

A Search for Displaced Leptons from Long-Lived Particle Decays with the ATLAS
Detector

Andrew Caldon Smith

Submitted in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy
under the Executive Committee
of the Graduate School of Arts and Sciences

COLUMBIA UNIVERSITY

2024



© 2024

Andrew Caldon Smith

All Rights Reserved

ABSTRACT

A Search for Displaced Leptons from Long-Lived Particle Decays with the ATLAS
Detector

Andrew Caldon Smith

This thesis presents a search for leptons displaced from the primary vertex with the ATLAS detector at the Large Hadron Collider at CERN. The search includes the full proton–proton collision dataset collected during Run 2 from 2015–2018 at $\sqrt{s} = 13$ TeV and a partial dataset collected during Run 3 in 2022–2023 at $\sqrt{s} = 13.6$ TeV, corresponding to integrated luminosities of 140 fb^{-1} and 56.3 fb^{-1} , respectively. The search is the first performed using proton–proton collision data collected by the ATLAS detector during Run 3, at the highest collision energy ever achieved at a collider. Final states with displaced electrons or muons are considered, and novel triggers introduced in Run 3 are employed that use large impact parameter tracking to reconstruct displaced tracks with low momentum. In addition, multivariate techniques and timing information from the ATLAS electromagnetic calorimeter are employed to broaden the sensitivity to channels with large background rates or highly displaced electrons. The results are consistent with the Standard Model background expectations and are used to set model-independent limits on the production of displaced electrons and muons. The analysis is also interpreted in the context of a gauge-mediated supersymmetry breaking model with pair-produced long-lived sleptons. The results include 95% CL exclusions of selectrons with lifetimes from 4 ps to 60 ns and a mass of 150 GeV, and exclusions of selectrons, smuons, and staus with a lifetime of 0.3 ns for masses up to 740 GeV, 830 GeV, and 440 GeV, respectively.

TABLE OF CONTENTS

Acknowledgments	xiii
Preface	xv
Dedication	xix
Introduction	1
I THEORY	4
1 The Standard Model	5
1.1 The Particle Content of the Standard Model	5
1.2 Symmetries in the Standard Model	8
1.3 Limitations of the Standard Model	12
1.3.1 Matter-Antimatter Asymmetry	13
1.3.2 Gravity	13
1.3.3 Dark Matter and Dark Energy	14
1.3.4 The Hierarchy Problem	14
2 Beyond the Standard Model	16
2.1 Supersymmetry	17
2.2 Minimal Supersymmetric Standard Model	18
2.3 Gauge Mediated Supersymmetry Breaking	20

2.4	Long-Lived Particles	22
II	EXPERIMENT	25
3	The Large Hadron Collider	26
3.1	LHC Accelerator Concepts	26
3.2	The CERN Accelerator Complex	28
3.3	Operating Conditions in Run 2 and Run 3	30
4	The ATLAS Detector	33
4.1	ATLAS Coordinate System	34
4.2	ATLAS Subsystem Overview	36
4.3	Inner Detector	36
4.4	Calorimetry	39
4.4.1	Liquid Argon Electromagnetic Calorimeter	39
4.4.2	Liquid Argon Hadronic Endcap and Forward Calorimeter	41
4.4.3	Liquid Argon Readout Electronics	42
4.4.4	Liquid Argon Phase-I Upgrade	42
4.4.5	Tile Hadronic Calorimeter	45
4.5	Muon Spectrometer	46
4.6	Forward Detectors	48
4.7	Magnet System	49
4.8	Trigger and Data Acquisition	50
5	Object Reconstruction	54
5.1	Primary Vertex Reconstruction	56
5.2	Calorimeter Clustering	56
5.3	Tracking	57
5.3.1	Inside-out Tracking	58

5.3.2	Outside-in Tracking	59
5.3.3	Large Radius Tracking	59
5.4	Electrons and Photons	61
5.4.1	Superclustering	62
5.4.2	Photon Identification	63
5.4.3	Electron Identification	65
5.5	Muons	66
5.6	Modified Lepton Identification	67
5.7	Jets	68
5.8	Isolation	70
5.9	Overlap Removal	71
5.10	Missing Transverse Energy	72
6	Boosted Decision Trees and Timing Calibration	74
6.1	Boosted Decision Trees	75
6.2	Displaced Electron and Photon BDTs	77
6.3	Electron BDT Input Variables	78
6.4	Photon BDT Inputs and Training	79
6.5	BDT Performance	81
6.6	Liquid Argon Timing	83
6.6.1	Time Measurement	83
6.6.2	Timing Calibration Procedure	85
III	SEARCH	88
7	Data and Simulated Samples	89
7.1	ATLAS Dataset	89
7.2	Simulated Samples	91

7.3	Triggers	93
8	Analysis Strategy	97
8.1	Event Selection	98
8.1.1	Baseline Object Definitions	98
8.1.2	Event Categorization	99
8.2	ABCD Strategy	102
8.3	EM-BDT Strategy	105
8.3.1	EM-BDT Selections	106
8.3.2	EM-BDT Region Definitions	107
8.3.3	EM-BDT Time Shifting	110
8.3.4	EM-BDT Background Estimation	113
8.3.5	EM-BDT Validation	114
9	Systematic Uncertainties	118
9.1	ABCD Background Systematic Uncertainties	119
9.2	EM-BDT Background Systematic Uncertainties	120
9.3	Signal Systematic Uncertainties	120
9.3.1	Luminosity Uncertainty	121
9.3.2	Theory Uncertainties	121
9.3.3	Missing Transverse Energy Systematic Uncertainty	122
9.3.4	Time Smearing Systematic Uncertainty	122
9.3.5	Custom Scale Factors and Displacement Uncertainties	123
9.3.6	Lifetime Reweighting Systematic Uncertainty	125
9.4	Summary	127
10	Statistical Analysis and Fit Procedure	130
10.1	Statistics Concepts	130
10.2	EM-BDT Bin Optimization	133

10.3	EM-BDT Background Subtraction	134
10.4	EM-BDT Signal Injection Tests	136
10.5	Background-only Fits	136
11	Results	139
11.1	EM-BDT Slepton Results	139
11.2	ABCD Slepton Results	142
11.3	Model Independent Results	143
11.4	Dark Chargino Interpretation	144
12	Conclusion	148
	References	152
	APPENDICES	159
A	BDT Hyperparameter Optimization	160
B	Lifetime Reweighting	165
B.1	Two-Dimensional Lifetime Reweighting	165
B.2	Acceptance Loss for Long-Lifetime Decays	166
B.3	One-Dimensional Lifetime Reweighting	167
B.4	Assessment of Lifetime Reweighting Accuracy	168
C	TRT Wedge Veto	172

LIST OF FIGURES

1.1	Particle content of the Standard Model	6
1.2	Diagram of the Higgs boson potential	11
1.3	Standard Model cross sections measured by ATLAS	12
2.1	Gauge couplings in the Standard Model and MSSM	18
2.2	Slepton pair-production simplified model diagram	21
2.3	A selection of Standard Model particle masses and lifetimes	23
3.1	The CERN accelerator complex	29
3.2	Three-dimensional rendering of LHC dipole magnets	30
3.3	LHC Run 2 and early Run 3 pileup profile	31
3.4	High-pileup event display	31
4.1	The ATLAS detector	34
4.2	ATLAS coordinate geometry	35
4.3	The ATLAS inner detector system	37
4.4	The ATLAS calorimeters	40
4.5	LAr calorimeter accordion geometry	41
4.6	LAr calorimeter pulse, before and after shaping	43

4.7	LAr trigger tower/super cell comparison	44
4.8	LAr Phase-I readout electronics	44
4.9	LAr Phase-I trigger performance	45
4.10	Slice of the Tile calorimeter	46
4.11	The ATLAS muon system	47
4.12	The ATLAS magnet system	49
4.13	The ATLAS TDAQ system	52
4.14	ATLAS Run 3 trigger performance following the Phase-I upgrade	52
5.1	Reconstructed physics objects in ATLAS	55
5.2	Diagram of a track reconstructed in the ATLAS inner detector system	58
5.3	ATLAS tracking algorithms	59
5.4	Large Radius Tracking performance	61
5.5	Superclustering algorithm for electrons and photons	63
5.6	Example electron identification likelihood distribution	67
5.7	Large Radius Tracking efficiency for slepton LLPs	69
5.8	Diagram of hadronization process	71
6.1	Example of a decision tree for classifying flowers	76
6.2	Electron BDT training inputs	80
6.3	Photon BDT training inputs	81
6.4	electron and photon BDT response distributions for signal and background	84
6.5	Timing resolution following LAr timing calibration procedure	86

7.1	LHC Run 2 and early Run 3 integrated luminosity	90
7.2	Large Radius Tracking lepton trigger efficiencies	94
7.3	Venn diagram of selectron signal sorted by trigger	95
8.1	Flowchart of event categorization by analysis region	100
8.2	ABCD region definitions	102
8.3	ABCD cosmic muon veto	104
8.4	Signal-sensitive EM-BDT spatial and timing variables	109
8.5	EM-BDT region definitions	110
8.6	Energy-dependent timing correction	111
8.7	Region-dependent timing correction	112
8.8	Validation region timing distribution after correction procedure	112
8.9	Validation region kinematic comparisons	115
8.10	Electron and photon BDT shape validation	116
8.11	BDT inputs compared to Standard Model simulation	117
9.1	Visualization of EM-BDT shape uncertainty	121
9.2	Visualization of time-smearing systematic uncertainty	123
9.3	Electron displacement uncertainty	125
10.1	Bin optimization procedure example	134
10.2	Signal contamination in electron BDT CR	135
10.3	EM-BDT signal injection tests	137
10.4	EM-BDT fit validation in the VR	138

11.1	Score distributions in the EM-BDT SR	140
11.2	Exclusion contours separated by analysis region	141
11.3	One-dimensional exclusion profiles	142
11.4	Exclusion contours for slepton pair-production models	143
11.5	Dark chargino model diagram	145
11.6	Dark chargino interpretation	146
11.7	EM-BDT dark chargino yields	147
12.1	ATLAS stau exclusion summary	149
A.1	Learning rate overtraining check for electron BDT	161
A.2	Tree depth overtraining check for electron BDT	162
A.3	Electron BDT performance for selected depths and learning rates	162
A.4	Overtraining check for photon BDT	163
A.5	Electron BDT performance across a range of hyperparameters	163
A.6	Optimal electron BDT parameters chosen via hyperparameter scan	164
B.1	Two-dimensional lifetime reweighting example	166
B.2	Lifetime reweighting validation in ABCD regions	167
B.3	Acceptance loss in simulated decays as a function of $ z $	168
B.4	One-dimensional lifetime reweighting example	169
B.5	Yield comparisons for one-dimensional reweighting procedure	170
B.6	Lifetime-reweighted electron BDT distribution comparisons	171
C.1	Diagram of TRT, showing disabled wedge during 2023	173

C.2	Number of TRT hits per reconstructed track	174
C.3	Number of electrons in 2023 data passing electron BDT selections	174
C.4	Impact of disabled TRT wedge on electron BDT score distributions	175
C.5	Impact of disabled TRT wedge on selected electron BDT inputs	175

LIST OF TABLES

1.1	Relative strength of the fundamental forces	8
2.1	Field content of the Minimal Supersymmetric Standard Model	19
4.1	ATLAS subsystem specifications	36
4.2	LAr calorimeter granularity	40
4.3	ATLAS muon system technologies	48
5.1	Large Radius Tracking and standard tracking comparison	60
5.2	Variables used in electron and photon identification	64
6.1	Electron BDT variable ranking	83
6.2	Photon BDT variable ranking	83
6.3	Intervals of validity for the 2022-2023 LAr timing calibration	85
7.1	Slepton masses and lifetimes generated	92
8.1	Overview of electron, photon, and muon baseline selections	99
8.2	Table of triggers used in the analysis, by region	101
8.3	Summary of selections for electron BDT analysis regions.	107
8.4	Summary of selections for photon BDT analysis regions.	107

9.1	ABCD non-closure uncertainties	119
9.2	Electron BDT systematic uncertainty ranking	128
9.3	Photon BDT systematic uncertainty ranking	129
11.1	Discovery region observed yields	145

ACKNOWLEDGMENTS

First and foremost, I would like to thank my advisor Professor John Parsons, for his enduring support, insight, and kindness throughout my time in New York and my time at CERN. He has guided me deftly and patiently through difficult physics problems, professional decisions, and personal pursuits. Truly, John has been a great mentor, and the work in this thesis would never have been possible without him.

I would also like to thank Julia Gonski, my former post-doc, who first asked the question “What if we fit the BDT shape?” in the summer of 2023. Now we have an answer. Julia is an exemplary physicist and a role model to me and every student that has the chance to work with her.

Thank you to Nikiforos Nikiforou, also known as Nikipedia. The nickname is well-earned. Every time I speak with Nikiforos, I seem to learn something new, and following his advice has always set me down the right path.

I would like to thank Professor Gustaaf Brooijmans and Professor Mike Tuts for sharing their knowledge and asking insightful questions during my time in the Columbia group. Thank you to Kiley Kennedy, who encouraged me to apply to Columbia, and later invited me on ski trips, and whose shining example I have always done my best to emulate in my graduate work. Thank you to Elena Busch, my nearest neighbor in the Columbia ATLAS group, always one step ahead and always offering essential tips. Thank you to Eleanor Woodward and Maria Bressan, whose contributions were critical to the completion of this work. Thank you to all of the students of the Columbia ATLAS group for your cama-

raderie; I had a great time working with and getting to know each of you. Thank you to the staff of the Columbia Physics Department, who strive to make Pupin a welcoming community. Thanks as well to Amy Garwood, Grace Ho, and the staff at Nevis Laboratories, who, among other things, made my move to Switzerland and my move back to New York smooth and comfortable.

To the displaced leptons team, including Professor Ben Hooberman and the students at UIUC, thank you for your dedication to producing the first Run 3 ATLAS $p - p$ search result. Thank you to Jonathan Long for your hard work on this project with UIUC *and* your continuing work with Columbia to see this result through to publication.

To the Liquid Argon operations team, and especially to Clement Camincher, Adriana Milic, Nikiforos Nikiforou, and Emmanuel Monnier, thank you for your passion and attention to detail in training me to be a capable and confident Run Coordinator. I greatly admire your dedication in keeping the detector running optimally and efficiently through a major upgrade.

I would also like to thank Professors Kerstin Perez, Mike Shaevitz, Robert Mawhinney, and Tim Andeen for reviewing this thesis and serving on my defense committee.

I could not have completed this work without the unconditional support of my friends and family. To my friends in Geneva: Antti, Sagar, Lauren, Dev, Kevin, Laura, Tom, and others, thank you for the January and July swims, the skiing, the "Sunday Classic" hikes, and, of course, the lunch discussions. To Julia, Anthony, and Rob, the friends I would look forward to seeing after finishing up in the Nevis Annex for the night, thank you for being a constant source of support for me in New York. Thank you to Maria – I never struggled with any problem alone, in Geneva or in New York, because I knew that I could always call you for support. I used this fact to my advantage countless times, and I would not have completed this work without you. Thank you to Mom and Dad for always encouraging me to pursue my interests, even when I told you that I would be moving to Switzerland for a second time.

PREFACE

I began my work on the ATLAS experiment in June of 2018. Following my undergraduate studies at the University of Chicago, I was hired by Professor Jim Pilcher to spend a year at the Large Hadron Collider (LHC) at CERN, working on the operations teams for the ATLAS Tile hadronic calorimeter. A week after my graduation, I was on a plane to Geneva, Switzerland.

Through the end of Run 2, in the Fall of 2018, I was responsible for the Charge Injection System (CIS) for calibrating the Tile calorimeter. In the CIS calibration, known charges spanning the calorimeter's dynamic range are injected into so-called *3-in-1 cards*, which shape and amplify the analog signal in two gain scales for digitization by an Analog to Digital Converter (ADC). The slope of the line relating injected signal to ADC output is the CIS calibration constant, applied in physics data taking when converting digital outputs in ADC counts to energies in GeV. Approximately two calibrations are taken per week during data taking, with up to 48 3-in-1 cards in each of the 256 on-detector Tile detector drawers; I was tasked with monitoring these calibrations to ensure good linearity for each calibration constant and to flag any channels exhibiting significant drift over time. Alongside my colleague Kathryn Chapman, I presented my findings weekly and updated the CIS constants in the ATLAS configuration database when required.

I also contributed to the operations team during this time as a "shifter" in the ATLAS control room during data taking, monitoring data quality for the ATLAS calorimeters and contacting experts (occasionally in the very early hours of the morning) if I noticed a

potential issue in their subsystem. Devin Mahon and Alex Emerman, both Columbia students at the time, trained me in this role. I would end up calling Devin in the middle of the night for help with data-taking issues quite a few times.

While at CERN I started thinking about graduate school, and I would hound every student I met at CERN about their graduate program. Kiley Kennedy, whom I met on shift in the control room, and my hiking buddy Elena Busch both had great things to say about the Columbia ATLAS group, and I decided to reach out to Professor John Parsons for a coffee to learn more.

At the end of Run 2, the second Long Shutdown (LS2) of the LHC began. By this time, myself and other Tile detector experts had compiled a list of drawers that required access in LS2 due to malfunctioning photomultiplier tubes (PMTs), 3-in-1 cards, or powering issues. With a hard hat, boots, and a dosimeter, I would meet with my team every morning at 9 AM to drive to the ATLAS site, take an elevator 100 m underground, and start testing/replacing Tile detector electronics. I was watching a colleague open up a Tile detector drawer in the ATLAS cavern when I received the email that my application to Columbia had been accepted. In the summer of 2019, I said goodbye to CERN after replacing radiation-damaged cooling connectors on all 256 Tile drawers, and moved to New York to start working on the ATLAS Liquid Argon (LAr) electromagnetic calorimeter.

In early 2020, just as my coursework began to give way to a more active research role, New York City was put on lockdown. The start of my career at Nevis, as a contributor to the LAr Phase-II electronics upgrade for the High-Luminosity LHC (HL-LHC), would be remote. My first task was to evaluate the performance of our custom radiation-hard ADC chip, designed by engineering collaborators at Columbia and UT Austin, known as the COLUTA ADC. The COLUTA is designed to digitize LAr ionization pulses on upgraded Front End Boards (FEB2s) to be installed during the Long Shutdown following Run 3. The ADC is implemented in 65 nm CMOS technology, with specifications including an effective number of bits (ENOB) > 11 , a 40 MSPS sampling rate, and a ≥ 14 -bit dynamic

range. John had set up a camera in the lab, and a programmable switch, to allow me to remotely send signals from an Arbitrary Wave Generator (AWG), mimicking those we would expect in data-taking, to each of the eight channels on the chip soldered to our test board. I measured the energy and timing resolution of the chip, as well as cross-talk, ENOB, linearity, and noise. As the spring progressed into summer, I started making the trip to Nevis, sitting in a van one-person-per-row, diagonally, with all windows open and a mask on. With access to our ADC chips, I then developed Quality Assurance/Quality Control (QA/QC) testing suite by performing a standardized set of measurements on a set of chips one-by-one in our socket-equipped test board. I propagated the results of these tests to a webpage, which we used to select the best-performing chips for soldering to a "Slice" of a LAr FEB2 for integration testing. After the first Slice Testboard arrived at Nevis, with eight ADC chips and eight corresponding pre-amplifier/shaper chips from our collaborators implemented in the full density needed for the final FEB2 boards, I took part in integration activities including correlated noise measurements among all chips and timing and energy resolution of the integrated shaper/ADC system. My work at Nevis informed the final design of the ADC chip and the FEB2, which at the time of writing is in preparation for full-scale production.

In the fall of 2021, I moved back to Geneva to join the LAr operations team as we prepared for the end of LS2 and the start of Run 3. Dr. Nikiforos Nikiforou guided my integration into the commissioning activities for the Phase-I Digital Trigger (DT) upgrade to be deployed at the start of the run. I assisted with the calculation and uploading to the ATLAS configuration database of calibration constants for the Phase-I ADC developed at Nevis, and I validated the off-detector to on-detector fiber mapping for the newly installed DT electronics against the mapping provided at production. With help from LAr Data Quality experts, I developed monitoring tools for Run 3 shifters to use in assessing the data coming from the new DT electronics during proton collisions. During the ramp-up to Run 3, I was trained as a LAr detector on-call expert and eventually as the LAr Run

Coordinator (RC). I was training for the role in the control room during the first "beam splashes" of Run 3, where a circulating proton beam is smashed into a closed aperture near ATLAS, effectively lighting up the full detector in a spray of particles used to check that all ATLAS instruments, workflows, machines, and software run as expected. I was fortunate enough to analyze this first Run 3 data and produce a subset of the first LAr calorimeter public plots of the run, validating the timing measurements coming from the DT [1]. As LAr RC, for week-long shifts at a time, I oversaw all activities of the LAr operations team and coordinated with experts of other ATLAS subsystems to ensure smooth data taking as the DT system was phased in. Instead of placing calls as a shifter, as RC I was responding to the calls in the wee hours to address hardware failures. My work on the LAr operations team culminated in my presentation of the status of the Phase-II LAr Upgrade on behalf of the LAr collaboration at DIS 2022 in Santiago, Spain [2].

I began searching in earnest for a thesis topic in early 2023. I knew that I wanted to search for a new physics process Beyond the Standard Model (BSM), and I knew that I wanted this search to make extensive use of the LAr calorimeter I now knew so much about. In parallel to my hardware and operations work, I contributed to the thesis projects of two graduate students I had met in my first year at CERN, Devin Mahon and Kiley Kennedy. Both projects were searches for BSM Long-Lived Particles (LLPs) with decays resulting in displaced and delayed photons measured by the LAr calorimeter [3,4]. Through my work on these projects, I learned how one might exploit timing and spatial measurements from the LAr calorimeter to design an LLP search. Using the Run 3 data that I had helped to collect as LAr RC in 2022, I would apply the lessons I learned from these searches to the first ATLAS Run 3 BSM search, targeting displaced and delayed leptons from LLP decays. I had the privilege of presenting the result in this thesis on behalf of the ATLAS Collaboration at ICHEP 2024 in Prague, Czech Republic [5], and our paper has been submitted to *Physical Review D* [6].

To my parents.

INTRODUCTION

This thesis presents a search for new phenomena that are not described by the Standard Model (SM) of particle physics, performed using the ATLAS detector [7] at the Large Hadron Collider (LHC) [8]. The search includes the full proton–proton collision dataset collected by ATLAS during Run 2 at $\sqrt{s} = 13$ TeV and a partial dataset collected during Run 3 in 2022–2023 at $\sqrt{s} = 13.6$ TeV, corresponding to integrated luminosities of 140 fb^{-1} and 56.3 fb^{-1} , respectively.

The SM is a theory describing the properties and interactions of the known elementary particles which make up the universe. The theory has had great success in accurately predicting phenomena later confirmed by experiment, culminating in the discovery of the Higgs Boson by the ATLAS and CMS experiments [9,10] in July 2012. However, despite its success, the theory is known to be an incomplete description of the universe, motivating the search for beyond the Standard Model (BSM) physics. To address the shortcomings of the SM, physicists have hypothesized the existence of yet-undiscovered BSM particles. Many of the theories predicting these BSM particles can be tested using the data collected at the LHC. At the time of writing, no evidence of BSM phenomena has been observed at the LHC, and the ATLAS and CMS search programs have placed heavy constraints on theories involving so-called *prompt* decays of new particles. In prompt decays, the decay products are measured with spatial and time coordinates indistinguishable from those of the initial proton-proton collision. However, the constraints on a particular set of well-motivated theories involving Long-Lived Particles (LLPs) which can travel a measurable

distance before decaying, remain relatively weak due to the experimental challenges associated with identifying displaced decays.

The search presented is for LLPs decaying to leptons, with a focus on displaced electrons. The analysis also includes electrons which, as a result of their large displacement, are identified as photons by ATLAS. The search follows a search for displaced leptons published by ATLAS using the full Run 2 dataset [11], and this thesis therefore highlights new techniques applied to the partial Run 3 dataset which improve substantially upon the previous result. The capability of the ATLAS Liquid Argon (LAr) calorimeter to determine the delay of the LLP signature is used alongside a novel algorithm in the ATLAS tracker which can provide spatial measurements for highly displaced signatures. A Machine Learning (ML) algorithm known as a Boosted Decision Tree (BDT) is developed to optimally exploit these techniques, as well as other calorimeter and tracker measurements.

This thesis is organized as follows. Chapter 1 gives an overview of the SM, including notable shortcomings of the theory. Chapter 2 introduces Supersymmetry (SUSY) as a well-motivated class of BSM theories that extend the SM to address some of these shortcomings, and describes the pair-produced slepton LLP scenario targeted as a benchmark model in this search. Chapter 3 briefly describes some fundamental principals of accelerator physics, and gives some details about the LHC accelerator complex. The ATLAS detector and its subsystems used to collect the data are introduced in Chapter 4. The signatures of particles traversing ATLAS are subjected to a set of standardized algorithms to reconstruct the particles, as described in Chapter 5. Chapter 6 describes the specialized variables – the LAr timing measurement and the BDT score – crucial to the analysis strategy. An overview of the data and simulated samples used in the analysis is given in Chapter 7. The analysis strategy implementation is described in Chapter 8. Chapter 9 discusses the uncertainties applied to the measured particles and simulated samples used in this result. Chapter 10 outlines the method employed to arrive at a quantitative statistical statement about the observed data, and the tests used to validate this method. Chapter 11

interprets the results of the statistical procedure through the lens of the benchmark SUSY model being studied, and presents an interpretation of the results in the absence of a particular signal model. Chapter 12 presents the results in the context of the ATLAS search program for BSM signatures, and briefly touches on possibilities for future work.

PART I:

THEORY

CHAPTER 1:

THE STANDARD MODEL

The Standard Model (SM) of particle physics is a local gauge invariant quantum field theory, which describes the elementary particles constituting the universe and interactions between them, via exchange of force-carrying particles. The theory represents the most complete understanding of particle physics, and since its conception in the 1970s has consistently predicted the existence and properties of fundamental particles with excellent accuracy. Section 1.1 introduces the particle content of the SM, while Section 1.2 highlights some key properties of the theory. The theory encompasses three of the four known fundamental forces: electromagnetism, the strong force, and the weak force. Notably, the SM does not provide a mechanism for the gravitational force, a significant limitation of the model. Section 1.3 discusses limitations of the SM, including a number of phenomena the model cannot explain, which motivate the need for a more complete theory.

1.1 THE PARTICLE CONTENT OF THE STANDARD MODEL

The SM describes the universe as composed of 12 types of matter *fermions*, elementary particles with spin $1/2$, 4 force-carrying *bosons* with spin 1, and the spin-0 (scalar) *Higgs boson*. The particles comprising the SM, as well as some key attributes for each particle, are summarized in Figure 1.1.




Figure 1.1: Particle content of the SM [12]. Mass, charge, and spin are included for each particle, as well as color charge for those particles participating in the strong interaction.

The fermions can be classified based on the interactions they participate in, or the forces which affect them. All fermions experience the weak force, and fermions carrying electric charge experience the electromagnetic force. A key subdivision of fermions is into the *quarks*, which carry a color charge and experience the strong force, and the *leptons*, which do not carry a color charge and do not participate in the strong interactions. The fermions are categorized into three generations. Each generation contains a charged lepton and corresponding neutral lepton neutrino (ν_l), and two quarks of electric charge $+\frac{2}{3}e$, $-\frac{1}{3}e$ ($\frac{q+2/3}{q-1/3}$).¹ As an example, the first generation consists of an electron (e , electric charge $-e$), a massless electron neutrino (ν_e , electrically neutral), an up quark (u , electric charge

¹The electric charge here is expressed in units of $e = 1.602 \times 10^{-19}$ C, the smallest unit of charge which can freely exist in nature.

$+\frac{2}{3}e$) and a down quark (d , electric charge $-\frac{1}{3}e$). The massive particles of the second and third generations differ from their first-generation counterparts only in that they are more massive. Grouped based on increasing mass, the three lepton generations are $(\begin{smallmatrix} e \\ \nu_e \end{smallmatrix})$, $(\begin{smallmatrix} \mu \\ \nu_\mu \end{smallmatrix})$, $(\begin{smallmatrix} \tau \\ \nu_\tau \end{smallmatrix})$, and the three quark generations as $(\begin{smallmatrix} u \\ d \end{smallmatrix})$, $(\begin{smallmatrix} c \\ s \end{smallmatrix})$, $(\begin{smallmatrix} t \\ b \end{smallmatrix})$. The names are verbalized as electron, muon, tau, up, down, charm, strange, top, bottom. Each matter fermion has an anti-matter counterpart, which can be obtained via the action of charge inversion (denoted C, performed by reversing the sign of charge of a particle), parity inversion (denoted P, performed by reversing the sign of the spatial coordinates of a dynamical system) and time reversal (denoted T, performed by reversing the sign of the time evolution of a dynamical system). These discrete symmetries play an essential role in particle physics, and performing all three simultaneously on a system is known as CPT inversion. Each antimatter particle therefore has the same mass, but electric charge of equal magnitude and opposite sign, to its particle "partner". As a concrete example, the antimatter counterpart of the second generation lepton μ^- is the anti-muon μ^+ with electric charge of $+e$. These 12 matter particles and 12 antimatter particles together comprise the fermion content of the SM.

The fermions of interest in the search presented are the charged leptons. Of these, the electron is stable, while a muon decays to a muon neutrino, an electron and an electron antineutrino with a mean lifetime of $2.2 \mu\text{s}$. With 17% probability, a τ will decay to an electron and electron antineutrino, and it will decay to a muon and muon antineutrino with the same probability. The other 66% of τ decays are via the strong force and result in quarks [13].

The spin-1 bosons of the SM are the force-carrying bosons which mediate the strong, weak, and electromagnetic interactions of the fermions. The photon (γ) is a massless, electrically neutral particle which mediates the electromagnetic interaction. The gluon (g) is a massless, electrically neutral particle which mediates the strong interaction. Unlike the photon, which has no electric charge, the gluon carries the charge of the force it mediates. The non-zero color charge of the gluon significantly impacts the structure of the strong

interaction, as it allows for gluon self-interactions in addition to interactions between the gluon and color-carrying fermions. Photon self-interaction in the electromagnetic interaction is by contrast not possible because the photon is electrically neutral. The W bosons ($W^\pm$, electric charge $\pm e$) and the electrically neutral Z boson (Z) are massive bosons which mediate the weak interaction. The relative strength of the forces, as well as the strength of gravity, are listed in Table 1.1. Note that these strengths are only indicative as the force strengths depend on the distance and energy scale under consideration, and that the graviton is a hypothesized particle in Beyond-the-Standard-Model (BSM) theories of quantum gravity. Note also that gravity is vanishingly weak compared to the forces mediated by SM gauge bosons at the length scales relevant to the study of elementary particle physics.

Force	Relative Strength	Boson(s)	Spin	Mass/GeV
gravity	10^{-37}	graviton G (theorized)	2	0
weak	10^{-8}	$W^\pm, Z$ bosons	1	80.4, 91.2
electromagnetic	10^{-3}	Photon γ	1	0
strong	1	Gluon g	1	0

Table 1.1: Strengths of the fundamental forces, given as approximate indicative values for two fundamental particles at a distance of roughly the proton radius ($\sim 10^{-15}$ m). Associated gauge bosons and some key properties are also listed [14].

The particle content of SM is completed by the electrically neutral spin-0 Higgs boson(H), which generates the masses of the SM particles. The next section will discuss some properties of SM theory.

1.2 SYMMETRIES IN THE STANDARD MODEL

Noether's Theorem, first published by mathematician Emmy Noether in 1918, states that every continuous symmetry corresponds to a conserved quantity [15]. The study of gauge symmetries and their consequences for general renormalizable Lagrangians invariant under the Poincaré group (the set of transformations respecting the spacetime interval

between events) forms the foundation of the SM.² The SM is a quantum field theory respecting the symmetries of the local gauge group:

$$G_{SM} = SU(3)_C \times SU(2)_L \times U(1)_Y \quad (1.1)$$

Quantum Chromodynamics (QCD) is the SM theory of the strong interactions, corresponding to the $SU(3)_C$ group of Equation 1.1. The label C denotes that QCD is invariant under rotations in color space, and the three conserved quantities under QCD are the color charges r, g, b which label the orthogonal states in color space. The eight generators of the $SU(3)$ symmetry correspond to the gluon fields which mediate the interactions between gluons and quarks.

The evolution of the strong coupling constant α_s as a function of energy leads to a phenomena unique to quarks known as *color confinement*. As two color-charged particles are separated, it becomes energetically favorable for a new quark-antiquark pair to appear rather than continuing to separate the particles to a longer distance. As a consequence, quarks cannot be observed individually in nature and are instead only observed in energetically favorable color-neutral states. These states manifest as either even combinations of quarks and antiquarks ($q\bar{q}, qq\bar{q}\bar{q}, \dots$) or odd combinations of quarks or antiquarks ($qqq, \bar{q}\bar{q}\bar{q}, qq\bar{q}\bar{q}\bar{q}, \dots$), collectively called *hadrons*. The even combinations are called *mesons* and the odd *baryons*. A proton is an example of a hadron consisting of three quarks, hence the name Large Hadron Collider for the accelerator complex described in Chapter 3. The "jets" described in Section 5.7 are a direct consequence of color confinement.

The electromagnetic and weak forces of the SM are low-energy manifestations of a single unified *electroweak* force at high energies [14]. The electroweak sector of the SM is represented by $SU(2)_L \times U(1)_Y$. The three generators of the $SU(2)_L$ subgroup correspond to three vector gauge bosons W^1, W^2, W^3 , and the conserved quantity corresponding to the

²A Lagrangian is an equation which describes the dynamics of a physical system via calculation of equations of motion.

symmetry of the group is referred to as the third component of weak isospin, I_3 . All SM particles have a definite value of I_3 ; fermions with half-integer I_3 can participate in the charged-current weak interactions mediated by the $W^\pm$ boson while those with $I_3 = 0$ cannot. The L subscript reflects the fact that right-handed chiral (RH) particles (and left-handed (LH) chiral antiparticles) are unaffected by $SU(2)$ gauge transformations due to the maximally parity-violating $V - A$ structure of the weak interaction; thus the isospin doublets transforming under $SU(2)_L$ consist only of LH(RH) (anti-)particles.

