessary to extrapolate the signal efficiency to other decay lengths to obtain limits as a function of $c\tau$. This was achieved by using a weighting method, applied separately to each signal sample, and described below. The upper limit at a given $c\tau$ was then obtained by scaling the limit at $c\tau_{\text{gen}}$ by the ratio of signal efficiencies at $c\tau$ and $c\tau_{\text{gen}}$.

4.6.1 Efficiency extrapolation

Generating a very large number of high-statistics samples spanning many orders of magnitude in $c\tau_{\text{gen}}$ was of course not feasible, so it was necessary to extrapolate the efficiency of the analysis to proper decay lengths other than those with which the MC samples presented in Table 4.2 were generated.

The extrapolation was performed starting from the simulated samples with lab-frame decay length of 5m (the *reference samples*). To extrapolate the signal efficiencies as a function of decay length, the existing samples were weighted to mimic different decay lengths of interest $c\tau_{\text{new}}$. The weight assigned to any given displaced jet i associated to an LLP generated with a decay length of $c\tau_{\text{gen}}$ is given as:

$$w_i(ct_i) = \frac{c\tau_{\text{gen}}}{\exp(-ct_i/c\tau_{\text{gen}})} \cdot \frac{\exp(-ct_i/c\tau_{\text{new}})}{c\tau_{\text{new}}}, \quad (4.2)$$

where the first factor is used to flatten the decay length distribution and the second factor is used to weight to the desired decay length. The quantity t_i is the proper decay time of the LLP that gives rise to displaced jet i and is calculated from the mass and momentum of the LLP. In the model considered the LLPs are pair produced, so the proper treatment required two weights to be applied per event, one for each LLP, using:

$$w'_i(ct_1, ct_2) = w_i(ct_1) \cdot w_i(ct_2). \quad (4.3)$$

[†]Individual LLP decay lengths in each sample are thus distributed following $\exp(-ct/c\tau_{\text{gen}})$.




Figure 4.17: The extrapolated signal efficiencies as a function of proper decay length of the S for all simulated samples in the low- E_T (left) and high- E_T (right) selections. The vertical error bars represent the statistical uncertainties.

When calculating the efficiency at $c\tau_{\text{new}}$, this weight was applied to all events passing the selection,

$$\epsilon_{\text{new}} = \frac{\sum_{i=\text{selected}} w'_i}{\sum_{\text{all}}} \quad (4.4)$$

where the summation in the numerator is over events that passed the selection and the summation in the denominator is over all events in the reference sample. This process was repeated for a range of values of $c\tau_{\text{new}}$ (including $c\tau_{\text{gen}}$) to produce the extrapolated signal efficiency curves in Figure 4.17.

4.6.2 Limit setting procedure

4.6.2.1 The likelihood-based ABCD method

The likelihood function used to derive the limit at one proper decay length is the product of four likelihoods, one for each region A, B, C, and D:

$$L(n_A, n_B, n_C, n_D | \mu, \theta_\mu) = \prod_{i=A,B,C,D} \frac{e^{-N_i} N_i^{n_i}}{n_i!} \quad (4.5)$$

where n_A, n_B, n_C , and n_D are the four observables that denote the number of events observed in each region in data. The signal strength, μ , is the ratio between the number of observed and expected events. A set of nuisance parameters, θ_μ , describes the systematic uncertainties. These are tied together by the definitions of N_i , which are the fit number of

events in each region (fit to the observed numbers, n_i , as well as possible):

$$N_A = \mu N_{A,expected}^S + N_A^B \quad (4.6)$$

$$N_B = \mu \epsilon_B N_{A,expected}^S + N_A^B \tau_B \quad (4.7)$$

$$N_C = \mu \epsilon_C N_{A,expected}^S + N_A^B \tau_C \quad (4.8)$$

$$N_D = \mu \epsilon_D N_{A,expected}^S + N_A^B \tau_B \tau_C \quad (4.9)$$

with N_A^S and N_A^B being estimated numbers of signal and background events in region A respectively, ϵ_i being the signal scaling factors derived from MC (such that $N_{B,expected}^S = \epsilon_B N_{A,expected}^S$ etc.) and τ_B and τ_C being the two nuisance parameters that relate the number of background events in region A to the other regions using the standard ABCD relation (i.e. $A = BC/D$).

4.6.2.2 Results

The hidden sector scalar models all have the proper decay length of the scalar as a free parameter, so the limits are presented as a function of LLP proper decay length. The observed and expected scaled limits for the low- m_Φ samples can be seen in Figures 4.18 and 4.19 and for the high- m_Φ samples in Figure 4.20. The lower bound for the range of lifetimes shown is defined by the point at which the efficiency curve for that particular signal model ceased to be visible in Figure 4.17. The observed limits are also summarised for each m_Φ in Figure 4.21.

4.6.2.3 Compatibility of limits with expected and observed background yields

It is worth noting that there is a subtlety associated with performing the ABCD background estimation and limit calculation in this way. In Figures 4.18–4.20 the observed limits typically appear to fall around the $\pm 1\sigma$ band of the expected limit in both the high- E_T and low- E_T analyses. However, when inspecting the unblinded yields in region A and comparing them to those predicted by the simple ABCD method in Table 4.7, a larger excess of events, of the order of 1.4σ and 2.5σ , appears to be present in the high- E_T and low- E_T analysis respectively.

This apparent contradiction is resolved by considering that the likelihood-based ABCD method considers signal and background simultaneously in the fit, and that there is more




Figure 4.18: The observed limits, expected limits and ± 1 and ± 2 sigma bands for the Φ masses of 125 GeV. The low- E_T selection was applied to the samples used to calculate these limits. Also shown are the SM Higgs boson cross-section for $m_H = 125$ GeV, assumed to be 48.58 pb at 13 TeV,^[97] and a comparison with the limits obtained for comparable models in the Run 1 analysis^[78] (scaled by the ratio of parton luminosities for gluon-gluon fusion between 13 TeV and 8 TeV for a particle of mass 125 GeV^[98]).




Figure 4.19: The observed limits, expected limits and ± 1 and 2 sigma bands for the Φ masses of 200 GeV. The low- E_T selection was applied to the samples used to calculate these limits.

information available to the fit after unblinding region A. In the likelihood-based ABCD method, the sum of signal and data events in a given region is constrained using Poisson distributions with rate parameter $\lambda = N^{\text{obs}}$ in each region. The expected background estimate in region A is obtained by running the fit with the signal strength, μ , set to be identically zero. The ABCD condition is enforced by the two parameters $\tau_B = N_B^{\text{exp}}/N_A^{\text{exp}}$ and $\tau_C = N_C^{\text{exp}}/N_A^{\text{exp}}$, by forcing N_D^{exp} to take the value $\tau_B \cdot \tau_C \cdot N_A^{\text{exp}}$.

The blinded background estimate is calculated by ignoring the number of events observed in region A (or equivalently, omitting the Poisson constraint on region A from the likelihood) when performing the fit. The values of the blinded background estimates were $6.2^{+3.2}_{-2.3}$ for the high- E_T selection and $2.5^{+2.5}_{-1.4}$ for the low- E_T selection. This is the nearest equivalent to the estimates obtained by applying the simple ABCD relation in Table 4.7, with which these values agree closely although with larger uncertainties since the ABCD uncertainty is now accounted for.

After unblinding, the background estimate must be re-evaluated to account for the actual observed data in region A; this is the unblinded estimate. In this case, the Poisson con-




Figure 4.20: The observed limits, expected limits and ± 1 and 2 sigma bands for the Φ masses of 400, 600 and 1000 GeV. The high- E_T selection was applied to the samples used to calculate these limits. For $m_\Phi = 600$ GeV, a comparison with the limits obtained in the Run 1 analysis^[78] is also shown (scaled by the ratio of parton luminosities for gluon-gluon fusion between 13 TeV and 8 TeV for a particle of mass 600 GeV^[98]).




Figure 4.21: The observed limits for the Φ masses of 125, 200, 400, 600 and 1000 GeV.

Table 4.10: The estimated background in region A, calculated in three ways. First, using the simple ABCD relation $N_A = N_B \times N_C / N_D$, as in Tables 4.4 and 4.7. In this case, the statistical uncertainties are simply propagated to the final value. Second, the background estimate calculated using the full statistical model, by setting the signal strength to zero, and ignoring the observed data in region A. This corresponds to the blinded background estimate and agrees well with the simple ABCD relation as expected, although the uncertainties are larger because the ABCD error is now accounted for. Finally, the unblinded background estimate calculated in the same way but including the observed data in region A and the Poisson constraint for that region.

Selection	Simple ABCD	Blinded	Unblinded
High- E_T selection	6.7 ± 2.3	$6.2^{+3.2}_{-2.3}$	$8.5^{+2.3}_{-2.0}$
Low- E_T selection	2.5 ± 1.8	$2.5^{+2.5}_{-1.4}$	$5.3^{+2.1}_{-1.6}$

straint on the number of events in region A is included in the likelihood. This allowed the fit to partially compensate for the excess of events with respect to the simple ABCD prediction by modifying the values of τ_B and τ_C within the range permitted by the Poisson constraints and other nuisance parameters. The values of the unblinded background estimates were $8.5^{+2.3}_{-2.0}$ for the high- E_T selection and $5.3^{+2.1}_{-1.6}$ for the low- E_T selection. The observed number of events in region A was approximately 1σ from these unblinded estimates in each case, which is consistent with what is observed in the limit plots, thus resolving the apparent discrepancy between Figures 4.18–4.20 and Table 4.7.

The blinded and unblinded background estimates for each selection are summarised in Table 4.10.

4.6.2.4 Limit calculation

To calculate the limits, an asymptotic approach was used to compute the CL_s value and the limits are defined by the region excluded at 95% confidence level. To confirm the validity of this method, the limit result was checked against the result obtained using a frequentist method based on toy Monte Carlo data. The results of this check are shown in Figure 4.22 for one low- m_Φ and one high- m_Φ sample, where it can be seen that the limits obtained from both methods agree well[†].

[†]In this case, the asymptotic approach was chosen over the frequentist approach because it was computationally faster.




Figure 4.22: Comparison of the observed and expected limits and $\pm 1\sigma$ bands obtained with the asymptotic method and the frequentist method for calculating the limit, for one low-mass and one high-mass sample. The two methods match very well, and therefore the asymptotic method (used to calculate all final limits) is valid.

4.6.2.5 Limit extrapolation

Once the limit had been determined at the sample's generated proper decay length using the fit described above, it was simply scaled to other lifetimes using the result of the efficiency extrapolation as follows,

$$(\sigma \times BR)_{new}^{95\%CL} = (\sigma \times BR)_{gen}^{95\%CL} \times \frac{n_A^{gen}}{n_A^{new}} \quad (4.10)$$

where *gen* and *new* are the generated and target proper decay lengths respectively, and n_A is the number of signal events in region A at a certain $c\tau$.

4.7 Discussion of results

In summary, the search for pair-produced LLPs decaying in the ATLAS calorimeter used data collected during pp collisions at the LHC in 2016. A benchmark hidden sector model was used to set limits, where the mediator's mass ranged between 125 and 1000 GeV, while the long-lived scalar's mass ranged between 5 and 400 GeV. The search selected events with two signal-like jets (which are typically narrow, trackless, and with a large fraction of their energy in the hadronic calorimeter) using machine learning techniques. Two signal regions were defined for the low- E_T and high- E_T datasets. The background estimation was performed using the data-driven ABCD method. No significant excess was observed in either signal region. The CL_s method was therefore used to set 95% CL limits on $\sigma_\Phi \times B_{\Phi \rightarrow SS}$ as a function of LLP decay length. For a mediator similar to the Higgs boson

and of mass $m_\Phi = 125$ GeV, decays of neutral LLPs with masses between 5 and 55 GeV were excluded for proper decay lengths between 5 cm and 5 m depending on the LLP mass (assuming a 10% branching ratio). For $m_\Phi = 200$ GeV, cross section times branching ratio values above 1 pb were ruled out between 5 cm and 7 m depending on the scalar mass. For models with $m_\Phi = 400$ GeV, $m_\Phi = 600$ GeV, and $m_\Phi = 1000$ GeV, values above 0.1 pb were ruled out between about 12 cm and 9 m, 7 cm and 20 m, and 4 cm and 35 m respectively, depending on the scalar masses.