The generator of the $U(1)_Y$ group, whose action preserves the quantity known as weak hypercharge Y , corresponds to the B boson. Adding explicit mass terms for the four bosons $W^{1,2,3}$ and B to the SM lagrangian would break the symmetry they generate and is therefore forbidden. In nature, however, two charged massive vector bosons $W^\pm$ and a neutral massive vector boson Z , as well as a massless boson (the photon, denoted γ) associated with the electroweak interaction have been experimentally observed.

At low energies the *Higgs mechanism* [16–19] spontaneously breaks the electroweak symmetry $SU(2)_L \times U(1)_Y \rightarrow U(1)_{\text{EM}}$, mixing the four massless vector boson fields and giving mass to the $W^\pm$, Z bosons. Upon symmetry breaking, weak hypercharge Y weak isospin I_W conservation are violated, while electric charge $Q = I_3 + \frac{Y}{2}$ remains conserved. The physical W , Z , and photon fields $W_\mu^\pm$, Z_μ , and A_μ are related to $W^{1,2,3}$ and B via the experimentally determined weak mixing angle θ_W as follows:

$$W_\mu^\pm = \frac{1}{\sqrt{2}}(W_\mu^1 \mp iW_\mu^2) \quad (1.2)$$

$$Z_\mu = -W_\mu^3 \sin \theta_W + B_\mu \cos \theta_W \quad (1.3)$$

$$A_\mu = W_\mu^3 \cos \theta_W + B_\mu \sin \theta_W \quad (1.4)$$

The coupling constants e, g_W for the electromagnetic and weak interactions, respectively, are related via the weak mixing angle as $e = g_W \sin \theta_W$. The weak mixing angle θ_W appearing in this equation can be written as $\tan \theta_W = \frac{g'}{g_W}$, where g' and g_W are the coupling constants corresponding to the strengths of the $U(1)_Y$ and $SU(2)_L$ gauge symmetries in the electroweak sector, respectively [14].

The Higgs mechanism is an example of spontaneous symmetry breaking, in which the complex scalar Higgs field ϕ acquires a non-zero vacuum expectation value (VEV) measured to be $v = 246$ GeV. This situation is shown diagrammatically in Figure 1.2. The field ϕ has four degrees of freedom, three of which are associated to the longitudinal polarizations of the physical $W^\pm$ and Z bosons, with the photon remaining massless. The remaining fourth degree of freedom is identified as an additional massive scalar boson known as the Higgs boson H . The mass of the Higgs boson has been measured experimentally by the ATLAS [9] and CMS [10] collaborations to be $m_H = 125$ GeV. The masses of the W and Z bosons are dependent on the VEV of the Higgs field, as $m_W = \frac{1}{2}g_W v$ and $m_Z = \frac{1}{2}v\sqrt{g_W^2 + g'^2}$ [14]. The masses of the fermions are also determined by their coupling to the Higgs field via their Yukawa coupling [20,21].




Figure 1.2: Diagram of the Higgs boson potential, $V(\phi)$, with $\phi = \frac{1}{\sqrt{2}}(\phi_1 + i\phi_2)$. The physical vacuum state of the field corresponds to a particular (non-zero) point on the dotted circle; by choosing this point the rotational symmetry of the field is broken [22].

1.3 LIMITATIONS OF THE STANDARD MODEL

The SM is laudable for its predictions of many physical phenomena, later confirmed to be both accurate and precise. As an example, the electron magnetic moment has been measured to agree with the SM prediction to within 1 part in 10^{12} [23]. The discovery of the charm quark [24,25], the top quark [26,27], the W and Z bosons [28,29], and the Higgs boson were all experimental confirmations of SM predictions. As shown in Figure 1.3, the ATLAS Collaboration at CERN has measured the production cross-section of an array of SM processes at multiple energy scales, with all results showing excellent agreement between experiment and theory.




Figure 1.3: A summary of a number of the SM cross sections measured by the ATLAS Collaboration as of June 2024. A comparison to theory prediction is shown in the right panel [30].

Despite these successes, the SM fails to describe some physical observations. For instance, the neutrinos of the SM are massless, but neutrinos have been observed to oscillate between flavor eigenstates, demonstrating that neutrinos have small nonzero mass values. The following sections briefly describe the most notable shortcomings of the SM.

1.3.1 MATTER-ANTIMATTER ASYMMETRY

In the early universe, baryons and antibaryons were in thermal equilibrium with high-energy photons via processes such as $\gamma + \gamma \rightleftharpoons p + \bar{p}$. Due to the cooling and expansion of the universe, both the forward and backward processes in this reaction eventually ceased. This process, known as Big Bang baryogenesis, effectively fixed the total number of baryons n_B and number of antibaryons $n_{\bar{B}}$ in the Universe [14]. SM calculations predict that the number of baryons and antibaryons should be the same, with $n_B = n_{\bar{B}} \sim 10^{-18} n_\gamma$, however observations of light isotope abundances are consistent instead with $\frac{n_B - n_{\bar{B}}}{n_\gamma} \sim 10^{-9}$, indicating a universe vastly dominated by matter over antimatter [14]. Violation of the symmetry of nature under inversion of both charge and parity, or CP violation in the SM, is a necessary condition to generate this observed matter-antimatter asymmetry. Although CP violation has been observed in the quark sector of the weak interaction, this alone is insufficient to explain the observed asymmetry and thus the SM does not provide an adequate explanation of this phenomenon [31].

1.3.2 GRAVITY

As alluded to at the start of this chapter, the SM encompasses electromagnetic, weak, and strong interactions but does not provide a mechanism for gravitational interactions. As shown in Table 1.1, gravity is many orders of magnitude weaker than the other forces, and is negligible at the energy scale and corresponding length scale probed at the LHC. However, at a very high energy scale known as the Planck scale of $O(10^{19})$ GeV, the effects of quantum gravity can no longer be neglected [32]. If the SM is a low-energy manifes-

tation of a so-called Grand Unified Theory (GUT) beyond the Planck scale, a theoretical extension to the SM to describe the quantum effect of gravity is necessary. Such theories have been proposed, but are typically difficult to test experimentally due to the extremely high energies (equivalently, small length scales) required to produce non-negligible measurable effects.

1.3.3 DARK MATTER AND DARK ENERGY

Cosmological observations imply that the matter content of the universe described by the SM comprises only approximately 5% of the total matter-energy content of the universe. Within the Λ CDM standard model of cosmology, the other 95% of energy-matter density in the universe consists of *dark matter* and *dark energy*. Dark matter is a hypothesized form of matter which only interacts weakly, and is required to explain the observed rotation curves of galaxies [33] and the anisotropies observed in the cosmic microwave background radiation spectrum [14], among other cosmological observations. The SM does not include a dark matter candidate, although weakly interacting massive particles (WIMPs) and axions are popular candidates for BSM particles that could extend the SM to explain cosmological observations [34]. Observations of supernovae indicate that the universe is expanding at an accelerated rate; dark energy is a proposed form of energy associated to a non-zero cosmological constant ($\Lambda \neq 0$) in Einstein's equations of general relativity, which could explain this acceleration [35]. As with dark matter, the SM does not provide a mechanism for dark energy.

1.3.4 THE HIERARCHY PROBLEM

The limitations listed so far pertain to physical observations that are not accounted for in the SM. The so-called *hierarchy problem* is by contrast a mathematical feature of the SM. The measured Higgs Boson mass of 125 GeV is the sum of a *bare mass* term appearing in the SM Lagrangian and quantum correction terms arising from interactions between

SM particles and the Higgs field. If no additional fundamental scale exists between the Higgs mass and the Planck scale of $O(10^{19})$ GeV, these quantum corrections should be many orders of magnitude larger than the Higgs mass. To result in such a light Higgs Boson compared to the Planck scale, the quantum correction terms are therefore required to all cancel one another other almost exactly. This *fine-tuning* of the Higgs mass is highly suggestive of SM extensions which include yet-undiscovered particles. In particular, as fermionic and bosonic quantum correction contributions to the Higgs mass calculation have opposite sign, an SM extension which introduces an additional symmetry between fermions and bosons could provide a natural mechanism for generating such a cancellation, explaining the large gap between the Higgs mass scale and the Planck scale. The class of extensions to the SM which introduce a symmetry between fermions and bosons, known as theories of *Supersymmetry*, are discussed in more detail in Chapter 2.

CHAPTER 2:

BEYOND THE STANDARD MODEL

Physics theories beyond the Standard Model (BSM) aim to address one or more of the SM shortcomings listed in Section 1.3, in many cases by proposing an extension to the particle content of the universe. If BSM physics processes arising from proton-proton collisions at the LHC manifest differently than SM physics, it is plausible that such processes would evade standard detection methods. As an example, an as-yet-undiscovered particle, decaying with an average lifetime in the range of tens of picoseconds to tens of nanoseconds to an SM lepton, could result in leptons measurable by the ATLAS detector with significant displacement and/or delay as compared to leptons typically produced at the LHC. A particle with these properties is classified as a Long-Lived Particle (LLP), and the search presented here targets BSM processes resulting in an LLP. This search is interpreted with a simplified Supersymmetry (SUSY) BSM model resulting in pairs of displaced electrons or muons, with a focus on the electron final state. An overview of SUSY is given in Section 2.1, and a well-motivated simple choice of SM extension via SUSY is introduced in Section 2.2. The benchmark model studied in this thesis is discussed in Section 2.3. An overview of LLPs is presented in Section 2.4.

2.1 SUPERSYMMETRY

SUSY [36–39] theories extend the SM by introducing an additional symmetry to those discussed in Chapter 1. The additional symmetry is one of rotations in spin space, which transforms fermions into bosons and vice versa [40]. A direct consequence of this additional symmetry is the introduction of so-called *superpartners* of the SM particles, which are identical to their SM counterparts except for a difference in spin of $1/2$. SM particle-superpartner pairs form groups known as *supermultiplets*. The bosonic superpartners of the fermions are called "sfermions", where the "s" denotes that these are scalar bosons. The *squarks* and *sleptons* are similarly named by adding "s" to the start of their fermion partner's name. For example, the sleptons are named the *selectron* $\tilde{e}$, *smuon* $\tilde{\mu}$, and the *stau* $\tilde{\tau}$. The fermionic superpartners of the gauge bosons are named by adding the suffix *-ino* to the end of the SM particle name, as in the *gluino*, $\tilde{g}$. An unbroken supersymmetry would result in superpartners of the same mass as their SM counterparts, and is easily ruled out by existing experimental observations. Plausible SUSY models therefore require a broken symmetry which produces superpartners much heavier than their SM partners.

SUSY theories can offer attractive solutions to the shortcomings of the SM. As an example, in the SM, the electromagnetic and weak forces are *unified* in that they are low-energy manifestations of a single electroweak force at high energies. Although it is not possible to simultaneously unify the strong, electromagnetic and weak interactions when the measured SM gauge couplings are extrapolated to high energy scales, this unification is possible in the MSSM dependent on the choice of sparticle masses as shown in Figure 2.1. Additionally, as alluded to at the end of the previous chapter, depending again on choice of sparticle masses, SUSY offers naturally-motivated cancellations to the Higgs mass corrections and could explain the hierarchy problem.




Figure 2.1: Strength of the forces as a function of energy scale Q for SM couplings (dashed black) and MSSM couplings with sparticle masses near 750 GeV (blue) and 2.5 TeV (red).

2.2 MINIMAL SUPERSYMMETRIC STANDARD MODEL

There is a wide choice of SUSY extensions to the SM, but a popular choice, due to its relative simplicity, is the Minimal Supersymmetric Standard Model (MSSM) [41], which extends the SM by the minimal particle content compatible with known SM particles and properties as measured by experiment. In the MSSM, the superpartners of the gluon, W , and B bosons of the SM are the gluino, wino, and bino, with each pair forming a gauge supermultiplet. Left-handed SM fermions form an $SU(2)_L$ doublet chiral supermultiplet with their superpartners, while right-handed SM fermions and their partners form singlet chiral supermultiplets. The Higgs sector of the MSSM requires two chiral supermultiplets, of hypercharge $Y = \pm \frac{1}{2}$, to ensure the cancellation of chiral anomalies arising from introducing a single superpartner for the charged component of the scalar field [42], and to lend masses to both up- and down-type fermions [43]. The supermultiplets of the MSSM are shown in Table 2.1.

Field Content of the MSSM

Super-multiplets	Super-field	Bosonic fields	Fermionic partners	SU(3)	SU(2)	U(1)
gluon/gluino	$\widehat{V}_8$	g	$\widetilde{g}$	8	1	0
gauge boson/ gaugino	$\widehat{V}$	$W^\pm, W^0$	$\widetilde{W}^\pm, \widetilde{W}^0$	1	3	0
	$\widehat{V}'$	B	$\widetilde{B}$	1	1	0
slepton/ lepton	$\widehat{L}$	$(\widetilde{\nu}_L, \widetilde{e}_L^-)$	$(\nu, e^-)_L$	1	2	-1
	$\widehat{E}^c$	$\widetilde{e}_R^+$	e_L^c	1	1	2
squark/ quark	$\widehat{Q}$	$(\widetilde{u}_L, \widetilde{d}_L)$	$(u, d)_L$	3	2	1/3
	$\widehat{U}^c$	$\widetilde{u}_R^*$	u_L^c	$\bar{3}$	1	-4/3
	$\widehat{D}^c$	$\widetilde{d}_R^*$	d_L^c	$\bar{3}$	1	2/3
Higgs/ higgsino	$\widehat{H}_d$	(H_d^0, H_d^-)	$(\widetilde{H}_d^0, \widetilde{H}_d^-)$	1	2	-1
	$\widehat{H}_u$	(H_u^+, H_u^0)	$(\widetilde{H}_u^+, \widetilde{H}_u^0)$	1	2	1

 Table 2.1: The supermultiplets of the MSSM, along with their $SU(3) \times SU(2) \times U(1)$ quantum numbers [13]. In the cases of quarks and leptons, only one generation is shown.

The Higgs scalar fields of the MSSM have eight degrees of freedom, three of which are absorbed into the longitudinal polarizations giving the $W^\pm$ and Z bosons mass following electroweak symmetry breaking. The other five degrees of freedom correspond to two neutral scalar Higgs bosons H^0, h^0 , a pair of charged particles $H^\pm$, and a pseudoscalar boson A , with h^0 identifying the SM 125 GeV Higgs boson [41]. Each of these five Higgs bosons has a superpartner, collectively referred to as *higgsinos*. The theoretically physically observable mass eigenstates known as *charginos* $\tilde{\chi}_1^\pm, \tilde{\chi}_2^\pm$ and *neutralinos* $\tilde{\chi}_1^0, \tilde{\chi}_2^0, \tilde{\chi}_3^0, \tilde{\chi}_4^0$, labeled in ascending mass order, result from the higgsinos mixing with the electroweak gauge boson superpartners. The MSSM also incorporates gravity via the hypothesized spin-2 graviton g and a spin- $\frac{3}{2}$ superpartner, the gravitino $\tilde{G}$.

In the SM, the stability of the proton is protected by the conservation of baryon number B and lepton number L . This agrees with experimental observation, which currently places a lower limit on the proton lifetime on the order of 10^{33} years [44]. In the MSSM, renormalizable terms are allowed in the Lagrangian which violate B and L conservation, and the couplings associated to these terms therefore must be very small to avoid con-

tradicting experimental observations. A new symmetry, known as R-parity P_R , which forbids such couplings, can be written as

$$P_R = (-1)^{3(B-L)+2s} \quad (2.1)$$

where s denotes the spin of the particle. With this definition, all SM particles have $P_R = +1$ and SUSY particles have $P_R = -1$; in the MSSM, R-parity is manifestly conserved. As a consequence, massive sparticles decay via chains of lighter sparticles and SM particles, resulting in a lightest supersymmetric particle (LSP) which then cannot decay to an SM particle. The resulting LSP is therefore a well-motivated candidate dark matter particle. Additionally, P_R requires that sparticles must be produced in pairs, as in the direct slepton pair-production model studied in this thesis.

2.3 GAUGE MEDIATED SUPERSYMMETRY BREAKING

The scenario used as a benchmark in this thesis belongs to a subset of SUSY models where SUSY is broken via Gauge Mediated Supersymmetry Breaking (GMSB) [43,45–49]. The benchmark model is useful in quantifying the sensitivity of the analysis method to a representative BSM scenario, although the search is designed for sensitivity to a general LLP decay resulting in displaced leptons. GMSB models are characterized by an *observable sector* containing the particles of the MSSM, a *secluded sector* in which SUSY is broken, and a *messenger sector*. The particles of the messenger sector are chiral supermultiplets that couple via $SU(3) \times SU(2) \times U(1)$ gauge interactions to the observable sector, and additionally couple to the secluded sector. The supersymmetry breaking in the secluded sector is thus "communicated" to the observable sector via the messengers. The mass scale of the messenger sector M and the mass-squared splittings in the messenger supermultiplets F are parameters which characterize the energy scale of the GMSB theory. Mass splittings arise as a result of interactions between the messenger and secluded sector, and the pa-

parameter F directly measures the magnitude of SUSY breaking in the messenger sector. Upon SUSY breaking, SM gauge boson and fermion masses are protected by gauge invariance while first(second) order loop corrections can generate gaugino(squark and slepton) masses [46].

Several distinct phenomenological characteristics arise in GMSB theories. Crucially, the LSP is always the gravitino $\tilde{G}$, and all SUSY decay chains result in the LSP and an SM particle in the visible sector. The gravitino mass is typically vanishingly small, as it is inversely proportional to the Planck scale of $O(10^{19})$ GeV. The next-to-lightest supersymmetric particle (NLSP) is either the lightest neutralino $\tilde{\chi}_1^0$, or the slepton $\tilde{l}$, with a small gravitational coupling allowing the NLSP to have a long lifetime (see Section 2.4).

The MSSM contains 120 free parameters, rendering a fully general search intractable. The search presented in this thesis makes use of a *simplified model*, which focuses on a subset of tuneable observable parameters. The model used is illustrated in Figure 2.2, where sleptons produced in $p - p$ collisions decay to their lepton counterpart and a nearly massless gravitino, $\tilde{l} \rightarrow l\tilde{G}$, resulting in a final-state characterized by two leptons and missing energy carried away by the LSPs.




Figure 2.2: Illustration of the simplified model decay topology considered in this search. In this GMSB-inspired configuration sleptons of all three generations ($\tilde{e}$, $\tilde{\mu}$, $\tilde{\tau}$) are pair-produced and subsequently decay to a lepton and a gravitino LSP.

The slepton NLSP lifetime τ can be expressed in terms of slepton mass $m_{\tilde{l}}$ and SUSY-breaking scale, $\sqrt{F}$ [50], as:

$$c\tau \approx 100\mu\text{m} \left(\frac{100\text{ GeV}}{m_{\tilde{l}}} \right)^5 \left(\frac{\sqrt{F}}{100\text{ TeV}} \right)^4 \quad (2.2)$$

SUSY breaking scales in the range $\sqrt{F} \sim 100\text{ TeV}$ and higher lead to measurably displaced slepton decay products at the LHC, depending on slepton mass.

In the $\tilde{\tau}$ case, the final state particles considered in this analysis are the muons and electrons which might result from subsequent leptonic decays of the τ . Therefore, while GMSB models often predict staus as the NLSP [50], the $\tilde{\tau}$ sensitivity is limited compared to other slepton flavors due to the requirement of an additional leptonic decay (the τ branching fraction to e or μ is 34%).

2.4 LONG-LIVED PARTICLES

LLPs occur in the SM, and are predicted in a wide variety of BSM scenarios, including both P_R -conserving [51–57] and P_R -violating supersymmetry (SUSY) [58, 59], split-SUSY [60, 61], GMSB [62–64], universal extra dimensions [65, 66], and Hidden Valley models [67].

The lifetime of a particle with mass m_X decaying to a set of final-state products $\{p_f\}$ is a function of the decay width Γ , given by

$$\Gamma = \frac{1}{2m_X} \int d\Pi_f |\mathcal{M}(m_X \rightarrow \{p_f\})|^2 \quad (2.3)$$

Here $d\Pi_f$ represents the Lorentz-invariant phase space available for the decay, and $\mathcal{M}$ is the matrix element encoding the probability of a decay to each final state p_f . Γ is the probability of a particle to decay per unit time, and the mean lifetime of a particle τ is its reciprocal $\tau = (\Gamma)^{-1}$. Thus, LLPs are characterized by small Γ , indicating either small $\mathcal{M}$, small phase space available for decay, or both.

In the SM, the b quark is an example of small $\mathcal{M}$ resulting in an LLP. The matrix elements dictating the rate of b decay into u or c quarks are highly suppressed, leading to b quarks from $p - p$ collisions at the LHC which typically travel a few millimeters at nearly the speed of light before decaying. The neutron decays via $n^0 \rightarrow p^+ + e^- + \bar{\nu}_e$, but is nearly mass-degenerate with the proton; much of the initial mass-energy in this process must go to the rest mass of the resulting proton, allowing very little remaining mass-energy for the electron and neutrino products. This small phase space available for neutron decay leads to a free neutron lifetime of nearly 15 minutes. Figure 2.3 gives a survey of SM particle mean lifetimes, showing the range of lifetimes for which an SM LLP typically decays within the ATLAS detector. In the GMSB scenario considered in this analysis, the slepton can acquire a long lifetime due to a small $\mathcal{M}$ dictating the coupling to the gravitino LSP.




Figure 2.3: A selection of SM particle masses and lifetimes [68]. Assuming each particle shown is produced moving near light speed, those which decay immediately compared to the scale of detector spatial and timing resolution are shown in the left shaded band. Those which, on average, escape the detector before decaying are shown in the right shaded band. The particles falling inside the band typically decay within ATLAS with some measurable displacement and/or delay.

The search presented in this thesis uses the GMSB slepton pair-production model of Section 2.3 as a benchmark to quantify the sensitivity of the search method to a particu-

lar BSM scenario. However, the search is built to be model-agnostic: it is designed to be generally sensitive to models which contain LLPs decaying leptonically in the tracker or before the electromagnetic calorimeter face of the ATLAS detector (see Chapter 4). To this end, a tool is developed (see Chapter 6) to measure the likelihood that an input electron reconstructed in an ATLAS event is consistent with having arisen from a BSM LLP decay. Additionally, the analysis strategy is designed to avoid the use of the missing energy carried by the gravitinos in the benchmark model when possible, as this is not a characteristic of all plausible BSM scenarios resulting in displaced leptons. Similarly, a search region is constructed which does not explicitly require more than one displaced lepton in the final state. In Chapter 11, the results of the search are interpreted via the GMSB slepton pair-production scenario presented, and as model-independent upper limits on BSM physics processes with displaced leptons in the final state.

PART II:

EXPERIMENT

CHAPTER 3:

THE LARGE HADRON COLLIDER

The Large Hadron Collider (LHC) accelerates protons to very near the speed of light before colliding them head-on to recreate conditions that existed very shortly after the Big Bang. The LHC is housed in a 27 km ring sitting 100 m underground at the border of France and Switzerland near the city of Geneva. The data analyzed in this thesis were delivered to the ATLAS experiment (see Chapter 4) by the LHC complex. The LHC is optimized to deliver as many collisions as possible to ATLAS and the other three main experiments at the CERN site: CMS, LHCb, and ALICE. Section 3.1 provides an overview of relevant collider physics concepts. Section 3.2 discusses the LHC complex and the design of the LHC. The operation of the LHC during the collection of data used in this thesis is discussed in Section 3.3.

3.1 LHC ACCELERATOR CONCEPTS

The LHC is designed to circulate two beams of protons p and/or ions A in three modes $(p - p)$, $(p - A)$, or $(A - A)$ in opposite directions before colliding them at interaction points (IPs) at one of the four main LHC experiments. Rather than circulating a continuous beam of particles, protons are sorted into discrete *bunches* using 16 radiofrequency (RF) cavities. Each RF cavity is a metallic chamber housing an electromagnetic field which changes direction to accelerate (decelerate) circulating protons at the falling (leading)

edge of a bunch until a bunch structure of protons accelerated to the desired beam energy is achieved. Dipole magnets are used to bend the circulating beams of protons and maintain their orbit around the LHC ring using a constant vertical magnetic field. The magnitude of the magnetic field from the dipoles is synchronized to the proton energy from RF acceleration to maintain a circular orbit. Quadrupole and higher-order multipole magnets are used for beam focusing, in particular near the IPs to squeeze the proton beams into as small an area as possible to maximize the probability of collisions.

In general, particle colliders such as the LHC aim to deliver the maximum possible sized data samples, in terms of number of collision events N_{events} , for participating physics experiments. This number is given by the equation

$$N_{\text{events}} = \sigma_{\text{events}} \int \mathcal{L}(t) dt \quad (3.1)$$

Here, the symbol σ_{events} represents the cross section of the event being studied, which is a direct function of collision energy. $\int \mathcal{L}(t) dt$ represents the *integrated luminosity* delivered by the beam over its operation time, while the instantaneous luminosity $\mathcal{L}$ is a function of the beam parameters. In particular, one can write the luminosity as

$$\mathcal{L} = \frac{N_b^2 n_b f_{\text{rev}} \gamma}{4\pi \epsilon_n \beta^*} F \quad (3.2)$$

In this equation, N_b is the number of particles in each bunch, n_b is the number of colliding bunches in each beam, f_{rev} is the frequency of beam revolution, γ is the relativistic Lorentz factor, ϵ_n or *normalized emittance* measures the beam spread transverse to the collision axis, β^* or the *amplitude function* is a measure of how much the beam is squeezed at the IP, and F is a geometric luminosity reduction factor determined by the crossing angle, beam size, and bunch length in the collision plane [8]. In general, the luminosity can be maximized by colliding many bunches of many protons at high frequency, and by focusing the beams at the IP such that β^* and ϵ_n are as small as achievable.

The LHC also seeks to maximize the center-of-mass energy $\sqrt{s}$ of the colliding beams, such that high-mass SM particles and potentially new massive BSM particles result from the collision and can be probed. As the LHC collides two beams with equal and opposite energy, conservation of energy dictates that the total energy of the colliding protons is twice the energy of each colliding beam. However, as protons are composite particles, the total energy in each collision depends on the momenta of the colliding constituent quarks or gluons within the proton, and is therefore unknown. The momentum and energy of collision products in the plane transverse to the beam, in which the total initial energy and momentum are zero, are the quantities of interest at a hadron collider. By contrast, in a point-particle collision (for instance, at an electron-positron collider) the total initial momentum and energy of the colliding particles are known, and the final total energy and momentum can therefore be measured precisely. The advantage of colliding protons rather than electrons, however, is the scaling of energy loss via synchrotron radiation (E_{synch}) for particles relativistically accelerated in a magnetic field. In particular, since the energy radiated $E_{\text{synch}} \propto \gamma^4$ and $\gamma = E/m_0c^2$, due to the proton-electron mass ratio $m_p/m_e \sim 10^3$ electrons lose energy in a circular accelerator at a rate $\sim 10^{12}$ times the energy loss rate for protons.

3.2 THE CERN ACCELERATOR COMPLEX

Before entering the LHC ring, the protons used for collisions are produced by stripping electrons from a bottle of hydrogen gas. These protons are then accelerated linearly by the Linac2 accelerator to an energy of 50 MeV, and injected into a Booster, where they are accelerated to 1.4 GeV. The protons are then injected into the Proton Synchrotron (PS) and accelerated to 25 GeV, followed by the Super Proton Synchrotron (SPS) where they are accelerated to 450 GeV. Finally, the beam is separated into clockwise and anticlockwise beams which circulate in the LHC ring. Here, they are accelerated to 6.5(6.8) TeV in Run

2(Run 3), and collide at the IPs at a bunch crossing frequency of 40 MHz. Figure 3.2 shows a diagram of the CERN accelerator complex showing the Linac2, Booster, PS, SPS and LHC, as well as the four main experiments and additional smaller experiments.




Figure 3.1: Diagram of the CERN accelerator complex, showing the positions of four major experiments on the LHC ring as well as the full LHC accelerator chain and the setup of other LHC experiments [69].

The LHC ring is divided into eight octants, each of which is independently powered and houses 154 individual dipole magnets connected in series in a continuous cryogenic system. The dipoles use the superconductive material niobium-titanium (NbTi, $T_C = 10$ K) cooled to 1.9 K which provide a magnetic field of 8.33 T. The dipoles for both beams employ a novel twin-bore technology which they share a cryostat to deal with the tight 3.7 m diameter constraint of the LHC tunnel [8]. Figure 3.2 shows an image of the LHC tunnel, with a 3D cutout of the beam pipe displaying the dipole magnets for both beams.




Figure 3.2: Image of the LHC tunnel , showing a 3D rendering of the dipole magnets for the circulating beams within the LHC beam pipe [70].

3.3 OPERATING CONDITIONS IN RUN 2 AND RUN 3

This result relies on data collected during the operating period of the LHC known as Run 2 (2015-2018) at $\sqrt{s} = 13$ TeV and the LHC's Run 3 (2022-2023) at $\sqrt{s} = 13.6$ TeV. Run 2 and Run 3 each consist of many shorter runs, or *LHC fills*. During a fill the machine first injects bunches into the LHC ring, ramps the beam energies to the nominal collision energy, focuses the beams at the IP, and finally collides the beams until n_b is depleted to a level that a re-fill is more efficient for luminosity maximization than continuing the run. The beam is then dumped and necessary calibrations are performed at the experiments as the next fill is prepared.

The circumference of the LHC ring allows for 3564 total bunches separated by 25 ns and traveling very near to the speed of light. However, in practice, approximately 2400 of those bunches are filled with protons, with the other bunches left empty. The filled bunches contain $O(10^{11})$ protons each at the start of each fill, with a peak instantaneous luminosity achieved in Run 2 and Run 3 of more than twice the design luminosity of $10^{34} \text{ cm}^{-2} \text{ s}^{-1}$.

Colliding dense, focused bunches of protons comes at the cost of many interactions occurring in a given bunch crossing, a phenomenon known as *in-time pileup*. The mean pileup $\langle\mu\rangle$ per bunch crossing in Run 2 and Run 3 is shown in Figure 3.3, with an average value of $\langle\mu\rangle = 33.7$ and $\langle\mu\rangle = 46.5$ in Run 2 and Run 3, respectively. This large pileup presents a challenge when identifying the underlying $p - p$ interaction of interest in a given bunch crossing, as demonstrated in Figure 3.4, where a $Z \rightarrow \mu\mu$ event is identified from a number of candidate vertices.




Figure 3.3: Mean number of interactions per bunch crossing delivered to ATLAS by the LHC for each year and on average (a) in Run 2 and (b) in early Run 3 [71].




Figure 3.4: ATLAS event display with a zoomed view of the tracker in the $R - z$ plane, showing a few cm along the beamline in each direction from the center of the detector. The many vertices (dots) and tracks (lines) shown all share a single bunch crossing. The vertex of interest in the event and corresponding tracks are shown as thick yellow lines [72].

The ATLAS detector and its event reconstruction software are optimized to measure the physical processes of interest in each collision event with excellent precision; these are the subjects of Chapter 4 and Chapter 5, respectively.

CHAPTER 4:

THE ATLAS DETECTOR

The ATLAS (A Toroidal LHC ApparatuS) experiment, weighing 7,000 tonnes with a diameter of 25 m and a length of 44 m, is the largest of the four main experiments located at an Interaction Point (IP) along the LHC ring. The detector design has been optimized to precisely measure the final state particles resulting from proton proton ($p-p$), proton-ion ($p-A$), and ion-ion ($A-A$) collisions. ATLAS is a general-purpose detector, with physics goals including and later probing the nature of the Higgs boson, measuring cross sections and branching fractions predicted in the standard model, and searching for physics Beyond the Standard Model (BSM). The ATLAS detector configuration as of 2023, with its subsystem components labeled, is depicted in Figure 4.1. Section 4.1 gives an overview of the ATLAS coordinate system, while Sections 4.3-4.5 describe the subsystems. More detailed information is given in Sections 4.4.1-4.4.4 to the Liquid Argon (LAr) electromagnetic (EM) calorimeter, as the analysis here makes extensive use of the measurement of electrons and photons by the LAr EM calorimeter.

The results in this thesis rely on data taken by the ATLAS detector during the LHC's Run 2 (2015-2018) and the early years of the LHC's Run 3 (2022-2023), which is planned to continue through 2025. During the LHC's second long shutdown (LS2) between Run 2 and Run 3, the so-called Phase-I upgrade of ATLAS took place, resulting in new configurations of the muon, EM calorimeter, and Trigger and Data Acquisition (TDAQ) sys-

tems [73] [74]. In each of these cases, the Run 3 configuration is highlighted in a dedicated subsection discussing the upgrade.




Figure 4.1: Diagram of the ATLAS detector configuration in Run 3, labeling subdetector components [74].

4.1 ATLAS COORDINATE SYSTEM

The ATLAS detector employs a right-handed coordinate system, with the origin centered at the IP. The $+x$ -axis points towards the center of the LHC ring, and the $+y$ -axis points upwards. The $+z$ -axis, in the direction of the "A-side" of the detector, points along the beamline. The symbol ϕ denotes the azimuthal angle around the beam axis, while θ denotes the polar angle from the beam axis. The pseudorapidity η is defined as

$$\eta \equiv -\ln \left[\tan \left(\frac{\theta}{2} \right) \right] \quad (4.1)$$

This coordinate is more convenient than θ because of its invariance under Lorentz boost in the approximation $v = c$, a reasonable assumption for the hyper-relativistic products




Figure 4.2: ATLAS coordinate geometry, (a) shown in a diagram with example measurements of $\vec{p}$, p_T and (b) showing the relationship between η and θ .

of $p - p$ collisions at the LHC. The distribution of particles arising from LHC collisions is also approximately uniform as a function of η at LHC collision energies. The ATLAS coordinate system, with a reference arrow to indicate the center of the LHC and an example transverse momentum (p_T) measurement, as well as a few reference values of $|\eta|$, is shown in Figure 4.2.

It is convenient to define a distance metric ΔR between objects measured in the detector in terms of η and ϕ , as

$$\Delta R \equiv \sqrt{\Delta\phi^2 + \Delta\eta^2} \quad (4.2)$$

The p_T and transverse energy (E_T) calculated for particle four vectors in each collision provide key measurements in characterizing the event. As the net momentum carried by the proton beams in the $x-y$ plane before each collision is zero, conservation of momentum dictates that the sum of p_T of collision products should also be zero. Certain particles, such as SM neutrinos or a host of BSM candidates, pass through ATLAS without interacting and can be detected indirectly through a measurement of missing transverse energy (E_T^{miss}), defined as

$$E_T^{\text{miss}} \equiv -|\Sigma \vec{p}_T| \quad (4.3)$$

where the sum runs over products of the collision.

4.2 ATLAS SUBSYSTEM OVERVIEW

The proceeding sections describe the various subsystems of the ATLAS detector. The detector is divided into three main subsystems, each exhibiting cylindrical symmetry, and forward-backward symmetry with respect to the IP. Starting closest to the beamline and working outwards, these are: the Inner Detector (ID) to measure the spatial coordinates and momenta of charged particles, the Calorimeter system to measure the energy of particles, and the Muon Spectrometer (MS) which is used to identify and measure muon kinematics. An overview of the performance requirements of each subsystem, as well as coverage in $|\eta|$ is shown in Table 4.1.

Detector Component	Required Resolution	η coverage	
		Measurement	Trigger
ID Tracking	$\sigma_{p_T}/p_T = 0.05\% p_T \oplus 1\%$	± 2.5	
EM Calorimetry	$\sigma_E/E = 10\%/\sqrt{E} \oplus 0.7\%$	± 3.2	± 2.5
Hadronic Calorimetry (jets) barrel and end-cap forward	$\sigma_E/E = 50\%/\sqrt{E} \oplus 3\%$	± 3.2	± 3.2
	$\sigma_E/E = 100\%/\sqrt{E} \oplus 10\%$	$3.1 < \eta < 4.9$	$3.1 < \eta < 4.9$
Muon Spectrometer	$\sigma_{p_T}/p_T = 10\% \text{ at } p_T = 1\text{GeV}$	± 2.7	± 2.4

Table 4.1: ATLAS performance requirements by subsystem, in terms of p_T , E resolution. Coverage in η at trigger level and in offline reconstruction is also shown [7].

4.3 INNER DETECTOR

On average $\langle \mu \rangle = 33.7$ ($\langle \mu \rangle = 47.7$) simultaneous inelastic $p - p$ -interactions occur alongside the hard proton scatter of interest in ATLAS in Run 2 (Run 3), with $p - p$ collisions occurring at a rate of 40 MHz (every 25 ns). This high-pileup environment results

in a high density of particles along the beamline during each $p - p$ collision, which necessitates precision vertex and momentum measurements with fine detector granularity to identify physics processes of interest. The ATLAS ID system, shown schematically in Figure 4.3, achieves this by measuring the trajectories of charged particles to allow determination of their charge and transverse momentum. The ID also helps to identify electrons for $|\eta| < 2.0$ with $0.5 < p_T/\text{GeV} < 150$. The inner two ID subsystems, the Pixel detector and the Semiconductor Tracker (SCT), are composed of discrete high-resolution semiconductor pixel and strip detectors, arranged in concentric cylinders in the barrel region and discs in the $x - y$ plane in the endcap regions. These layers are surrounded by the straw-tube Transition Radiation Tracker (TRT).




Figure 4.3: Diagram of the ATLAS inner detector, with individual components labeled [74].

A 2T magnetic field along the $+z$ axis of ATLAS is provided by a solenoid surrounding the ID (see Section 4.7), with a length of 5.3 m and diameter of 2.5 m. The magnetic field bends the trajectories of charged particles to allow measurement of their charge and transverse momentum. Pattern recognition algorithms are used to reconstruct individual particle tracks by connecting "space-points" from the pixel and SCT detectors, and then

extending these tracks into measurements from the TRT. Track intersections are then used to identify candidates for the position of the hardest scattering process in a bunch crossing, known as the Primary Vertex (PV).

The pixel detector is segmented in $R - \phi$ and z and has approximately 80.4 million readout channels. The innermost pixel layer, the Insertable B-layer (IBL) sits at a radius of $R = 33$ mm from the beamline, with a sensor size of $50 \times 250 \mu\text{m}$ ².¹ The other three barrel layers, with radii from 50.5 – 122.5 mm, as well as the endcap disks, with $|z| = 495 - 650$ mm, have a sensor size of $50 \times 400 \mu\text{m}$ ². The barrel layers have an intrinsic accuracy of $10 \mu\text{m}$ ($R - \phi$) and $115 \mu\text{m}$ (z), while the endcap disks have an intrinsic accuracy of $10 \mu\text{m}$ ($R - \phi$) and $115 \mu\text{m}$ (R).

The SCT, with approximately 6.3 million readout channels, employs the same silicon-based technology as the pixel detector but instead uses 64 mm long strip sensors with a pitch of $80 \mu\text{m}$ to provide measurements in $R - \phi$. The barrel SCT consists of four layers of strips at $R = 299 - 514$ mm, each layer with one strip parallel to the beam and another "stereo" layer at a small angle of 40 mrad. The endcap SCT consists of nine disks of stereo strips, each with one disk in the $R - \phi$ plane and one at 40 mrad, at $|z| = 853 - 2710$ mm. The barrel layers have an intrinsic accuracy of $17 \mu\text{m}$ ($R - \phi$) and $580 \mu\text{m}$ (z), while the endcap disks have an intrinsic accuracy of $17 \mu\text{m}$ ($R - \phi$) and $580 \mu\text{m}$ (R). The spatial resolution in the $z(R)$ direction in the barrel (endcap) case results from the stereo structure of the layers.

The outermost ID layer, the TRT with approximately 360,000 readout channels, uses 4 mm diameter straw tubes to provide on average 36 hits per track, providing an intrinsic accuracy of $130 \mu\text{m}$ per straw in $R - \phi$ for tracking up to $|\eta| < 2.0$. In the barrel the TRT straws are 1,440 mm long and oriented along the z axis. In the endcap the straws are 370 mm long and oriented radially in wheels. The xenon-based gas mixture in the straw tubes

¹The IBL is named for its use in identifying displaced vertices from B-hadrons, which have lifetimes $O(\text{ps})$ and therefore typically travel a few mm before decaying.

allows detection of transition-radiation photons which aids in the ID's electron identification abilities.

4.4 CALORIMETRY

The ATLAS calorimeter system, shown in Figure 4.4, provides measurements of the energies of particles in the range $|\eta| < 4.9$. The EM calorimeter, consisting of the Liquid Argon (LAr) barrel (EMB) and endcap (EMEC) offers precision measurements of electrons and photons. The LAr forward calorimeter (FCAL) and hadronic endcap (HEC) and the Tile hadronic calorimeter (surrounding LAr in the barrel and endcap) have coarser granularities and are primarily tasked with measurements of jets and E_T^{miss} . The calorimeter depths, in terms of radiation lengths (X_0) for EM showers and interaction lengths (λ) for hadronic showers, are designed carefully to contain showers and provide a good E_T^{miss} measurement, preventing EM punch-through into the hadronic calorimeter and hadronic punch-through into the muon system. The EM calorimeter in the barrel(endcap) has a thickness of $X_0 > 22$ ($X_0 > 24$). The total thickness of the calorimeters is 11λ at $\eta = 0$.

4.4.1 LIQUID ARGON ELECTROMAGNETIC CALORIMETER

The LAr calorimeter EMB provides coverage for $|\eta| < 1.475$, and consists of two identical half-barrels separated by a 4mm gap at $z = 0$. Each of the two LAr EMECs contain an inner wheel covering $1.375 < |\eta| < 2.5$ and an outer wheel covering $2.5 < |\eta| < 3.2$. The lead absorbers and copper-kapton electrodes of the LAr EM calorimeters are arranged in an accordion geometry, as shown in Figure 4.5. The accordion waves run in ϕ in the barrel, and run parallel to the radial direction in the endcap. This choice allows full coverage in ϕ , as opposed to a radially symmetric structure which might introduce gaps at fixed ϕ .

In the EMB, the calorimeter is segmented in (R, η, ϕ) into a presampler followed by front, middle, and back layers. The cell granularity in $\Delta\eta \times \Delta\phi$ in each layer as a function




Figure 4.4: Diagram of the ATLAS calorimeter system, with individual components labeled [74].

of $|\eta|$, and the total number of channels, is shown in Table 4.2. The combination of energy and spatial measurements in each layer are used to calculate the angle at which incoming particles traverse the calorimeter, and their shower shapes. The presampler, for $|\eta| < 1.8$, accounts for energy losses before a particle reaches the EM calorimeter. The front layer is finely segmented; this is useful for instance in identifying $\pi^0 \rightarrow \gamma\gamma$ decays in which EM showers are expected to have two collimated maxima. Most of the energy from a shower is captured in the middle layer, and the back layer mitigates energy losses from high-energy EM showers in LAr.

$ \eta $ range	0 to 1.4	1.4 to 1.8	1.8 to 2.0	2.0 to 2.5	2.5 to 3.2
Presampler	0.025×0.1	0.025×0.1			
Front Layer	0.003×0.1	0.003×0.1	0.004×0.1	0.006×0.1	0.1×0.1
Middle Layer	0.025×0.025	0.025×0.025	0.025×0.025	0.025×0.025	0.1×0.1
Back Layer	0.050×0.025	0.050×0.025	0.050×0.025	0.050×0.025	
Trigger	0.1×0.1	0.1×0.1	0.1×0.1	0.1×0.1	0.2×0.2
No. Channel	110,208	25,600	12,288	24,064	1,792

Table 4.2: LAr EM calorimeter cell granularity ($\Delta\eta \times \Delta\phi$) and number of readout channels as a function of eta [76].




Figure 4.5: Diagram of LAr EMB, showcasing the accordion structure as well as the arrangement of the electrodes and absorbers across the Liquid Argon gap [75].

4.4.2 LIQUID ARGON HADRONIC ENDCAP AND FORWARD CALORIMETER

The LAr subsystem also provides calorimetry targeted primarily towards hadronic interactions in the forward region. The HEC covers $1.375 < |\eta| < 3.2$, and consists of an inner and outer wheel, each with two segments in depth. Each wheel is projective in ϕ , built from 32 identical wedge modules. The inner wheel is built from 25mm copper plates, while the outer wheel is built from 50mm copper plates, in each case interleaved with 8.5mm LAr gaps analogous to the LAr EM calorimeters. The FCAL aims to provide calorimeter coverage at very large $|\eta|$, as well as reduce radiation background penetrating through to the MS. The FCAL is subdivided into three modules, the innermost of which is made of copper and optimized for EM calorimetry. The outer two layers are made of

tungsten and are designed primarily for hadronic calorimetry. Regularly spaced concentric rods and tubes run along the z axis, with LAr as the active medium filling the gap between the rods and tubes. The FCAL interaction length of $\lambda = 10$ mitigates forward leakage into the muon system, while the rod-tube geometry allows for excellent control of the gaps to avoid issues with ion buildup.

4.4.3 LIQUID ARGON READOUT ELECTRONICS

When a particle passes through the calorimeter and produces an electromagnetic shower in the absorbers, the electrons in the shower ionize the Liquid Argon, and the ionization charges then drift across the LAr gap to the readout electrodes. The ionization signal is then passed to the on-detector LAr readout Front End Boards (FEBs) housed in Front End Crates (FECs), which shape, amplify, and sample the signal in three overlapping gain scales at 40 MSPS. The bipolar $CR - (RC)^2$ shaping applied to the analog triangular pulse in each gain to optimize signal-to-noise ratio is shown in Figure 4.6. If an event is accepted by the Level-1 trigger (L1A), four samples of the shaped pulse are digitized corresponding to bunch crossings around the triggered collision and transmitted off-detector on optical fibers to Readout Drivers (RODs) which use Digital Signal Processors (DSPs) to calculate the energy and peaking time of the pulse in each calorimeter cell using a technique known as Optimal Filtering [77] designed to optimize energy and time resolution. The calculated energies are then sent to the ReadOut System (ROS) for inclusion in a stored ATLAS event.²

4.4.4 LIQUID ARGON PHASE-I UPGRADE

During LS2, the LAr system was upgraded to provide finer-granularity, better-resolution and longitudinal shower information to the upgraded Level-1 trigger system to deal with

²Peaking time, a pulse quality factor, and raw Analog-to-Digital Converter (ADC) samples, are also sent, depending on whether the cell energy falls above a programmable threshold




Figure 4.6: Liquid Argon calorimeter pulse, before and after shaping is applied. The shaped pulse shows the samples taken, separated by 25 ns [7].

higher pileup conditions in Run 3 (see Section ??). The legacy "trigger towers" were replaced by "super cells" (SCs), which provide the L1 trigger with per-layer information as well as finer segmentation in $\eta - \phi$ space, as shown in Figure 4.7. The upgrade was implemented with a new trigger front-end board called the LAr Trigger Digitizer Board (LTDB), which transmits the higher-granularity trigger signals to a new off-detector read-out "digital trigger" path added in parallel to the legacy system. On the LTDB, the pulse for each Super Cell is sampled at 40 MSPS and digitised with a custom-designed 12-bit, 4-channel ADC implemented in a 130 nm CMOS technology.

The LTDBs are supported by new baseplanes which route the digital trigger signal and the legacy trigger signal to new off-detector electronics. An optimal filtering procedure, similar to the one used for the precision main-readout, is used by the new off-detector LAr Digital Processor Blades (LDPBs) to calculate energy and time in each supercell at 40 MHz for delivery to the L1 calorimeter system. The updated LAr electronics scheme is shown in Figure 4.8, with the new path highlighted in red. The finer granularity provided by LAr at L1 allows higher background rejection at a given electron efficiency in the trigger.




Figure 4.7: An electron (with 70 GeV of transverse energy) as seen by (a) the legacy Level-1 Calorimeter trigger electronics and (b) by the Phase-I trigger electronics. [78]




Figure 4.8: LAr electronics schematic following the Phase-I upgrade (current configuration in Run 3). Components added during the upgrade are highlighted in red. [78].

This improvement is demonstrated in Figure 4.9, where for $\langle \mu \rangle = 80$, at a fixed electron efficiency of 95%, lower trigger rates are possible with the Phase-I system compared to the Run 2 configuration. In this way, the Phase-I upgrade to LAr allows ATLAS to deal with the higher pileup conditions delivered by the LHC during Run 3.




Figure 4.9: Trigger rate comparison in simulation between Run2 Level-1 configuration (blue points) and Run-3 upgraded configuration (black and green triangles). For a fixed electron efficiency of 95%, the L1 rate with the Phase-I trigger is substantially lower for EM signatures with $E_T < 50$ GeV than the rate for the Run 2 trigger configuration [78].

4.4.5 TILE HADRONIC CALORIMETER

The Tile Calorimeter (TileCal) [79] sits directly outside of the LAr EM calorimeter. The TileCal covers $|\eta| < 1.0$ in the barrel region, and $0.8 < |\eta| < 1.7$ with its extended barrel region. The TileCal is a sampling calorimeter with steel as the absorber and scintillating tiles as the active medium, whose signals are read out by wavelength-shifting fibers into photomultiplier tubes. A schematic view of a TileCal slice is shown in Figure 4.10 The TileCal barrel and extended barrel sections are composed of 64 identical modules spaced evenly in the ϕ direction. Each Tile slice is segmented into three longitudinal layers, the first two of which have a granularity of $\Delta\eta \times \Delta\phi = 0.1 \times 0.1$, and the outermost of which has $\Delta\eta \times \Delta\phi = 0.2 \times 0.1$. The finer granularity in the innermost tile segments help resolve leakage from EM showers escaping the LAr calorimeter, which is beneficial for electron and photon identification.




Figure 4.10: A slice of the tile calorimeter, showing the interleaved steel and plastic scintillators and wavelength-shifting optical fibers. [7].

4.5 MUON SPECTROMETER

The Muon Spectrometer (MS), shown in Figure 4.11, is the outermost subsystem of the ATLAS detector, tasked with detecting muons which penetrate the barrel and end-cap calorimeters. The MS is submerged in a toroidal magnetic field which bends particle tracks, and provides measurements of momentum in the range $|\eta| < 2.7$, and triggering for $|\eta| < 2.4$. The MS consists of a barrel with three cylindrical stations at radii of $R = 5\text{m}$, 7.5m , and 10m respectively, and an endcap with inner, middle, and outer disks referred to as "wheels" as well as an extended endcap at distances of $|z| = 7.4\text{m}$, 10.8m , 14m , and 21.5m .

The three barrel stations use multilayered Monitored Drift Tube (MDT) chambers to provide two precision measurements each of the bending coordinate, with the outer and middle stations also using Resistive Plate Chambers (RPCs) to measure the azimuthal coordinate and for triggering.

In the endcap region, the outer wheel uses MDTs and the middle wheel uses both MDTs for precision tracking in the bending coordinate and Thin Gap Chambers (TGCs) for measuring the azimuthal coordinate. The inner wheels, also called "small wheels", sit within the barrel toroid and were completely replaced during LS2 by the New Small Wheel (NSW) system as part of the Phase-I muon upgrade.




Figure 4.11: Diagram of the ATLAS muon system, including the New Small Wheels for Run 3 [74].

Analagous to the LAr Phase-I upgrade, the objective of the MS upgrade during LS2 was to increase muon trigger background rejection to keep L1 rates manageable in the high-pileup environment of Run 3 without significantly raising muon trigger p_T thresholds. The upgrade was accomplished with the installation of the NSWs during LS2, which use two technologies: small-strip TGCs (sTGCs) and a micro-mesh gaseous structure (MicroMegas). Both technologies are designed with fast enough readout for use by the L1

trigger system, and give precision tracking in the bending direction as well as better resolution in the azimuthal direction with respect to the previous small wheels. These upgrades result in good p_T resolution in the trigger with increased background rejection, allowing for a low muon p_T threshold and manageable Level-1 trigger rate. A summary of the technologies used in the muon system and their function and coverage is given in Table 4.3.

MDT	
$ \eta $ coverage	$ \eta < 2.7$ (innermost layer: $ \eta < 1.3$)
Function	Precision tracking
sTGC	
$ \eta $ coverage	$1.3 < \eta < 2.7$ (2.4 for trigger)
Function	Trigger, precision tracking, 2nd coordinate
MicroMega	
$ \eta $ coverage	$1.3 < \eta < 2.7$ (2.4 for trigger)
Function	Trigger, precision tracking, 2nd coordinate
RPC	
$ \eta $ coverage	$-1.05 < \eta < 1.3$
Function	Trigger, 2nd coordinate
TGC	
$ \eta $ coverage	$1.05 < \eta < 2.7$ (2.4 for trigger)
Function	Trigger, 2nd coordinate

Table 4.3: Main parameters of the muon spectrometer [74].

4.6 FORWARD DETECTORS

In addition to the main detector subsystems described above, three smaller detectors sit in the forward region of ATLAS. Located ± 17 m from the interaction point, the LUCID (Luminosity measurement using Cerenkov Integrating Detector) [80] is used to detect inelastic $p - p$ scattering in the forward direction, providing a relative luminosity measurement. Located ± 240 m from from the interaction point, the ALFA (Absolute Luminosity For ATLAS) detector is made of scintillating fibers inside Roman Pots which approach as close as 1 mm to the LHC beam, and is also tasked with providing a luminosity measure-

ment. Finally, at ± 140 the ZDC (Zero-Degree Calorimeter) is used primarily for heavy-ion physics data taking. Alternating layers of quartz rods and tungsten plates measure neutral particles at $|\eta| > 8.2$ to determine the impact parameter of heavy-ion collisions.

4.7 MAGNET SYSTEM

The geometric layout of the ATLAS magnet system is shown in Figure 4.12. The magnet system is made of a solenoid magnet in the barrel region, whose field bends the trajectories of charged particles passing through the ID, as well as eight toroid magnets in the barrel and a set of eight toroid magnets in each endcap. The toroid fields bend the trajectories of charged particles passing through the MS.




Figure 4.12: The ATLAS magnet system, including barrel solenoid, eight barrel toroids, and two sets of eight endcap toroids [7].

The solenoidal field, provided by a single-layer coil wound with an Al-stabilised NbTi conductor, submerges the ID in a 2 T magnetic field pointing along the $+z$ axis. The solenoid magnet was carefully constructed to contribute only 0.66 radiation lengths at normal incidence before the face of the calorimeters. The construction was achieved, in

part, by housing the solenoid inside the same cryostat as the LAr calorimeter so that only a single vacuum is required, removing the need for two additional vacuum walls.

The toroid magnets provide a field of 1 T in the barrel and .5 T in the endcap. An Al-stabilised Nb/Ti/Cu conductor is wound into pancake-shaped coils and placed in a vacuum. The barrel toroids extend for 25.3 m along the z axis and have an inner diameter of 9.4 m and outer diameter of 20.1 m. The endcap toroids, each at $|z| = 5$ m have an inner diameter of 1.65 m and outer diameter of 10.7 m.

4.8 TRIGGER AND DATA ACQUISITION

The ATLAS detector was designed with the objective of probing rare physics processes. As an example, the Higgs boson production rate at $\sqrt{s} = 13.6$ TeV is 6 – 8 orders of magnitude smaller than for soft QCD processes, among other physics processes which one typically does not consider "interesting". The ATLAS trigger system is designed with the intent of identifying "interesting" physics to be collected by the Data Acquisition (DAQ) system for storage and later analysis, rejecting a large fraction of uninteresting $p - p$ -collision events.

The trigger system is two-layered, beginning with a Level-1 (L1) hardware-based trigger implemented on custom electronics, followed by a High-Level trigger (HLT) software-based system. The L1 system reduces the event rate from the $p - p$ -collision rate of 40 MHz to 100 kHz, with a fixed latency of $2.5\mu\text{s}/\text{event}$. The HLT further reduces the event rate from 100 kHz to 3 kHz averaged over an LHC fill, with a processing time of 400 ms for runs with a peak luminosity of $2 \times 10^{-34} \text{ cm}^{-2} \text{ s}^{-1}$.

The L1 trigger is further subdivided into two independent systems based on signals from either the calorimeter system (L1Calo) or muon detectors (L1Muon). The topological processor (L1Topo) uses kinematic information from physics objects reconstructed in L1Calo or L1Muon to make selections.

The TDAQ system is accompanied by a *trigger menu*, a list of rules each of which typically specify a physics object multiplicity, some quality cuts and p_T thresholds above which an event will be saved by the DAQ system. The menu is curated as a function of luminosity to maintain L1 and HLT rates at the levels specified above, while maintaining a varied mixture of "interesting" events for later analysis. An important charge of the TDAQ system in Run 3 is to maintain the trigger rates and performance in Run 2 while dealing with considerably higher pileup, which is achieved by the Phase-I upgrade briefly discussed in the next section.

A diagram of the TDAQ system in its Run 3 configuration is shown in Figure 4.13. The L1Calo system has been upgraded to handle the finer granularity inputs provided by LAr SCs described in Section 4.4.4. SC information, now fully digitized on-detector by LAr electronics, is passed to L1Calo EM Feature Extraction (eFEX) processors. Finer granularity inputs and shower shape variables allow configurable background rejection rates for the eFEX. The upgraded L1Calo system also includes jet (jFEX) and global (gFEX) feature extraction processors, which improve jet and E_T^{miss} trigger rates, purity, and efficiency with respect to Run 2.

The L1Muon system in Run 3 benefits from the replacement of the Small Wheels with the NSWs, and an upgrade to the L1 endcap trigger processor board which handles input from the muon system TGCs. The upgraded Endcap Sector Logic increases the number of TGC thresholds from 6 to 15 and supports attaching up to four quality flags to muon candidates, giving better resolution for the invariant mass of two muons by the L1Topo system. The results of the upgrade on electron and muon triggers in Run 3 are shown in Figure 4.14.

After commissioning and tuning of the electromagnetic objects for the new Run-3 L1Calo trigger in 2022 for use in 2023, a significant reduction in the L1 trigger rate of up to 10 kHz was achieved, as well as an improved efficiency for electron triggers requiring isolation. For triggering on single electrons with $p_T \geq 26$ GeV in the barrel, the eFEX with electron




Figure 4.13: Diagram of the ATLAS TDAQ system in Run 3, including the newly introduced eFEX, gFEX, and jFEX processors [73].




Figure 4.14: Trigger performance for electrons and muons following Phase-I upgrades (a) Efficiency of the L1 single-object EM trigger in the inner EM barrel within $|\eta| < 0.8$. The isolated legacy trigger and eFEX trigger with identification have the same rate. (b) Efficiency of L1 muon triggers in the endcap region for various p_T thresholds. [73]

identification implemented has the same rate as the legacy trigger with $p_T \geq 22$ GeV and an isolation requirement implemented. The "turn-on curve" for the eFEX trigger is much steeper, with a quick rise to 100% efficiency while the legacy trigger with isolation does not achieve the same level of efficiency for electrons below 50 GeV.

CHAPTER 5:

OBJECT RECONSTRUCTION

The subsystems of ATLAS described in the previous chapters record the passage of particles through the detector as a collection of many granular electrical signals. A complex process then takes place to combine these signals such that the physical particles created in each $p - p$ collision can be *reconstructed* for use in physics analysis. Figure 5.1 shows an example of a few *physics objects* (such as electrons, photons, muons, E_T^{miss} , jets) being reconstructed by combining signatures of particles passing through the ATLAS detector. The algorithms responsible for the reconstruction process are the culmination of many years of work by the ATLAS collaboration to ensure high-quality physics objects are available for analysis. These algorithms are the focus of this chapter.

The projective geometry of the ATLAS detector is optimized for particles which are created nearly instantaneously after the $p - p$ collision has occurred near the center of the detector. Analogously, standard ATLAS object reconstruction algorithms are typically biased towards providing quality reconstruction of objects whose parameters for tracks/shower shapes are consistent with originating very near to the hardest $p - p$ scatter in a given bunch crossing.

The results presented in this thesis target electrons and muons from displaced decays of LLPs. Specialized reconstruction is necessary to avoid the bias introduced in standard reconstruction. In addition, if an electron is sufficiently displaced, ATLAS may recon-




Figure 5.1: Diagram demonstrating the reconstruction of a number of physics objects through a combination of signatures from various subsystems of ATLAS [81].

struct a photon rather than an electron object; the results in this thesis also treat this case. This chapter is structured as follows. Section 5.1 discusses reconstruction of the primary vertex (PV) which defines the position of the $p - p$ collision of interest in each event. In Sections 5.2 - 5.3, the calorimeter *clusters* and *tracks* used to form physics objects are described. A special extension to standard tracking known as Large-Radius Tracking (LRT), used to boost sensitivity for LLP decays to charged particles in the tracker, is discussed in Section 5.3.3. A description of reconstruction algorithms used for each physics object used in the analysis is given, in Section 5.4 (Electrons and Photons), Section 5.5 (Muons), Section 5.7 (Jets), and Section 5.10 (E_T^{miss}).

5.1 PRIMARY VERTEX RECONSTRUCTION

The PV is the vertex corresponding to the hardest scattering $p - p$ interaction in a given bunch crossing. The PV is a defining characteristic of an ATLAS event: most "prompt" ATLAS measurements and searches aim to reconstruct and identify physics objects which originate very near to the PV, while this analysis searches for new particles which travel a measurable distance from the PV before they decay, resulting in an electron or muon. To reconstruct the PV, ID tracks (discussed in detail in Section 5.3) are first used in a vertex-finding algorithm [82] which associates groups of tracks to PV candidates. The PV is selected as the candidate with the largest value of $\sum p_T^2$, summed over the associated tracks. The remaining PV candidates typically correspond to inelastic proton scatters and are referred to as *pileup vertices*.

5.2 CALORIMETER CLUSTERING

As it traverses the ATLAS detector A single particle created in a collision event may deposit energy in many calorimeter cells longitudinally and laterally. For reconstruction of both EM and hadronic objects A topological clustering algorithm [83] is applied to create groups of cells and to sum the total energy deposited in a given cluster. The goal of the algorithm is to form groups from neighboring clusters that each have significant energy (E_{cell}) compared to the expected level of electronic noise and pileup in a calorimeter cell ($\sigma_{\text{cell}} \equiv \sigma_{\text{pileup}} \oplus \sigma_{\text{noise}}$). The defining parameter dictating the growth of a cluster is the *cell significance* $\zeta_{\text{cell}} \equiv |E_{\text{cell}}|/|\sigma_{\text{cell}}|$.

Topo-cluster building for an event begins with the formation of *proto-clusters*, each seeded by a cell with $\zeta_{\text{cell}} \geq 4$. In the growth iterations which follow, cells neighboring the proto-cluster are added if they have $\zeta_{\text{cell}} \geq 2$. During this process, if proto-cluster overlap, they are merged. Once no neighboring cells with $\zeta_{\text{cell}} \geq 2$ remain, cells at the

perimeter of each cluster are added if $\zeta_{\text{cell}} \geq 0$. The algorithm is often referred to as 4-2-0 topo-clustering due to the choice of seed-growth-perimeter parameter values, optimized to efficiently find low-energy clusters without being overwhelmed by noise.

Following this process, the resulting proto-clusters are checked for multiple local maxima. A local maximum is defined here as a cell with $E_{\text{cell}} > 500$ MeV and at least four neighbors also within the proto-cluster, none of which have a larger E_{cell} . Multiple local maxima could indicate that the proto-cluster contains multiple close-by physics objects. In such a case, the proto-cluster is split into multiple clusters, which can in principle share a cell. The energy of any shared cell is split among the clusters based on their relative energies. The proto-clusters remaining after this algorithm are the topo-clusters which are used in downstream physics object reconstruction. For each topo-cluster, the fraction of energy deposited in the LAr calorimeter with respect to the total cluster energy calculated from both the Tile and LAr calorimeters (f_{EM}) is calculated. This is later used as a discriminant when selecting topo-clusters for electron and photon reconstruction, as described in Section 5.4.1.

5.3 TRACKING

Tracks are formed by algorithms that run on hits left by charged particles traversing the ID. Tracks provide precision measurements of the momenta and direction of charged particles, and are uniquely determined by five parameters. The transverse and longitudinal impact parameters (d_0 and z_0 , respectively) denote the distance of closest approach in the transverse and longitudinal plane to the beamspot position. The other three parameters are the longitudinal and polar angle of the track momentum (ϕ, θ) and the ratio of the track charge to its magnitude of momentum ($q/|\vec{p}|$). A diagram of a track including these parameters is shown in Figure 5.2.




Figure 5.2: Diagram of a track object formed by ATLAS reconstruction software, showing defining parameters $(d_0, z_0, \phi, \theta, \vec{p})$. The charge q , determined from the curvature of the track, is also a defining parameter [84].

Tracking in ATLAS proceeds in two steps [85]: *inside-out* tracking which uses 3D space-points (SPs) from measurements in either the Pixel or SCT detectors, followed by *outside-in* tracking which attempts to extend the tracks from the inside-out pass using information from the TRT and the LAr calorimeter.

5.3.1 INSIDE-OUT TRACKING

The first tracking pass begins with the formation of track seeds formed from three hits in different layers of the pixel and/or SCT detectors which are consistent with having originated from a charged particle track. Cuts on d_0 and z_0 are made at this stage to reduce computational time, which introduces a bias towards the hypothesis that the candidate particle originated near the PV. A combinatorial Kalman Filter [86] is applied to extend these seeds inwards and outwards radially, within a narrow ΔR window. Seeds failing this step are checked for ΔR compatibility with calorimeter clusters. For compatible clusters, the procedure is redone allowing for a kink in the track. With this procedure, some electrons whose tracks undergo a momentum and direction change due to Bremsstrahlung radiation can be recovered.




Figure 5.3: ATLAS track reconstruction (a) inside-out pass, showing track seed, extension window, and result of global fit [84] (b) outside-in pass, showing ROI based on calorimeter signatures (green bands) encompassing TRT hits and search windows for extension into SCT [87].

The surviving track candidates then undergo an ambiguity resolution step, which scores tracks on a number of quality criterion and rejects low-quality tracks sharing hits with tracks that are higher-ranked. Following this overlap removal, a global χ^2 fit is performed to obtain high-precision track parameters.

5.3.2 OUTSIDE-IN TRACKING

Following the primary track reconstruction pass described above, a secondary pass begins in Regions of Interest (ROIs) which are windows of ΔR defined by energy deposits in the EM calorimeter. Tracks are seeded by segments in the TRT consistent with originating in an ROI. The seeds are extended into the SCT, using the same Kamlan Filter, ambiguity resolver, and global χ^2 fit as in the inside-out pass. Diagrams describing the standard tracking procedure for both inside-out and outside-in are shown in Figure 5.3.

5.3.3 LARGE RADIUS TRACKING

Large Radius Tracking (LRT) is an additional tracking step, run on hits leftover from the standard tracking passes, which aims to improve sensitivity with respect to standard

tracking for LLP scenarios [88]. This algorithm re-runs the inside-out track reconstruction pass on ID hits which remain unassociated to a track after the standard tracking pass. A number of parameters are changed in this LRT tracking pass, as summarized in Table 5.1. Notably, the allowed range of $|d_0|$ is extended from 5 mm to 300 mm, and $|z_0|$ from 200 mm to 500 mm.

Selection criteria	Primary (<i>inside-out</i>)	LRT
max. $ d_0 $ [mm]	5	300
max. $ z_0 $ [mm]	200	500
min. p_T [GeV]	0.5	1
max. $ \eta $	2.7	3.0
max. silicon holes	2	1
max. double holes	1	0
max. holes gap	2	1
road width [mm]	12	5
seeding	Pixels and SCT	SCT only
max. seeds per middle Pixel SP	1	—
max. seeds per middle SCT SP	5	1
Common selection criteria		
min. silicon hits	8	
min. unshared silicon hits	6	
max. track χ^2/n_{DoF}	9	
keep all confirmed seeds	true	

Table 5.1: Comparison of parameter values used in standard inside-out tracking and LRT algorithms [88].

The specialized LRT track reconstruction is a key input to this result, as the use of the LRT track collection to reconstruct electrons and muon objects gives vastly improved efficiency with respect to standard tracking for displaced leptons. However, the relaxation of $|d_0|$ and $|z_0|$ cuts also introduces large backgrounds, as this greatly increases the number of fake tracks that can be reconstructed from random noise and pileup hits in the ID. The associated increase in computational complexity and event size prevented LRT tracking from being run as a default tracking algorithm during Run 2. Instead, in the Run 2 data used in this result, a trigger and event topology filter is applied to preselect approximately 10% of events on which to run LRT tracking. Between Run 2 and Run 3, improvements were made to the LRT tracking algorithm which allowed for an approximately 4000% re-




Figure 5.4: ATLAS LRT tracking performance improvements for Run 3. (a) Disk space reduction in updated LRT tracking reconstruction for Run 3 compared to Run 2. (b) Efficiency for benchmark LLP signal and fake rejection rate compared between Run 2 LRT tracking configuration and Run 3, showing drastic improvement at large radii in the new configuration [88].

duction in event size, and 95% improvement in background rejection at large radii, without any sacrifice in efficiency for a benchmark LLP scenario, as shown in Figure 5.4.