Compared with the CalRatio search that was performed in Run 1,^[78] this analysis benefitted from the additional power of machine learning techniques for discriminating signal from background. This advantage was quantified by defining a set of selections to reproduce, as closely as possible, the analysis requirements from the Run 1 search. The $S/\sqrt{B}$ ratio in two representative signal samples for each of the low- E_T and high- E_T analyses was then compared between the emulated Run 1 selections and the Run 2 selections. For the Run 1 selection where the low- E_T trigger was used, the value of $S/\sqrt{B}$ at $c\tau_{\text{gen}}$ was 25.5, compared to the Run 2 low- E_T selection's value of 102. On the other hand, for the Run 1 selections based on the high- E_T trigger, the value of $S/\sqrt{B}$ at $c\tau_{\text{gen}}$ was 488, compared to the Run 2 high- E_T selection's value of 9591. Therefore the Run 2 CalRatio search, which benefitted from the use of machine learning techniques, performed around four times better in the case of the low- E_T analysis, and around twenty times better in the case of the high- E_T analysis. However it should be noted that the lower luminosity in the low- E_T data sample (10.8 fb^{-1}) meant that this difference is not reflected in the limits when comparing the Run 2 result with the Run 1 result at $c\tau_{\text{gen}}$.

This result may also be compared with other ATLAS searches for neutral LLPs decaying hadronically, which were described in §4.1.2. An additional analysis was completed following the publication of the CalRatio analysis,^[99] which used displaced muon triggers to search for one displaced vertex in the MS and one in the ID,^[100] in 33.0 fb^{-1} of Run 2 data. A summary of the proper decay lengths excluded in each of these searches is presented in Table 4.11, where two HS models considered by all the searches are used as a benchmark. The results from the analysis in this chapter were also combined with those from the Run 2 MS search, which is the subject of Chapter 5. A reinterpretation of the CalRatio search

Table 4.11: Proper decay lengths excluded by the ATLAS searches that considered the same HS model used in the Run 2 CalRatio analysis^[99] described in this chapter, for two signal models common to them all. The Run 1 MS+ID analysis used either a displaced muon or single-jet plus E_T^{miss} trigger to search for two displaced vertices in the ID, two displaced vertices in the MS or one displaced vertex in each of the ID and MS subdetectors.^[77] The Run 1 CalRatio search used a CalRatio trigger to select events with displaced jets originating in the HCal, via a cut-based analysis.^[78] The Run 2 MS analysis searched for either two displaced vertices in the MS, or one displaced vertex in the MS and E_T^{miss} .^[79] The Run 2 ID+MS analysis, completed after the Run 2 CalRatio analysis, used displaced muon triggers to search for one displaced vertex in the MS and one in the ID.^[100] The analyses are listed in the order they were published.

Analysis	Data	$m_\Phi = 125 \text{ GeV}$ $m_S = 25 \text{ GeV}$ SM Higgs BR 10%	$m_\Phi = 600 \text{ GeV}$ $m_S = 150 \text{ GeV}$ $\sigma \times B_{\Phi \rightarrow ss} = 1 \text{ pb}$
Run 1 CalRatio	$\sqrt{s} = 8 \text{ TeV}, 20.8 \text{ fb}^{-1}$	$c\tau = 0.4 - 8 \text{ m}$	$c\tau = 0.6 - 10 \text{ m}$
Run 1 MS+ID	$\sqrt{s} = 8 \text{ TeV}, 20.8 \text{ fb}^{-1}$	$c\tau = 0.4 - 20 \text{ m}$	$c\tau = 0.4 - 80 \text{ m}$
Run 2 MS	$\sqrt{s} = 13 \text{ TeV}, 36.0 \text{ fb}^{-1}$	$c\tau = 0.2 - 150 \text{ m}$	$c\tau = 0.5 - 60 \text{ m}$
Run 2 CalRatio	$\sqrt{s} = 13 \text{ TeV}, 10.8 \text{ or } 33.0 \text{ fb}^{-1}$	$c\tau = 0.16 - 3 \text{ m}$	$c\tau = 0.1 - 60 \text{ m}$
Run 2 ID+MS	$\sqrt{s} = 13 \text{ TeV}, 33.0 \text{ fb}^{-1}$	$c\tau = 0.12 - 6 \text{ m}$	$c\tau = 0.2 - 10 \text{ m}$

with respect to three further LLP signal models followed the publication of this search, and is presented in Chapter 6.

It is intended that the analysis described in this chapter, which used 2016 Run 2 data, will be improved upon in future iterations. Specifically, the next iteration will consider the full Run 2 dataset, amounting to almost 139 fb^{-1} . The per-jet identification described in §4.3.1 will be developed further as a standalone CalRatio jet tagging tool, which may introduce the possibility of considering other signatures for a full Run 2 search, such as a single CalRatio jet with prompt jet activity, or a single CalRatio jet with either a displaced vertex in the MS or E_T^{miss} . Additionally, the 2016 analysis was somewhat limited by the relatively high pile-up levels in the detector during Run 2, which particularly affected the signal and background discrimination at low- p_T (i.e. for low- m_Φ models). Therefore improvements are under development for reducing the effect of pile-up at the trigger level in Run 3 and beyond, since these low- m_Φ models are of special interest to the theory LLP community.^[101]

Chapter 5

Combination with search for displaced vertices in the muon spectrometer

In this chapter, the combination of the results from the CalRatio displaced jets search discussed in Chapter 4 with those from another ATLAS displaced jets search^[79] is presented. This analysis considered neutral LLP decays in the muon spectrometer (MS), with the same benchmark HS signal model as the CalRatio displaced jets search, so a combination of the results was performed to provide improved limits for the HS model overall. Since the selections used in the two searches were orthogonal and the same signal models were considered, a statistical combination of the limits was performed using a simultaneous fit of each analysis' likelihood function.

The two relevant MS search strategies are summarised in §5.1, and the orthogonality of the final selections in the MS and CalRatio searches are discussed in §5.2. The procedure for the combination is described in §5.3 and the results are presented in §5.4.

5.1 Search for displaced jets in the muon spectrometer

The MS displaced jets search considered three search strategies, but only two were applied to the HS model the CalRatio displaced jets search considered, and so only these two are described here. The first of these was designed to select two displaced vertices (2Vx); the second was designed to select a single displaced vertex (1Vx) with minimal additional re-

Table 5.1: A summary of the selections applied in the MS vertex search strategies described here.^[79]

Common baseline selections	
Event passes muon cluster trigger ^[53] Event passes cleaning criteria Event has a good primary vertex Event has 1 good MS vertex (Vx1) Vx1 matched to triggering cluster	
Two-vertex search selections	Single-vertex search selections
Second good MS vertex (Vx2) Vx2 matches 2nd triggering cluster (if present) Vx1 and Vx2 isolated	No second MS vertex Vx1 in fiducial volume Vx1 isolated $E_T^{\text{miss}} > 30 \text{ GeV}$ $ \Delta\phi(E_T^{\text{miss}}, Vx1) < 1.2$ Barrel: $n_{\text{hits}} > 1200$ Endcaps: $n_{\text{hits}} > 1500$ $n_{\text{tracklets}} \geq 5$

quirements. The results from the 2Vx search were better than those from the 1Vx search across the full proper decay length range for $m_\Phi > 125 \text{ GeV}$, whereas for $m_\Phi = 125 \text{ GeV}$ both searches gave complementary results and were therefore combined to give a single limit per m_S . The selections applied in each search strategy are briefly summarised in Table 5.1.

Both strategies considered the reconstruction of displaced vertices in the MS barrel and endcap regions separately, because both the muon cluster trigger and the vertex reconstruction had different requirements in each region. Data-driven background estimation methods were used, and the signal efficiency was extrapolated to other lifetimes using a toys-based method, different to that described in §4.6.1. For each LLP in the simulated sample, a random lifetime was generated, sampled from an exponential distribution with the target mean lifetime. The physical decay position in the detector was then calculated for each LLP using its four-momentum from the simulated sample and this lifetime. The overall probability that the event would satisfy the selection criteria was then evaluated from the efficiencies of each selection criterion, which were parameterised as a function of the LLP decay position. Limits were set for both search strategies using the CL_s method,

and in the case of $m_\Phi = 125$ GeV the limits from the 2Vx and 1Vx strategies were combined.

Two-vertex search

The first relevant strategy was an inclusive search for two isolated displaced vertices in the MS. This was designed to be sensitive for models where the LLP is pair-produced, as was the case in the HS model, and decays hadronically between the outer region of the HCal and the middle station of the MS. Requiring two displaced vertices significantly reduced the expected background[†].

A single search region was defined based on the selections listed in Table 5.1, and a data-driven method was used to estimate the total background contribution in this region (0.027 ± 0.011) which was compared to the observed number of events (0). Given the absence of an excess, 95% CL limits were set with the CL_s method, with a counting model based on Poisson statistics used as the likelihood function.^[79]

Single-vertex search

The second search strategy considered a single MS vertex with minimal additional requirements, in order to increase the sensitivity in the regime of long lifetimes. For models with two LLPs, the probability of both LLPs decaying inside the detector decreases for mean lab-frame decay lengths greater than ~ 5 m, so the sensitivity of a search for two displaced vertices scales as $1/(c\tau)^2$ at large $c\tau$. Thus, extending sensitivity to longer proper decay lengths for a given model also requires a strategy of using only one reconstructed displaced decay, for which the sensitivity instead scales as $1/(c\tau)$ at large $c\tau$. Such a strategy has unique sensitivity compared to other displaced searches, although it is affected by higher levels of background. Furthermore, the loose additional requirements allow for sensitivity to lower- p_T events—especially useful when considering the HS model with $m_\Phi = 125$ GeV, since the p_T of the Higgs boson is typically only $\mathcal{O}(10 \text{ GeV})$.

The likelihood-based ABCD method was used for background estimation in the barrel and endcap regions separately—the same method used in the CalRatio displaced jets search

[†]This is because the probability for the existence of a single displaced vertex from any background source is very small, and the existence of a second in the same event is uncorrelated and thus proportional to the square of an already very small number.

Table 5.2: Event counts in each of the four regions of the ABCD plane for the MS 1Vx search described here..^[79]

Region	A	Expected background	B	C	D
Barrel	224	243 ± 38 (stat.) ± 29 (syst.)	42	132,000	22,800
Endcaps	489	497 ± 51 (stat.) ± 30 (syst.)	94	165,800	31,390

(described in §4.4), but with a different plane definition. The results of the search are summarised in Table 5.2. No excess of events was seen in the search regions, so 95% CL limits were set with the CL_s method, using a similar likelihood function to the one described in §4.6.2.^[79]

The limits obtained for $m_\Phi = 125$ GeV from the 1Vx search were combined with those from the 2Vx search using the CL_s method on the product of the counting model based on Poisson statistics (from the 2Vx search) and the likelihood of the simultaneous fit (from the 1Vx search).

5.2 Orthogonality of searches

The orthogonality of the CalRatio and MS analyses was checked in both data and signal samples to ensure the final selections were statistically independent, and that a purely statistical combination of the results would therefore be valid.

In data, the run and event numbers of the events entering the analyses' search regions were compared, and no overlap was found.

The overlap in signal was studied by investigating the fraction of events passing the following requirements: (a) passing the CalRatio trigger with a clean jet matched to the triggering jet; (b) passing the MS trigger. The number of events which pass the CalRatio requirement that also pass the MS requirement, $n(b|a) = n(a \cap b)/n(a)$, was 8% and 1% in a high- m_Φ sample and low- m_Φ sample respectively. Meanwhile, the number of events which pass the MS requirement that also pass the CalRatio requirement, $n(a|b) = n(a \cap b)/n(b)$, was 0.8% and 0.03% in a high- m_Φ sample and low- m_Φ sample respectively. This overlap was small enough that a combination was possible. Furthermore, since these numbers were calculated before the full selection for each analysis, the final overlap was likely to be very much smaller.

5.3 Statistical combination procedure

The combination was performed using a simultaneous fit of the analyses' likelihood functions, which are defined as follows:

- *CalRatio (CR) search likelihood*: As described in §4.6.2, the likelihood function is constructed from the product of the Poisson probabilities of observing the number of events N_R^{obs} , given an expectation N_R^{exp} , in each region R , where $R = A, B, C, D$.
 - The value of N_R^{exp} in each region is the sum of the expected signal yield N_R^{sig} in region R multiplied by the signal strength μ (the parameter of interest), and the expected background yield N_R^{bkg} .
 - The background expectations are constrained by $N_A^{bkg} = (N_B^{bkg} \cdot N_C^{bkg})/N_D^{bkg}$.
 - Systematic uncertainties are represented by nuisance parameters, each assigned a Gaussian constraint of width (σ) corresponding to the size of the uncertainty.
- *Single-vertex (MS1) search likelihoods*: Two ABCD planes were used in this search (one for events in the barrel and one for events in the endcaps), and in each case the form of the likelihood function is identical to that of the CR likelihood.
- *Two-vertex (MS2) search likelihood*: The likelihood is constructed from the product of a Poisson constraint in the signal region with Gaussian constraints for the nuisance parameters.