The improvements to the LRT tracking algorithm allow for LRT tracking to be run as part of the default reconstruction chain in ATLAS, therefore allowing LRT to run on all events in Run 3, rather than the subset used in Run 2. Crucially, LRT tracking can now be run at the ATLAS trigger level. New triggers in Run 3 for electrons and muons reconstructed with LRT tracks are a key component of this analysis, and are discussed in more detail in Chapter 8.

5.4 ELECTRONS AND PHOTONS

Electrons and photons leave similar showers in the LAr calorimeter. The main discriminant between their signatures is the presence of tracks in the ID in the case of an electron. Complications in this simple classification arise due to the conversion of pho-

tons to electron-positron pairs in the ID, resulting in a *conversion vertex*, and the radiation of photons from electrons due to Bremsstrahlung radiation. The following sections describe the *superclustering* algorithm which runs on input topo-clusters to reconstruct photons and electrons, and the subsequent algorithms for identification of electrons and photons. A detailed description is given for the electron identification algorithm, which requires modification in this analysis to increase the efficiency for displaced electrons.

5.4.1 SUPERCLUSTERING

Electron and photon reconstruction begins with topo-clusters that have energy in the EM calorimeter greater than 400 MeV and f_{EM} greater than 0.5. The full list of topo-clusters matching this criterion are tested one-by-one in order of descending E_{T} for use as supercluster seeds. Supercluster reconstruction proceeds separately for electrons and photons, with different minimum E_{T} requirements and the key requirement in the electron case of a ΔR -matched track with at least four silicon hits.

An attempt is then made to add *satellite clusters* to each seed cluster. For both electrons and photons, all clusters within a window of $\Delta\eta \times \Delta\phi = 0.075 \times 0.125$ of the supercluster are added. For electrons only, clusters matched to the same track as the seed within $\Delta\eta \times \Delta\phi = 0.075 \times 0.300$ are added. For photon superclusters matched to conversion vertices only, clusters matched to the same vertex are added. The algorithm is represented diagrammatically in Figure 5.5.

It is possible at this step to have a photon supercluster which shares a seed cluster with an electron supercluster. An ambiguity resolution step is therefore applied: if an object has no good track attached, it is unambiguously labeled as a photon, and if it has a high-quality track attached and no conversion vertex attached, it is unambiguously labeled as an electron. In any other case, both an electron and a photon are created and tagged as ambiguous for later classification. In this analysis, electrons are favored over photons during the overlap-removal phase described in Section 5.9, as true physical leptons are




Figure 5.5: Diagram of the superclustering algorithm for e/γ object reconstruction [89].

the target of the search (photon objects are used only to recover acceptance for highly displaced electrons).

In addition to the electrons reconstructed using the collection of tracks output by the standard tracking algorithm, this analysis also makes use of electrons reconstructed using the track collection output by the LRT algorithm. The track-cluster matching is performed in the same way as for the standard electron reconstruction described above, and uses the same set of superclusters. The reconstructed electrons resulting from this algorithm are referred to as LRT electrons.

5.4.2 PHOTON IDENTIFICATION

This analysis makes use of standard photon objects reconstructed by ATLAS to increase acceptance for electron objects with sufficient displacement that an ID track is not reconstructed. The identification of photons in ATLAS is well-studied and three standardized working points (WPs) which give fixed efficiencies for benchmark samples are provided: Loose, Medium, and Tight. The WPs are defined by a set of rectangular cuts on shower shape variables listed in Table 5.2. These variables can discriminate photons originating from the PV against non-prompt photons from neutral hadron decays within jets, as well as from jets with a large fraction of energy in the EM calorimeter which might "fake" pho-

Category	Description	Name	Usage
Hadronic leakage	Ratio of E_T in the first layer of the hadronic calorimeter to E_T of the EM cluster (used over the ranges $ \eta < 0.8$ and $ \eta > 1.37$)	R_{had1}	e/γ
	Ratio of E_T in the hadronic calorimeter to E_T of the EM cluster (used over the range $0.8 < \eta < 1.37$)	R_{had}	e/γ
EM third layer	Ratio of the energy in the third layer to the total energy in the EM calorimeter	f_3	e
EM second layer	Ratio of the sum of the energies of the cells contained in a $3 \times 7 \eta \times \phi$ rectangle (measured in cell units) to the sum of the cell energies in a 7×7 rectangle, both centred around the most energetic cell	R_η	e/γ
	Lateral shower width, $\sqrt{(\sum E_i \eta_i^2)/(\sum E_i) - ((\sum E_i \eta_i)/(\sum E_i))^2}$, where E_i is the energy and η_i is the pseudorapidity of cell i and the sum is calculated within a window of 3×5 cells	$w_{\eta 2}$	e/γ
EM first layer	Ratio of the sum of the energies of the cells contained in a $3 \times 3 \eta \times \phi$ rectangle (measured in cell units) to the sum of the cell energies in a 3×7 rectangle, both centred around the most energetic cell	R_ϕ	e/γ
	Total lateral shower width, $\sqrt{(\sum E_i (i - i_{max})^2)/(\sum E_i)}$, where i runs over all cells in a window of $\Delta\eta \approx 0.0625$ and i_{max} is the index of the highest-energy cell	$w_{s\ tot}$	e/γ
	Lateral shower width, $\sqrt{(\sum E_i (i - i_{max})^2)/(\sum E_i)}$, where i runs over all cells in a window of 3 cells around the highest-energy cell	$w_{s\ 3}$	γ
	Energy fraction outside core of three central cells, within seven cells	f_{side}	γ
	Difference between the energy of the cell associated with the second maximum, and the energy reconstructed in the cell with the smallest value found between the first and second maxima	ΔE_s	γ
	Ratio of the energy difference between the maximum energy deposit and the energy deposit in a secondary maximum in the cluster to the sum of these energies	E_{ratio}	e/γ
	Ratio of the energy measured in the first layer of the electromagnetic calorimeter to the total energy of the EM cluster	f_1	e/γ
	Ratio of the energy measured in the first layer of the electromagnetic calorimeter to the total energy of the EM cluster	f_1	e/γ
Track conditions	Number of hits in the innermost pixel layer	$n_{innermost}$	e
	Number of hits in the pixel detector	n_{pixel}	e
	Total number of hits in the pixel and SCT detectors	n_{Si}	e
	Transverse impact parameter relative to the beam-line	d_0	e
	Significance of transverse impact parameter defined as the ratio of d_0 to its uncertainty	$ d_0/\sigma(d_0) $	e
Track-cluster matching	Momentum lost by the track between the perigee and the last measurement point divided by the momentum at perigee	$\Delta p/p$	e
	Likelihood probability based on transition radiation in the TRT	eProbabilityHT	e
	$\Delta\eta$ between the cluster position in the first layer of the EM calorimeter and the extrapolated track	$\Delta\eta_1$	e
	$\Delta\phi$ between the cluster position in the second layer of the EM calorimeter and the momentum-rescaled track, extrapolated from the perigee, times the charge q	$\Delta\phi_{res}$	e
	Ratio of the cluster energy to the measured track momentum	E/p	e

Table 5.2: Discriminating variables used for electron and photon identification. The "Usage" column indicates if the variables are used for the identification of electrons, photons, or both. [89].

tons. Loose ID uses R_{had1} , R_{had} , R_η , R_ϕ , $w_{\eta 2}$ (see Table 5.2 for definitions) variables calculated using the LAr calorimeter's second layer. Medium ID adds a cut on E_{ratio} , used to reject showers with multiple maxima in the strip layer indicative of $\pi^0 \rightarrow \gamma\gamma$ decays. Tight ID adds a number of additional cuts on variables calculated in the strip layer (the remaining listed variables in Table 5.2 with "Usage" indicating photon ID). The cuts used in Tight ID have been found previously [90–92] to significantly bias against displaced objects due to the unconventional shower shapes expected to arise in displaced decays. Therefore Medium

ID is selected for use in this analysis, to achieve a reasonably pure photon sample without sacrificing efficiency for displaced signal electrons reconstructed as photon objects.

5.4.3 ELECTRON IDENTIFICATION

Electrons objects are classified based on quality using a likelihood (LH) based identification [93]. The LH function is constructed from probability distribution functions (PDFs) of a collection of track and calorimeter variables for signal and background samples simulated in MC. The list of inputs is included as those labeled "LH" in Table 5.3. Rectangular cuts on a few variables are also used in electron identification; these are labeled in Table 5.3 as "C". An LH distribution is constructed for a signal sample ($\mathcal{L}_S$) and a background sample ($\mathcal{L}_B$). The signal sample consists of prompt electrons from simulated $Z \rightarrow ee$ events, and the background contains prompt electrons resulting from mis-reconstruction of hadronic jets or photons. The LH distribution is constructed as:

$$\mathcal{L}_{S(B)}(\vec{x}) = \prod_{i=1}^N P_{S(B),i}(x_i) \quad (5.1)$$

That is, for a collection of input values $\vec{x}$, the value of the LH is the multiplication of the PDF output $P_i(x_i)$ for each input variable ($i \dots N$). The signal and background LH distributions are used to construct a discriminant $d_{\mathcal{L}} = \mathcal{L}_S / (\mathcal{L}_S + \mathcal{L}_B)$ which peaks at 1(0) for signal(background)-like electrons. This sharply peaked function is then transformed with a sigmoid to produce the final identification output variable $d'_{\mathcal{L}}$. This procedure is performed in bins of $(E_T, |\eta|)$. An example $d'_{\mathcal{L}}$ distribution is shown in Figure 5.6. Standardized WPs are defined based on cuts on the transformed likelihood, corresponding to fixed efficiencies for identifying an electron with $E_T = 40$ GeV: Loose (93%), Medium(88%), Tight(80%). A VeryLoose WP is also defined for specialized trigger studies. A modified version of the VeryLoose working point, described in Section 5.6, is used in this analysis.

Type	Description	Name	Rejects			Usage
			LF	γ	HF	
Hadronic leakage	Ratio of E_T in the first layer of the hadronic calorimeter to E_T of the EM cluster (used over the range $ \eta < 0.8$ or $ \eta > 1.37$)	R_{had1}	x	x		LH
	Ratio of E_T in the hadronic calorimeter to E_T of the EM cluster (used over the range $0.8 < \eta < 1.37$)	R_{had}	x	x		LH
Third layer of EM calorimeter	Ratio of the energy in the third layer to the total energy in the EM calorimeter. This variable is only used for $E_T < 80 \text{ GeV}$, due to inefficiencies at high E_T , and is also removed from the LH for $ \eta > 2.37$, where it is poorly modelled by the simulation.	f_3	x			LH
Second layer of EM calorimeter	Lateral shower width, $\sqrt{(\Sigma E_i \eta_i^2)/(\Sigma E_i) - ((\Sigma E_i \eta_i)/(\Sigma E_i))^2}$, where E_i is the energy and η_i is the pseudorapidity of cell i and the sum is calculated within a window of 3×5 cells	$w_{\eta 2}$	x	x		LH
	Ratio of the energy in 3×3 cells over the energy in 3×7 cells centred at the electron cluster position	R_ϕ	x	x		LH
	Ratio of the energy in 3×7 cells over the energy in 7×7 cells centred at the electron cluster position	R_η	x	x	x	LH
First layer of EM calorimeter	Shower width, $\sqrt{(\Sigma E_i (i - i_{\text{max}})^2)/(\Sigma E_i)}$, where i runs over all strips in a window of $\Delta\eta \times \Delta\phi \approx 0.0625 \times 0.2$, corresponding typically to 20 strips in η , and i_{max} is the index of the highest-energy strip, used for $E_T > 150 \text{ GeV}$ only	w_{atot}	x	x	x	C
	Ratio of the energy difference between the maximum energy deposit and the energy deposit in a secondary maximum in the cluster to the sum of these energies	E_{ratio}	x	x		LH
	Ratio of the energy in the first layer to the total energy in the EM calorimeter	f_1	x			LH
Track conditions	Number of hits in the innermost pixel layer	n_{Blayer}		x		C
	Number of hits in the pixel detector	n_{Pixel}		x		C
	Total number of hits in the pixel and SCT detectors	n_{Si}		x		C
	Transverse impact parameter relative to the beam-line	d_0		x	x	LH
	Significance of transverse impact parameter defined as the ratio of d_0 to its uncertainty	$-d_0/\sigma(d_0)$		x	x	LH
	Momentum lost by the track between the perigee and the last measurement point divided by the momentum at perigee	$\Delta p/p$	x			LH
TRT	Likelihood probability based on transition radiation in the TRT	eProbabilityHT	x			LH
Track-cluster matching	$\Delta\eta$ between the cluster position in the first layer and the extrapolated track	$\Delta\eta_1$	x	x		LH
	$\Delta\phi$ between the cluster position in the second layer of the EM calorimeter and the momentum-rescaled track, extrapolated from the perigee, times the charge q	$\Delta\phi_{\text{res}}$	x	x		LH
	Ratio of the cluster energy to the track momentum, used for $E_T > 150 \text{ GeV}$ only	E/p	x	x		C

Table 5.3: Type and description of the quantities used in the electron identification. The columns labelled "Rejects" indicate whether a quantity has significant discrimination power between prompt electrons and light-flavour (LF) jets, photon conversions (γ), or non-prompt electrons from the semileptonic decay of hadrons containing heavy-flavour (HF) quarks (b- or c-quarks). In the column labelled "Usage," an "LH" indicates that the pdf of this quantity is used in forming the likelihood and a "C" indicates that this quantity is cut on directly as a selection criterion [93].

5.5 MUONS

Muons are reconstructed using a combination of information from one or more sub-systems: the MS, the calorimeters, and the ID [94]. The only type of reconstructed muon object relevant to this analysis is a Combined (CB) Muon, reconstructed by combining MS information with ID tracks.

Hits in the MS are combined using a Hough Transform [95] to form *muon segments* within each MS layer. An attempt is then made to combine segments from different MS




Figure 5.6: The standard LH-based identification discriminant for reconstructed electron candidates with good quality tracks with $30 \text{ GeV} < E_T < 35 \text{ GeV}$ and $|\eta| < 0.6$. The black histogram is for prompt electrons in a $Z \rightarrow ee$ simulation sample, and the red (dashed-line) histogram is for backgrounds that may mimic the signature of prompt electrons. This includes simulation samples of multijet production ($qg \rightarrow q\gamma$ or $q\bar{q} \rightarrow g\gamma$) and top-quark production [93].

layers to form MS tracks. An ambiguity resolution is performed, followed by a global χ^2 fit to form the final MS track collection. ID tracks within a narrow ΔR window of an MS track are combined with the track, and another global χ^2 fit is performed to account for ionization loss in the calorimeters. The resulting ID-MS track combination is a CB muon object.

Analogous to LRT electrons, this analysis also makes use of muons reconstructed using the track collection output by the LRT algorithm. All MS hits in an event are used alongside the LRT track collection, to increase sensitivity to LLPs decays resulting in displaced muons. The resulting muons are referred to as LRT muons.

5.6 MODIFIED LEPTON IDENTIFICATION

In addition to LRT tracking as an input to electron and muon reconstruction, this analysis uses modifications to the standardized ATLAS identification criteria for both electrons

and muons. In both cases, this is done to avoid the bias in typical ATLAS reconstruction towards prompt physics objects which originate very near to the PV.

The rectangular cuts and LH used in standard electron identification described in Section 5.4.3 introduce significant bias towards prompt electrons. To reduce this bias, the requirement that a pixel hit be present in the electron track is removed, as sufficiently displaced signal electrons may arise from decays beyond the outermost pixel layer. All variables involving $|d_0|$ are removed from the electron LH inputs, as the inclusion of $|d_0|$ causes the LH algorithm to "learn" that high-quality electrons should have very small $|d_0|$ values, which is not the case for high-quality displaced electrons.

The muon identification WP used in this analysis is a modification of the standardized Medium WP. The Medium WP accepts only CB muons for $|\eta| < 2.5$ with measurements in at least two MDT stations. For $|\eta| < 0.1$, muons with only one precision station are accepted if they have at most one *hole*, or missing layer in the track. In the ID, at least one pixel hit and at least 5 SCT hits are required on the ID track. Fewer than 3 holes are required in the SCT. A compatibility also is required between the $q/|\vec{p}|$ measurement from the independent track fits in the MS and the ID. For the purposes of this analysis, to remove the bias against displaced decays, the Medium requirement of at least one pixel hit on the muon track is removed.

The result of adding the LRT collection to lepton reconstruction and modifying lepton identification criteria is demonstrated for a benchmark LLP signal sample in Figure 5.7. Significant improvement in reconstruction and identification is observed, in particular for simulated electrons and muons with $|d_0| > 5$ mm.

5.7 JETS

A *jet* is a collimated spray of particles arising from the fragmentation and subsequent hadronization of quarks or gluons. A diagram of the hadronization process is shown in




Figure 5.7: Efficiency to reconstruct and identify electrons (left) and muons (right) using LRT and standard tracking (Std.) with respect to MC generator-level leptons in a simulated GMSB slepton signal models, namely with a slepton mass and lifetime of 300 GeV and 1 ns, respectively. The efficiencies for the various LLP WP are shown as a function of the true transverse impact parameter $|d_0|$ of the daughter lepton.

Figure 5.8. Jets can contain many constituent particles, which can produce tracks in the ID and deposit energy in the calorimeter system. Jet objects are used only indirectly in this analysis, in the object overlap removal procedure described in Section 5.9 and in the calculation of E_T^{miss} described in Section 5.10.

Jets are seeded by topo-clusters that may be matched to ID tracks. Topo-clusters and tracks provide complementary inputs to jet reconstruction: tracks are only formed by charged particles, while topo-clusters result from both neutral and charged particles. The tracker covers a smaller solid angle ($|\eta| < 2.47$) compared to the calorimeter ($|\eta| < 4.9$), but gives better spatial resolution than the calorimeters, as well as better energy resolution in the low-energy regime. A particle flow (PFlow) algorithm [96] is used to exploit this complementarity, combining calorimeter and tracking information when reconstructing jets. The algorithm begins with a set of nearby tracks and topo-clusters. For selected tracks, the algorithm predicts the expected single topo-cluster energy using track momentum

and track-cluster ΔR . Successfully matched clusters will have their energies subtracted from the total jet energy, with the goal of creating a set of tracks from charged particles and clusters from neutral particles only. The resulting object is known as a Particle Flow Object (PFO).

The PFOs in each event are used as inputs to a sequential recombination algorithm known as the anti- k_T [97] algorithm. The algorithm iterates over pairs of PFOs, using their p_T and ΔR to evaluate whether they should be merged into a single jet, with the parameter controlling the radius of the resulting jet fixed to $R = 0.4$. The algorithm is biased against combining low-energy PFOs with one another, preferring to combine low-energy PFOs with high-energy PFOs, an advantage which results in collimated jets surrounding the highest-energy particles formed from the hadronization process.

5.8 ISOLATION

The photon, electron, and muon objects used in this analysis are required to be *isolated* from nearby tracker or calorimeter activity. This requirement serves primarily to reduce backgrounds from heavy flavor decays. For the calorimeter, the variable $E_T^{\text{cone}(X)}$, is defined as the net E_T deposited with $\Delta R < X/100$ from a given object. For the tracker, $p_T^{\text{varcone}(X)}$ is the sum of momenta of $p_T > 1$ GeV tracks within $\Delta R < \min(10 \text{ GeV}/p_T, X/100)$ from a given object. The "var" in the term "varcone" indicates that the isolation cone has a variable radius; for low- p_T objects it has a maximum size of $\Delta R = 0.3$, which shrinks with increasing p_T . Fixed-radius calorimeter isolation, denoted as $p_T^{\text{cone}(X)}$, is also defined. A PFO based isolation variable, labelled $E_T^{\text{neflow}(X)}$, is defined as the sum of transverse energy of neutral PFOs with $\Delta R < (X/100)$ from a given physics object. The isolation requirements used here are all based on standardized ATLAS WPs, and are as follows:

- **Electrons** $E_T^{\text{cone20}}/p_T < 0.06$ and $p_T^{\text{varcone30}} p_T < 0.06$ (Tight_VarRad).
- **Photons** $E_T^{\text{cone20}}/p_T < 0.065$ and $p_T^{\text{cone20}} p_T < 0.05$ (FixedCutLoose).




Figure 5.8: A quark produced in the $p - p$ collision undergoes fragmentation and hadronization, resulting in a jet which can be measured by ATLAS [98].

- **Muons** $(p_T^{\text{varcone30}} + 0.4 \times E_T^{\text{neflow20}})/p_T < 0.16$ (PflowLoose_VarRad).

In all cases, the isolation WP is chosen to reduce backgrounds as much as possible while maintaining near-100% signal efficiency.

5.9 OVERLAP REMOVAL

The algorithms discussed thus far all run independently; therefore the objects which are produced could share track hits or calorimeter clusters and as a result have very small ΔR separation. A standard overlap removal procedure is performed to ensure that physics objects are not double-counted. This overlap removal must be done upstream of the E_T^{miss} calculation, again to ensure that the resulting E_T^{miss} value is not affected by double-counting of transverse energies of measured objects in the event.

Due to the inclusion of LRT electrons and LRT muons in this result, an additional specialized overlap removal was designed. In the case of both electrons and muons, two collections are created by the algorithms described above, one containing physics objects associated to standard tracks and one containing physics objects associated to tracks reconstructed with LRT. Standard tracking and LRT are by definition run on different sets of hits and therefore contain orthogonal sets of tracks, however the track-cluster matching to reconstruct electrons runs on each of these collections separately. As a result, it is

possible to have two electrons from different collections in an event share a cluster. If a standard and large radius track both point to the same EM cluster, the electron passing the stricter identification (modified as in Section 5.6) requirement is kept. If both candidate electrons pass the same level of identification, the electron reconstructed with a standard track is kept. This approach also applies to muon reconstruction, where standard and LRT muons cannot share ID tracks by definition, but can share MS track segments. If a standard and LRT track both point to the same muon segment and lead to the reconstruction of a standard and LRT muon with the same segment, the duplication is resolved by keeping the muon with the smaller difference in η between the ID track and the muon track segment extrapolated to the PV.

The recommended overlap removal by the ATLAS collaboration for analyses using photons and boosted leptons, using the merged electron and muon collections as input, then proceeds as follows [99]. Electrons sharing an ID track with a muon are removed. Photons within $\Delta R < 0.4$ of an electron or muon are removed. Jets residing within $\Delta R < 0.2$ of electrons are removed. Electrons within $\Delta R < \min(0.4, 0.04 + 10 \text{ GeV}/p_T)$ of remaining jets are then removed. Jets within $\Delta R < 0.2$ of a muon with less than three matched ID tracks are removed. Muons are then removed if they are within $\Delta R < \min(0.4, 0.04 + 10 \text{ GeV}/p_T)$ of remaining jets. Jets within $\Delta R < 0.4$ of a photon are then removed.

5.10 MISSING TRANSVERSE ENERGY

The electron, photon, muon, and jet objects that remain after the overlap removal procedure contribute to the calculation of $E_T^{\text{miss}} \equiv -|\Sigma \vec{p}_T|$, the negative vector sum of energy measured in the transverse plane first described in Section 4.1. Neutrinos, which pass through the detector without interacting, are an SM source of E_T^{miss} . Theorized BSM particles such as the gravitino in the benchmark SUSY model studied here are also potential

sources of E_T^{miss} . In practice, E_T^{miss} can also arise from imperfect calibrations, holes in detector coverage, or other instrumental effects.

The so-called standard ATLAS track soft term (TST) based E_T^{miss} calculation [100] is used in this analysis. In addition to the "hard term" from the vector sum of the p_T of each group of objects surviving the overlap removal, a "soft term" from unused tracks is added:

$$E_T^{\text{miss}} \equiv - \left| \underbrace{\Sigma \vec{p}_T^\gamma + \Sigma \vec{p}_T^\mu + \Sigma \vec{p}_T^e + \Sigma \vec{p}_T^{\text{jets}}}_{\text{hard term}} + \underbrace{\Sigma \vec{p}_T^{\text{unused tracks}}}_{\text{soft term}} \right| \quad (5.2)$$

An ambiguity resolution procedure is used to prevent double counting between the soft term and the terms from reconstructed objects.

CHAPTER 6:

BOOSTED DECISION TREES AND TIMING CALIBRATION

As motivated in Chapter 5, searches for LLPs often require dedicated reconstruction algorithms to provide necessary measurements of unique signatures. However, specialized algorithms such as LRT tracking typically trade increased LLP efficiency for large increases in background (see Section 5.3.3) which require custom variables to suppress. This chapter focuses on the development of two custom variables essential to the displaced electron search in the results presented.

To isolate the signature of a displaced and delayed electron from prompt backgrounds, the ID provides precision spatial information while the LAr calorimeter provides timing and additional spatial information. Careful synthesis of LLP-sensitive variables from the standard set of ID and LAr measurements provided by ATLAS is required to ensure optimal sensitivity to the signature under study. To achieve this goal, a multivariate analysis (MVA) method known as a Boosted Decision Tree (BDT) is used to create a hyper-variable which provides more sensitivity than rectangular cuts on individual ID and LAr variables. Section 6.1 gives a general overview of BDTs, and Sections 6.2 - 6.5 are dedicated to the BDTs implemented for this search.

In addition, taking inspiration from Refs. [3, 4, 90, 91], a calibration procedure is performed on the timing measurement provided by the LAr calorimeter, resulting in a res-

olution capable of distinguishing delayed EM energy clusters from prompt background. In this analysis, the timing calibration procedure is applied to ATLAS data from Run 3 for the first time. Also for the first time, the timing calibration is used to search for a signature including displaced electrons in the tracker, whereas previous searches looked only for photons or electrons with sufficient displacement to be reconstructed by ATLAS as photons. Section 6.6.2 describes the timing calibration procedure and the resulting LLP-sensitive calibrated timing measurement.

6.1 BOOSTED DECISION TREES

A BDT is a Machine Learning (ML) algorithm applicable to a wide class of classification and regression tasks. The overview given here of decision trees, and the practice of *boosting*, are summarized from the discussion in Ref. [101].

The basic building block of a BDT is a classification decision tree. Such a decision tree takes as input an instance of data with any number of attributes. Based on these attributes, the algorithm will return a "score" corresponding to the probability that the instance belongs to one or more classes. In the case of this thesis, each instance corresponds to one event and the associated attributes are either electron, photon, or event-level variables. The classification trees in this thesis assign a score based on two classes: *signal-like* events are those with properties consistent with LP decays to electrons, and *background-like* events are consistent with SM processes.

Like many ML algorithms, a decision tree is "grown" during a *training* phase, where an attempt is made to classify instances by minimizing an objective function known as *impurity*. The tree will begin at a *root node*, and based on one of the attributes of an input instance will choose a split value that creates maximum separation between classes, therefore minimizing impurity.¹ Based on this split value, two more nodes are created

¹Typically the choice of attribute is optimized using a Classification and Regression Tree (CART) algorithm or analogous algorithm, such that the attribute and split used at each node minimize impurity.

and the process is repeated with a new attribute. The process continues until a maximum tree depth or similar termination condition specified by the user is reached, at which point the nodes that have not been split are referred to as *leaf nodes*. A popular choice for the impurity function, and the one used here, is the *gini* impurity, defined as $G_i = 1 - \sum_{k=1}^n p_{i,k}^2$, where i labels a node, and $p_{i,k}$ is the fraction of instances of class k classified into node i . An example of a simple decision tree to classify 150 flower instances into *setosa*, *versicolor* and *virginica* classes based on petal width and petal length attributes is shown in Figure 6.1.




Figure 6.1: Visualization of a simple decision tree to classify flowers, showing (a) node structure, along with gini score and attribute/split choices and (b) a plot of instances in the 2D attribute space, showing the cuts made on the attributes at each decision tree depth. Good separation is achieved among the instances of different classes (represented by marker color) in the leaf nodes (represented by transparent box color) [101].

The BDTs trained in this thesis use a method known as *Adaptive Boosting*, or AdaBoost. A sequence of decision trees are trained iteratively. Once a tree is trained, the instances used in training are assigned weights based on the degree of mis-classification. These weights instruct the successive tree to assign more importance when attempting to classify previously poorly classified instances. In this way, each successive tree "pays closer attention" to the "mistakes" of the previous tree, resulting in an ensemble of decision trees which in succession "learn" from one another. The user is free to tune a number of *hyperparameters* when training/optimizing the BDT ensemble, such as the total number of iterations and the *learning rate*, which determines the magnitude by which the weights are updated between iterations. The final classifier is compiled via *voting* of the ensemble

of trees, where the vote of each tree is assigned a weight based on the tree's accuracy in classifying the weighted training set. The final classifier is known as a BDT.

6.2 DISPLACED ELECTRON AND PHOTON BDTs

For this analysis, two BDTs are trained to optimally distinguish displaced e/γ objects from LLP decays against prompt, pointing e/γ background. The variables available to distinguish displaced from prompt e/γ objects vary between electrons and photons, due to the presence of tracks in the former.

The previous ATLAS pair-produced slepton search [11] targeted events with two displaced leptons in the final state. For example, for pair-produced selectron models the search required two electrons with $|d_0| > 3$ mm in each event. This analysis includes the same approach, but also extends the search with BDTs that target events where ATLAS does not reconstruct both high- d_0 electrons. If an LLP decay vertex is far enough into the tracker that standard and LRT tracking algorithms do not have sufficient hits available to reconstruct an electron track, or if the decay occurs between the tracker and LAr calorimeter face, an electron resulting from the decay can be reconstructed as a photon. If an LLP decay occurs beyond the LAr calorimeter face, an electron resulting from the decay will not be reconstructed at all. If one LLP decay in the event occurs promptly, ATLAS can reconstruct one electron with $|d_0| > 3$ mm and another with $|d_0| < 3$ mm. All of these scenarios result in reconstructed final states with one or more displaced electrons or photons that were not studied in the previous search.

The first BDT, referred to in this text as the $eBDT$, functions as a object-level "tagger" which can be thought of as an extension to standard electron identification, tuned specifically to identify displaced and delayed electrons. The electron BDT provides new sensitivity to events where exactly one electron and one or more photons are reconstructed by ATLAS. This BDT also allows the study of events with two or more reconstructed elec-

trons where one of the electrons fails to meet the requirements of other search regions (see Section 8.2) in the analysis. This includes, for example, events where one electron has $|d_0| < 3$ mm.

The second BDT, referred to in this text as the γ BDT, is an event-level algorithm that provides additional sensitivity for the case where both decays occur at large enough radii that the number of hits in the tracker is insufficient to reconstruct an electron, but both decays still occur before the face of the LAr calorimeter. In these cases, two photons can be reconstructed by ATLAS from the two signal electrons in the pair-produced selectron signal model. The γ BDT also targets events where one signal electron is reconstructed as a photon and the other is reconstructed as an electron which fails to meet the requirements of other search regions. In this case the reconstructed electron is treated as a photon by the γ BDT.

Both BDTs are trained with simulated signal samples in which the arrival time to the calorimeter t of the e/γ objects used in the BDT is required to be greater than zero, and data requiring objects with $t < 0$. This choice of training set is motivated in the analysis strategy discussion of Chapter 8 (see, for instance, Section 8.3.4).

To make use of the entire sample of the MC simulation and the CR data in training each BDT, the signal sample and data used for training are split randomly into three “folds” and training is performed on each group individually. It is then required that the score assigned to a particular event during evaluation is determined using a fold other than the one assigned to that event in training.

6.3 ELECTRON BDT INPUT VARIABLES

The eBDT is trained over the track variables $\Delta p_T/p_T$ (which rejects fakes based on agreement between track and calo-measured p_T), the track fit quality χ^2 , the number of Pixel hits in the track, the number of Pixel or SCT layers which are missing a track

hit beyond the innermost ID layer containing a hit (referred to as the number of missing layers), and the transverse and longitudinal displacements $|d_0|$ and $|z_0|$. The natural logarithm transformation is applied to some variables as a *regularization* to ensure the BDT algorithm does not assign inflated significance to a small number of outlier events. The eBDT is also trained over variables calculated from the electron calorimeter clusters: the LAr cluster time $|t|$, the LAr pointing $|z|$, the difference between LAr pointing and track pointing $|z - z_0|$ (similar to $\Delta p_T/p_T$, this is meant to reject fakes) and two LAr shower shape variables, the fraction of energy in layers 1 and 3 (f_1 and f_3 respectively). It is found that the eBDT is able to distinguish signal based on sensitive inputs without significant sculpting as a function of object (p_T, η) . The training signal and background distributions for all eBDT input variables are shown in Figure 6.2.

6.4 PHOTON BDT INPUTS AND TRAINING

As tracks are unavailable for electrons reconstructed as photons, the object-level variables used by the γ BDT are calculated from the LAr calorimeter measurements $|z|$ and $|t|$. The event-level variable E_T^{miss} is also used. The γ BDT represents a more model-dependent approach in which two e/γ objects (at least one of which is required to be a photon) are required to be reconstructed in the event. It was found that this requirement on event topology drastically reduces backgrounds by rejecting events with only one photon, and improves performance by allowing the algorithm to classify events based on two independent decays, each of which is equally likely to exhibit displaced and delayed kinematics in the signal model. The goal of the γ BDT is primarily to extend the exclusion depth from the eBDT for the slepton pair-production models considered in this analysis. As such, the γ BDT relies on the model-dependent variable E_T^{miss} for increased performance. In the pair-produced slepton model, a high value of E_T^{miss} is expected which grows with increasing slepton mass. A nearly-massless gravitino which does not interact with ATLAS is created




Figure 6.2: eBDT training distributions for all input variables. Signal events (red) are from simulated selectron events with 200-600 GeV mass and 10-30 ns lifetimes. Prompt background events (blue) are from ATLAS dataset in 2022-2023.

in each slepton decay alongside an SM lepton. The slepton typically imparts a significant fraction of its energy to the gravitino, which is measured by ATLAS as E_T^{miss} . In addition, the calorimeter variables listed above, as well as the sum and differences for calorimeter pointing and timing calculated for the two objects, are used as inputs to the γ BDT. As with the eBDT, the natural logarithm transformation is applied to some variables as a regularization. The training signal and background distributions for all γ BDT input variables are shown in Figure 6.3.




Figure 6.3: γ BDT training distributions for all input variables. Signal events (red) are from simulated selectron events with 200-700 GeV mass and 10-30 ns lifetimes. Prompt background events (blue) are from ATLAS dataset in 2022-2023.

6.5 BDT PERFORMANCE

Each input variable is assigned an importance by the BDT, corresponding to the separation achieved using that variable. The importance ranking of the BDT input variables

is derived by counting how often the variables are used to split decision tree nodes, and by weighting each split occurrence by the separation gain-squared it has achieved and by the number of events in the node. The separation gain-squared is defined in the Toolkit for Multivariate Data Analysis (TMVA) user's guide [102] as

$$\langle S^2 \rangle = \frac{1}{2} \int \frac{(\hat{y}_S(y) - \hat{y}_B(y))^2}{\hat{y}_S(y) + \hat{y}_B(y)} dy$$

where $\hat{y}_S$ and $\hat{y}_B$ are the signal and background Probability Density Functions (PDFs) of y , respectively. This variable returns 0 for identical signal shapes, and 1 for shapes with no overlap. Table 6.1 and Table 6.2 show the importance ranking of the input variables for the eBDT and γ BDT, respectively. For the eBDT, $\Delta p_T/p_T$ and $|d_0|$, which when combined can distinguish clusters with displaced tracks whose p_T measurement in the calorimeter and tracker are consistent, are the highest-ranked. For the γ BDT, the event-level variable E_T^{miss} , which is expected to have a true value of zero for background and a large value in the benchmark signal model due to the momentum imparted from each selectron to an unmeasured gravitino, is the highest ranked. In general, the calorimeter inputs $|t|, |z|$ are of greater significance for the γ BDT than the eBDT. The requirement to reconstruct a displaced electron as a photon necessitates that the electron does not leave enough ID hits for ATLAS to reconstruct a track. This requirement introduces a heavy bias towards the tails of the exponential mean lifetime spectra of the selectrons where large t, z of decay products are expected.

A hyperparameter scan was performed to ensure optimal performance of the BDTs and to mitigate *overtraining*, which occurs when an ML algorithm "learns" to classify events based on features of the training set that are not general features of the dataset from which the training events are sampled. Overtraining manifests as weak performance of BDTs when applied to data outside of the training set. A description of the scan procedure and the resulting BDT configuration, such as learning rate, maximum tree depth, and

loss function, can be found in Appendix A. Each BDT achieves a final Area Under the Curve (AUC) for its Receiver Operating Characteristic (ROC) curve greater than 0.98. The AUC of a ROC curve quantifies the ability of a classifier to correctly classify signal and background events. An AUC value near 0.5 indicates random classification and a value close to 1 indicates very strong classification performance. Figure 6.4 shows the BDT score distribution on training signal and background for the eBDT and γ BDT. The score distribution is central to the analysis strategy, as it is used to predict the background in the signal region (Chapter 8) and is the subject of the fit performed to determine whether a BSM signal is present in the ATLAS dataset (Chapter 10).

Table 6.1: eBDT input variable ranking.

eBDT Input	Importance
$\Delta p_T/p_T$	.76
$\ln d_0 $	.41
$\ln z - z_0 $	.18
$\ln z $	.14
χ^2/N_{DoF}	.13
f_1	.13
$N_{\text{Pixel hits}}$	.12
$\ln z_0 $	.11
$\ln t_e $	.07
$N_{\text{Missing Layers}}$	.04
f_3	.01

Table 6.2: γ BDT input variable ranking.

γ BDT Input	Importance
E_T^{miss}	.61
$\ln z_1 $	.40
$\ln z_1 - z_2 $	.37
$\ln z_1 + z_2 $	.34
$\ln t_1 + t_2 $	.33
$\ln z_2 $	.30
$\ln t_2 $	.24
$\ln t_1 - t_2 $	.22
$\ln t_1 $	.20

6.6 LIQUID ARGON TIMING

6.6.1 TIME MEASUREMENT

The LAr calorimeter has excellent intrinsic timing resolution, with the measurement taken from the middle-layer calorimeter cell with the maximum energy within the cluster ($E_{\text{cell}}^{\text{max}}$). In each LAr cell, an Optimal Filtering [77] technique is performed on the four samples of the LAr pulse (See Figure 4.6) spaced by 25 ns around and including the triggering




Figure 6.4: BDT response for training signal (red) and background (blue) for (a) eBDT and (b) γ BDT.

bunch crossing. Given a template pulse shape and autocorrelation function, Optimal Filtering Coefficients (OFCs) labeled a_i and b_i , where the index i runs over the four samples, are calculated. These OFCs are then used along with the samples S_i to calculate the pulse amplitude (A , proportional to pulse energy) and peak time (t) as

$$A = \sum a_i S_i \quad (6.1)$$

$$t = \frac{1}{A} \sum b_i S_i \quad (6.2)$$

During data taking, the measured t from the maximum energy middle-layer cell is used to veto *out-of-time pileup*, reconstructed objects due to particles from proton collisions in bunch crossings neighboring the proton collision under study. The energy in this cell, $E_{\text{cell}}^{\text{max}}$, is typically close to one-third of the total energy deposited in an EM cluster. The online timing resolution is roughly 1 ns, which is more than sufficient to associate prompt e/γ clusters to the correct bunch crossing. An offline timing calibration procedure is able to improve this resolution by a factor of about 4 for high $E_{\text{cell}}^{\text{max}}$ clusters. Irreducible factors affecting the resolution include the $O(20 \text{ ps})$ clock jitter of LAr electronics and the lim-

iting $O(190 \text{ ps})$ intrinsic resolution from the finite spread of the proton beams along the collision plane.

6.6.2 TIMING CALIBRATION PROCEDURE

The online timing measurement is subject to an offline calibration procedure, which uses high-statistic samples of $W \rightarrow e\nu$ events to derive calibration constants and is tested in $Z \rightarrow ee$ events. For the Run 3 dataset, the calibration is performed in three so-called Intervals of Validity (IOVs) across the ATLAS dataset taken in 2022 and 2023, corresponding to changes in LAr online OFCs on one or more FEB boards. The run number range, luminosity, and fractional luminosity of the full 2022-2023 dataset for each IOV are shown in Table 6.3.

Year	IOV	N Runs	Lumi. (fb^{-1})	Pct. Total Lumi.
2022	1	428648-430896	3.46	5.90
	2	431037-440613	27.94	47.65
2023	3	451587-456749	27.24	46.45

Table 6.3: Intervals of validity (IOVs) for the 2022-23 LAr timing calibration.

The calibration procedure closely follows the procedure used in Run 2; details of the method used to derive and validate constants for each pass can be found in Ref. [103]. Within each IOV, a time-of flight (ToF) correction is applied using the distance from the PV to the calorimeter cell in which the EM object was measured, and the speed of light. Subsequent corrections are then calculated per-gain (for high and medium gain) in *passes* to derive corrections per-run number, per-FEB number, per-individual channel number, for the energy in the middle-layer cell used to calculate the time, and for crosstalk effects. At each pass, a constant is calculated which bring the average reconstructed time for electrons to zero as a function of the variable chosen. The constant from the previous pass is applied upon all subsequent passes, with the overall effect of collapsing many approximately Gaussian distributions with different means onto a single narrower distribution

with a mean of approximately $t = 0$ ns. The calibration is applied offline with a centrally maintained software package.² Figure 6.5 shows the result of applying the calibration procedure to a selection of $W \rightarrow e\nu$ decays from the 2022 ATLAS dataset.




Figure 6.5: Results of the timing calibration for IOV2, High gain. a) Pre- and post-calibration timing distribution, with a Gaussian fit to the core of the distribution in the calibration dataset. b) Timing resolution in the validation set of $Z \rightarrow ee$ candidate events, measured as a fit to the function $\sigma_t = \frac{p_0}{E} \oplus p_1$, where the fitted value of p_1 represents the timing resolution for large pulses.

After the calibration is applied, the resulting distribution is described by a Gaussian core with non-Gaussian tails due to nonlinear detector effects, imperfect calibrations, etc., approximately centered about $t = 0$ ns, as expected. In Figure 6.5(a), a Gaussian fit to the core of the distribution returns a resolution of $\sigma = 315$ ps, integrated over bins of $E_{\text{cell}}^{\text{max}}$. As shown in Figure 6.5(b), the application of the calibration constants to the $Z \rightarrow ee$ validation set shows that the resolution varies as a function of $E_{\text{cell}}^{\text{max}}$, with the best resolution achieved of $\mathcal{O}(250$ ps) for e/γ objects with $E_{\text{cell}}^{\text{max}} > 30$ GeV. This precision is sufficient to distinguish e/γ objects resulting LLP decays against prompt background from SM processes, in particular when $\tau(\text{LLP}) \geq 1$ ns. A critical feature of the calibrated timing distribution for prompt SM processes is its symmetry about $t = 0$ ns. These properties make the calibrated

²<https://gitlab.cern.ch/ancsmith/OfflineTimingTool/>

LAr calorimeter time a key input to the BDT analysis strategy for e/γ objects described in Chapter 8.

PART III:

SEARCH

CHAPTER 7:

DATA AND SIMULATED SAMPLES

The result presented is a search for evidence of LLPs decaying to leptons, using the $p - p$ dataset collected by the ATLAS detector between 2015-2023. After it is collected, the ATLAS dataset is subjected to Data Quality (DQ) criterion which reject events with defects. The amount and quality of the data collected are the subject of Section 7.1. A Monte Carlo (MC) simulation of slepton pair-production events in a simulated ATLAS detector is used to quantify the sensitivity of the analysis method to detect evidence of LLP decays in the dataset. The MC samples are discussed in Section 7.2. Both the data and simulated samples are filtered through a set of triggers which enrich the dataset in leptonic and photonic final states, including novel LRT triggers critical for the search strategy employed on events collected in Run 3, as described in Section 7.3.

7.1 ATLAS DATASET

The Run 2 $p - p$ dataset used in this result consists of data collected between 2015-2018 at $\sqrt{s} = 13$ TeV. The Run 3 $p - p$ dataset used consists of data collected during 2022-2023 at a slightly higher center-of-mass energy of $\sqrt{s} = 13.6$ TeV. The size of each of these datasets is measured as *integrated luminosity*, the integral over time of the instantaneous luminosity described in Chapter 3. This is typically written as a unit of inverse area known as an inverse femtobarn ($1 \text{ fb}^{-1} = 10^{-43} \text{ m}^{-2}$), which when multiplied by the cross section σ of

a process of interest gives the total number of events expected. As an example, the cross section of proton collisions, colloquially the rate of "protons to anything", at $\sqrt{s} = 13$ TeV is approximately 73 trillion fb, or 73 trillion events per fb^{-1} delivered.

Figure 7.1 shows the integrated luminosity delivered by the LHC (156 fb^{-1} in Run 2, 70 fb^{-1} in Run 3) and collected by ATLAS (147 fb^{-1} in Run 2, 66 fb^{-1} in Run 3). Events rejected online and offline due to data corruption, as well as the so-called "warm-start" in which the inner tracking detectors ramp up high voltage after Stable Beams are declared, contribute to the 6% inefficiency in ATLAS recording. From the recorded data, the ATLAS Data Quality team certifies lists of runs which are "good for physics", known as Good Runs Lists (GRLs). A GRL can depend on the triggers used in a particular data period, and also excludes events flagged by the DQ team as containing corrupted data, large bursts of correlated noise in a region of the LAr calorimeter, among other defects. For this analysis, 140.0 fb^{-1} and 56.3 fb^{-1} of data in 2022-2023 are collected in the GRL, corresponding to a total dataset size of 196.3 fb^{-1} . It is important to note that a search for displaced leptons has already been published by ATLAS using the Run 2 dataset [11], and that this result aims to add to that result by combining new data in Run 3 and novel analysis methods in Run 3.




Figure 7.1: Integrated luminosity delivered by the LHC and collected by ATLAS in (a) 2015-2018 and (b) 2022-2023 [71].

7.2 SIMULATED SAMPLES

The GMSB pair-produced slepton model introduced in Chapter ?? is used to optimize the analysis for sensitivity to LLP decays to leptons and quantitatively interpret the results. Signal model MC simulation events with up to two additional partons at leading order are generated using MADGRAPH5_aMC@NLOv2.6.1 [104] interfaced to PYTHIA8.230 [105] with the A14 parameter set (tune) [106] along with the NNPDF2.3LO parton distribution function (PDF) set [107]. GEANT [108] is used to simulate the propagation and decay of the slepton without preserving information about its chirality. To model the impact of out-of-time and in-time pile-up, each hard-scattering event is overlaid [109] with simulated minimum bias events generated by EPOS 2.0.1.4 [110, 111] with the EPOS LHC tune and PYTHIA8.307 with the NNPDF2.3 LO set of PDFs and A3 tune [112]. The 13 TeV slepton-pair production cross sections and uncertainties were calculated at next-to-leading-order (NLO) in α_s , with soft-gluon emission effects at next-to-leading-logarithm. At 13.6 TeV, these were calculated at approximate next-to-next-to-leading-order in α_s matched with threshold resummation at the next-to-next-to-leading logarithmic accuracy (aNNLO+NNLL) using an envelope of predictions with different factorization and renormalization scales and PDF sets [113]. The superpartners of the left- and right-handed leptons are assumed to have equal mass.

In the models generated, the signature of each LLP decay manifests as E_T^{miss} from the gravitino and a lepton whose displacement from the primary vertex is correlated to the $\tilde{l}$ mean lifetime. Final states with a range of properties can therefore be scanned by considering range of possible masses $m(\tilde{l})$ and lifetimes $\tau(\tilde{l})$ for the slepton, with larger $m(\tilde{l})$ leading to larger E_T^{miss} values and larger $\tau(\tilde{l})$ leading to larger lepton displacement and delay. The generated signal grid which spans this phase space is summarized in Table 7.1.

Due to lepton flavor universality, the cross sections shown in Table 7.1 are the same for all three slepton generations. In the case of selectron(smuon) models, 100% of decays

	1 ps	10 ps	100 ps	1 ns	10 ns	30 ns	60 ns	$\sigma(\text{pp} \rightarrow \tilde{l}\tilde{l}) @ 13.6 \text{ TeV}$
50 GeV	x	x	x	x	x	x	x	6037 fb
100 GeV	x	x	x	x	x	x	x	407.1 fb
200 GeV	x	x	x	x	x	x	x	33.655 fb
300 GeV	x	x	x	x	x	x		6.968 fb
400 GeV	x	x	x	x	x	x		2.0858 fb
500 GeV	x	x	x	x	x	x		0.7631 fb
600 GeV	x	x	x	x	x	x		0.31692 fb
700 GeV	x	x	x	x	x	x		0.1436 fb
800 GeV		x	x	x	x	x		0.06934 fb
900 GeV		x	x	x	x			0.035141 fb
1000 GeV		x	x	x	x			0.018522 fb

Table 7.1: The slepton mass and lifetime points generated across the grid of simulated signal models. The expected cross section from pp collisions at $\sqrt{s} = 13.6 \text{ TeV}$ is also listed.

are assumed to be to electrons(muons) and a gravitino, while in the stau case 100% are to taus which then decay promptly with 34% probability to an electron or muon. The cross sections fall rapidly with increasing slepton mass. In the full $140 \text{ fb}^{-1} + 56.3 \text{ fb}^{-1}$ of data collected in Run 2 + Run 3, theory would predict approximately 1.2M events for a selectron with a mass of 50 GeV and approximately 4 events for a selectron with a mass of 1000 GeV.

Signal models with large $m(\tilde{l})$ and $\tau(\tilde{l})$ are kinematically more distinct from SM backgrounds, and are therefore generally easier for a sensitive variable such as a BDT score or calibrated calorimeter timing to separate from background. However, large $m(\tilde{l})$ is correlated to lower event yields due to decreasing cross section, and large $\tau(\tilde{l})$ leads to decays at such large radii that decay products escape detection by ATLAS. The signal grid is chosen to probe the range of $m(\tilde{l}), \tau(\tilde{l})$ where a sufficient number of signal-like slepton events exist in the ATLAS dataset (in the signal hypothesis scenario) to allow for a quantitative statement about the likelihood of the dataset containing such events.

7.3 TRIGGERS

Based on studies with the pair-produced stau, smuon, and selectron simulated samples discussed in the previous section, a handful of muon, electron, and photon triggers from the Run 2 and Run 3 ATLAS trigger menu (See Section 4.8) were selected as having high efficiency for displaced leptons. For a pair of muons, the $2\mu 50$ trigger is used. This nomenclature first lists the number of objects (2), the type of object (muon), and the minimum p_T of the objects needed to fire the trigger (50 GeV); this scheme is common to the naming of all triggers used.¹ For a muon and a displaced electron reconstructed at trigger level as a photon, $\gamma 40\mu 40$ is used. For both of these triggers, the muon is identified using MS tracks only (the ID is omitted) which avoids bias against displaced muons. For one or more electrons reconstructed at trigger level as a photon, $1\gamma 140$ and $2\gamma 50$ are used. As the trigger-level object is a photon, again the ID is not used and a bias against lepton displacement is therefore avoided.

As mentioned in Section 5.3.3, new triggers for LRT electrons and muons, dedicated for LLP searches, were developed for Run 3 and were used in 2022-2023 datataking. These triggers provide a key advantage of lower p_T thresholds with respect to the muon and photon triggers used in the previous full Run 2 result, as well as the ability to trigger on electron objects directly. Figure 7.2 shows the expected trigger efficiencies of the new LRT triggers for the displaced stau model used in this analysis, merging several mass points. The new triggers achieve lower p_T thresholds of 20(30) GeV and a d_0 requirement of 2(3) mm for muons(electrons). The LRT electron trigger uses the LLP loose ID WP described in Section 5.4.3 and has no isolation requirement. The triggers required to select a given event for use in the analysis are dependent on the region the event is categorized into; more detail is given in Section 8.1.2 and Table 8.2.

¹Triggers used also require a level-1 seed, an ID WP, and in some cases additional kinematic cuts. These are omitted here for brevity, in all cases the ID WP used in the trigger is at least as tight as the requirement for the offline reconstructed physics object.




Figure 7.2: Electron (a,b) and muon (c,d) trigger efficiencies for LRT triggers (green) and standard single electron and muon triggers (blue) as a function of offline reconstructed lepton production radius (a,c) and p_T (b,d). The logical OR between the LRT and standard triggers is also shown (red circles). The LRT electron(muon) trigger runs in a Region of Interest (RoI) defined by the primary prompt L1 electron(muon) trigger chain, and requirements are placed on electrons(muons) reconstructed in the RoI offline. Results are shown for stau models with masses 100-500 GeV and a lifetime of 1 ns [73].

Another key advantage of the LRT electron trigger for Run 3 is shown in Figure 7.3, which compares a set of electron and photon trigger rates for events where ATLAS reconstructs only one electron candidate passing baseline analysis criteria (see Section 8.1.1).

N_{total} : 862.19; N_{fiduc} :229.72; (26.6%); N_{trig} :163.03; (18.9%)




Figure 7.3: Venn diagram showing the overlaps among triggers fired by a 200 GeV selectron MC sample with a 10 ns mean lifetime, in events where ATLAS reconstructs only one electron. Of the total weighted events generated, 26.6% fall into the fiducial volume for ATLAS to have a chance of reconstructing an electron, and 18.9% of these events fire a trigger on the signal electron. Note that the e60-lhmedium trigger adds no standalone sensitivity for this model when combined with the other triggers shown.

The LRT trigger is listed in the diagram as e30-lhloose-nopix-lrtmedium. For a 200 GeV 1 ns selectron sample, 38.2% of all triggered events fire the LRT trigger and are missed by all other candidate triggers, including the $1\gamma 140$ trigger used in the previous full Run 2 analysis. The near-doubling of signal acceptance from the LRT trigger in the 1-10ns LLP lifetime range brings crucial sensitivity to a novel analysis strategy for events with only one reconstructed electron. In the previously published ATLAS Run 2 result [11], two reconstructed leptons were required for all studied events. This choice drastically reduces backgrounds but limits the available LLP phase space that can be studied, in particular at lifetimes of 1ns and above. The novel BDT-based method allowed by the new phase space

accessible with the LRT electron trigger is the crux of this thesis, and is a primary subject of Chapter 8.

CHAPTER 8:

ANALYSIS STRATEGY

This search targets displaced electrons and muons arising from a BSM physics scenario, using the slepton pair production model introduced in Section 2.3 as a benchmark. The analysis strategy is divided into two orthogonal components. One component, referred to in the text as the *ABCD* search, requires at least two electrons, two muons, or an electron and a muon. The *ABCD* search exploits the Run 3 LRT triggers to build upon techniques developed in a previous search for this signature performed using ATLAS Run 2 data [11]. The other component, referred to as the *EM-BDT* search, also exploits the Run 3 LRT triggers for electrons, and leverages precision tracking and EM calorimeter information to improve the displaced e/γ selection. This component adds a single-electron search region and a multi-photon search region, and uses the precision LAr spatial and timing measurements and the BDTs discussed in Chapter 6. The techniques used in the *EM-BDT* search take inspiration from earlier searches which exploit LAr measurements to search for displaced photons from BSM models [3,4,90,91]. In both components of the analysis, background estimates are determined with entirely data-driven methods. Both the *ABCD* search and the *EM-BDT* search use 56.3 fb^{-1} of Run 3 data collected at $\sqrt{s} = 13.6 \text{ TeV}$ in 2022–2023, while the *ABCD* search additionally considers 140 fb^{-1} of data from the full Run 2 dataset collected at $\sqrt{s} = 13 \text{ TeV}$.

A common event selection (Section 8.1) is applied to the dataset to determine whether an event is categorized into an ABCD region, an EM-BDT region, or neither. The focus of this thesis is the novel analysis methods developed and executed in the EM-BDT portion of the search. As the selections applied to the EM-BDT take the ABCD selections into account, and the final result relies on the combination of ABCD and EM-BDT analysis channels, the ABCD strategy is first discussed in Section 8.2. The remainder of the chapter (Section 8.3) is dedicated to the EM-BDT strategy.

8.1 EVENT SELECTION

The selections which determine the analysis region of each event are based on reconstructed electron, muon, and photon objects, as described in Chapter 5. These objects are first subject to a set of common baseline selections to define *baseline object* collections, described in Section 8.1.1. The baseline objects are used to categorize events into an ABCD search region, an EM-BDT search region, or neither, based on p_T , $|\eta|$, $|d_0|$, and object multiplicity, as described in Section 8.1.2. For ABCD and EM-BDT regions, events are then required to pass a selection of the triggers based on year of data collection and event topology.

8.1.1 BASELINE OBJECT DEFINITIONS

The requirements which define baseline objects are listed in Table 8.1. Note that for photons, which have no tracks, there are no requirements on $|d_0|$ and $|z_0|$. For e/γ objects, the region $1.37 < |\eta| < 1.52$ is rejected, this requirement is referred to as the *crack veto* because it removes the poorly covered fiducial volume between the LAr barrel and endcaps, where very large fake rates of electron and photon objects are expected due to incomplete shower information. The ID Quality requirements are described in more detail in Section 5.6 (electrons) Section 5.4.2 (photons) and Section 5.5 (muons). For muons, a

requirement at $|d_0| > 2$ mm is used to reject prompt muon backgrounds; this requirement is omitted for electrons to allow the EM-BDT to probe the case where one LLP decay is prompt (low $|d_0|$) and the other is delayed/displaced.

Parameter	Electrons	Muons	Photons
p_T [GeV]	> 10	> 21	> 30
$ \eta $	$ \eta < 1.37$ OR $1.52 < \eta < 2.47$	< 2.5	$ \eta < 1.37$ OR $1.52 < \eta < 2.37$
$ d_0 $ [mm]	< 300	$2 < d_0 < 300$	-
$ z_0 $ [mm]	< 500	< 500	-
ID Quality	LLP VeryLooseLLH	LLP Medium	Loose

Table 8.1: Overview of electron, photon, and muon baseline selections.

In addition to the requirements in Table 8.1, which are chosen to reject fake background while maintaining high efficiency for LLPs decaying to leptons, an additional requirement on a small detector volume is added due to the TRT operational status in Run 3. Leptons that fall into the region $(-1.759 < \phi < -1.257, |\eta| < 1.37)$ pass through a part of the TRT that was turned off in 2023 (the so-called *TRT hole*) resulting in no TRT hits on the tracks. As discussed in more detail in Appendix C, the BDT response to leptons in this region of $(|\eta|, \phi)$ is artificially biased towards signal-like classification. Events collected during 2023 with leptons falling into that hole in 2023 are therefore vetoed in all SRs.

8.1.2 EVENT CATEGORIZATION

The analyses requires at least two displaced baseline leptons, and failing that, at least one baseline electron or photon. The categorization of events into ABCD, EM-BDT, or neither is illustrated in the flowchart of Figure 8.1. The Run 2 dataset is used to repeat the previously published ATLAS Run 2 result [11]; events with at least 2 leptons, both with $p_T > 65$ GeV and $|d_0| > 3$ mm, are placed into a category called the *ABCD high- p_T region*. With the Run 3 dataset, an analogous region is created. If an event fails this selection, but has at least 2 leptons both with $p_T > 31(21)$ GeV and $|d_0| > 3(2)$ mm for electrons(muons), the event falls into the ABCD LRT region, a novel region which exploits the lower p_T

thresholds of the Run 3 LRT triggers. If an event fails this selection, but contains at least one barrel electron with $p_T > 10$ GeV, $|d_0| > 3$ mm, or at least two barrel photons with $p_T > 30$ GeV, it falls into the EM-BDT region. Of the EM-BDT events, those with an electron passing the listed selections are used as input to the eBDT defined in Section 6.2, while the others are passed to the γ BDT. Precedence is given to events with an electron passing the listed selections, as a displaced electron is the target signature of this search, while the γ BDT region is designed to supplement the analysis sensitivity when no barrel electrons passing the requirements of upstream regions are reconstructed.




Figure 8.1: Flowchart of the event categorization into the ABCD high- p_T , ABCD LRT, and EM-BDT SRs [6]. The Run 2 and Run 3 ABCD high- p_T SRs are separate.

All events are required to pass at least one of the triggers listed in Section 7.3, with specific trigger requirements depending on the final state objects in the signal region. These requirements are listed in Table 8.2 and in each case reflect the object multiplicity/ p_T requirements of the corresponding event category. For example, the EM-BDT requires $1e30$ (LRT) or $1\gamma140$ in the case of exactly one baseline electron, as this electron could be reconstructed as an LRT electron or as a photon at trigger level, but in the case of 2 or more e/γ objects events passing baseline selection, the $2\gamma50$ trigger is also permitted. In all cases a *trigger match* is required to ensure the offline reconstructed object is the same that fired

Topology	Trigger Thresholds [GeV]
ABCD High- p_T	
if $\geq 1e$, $p_T > 160$ GeV	1 γ 140
else if $\geq 2e$, $p_T > 60$ GeV	2 γ 50
else if $\geq 1\mu$, $p_T > 60$ GeV, $ \eta < 1.07$	1 μ 60 (MS-only)
else if $e\mu$ (Run 3 only)	γ 40 μ 40 (MS-only)
else if $\mu\mu$ (Run 3 only)	2 μ 50 (MS-only)
ABCD LRT	
e , $p_T > 31$ GeV	1 e 30 (LRT)
μ , $p_T > 21$ GeV	1 μ 20 (LRT)
EM-BDT	
$= 1e, 0\gamma$	1 e 30 (LRT) OR 1 γ 140
$\geq 2e$ OR (1 e , 1 γ)	1 e 30 (LRT) OR 2 γ 50 OR 1 γ 140
$\geq 2\gamma$	2 γ 50 OR 1 γ 140

Table 8.2: Summary of trigger thresholds used for the ABCD and EM-BDT analyses, based on the object topology. Trigger objects are denoted as γ for photons, e for electrons, and μ for muons, preceded by the multiplicity and followed by the p_T threshold in GeV. MS-only indicates triggers that only rely on the Muon Spectrometer (MS), and not the inner detector. Trigger names contain the number of objects, the type of the object, and its p_T threshold. For example, 2 γ 50 means a trigger requiring the presence of two photons, each with $p_T > 50$ GeV. The 1 μ 60 MS-only trigger has a restriction of $|\eta| < 1.05$. The LRT electron (muon) trigger has a $|d_0| > 3(2)$ mm requirement. None of the triggers have an isolation requirement. "OR" indicates a logical OR of multiple triggers.

the trigger.¹ This requirement is necessary for the use of trigger Scale Factors (SFs) which provide a correction and associated uncertainty for trigger efficiency differences between data and simulation (see Chapter 9).

After event categorization and trigger matching, a strategy is implemented in each region to estimate the number of events expected in a signal-enriched *Signal Region* (SR) in the case of the signal hypothesis and the case of the background-only hypothesis. The background estimate is derived in each case from a signal-depleted *Control Region* (CR) and a consistency check is performed in a signal-depleted *Validation Region* (VR). The

¹This is excepting the case of offline electrons matched to photon triggers, where trigger-level cluster information necessary for matching is not readily accessible offline.

construction and methodology of the CRs, VRs, and SRs are discussed in the following sections.

8.2 ABCD STRATEGY

As shown in Figure 8.1, the ABCD analysis consists of nine orthogonal regions. There is one SR for each lepton flavor combination ee , $e\mu$, and $\mu\mu$ in each of Run 2 high- p_T , Run 3 high- p_T , and Run 3 LRT regions. All of these regions employ a data-driven background estimation strategy commonly known in experimental particle physics as the *ABCD method*. The method relies on at least two uncorrelated variables to construct four orthogonal regions, typically labeled A, B, C, and D, and allows one to predict the expected number of events in region A given knowledge of the yields in B, C, and D. In the case of this analysis, the uncorrelated variables chosen are attributes of the two highest- p_T leptons in each event. Based on these variables, pass/fail criterion can be used to construct regions B, C, and D, for each lepton flavor combination, as represented in Figure 8.2.




Figure 8.2: Definition of A, B, C, and D regions for background estimate in the case of (a) two electrons (b) electron and muon (c) two muons. The A regions, where both leptons pass quality requirements, define the signal-enriched SRs.

Under the assumption that the variables used to construct the pass/fail criterion are uncorrelated, one expects $N_A/N_C = N_B/N_D$, and one can therefore predict the event yield in SR region A as

$$N_A^{\text{predicted}} = \frac{N_B \times N_C}{N_D} \quad (8.1)$$

For electrons, the quality requirements used to define the pass/fail criteria are ID track fit $\chi_{\text{ID}}^2/\text{n.d.f.} < 2$, number of missing track hits $N_{\text{miss}} \leq 1$, and $\Delta p_T/p_T \geq -0.5$. An electron which meets all of these criteria passes, while an electron that does not meet one or more of these criteria fails. Muons that have the number of precision tracking layer hits $N_{\text{prec}} > 3$, the presence of a high-precision ϕ measurement $N_\phi > 0$, ID track fit $\chi_{\text{ID}}^2/\text{n.d.f.} < 2$, the combined muon fit $\chi_{\text{CB}}^2/\text{n.d.f.} < 3$, and the number of missing track hits $N_{\text{miss}} \leq 1$ pass, while muons that do not meet one or more of these criteria fail. These requirements aim to distinguish real displaced leptons from combinatoric fake tracks, random overlaps of hadronic tracks with EM clusters or MS tracks, photon conversions, and decays of heavy-flavor bottom and charm mesons. The background estimate derived from this method is therefore referred to as the Fake and Heavy Flavor (FHF) estimate. Note that, for each of the nine estimates, region A is populated by events with two high-quality displaced leptons; in each case region A is the signal-enriched SR while B, C, and D collectively form the signal-depleted CR.

To validate the assumption that the quality criteria of the two FHF leptons are uncorrelated with each other and account for other potential systematic biases, two VRs are defined for each of the nine regions. The VR-fake regions are enriched in fake lepton backgrounds by inverting the requirement on $\Delta p_T/p_T$ for electrons and the CB track $\chi_{\text{CB}}^2/\text{n.d.f.}$ requirement for muons. The VR-Heavy-Flavor (VR-HF) regions are enriched in heavy-flavor leptons by inverting the lepton isolation criteria (see Section 5.8). The observed yields in the nine VR-fake and VR-HF regions corresponding to the nine SRs are consistent with the background expectations obtained by performing the ABCD estimate in these VRs.

The ABCD $\mu\mu$ regions require some additional treatment. In $\mu\mu$ regions, it is possible that a cosmic muon traverses ATLAS and passes near to the PV as an event is being

recorded. Such a muon has the potential to be reconstructed by ATLAS as two back-to-back muons originating near the PV, and is therefore a source of background in $\mu\mu$ regions. In $\mu\mu$ SRs, a cosmic ray veto rejects events with back-to-back fully-reconstructed muons, μ_1 and μ_2 , with $|\eta(\mu_1) + \eta(\mu_2)| < 0.02$ and $\pi - |\phi(\mu_1) - \phi(\mu_2)| < 0.01$. As shown in Figure 8.3, this veto eliminates almost all events in these regions, rendering the use of an ABCD estimate after application of the veto impossible.




Figure 8.3: A two-dimensional plane of the variables $|\Sigma\eta| \equiv |\eta(\mu_1) + \eta(\mu_2)|$ and $|\Delta\phi| \equiv |\phi(\mu_1) - \phi(\mu_2)|$ for events in (a) the cosmic muon VR regions A,B,C, and D, and (b) the $\mu\mu$ CR regions B, C, and D.

Instead, an efficiency transfer factor for the veto of $T = (3.7 \pm 2.6) \times 10^{-3}$ is calculated using the superset of events in the six $\mu\mu$ high- p_T and LRT regions B, C, and D. An ABCD estimate using pass/fail criteria defined by (1) a cosmic muon veto relying on MS segments and (2) the same quality criteria as the FHF estimate is then performed in each $\mu\mu$ region. Pass/fail criteria (1) and (2) are applied to the *top muon* in the event, or the muon reconstructed in the top half of ATLAS. The estimate is validated in VRs enriched in cosmic ray muons by instead applying the same ABCD estimate to the bottom muon, and observed VR yields are consistent with VR background predictions. $N_A^{\text{predicted}}$ for each of

the three $\mu\mu$ regions is scaled by T to arrive at a final cosmic muon background estimate, which is then added to the FHF background estimate.