In the simultaneous fit, the signal strength as well as the nuisance parameters for the theory and luminosity uncertainties are common between the CalRatio and MS1 and MS2 likelihoods. The other uncertainties were taken to be uncorrelated, and in order to validate this choice, it was confirmed that fully correlating the uncertainties produced a negligible effect on the final result compared to fully uncorrelating them.

5.4 Results

The combined limits are shown in Figures 5.1–5.5. For the models with $m_\Phi = 125$ GeV, the limits are presented divided by the SM Higgs boson cross-section, assumed to be 48.58 pb at 13 TeV.^[97]

The CalRatio limits for models with $m_\Phi = 125 \text{ GeV}$ [†] were combined with the results from the MS 1Vx and 2Vx searches. The MS 1Vx search did not set limits for models with $m_\Phi \geq 200 \text{ GeV}$ so the CalRatio limits for these models were combined with the MS 2Vx results only. As in the individual searches, an asymptotic approach was used to compute the CL_s value, and the limits were defined by the region excluded at 95% CL. The limits were calculated using a global fit, where the overall likelihood function was the product of the individual likelihood functions of the searches to be combined. In each case the signal efficiencies were scaled by the result of the lifetime extrapolations, and the limits were calculated separately at each point in the $c\tau$ range of interest.

In the case of $m_\Phi = 125 \text{ GeV}$, the MS analysis had higher sensitivity than the CalRatio analysis at large decay lengths. For short decay lengths ($< 10 \text{ cm}$) the sensitivities of both analyses are comparable and the combination of their limits slightly improves those of the MS-only analysis. The limits for intermediate masses, $m_\Phi = 200 \text{ GeV}$ and 400 GeV , show a clear complementarity of both analyses: the CalRatio limits, which improve as m_Φ increases, are stronger at shorter decay lengths while the MS search sets stronger limits at large decay lengths. In this case the combination improves the limits across the full decay length range. For higher masses, $m_\Phi \geq 600 \text{ GeV}$, the CalRatio analysis was more sensitive than the MS analysis in general. Even in this case, the combination provides a modest improvement to the CalRatio-only limit at long decay lengths.

The combined limits for $m_\Phi = 125 \text{ GeV}$ are also shown alongside the results of another ATLAS search for a similar process^[82] in Figure 5.6. This search considered exotic decays of the Higgs boson into a pair of scalar particles which decay into b-quarks promptly or with a mean proper decay length up to 5 mm , considering a process where the Higgs boson is produced in association with a vector boson to provide an object to trigger on in the challenging nearly-prompt decay length regime. The limits set by the displaced jet searches are stronger than those from the nearly-prompt jet search due to the lower background rates when considering decays far from the interaction point. However, the complementarity in decay length coverage between the two different types of searches can be clearly seen.

[†]There is no combined limit for the HS model with $m_\Phi = 125 \text{ GeV}$ and $m_S = 55 \text{ GeV}$ since it was not considered in the MS search.

In addition, following the publication of the CalRatio analysis^[99] which included the combined limits presented here, another analysis using displaced muon triggers to search for one displaced vertex in the MS and one in the ID in 33.0 fb^{-1} of Run 2 data was completed.^[100] The results of this analysis were also combined with the results from the CalRatio and MS displaced jets searches since the same HS signal models were studied. For $m_\Phi \leq 200 \text{ GeV}$, these combined limits improved further on those shown in this Chapter for proper decay lengths smaller than a few tens of centimetres. These combinations provide a useful summary of the overall ATLAS sensitivity to models with pair-produced, hadronically-decaying, neutral LLPs across the full range of decay lengths accessible to the ATLAS experiment.




Figure 5.1: Combined limits for models with $m_\phi = 125$ GeV from the CalRatio analysis (CR) and the MS analysis (MS1+MS2). The observed, expected and $\pm 1\sigma$ limits are shown in all three cases.




Figure 5.2: Combined limits for models with $m_\phi = 200$ GeV from the CalRatio analysis (CR) and the MS analysis (MS2). The observed, expected and $\pm 1\sigma$ limits are shown in all three cases.




Figure 5.3: Combined limits for models with $m_\phi = 400$ GeV from the CalRatio analysis (CR) and the MS analysis (MS2). The observed, expected and $\pm 1\sigma$ limits are shown in all three cases.




Figure 5.4: Combined limits for models with $m_\phi = 600$ GeV from the CalRatio analysis (CR) and the MS analysis (MS2). The observed, expected and $\pm 1\sigma$ limits are shown in all three cases.




Figure 5.5: Combined limits for models with $m_\phi = 1000$ GeV from the CalRatio analysis (CR) and the MS analysis (MS2). The observed, expected and $\pm 1\sigma$ limits are shown in all three cases.




Figure 5.6: Summary of combined limits from the CalRatio and MS analyses for models with $m_\phi = 125 \text{ GeV}$, shown alongside limits from another ATLAS search for exotic Higgs boson decays into a pair of scalar particles which decay into b-quarks promptly or with a short lifetime.^[82]

Chapter 6

Reinterpretation of search for displaced jets using RECAST

This chapter presents the integration of the search for displaced jets in the ATLAS hadronic calorimeter (described in Chapter 4) into the RECAST^[102] framework, and its immediate use therein to test alternate signal hypotheses. An overview of reinterpretation at the LHC is given in §6.1. A description of the RECAST framework is provided in §6.2, while §6.3 details the preservation of the CalRatio analysis workflow. The results are shown in §6.4, which presents new constraints for three signal models not previously considered in the analysis, and compares these to existing constraints from other ATLAS searches. A discussion of the future of the RECAST project is given in §6.5.

6.1 Reinterpretation

Experimental searches for BSM physics may be sensitive to a broad range of models. Usually, however, only a few models are considered in the interpretation of results from experimental searches. The number of models that can be considered for interpretation are only growing over time, and searches are a significant investment of experimental physicists' time and resources. In order to determine the implications of LHC data for as broad a range of theories as possible, experimental analyses should therefore be *reinterpretable* in terms of theories not considered in the original analysis publication.

Performing reinterpretations makes it possible for phenomenologists to give detailed feedback on the original analyses, and to better suggest promising new avenues for further

experimental analysis. In addition, ensuring the reinterpretability of a data analysis preserves its shelf life, and increases its scientific impact through the subsequent work derived from it. The preparation of the material to facilitate reinterpretation also invariably makes internal reproduction easier for experimental collaborations, for example after a student responsible for an analysis has moved to a different position or field. Finally, the provision of such data helps meet the increasingly stringent requirements from funders for openness and reproducibility of publicly-funded research.

In simple terms, a reinterpretation is interpreting an existing experimental result in the context of an alternative BSM physics scenario, by recording the data and backgrounds from an analysis and comparing these to a new signal model. As the data and backgrounds are unchanged, reinterpretations do not provide any scope for making a discovery, but rather allow limits to be set on models not considered in the original search. The methods for reinterpretation can be separated into two broad classes:

- **Unfolded reinterpretations:** Measured distributions from reconstructed physics objects are distorted by the finite resolution and limited acceptance of the detectors, and inefficiencies in the reconstruction; the transformation to the underlying true distribution by correcting the data and backgrounds for detector and reconstruction effects is called *unfolding*, shown in Figure 6.1. In an unfolded reinterpretation, the unfolded data and backgrounds are compared to the new signal model at particle-level.
- **Folded reinterpretations:** In this case, the signal model is *folded* either via a fast detector simulation or the full detector simulation and reconstruction, and compared to the reconstruction-level data and backgrounds obtained in the original search.

Unfolded reinterpretations avoid the need to pass the signal model through the computationally-expensive detector simulation, and unfolding is widely performed for SM measurements. For exotic searches, however, unfolding is difficult because the detector effects tend to be more important and non-regular, so they are not easily described by a simple convolution, for example. Unfolding also works best in high-statistics regions where the folding matrix can be learnt and inverted well, while exotics searches are often statistically-limited. This is the case in the search described in Chapter 4 so the reinter-




Figure 6.1: Schematic showing reconstruction and unfolding processes to convert between particle-, detector- and reconstruction-level information.

pretation of that search, described in this chapter, was a folded reinterpretation. There are two options for performing folded reinterpretations:

- Experiments can provide extensive information about the analysis details, such as parameterised efficiencies and acceptances, to phenomenologists. This allows the phenomenologists to quickly check a new model, and broadly assess interesting parameter space regions. Usually this is via a fast detector simulation or by smearing the physics objects using the parameterised efficiencies to go directly from particle-level to reconstruction-level information, as shown in Figure 6.1. These reinterpretations are necessarily performed after the information is publicised (for example, on HEPData^[103]) and carried out independently of the experimental collaboration.
- Experiments can preserve their analyses in full fidelity internally. Then reinterpretations may be performed on-demand for models that are deemed sufficiently promising, with the work done within or by the collaborations and made public after an approval process. This is slower than the previous option and is computationally more involved with the full detector simulation and reconstruction required to go from particle- to detector- to reconstruction-level information as shown in Figure 6.1. It produces a publication-quality official result by the experiment, however, where the selection of events and the statistical treatment of the result in the context of a new model are accurate.

It is important to note some challenges specific to the reinterpretation of LLP searches, and the search presented in Chapter 4 in particular, which make full-fidelity internal reinterpretations the only reliable option. Compared with BSM searches with promptly-

decaying particles, more information is needed to capture the effects of the non-standard physics objects used in LLP searches, such as displaced jets or highly-ionising tracks. Custom trigger and reconstruction techniques are often used to target these signatures, and efficiencies may be functions of several new and potentially model-dependent parameters (e.g. the trigger efficiency strongly depends on the reconstructed LLP decay position in the displaced jets analysis). Efficiency parameterisations published by LLP searches cannot therefore be used to accurately constrain models with event topologies sufficiently different from that of the original benchmark model. Fast detector simulation, although sufficient for simpler collider signatures, is also not reliable because seemingly slight deviations in the simulation could lead to substantial changes in the expected signal efficiency. Thus, without the resources available to collaboration members such as the full detector simulation, reinterpretation of an LLP search is often impossible.^[101] Moreover, the displaced jets search considered here used non-trivial selection criteria involving trained machine learning algorithms to distinguish signal and background events. These algorithms, partly based on detector-level variables, add further barriers to third-party reinterpretation frameworks such as Rivet^[104] since the particle-level selections on which they rely are not well defined. As a result, the only option available for the displaced jets search (and often other LLP searches) is to perform a full-fidelity internal reinterpretation.

Several ‘official’ reinterpretations have already been performed within the LHC experiments, but this is a very manual process. It is necessary to find people that know how to run the analysis, and have the time to do so. For exotics searches, and especially LLP searches, this is often a problem because the analysis teams are usually small. There is clearly a need for a system that can preserve the analyses once (at publication time) and re-execute them independently of the original collaboration members—the RECAST framework is intended to be such a system.

6.2 The RECAST framework

RECAST^[102] is a software framework that aims to provide the infrastructure necessary to organise systematic reinterpretations of analyses performed in high energy physics experiments. It is designed to reuse background estimates and observations in the data from searches to test alternative signal hypotheses using only a fraction of the effort required to

perform a new search. The analysis preservation strategies employed by RECAST have also widely influenced the overall CERN analysis preservation efforts.^[105]

In general, a search analysis attempts to identify a phase space region that is efficient for the signal model while rejecting as much SM background as possible. The definition of such a region through construction of suitable observables and development of the event selection constitutes a large fraction of the analysis work. The search region is examined for reconstructed data from the experiment. In addition, the expected contribution from BSM signal (and often SM background) processes must be determined—these processes must be simulated, reconstructed, and passed through the event selections to achieve this. The response to systematic variations of the simulation or reconstruction algorithms must also be determined. All these inputs then form the basis for a statistical analysis procedure. This general analysis workflow is shown diagrammatically in Figure 6.2.

The resources required for a reinterpretation with RECAST are far fewer than those necessary for the construction of a new analysis. As the event selection of an analysis is fixed, the background estimates and observed data distributions do not change in the context of a reinterpretation. Only the simulation of the new BSM signal input and the execution of the analysis chain on this input is required, as shown in Figure 6.3. With this in mind, there is the possibility of the new signal model having an appreciable yield across multiple regions of the analysis beyond the intended signal region in the original analysis. This may affect the background estimates of the analysis and in turn, compromise the validity of the reinterpretation or influence the quoted sensitivity from it. Reinterpretation with the full analysis enables a complete evaluation of the magnitude of such signal contributions and even the ability to incorporate them consistently through a joint statistical analysis, thereby taking them into account properly. It is important to note, however, that prior to performing such a statistical interpretation, the extent of the signal contamination should be inspected carefully to evaluate its effect on the final result.