Additionally, in the $\mu\mu$ high- p_T regions only, the observed yields in the ABCD CRs are too small to provide individual predictions for Run 2 and Run 3. Therefore the yields in the ABCD regions for these two datasets are combined and the predicted FHF background is scaled by the fractions of the total luminosity fractions of $\frac{140.0 \text{ fb}^{-1}}{196.3 \text{ fb}^{-1}} = 0.71$ for Run 2 and $\frac{56.3 \text{ fb}^{-1}}{196.3 \text{ fb}^{-1}} = 0.29$ for Run 3.

Performing the described ABCD method on the nine ABCD SRs returns very small predicted SM background of less than 0.2 expected events in each case. Systematic uncertainties associated with the ABCD regions are discussed in Section 9.1 and the results of the search in these regions are discussed in Section 11.3.

8.3 EM-BDT STRATEGY

As shown in Figure 8.1, the EM-BDT analysis is used on Run 3 data only, and makes use of the two BDTs introduced in Chapter 6, with one SR corresponding to each BDT. The background to the EM-BDT analysis is dominated by prompt particles which are mismeasured or poorly reconstructed, requiring a fully data-driven background estimation method. Following the selections in Section 8.1.2, some additional selections (Section 8.3.1) and corrections to the calibrated timing distributions (Section 8.3.3) are performed. CR, VR, and SRs are defined based on lepton kinematics (Section 8.3.2), and a background estimation is derived using data in the CR (Section 8.3.4) and validated in the VR (Section 8.3.5).

8.3.1 EM-BDT SELECTIONS

Selections are made on e/γ objects in the EM-BDT regions, whose input events are by the definitions in Figure 8.1 orthogonal to ABCD regions, to ensure high-quality objects are used as BDT inputs.

The eBDT introduced in Section 6.3 is used to analyze events with at least one electron in the LAr calorimeter barrel ($|\eta| < 1.37$). Events are required to have at least one electron, and the triggers required are based on object multiplicity (in the case of 2 objects, events passing the $2\gamma 50$ trigger are allowed). Triggers are designed to ideally accept only objects above their p_T threshold, and reject objects below this threshold; the finite resolution of the detector leads to a non-ideal s -shape for each trigger efficiency curve (rather than an ideal step-function shape) near the p_T threshold. To account for this non-ideality, when a trigger is fired, the trigger-matched object must have a p_T of a few GeV above the trigger threshold (dependent on the object, trigger, and region). This requirement allows for easy application of scale factors (calculated in bins of $p_T, |\eta|$ using MC objects with p_T greater than the trigger threshold) to account for different trigger efficiencies between simulation and data. The highest- p_T barrel electron in the event is chosen as input to the BDT, as central e/γ objects are better resolved than objects in the endcaps. The chosen electron is required to have $E_{\text{cell}}^{\text{max}} > 5$ GeV to reduce the impact of noise on the calorimeter timing, and an absolute value of its calorimeter time less than 12.5 ns, to reject backgrounds from other bunch crossings. These requirements on electrons input to the eBDT are listed in Table 8.3.

Given the large backgrounds in single-photon event samples as well as the lack of tracking information for photons, the γ BDT is used to analyze only events with at least two reconstructed e/γ objects. At least one photon passing the baseline selections is required, in addition to another baseline photon or a baseline electron which fails the eBDT requirements. At least one EM object must be in the LAr calorimeter barrel ($|\eta| < 1.37$).

Parameter	Selection requirements	
e multiplicity	1	≥ 2
trigger	LRTe30 OR g140	LRTe30 OR g140 OR 2g50
p_T (e,trig) [GeV]	3 GeV above threshold for e trig, 10GeV for g140, 2g50	
η	Choose leading e with $ \eta < 1.37$	
ID Quality	Modified Loose (no $ d_0 $ in LLH, no Pixel cut)	
$ d_0 $ [mm], p_T [GeV]	$ d_0 > 3$, $p_T > 10$	
t [ns]	$t \in [-12.5, 12.5]$	
$E_{\text{cell}}^{\text{max}}$ [GeV]	leading barrel e $E_{\text{cell}}^{\text{max}} > 5$	

Table 8.3: Summary of selections for electron BDT analysis regions.

Analogous to the electron objects selected for the eBDT, EM objects are required to have $E_{\text{cell}}^{\text{max}} > 5$ GeV and absolute value of calorimeter time less than 12.5 ns. A requirement of $|z| < 2000$ mm for both EM objects eliminates beam-induced backgrounds which can fake the signature of a photon. These requirements on e/ γ objects input to the γ BDT are listed in Table 8.4.

Parameter	Selection requirements	
e, γ multiplicity	0e, $\geq 2 \gamma$	$\geq 1 \gamma$, 1e
trigger	g140 OR 2g50	
p_T (γ ,trig) [GeV]	5 GeV above trigger threshold	
η	At least one γ with $ \eta < 1.37$	
ID Quality	Medium for γ , modVeryLoose for e	
t [ns]	$t \in [-12.5, 12.5]$	
$E_{\text{cell}}^{\text{max}}$ [GeV]	$E_{\text{cell}}^{\text{max}} > 5$	
$ z $ (pointing)	$ z < 2000$ mm	

Table 8.4: Summary of selections for photon BDT analysis regions.

8.3.2 EM-BDT REGION DEFINITIONS

Following the event selections of the previous section, the BDT regions are each separated into a signal-enriched SR and a signal-depleted CR and VR. The primary observable used to define these regions is the calibrated LAr timing measurement described in Section 6.6.1. After the calibration procedure is performed, the distribution of the measured LAr timing t for prompt SM background is a symmetric distribution about $t = 0$ with a

gaussian core and a resolution of $O(250 \text{ ps})$. For selectron signal samples with $\tau(\tilde{e}) \geq 1 \text{ ns}$ and masses such that $\bar{\beta} = \frac{\bar{v}}{c} < 1$ (masses for which, on average, the selectron moves significantly slower than the speed of light), the time-of-flight distribution of the signal electron measured at the calorimeter features few events with $t < 0$, and a substantial tail with $t > 0$. The signal contamination for events with $t < 0$ is therefore small. The LAr calibrated timing distribution shapes for data and a few representative signal samples are shown in Figure 8.4.

As data events selected for use in the eBDT with electron time-of-flight measurement $t < 0$ are expected to be signal-depleted, this subset of events is designated as the CR for use in analysis optimization and background estimate development. Events selected for the eBDT region with $t > 0$ are classified as SR events. In events selected for the γ BDT analysis, both e/γ objects are required to have $t < 0$ ($t > 0$) to be classified into the CR(SR).

From events selected for the CR and SR, those consistent with an SM W -boson or Z -boson decay are reserved instead for use in the signal-depleted VR, a region used to validate the EM-BDT search strategy prior to unblinding. To select W -boson enriched events, the transverse mass m_T is calculated from E_T^{miss} and p_T of the electron in events containing one electron, and required to satisfy $60 < m_T/\text{GeV} < 100$. To select Z -boson enriched events, the di-electron invariant mass m_{ee} is calculated from the p_T of the two leading electrons in events with more than one electron, and required to satisfy $80 < m_{ee}/\text{GeV} < 100$. To ensure that the three regions are orthogonal, any events that fall within the VR selection are vetoed from the CR and SR selections. For all slepton masses studied, the fraction of events falling into the VR is a fraction of a percentage of the data events in the VR, with the VR signal contamination maximal for lower slepton masses. Figure 8.5 provides a cartoon indicating the selections which define the SR, VR, and CR used for the analysis.

The SR was kept *blinded* to avoid bias while the analysis strategy was developed; the yields and distribution shapes of sensitive variables in the SR were kept as unknown quan-




Figure 8.4: Representative data and signal model distribution for the (top) LAr timing distribution and (bottom) $|d_0|$ using the EM-BDT search regions, for Run 3 data and selected signal samples. The data events with negative times (gray squares) represent the signal-depleted control region (CR) used to determine the background estimate, while the data events with positive times (black circles) represent the unblinded signal region (SR) events on which the final search is performed.)

tities until the analysis strategy was finalized. The results presented in Chapter 11 are based on a statistical fit procedure performed on *unblinded* SR distributions.




Figure 8.5: Representation of the EM-BDT SR, CR, and VR definitions. The VR definition of $m_T \in [60, 100]$ GeV (1 e/γ object) or $m_{ll} \in [80, 100]$ GeV (2 e/γ objects) is chosen to select a set of events highly enriched in SM W or Z bosons respectively. The CR(SR) is defined by requiring a negative(positive) calibrated LAr calorimeter time of flight measurement (for both objects in the case of 2 e/γ objects).

8.3.3 EM-BDT TIME SHIFTING

A procedure is preformed on data distributions in each EM-BDT region to correct for small residuals to the means of e/γ calibrated timing distributions following the offline timing calibration procedure. Due to finite detector resolution, the timing distribution in each analysis region is centered about a random nonzero value on the order of picoseconds. This effect is primarily driven by low-energy EM clusters, presumably reflecting differences in the EM shower development of electrons and photons at low energies. The correction for this random shift improves the symmetry about $t = 0$ for the subset of the dataset selected for each analysis region. Therefore, a two-step correction procedure is applied, first in bins of $E_{\text{cell}}^{\text{max}}$ and then per analysis region.

First, in bins of $E_{\text{cell}}^{\text{max}}$ an iterative Gaussian fit (in the range $\pm 1.5\sigma$) is performed on e/γ timing distributions to calculate a mean correction.² The correction is subtracted from time measured for all objects in each $E_{\text{cell}}^{\text{max}}$ bin to shift the mean of the timing distribution to $t = 0$. This step is demonstrated in Figure 8.6 for the timing distribution of the electron in each event used as eBDT input. Second, the resulting corrected distributions are separated by CR/SR/VR. A second iterative Gaussian fit (in the range $\pm 1.5\sigma$) is performed to derive an additional shift for each analysis region. This step is demonstrated in Figure 8.7.




Figure 8.6: Demonstration of $E_{\text{cell}}^{\text{max}}$ dependent step in correction procedure for the timing distribution. (left) The timing distribution for the leading barrel electron in all analysis regions is binned in $E_{\text{cell}}^{\text{max}}$ and a correction is calculated as the mean returned by an iterative Gaussian fit to the core of each distribution. (right) The correction is subtracted from each distribution, aligning each time distribution about ($t_e = 0$), as validated by performing a second iterative Gaussian fit, which returns 0 in each bin.

The overall shifted time t' for an event in a particular ($E_{\text{cell}}^{\text{max}}$, region) bin typically has a magnitude of $O(10-20)$ ps. Figure 8.8 shows the timing distribution before and after this shifting procedure is applied in the VR, demonstrating an overall effect of a shift on the order of a few picoseconds.

²The fit range is allowed to grow until the returned σ upon successive iterations does not change by more than a user-specified tolerance.




Figure 8.7: Demonstration of analysis-region dependent step in correction procedure for the timing distribution, for electrons in the VR. After calculating and applying the $E_{\text{cell}}^{\text{max}}$ dependent shift, an additional correction is calculated (for the case shown, +2.7ps) via an iterative Gaussian fit to all VR electrons. The same procedure is performed for electrons in the CR_e/SR_e .




Figure 8.8: Calibrated timing distribution before and after the mean shift procedure for the VR, showing the overall effect of an $\mathcal{O}(\text{ps})$ shift to the timing distribution.

8.3.4 EM-BDT BACKGROUND ESTIMATION

In contrast to the ABCD regions, which rely on an unbinned background yield prediction in each SR, for each EM-BDT region a fit to the shape of the BDT response for SR data is performed to quantify the likelihood that the observed shape is consistent with an LLP signal hypothesis. The background in the EM-BDT analysis is dominated by prompt particles that are mismeasured or poorly reconstructed. These effects appear as non-gaussian tails in distributions of interest for this search, and cannot be simulated to a required level of precision. For this reason, the background estimation in the EM-BDT regions is fully data-driven. The aim is therefore to use the events in the EM-BDT CR to produce an accurate estimate of the shape of the BDT response for background events in the SR for both the eBDT and γ BDT regions, prior to unblinding.

The BDT background estimation strategy critically relies on the assumption that the calibrated LAr time distribution for SM background events is approximately symmetric about $t = 0$. The claim is physically motivated as follows: the background events result from prompt SM decays with a true time-of-flight of $t = 0$, and the nonzero extent of the measured timing distribution is due to finite detector resolution.³ This symmetry is depicted visually in Figure 8.4. Given that the finite extent of the time distribution is a resolution effect, the additional inputs to each BDT will be uncorrelated to the sign of the measured time of prompt background objects, and the resulting BDT response for background is therefore expected to be uncorrelated to the sign of the calorimeter time-of-flight. The independence of the sign of measured time of flight from other BDT inputs for background allows for the eBDT and γ BDT described in Chapter 6 to be trained using the unsigned calorimeter time of flight as an input, with $t < 0$ data events labeled as background and $t > 0$ simulated selectron pair-production events labeled as signal. As the timing distribution for SM background events in the SR is expected to mirror that of CR

³The time-of-flight referred to here is following the correction for the finite speed of light is applied during the calibration procedure

background events, each trained BDT is expected to have an identical response to SM background events and CR background events. Therefore, the shape of the SM background in the signal region can be estimated by using the shape of the BDT response to CR background as a template. The EM-BDT background estimation strategy implicitly assumes that the variables used as input to each BDT (see Chapter 6) are sufficiently well-modeled in simulation that the BDT "learns" to distinguish between signal and background based on physical differences in signal and background events, rather than based on some systematic mis-modeling of physical quantities present in simulated events but not in the dataset. This assumption, as well as the claim that the BDT response distribution for background events is symmetric upon reflection about $t = 0$, are validated using VR events in the next section.

8.3.5 EM-BDT VALIDATION

The VR described in the previous section is largely populated by events containing prompt electrons from SM W or Z boson decays. To investigate the assumption that the measured properties of a prompt SM electron input to the BDT are uncorrelated to the sign of the measured time-of-flight, the VR data is divided into two subsets based on the sign of the measured calorimeter time-of-flight, and the shapes of key variables are compared in Figure 8.9. Reasonable shape agreement is observed, giving confidence that a BDT trained on negative-timed background events only will have a similar response distribution for positively or negatively timed background events.

The strategy is further validated by applying the eBDT and γ BDT trained on negative-timed CR events directly to both positive-timed and negative-timed events in the VR. The result of this comparison is shown in Figure 8.10, where good agreement is observed for more background-like BDT scores in both the eBDT and γ BDT, and disagreement of 33%(84%) observed for the most signal-like eBDT (γ BDT) scores. This disagreement is due both to the finite statistics available (in particular in the γ BDT case, where very few




Figure 8.9: Distributions of p_T , $|d_0|$, $|z|$, and $\Delta p_T/p_T$ for VR events with positive (red) and negative (blue) calorimeter time-of-flight measurements, showing reasonable shape agreement.

events with multiple electrons passing VR selections are available), as well as detector effects in the tails of the timing distributions that remain after calibration which are potentially correlated with other BDT inputs. The observed differences are incorporated into the fit as non-closure systematic uncertainties (see Chapter 9) affecting the shape of the background template. The statistical uncertainty on the low event yields in the most signal-like BDT bins is the dominant source of uncertainty on the results of this thesis.

A comparison between $Z \rightarrow ee$ data events falling into the VR and a simulated sample of $Z \rightarrow ee$ events falling into the VR is performed to validate that the BDT is not "learning" to distinguish differences between simulation and measurement of input distributions when classifying signal-like and background-like events. Figure 8.11 shows the results of this comparison, for a selection of sensitive BDT inputs and for the BDT response distri-




Figure 8.10: eBDT (γ BDT) response comparison between events with negative (light green squares) and positive (dark green circles) calorimeter time-of-flight, in the W/Z enriched Validation Region. A non-closure systematic uncertainty is assigned as the shape difference observed between positive and negative data events.

bution. The comparison is statistically limited, as many simulated and measured $Z \rightarrow ee$ fail to pass the large- $|d_0|$ requirement applied across analysis regions. In fact, this requirement biases data towards a regime where it is poorly measured, and the accuracy of simulation is therefore expected to degrade. Despite these caveats, reasonable agreement is observed between the simulated SM events and the SM background, and the BDT input distributions for the representative signal samples overlaid show clearly distinct shapes with respect to the distributions for simulated $Z \rightarrow ee$ events and $Z \rightarrow ee$ -enriched data.

The ABCD and EM-BDT strategies presented in this chapter are subject to systematic uncertainties affecting the predicted yields and shapes of signal and background distributions in all SRs, as described in Chapter 9. These uncertainties are taken into account when performing the combined fit described in Chapter 10 to quantify the likelihood that the collected data contains an LLP signal.




Figure 8.11: Comparison of $Z \rightarrow ee$ MC and $Z \rightarrow ee$ data for the eBDT score distribution and a selection of the eBDT input variables. Selectron signal samples with SR selections applied are overlaid to demonstrate the difference in shape between $Z \rightarrow ee$ events and signal events. Reasonable agreement is observed between $Z \rightarrow ee$ MC and $Z \rightarrow ee$ data for individual inputs and in the eBDT score.

CHAPTER 9:

SYSTEMATIC UNCERTAINTIES

In the ABCD and EM-BDT analysis strategies presented in the previous chapter, the quantities of interest are the number of predicted background events and signal events (for a particular signal hypothesis) in the region being studied. The ABCD analysis predicts the number of signal and background events expected in region A for the nine SRs constructed, and the EM-BDT analysis estimates the number of signal and background events across two correlated groups of five bins for each BDT. Each of these estimates comes with a *statistical uncertainty* obeying Poisson statistics, due to inherent fluctuations in measurements made in counting experiments. Each estimate also has associated *systematic uncertainties* due to limitations and potential biases in the experimental method. Each source of systematic uncertainty is quantified and included in the fit procedure of Chapter 10.

Systematic uncertainties are incorporated into the fit as either *normalization* or "flat" uncertainties, which affect the overall yields of signal or background predictions, or *shape* uncertainties which affect the shape of the predicted signal or background distributions in each SR. In the context of this result, shape uncertainties only apply to EM-BDT regions, as ABCD regions all involve predictions in only a single bin. Uncertainties on ABCD and EM-BDT background estimates are first presented in Sections 9.1– 9.2, followed by a description of uncertainties on signal estimates in Section 9.3.

9.1 ABCD BACKGROUND SYSTEMATIC UNCERTAINTIES

To assign a systematic uncertainty to the FHF background estimate in the ABCD SRs, a *non-closure* uncertainty is calculated from each VR-HF and VR-fake by comparing $N_{A(VR)}^{\text{observed}}$ to $N_{A(VR)}^{\text{predicted}}$. The magnitude of each observed non-closure is shown in Table 9.1. For each region, the maximum of two quantities is taken as the systematic uncertainty: the VR non-closure, or the statistical uncertainty in the expected background. The values are extracted from VR-fake and VR-HF separately and then added in quadrature. The magnitude of the uncertainties range from 20 – 240%.

		VR			VR-HF		
		$N_A^{\text{predicted}}$	N_A^{observed}	non-closure	$N_A^{\text{predicted}}$	N_A^{observed}	non-closure
ee	high p_T	108.6	89	18.07%	61.2	69	12.8%
	LRT	13.5	19	40.74%	13.8	19	38.2%
$e\mu$	high p_T	1	0	100%	2	2	0%
	LRT	5.14	2	61.82%	5.89	1	71.6%
$\mu\mu$	high p_T	-	-	-	-	-	-
	LRT	0.30	0	100%	1.17	2	71.4%

Table 9.1: The non-closure in VR and VR-HF regions for ee , $e\mu$ and $\mu\mu$ channels. The $\mu\mu$ high- p_T VR region is almost empty and thus the non-closure calculation is omitted. In the table, $N_A^{\text{predicted}}$ means the estimation of region A and N_A^{observed} means the actual observation in region A. The non-closure is the percentage ratio between the deviation and the estimation.

For the $\mu\mu$ channel, an additional systematic uncertainty is assessed in two ways on the cosmic background estimation. As described in Section 8.2, the cosmic background estimation strategy uses the muon in the top half of the ATLAS detector to define the ABCD regions. A non-closure uncertainty is calculated from comparing observed and predicted yields in region A using the bottom muon, amounting to a 41% uncertainty for the LRT region and a 2% systematic in the high- p_T region. In the $\mu\mu$ channel, regions A,B,C, and D are defined using the variables N_{prec} , N_ϕ , and χ_{CB}^2 related to muon quality. A non-closure is calculated three times by inverting two out of three of these criteria, and defining A,B,C, and D regions by inverting the third. The largest difference in these three

non-closures is taken as an additional systematic uncertainty, amounting to 31.7% in the LRT region and 38.4% in the high- p_T region.

9.2 EM-BDT BACKGROUND SYSTEMATIC UNCERTAINTIES

There is an expected uncertainty on the final shape of the background estimation in the SR due to the extrapolation of each BDT from the CR to the positive-timed SR. As described in section 8.3.5, the eBDT and γ BDT responses in the VR are approximately symmetric about $t = 0$, with some disagreement observed at large BDT score values. An uncertainty is assessed using the non-closure of each BDT score shape under time reflection in the VR. The aim is to account for differences that remain across positive and negative times following the time-shifting procedure, caused by detector effects in the tails of the timing distributions after calibration and correlations with other BDT inputs. A systematic uncertainty is applied to the CR template by symmetrizing the differences observed in VR BDT score shapes shown in Figure 8.9. The shape and magnitude of the systematic uncertainty, which is visualized in Figure 9.1, is maximally a 33%(84%) uncertainty in the most signal-like eBDT(γ BDT) score bin. Uncertainty of the EM-BDT background template is nonetheless dominated by statistical uncertainties in the most sensitive bin of the two BDTs, each with an estimate of one background event.

9.3 SIGNAL SYSTEMATIC UNCERTAINTIES

Contributions to systematics uncertainties on signal yields come from parameters used in signal generation, known as theory uncertainties, and experimental uncertainties due to the intrinsic resolution of measurements with ATLAS and potential mis-modeling of the detector response in simulation. A description of the dominant signal uncertainties, with a focus on those affecting the EM-BDT regions, is given in the following sections.




Figure 9.1: CR background template for (a) eBDT (b) γ BDT score, with the shape uncertainty envelope calculated from VR non-closure applied. A 300 GeV 1 ns selectron sample is overlaid.

9.3.1 LUMINOSITY UNCERTAINTY

The total number of events in the dataset is calculated using the LUCID-2 detector [114] for the primary luminosity measurements, complemented by measurements using the inner detector and calorimeters. A flat uncertainty on this measurement is applied to the signal normalization. The uncertainty in the Run 2 integrated luminosity is 0.83% [115], for 2022 it is 2.2% [116], and for 2023 it is 2.0% [117], calculated using the methodology in Ref. [115].

9.3.2 THEORY UNCERTAINTIES

The total signal yields are potentially sensitive to the choice of theory parameters used for the signal generation described in Section 7.2. To quantify this dependence, the signal is re-generated under variation of each of the following parameters: the strong renormalization factor μ_R , the strong factorization scale μ_F , and the choice of Parton Distribution Function (PDF) to model the relative momenta of the proton constituents. The variations have little effect on the shape of the BDT score distribution, and changes in signal yields

under changing the theory parameters are therefore taken as a source of normalization uncertainty. The parameters μ_F and μ_R are each doubled and halved in a correlated and uncorrelated way, and the maximum signal yield difference of 12% from these variations is the dominant source of theory uncertainty.

9.3.3 MISSING TRANSVERSE ENERGY SYSTEMATIC UNCERTAINTY

Uncertainties in reconstructed objects are propagated to the reconstruction of E_T^{miss} . Additional uncertainties in E_T^{miss} come from varying the scale and resolution of the soft term (defined in Section 5.10) in the plane parallel and perpendicular to the hard term [100]. As E_T^{miss} is used in selecting events with enhanced SM W production to define the EM-BDT VR, a systematic uncertainty is assessed on the impact of E_T^{miss} -reconstruction uncertainties in the fitted EM-BDT score distributions. The migration of events from EM-BDT SRs to the EM-BDT VR was studied after redefining the VR using variations in E_T^{miss} . Based on this, a 2% uncertainty in the signal yields in the EM-BDT SRs was assigned. For the γ BDT, which uses E_T^{miss} as an input variable, an additional uncertainty is evaluated by propagating the E_T^{miss} systematic uncertainties into the γ BDT score calculation and re-evaluating the γ BDT shape; these variations are treated as a shape uncertainty and affect the yields in the γ BDT by only a fraction of a percent.

9.3.4 TIME SMEARING SYSTEMATIC UNCERTAINTY

A time-smearing procedure is applied to signal samples to model the resolution of the calibrated LAr timing in data. In this procedure, the smearing consists of sampling two gaussians. The first is a Gaussian centered at zero with a standard deviation of $\sigma = 180$ ps correlated among objects on a per-event level. This Gaussian, when added in quadrature with the intrinsic $\sigma = 55$ ps intrinsic correlated resolution observed in simulated samples, models the 190 ps resolution effect of the beam-spread in data. The second Gaussian is applied per-object, and models electronic and pileup noise. The standard deviation of this

Gaussian is computed by comparing of $Z \rightarrow ll\gamma$ simulated samples and data in bins of η after applying the correlated smearing to simulated samples, and typically has a value in the range of $\sigma \sim 100 - 150$ ps. To derive a systematic uncertainty on this procedure, the uncorrelated component in each bin is varied up and down, typically on the order of 15 – 30%, and resampled. The effect of varying the correlated component is negligible by comparison. This systematic uncertainty is applied in the analysis by supplying each BDT with the resulting up-and-down timing variations, rather than the nominal time returned by the smearing procedure. The BDT response is found to be insensitive to variations in this systematic uncertainty, as shown in Figure 9.2 for the eBDT across a mixture of selectron signal masses and lifetimes.




Figure 9.2: eBDT response to nominal versus up-down variations in time smearing systematic uncertainty. The ratio panel shows that the BDT has very little sensitivity to this source of uncertainty.

9.3.5 CUSTOM SCALE FACTORS AND DISPLACEMENT UNCERTAINTIES

To avoid bias against displaced objects the analysis makes use of specialized LRT track reconstruction, LRT track-based triggers, and custom lepton identification. As the objects are non-standard, the ATLAS collaboration does not provide recommendations for quan-

tifying simulation efficiency with respect to data and associated uncertainties. Custom systematic uncertainties are therefore assessed for the use of LRT objects and non-standard object ID.

The uncertainty in the modeling of LRT tracking in offline reconstruction was determined from comparisons of data and simulation using K_S^0 vertices [118]. In bins of p_T and radius of the first hit on the track, tracks are removed to enforce that the simulation efficiency matches the efficiency observed in data. The resulting decrease in simulation efficiency ranges 0.3–16.8%, depending on p_T and decay radius.

Corrections to the modeling of identification and isolation requirements for leptons were computed via comparisons between simulation and data distributions for prompt $Z \rightarrow \ell\ell$ events, in bins of p_T and η . This was done using the same methods typically used by the ATLAS collaboration for standard object reconstruction [89, 119]. Uncertainties in the corrections for electrons are generally less than 11% for $p_T > 30$ GeV; they rise to 70% for electrons with $p_T < 20$ GeV. For muons, the uncertainties in the corrections are generally less than 3%.

The change in baseline selection efficiency for signal electrons with increasing displacement is assessed in simulation, with a systematic uncertainty calculated as the relative decrease in efficiency in increasing bins of $|d_0|$, compared to efficiency for electrons with $|d_0| < 5$ mm. In order to separate the baseline selection efficiency from the tracking efficiency (considered elsewhere), the efficiencies are calculated for tracks, baseline, and individual electron ID working points; first with respect to the true simulation efficiency (that is, with respect to all signal leptons with generated $p_T > 10$ GeV which decay within a radius of 300 mm), and then with respect to truth-matched tracks. The latter efficiencies are calculated by dividing the baseline and ID efficiencies by the tracking efficiency in each bin. The resulting displacement systematic uncertainty on each working point is computed as one minus the calculated efficiency ratio. These uncertainties, applied in

bins of $|d_0|$, are presented in Figure 9.3, and are maximally 14% for the modified VeryLoose electron ID and 16% for the modified Loose electron ID.




Figure 9.3: Electron displacement uncertainty for modified VeryLoose (left) and modified Loose (right) electron ID working points used in the analysis.

9.3.6 LIFETIME REWEIGHTING SYSTEMATIC UNCERTAINTY

To conserve computing resources, a discrete grid of signal model points are generated in a region of phase space where the analysis method is projected to be sensitive. The analysis method can be evaluated beyond this grid by reweighting the lifetime spectrum simulated for the parent particle to a distribution from a different lifetime, and using this to calculate a per-event weight. This method is particularly useful for the case of very long lifetimes, where many decays occur outside beyond the ATLAS tracker and/or calorimeter. For such models, the signal acceptance is low enough that the computing cost of generating signal events is quite high; generation of many signal events is then needed to yield sufficient statistics for analysis. To avoid this, one can reweight a generated model point to a longer lifetime and the exclusion power of the analysis as usual, with a systematic uncertainty assigned to reflect the accuracy of the reweighting procedure.

In the case of slepton pair production, an event consists of two independent decays. Taking selectrons as an example, the probability of a decay at time $t_{\tilde{e}}$ is given by the mean lifetime equation

$$P(t_{\tilde{e}}) = \tau^{-1} e^{-t_{\tilde{e}}/\tau}$$

To reweight a particular event with leading and subleading slepton decays at $(t_{\tilde{e}1}, t_{\tilde{e}2})$, a weight is calculated to rescale the joint probability distribution from mean lifetime τ_1 to τ_2 as

$$\begin{aligned} wt(\tau_1 \rightarrow \tau_2) &= \frac{(\tau_1^{-1})e^{-t_{\tilde{e}1}/\tau_1} \times (\tau_1^{-1})e^{-t_{\tilde{e}2}/\tau_1}}{(\tau_2^{-1})e^{-t_{\tilde{e}1}/\tau_2} \times (\tau_2^{-1})e^{-t_{\tilde{e}2}/\tau_2}} \\ &= \left(\frac{\tau_2}{\tau_1}\right)^2 \exp \left[(-t_{\tilde{e}1} - t_{\tilde{e}2}) \times \left(\frac{1}{\tau_1} - \frac{1}{\tau_2} \right) \right] \end{aligned}$$

Note that when $\tau_2 = \tau_1$ the weights returned are identically one. This reweighting procedure is always possible for the ABCD analysis regions, where two leptons are required in each SR. Here, the position of the simulated decay is stored in the simulation's *truth record* and can be accessed to determine $t_{\tilde{e}}$. However, the EM-BDT analysis is sensitive to events where exactly one selectron is sufficiently displaced such that the decay occurs beyond the extent of the ATLAS detector. In these events, the decay position is not recorded in the truth record. Therefore, a one-dimensional lifetime reweighting is implemented using the selectron which is truth-matched to the electron used in the BDT, integrating over the decay spectrum of the unused selectron. The one-dimensional reweighting equation becomes

$$\begin{aligned} wt(\tau_1 \rightarrow \tau_2) &= \frac{(\tau_1^{-1})e^{-t_{\tilde{e}}/\tau_1}}{(\tau_2^{-1})e^{-t_{\tilde{e}}/\tau_2}} \\ &= \left(\frac{\tau_2}{\tau_1}\right) \exp \left[(-t_{\tilde{e}}) \times \left(\frac{1}{\tau_1} - \frac{1}{\tau_2} \right) \right] \end{aligned}$$

Given the finite simulation statistics and the finite extent of the ATLAS detector, both reweighting procedures above are only valid for a finite range of source-lifetime targets. As an example, the shape of a distribution arising from a 1 ns model of any mass will suffer from lack of statistics in the tails when reweighting to a lifetime of 60 ns. To assess the validity of the reweighting from a given source lifetime to a given target lifetime, the ratio of the normalization between reweighted samples and generated samples at a given target lifetime are compared (see Appendix B for details). The 10 ns samples at each mass are used for reweighting, as these model points have reasonable detector acceptance and reasonably long lifetime, and a 20% normalization uncertainty is assigned for reweighting to lifetimes less than 100 ns, with a 50% uncertainty assigned for reweighting to lifetimes greater than 100 ns. This uncertainty is chosen to give a conservative estimate on the exclusion for very high lifetimes. When lifetime reweighting is applied, it is one of the dominant systematic uncertainties affecting signal yields.

9.4 SUMMARY

The sensitivity of the combined EM-BDT and ABCD analyses is dominated by statistical uncertainties in the most signal-like SR bins. Nonetheless, the systematic uncertainties described in this chapter are included in the background and signal templates in the fit performed in Chapter 10. A ranking of the systematic uncertainties in the eBDT (γ BDT) fit are shown in Table 9.2 (Table 9.3), for the background template and for a representative signal sample. For the eBDT, a 10 ns sample that has been reweighted to a lifetime of 50 ns is used, to show the relative impact of lifetime reweighting. The lifetime reweighting is a dominant source of uncertainty when applicable, comparable to the theory systematic uncertainty from variations in μ_F and μ_R and the uncertainty in signal reconstruction efficiency for large $|d_0|$. In addition to the uncertainties discussed in this chapter, uncertainties on resolution and scale corrections for simulated e/γ objects which are standardized

Uncertainty $\pm 1\sigma$ [%]	$\tilde{e}\tilde{e}$, 100 GeV 50 ns	eBDT background template
Total background expectation	52.76	2603.00
Total statistical ($\sqrt{N_{\text{exp}}}$)	± 7.26	± 51.02
Total systematic	± 15.92 [30.18%]	± 182.05 [6.99%]
Lifetime reweighting	± 9.67 [18.3%]	-
large- $ d_0 $ efficiency	± 8.52 [16.1%]	-
μ_F, μ_R	± 6.33 [12.0%]	-
e ID scale	± 5.51 [10.5%]	-
Large Radius Tracking	± 2.65 [5.0%]	-
Luminosity	± 1.16 [2.2%]	-
PDF variations	± 1.06 [2.0%]	-
E_T^{miss}	± 1.06 [2.0%]	-
Pileup reweighting	± 0.83 [1.6%]	-
e/γ scale	± 0.30 [0.56%]	-
e reconstruction	± 0.16 [0.30%]	-
e/γ resolution	± 0.05 [0.10%]	-
e trigger scale	± 0.02 [0.03%]	-
Time Smearing	± 0.00 [0.00%]	-

Table 9.2: Breakdown of the dominant systematic uncertainties on the EM-BDT background template and signal yields in the eBDT region, for a lifetime-reweighted 100 GeV 50 ns selectron sample. The systematic uncertainty on the background template is due only to the non-closure shape uncertainty. Note that the individual uncertainties can be correlated, and do not necessarily add up quadratically to the total background uncertainty. The percentages show the size of the uncertainty relative to the total expected background. For shape systematic uncertainties, the percentage represents the average relative uncertainty per bin.

across simulations used for ATLAS analyses are included in both tables. An uncertainty on the reweighting of simulated samples to match the observed pileup profile in data is also listed. These additional variations have less than 1% impact on signal yields and therefore negligibly impact the final result.