Figure 6.2: A full analysis workflow. Three main paths exist: the determination of the data counts observed in a phase space region of the ATLAS data, the SM background estimation, and the BSM signal estimation. The data path begins with the physics object reconstruction in the recorded data, followed by the event selection to obtain the observed data. Both the background and signal paths include the particle generation where the hard scattering process and subsequent hadronisation are simulated, the detector simulation, the physics object reconstruction and the event selection to obtain an estimate of background or signal. All three are then combined into a statistical model, from which an inference result such as a limit is derived.




Figure 6.3: A reinterpretation workflow. As the background and data yields are already computed by the original analysis, they can be read from an archive and only the signal path must be executed, followed by the statistical analysis. The signal processing is typically very fast as it is usually only computed at leading order and the sample sizes are smaller than those for e.g. QCD processes.

6.2.1 Analysis preservation in RECAST

Systematic reinterpretation hinges on preserving analyses in a way that allows new collaborators to re-execute them independently. This is achieved in RECAST through the use of declarative specifications and modern cloud computing technologies. A catalogue of preserved analyses that can be re-executed by the corresponding experimental collaborations for a new signal model may then provide the basis of a streamlined reinterpretation program within the high energy physics community.^[106]

For the remainder of this chapter, *analysis* will refer to the portion of the ATLAS data processing which is not handled by centrally-provided infrastructure. The generation of a new signal in the format required by the analysis (i.e. derived analysis data) may be performed using standard tools implemented within the experiments and thus is fairly independent of the analysis team. In contrast, the ability to execute event processing and statistical analysis stages are usually only available to the original authors of the analysis. ATLAS analyses to be implemented in RECAST begin with the processing of events provided by the collaboration in the Derived Analysis Object Data (DxAOD) format.^[107] The analysis *workflow* can be efficiently modelled as a directed acyclic graph, in which nodes represent individual processing steps and edges denote step dependencies. This workflow must be preserved so that it can be re-executed on new inputs. This means that the software for the individual processing steps must be archived, the correct usage of the software must be preserved in parameterised job templates, and the workflow graph structure must be recorded.

To be reusable, the software of an analysis must be preserved in a manner that is portable and deployable on a wide range of computing infrastructure. This is primarily achieved through software archival in the form of Docker container images.^[108,109] Container images provide the full root filesystem (i.e. all software dependencies) necessary for the execution of a given process, and provide a suitable balance between isolation and efficient deployment on diverse infrastructure. ATLAS provides suitable container base images, which hold the ATLAS analysis software release as well as its dependencies.^[110] The analysis release provides the necessary software libraries to read the ATLAS data files and implement calibrations of the selected physics objects and thus is a key dependency of the user analysis code. The analysis' Docker image is then built on top of this base image, and includes the code required for the event selection and any other processing steps.

Simply preserving the software environment of an analysis is not sufficient to enable its re-execution. A description is also needed to document the correct usage of the command line programs provided within these environments. As reinterpretation is concerned not only with reproducing a given result but also reusing the software to derive new results, this description must be parameterised to allow the specification of new input data. Thus, the individual processing steps of the analysis must not be static reproductions of the commands executed to obtain the original result, but rather be job templates with parameters that are referenced in a workflow specification. Given a new input, the templates can then be translated to produce concrete job descriptions. A job description therefore needs to provide three basic pieces of information: a description of the processing step, often in the form of a command line string; the container image to be used for the process; and an identification of the relevant outputs in order to assemble multiple steps into a workflow. In this case, yadage was used, with the job templates and the workflow structure itself both captured in YAML documents, a format chosen for its simplicity and ubiquity. The parameterised job templates specify the commands required to execute the user analysis code for performing specific tasks, such as physics object calibration, event selection, estimation of systematic uncertainties, and statistical inference. The workflow orchestrates the individual steps of the analysis by specifying individual steps using templates and their input parameters.

In summary, a reinterpretation using RECAST proceeds as follows:

- The original analysis is preserved using Docker containers for the software, job templates for individual processing steps, and a workflow to link the steps together.
- A new signal model is proposed for a reinterpretation, and the central ATLAS data management system generates the sample required by the analysis.
- The preserved analysis is re-executed on the new signal sample, and a new limit is produced.

6.3 Preservation of CalRatio displaced jets analysis

The CalRatio displaced jets search was expected to be sensitive to models containing displaced hadronic jets from neutral LLPs other than the one used as a benchmark, so the

analysis used in the Run 2 33.0 fb⁻¹ result^[99] was preserved in the RECAST framework as follows.

Software preservation

The software used in the CalRatio displaced jets analysis is organised across two separate software repositories; thus, two container images were produced. The first repository (called *DiVertAnalysisCode*) provides a command line interface to process and calibrate the event information extracted from centrally-produced DxAOD files and create ROOT^[111] files containing the necessary information for the remainder of the analysis. The second repository (named *CalRatioExtrapolation*) contains code to apply the high- E_T and low- E_T selections of the CalRatio displaced jets analysis, as well as the code to perform the statistical analysis and lifetime extrapolation. For both repositories, the image building is incorporated into the continuous integration configuration provided by the GitLab^[112] version control platform such that the software is preserved automatically each time the source code is modified. This feature was useful while the preserved analysis software was still under development to remove the need for any hard-coded information.

Processing steps and workflow preservation

The workflow for the CalRatio displaced jets analysis consists of nine individual steps: one data processing step, followed by two parallel streams of four steps (event selection, systematic calculation, extrapolation and limit setting), for the low- E_T and high- E_T selections. These are described in more detail below.

Data processing: The data processing step handles the calibration of physics objects, the calculation of the MLP-predicted decay positions and BDT scores, and the variation of parameters in the simulated events to account for systematic uncertainties. This step is performed within the *DiVertAnalysisCode* container image described above. The main input to the processing step is a DxAOD file, and the output is a single ROOT TTree containing the information required for the remainder of the analysis.

Event selection: Two event selection steps apply either the high- E_T or low- E_T selection to a signal sample, within the *CalRatioExtrapolation* container image described above. They

each result in a file containing the key information required for the later analysis stages, including the distribution of signal events in the ABCD plane used for the background estimation, and the truth-level variables required to perform the extrapolation as a function of decay length.

Extrapolation: A further step performs the lifetime extrapolation. This uses the distribution of signal events in the ABCD plane and their properties to extrapolate the analysis efficiency as a function of the long-lived particle decay length with the function described in §4.6.1. This stage outputs the extrapolated efficiency curves in each of the four ABCD regions, which are used to scale the final limit to lifetimes other than the one the sample was generated with.

Systematics: Two more stages evaluate the systematic uncertainties related to the signal models by applying either the high- or low- E_T selection to the systematically-varied branches in the ROOT TTree. The uncertainties calculated within these stages are the standard jet energy resolution and scale uncertainties (*JER+JES*), the re-derived jet energy scale uncertainty for CalRatio jets (*CalRatio JES*), the uncertainty in the additional number of interactions per event (*Pileup*), the trigger uncertainties and the uncertainties from the signal model generation (*Theory*). These are calculated according to the methods used in the original analysis, described in §4.5. The uncertainties relating to the luminosity and the background estimation method are not affected by changing the signal model considered and therefore were hard-coded in the step definition itself. It should be noted that the uncertainty involving the modelling of the ML input variables was not recalculated as part of the workflow. The original procedure to calculate this systematic required re-running large parts of the analysis workflow and performing a sampling technique, described in §4.5 which was deemed to be unnecessary to implement in this workflow, especially in light of the fact that its value was always very small ($<2\%$) compared to the size of the total systematic uncertainty. A conservative 5% was therefore assigned for all new models instead.

Limit setting: The final stage performs the statistical analysis to calculate the final cross-section limits. This involves constructing the modified ABCD method likelihood described in §4.6.2.1. In order to build the new likelihood, the (archived) observed data




Figure 6.4: A schematic of the workflow used in the analysis preservation. Individual steps (in blue) are defined with dependencies from other steps (indicated by the arrows).

yields in the four ABCD regions for both low- E_T and high- E_T selections are stored in the job template itself, and can then be combined with the newly-derived signal estimates and their uncertainties to set upper limits on the signal cross-section times branching ratio.

Similar to the job templates, the full workflow is specified in a YAML document. The workflow references the job templates and declares the dependencies between them. The structure of this workflow is shown diagrammatically in Figure 6.4, where the dependencies between steps are indicated by the arrows and the low- E_T and high- E_T selections are denoted by `_high` and `_low` respectively. A more detailed schematic, showing the required inputs for each step and the results of each step, is shown in Figure 6.5.

6.3.1 Validation with signal of original analysis

The workflow was validated by running on signal samples used in the original analysis. Two samples were chosen to perform the validation (one for each of the low- E_T and high- E_T analyses), and the results from these were compared with those from the original published search.^[99] Figure 6.6 show the validation results. The original limit at $c\tau_{\text{gen}}$ was reproduced to within 0.15σ with the RECAST workflow.




Figure 6.5: A more detailed view of the workflow used in the analysis preservation. Individual steps (blue boxes) have inputs (red ellipses) which are either set initially by the user, or result from previous steps. The arrows indicate the direction of information flow between the steps. The initialisation step is denoted by the diamond at the top of the figure, while `_high` and `_low` refer to the high- and low- E_T selections respectively.




Figure 6.6: Comparison of the original published limits (red) with those obtained from running the preserved analysis workflow (blue) with the low- E_T selection for $m_\phi = 200 \text{ GeV}$, $m_S = 8 \text{ GeV}$ and the high- E_T selection for $m_\phi = 600 \text{ GeV}$, $m_S = 150 \text{ GeV}$. In each case, the limit was extrapolated from the generated LLP decay length ($c\tau_{\text{gen}}$) indicated by the dashed vertical line.

6.4 Constraints on new models

With the CalRatio displaced jets analysis preserved as described in §6.3, and given that no significant excesses were observed in data in either of the analysis' signal regions, cross-section limits can be produced relatively simply for other benchmark models. New cross-section limits on three benchmark models probed in recent ATLAS papers, and which were described in Chapter 2, are shown in the following sections. In all cases, the high- E_T analysis selection gave a better sensitivity (overall, or in the regions where this search was complementary to the existing constraints) and so only the results from the high- E_T selection are shown.

For each new signal model considered, the signal contamination in the B, C, and D regions was studied to ensure the simultaneous fit to background and signal during the limit setting procedure would not break down in cases where the signal was distributed significantly differently to that in the original analysis. It was determined that the level of signal contamination was acceptable in all cases and that the simultaneous fit was robust against this.

6.4.1 Stealth SUSY model results

The first model considered was a Stealth SUSY model, as described in §2.4.3.1. The masses and decay lengths considered in this model are detailed in Table 6.1. Previously, the viability of this SUSY model was tested by the search for displaced hadronic jets in the ATLAS muon spectrometer.^[79] This search also constrained the HS model originally considered by the CalRatio analysis, for which a combination of the limits was performed as described in Chapter 5. The results of this combination proved that the two searches were sensitive to similar topologies in complementary decay length regions, so a reinterpretation of the CalRatio jets search was also expected to provide good sensitivity for the other models considered in the muon spectrometer search.

In the original search, referred to as the *MS search*, limits were set using a two-vertex signal region (2Vx) or the statistical combination of this region with a one-vertex plus additional activity signal region (*Comb*) that improved the sensitivity for longer lifetimes. Compared with the final state of the HS model previously considered in the CalRatio displaced jets

Table 6.1: Summary of the masses and proper decay lengths ($c\tau$) simulated for the Stealth SUSY model.

$m_{\tilde{g}}$ [GeV]	$m_{\tilde{S}}$ [GeV]	m_S [GeV]	$c\tau_{\tilde{S}}$ [m]
250	100	90	0.96
500			0.76
800			0.62
1200			0.50
1500			0.45
2000			0.37

Table 6.2: The dominant systematic uncertainties on the Stealth SUSY model shown for all generated mass points. In all cases, the trigger uncertainties were negligible in size compared to the total, and the BDT input variable modelling uncertainty was assigned a value of 5%. The background estimation systematic error from the original analysis was 22% in the high- E_T signal region.