Uncertainty of channel	$\tilde{e}\tilde{e}$, 300 GeV 10 ns	γ BDT background template
Total background expectation	7.38	11766.00
Total statistical ($\sqrt{N_{\text{exp}}}$)	± 2.72	± 108.47
Total systematic	± 1.27 [17.25%]	± 5255.04 [44.66%]
μ_F, μ_R	± 0.89 [12.0%]	-
Luminosity	± 0.16 [2.2%]	-
PDF variations	± 0.15 [2.0%]	-
E_T^{miss}	± 0.15 [2.0%]	-
γ ID scale	± 0.12 [1.7%]	-
Large Radius Tracking	± 0.11 [1.4%]	-
γ Isolation scale	± 0.07 [0.99%]	-
E_T^{miss} resolution $\parallel$	± 0.01 [0.17%]	-
E_T^{miss} resolution $\perp$	± 0.01 [0.17%]	-
Pileup reweighting	± 0.01 [0.15%]	-
e/γ scale	± 0.00 [0.05%]	-
Time smearing	± 0.00 [0.05%]	-
E_T^{miss} scale	± 0.00 [0.03%]	-
e/γ resolution	± 0.00 [0.01%]	-
γ trigger scale	± 0.00 [0.00%]	-

Table 9.3: Breakdown of the dominant systematic uncertainties on the EM-BDT background template and signal yields in the γ BDT region, for a lifetime-reweighted 300 GeV 10 ns selectron sample. The systematic uncertainty on the background template is due only to the non-closure shape uncertainty. Note that the individual uncertainties can be correlated, and do not necessarily add up quadratically to the total background uncertainty. For shape systematic uncertainties, the percentage represents the average relative uncertainty per bin.

CHAPTER 10:

STATISTICAL ANALYSIS AND FIT

PROCEDURE

The goal of this thesis is to combine the background estimates, simulated signal samples, systematic uncertainties, and collected data described thus far into a quantitative statement about the presence or absence of an LLP signal in the observed data upon unblinding the SRs. Such a statement can be made using *hypothesis testing* via a *profile likelihood ratio*, statistics concepts which are described in Section 10.1. The remainder of the chapter is dedicated to the preparation of the EM-BDT regions for use in the fit procedure, and validations performed before the SRs were unblinded to ensure the robustness of the fit framework.

10.1 STATISTICS CONCEPTS

A hypothesis test is constructed to determine the compatibility of the observed data with two possible hypotheses. The first, known as the null hypothesis H_0 , is a *background-only* hypothesis which supposes that the signal is not present in the dataset, and therefore the only observed data in the SRs are due to (expected) background events. The alternate hypothesis H_1 is a *signal-plus-background* (S+B) which supposes that the signal is present in the dataset, in addition to the background events.

A theoretical cross-section σ_{theory} is calculated for each slepton signal model studied. These cross sections are listed in Table 7.1, and vary as a function of slepton mass and collision center-of-mass energy. The number of signal events expected for a particular model in each SR, for a given luminosity, is directly proportional to the theoretical cross section. Each model can be generalized to include a continuous range of observed cross-sections (σ_{obs}), and therefore a continuous range of observed signal events, via the introduction of the *signal strength* parameter $\mu = \sigma_{\text{obs}}/\sigma_{\text{theory}}$. The total number of events predicted in a single SR bin can be written as

$$N_{\text{pred}} = \mu \cdot S + B \quad (10.1)$$

Here, S is the signal yield predicted by theory in this bin, B is the background yields predicted by the background estimations in this bin, and μ can vary smoothly from 0 to 1 to modify the number of expected signal events. H_0 corresponds to $\mu = 0$, and H_1 corresponds to $\mu = 1$. It is assumed that the observed data follows a Poisson distribution. Therefore, the likelihood that the number of data events observed in this bin corresponds to a model with signal strength μ can be written as

$$\mathcal{L}(\mu) = \frac{(\mu \cdot S + B)^{N_{\text{obs}}}}{N_{\text{obs}}!} \exp[-(\mu \cdot S + B)] \quad (10.2)$$

Systematic and statistical uncertainties on the predicted signal(background) yields are introduced into the likelihood function $\mathcal{L}(\mu, \theta)$ as *Nuisance Parameters* (NPs) denoted $\theta_{S(B)}$. Each NP is modeled using a Gaussian distribution encoding the central value and uncertainty described by the NP. The likelihood function across all bins is the joint distribution of the likelihood in each bin:

$$\mathcal{L}(\mu, \theta) = \prod_i \prod_j \prod_k \mathcal{P}(S^i, B^i, N_{\text{obs}}^i) \mathcal{G}(\tilde{\theta}_S^j | \theta_S^j, \sigma_S^j) \mathcal{G}(\tilde{\theta}_B^k | \theta_B^k, \sigma_B^k) \quad (10.3)$$

Here, the index (i) runs over the SR bins, $\mathcal{P}$ are Poisson PDFs of equation 10.2, $\mathcal{G}$ are Gaussian PDFs for each of the (j) signal systematic uncertainties and (k) background systematic uncertainties, and $\tilde{\theta}$ is the measured value of the NP given best estimate θ with width σ . Before fitting, each NP is rescaled $\theta \rightarrow \alpha = (\theta - \tilde{\theta})/\sigma$ to have a mean of zero and unity width. The negative natural logarithm of the likelihood function $(-\ln \mathcal{L})$ is then subject to a global minimization, allowing the signal strength and the NPs to vary to arrive at fitted values $\hat{\mu}, \hat{\alpha}$. The function $-\ln \mathcal{L}$ is also subject to a conditional minimization to return NP values $\hat{\alpha}$ for a fixed strength μ . The profile likelihood $\lambda(\mu)$ is then defined as the ratio of the likelihoods corresponding to these minima [120]:

$$\lambda(\mu) = \frac{\mathcal{L}(\mu, \hat{\alpha})}{\mathcal{L}(\hat{\mu}, \hat{\alpha})} \quad (10.4)$$

This ratio is used to calculate a p -value corresponding to a given signal strength μ [121], with the calculated value p_μ corresponding to the probability of a signal with strength μ , with p_0 corresponding to the probability of the background-only hypothesis. A small value of p_0 indicates a very small probability that the data would have been observed given that H_0 is true, while a value of p_0 near 0.5 indicates good agreement between observed data and the background-only hypothesis. From these p -values, a *test statistic* CL_s is used to define a Confidence Level (CL) for a signal hypothesis via the so-called CL_s method [122,123]:

$$\text{CL}_s = \frac{p_\mu}{1 - p_0} \quad (10.5)$$

The division by $(1 - p_0)$ here serves to avoid claiming exclusion of a signal hypotheses where the experiment has little sensitivity, by inflating the computed CL_s when the p -value distributions for the background-only hypothesis and signal hypotheses are nearly indistinguishable. Note that these distributions are guaranteed to approach one another as the signal strength approaches zero.

The statistical fit described above is implemented in the HistFitter framework developed by the ATLAS Collaboration [124], with a shared background normalization and an independent shared signal normalization across all SR bins allowed to float. Exclusion is claimed at 95% CL for signal hypotheses with strength μ yielding $\text{CL}_s < 0.05$. As μ is directly proportional to the cross section, an upper limit on the cross-section for each excluded signal model can therefore be computed with 95% CL. In the limit that zero background events are expected (as in all ABCD SRs constructed in this analysis) the number of observed events necessary to claim exclusion with 95% confidence using the CL_s method is three [125]. The remainder of this chapter is dedicated to the preparation of the EM-BDT region for use in the fit framework, and tests performed to validate the robustness of the framework.

10.2 EM-BDT BIN OPTIMIZATION

The number of bins, and bin widths, of the BDT response distributions are chosen based on an optimization procedure that considers the signal sensitivity and the available statistics in the background template. For each BDT region, a scan of bin numbers varying from 1 to 6 total bins is applied using CR data and SR signal simulation. The chosen bins are required to have monotonically increasing widths, with a minimum width of 0.01 and at least 1 background event in each bin. The parameter optimized is a sum in quadrature across bins of sensitivity $Z = \sqrt{\sum_i Z_i^2}$, where

$$Z_i = \left(\frac{-1}{2 \left(B_i \log \left(1.0 + \frac{S_i}{B_i} \right) - S_i \right)} + \frac{B_i}{S_i^2} + \frac{1}{S_i} \right)^{-\frac{1}{2}} \quad (10.6)$$

and B_i , S_i are the values of background and signal in bin i . This formula is an extension of the typical asymptotic $Z = S/\sqrt{B}$, with added shape terms corresponding to statistical uncertainties of the background and signal shape in each bin for the case of low-statistics

bins. The first bin edge is fixed at zero. An example of the optimization procedure for the eBDT, showing the Z-score calculated in each BDT score bin and Z-score summed across bins for a representative selectron signal sample with 300 GeV mass and 10 ns lifetime, is shown in Figure 10.1.




Figure 10.1: eBDT binning resulting from the bin optimization procedure for a selectron signal sample with 300 GeV mass and 10 ns lifetime. The number of events (normalized to 56.3 fb^{-1} for signal) in each bin are shown. The ratio panel shows the Z-score calculated from Equation 10.6 in each bin, with the dominant sensitivity coming from the most signal-like bin.

Events with BDT response less than zero are very background-like and are found to be insensitive to signal, such that the result of the fit does not change significantly when these events are omitted. Therefore, the bin $\in [-1, 0]$ is not used in the fit. The optimal bins chosen via this procedure for the eBDT are $[0, 0.11, 0.22, 0.33, 0.46, 1]$ and for the γ BDT are $[0, 0.17, 0.34, 0.52, 0.72, 1]$. In the case of both BDTs, exactly one event is predicted by the background template in the most signal-like BDT score bin.

10.3 EM-BDT BACKGROUND SUBTRACTION

For selectron signal models with $\tau(\tilde{e}) \leq 1 \text{ ns}$, and in particular for models with smaller selectron masses, a significant number of signal events have a measured calorimeter time of $t < 0 \text{ ns}$, as the time-of-flight delay due to the mass and lifetime of the selectron is

comparable to the scale of the detector resolution. The impact across the selectron signal grid, and for a representative model point, are shown in Figure 10.2.

The contribution of signal events for these signal models to the shape of the background template in the case that H_1 is true, known as *signal contamination* of the CR, must be considered. This is done by modifying the signal template shape when fitting, such that the template used in the fit is the result of subtracting the distribution for signal events with $t < 0$ from the distribution for signal events with $t > 0$. This procedure accounts for the potential contribution of signal to the background template in the case $\mu \neq 0$, and functionally serves to weaken the signal sensitivity at lower lifetimes by decreasing the signal template normalization in the signal-like BDT bins. In the low-lifetime regime where the effect is non-negligible, the signal sensitivity is dominated by the ABCD strategy, mitigating the impact on the results. Conversely, the effect is negligible at higher lifetimes, where sensitivity to LLP models is driven by the BDT strategy.




Figure 10.2: Contamination of selectron events in the eBDT CR shown (left) as a fraction of the total number of SR and CR events across the selectron signal grid, and (right) in optimized eBDT score bins for a 300 GeV, 1 ns selectron signal point. The signal template shape used in the fit for this signal point is the result of subtracting the negatively-timed signal (pink) from the positively-timed signal (green) to account for the signal events present in the background template if the signal strength is non-zero.

10.4 EM-BDT SIGNAL INJECTION TESTS

The background estimation and fit procedure is required to be sufficiently robust that a quantitative statistical statement can be made about a dataset if H_1 is found to be true (that is, if a signal is present in the observed data). To ensure that this is the case for the EM-BDT regions, a check known as *signal injection* is performed. A signal model is injected into the observed data at a number of signal strengths, and a fit is performed in each case to determine whether the fit procedure is able to identify a signal of the injected strength in that modified data.

Figure 10.3 shows the results of signal injection tests using example selectron model points with $m(\tilde{e}), \tau(\tilde{e}) = 200 \text{ GeV}, 1 \text{ ns}$ ($300 \text{ GeV}, 10 \text{ ns}$) for an injected signal strength ranging from $(\mu = 0.2)$ to $(\mu = 1.0)$. The signal is injected into the CR data and an S+B fit is performed with the BDT shape systematic implemented on the background template, to verify that the magnitude and shape of the systematic uncertainty does not inhibit the ability of the fit framework to find a signal. Good linearity is found between injected and fitted signal strength for both models tested, indicating that the fit procedure is sensitive to the signal models under search, even when a signal is present with a faint strength. The larger uncertainties for 300 GeV, 10 ns model result primarily from the smaller predicted number of signal events in the SRs for this model with respect to the 200 GeV, 1 ns model.

10.5 BACKGROUND-ONLY FITS

A test of the fit framework is performed by fitting the CR template to data in the W/Z VRs, in both background-only and S+B scenarios. Figure 10.4 shows an example of pre-fit and post-fit S+B fit distributions of the eBDT score in negative-timed VR events, along with fit pulls for a selection of dominant signal and background systematic uncertainties.




Figure 10.3: Signal injection tests using the EM-BDT VR for selectron signal points with (left) $m(\tilde{e}), \tau(\tilde{e}) = 200 \text{ GeV}, 1 \text{ ns}$ and (right) $m(\tilde{e}), \tau(\tilde{e}) = 300 \text{ GeV}, 10 \text{ ns}$. The line ($y = x$) is drawn for comparison to indicate the ideal relationship between injected and fitted signal strength. For both models, the injected signal agrees well with the signal strength resulting from the fit.

A selectron model with $m(\tilde{e}), \tau(\tilde{e}) = 300 \text{ GeV}, 10 \text{ ns}$ is used as the signal. The fit pulls the signal strength μ_S down, while μ_B is pulled slightly up as a result of slightly higher statistics in the negatively-timed VR with respect to the CR. The pulls on the systematic uncertainties affecting the signal sample are minimal. The pull on the BDT shape systematic uncertainty $\alpha_{\text{BDT shape}}$ reflects the different BDT score shape expected between the VR and CR, as the W/Z -enriched VR contains events with different kinematics than the CR. These kinematic differences are not expected in the final fit to EM-BDT SR data using the CR-derived background template, as the EM-BDT SR and CR backgrounds differ only in the measured sign of the calorimeter time of flight. The pull remains within 1σ from the central value of the systematic uncertainty, with a reasonable width. Overall, the test shows good modeling of background from the CR-derived template, and reasonable estimates of the widths and shapes of signal and background systematic uncertainties, indicating a robust fit framework.




Figure 10.4: Pre-fit (top left) and post-fit (top right) S+B fit distributions of the eBDT score, using the background template derived in the CR to fit negative-timed VR data. A selectron model with $m(\tilde{e})$, $\tau(\tilde{e}) = 300$ GeV, 10 ns is chosen as the test signal sample. The red distribution corresponds to the ratio between the observed data and the pre-fit S+B hypothesis. Pulls on systematic uncertainties resulting from the fit are also shown (bottom).

CHAPTER 11:

RESULTS

Following the successful tests of the background estimation method and fit framework using the VRs in Chapter 10, the data in the SRs were unblinded. In this chapter, the results obtained from applying the fit framework to this unblinded data presented. No significant deviation from expected SM backgrounds are observed, and the results are therefore interpreted to set 95% CL upper limits on the pair-produced slepton cross-section, $\sigma(pp \rightarrow \tilde{l}\tilde{l})$. Limits are presented as a function of $m(\tilde{l})$ and $\tau(\tilde{l})$ for all three slepton generations, and extend the mass and lifetimes excluded with respect to the previous result [11]. Model-independent limits are also presented using event yields in designated *discovery regions*. In Section 11.4, a Dark Chargino model probed only by the ABCD analysis regions is discussed as a potential area of future study for the EM-BDT analysis.

11.1 EM-BDT SLEPTON RESULTS

For selectron models, the novel strategy employed in the EM-BDT regions provide sensitivity to selectron models with lifetimes $\tau(\tilde{e}) \sim 10$ ns and above which are inaccessible using only the ABCD methods. The BDT score distributions for SR data in both EM-BDT regions, overlaid against the pre-fit background template including systematic and statistical uncertainties, and the pre-fit BDT score distribution for a few representative signal samples, are shown in Figure 11.1. The ratio panel shows agreement between the ob-

served data and the background template within the uncertainties assigned. The overlaid signal score distributions illustrate the clear shape difference expected between the observed data and background template in the signal-plus-background case, in particular in the more signal-like BDT score bins.




Figure 11.1: BDT score distributions for the (a) eBDT and (b) γ BDT, comparing the expected yields from the control region to the observed data. Selectron pair-production models with a mass of 300 GeV and lifetime of 1 ns, and mass of 200 GeV with lifetime of 1 ns or 10 ns are overlaid with 10 times the expected yield. The lower panel is the ratio of the observed data over the expected background. The shaded blue bands represent the quadrature sum of the systematic and statistical uncertainties on the predicted background event yields in each bin.

As no significant excess with respect to predicted SM backgrounds is observed, the results are used to set limits on the pair-produced slepton cross-section $\sigma(pp \rightarrow \tilde{l}\tilde{l})$ with the CL_s method [122,123] under the asymptotic approximation [31]. The 95% CL combined exclusion contours for selectron models are the result of fitting the EM-BDT regions and all ABCD regions simultaneously, again requiring a shared signal strength parameter μ . The result of the combined fit is shown in Figure 11.2, as well as the contribution to the final result due to the ABCD high- p_T , ABCD LRT, and EM-BDT analysis regions individually. Selectrons with a lifetime of 0.3 ns are excluded up to 740 GeV by the combined fit. The improvement over the previous ATLAS run 2 result [11] for longer lifetimes is driven entirely by the EM-BDT analysis. The EM-BDT contour, shown in purple, allows new exclusion for selectron models with lifetimes in the $\tau(\tilde{e}) = 10\text{-}100$ ns range for $m(\tilde{e})$

up to 400 GeV which are not excluded using the ABCD strategy. For example, selectrons with a mass of 150 GeV are excluded by the EM-BDT analysis for lifetime values ranging from 4 ps to 60 ns, whereas the exclusion for 150 GeV selectrons from the previous ATLAS result reaches only up to 20 ns.




Figure 11.2: Expected and observed 95% CL exclusion contours for the selectron model from the combined fit of all analysis regions and from individual fits to ABCD high- p_T , ABCD LRT, and EM-BDT analysis regions. The dashed lines indicate the expected limits. The gray shaded area illustrates the observed exclusion obtained in the prior ATLAS Run 2 analysis [11].

The *depth* of exclusion given by the EM-BDT contour, comparing the theory cross section for signal models to the maximum cross section allowed for each model by the fit, is visualized as a function of $m(\tilde{e})$ and $\tau(\tilde{e})$ in Figure 11.3 for a few representative selectron masses and lifetimes. Masses(lifetimes) for which the colored band falls below the corresponding theory cross-section curve are excluded by the EM-BDT standalone fit, with the strongest lifetime limits set at $\tau(\tilde{e}) \sim 1$ ns for a 300 GeV selectron model.




Figure 11.3: 95% CL limits on $\sigma(pp \rightarrow \tilde{e}\tilde{e})$ as a function of (left) $m(\tilde{e})$ and (right) $\tau(\tilde{e})$ for a few representative selectron models. The theory cross-section corresponding to a signal strength of $\mu = 1$, which varies as a function of selectron mass, is also shown.

11.2 ABCD SLEPTON RESULTS

The smuon models considered produce only final states where the displaced lepton is a muon. For stau models, final states with one or two electrons or muons are possible. However, due to the large hadronic branching fraction of the τ lepton, the fraction of events containing displaced electrons in the final state is sufficiently small so as to render the EM-BDT analysis relatively insensitive to stau models when compared to the ABCD analysis. Therefore, the exclusion limits as a function of slepton mass and lifetime smuon and stau scenarios considered are driven entirely by the ABCD regions.¹ The 95% observed and expected limits on $\sigma(pp \rightarrow \tilde{\mu}\tilde{\mu})$ and $\sigma(pp \rightarrow \tilde{\tau}\tilde{\tau})$ are shown in Figure 11.4. The expected limits and corresponding $\pm 1\sigma$ band result from performing a simultaneous background-only fit on pseudo-datasets generated from the background predictions in the SRs. The observed limits are the result of applying a simultaneous fit to the observed yields in the SRs, constraining the signal yields across SRs by requiring a shared signal strength parameter μ . The 95% CL upper limits set by the previous ATLAS analysis [11] are significantly outperformed by this result, as a function of both $m(\tilde{l})$ and $\tau(\tilde{l})$. Smuons

¹For the smuon interpretation, the EM-BDT signal regions are not included in the combined fit. For the stau interpretation, the EM-BDT regions are included by minimally impact the exclusion reach and depth.

and staus with a lifetime of 0.3 ns are excluded up to 830 GeV, and 440 GeV, respectively. For smuons, in addition to the LRT-based triggers, the improved sensitivity is due to better trigger acceptance for muons in the forward region in the Run 3 dataset. The sensitivity to stau models also benefits from the new LRT triggers, since the final-state leptons tend to have lower p_T values accessible only by the LRT triggers.




Figure 11.4: Expected and observed 95% CL exclusion contours for the smuon (left) and stau (right) models. The dashed line indicates the expected limit and the surrounding band shows the 1σ variation of the expected limit due to uncertainties on the background prediction and experimental uncertainties on the signal modeling. The gray shaded area illustrates the observed exclusion obtained in the prior ATLAS Run 2 analysis [11].

11.3 MODEL INDEPENDENT RESULTS

A generic, model-independent statement about the presence of displaced lepton events in the dataset can be made in so-called discovery regions by comparing unblinded event yields to background predictions, without a comparison to a simulated signal model. For the ABCD analysis, each of the nine SRs already defines a single bin constructed to target displaced lepton events, and is therefore agnostic to any signal hypothesis motivated by a specific model. Each of these nine bins is used as a discovery region. The γ BDT explicitly relies on a signature to have a large value of E_T^{miss} and multiple displaced/delayed leptons; these choices are quite model-dependent and the γ BDT is therefore not considered

when constructing discovery regions. The eBDT, by contrast, is explicitly object-based, as it requires only the measured properties of a single electron per event and does not rely on event-level information when classifying events as signal-like or background-like. The most signal-like bin (eBDT score > 0.46), is therefore used as a discovery region in the model-dependent interpretation. Exactly one background event with a symmetric systematic uncertainty of 33% is expected in the EM-BDT discovery region.

The unblinded model-independent results are shown in Table 11.1. In eight of the nine ABCD discovery regions, zero events are observed, consistent with the near-zero background predictions. In the ee -LRT signal region, 1 event is observed, corresponding to a local significance of 2.2σ . In the EM-BDT discovery region, three events are observed where one is predicted, corresponding to a local significance of 0.9σ . As no significant signal excess is observed in the SRs, the results are used in the following sections to set 95% CL exclusion limits on the signal production cross section $\sigma(pp \rightarrow \tilde{l}\tilde{l})$ for each of the three slepton generations via the CLs technique [122, 123] under the asymptotic approximation [121].

11.4 DARK CHARGINO INTERPRETATION

For the ABCD regions, interpretations are also performed using a Higgs portal model that introduces a dark sector consisting of a chargino ($\tilde{\chi}^\pm$) and neutralino ($\tilde{\chi}^0$) with mass difference $\Delta m \equiv m(\tilde{\chi}^\pm) - m(\tilde{\chi}^0)$ [126], which can produce the observed dark matter relic density for $10 \lesssim \Delta m \lesssim 40$ GeV [127]. For the signal process $\tilde{\chi}^+\tilde{\chi}^- \rightarrow W^+W^-\tilde{\chi}^0\tilde{\chi}^0$ shown in Figure 11.5, the charged state can acquire a non-negligible lifetime due to the weak coupling between the dark states and the W boson.

If the W bosons decay leptonically, this model results in events with two displaced leptons, the same final state as with the GMSB slepton model. However, with this model the

Discovery Region	Total Bkg.	N_{obs}	$\langle A\epsilon\sigma \rangle_{\text{obs}}^{95}$ [fb]	S_{obs}^{95}	S_{expected}^{95}	CL_b	$p(s=0)$ (Z)
ee high- p_T (Run 2)	0.031 ± 0.031	0	0.02	3.0	$3.0_{-0.0}^{+0.0}$	0.46	0.5 (0)
ee high- p_T (Run 3)	0.06 ± 0.05	0	0.05	3.0	$3.0_{-0.1}^{+0.0}$	0.45	0.5 (0)
ee LRT	$0.0016_{-0.0016}^{+0.0029}$	1	0.07	4.1	$3.0_{-0.1}^{+0.0}$	0.97	0.01 (2.2)
$\mu\mu$ high- p_T (Run 2)	$0.02_{-0.02}^{+0.22}$	0	0.02	3.0	$3.0_{-0.0}^{+0.0}$	0.49	0.5 (0)
$\mu\mu$ high- p_T (Run 3)	$0.01_{-0.01}^{+0.11}$	0	0.05	3.0	$3.0_{-0.0}^{+0.0}$	0.48	0.5 (0)
$\mu\mu$ LRT	$0.02_{-0.02}^{+0.04}$	0	0.05	3.0	$3.0_{-0.0}^{+0.0}$	0.49	0.5 (0)
$e\mu$ high- p_T (Run 2)	$0.0016_{-0.0016}^{+0.0033}$	0	0.02	3.0	$3.0_{-0.1}^{+0.0}$	0.50	0.5 (0)
$e\mu$ high- p_T (Run 3)	$0.004_{-0.004}^{+0.010}$	0	0.05	3.0	$3.0_{-0.0}^{+0.0}$	0.50	0.5 (0)
$e\mu$ LRT	$0.2_{-0.2}^{+0.4}$	0	0.05	3.0	$3.0_{-0.1}^{+0.1}$	0.45	0.5 (0)
EM-BDT	$1.0_{-0.9}^{+2.3}$	3	0.13	7.1	$5.7_{-0.8}^{+1.9}$	0.77	0.2 (0.9)

Table 11.1: Expected and observed yields in the signal regions and model-independent upper limits derived using pseudo-experiments. Left to right: total expected background after a background-only fit, with the combined statistical and systematic uncertainties; number of events observed in data; 95% CL upper limits on the visible cross-section ($\langle A\epsilon\sigma \rangle_{\text{obs}}^{95}$) and on the number of signal events (S_{obs}^{95}). The sixth column (S_{expected}^{95}) shows the expected 95% CL upper limit on the number of signal events, given the expected number (and $\pm 1\sigma$ excursions of the expectation) of background events. The last two columns indicate the confidence level of the background-only hypothesis (CL_b) and discovery p -value with the corresponding Gaussian significance ($Z(s=0)$). CL_b provides a measure of compatibility of the observed data with the signal strength hypothesis at the 95% CL limit relative to fluctuations of the background, and $p(s=0)$ measures compatibility of the observed data with the background-only hypothesis relative to fluctuations of the background. The p -value is capped at 0.5.




Figure 11.5: Feynman diagram of dark chargino model [6].

leptons tend to have lower p_T values since the lighter neutral state is massive, highlighting the need for improved low-momentum lepton acceptance.

Figure 11.6 shows the 95% CL exclusion contours for the dark chargino model as a function of $m(\tilde{\chi}^0)$, at various mass splittings for fixed lifetime of $\tau = 0.3$ ns, and for various lifetimes at a fixed mass splitting of $\Delta m = 40$ GeV. For a lifetime of $\tau = 0.3$ ns, mass differences down to $\Delta m = 16$ GeV are excluded for a chargino mass of 100 GeV, and charginos with masses up to 900 GeV are excluded for large values of Δm . For a mass difference of $\Delta m = 40$ GeV, charginos with masses up to 380 GeV are excluded and chargino lifetimes of $0.006 < \tau < 10$ ns for a chargino mass of 100 GeV are excluded.




Figure 11.6: Expected and observed 95% CL exclusion contours for the dark chargino model with a (left) fixed lifetime and (right) fixed mass splitting.

The EM-BDT regions are not used for the chargino interpretation due to low EM-BDT SR yields observed for generated signals at fixed lifetime and fixed mass splitting, as shown in Figure 11.7. However, the generated signal samples require both W bosons to decay leptonically, as the ABCD analysis signal regions require two displaced leptons in each SR. A W decay will result in an electron, muon, or tau in approximately 33% of events, and fully leptonic final states therefore account for approximately 10% of events in processes with two independent W boson decays. Only this fraction of final states are considered in the yields of Figure 11.7.




Figure 11.7: Expected eBDT SR event yields scaled to 56.3 fb^{-1} for dark chargino models with a fixed chargino lifetime of 0.3 ns (left) and a fixed chargino-neutralino mass difference of 40 GeV (right).

A W boson will decay hadronically in 68% of events, therefore events resulting in one leptonic decay and hadronic decay will occur approximately 45% of the time. Thus, a future analysis might consider generating signal events including hadronic decay modes of exactly one W , leading to an expected factor of 4.5 in signal yields in the EM-BDT region. This factor is conservative, as the event yields in Figure 11.7 are required to be orthogonal to those events in the ABCD regions, but events featuring one hadronic decay will by definition meet this requirement as the ABCD regions require two leptons. The increased signal acceptance could allow for competitive sensitivity in EM-BDT regions for models with 1-10 ns chargino lifetimes.

CHAPTER 12:

CONCLUSION

This thesis has presented a search for displaced and delayed leptons using the full ATLAS Run 2 and partial Run 3 dataset. This work represents the first ATLAS search at the highest center-of-mass energy of 13.6 TeV ever achieved for $p - p$ collisions. No significant deviation from expected SM backgrounds are observed in the SRs, and the predicted and observed event yields are therefore used to set limits on LLP decays resulting in leptons in the final state. The result of the search is interpreted in the context of a GMSB slepton pair-production model, improving upon the previous ATLAS results due to the addition of more data, the use of novel triggers and reconstruction algorithms to target displaced objects, and the incorporation of LAr calorimeter timing information into a novel analysis strategy. For example, selectrons with a mass of 150 GeV are excluded for lifetime values up to 60 ns, improving on the previous exclusion reach for this selectron mass by a factor of approximately three. Smuons(staus) with a lifetime of 0.3 ns are excluded up to 830 GeV(440 GeV) respectively, whereas in the previous result the highest mass excluded for this lifetime was 700 GeV(310 GeV). The ABCD regions defining the two-lepton SRs, and the eBDT ML algorithm developed to identify displaced and delayed electrons, are largely model-independent and are also used to set limits on generic LLP decays resulting in displaced leptons.

The contribution of this result to the ATLAS LLP search program is visualized for $\tilde{\tau} \rightarrow \tau \tilde{G}$ models in Figure 12.1, which also includes ATLAS searches for a detector-stable stau [128, 129] and a stau which decays promptly [130]. Much of the $\tilde{\tau}$ lifetime regime between the prompt and stable searches is covered by this result, with some notable gaps still uncovered for potential future searches.




Figure 12.1: Summary of 95% CL exclusion limits based on 13 and 13.6 TeV data in the (mass, lifetime) plane for different analyses probing stau pair production, with decays $\tilde{\tau} \rightarrow \tau \tilde{G}$, in the GMSB scenario. The displaced leptons result presented in this thesis nearly bridges the full lifetime gap between previous ATLAS searches for prompt and stable $\tilde{\tau}$ models with masses less than 500 GeV.

The model independence of this analysis is exemplified by the re-interpretation of the result in the ABCD regions in a dark chargino scenario in Section 11.4. A straightforward continuation of the work presented could use the EM-BDT regions to probe the same scenario, including the cases where one W boson decays hadronically. As demonstrated with the slepton pair-production model, the addition of the EM-BDT to the search is expected to expand the search sensitivity for lifetimes on the order of 1 ns and greater. More generally, the eBDT algorithm in particular is potentially applicable to a wide range of LLP scenarios. The tool essentially functions as an extension of the electron ID for prompt electrons stan-

standardized across the ATLAS collaboration, tuned specifically to tag high-quality displaced and delayed electrons. Incorporation of this tool into a centralized ATLAS framework, a retraining over a wider range of models, and a subsequent re-tuning of the architecture and inputs would make the algorithm more generalizable.

The γ BDT developed does not benefit from the same model independence as the eBDT. In principle, the treatment of photons as displaced electrons should increase the range of lifetimes accessible in LLP searches; an electron reconstructed as a photon is likely to have very few ID hits and therefore is likely to have resulted from a more displaced/delayed decay than an electron reconstructed as an electron using standard or LRT tracking algorithms. The current γ BDT implementation relies on E_T^{miss} and requires multiple photons to combat the large background of fake single and multi-photon events in the signal-enriched search regions. Future work might explore the use of additional variables describing the photon shower shape in the electromagnetic calorimeter as BDT inputs, to increase signal-to-background discrimination and allow the analysis to relax requirements on E_T^{miss} or photon multiplicity. The development of an object-level displaced photon tagger analogous to the eBDT is in this way potentially feasible.

This analysis was performed with a partial Run 3 dataset amounting to 56.3 fb^{-1} of $p-p$ collision data. At the time of writing, the data collected by ATLAS in 2024 currently available for analysis represents an additional 110.0 fb^{-1} , and a similar additional luminosity is expected to be delivered in both 2025 and 2026. Adding this data to the existing result will immediately improve the depth and breadth of model-dependent exclusion limits, in particular at high slepton masses where the analysis begins to lose sensitivity due to rapidly decreasing cross sections.

Following Run 3, the LHC will be upgraded to deliver an instantaneous luminosity between 5 and 7.5 times its nominal design value, marking the start of the so-called High Luminosity LHC (HL-LHC) era. The expected integrated luminosity over the HL-LHC lifetime is 3 ab^{-1} , more than 50 times the size of the Run 3 dataset used in this thesis. The

addition of HL-LHC data to this search will improve the reach of the result, but as the HL-LHC progresses the luminosity doubling time of the LHC will experience diminishing returns. In this sense, the relative scarcity of new data used in this search represented a challenge as well as an opportunity to showcase the efficacy of novel techniques. Targeted LLP triggers, novel reconstruction methods, and search strategies relying on the exploitation of specialized variables and ML techniques, will become crucial in maximizing the potential of the ATLAS Collaboration's LLP search program for the remainder of Run 3 and in the HL-LHC era.

REFERENCES

- [1] ATLAS Collaboration, “LAr Calorimeter pilot beam results, Oct 2021.” <https://twiki.cern.ch/twiki/bin/view/AtlasPublic/LArCaloPublicPilotBeam2021>.
- [2] A. C. Smith, tech. rep., CERN, Geneva, 2022.
<https://cds.cern.ch/record/2807338>.
- [3] ATLAS Collaboration, tech. rep., CERN, Geneva, Apr. 2022.
<https://cds.cern.ch/record/2805799>.
- [4] ATLAS Collaboration, Phys. Rev. D **108** no. 1, (2023) 012012, arXiv:2304.12885.
<https://cds.cern.ch/record/2856888>.
- [5] A. C. Smith, tech. rep., CERN, Geneva, 2024.
<https://cds.cern.ch/record/2903758>. ICHEP, July 17th 2024.
- [6] ATLAS Collaboration, arXiv:2410.16835 [hep-ex].
<https://arxiv.org/abs/2410.16835>.
- [7] ATLAS Collaboration, JINST **3** (2008) S08003.
<http://stacks.iop.org/1748-0221/3/i=08/a=S08003>.
- [8] L. Evans and P. Bryant, JINST **3** (2008) S08001.
- [9] ATLAS Collaboration, Phys. Lett. B **716** no. 1, (2012) 1 – 29.
<http://www.sciencedirect.com/science/article/pii/S037026931200857X>.
- [10] CMS Collaboration, Phys. Lett. B **716** (2012) 30–61, arXiv:1207.7235 [hep-ex].
- [11] ATLAS Collaboration, Tech. rep. 5, 2021. arXiv:2011.07812.
<https://cds.cern.ch/record/2744332>.
- [12] A. Purcell, <https://cds.cern.ch/record/1473657>.
- [13] Particle Data Group, Progress of Theoretical and Experimental Physics **2020** no. 8, (2020). <https://doi.org/10.1093/ptep/ptaa104.083C01>.
- [14] M. Thomson, Modern particle physics. Cambridge University Press, New York, 2013.
- [15] E. Noether, Nachrichten von der Gesellschaft der Wissenschaften zu Göttingen, Mathematisch-Physikalische Klasse **1918** (1918) 235–257.
<http://eudml.org/doc/59024>.

- [16] F. Englert and R. Brout, Phys. Rev. Lett. **13** (1964) 321–323.
- [17] P. W. Higgs, Phys. Lett. **12** (1964) 132–133.
- [18] P. W. Higgs, Phys. Rev. Lett. **13** (1964) 508–509.
- [19] P. W. Higgs, Phys. Rev. **145** (1966) 1156–1163.
- [20] Y. Nambu and G. Jona-Lasinio, Phys. Rev. **122** (1961) 345–358.
- [21] S. Weinberg, Phys. Rev. Lett. **19** (1967) 1264–1266.
- [22] G. Costa and G. Fogli, Symmetries and group theory in particle physics. An introduction to space-time and internal symmetries, vol. 823. Springer Berlin, Heidelberg, 2012.
- [23] X. Fan, T. Myers, B. Sukra, and G. Gabrielse, Physical Review Letters **130** no. 7, (2023). <http://dx.doi.org/10.1103/PhysRevLett.130.071801>.
- [24] E598 Collaboration, J. J. Aubert et al., Phys. Rev. Lett. **33** (1974) 1404–1406.
- [25] SLAC-SP-017 Collaboration, J. E. Augustin et al., Phys. Rev. Lett. **33** (1974) 1406–1408.
- [26] D0 Collaboration, S. Abachi et al., Phys. Rev. Lett. **74** (1995) 2632–2637. <https://link.aps.org/doi/10.1103/PhysRevLett.74.2632>.
- [27] CDF Collaboration, F. Abe et al., Phys. Rev. Lett. **74** (1995) 2626–2631. <https://link.aps.org/doi/10.1103/PhysRevLett.74.2626>.
- [28] UA1 Collaboration, G. Arnison et al., Phys. Lett. B **122** (1983) 103–116. 31 p. <https://cds.cern.ch/record/142059>.
- [29] UA1 Collaboration, G. Arnison et al., Phys. Lett. B **126** no. 5, (1983) 398–410. <https://www.sciencedirect.com/science/article/pii/0370269383901880>.
- [30] ATLAS Collaboration, “Standard Model summary plots.” <https://atlas.web.cern.ch/Atlas/GROUPS/PHYSICS/PUBNOTES/ATL-PHYS-PUB-2022-009>.
- [31] M. Dine and A. Kusenko, Reviews of Modern Physics **76** no. 1, (2003) 1–30. <http://dx.doi.org/10.1103/RevModPhys.76.1>.
- [32] R. J. Adler, Am. J. Phys. **78** no. 9, (2010) 925–932. <http://dx.doi.org/10.1119/1.3439650>.
- [33] K. G. Begeman, A. H. Broeils, and R. H. Sanders, Mon. Notices Royal Astron. Soc. **249** no. 3, (1991) 523–537. <https://doi.org/10.1093/mnras/249.3.523>.
- [34] G. Bertone, D. Hooper, and J. Silk, Physics Reports **405** no. 5–6, (2005) 279–390. <http://dx.doi.org/10.1016/j.physrep.2004.08.031>.

- [35] The Supernova Cosmology Project Collaboration, S. Perlmutter et al., *Astrophys. J* **517** no. 2, (1999) 565–586. <http://dx.doi.org/10.1086/307221>.
- [36] J. L. Gervais and B. Sakita, *Nucl. Phys. B* (1971) 632–639.
https://www.worldscientific.com/doi/abs/10.1142/9789812795700_0019.
- [37] P. Ramond, *Phys. Rev. D* **3** no. 10, (1971) 2415–2418.
<https://cds.cern.ch/record/432382>.
- [38] H. Miyazawa, *Prog. Theor. Phys.* **36** no. 6, (1966) 1266–1276.
- [39] J. Wess and B. Zumino, *Nucl. Phys. B* **70** no. 1, (1974) 39–50.
<https://www.sciencedirect.com/science/article/pii/0550321374903551>.
- [40] G. L. Kane and M. Shifman, The Supersymmetric World: The Beginnings of The Theory. 2001. [arXiv:hep-ph/0102298](https://arxiv.org/abs/hep-ph/0102298) [hep-ph].
<https://arxiv.org/abs/hep-ph/0102298>.
- [41] S. P. Martin, *Adv. Ser. Direct. High Energy Phys.* **21** (2010) 1–153,
[arXiv:hep-ph/9709356](https://arxiv.org/abs/hep-ph/9709356).
- [42] L. Alvarez-Gaume and E. Witten, *Nucl. Phys. B* **234** (1984) 269.
- [43] S. Dimopoulos and H. Georgi, *Nucl. Phys. B* **193** (1981) 150–162.
- [44] Super-Kamiokande Collaboration, R. Matsumoto et al., [arXiv:2208.13188](https://arxiv.org/abs/2208.13188) [hep-ex]. <https://arxiv.org/abs/2208.13188>.
- [45] G. Sguazzoni, <https://cds.cern.ch/record/2296652>.
- [46] G. Giudice and R. Rattazzi, *Physics Reports* **322** no. 6, (1999) 419–499.
[http://dx.doi.org/10.1016/S0370-1573\(99\)00042-3](http://dx.doi.org/10.1016/S0370-1573(99)00042-3).
- [47] J. Alwall, M.-P. Le, M. Lisanti, and J. G. Wacker, *Phys. Lett. B* **666** (2008) 34,
[arXiv:0803.0019](https://arxiv.org/abs/0803.0019) [hep-ph].
- [48] J. Alwall, P. Schuster, and N. Toro, *Phys. Rev. D* **79** (2009) 075020,
[arXiv:0810.3921](https://arxiv.org/abs/0810.3921) [hep-ph].
- [49] D. Alves et al., *J. Phys. G* **39** (2012) 105005, [arXiv:1105.2838](https://arxiv.org/abs/1105.2838) [hep-ph].
- [50] J. A. Evans and J. Shelton, *JHEP* **04** no. 056, (2016) 056.
[http://dx.doi.org/10.1007/JHEP04\(2016\)056](http://dx.doi.org/10.1007/JHEP04(2016)056).
- [51] G. R. Farrar and P. Fayet, *Phys. Lett. B* **76** (1978) 575.
- [52] Y. Golfand and E. Likhtman, *JETP Lett.* **13** (1971) 323. [*Pisma Zh. Eksp. Teor. Fiz.* **13** (1971) 452].
- [53] D. Volkov and V. Akulov, *Phys. Lett. B* **46** (1973) 109.

- [54] J. Wess and B. Zumino, Nucl. Phys. B **70** (1974) 39.
- [55] J. Wess and B. Zumino, Nucl. Phys. B **78** (1974) 1.
- [56] S. Ferrara and B. Zumino, Nucl. Phys. B **79** (1974) 413.
- [57] A. Salam and J. Strathdee, Phys. Lett. B **51** (1974) 353.
- [58] R. Barbier et al., Phys. Rept. **420** (2005) 1–202, arXiv:hep-ph/0406039 [hep-ph].
- [59] B. C. Allanach et al., Phys. Rev. D **75** (2007) 035002, arXiv:hep-ph/0609263.
- [60] G. Giudice and A. Romanino, Nucl. Phys. B **699** (2004) 65, arXiv:hep-ph/0406088.
- [61] N. Arkani-Hamed and S. Dimopoulos, JHEP **06** (2005) 073, arXiv:hep-th/0405159.
- [62] M. Dine and W. Fischler, Phys. Lett. B **110** (1982) 227.
- [63] L. Alvarez-Gaume, M. Claudson, and M. B. Wise, Nucl. Phys. B **207** (1982) 96.
- [64] C. R. Nappi and B. A. Ovrut, Phys. Lett. B **113** (1982) 175.
- [65] T. Appelquist et al., Phys. Rev. D **64** (2001) 035002, arXiv:hep-ph/0012100 [hep-ph].
- [66] H.-C. Cheng, K. T. Matchev, and M. Schmaltz, Phys. Rev. D **66** (2002) 056006, arXiv:hep-ph/0205314 [hep-ph].
- [67] M. J. Strassler, arXiv:hep-ph/0607160.
- [68] L. Lee et al., Progress in Particle and Nuclear Physics **106** (2019) 210–255. <https://www.sciencedirect.com/science/article/pii/S0146641019300109>.
- [69] E. Mobs, <https://cds.cern.ch/record/2636343>.
- [70] D. Dominguez, <https://cds.cern.ch/record/1741036>. General Photo.
- [71] ATLAS Collaboration, “Luminosity public results run 3.” <https://twiki.cern.ch/twiki/bin/view/AtlasPublic/LuminosityPublicResultsRun3>.
- [72] ATLAS Collaboration, “ATLAS stand-alone event displays.” <https://twiki.cern.ch/twiki/bin/view/AtlasPublic/EventDisplayStandAlone>.
- [73] ATLAS Collaboration, JINST **19** (2024) P06029, arXiv:2401.06630. <https://cds.cern.ch/record/2886407>.
- [74] ATLAS Collaboration, JINST **19** no. 05, (2024) P05063, arXiv:2305.16623. <https://cds.cern.ch/record/2859916>.
- [75] ATLAS Collaboration, tech. rep., CERN, Geneva, 2017. <https://cds.cern.ch/record/2285582>.

- [76] ATLAS Collaboration, Tech. rep. CERN-LHCC-96-041, CERN, Geneva, 1996.
<https://cds.cern.ch/record/331061>.
- [77] W. Cleland and E. Stern, NIM A **338** no. 2, (1994) 467 – 497.
<http://www.sciencedirect.com/science/article/pii/0168900294913323>.
- [78] M. Aleksa et al., tech. rep., 2013. <https://cds.cern.ch/record/1602230>.
- [79] ATLAS Collaboration, tech. rep., Geneva, 1996.
<https://cds.cern.ch/record/331062>.
- [80] G. Avoni et al., JINST **13** no. 07, (2018) P07017–P07017.
<https://doi.org/10.1088%2F1748-0221%2F13%2F07%2Fp07017>.
- [81] J. Pequeno, “Event Cross Section in a computer generated image of the ATLAS detector.” 2008. <https://cds.cern.ch/record/1096081>.
- [82] ATLAS Collaboration, Eur. Phys. J. C **77** no. 5, (2017).
<https://doi.org/10.1140%2Fepjc%2Fs10052-017-4887-5>.
- [83] ATLAS Collaboration, Eur. Phys. J. C **77** (2016) 490. 87 p, arXiv:1603.02934.
<https://cds.cern.ch/record/2138166>.
- [84] ATLAS Tracking CP, “ATLAS Tracking Software Tutorial.” Accessed: 2024-08-22, Jul. 2021. <https://atlassoftwaredocs.web.cern.ch/internal-links/tracking-tutorial/>.
- [85] T. Cornelissen et al., Journal of Physics: Conference Series **119** no. 3, (2008) 032013. <https://doi.org/10.1088/1742-6596/119/3/032013>.
- [86] R. Frühwirth, NIM A **262** no. 2, (1987) 444–450.
<https://www.sciencedirect.com/science/article/pii/0168900287908874>.
- [87] ATLAS Collaboration, “ATLAS Experiment Implements Heterogeneous Particle Reconstruction with Intel oneAPI Tools.” General Photo, 2023.
<https://cds.cern.ch/record/2879387>.
- [88] ATLAS Collaboration, Eur. Phys. J. C **83** (2023) 1081, arXiv:2304.12867.
<https://cds.cern.ch/record/2856887>.
- [89] ATLAS Collaboration, JINST **14** (2019) P12006, arXiv:1908.00005 [hep-ex].
- [90] ATLAS Collaboration, Phys. Rev. D **88** (2013) 012001.
<https://link.aps.org/doi/10.1103/PhysRevD.88.012001>.
- [91] ATLAS Collaboration, Phys. Rev. D **90** (2014) 112005.
<https://link.aps.org/doi/10.1103/PhysRevD.90.112005>.
- [92] ATLAS Collaboration, Phys. Rev. D **108** no. 1, (2023) 012012, arXiv:2304.12885.
<https://cds.cern.ch/record/2856888>.

- [93] ATLAS Collaboration, Eur. Phys. J. C **79** no. 8, (2019).
<https://doi.org/10.1140%2Fepjc%2Fs10052-019-7140-6>.
- [94] ATLAS Collaboration, Eur. Phys. J. C **76** (2016) 292, arXiv:1603.05598 [hep-ex].
- [95] J. Illingworth and J. Kittler, Comput. Gr. Image Process. **44** no. 1, (1988) 87–116.
<https://www.sciencedirect.com/science/article/pii/S0734189X88800331>.
- [96] ATLAS Collaboration, Eur. Phys. J. C **77** (2017) 466, arXiv:1703.10485 [hep-ex].
- [97] M. Cacciari et al., JHEP **04** (2008) 063, arXiv:0802.1189 [hep-ph].
- [98] E. Metodiev, “The Fractal Lives of Jets.” Accessed: 2024-08-22, 2019.
<https://www.ericmetodiev.com/post/jetformation/>.
- [99] ATLAS Collaboration, tech. rep., CERN, Geneva, 2014.
<https://cds.cern.ch/record/1700874>.
- [100] ATLAS Collaboration, arXiv:2402.05858 [hep-ex].
- [101] A. Geron, Hands-On Machine Learning with Scikit-Learn, Keras, and TensorFlow: Concepts, Tools, and Techniques to Build Intelligent Systems. O’Reilly Media, Inc., 2nd ed., 2019.
- [102] A. Hoecker et al., tech. rep., CERN, Geneva, 2007.
<https://cds.cern.ch/record/1019880>. TMVA-v4 Users Guide.
- [103] J. Chen, D. J. Mahon, and J. Parsons, tech. rep., CERN, Geneva, 2020.
<https://cds.cern.ch/record/2707701>.
- [104] J. Alwall et al., JHEP **07** (2014) 079, arXiv:1405.0301 [hep-ph].
- [105] T. Sjöstrand, S. Mrenna, and P. Skands, JHEP **05** (2006) 026,
arXiv:hep-ph/0603175.
- [106] ATLAS Collaboration,, “ATLAS Pythia 8 tunes to 7 TeV data.”
ATL-PHYS-PUB-2014-021, 2014. <https://cds.cern.ch/record/1966419>.
- [107] R. D. Ball et al., Nucl. Phys. **B867** (2013) 244–289, arXiv:1207.1303 [hep-ph].
- [108] S. Agostinelli et al., Nucl. Instrum. Meth. A **506** (2003) 250.
- [109] ATLAS Collaboration, Comput. Softw. Big Sci. **6** (2022) 3, arXiv:2102.09495 [hep-ex].
- [110] S. Porteboeuf, T. Pierog, and K. Werner, pp. , 135–140. Gioi Publishers, 2010.
arXiv:1006.2967 [hep-ph].
- [111] T. Pierog, I. Karpenko, J. M. Katzy, E. Yatsenko, and K. Werner, Phys. Rev. C **92** no. 3, (2015) 034906, arXiv:1306.0121 [hep-ph].

- [112] ATLAS Collaboration, tech. rep., 2016. <https://cds.cern.ch/record/2206965>.
- [113] J. Fiaschi et al., Eur. Phys. J. C **83** no. 8, (2023) 707, arXiv:2304.11915 [hep-ph].
- [114] G. Avoni et al., JINST **13** (2018) P07017.
- [115] ATLAS Collaboration, Eur. Phys. J. C **83** (2023) 982, arXiv:2212.09379 [hep-ex].
- [116] ATLAS Collaboration, tech. rep., 2023. <https://cds.cern.ch/record/2853525>.
- [117] ATLAS Collaboration, tech. rep., 2024. <https://cds.cern.ch/record/2900949>.
- [118] ATLAS Collaboration, Eur. Phys. J. C **83** (2023) 1081, arXiv:2304.12867 [hep-ex].
- [119] ATLAS Collaboration, Eur. Phys. J. C **81** (2021) 578, arXiv:2012.00578 [hep-ex].
- [120] J. Neyman and E. S. Pearson, On the Problem of the Most Efficient Tests of Statistical Hypotheses. Springer New York, New York, NY, 1992.
https://doi.org/10.1007/978-1-4612-0919-5_6.
- [121] G. Cowan, K. Cranmer, E. Gross, and O. Vitells, Eur. Phys. J. C **71** no. 2, (2011).
<https://doi.org/10.1140/epjc%2Fs10052-011-1554-0>.
- [122] A. L. Read, J. Phys. G **28** (2002) 2693–2704.
- [123] T. Junk, NIM A **434** no. 2, (1999) 435–443.
<https://www.sciencedirect.com/science/article/pii/S0168900299004982>.
- [124] G. J. Besjes et al., J. Phys.: Conf. Ser. **664** no. 7, (2015) 072004.
<https://cds.cern.ch/record/2134618>.
- [125] K. Bjoerke et al., tech. rep., CERN, Geneva, 2017.
<https://cds.cern.ch/record/2280679>.
- [126] A. Filimonova and S. Westhoff, JHEP **02** (2019) 140, arXiv:1812.04628 [hep-ph].
- [127] F. Blekman et al., JHEP **11** (2020) 112, arXiv:2007.03708 [hep-ph].
- [128] ATLAS Collaboration, <https://cds.cern.ch/record/2870112>.
- [129] ATLAS Collaboration, JHEP **2306** (2023) 158, arXiv:2205.06013.
<https://cds.cern.ch/record/2809388>.
- [130] ATLAS Collaboration, JHEP **2405** (2024) 150, arXiv:2402.00603.
<https://cds.cern.ch/record/2888091>.

APPENDICES

APPENDIX A:

BDT HYPERPARAMETER OPTIMIZATION

When designing the electron BDT algorithm, a number of learning rates and maximum tree depths were scanned to check for so-called *overtraining* by comparing test and train classifier distributions in each fold of the cross-fold validation. Reducing the learning rate from the TMVA default of 0.5 to 0.1, as shown in Figure A.1 and reducing the maximum tree depth (`maxDepth`) from 5 to 3, as shown in Figure A.2, was found to increase agreement between test and train eBDT score distributions, mitigating overtraining effects. The degree of overtraining is quantified by comparing the AUC score of the ROC between training and test datasets used in the BDT folds, as shown in Figure A.3. Based on this study, a learning rate of 0.1 and a maximum tree depth of 3 were chosen to mitigate overtraining. The optimal hyperparameters chosen for the eBDT are also chosen for the γ BDT, and found to give good discrimination between signal and background while avoiding overtraining. This is quantified by the small differences observed between the training and test samples for the three folds used in the γ BDT as shown in Figure A.4.

A hyperparameter scan was performed to choose optimal values for eBDT inputs `Ntrees`, `separationType`, `minNodeSize`, `nCuts`. The eBDT was trained 100 times with a different `randomSplit` seeds for the folds, to estimate an uncertainty on the output AUC score in each fold. The score was observed to be relatively stable under variations of these hyperparameters, with an RMS of ± 0.0001 over the 100 trials. Each of the hyperparameters

was then varied individually, holding learning rate and maxDepth at the values selected above to assess the impact on AUC, shown in Figure A.5. The optimal values chosen are summarized in Figure A.6.




Figure A.1: eBDT score distribution, test (line) and train (dots) overlaid for learning rate of 0.1 (top row) and 0.5 (bottom row). Agreement between test and train, in particular between eBDT response 0.0 – 0.3, is better for learning rate of 0.1, implying less overtraining.

Folds for Max Depth = 5, 3




Figure A.2: eBDT output distribution, test (line) and train (dots) overlaid for maximum tree depth of 5 (top row) and 3 (bottom row). Agreement between test and train, in particular for eBDT response in the range $(-0.2, 0.4)$, is better for a maximum tree depth of 3.

nTrees=1200, separationType=GiniIndex, minNode=0.6%, nCuts=70:




Figure A.3: Example ROC curves from 3 folds of BDT training for a number of tree depths and learning rates. While higher learning rates and maxDepth settings show slightly better training performance, It is found that the ROC score between test and training sample show the best agreement for learning rate of 0.1 and maxDepth of 3 (far left).




Figure A.4: γ BDT response for the three folds and corresponding test/training ROCs for the three folds of the γ BDT. Based on the ROC comparisons, minimal evidence of overtraining is observed.

Max Tree Depth	Av ROC Score	Learning Rate	Av ROC Score	Separation Type	Av ROC Score
3	0.99778	0.05	0.99791	GiniIndex	0.99795
4	0.99799	0.07	0.99795	CrossEntropy	0.99794
5	0.99806	0.1	0.997982		
Min Node Size	Av ROC Score	Number of Trees	Av ROC Score	Number of Cuts	Av ROC Score
0.2%	0.99868	600	0.99790	10	0.99569
0.6%	0.99872	800	0.99794	30	0.99866
1%	0.99872	1000	0.99796	50	0.99871
4%	0.99865	1200	0.99798	70	0.99872
7%	0.99496				

Figure A.5: ROC AUC scores for a hyperparameter scan over eBDT inputs Ntrees, separationType,minNodeSize,nCuts.

numTrees	800
separationType	GiniIndex
minNodeSize	1%
nCuts	50
Learning Rate	0.1
maxDepth	3

Figure A.6: Optimal eBDT parameters chosen via hyperparameter scan. This combination of parameters is found to maximize performance while minimizing overtraining as quantified via AUC difference between train/test sets among k-folds. The average AUC score returned for the three folds is 99.82%. The same set of hyperparameters are chosen for the γ BDT.

APPENDIX B:

LIFETIME REWEIGHTING

As described in Section 9.3.6, a so-called *lifetime reweighting* procedure is used to interpret the result for slepton models with lifetimes that were not generated, while preserving computing resources. The procedure takes as input a generated *source* sample and reweights to a *target* sample with a new lifetime. In particular, when reweighting to target lifetimes higher than those generated, the lack of statistics in long-lifetime tails of source distributions renders the procedure valid only up to a finite range of target lifetimes. A systematic uncertainty is assessed for the reweighting procedure based on observed disagreements between source-target pairs.

B.1 TWO-DIMENSIONAL LIFETIME REWEIGHTING

For the ABCD analysis, the two-dimensional lifetime reweighting equation given in Section 9.3.6 is used. Figure B.1 gives an example of the reweighting procedure for a source sample with a lifetime of 10 ns reweighted to 1 ns and 30 ns, with a narrower(wider) decay spectrum shown for the 1 ns(30 ns) target compared to the source, as expected. To validate the two-dimensional lifetime reweighting procedure, truth particle lifetime distributions are compared between the reweighted and generated samples. A simulated selectron sample with a source lifetime of 1 ns, and a sample with a lifetime of 10 ns are reweighted to a target lifetime of 30 ns. Figure B.2 shows a comparison between the source and target

lifetimes distributions, showing good agreement and indicating that the reweighting procedure is valid in this range. Note that in ABCD regions, the reweighting of samples to lifetimes longer than those generated contributes negligibly to the sensitivity of the analysis (see, for instance, Figure 11.2) as the EMBDT regions dominate the result at lifetimes greater than approximately 10-15 ns. The ABCD regions dominate the sensitivity of the result when reweighting to lifetimes *lower* than those generated, where the lack of statistics in long-lifetime tails does not impact the validity of the reweighting procedure.




Figure B.1: Example of two-dimensional lifetime reweighting of the decay spectrum of the subleading selectron in each event, obtained from reweighting a 100 GeV, 10 ns selectron sample to 1, 10, and 30 ns.

B.2 ACCEPTANCE LOSS FOR LONG-LIFETIME DECAYS

When investigating the reweighting procedure for EM-BDT regions, it was found that the spatial coordinates of slepton decays occurring outside of the ATLAS detector are not saved to the so-called *truth record*, meaning that this information is no longer accessible after the samples are generated. This effect is demonstrated in Figure B.3. In the truth record, 99.8% of decays in 30,000 events are saved for a 100 GeV 1 ns selectron sample (containing 60,000 total slepton decays), where decays occurring at $|z|$ values beyond the ATLAS detector are lost. The resulting acceptance loss grows with selectron lifetime, and is 99.8%, 82.4%, 53.7%, for samples generated with lifetimes of 1, 10, and 30 ns respectively.




Figure B.2: The distributions of truth particle lifetimes for leading (left) and subleading (right) selectron samples in ABCD regions. In each plot, the lifetime distributions for the source and the target simulated samples are shown. The ratio panel shows reasonable agreement between source and target distributions.

This loss of truth-level information does not directly impact the analysis strategy, as the decays are outside of ATLAS and therefore the resulting electrons cannot be reconstructed for use in signal regions. However, in the two-dimensional lifetime reweighting equation of Section 9.3.6 the decay coordinates for both SUSY particles are needed to calculate a valid weight for each event.

B.3 ONE-DIMENSIONAL LIFETIME REWEIGHTING

For EM-BDT regions, where two leptons are not required and one decay therefore often occurs outside of ATLAS, a one-dimensional lifetime reweighting is implemented. This takes the lifetime of the selectron which is truth-matched to the electron that is used in the




Figure B.3: Distribution of $|z|$ for the leading selectron decay in selectron samples with 1ns, 10ns, 30ns mean lifetimes, showing a cutoff at the edge of the ATLAS detector beyond which selectrons are not decayed in simulation. This leads to an acceptance loss for lifetime reweighting, which grows with lifetime.

BDT, effectively "integrating out" the decay spectrum of the unused slepton in each event.

The one-dimensional reweighting equation becomes

$$wt(\tau_1 \rightarrow \tau_2) = \left(\frac{\tau_2}{\tau_1} \right) \exp \left[(-t_{\tilde{e}}) \times \left(\frac{1}{\tau_1} - \frac{1}{\tau_2} \right) \right]$$

Figure B.4 shows an example of the selectron time distribution for a reweighted sample overlaid with the original lifetime and the target lifetime.

B.4 ASSESSMENT OF LIFETIME REWEIGHTING ACCURACY

Given the finite statistics in the simulated samples, and the finite extent of the ATLAS detector, the reweighting procedure (which relies on the ratio of exponential decay distributions with different time constants) is only valid over a finite range of source-target pairs. As an example, the shape of a distribution arising from a slepton sample with a mean lifetime 1 ns sample will suffer from lack of statistics in the tails when reweighting to 60 ns distribution (independent of the simulated slepton mass). This will lead to disagreement in shape and statistics when comparing reweighted distributions to a generated 60 ns model point at the same mass as the source lifetime. To assess the validity of




Figure B.4: Example of one-dimensional lifetime reweighting for the selectron truth-matched to the electron used in the BDT. The distribution of the selectron time is shown for the sample reweighted from 60 ns to 1 ns (cyan), as well as the original 1 ns (black) and original 60 ns (magenta) samples. The reweighted sample agrees well with the original target lifetime sample.

reweighting from a given source lifetime to a given target lifetime, the ratio of the normalization between reweighted samples and generated samples over a range of target lifetimes are compared. The normalization ratios for the one-dimensional reweighting procedure are summarized in Figure B.5. Based on the table shown in this figure, the 10 ns samples at each mass were used for reweighting (given reasonable number of generated events, and reasonably long lifetime) and a 20% normalization uncertainty was assessed for reweighting to lifetimes less than 100 ns, with a 50% uncertainty assessed for reweighting to lifetimes greater than 100 ns. This uncertainty is meant to give a conservative estimate on the exclusion for very high lifetimes which are not generated due to low efficiency of accepted events. To verify that the one-dimensional reweighting procedure is valid when applied to EMBDT regions, the BDT score between source and target pairs is compared in Figure B.6, with reasonable agreement observed.




Figure B.5: Ratio of expected signal yields in eBDT SRs between generated lifetimes and reweighted lifetimes. For 50 GeV (left), 100 GeV (middle) and 200 GeV (right) selectron samples, yields of generated samples at various lifetimes (table columns) are compared to yields for a 10 ns sample reweighted to each target lifetime (table rows). When a 1 ns sample at a mass of 50 GeV is reweighted to a 60 ns lifetime, the resulting BDT shape is similar to that of a generated 60 ns sample, but the predicted yield is 56% of that in the generated sample.




Figure B.6: BDT distribution in the eBDT SR for multiple reweighted samples. The reweighted distribution is overlaid with the original and target samples. The source lifetime (magenta) and reweighted (cyan) distributions are very similar, as the reweighted distribution lies directly over the original curve. Reasonable agreement is observed between the reweighted distributions (cyan) and the original target lifetime distributions (black).

APPENDIX C:

TRT WEDGE VETO

During 2023 datataking, an entire wedge of the Transition Radiation Tracker (TRT) corresponding to a region of ϕ in the barrel of the detector ($|\eta| < 1.37$) was disabled due to cooling issues, as shown in Figure C.1. This led to a large number of standard tracks without TRT extensions for $|\eta| < 1.37$, $\phi \in (-1.257, -1.759)$ in the 2023 dataset. Typically, one expects approximately thirty TRT hits per track, but the average number of hits per track drops rapidly to zero in this area of the detector, as shown in Figure C.2. The number of tracks without TRT extension is 10-20% for most of $(|\eta|, \phi)$ range covered by the ATLAS ID, but this drops to 100% in this region, referred to as the *TRT hole*. The effect is amplified for LRT tracking objects with respect to those reconstructed via standard tracking, due to relaxed requirements for LRT tracking (in particular, the removal of the typical requirement in electron ID that the track have at least one pixel hit). The effect on electrons falling into the eBDT CR/VR/SR is shown in Figure C.3 for data and simulated signal used in BDT training, where a significant spike in the number of LRT tracking electrons is seen in 2023 data in this region of ϕ , but not in simulated signal. The effect was masked in the $(|\eta|, \phi)$ coverage map constructed before unblinding (such plots are typically required for all ATLAS SUSY analyses), as the plots combined data from 2022 and 2023, and the logarithmic scale used in the plots hides the impact of the TRT hole. The effect of the TRT hole on electrons used in the BDT analysis was found when eBDT and γ BDT channels were

unblinded. In the discovery region, 5 events were found, where 1 event was expected. Upon further scrutiny, it was found that in two of the five events the signal electron had zero Pixel hits, zero TRT hits (i.e. no TRT extension) and eight SCT hits. In addition, the electron in each of these two events had coordinates within the $(|\eta|, \phi)$ region affected by the TRT hole.




Figure C.1: Section of TRT disabled in 2023 due to cooling issues.

After investigating the BDT score and BDT inputs for barrel electrons both inside and outside $|\eta| < 1.37, \phi \in (-1.257, -1.759)$, it was found that a significant "signal-like" bias in BDT score is observed for data electrons in the CR and SR in the TRT hole region. This bias is shown in Figure C.4.

The input variables to the BDT score were also studied inside and outside of the TRT hole to determine the most significant source of the bias. It was found that most of the variables had no discernable sculpting in the TRT hole, with the exception of $\Delta p_T/p_T$ and track χ^2/DoF . It is likely that a lack of TRT hits on the track biases the track p_T substantially, and the lack of TRT hits biases the number of degrees of freedom on the track fit as well as the possible χ^2 values. The biases on these two inputs are shown in Figure C.5.

Number of TRT hits per track, eta-phi profile




Figure C.2: Number of TRT hits per track as a function of η , ϕ from example run 452241 in 2023. For $\phi \in (-1.257, -1.759)$ there is a significant drop in efficiency of TRT hits per track.




Figure C.3: In 2023 eBDT CR data (right), a spike in number of electrons in the region where TRT is disabled is observed. This structure is not observed in the simulated signal (left) used to train the eBDT.




Figure C.4: eBDT score distribution for CR data (left) inside and outside of region where TRT is disabled in 2023. CR+SR eBDT score distribution (right) inclusive of TRT hole compared to compared to distribution in this region. Both distributions show a clear "signal-like" bias in the TRT-disabled region compared to the bulk of $(|\eta|, \phi)$ space.




Figure C.5: Shape of $\Delta p_T/p_T$ and track χ^2/DoF distributions inside the region where TRT is disabled, compared to the bulk of $(|\eta|, \phi)$.

Due to the shape bias introduced to background events in the TRT hole region, the BDT score of objects falling into the TRT hole is an unreliable predictor and therefore events in the 2023 dataset with $|\eta| < 1.37, \phi \in (-1.257, -1.759)$ are vetoed from the eBDT before the fit is performed. This veto is applied to simulated signal and to data. This results in an acceptance loss of approximately 5.4% in simulated signal, and removes 138 events from the eBDT CR and 154 events from the eBDT SR in data. The veto is additionally applied to electrons in the ABCD regions of the analysis for consistency, and to muons to avoid biases. However, the veto has little impact on the expected background for the ABCD SRs.