$m_{\tilde{g}}$ [GeV]	JER+JES	CalRatio JES	Pileup	Theory	Total
250	1.9%	3.8%	3.2%	8.4%	11.3%
500	1.3%	1.6%	2.9%	12.5%	14.0%
800	1.3%	0.7%	1.9%	16.3%	17.3%
1200	1.3%	0.3%	1.6%	20.3%	21.1%
1500	1.4%	0.1%	1.0%	22.7%	23.4%
2000	1.3%	0.1%	1.3%	25.8%	26.4%

search, the final state of the Stealth SUSY model includes an additional two prompt jets in each event. Nevertheless, the similarity of the experimental signatures made this an excellent candidate model to be constrained by the preserved CalRatio displaced jets search.

The systematic uncertainties evaluated within the preserved workflow are given in Table 6.2. The 95% CL limits on $\sigma_{\tilde{g}\tilde{g}} \times B_{\tilde{g} \rightarrow g\tilde{S}}$ as a function of LLP decay length obtained from the reinterpretation are presented and compared to the previous results from the MS search in Figures 6.7 and 6.8 for the six gluino masses considered in the Stealth SUSY model.

In all cases, the existing constraints were improved and extended at shorter decay lengths, with the size of the improvement depending on the gluino mass. The excluded singlino $c\tau$ ranges are summarised in Table 6.3. The limits obtained are complementary to the




Figure 6.7: A comparison of the constraints for the Stealth SUSY model with $m_{\tilde{g}} = 250, 500$ and 800 GeV from the high- E_T dataset of the newly-preserved CalRatio jet analysis (pink/blue) and the results from the MS search (green/yellow). The MS search results are taken from the combination of both signal regions (Comb) for all masses, except $m_{\tilde{g}} = 250$ GeV where only the 2Vx region was used. Various $\sigma_{\tilde{g}\tilde{g}} \times B_{\tilde{g} \rightarrow g\tilde{S}}$ values are also indicated with dashed blue lines.




Figure 6.8: A comparison of the constraints for the Stealth SUSY model with $m_{\tilde{g}} = 1200, 1500$, and 2000 GeV from the high- E_T dataset of the newly-preserved CalRatio jet analysis (pink/blue) and the results from the MS search (green/yellow). The MS search results are taken from the combination of both signal regions (Comb). Various $\sigma_{\tilde{g}\tilde{g}} \times B_{\tilde{g} \rightarrow g\tilde{S}}$ values are also indicated with dashed blue lines.

Table 6.3: Ranges of singlino proper decay lengths excluded at 95% CL assuming various $\sigma_{\tilde{g}\tilde{g}} \times B_{\tilde{g} \rightarrow g\tilde{S}}$ values for the Stealth SUSY model^[113] following the reinterpretation of the CalRatio displaced jets search. The existing minimum and maximum decay lengths excluded by the MS search are denoted by $c\tau_{\min}^{\text{old}}$ and $c\tau_{\max}$. The maximum decay lengths excluded were unchanged following the reinterpretation, but the extended minimum decay lengths are indicated by $c\tau_{\min}^{\text{new}}$. A cell labelled as *n.s.* means *no sensitivity*.

$m_{\tilde{g}}$ [GeV]	B	$c\tau_{\min}^{\text{new}}$ [m]	$c\tau_{\min}^{\text{old}}$ [m]	$c\tau_{\max}$ [m]
250	0.01%	0.15	0.5	18
500	1%	0.08	0.15	64
800	10%	0.07	0.14	42
1200	100%	0.07	0.12	22
1500	100%	0.11	0.31	1.7
2000	100%	<i>n.s.</i>	<i>n.s.</i>	<i>n.s.</i>

existing results, and highlight the impact of the RECAST reinterpretation framework, since new regions of phase space are explored here at relatively little cost.

6.4.2 Higgs-portal baryogenesis model results

The next model considered was a Higgs-portal baryogenesis model, as described in §2.4.4.1. The masses and decay lengths considered in this model are detailed in Table 6.4, and four decay modes (and their charge conjugates) were simulated: $\chi \rightarrow \tau^+\tau^-\nu_\ell$, $\chi \rightarrow cbs$, $\chi \rightarrow \ell^\pm cb$, and $\chi \rightarrow \nu b\bar{b}$. As in the case of the Stealth SUSY model, the viability of this model was previously tested by the MS search,^[79] using either the 2Vx or Comb signal region. Again, the prior combination with the CalRatio analysis described in Chapter 5 showed that the two searches were sensitive to similar final states in complementary decay length regions. Together with the similarity of the signature of this model to the original HS model (i.e. two displaced jets), this complementarity made this another excellent candidate model to be constrained using the RECAST framework.

The systematic uncertainties evaluated within the preserved workflow are given in Table 6.5. The 95% CL limits on $\sigma_h \times B_{h \rightarrow \chi\chi}$ from the high- E_T dataset are compared to the published results from the MS search in Figures 6.9–6.12 for the four decay modes in the Higgs-portal baryogenesis model. The existing limits are improved (and extended) at shorter decay lengths. There were two exceptions for the $\chi \rightarrow \tau\tau\nu$ decay mode. The

Table 6.4: Summary of the masses and proper decay lengths ($c\tau$) simulated for the Higgs-portal baryogenesis model.

m_h [GeV]	m_χ [GeV]	$c\tau_\chi$ [m]
125	10	0.92
	30	2.75
	50	5.55
	100	3.30

Table 6.5: The dominant systematic uncertainties on the Higgs-portal baryogenesis model shown for all generated mass points, averaged over all four decay channels. In all cases, the trigger uncertainties were negligible in size compared to the total, and the BDT input variable modelling uncertainty was assigned a value of 5%. The background estimation systematic error from the original analysis was 22% in the high- E_T signal region.

m_χ [GeV]	JER+JES	CalRatio JES	Pileup	Theory	Total
10	10.6%	16.9%	3.0%	2.5%	20.9%
30	6.3%	20.4%	6.9%	2.5%	23.1%
50	5.4%	14.2%	8.8%	2.7%	18.5%
100	2.1%	10.2%	3.6%	4.5%	12.9%




Figure 6.9: A comparison of the constraints on the $\chi \rightarrow \tau\tau\nu$ decay mode of the Higgs-portal baryogenesis model from the high- E_T dataset of the preserved CalRatio displaced jet analysis (pink/blue) and the results from the MS search (green/yellow). The dashed blue lines indicate $\sigma_h \times B_{h \rightarrow \chi\chi}$ with 100% or 10% branching ratios and the SM Higgs gluon-gluon fusion production cross-section $\sigma_h = 48.58 \text{ pb}^{[97]}$ assumed for on-shell χ production, and $\sigma_h \times B_{h \rightarrow \chi\chi}(100\%) = 0.007 \text{ pb}$ assumed for off-shell production ($m_\chi = 100 \text{ GeV}$).




Figure 6.10: A comparison of the constraints on the $\chi \rightarrow \nu b \bar{b}$ decay mode of the Higgs-portal baryogenesis model from the high- E_T dataset of the preserved CalRatio displaced jet analysis (pink/blue) and the results from the MS search (green/yellow). The dashed blue lines indicate $\sigma_h \times B_{h \rightarrow \chi\chi}$ with 100% or 10% branching ratios and the SM Higgs gluon-gluon fusion production cross-section $\sigma_h = 48.58 \text{ pb}^{[97]}$ assumed for on-shell χ production, and $\sigma_h \times B_{h \rightarrow \chi\chi}(100\%) = 0.007 \text{ pb}$ assumed for off-shell production ($m_\chi = 100 \text{ GeV}$).

first is for $m_\chi = 10 \text{ GeV}$, where the existing limit was not improved upon. The second is for $m_\chi = 30 \text{ GeV}$, where the calculated systematic uncertainties were too high to be able to set a limit at all. The limit plots for these models are therefore not shown. The excluded $\chi c\tau$ ranges are summarised in Table 6.6. Again, the improved limits highlight the impact of the RECAST reinterpretation framework.




Figure 6.11: A comparison of the constraints on the $\chi \rightarrow cbs$ decay mode of the Higgs-portal baryogenesis model from the high- E_T dataset of the preserved CalRatio displaced jet analysis (pink/blue) and the results from the MS search (green/yellow). The dashed blue lines indicate $\sigma_h \times B_{h \rightarrow \chi\chi}$ with 100% or 10% branching ratios and the SM Higgs gluon-gluon fusion production cross-section $\sigma_h = 48.58 \text{ pb}^{[97]}$ assumed for on-shell χ production, and $\sigma_h \times B_{h \rightarrow \chi\chi}(100\%) = 0.007 \text{ pb}$ assumed for off-shell production ($m_\chi = 100 \text{ GeV}$). The fluctuation near $c\tau = 12 \text{ m}$ for the $m_\chi = 30 \text{ GeV}$ model is due to a statistical fluctuation involving one event with very high proper decay length passing the selection.




Figure 6.12: A comparison of the constraints on the $\chi \rightarrow \ell cb$ decay mode of the Higgs-portal baryogenesis model from the high- E_T dataset of the preserved CalRatio displaced jet analysis (pink/blue) and the results from the MS search (green/yellow). The dashed blue lines indicate $\sigma_h \times B_{h \rightarrow \chi\chi}$ with 100% or 10% branching ratios and the SM Higgs gluon-gluon fusion production cross-section $\sigma_h = 48.58$ pb^[97] assumed for on-shell χ production, and $\sigma_h \times B_{h \rightarrow \chi\chi}(100\%) = 0.007$ pb assumed for off-shell production ($m_\chi = 100$ GeV).

Table 6.6: The Higgs-portal baryogenesis χ proper decay lengths excluded at 95% CL for $\sigma_h \times B_{h \rightarrow \chi\chi}$, following the reinterpretation of the CalRatio displaced jets search. Branching ratios of 100% or 10% and the SM Higgs gluon-gluon fusion production cross-section $\sigma_h = 48.58 \text{ pb}^{[97]}$ are assumed for on-shell χ production, and $\sigma \times B_{h \rightarrow \chi\chi}(100\%) = 0.007 \text{ pb}$ is assumed for off-shell production ($m_\chi = 100 \text{ GeV}$). The existing minimum and maximum decay lengths excluded by the MS search are denoted by $c\tau_{\min}^{\text{old}}$ and $c\tau_{\max}$. The maximum decay lengths excluded are unchanged following the reinterpretation, but the extended minimum decay lengths are indicated by $c\tau_{\min}^{\text{new}}$. A cell labelled as *n.s.* means no sensitivity, and *n.e.* means that no extension of the excluded range was obtained from the reinterpretation.

Channel	m_χ [GeV]	$B_{h \rightarrow \chi\chi}$	$c\tau_{\min}^{\text{new}}$ [m]	$c\tau_{\min}^{\text{old}}$ [m]	$c\tau_{\max}$ [m]
$\chi \rightarrow \tau\tau\nu$	10	100%	<i>n.e.</i>	0.15	4.4
	30	100%	<i>n.s.</i>	0.66	11
	50	100%	0.4	1.3	18
	100	100%	<i>n.s.</i>	<i>n.s.</i>	<i>n.s.</i>
$\chi \rightarrow \nu b\bar{b}$	10	10%	0.05	0.08	31
	30	100%	0.1	0.18	371
	50	100%	0.2	0.45	421
	100	100%	<i>n.s.</i>	<i>n.s.</i>	<i>n.s.</i>
$\chi \rightarrow cbs$	10	10%	0.04	0.08	30
	30	10%	0.1	0.26	116
	50	100%	0.2	0.36	776
	100	100%	<i>n.s.</i>	<i>n.s.</i>	<i>n.s.</i>
$\chi \rightarrow \ell cb$	10	10%	0.03	0.09	24
	30	100%	0.09	0.17	457
	50	100%	0.3	0.45	426
	100	100%	<i>n.s.</i>	<i>n.s.</i>	<i>n.s.</i>

6.4.3 Dark photon model results

The final model considered in the reinterpretation is a hidden sector model described in §2.4.2.2, where the final state is two displaced lepton-jets (LJs) formed from the decay products of either two or four dark photons. Table 6.7 contains the simulated mass and lifetime parameters used in this dark photon model. This model was previously considered by a search for events compatible with the decay of an LLP into collimated leptons or light hadrons in the ATLAS detector.^[80] The search defined three separate search regions, designed to correspond to both LLPs decaying leptonically ($\mu\text{LJ}-\mu\text{LJ}$), both decaying hadronically ($\text{hLJ}-\text{hLJ}$), and for the mixed case ($\mu\text{LJ}-\text{hLJ}$). Below, the new constraints are compared to those obtained from the $\text{hLJ}-\text{hLJ}$ selection in the previous search, since this provides the fairest comparison of the different analyses' sensitivities. Table 6.8 presents the systematic uncertainties evaluated within the preserved workflow.

In addition to the limit comparisons, the extrapolated efficiencies for the low- and high- E_{T} selections are compared to those from the $\text{hLJ}-\text{hLJ}$ selection of the dedicated displaced

Table 6.7: Summary of the masses and proper decay lengths ($c\tau$) simulated for the dark photon model.

Process	m_H [GeV]	$m_{f_{d_2}}$ [GeV]	$m_{s_{d_1}}$ [GeV]	m_{γ_d} [GeV]	$c\tau_{\gamma_d}$ [m]
$H \rightarrow 2\gamma_d + X$	125	5.0	–	0.4	0.04923
	800				0.01176
$H \rightarrow 4\gamma_d + X$	125	5.0	2.0	0.4	0.08240
	800				0.02104

Table 6.8: The dominant systematic uncertainties on the dark photon model shown for all generated mass points. In all cases, the trigger uncertainties were negligible in size compared to the total, and the BDT input variable modelling uncertainty was assigned a value of 5%. The background estimation systematic error from the original analysis was 22% in the high- E_{T} signal region.

Process	m_H [GeV]	JER+JES	CalRatio JES	Pileup	Theory	Total
$H \rightarrow 2\gamma_d + X$	125	17.8%	12.7%	9.5%	2.6%	30.3%
	800	2.7%	6.1%	0.4%	7.2%	11.3%
$H \rightarrow 4\gamma_d + X$	125	16.6%	10.9%	7.8%	2.8%	27.7%
	800	3.0%	9.6%	5.3%	7.5%	14.8%




Figure 6.13: A comparison of the extrapolated efficiency from the preserved CalRatio displaced jet analysis for two values of m_H for the two dark photon final state (left) and the four dark photon final state (right), and for $m_H = 125$ GeV (top) and $m_H = 800$ GeV (bottom). Where applicable, these efficiencies are compared to those from the hLJ-hLJ selection in the dedicated displaced lepton-jets search. The shaded region represents the statistical uncertainty on the efficiency (calculated at $c\tau_{\text{gen}}$ and extrapolated to other lifetimes).

lepton-jets search (where applicable) in Figure 6.13 for the two values of m_H . In the case with $m_H = 125$ GeV, the dedicated analysis did not have enough sensitivity in the hLJ-hLJ channel to set limits for either the two or four dark photon final states. In the two dark photon final state, however, the dedicated analysis selected enough events to calculate the signal efficiency as a function of the dark photon lifetime. The signal efficiencies obtained in the dedicated search and the RECAST reinterpretation are therefore presented here to offer some comparison between the two results.

The archived data and background estimation in combination with the dark photon signal hypothesis were used to set 95% CL limits on $\sigma_H \times B_{H \rightarrow N\gamma_d + X}$ as a function of LLP decay length.

The constraints from the high- E_T dataset are compared to the results from the dedicated displaced lepton-jets search in Figure 6.14 for the two values of m_H , where applicable. The




Figure 6.14: The constraints from the preserved CalRatio displaced jet analysis (compared to those from the hLJ-hLJ selection of the dedicated displaced lepton-jets search where applicable) for two values of m_H for the two dark photon final state (left) and the four dark photon final state (right), and for $m_H = 125$ GeV (top) and $m_H = 800$ GeV (bottom). The dashed blue lines indicate $\sigma_H \times B_{H \rightarrow N\gamma_d + X}$ with 100% or 10% branching ratios and the SM Higgs gluon-gluon fusion production cross-section $\sigma_H = 48.58$ pb $^{[97]}$ assumed for $m_H = 125$ GeV, and $\sigma_H \times B_{H \rightarrow N\gamma_d + X}(100\%) = 5$ pb assumed for $m_H = 800$ GeV.

excluded γ_d $c\tau$ ranges are summarised in Tables 6.9 and 6.10. In the $m_H = 800$ GeV case, the limit from the high- E_T selection covers a region complementary to the existing constraints. The existing limits are especially improved and extended at shorter decay lengths. For $m_H = 125$ GeV, no limits were set in the dedicated displaced lepton-jets search for the hLJ-hLJ selection due to insufficient statistics and efficiency. The limits from this re-interpretation are therefore the first to be set in this channel using ATLAS data, serving as another example of the impact of the RECAST framework.

Table 6.9: The γ_d proper decay lengths excluded at 95% CL following the reinterpretation of the CalRatio displaced jets search for $\sigma_H \times B_{H \rightarrow N\gamma_d+X}$ assuming 100% branching ratio for the dark photon model with $m_H = 125$ GeV, where $\sigma_H = 48.58$ pb is the SM Higgs gluon-gluon fusion production cross-section.^[97] The minimum and maximum decay lengths excluded are denoted by $c\tau_{\min}$ and $c\tau_{\max}$.

Process	B	$c\tau_{\min}$ [m]	$c\tau_{\max}$ [m]
$H \rightarrow 2\gamma_d + X$	100%	0.001	0.05
$H \rightarrow 4\gamma_d + X$	100%	0.02	0.1

Table 6.10: Extension of the γ_d proper decay lengths excluded at 95% CL following the reinterpretation of the CalRatio displaced jets search, assuming $\sigma_H \times B_{H \rightarrow N\gamma_d+X} = 5$ pb, for the dark photon model with $m_H = 800$ GeV. The existing minimum and maximum decay lengths excluded using the hLJ-hLJ selection are denoted by $c\tau_{\min}^{\text{old}}$ and $c\tau_{\max}$. The maximum decay lengths excluded are unchanged following the reinterpretation, but the extended minimum decay lengths are indicated by $c\tau_{\min}^{\text{new}}$.

Process	$\sigma_H \times B_{H \rightarrow N\gamma_d+X}$ [pb]	$c\tau_{\min}^{\text{new}}$ [m]	$c\tau_{\min}^{\text{old}}$ [m]	$c\tau_{\max}$ [m]
$H \rightarrow 2\gamma_d + X$	5	0.00046	0.0073	1.298
$H \rightarrow 4\gamma_d + X$	5	0.0006	0.0136	0.231

6.5 Further work

The work described in this chapter represents a significant increase in the impact of the CalRatio displaced jets search, and the results and methodology were publicised to the wider LLP community via an ATLAS public note.^[114] There are, however, a number of improvements that could be made to this RECAST workflow:

- Providing the file required to correct the simulated number of interactions per collision for each signal as an input to the workflow (or producing this file as part of the workflow). Currently, the specific ones used are stored in the container image, and a ‘default’ file is available for other models, but this file is not guaranteed to always correctly account for pile-up.
- Including the functionality to specify which generator-level particles correspond to the LLPs, their decay products, and other particles in the model. This is required to correctly identify the LLPs so that the extrapolation of the analysis’ sensitivity can be correctly performed. Currently this is hard-coded in the DiVertAnalysisCode

container, which is clearly not an ideal solution when considering future reinterpretations of this analysis.

- Adding a workflow definition in which the analysis can be executed without the extrapolation, to obtain a limit for models where only a specific generated lifetime is of interest.
- A better implementation of the BDT mis-modelling uncertainty calculation procedure. All results from the current workflow include the caveat that this is fixed at 5%, which should be sufficient for reasonably similar signatures to the original model's, but may not be in cases where the signatures are very different.
- Avoiding the dependency on CERN Virtual Mounted File System (cvmfs) to run the containers. Currently, the DiVertAnalysisCode container requires that cvmfs is mounted on the system where it is run—necessary in this case because the software release used was not officially supported with a dedicated ATLAS software container. In contrast, the software releases used for full Run 2 searches are fully supported, so this will be resolved in future workflows.

The CalRatio analysis of 2016 data is already being extended to include the full Run 2 dataset, i.e. adding 2017 and 2018 data. In the ATLAS Exotics working group the preservation of analyses in RECAST is now a requirement for publication, so the full Run 2 CalRatio analysis will be designed in such a way as to facilitate accurate preservation with a view to reinterpretation from the outset. The improvements listed above will therefore not be made to the 2016 CalRatio analysis workflow, but will rather be considered when designing future workflows.

The full Run 2 CalRatio analysis has a number of improvements planned with respect to the 2016 analysis, as discussed in §4.7. The reinterpretations presented in this chapter have already influenced this full Run 2 CalRatio search strategy. It is clear from the results of the reinterpretation that the search was previously quite tailored to one specific signal model, especially in the case of the low- E_T analysis selection, where sensitivity to the models with low mediator masses was lacking. The machine learning algorithms for the low- E_T analysis were trained with the low- m_Φ signal samples only, in contrast to the high- E_T version of the algorithms which were trained on a mixture of high- and low- m_Φ models. Therefore,

a broader mixture of masses, and a greater variety of signal types, are being considered for the next iteration of the analysis' machine learning trainings.

This study has also indicated that there is potential to be sensitive to a wider range of signatures than previously thought, with overlaps in sensitivity between different LLP searches. This successful reinterpretation should influence ATLAS LLP searches in general, encouraging more sharing of signal models and greater collaboration between similar analyses, so that the efforts of what is often a very small number of analysers can be most efficiently used.

In addition, the sensitivity to new models demonstrated here could be further extended in future by preserving the statistical combination of the CalRatio displaced jets search with that performed in the ATLAS muon spectrometer,^[79] described in Chapter 5. More recently, a combination of a search for displaced vertices in the inner detector and muon spectrometer^[100] with the existing CalRatio and muon spectrometer search combination was performed; the preservation of this search could also extend the sensitivity to new models, especially at shorter lifetimes.

Chapter 7

Summary

Long-lived particles (LLPs) occur in many extensions to the Standard Model and may elude searches for new promptly-decaying particles. This thesis presented a search for neutral LLPs decaying into hadronic states in the calorimeter of the ATLAS detector at the LHC, using data collected in 2016 at a centre-of-mass energy of 13 TeV,^[99] described in Chapter 4. The signature targeted was two narrow, trackless jets with a high ratio of energy deposited in the HCal to energy deposited in the ECal (CalRatio displaced jets); a hidden sector model was used as a benchmark, where the mediator’s mass ranged between 125 and 1000 GeV, while the LLP’s mass ranged between 5 and 400 GeV. Two dedicated triggers and machine learning techniques were employed to identify candidate events containing displaced jets, and two search regions were defined to target models with high and low mediator masses separately. A data-driven estimate of the background in each of these regions was performed using the likelihood-based ABCD method, and limits were set on the production cross section times branching ratio as a function of the LLP decay length. Compared with the previous CalRatio displaced jets search performed in Run 1,^[78] this analysis extended and improved the limits as a function of LLP proper decay length, most notably for models with high-mass mediators.

The results from the CalRatio displaced jets search were also combined with those from an ATLAS MS displaced jets search,^[79] presented in Chapter 5. For models with lower-mass mediators, at large decay lengths the MS analysis had higher sensitivity than the CalRatio analysis, while at shorter decay lengths the sensitivities of both analyses were comparable and the combination of their limits slightly improved those of the MS-only analysis. For models with high-mass mediators, the CalRatio search was more sensitive than the

MS search in general. Even in this case, the combination provided a modest improvement to the limit set by the CalRatio search at long decay lengths. Since the publication of the CalRatio displaced jets search, another search for one displaced vertex in the MS and one in the ID was completed^[100] using the same benchmark signal model, and its results were combined with the results from the CalRatio and MS displaced jets searches. These combinations provide a useful summary of the overall ATLAS sensitivity to models with pair-produced, hadronically-decaying, neutral LLPs across the full range of decay lengths accessible to the ATLAS experiment.

As the time required to double the luminosity collected by LHC experiments increases to years, the importance of reliable and easy-to-use tools for reinterpretations is growing. In the case of LLP searches, the RECAST framework is critical as no other tool allows for high fidelity reinterpretations of published results. For a small fraction of the effort necessary to resurrect a published result, let alone to perform a new search, the RECAST framework was used to improve the scope and depth of the ATLAS search program, in the work described in Chapter 6. The preservation and subsequent reinterpretation of the CalRatio displaced jets search using the RECAST framework^[114] resulted in an accurate and efficient reinterpretation of the published result in terms of three new physics models not considered in the original search. A Stealth SUSY model and a Higgs-portal baryogenesis model were probed, both previously targeted by the MS displaced jets search.^[79] Additionally, a dark sector model predicting Higgs and heavy boson decays to hadrons via long-lived dark photons was studied, which had previously been considered in an ATLAS search for LLP decays to collimated leptons or light hadrons.^[80] In all three cases, the existing limits from the dedicated searches were extended to shorter decay lengths. This work has already influenced the next CalRatio search strategy which will consider the full Run 2 dataset, and has indicated that there is the potential for the search to be sensitive to a wider range of signatures than previously thought. This successful reinterpretation using the RECAST framework has been an example for other LLP searches to follow in future, encouraging more sharing of signal models and greater collaboration between similar analyses, thereby allowing the data recorded by the ATLAS experiment to remain useful for as long as possible.

Bibliography

- [1] ATLAS Collaboration, *Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC*, Phys. Lett. B **716** (2012) 1, arXiv: 1207.7214 (cit. on pp. 34, 38).
- [2] CMS Collaboration, *Observation of a new boson at a mass of 125 GeV with the CMS experiment at the LHC*, Phys. Lett. B **716** (2012) 30, arXiv: 1207.7235 (cit. on pp. 34, 38).
- [3] R. Davis Jr., D. S. Harmer and K. C. Hoffman, *Search for neutrinos from the sun*, Phys. Rev. Lett. **20** (1968) 1205 (cit. on p. 35).
- [4] SNO Collaboration, *Measurement of the rate of $\nu_e + d \rightarrow p + p + e^-$ interactions produced by 8B solar neutrinos at the Sudbury Neutrino Observatory*, Phys. Rev. Lett. **87** (2001) 71301, arXiv: 0106015 [nucl-ex] (cit. on p. 35).
- [5] LHCb collaboration, *Observation of $J/\psi\phi$ structures consistent with exotic states from amplitude analysis of $B^+ \rightarrow J/\psi\phi K^+$ decays*, Phys. Rev. Lett. **118** (2016) 22003, arXiv: 1606.07895 (cit. on p. 37).
- [6] LHCb collaboration, *Amplitude analysis of $B^+ \rightarrow J/\psi\phi K^+$ decays*, Phys. Rev. D **95** (2016) 012002, arXiv: 1606.07898 (cit. on p. 37).
- [7] BESIII Collaboration, *Observation of a charged charmoniumlike structure in $e^+e^- \rightarrow (D^*\bar{D}^*)^\pm\pi^\mp$ at $\sqrt{s} = 4.26$ GeV*, Phys. Rev. Lett. **112** (2014) 132001, arXiv: 1308.2760 (cit. on p. 37).
- [8] Belle Collaboration, *Study of $e^+e^- \rightarrow \pi^+\pi^-J/\psi$ and observation of a charged charmonium-like state at Belle*, Phys. Rev. Lett. **110** (2013) 252002, arXiv: 1304.0121 (cit. on p. 37).

- [9] LHCb Collaboration, *Observation of $J/\psi p$ resonances consistent with pentaquark states in $\Lambda_b^0 \rightarrow J/\psi K^- p$ decays*, Phys. Rev. Lett. **115** (2015) 72001, arXiv: 1507.03414 (cit. on p. 37).
- [10] M. Tanabashi (Particle Data Group), *Review of Particle Physics*, Phys. Rev. D **98** (2018) 030001 (cit. on pp. 37, 76).
- [11] L. A. Harland-Lang et al., *Parton distributions in the LHC era: MMHT 2014 PDFs*, Eur. Phys. J. C **75** (2015) 204, arXiv: 1412.3989v2 (cit. on p. 39).
- [12] F. Zwicky, *On the masses of nebulae and of clusters of nebulae*, Astrophysical Journal **86** (1937) 217 (cit. on p. 40).
- [13] D. Walsh, R. F. Carswell and R. J. Weymann, *0957 + 561 A, B: twin quasistellar objects or gravitational lens?*, Nature **279** (1979) 381 (cit. on p. 40).
- [14] D. Clowe et al., *A direct empirical proof of the existence of dark matter*, Astrophysical Journal **648** (2006) L109, arXiv: 0608407 [astro-ph] (cit. on p. 40).
- [15] G. R. Blumenthal et al., *Formation of galaxies and large-scale structure with cold dark matter*, Nature **311** (1984) 517 (cit. on p. 40).
- [16] Planck Collaboration, *Planck 2018 results. VI. Cosmological parameters*, 2018, arXiv: 1807.06209 (cit. on p. 40).
- [17] S. Chang, P. J. Fox and N. Weiner, *Naturalness and Higgs decays in the MSSM with a singlet*, JHEP **08** (2006) 68, arXiv: 0511250 (cit. on p. 43).
- [18] S. Chang et al., *Nonstandard Higgs boson decays*, Ann. Rev. Nucl. Part. Sci. **58** (2008) 75, arXiv: 0801.4554 (cit. on p. 43).
- [19] M. J. Strassler and K. M. Zurek, *Echoes of a hidden valley at hadron colliders*, Phys. Lett. B **651** (2007) 374, arXiv: 0604261 (cit. on p. 43).
- [20] M. J. Strassler and K. M. Zurek, *Discovering the Higgs through highly-displaced vertices*, Phys. Lett. B **661** (2008) 263, arXiv: 0605193 (cit. on p. 43).
- [21] D. Curtin et al., *Exotic decays of the 125 GeV Higgs boson*, Phys. Rev. D **90** (2014) 75004, arXiv: 1312.4992 (cit. on pp. 43, 45, 48).
- [22] A. Falkowski et al., *Hidden Higgs decaying to lepton jets*, JHEP **05** (2010) 77, arXiv: 1002.2952 (cit. on p. 44).

- [23] A. Falkowski et al., *Discovering Higgs decays to lepton jets at hadron colliders*, Phys. Rev. Lett. **105** (2010) 241801, arXiv: 1007.3496 (cit. on p. 44).
- [24] C. Cheung et al., *Lepton jets in (supersymmetric) electroweak processes*, JHEP **04** (2010) 116, arXiv: 0909.0290 (cit. on p. 44).
- [25] P. Meade, M. Papucci and T. Volansky, *Dark Matter sees the light*, JHEP **12** (2009) 52, arXiv: 0901.2925 (cit. on p. 44).
- [26] B. Batell, M. Pospelov and A. Ritz, *Probing a secluded $U(1)$ at B-factories*, Phys. Rev. D **79** (2009) 115008, arXiv: 0903.0363 (cit. on p. 44).
- [27] J. Fan, M. Reece and J. T. Ruderman, *Stealth supersymmetry*, JHEP **11** (2011) 12, arXiv: 1105.5135 (cit. on pp. 46, 47).
- [28] J. Fan, M. Reece and J. T. Ruderman, *A stealth supersymmetry sampler*, JHEP **07** (2012) 196, arXiv: 1201.4875 (cit. on p. 46).
- [29] Yu. A. Golfand and E. P. Likhtman, ‘Extension of the algebra of Poincare group generators and violation of P-invariance’, *Supergravities in Diverse Dimensions*, vol. 13, World Scientific Publishing Company, 1989 20 (cit. on p. 46).
- [30] D. V. Volkov and V. P. Akulov, *Is the neutrino a Goldstone particle?*, Phys. Lett. B **46** (1973) 109 (cit. on p. 46).
- [31] J. Wess and B. Zumino, *Supergauge transformations in four-dimensions*, Nucl. Phys. B **70** (1974) 39 (cit. on p. 46).
- [32] J. Wess and B. Zumino, *Supergauge invariant extension of quantum electrodynamics*, Nucl. Phys. B **78** (1974) 1 (cit. on p. 46).
- [33] S. Ferrara and B. Zumino, *Supergauge invariant Yang-Mills theories*, Nucl. Phys. B **79** (1974) 413 (cit. on p. 46).
- [34] A. Salam and J. A. Strathdee, *Supersymmetry and nonabelian gauges*, Phys. Lett. B **51** (1974) 353 (cit. on p. 46).
- [35] A. D. Sakharov, *Violation of CP Invariance, C asymmetry, and baryon asymmetry of the Universe*, JETP Lett. **5** (1967) 24 (cit. on p. 48).
- [36] C. S. Fong et al., *New ways to TeV scale leptogenesis*, JHEP **08** (2013) 104, arXiv: 1305.6312 (cit. on p. 48).

- [37] S. Blanchet et al., *Probing resonant leptogenesis at the LHC*, Phys. Rev. D **82** (2010) 76008, arXiv: 0904.2174 (cit. on p. 48).
- [38] H. An and Y. Zhang, *Direct detection of baryogenesis mechanism from squark decays at LHC*, Phys. Rev. D **89** (2014) 71902, arXiv: 1310.2608 (cit. on p. 48).
- [39] Y. Cui and B. Shuve, *Probing baryogenesis with displaced vertices at the LHC*, JHEP **02** (2015) 49, arXiv: 1409.6729 (cit. on p. 48).
- [40] ALICE Collaboration, *The ALICE experiment at the CERN LHC*, JINST **3** (2008) S08002 (cit. on p. 52).
- [41] CMS Collaboration, *The CMS experiment at the CERN LHC*, JINST **3** (2008) S08004 (cit. on p. 52).
- [42] LHCb Collaboration, *The LHCb detector at the LHC*, JINST **3** (2008) S08005 (cit. on p. 52).
- [43] ATLAS Collaboration, *The ATLAS experiment at the CERN Large Hadron Collider*, JINST **3** (2008) S08003 (cit. on pp. 52, 56, 57, 60, 61).
- [44] E. Mobs, *The CERN accelerator complex - August 2018*, 2018, URL: <https://cds.cern.ch/record/2636343> (cit. on p. 52).
- [45] L. Evans and P. Bryant, *LHC machine*, JINST **3** (2008) S08001 (cit. on p. 53).
- [46] G. Avoni et al., *The new LUCID-2 detector for luminosity measurement and monitoring in ATLAS*, JINST **13** (2018) P07017 (cit. on p. 54).
- [47] ATLAS Collaboration, *Luminosity determination in pp collisions at $\sqrt{s} = 8$ TeV using the ATLAS detector at the LHC*, Eur. Phys. J. C **76** (2016), arXiv: 1608.03953 (cit. on p. 54).
- [48] ATLAS Collaboration, *Luminosity public results Run 2*, URL: <https://twiki.cern.ch/twiki/bin/view/AtlasPublic/LuminosityPublicResultsRun2> (cit. on pp. 55, 56).
- [49] ATLAS Collaboration, *Track reconstruction performance of the ATLAS inner detector at $\sqrt{s} = 13$ TeV*, 2015, URL: <https://cds.cern.ch/record/2037683> (cit. on p. 58).
- [50] ATLAS Collaboration, *ATLAS Insertable B-Layer technical design report*, 2010, URL: <https://cds.cern.ch/record/1291633> (cit. on p. 58).

- [51] J. Pequeno, *Computer generated image of the ATLAS muons subsystem*, 2008, URL: <https://cds.cern.ch/record/1095929> (cit. on p. 62).
- [52] ATLAS Collaboration, *Muon reconstruction performance of the ATLAS detector in proton–proton collision data at $\sqrt{s} = 13$ TeV*, Eur. Phys. J. C **76** (2016) 292, arXiv: 1603.05598 (cit. on pp. 63, 71).
- [53] ATLAS Collaboration, *Triggers for displaced decays of long-lived neutral particles in the ATLAS detector*, JINST **8** (2013) P07015, arXiv: 1305.2284 (cit. on pp. 65, 86, 140).
- [54] ATLAS Collaboration, *Performance of the ATLAS track reconstruction algorithms in dense environments in LHC Run 2*, Eur. Phys. J. C **77** (2017) 673, arXiv: 1704.07983 (cit. on p. 68).
- [55] ATLAS Collaboration, *Reconstruction of primary vertices at the ATLAS experiment in Run 1 proton–proton collisions at the LHC*, Eur. Phys. J. C **77** (2017) 332, arXiv: 1611.10235 (cit. on p. 68).
- [56] ATLAS Collaboration, *Number of reconstructed vertices versus number of interactions in 2016 data and simulation*, 2016, URL: <https://atlas.web.cern.ch/Atlas/GROUPS/PHYSICS/PLOTS/IDTR-2016-013/> (cit. on p. 69).
- [57] ATLAS Collaboration, *Performance of the reconstruction of large impact parameter tracks in the inner detector of ATLAS*, 2017, URL: <https://cds.cern.ch/record/2275635> (cit. on p. 70).
- [58] ATLAS Collaboration, *Electron reconstruction and identification in the ATLAS experiment using the 2015 and 2016 LHC proton–proton collision data at $\sqrt{s} = 13$ TeV*, Eur. Phys. J. C **79** (2019) 639, arXiv: 1902.04655 (cit. on p. 71).
- [59] ATLAS Collaboration, *Measurement of the photon identification efficiencies with the ATLAS detector using LHC Run 2 data collected in 2015 and 2016*, Eur. Phys. J. C **79** (2018) 205, arXiv: 1810.05087 (cit. on p. 71).
- [60] M. Cacciari, G. P. Salam and G. Soyez, *The anti- k_t jet clustering algorithm*, JHEP **04** (2008) 63, arXiv: 0802.1189 (cit. on p. 73).
- [61] Y. L. Dokshitzer et al., *Better jet clustering algorithms*, JHEP **1997** (1997) 001, arXiv: 9707323v2 [hep-ph] (cit. on p. 73).

- [62] M. Wobisch and T. Wengler, ‘Hadronization corrections to jet cross sections in deep-inelastic scattering’, *Proceedings of the Workshop on Monte Carlo Generators for HERA Physics*, 1999, arXiv: 9907280v1 [hep-ph] (cit. on p. 73).
- [63] ATLAS Collaboration, *Jet kinematic distributions in proton–proton collisions at $\sqrt{s} = 900$ GeV with the ATLAS detector*, 2010, URL: <https://cds.cern.ch/record/1273173> (cit. on p. 74).
- [64] ATLAS Collaboration, *Data-quality requirements and event cleaning for jets and missing transverse energy reconstruction with the ATLAS detector in proton–proton collisions at a center-of-mass energy of $\sqrt{s} = 7$ TeV*, 2010, URL: <https://cds.cern.ch/record/1277678> (cit. on p. 74).
- [65] ATLAS Collaboration, *Reconstruction, energy calibration, and identification of hadronically decaying tau leptons in the ATLAS experiment for Run 2 of the LHC*, 2015, URL: <https://cds.cern.ch/record/2064383> (cit. on p. 76).
- [66] ATLAS Collaboration, *Identification and energy calibration of hadronically decaying tau leptons with the ATLAS experiment in pp collisions at $\sqrt{s} = 8$ TeV*, Eur. Phys. J. C 75 (2015) 303, arXiv: 1412.7086 (cit. on p. 76).
- [67] T. Gleisberg et al., *Event generation with SHERPA 1.1*, JHEP 02 (2009) 7, arXiv: 0811.4622 (cit. on p. 79).
- [68] ATLAS Collaboration, *The ATLAS simulation infrastructure*, Eur. Phys. J. C 7 (2010) 823, arXiv: 1005.4568 (cit. on pp. 79, 92).
- [69] S. Agostinelli et al., *GEANT4: A simulation toolkit*, Nucl. Instrum. Meth. A 506 (2003) 250 (cit. on pp. 79, 92).
- [70] ATLAS Collaboration, *The simulation principle and performance of the ATLAS fast calorimeter simulation FastCaloSim*, 2010, URL: <https://cds.cern.ch/record/1300517> (cit. on p. 79).
- [71] A. Haas et al., *Looking for milli-charged particles with a new experiment at the LHC*, Phys. Lett. B 746 (2015) 117, arXiv: 1410.6816 [hep-ph] (cit. on p. 81).
- [72] MoEDAL Collaboration, *The Physics Programme Of The MoEDAL Experiment At The LHC*, Int. J. Mod. Phys. A 29 (2014) 1430050, arXiv: 1405.7662 [hep-ph] (cit. on p. 81).

- [73] J. P. Chou, D. Curtin and H. J. Lubatti, *New Detectors to Explore the Lifetime Frontier*, Phys. Lett. B **767** (2017) 29, arXiv: 1606.06298 [hep-ph] (cit. on p. 81).
- [74] S. Alekhin et al., *A facility to Search for Hidden Particles at the CERN SPS: the SHiP physics case*, Rept. Prog. Phys. **79** (2016) 124201, arXiv: 1504.04855 [hep-ph] (cit. on p. 81).
- [75] V. V. Gligorov et al., *Searching for long-lived particles: a Compact Detector for Exotics at LHCb*, Phys. Rev. D **97** (2018) 15023, arXiv: 1708.09395 [hep-ph] (cit. on p. 81).
- [76] J. L. Feng et al., *ForwArd Search ExpeRiment at the LHC*, Phys. Rev. D **97** (2018) 35001, arXiv: 1708.09389 [hep-ph] (cit. on p. 81).
- [77] ATLAS Collaboration, *Search for long-lived, weakly interacting particles that decay to displaced hadronic jets in proton–proton collisions at $\sqrt{s} = 8$ TeV with the ATLAS detector*, Phys. Rev. D **92** (2015) 12010, arXiv: 1504.03634 [hep-ex] (cit. on pp. 84, 137).
- [78] ATLAS Collaboration, *Search for pair-produced long-lived neutral particles decaying to jets in the ATLAS hadronic calorimeter in pp collisions at $\sqrt{s} = 8$ TeV*, Phys. Lett. B **743** (2015) 15, arXiv: 1501.04020 (cit. on pp. 84, 130, 132, 136, 137, 181).
- [79] ATLAS Collaboration, *Search for long-lived particles produced in pp collisions at $\sqrt{s} = 13$ TeV that decay into displaced hadronic jets in the ATLAS muon spectrometer*, Phys. Rev. D **99** (2019) 052005, arXiv: 1811.07370 (cit. on pp. 85, 137, 139–142, 164, 168, 179, 181, 182).
- [80] ATLAS Collaboration, *Search for light long-lived neutral particles produced in pp collisions at $\sqrt{s} = 13$ TeV and decaying into collimated leptons or light hadrons with the ATLAS detector*, 5, 2020 450, arXiv: 1909.01246 (cit. on pp. 85, 174, 182).
- [81] ATLAS Collaboration, *Search for the production of a long-lived neutral particle decaying within the ATLAS hadronic calorimeter in association with a Z boson from pp collisions at $\sqrt{s} = 13$ TeV*, Phys. Rev. Lett. **122** (2019) 151801, arXiv: 1811.02542 [hep-ex] (cit. on p. 85).

- [82] ATLAS Collaboration, *Search for the Higgs boson produced in association with a vector boson and decaying into two spin-zero particles in the $H \rightarrow aa \rightarrow 4b$ channel in pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector*, JHEP **10** (2018) 31, arXiv: 1806.07355 [hep-ex] (cit. on pp. 85, 144, 149).
- [83] ATLAS Collaboration, *Studies of the performance of the ATLAS detector using cosmic-ray muons*, Eur. Phys. J. C **71** (2011) 1593, arXiv: 1011.6665 (cit. on p. 87).
- [84] J. Alwall et al., *The automated computation of tree-level and next-to-leading order differential cross sections, and their matching to parton shower simulations*, JHEP **07** (2014) 79, arXiv: 1405.0301 (cit. on p. 91).
- [85] R. D. Ball et al., *Parton distributions with LHC data*, Nucl. Phys. B **867** (2013) 244, arXiv: 1207.1303 (cit. on p. 91).
- [86] T. Sjöstrand et al., *An introduction to PYTHIA 8.2*, Comput. Phys. Commun. **191** (2015) 159, arXiv: 1410.3012 (cit. on p. 91).
- [87] ATLAS Collaboration, *ATLAS Pythia 8 tunes to 7 TeV data*, 2014, URL: <https://cds.cern.ch/record/1966419> (cit. on p. 91).
- [88] T. Sjöstrand, S. Mrenna and P. Z. Skands, *A brief introduction to PYTHIA 8.1*, Comput. Phys. Commun. **178** (2008) 852, arXiv: 0710.3820 (cit. on p. 91).
- [89] ATLAS Collaboration, *Summary of ATLAS Pythia 8 tunes*, 2012, URL: <https://cds.cern.ch/record/1474107> (cit. on p. 91).
- [90] A. D. Martin et al., *Parton distributions for the LHC*, Eur. Phys. J. C **63** (2009) 189, arXiv: 0901.0002 (cit. on p. 91).
- [91] A. Hoecker et al., *TMVA — Toolkit for multivariate data analysis*, 2007, arXiv: 0703039 [physics] (cit. on p. 95).
- [92] F. Pedregosa et al., *Scikit-learn: Machine learning in Python*, J. Mach. Learn. Res. **12** (2011) 2825, arXiv: 1201.0490 (cit. on p. 98).
- [93] ATLAS Collaboration, *Updated luminosity determination in pp collisions at $\sqrt{s} = 7$ TeV using the ATLAS detector*, 2011, URL: <https://cds.cern.ch/record/1334563> (cit. on p. 124).

- [94] ATLAS Collaboration, *Improved luminosity determination in pp collisions at $\sqrt{s} = 7$ TeV using the ATLAS detector at the LHC*, Eur. Phys. J. C **73** (2013) 2518, arXiv: 1302.4393 (cit. on p. 124).
- [95] ATLAS Collaboration, *In situ jet pseudorapidity intercalibration of the ATLAS detector using dijet events in $\sqrt{s} = 7 \sim 7$ TeV proton-proton 2011 data*, 2012, URL: <http://cds.cern.ch/record/1474490> (cit. on p. 124).
- [96] A. L. Read, *Presentation of search results: the CL_S technique*, J. Phys. G **28** (2002) 2693 (cit. on p. 127).
- [97] D. de Florian et al., *Handbook of LHC Higgs cross sections: 4. Deciphering the nature of the Higgs sector*, 2016, arXiv: 1610.07922 (cit. on pp. 130, 143, 169–173, 176, 177).
- [98] W. J. Stirling, *Private communication* (cit. on pp. 130, 132).
- [99] ATLAS Collaboration, *Search for long-lived neutral particles in pp collisions at $\sqrt{s} = 13$ TeV that decay into displaced hadronic jets in the ATLAS calorimeter*, Eur. Phys. J. C **79** (2019) 481, arXiv: 1902.03094 (cit. on pp. 136, 137, 145, 159, 161, 181).
- [100] ATLAS Collaboration, *Search for long-lived neutral particles produced in pp collisions at $\sqrt{s} = 13$ TeV decaying into displaced hadronic jets in the ATLAS inner detector and muon spectrometer*, Phys. Rev. D **101** (2020) 052013, arXiv: 1911.12575 (cit. on pp. 136, 137, 145, 179, 182).
- [101] J. Alimena et al., *Searching for long-lived particles beyond the Standard Model at the Large Hadron Collider*, 2019, arXiv: 1903.04497 (cit. on pp. 137, 154).
- [102] K. Cranmer and I. Yavin, *RECAST: Extending the impact of existing analyses*, JHEP **04** (2011) 38, arXiv: 1010.2506 (cit. on pp. 151, 154).
- [103] E. Maguire, L. Heinrich and G. Watt, *HEPData: a repository for high energy physics data*, J. Phys. Conf. Ser. **898** (2017) 102006, arXiv: 1704.05473 [hep-ex] (cit. on p. 153).
- [104] C. Bierlich et al., *Robust Independent Validation of Experiment and Theory: Rivet version 3*, SciPost Phys. **8** (2020) 026, arXiv: 1912.05451 (cit. on p. 154).
- [105] X. Chen et al., *Open is not enough*, Nature Phys. **15** (2018) 113 (cit. on p. 155).

- [106] K. Cranmer and L. Heinrich, *Analysis preservation and systematic reinterpretation within the ATLAS experiment*, J. Phys. Conf. Ser. **1085** (2018) 42011 (cit. on p. 157).
- [107] J. Catmore et al., *A new petabyte-scale data derivation framework for ATLAS*, J. Phys. Conf. Ser. **664** (2015) 72007 (cit. on p. 157).
- [108] D. Bernstein, *Containers and cloud: from LXC to Docker to Kubernetes*, IEEE Cloud Computing **1** (2014) 81 (cit. on p. 157).
- [109] D. Merkel, *Docker: lightweight Linux containers for consistent development and deployment*, Linux J. **2014** (2014) (cit. on p. 157).
- [110] L. Heinrich et al., *Analysis streamlining in ATLAS*, 2018, URL: <https://cds.cern.ch/record/2627651> (cit. on p. 157).
- [111] R. Brun and F. Rademakers, *ROOT — An object oriented data analysis framework*, Nucl. Instrum. Meth. A **389** (1997) 81 (cit. on p. 159).
- [112] J. Hethey, *GitLab repository management*, Packt Publishing, 2013 (cit. on p. 159).
- [113] C. Borschensky et al., *Squark and gluino production cross sections in pp collisions at $\sqrt{s} = 13, 14, 33$ and 100 TeV*, Eur. Phys. J. C **74** (2014) 3174, arXiv: 1407.5066 (cit. on p. 168).
- [114] ATLAS Collaboration, *Reinterpretation of the ATLAS search for displaced hadronic jets with the RECAST framework*, 2020, URL: <https://cds.cern.ch/record/2714064> (cit. on pp. 177, 182).